

WSARE: What's Strange About Recent Events?

Weng-Keen Wong, M.Sc.
Department of Computer Science
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA, 15213
Phone: (412)268-3052
Fax: (412)268-7350
Email: wkw@cs.cmu.edu

Andrew Moore, Ph.D.
Department of Computer Science
Carnegie Mellon University

Gregory Cooper, M.D., Ph.D.
Center for Biomedical Informatics
University of Pittsburgh

Michael Wagner, M.D., Ph.D.
Center for Biomedical Informatics
University of Pittsburgh

This research was funded by a grant from the DARPA IAO Biosurveillance Program and by the
State of Pennsylvania

Abstract

This paper, which is a shortened version of (1), presents an algorithm for performing early detection of disease outbreaks by searching a database of emergency department cases for anomalous patterns. Traditional techniques for anomaly detection are unsatisfactory for this problem because they identify individual data points that are rare due to particular combinations of features. Thus, these traditional algorithms discover isolated outliers of particularly strange events, such as someone accidentally shooting their ear, that are not indicative of a new outbreak. Instead, we would like to detect groups with specific characteristics whose recent pattern of illness is anomalous relative to historical patterns. We propose using an anomaly detection algorithm that would characterize each anomalous pattern with a rule. The significance of each rule would be carefully evaluated using Fisher's Exact Test and a randomization test. In this study, we compared our algorithm with a standard detection algorithm by measuring the number of false positives and the timeliness of detection. Simulated data, produced by a simulator that creates the effects of an epidemic on a city, were used for evaluation. The results indicate that our algorithm has significantly better detection times for common significance thresholds while having a slightly higher false positive rate.

Abstract word count: 202

Text word count: 3660

Keywords: Syndromic surveillance, data mining, anomaly detection, detection algorithm, multiple hypothesis testing

Introduction

This paper is a shortened version of (1). Detection systems typically monitor multidimensional temporal data for anomalies and raise an alert upon discovery of any deviations from the norm. For example, in the case of an intrusion detection system, an anomaly would indicate a possible breach of security (2, 3, 4). Although, early disease outbreak detection appears to be similar to traditional anomaly detection systems, shortcomings in these systems, which we will illustrate, limit their usefulness in early disease outbreak detection.

In our database of emergency department (ED) cases from several hospitals in a city, each record contains information about the individual who was admitted to the ED. This information includes age, gender, symptoms exhibited, home location, work location, and time admitted. (To maintain patient confidentiality, personal identifying information, such as patient's names, addresses, and identification numbers, were not in the dataset used in this research.) Clearly, when an epidemic sweeps through a region, there will be extreme increases in the number of ED visits. While these dramatic upswings are readily apparent during the late stages of an epidemic, the challenge is to detect the outbreak during its early stages and mitigate its effects. Different diseases cause different signals to appear in temporal, spatial, and demographic data (5). In order for any anomaly detection algorithm to be successful in early detection of disease outbreaks, it must be able to detect abnormalities in these three aspects of ED data.

A simplistic first approach would be to report an ED case as an anomaly if it has a rare value for some attribute. For example, we would signal an anomaly if we encountered a patient more than 100 years old. While this method detects the outliers for a single attribute, it fails to identify anomalies that occur due to combinations of features that by themselves might not be abnormal but together would be unusual. For instance, the first technique would not find anomalies in cases where the patients were male and under the age of 30 but exhibited symptoms

associated with a disease that affects primarily female senior citizens. Fortunately, there are plenty of anomaly detection algorithms that can identify outliers in multidimensional feature space. Typically these detection algorithms build a probabilistic model of the “normal” data using a variety of techniques such as neural nets (6) and a mixture of naive Bayes submodels (7).

Even that kind of sophisticated outlier detection, however, is insufficient for our purposes. Outlier detection succeeds at finding data points that are rare based on the underlying density, but these data points are treated in isolation from each other. Early epidemic detection, on the other hand, hinges on identifying anomalous groups, which we will refer to as *anomalous patterns*. Specifically, we want to know if the recent proportion of a group with specific characteristics is anomalous based on what the proportion is normally. This approach is closely related to the work done by Bay and Pazzani (8) in mining data from contrast sets and is somewhat similar to itemset mining (9). Traditional outlier detection, on the other hand, will likely return isolated irregularities that are insignificant to the early detection system.

We might, then, argue that aggregate daily counts of a single attribute or combination of attributes should be monitored in order to detect an anomalous group. For instance, we could monitor the daily number of people appearing in the ED with respiratory problems. A naive detector would determine the mean and variance of the monitored signal over a training set that is assumed to capture the normal behavior of the system. Then, a threshold would be established based on these values. Whenever the daily count exceeds this threshold, an alert would be raised. This technique works well if the monitored features are known. The spatial, temporal, and demographic signatures of diseases are simply too extensive, however, for us to know a priori which features to monitor. We could well miss some combination of features that would indicate an outbreak of a particular disease. Thus, we need an algorithm that is able to detect anomalous patterns rather than pre-defined anomalies.

Our approach to this problem uses a rule-based anomaly pattern detector. Each anomalous pattern is summarized by a rule, which in our current implementation consists of one or two components. Each component takes the form $X_i = V_j^i$, where X_i is the i th feature and V_j^i is the j th value of that feature. Multiple components are joined by a logical AND. For example, a two component rule would be Gender = Male AND Age Decile = 4. One benefit to a rule-based system is that the rules are easily understood by a non-statistician.

We need to be wary, however, of the pitfalls of rule-based anomaly pattern detection. Since we are finding anomalous patterns rather than isolated anomalies, we will be performing multiple hypothesis tests. When multiple hypothesis tests are performed, the probability of a false positive becomes inflated unless a correction is made (10). In addition, as we add more components to a rule, overfitting becomes a serious concern. Thus, a careful evaluation of significance is clearly needed. Furthermore, temporal healthcare data used for disease outbreak detection, are frequently subject to “seasonal” variations. As an example, the number of influenza cases is typically higher during winter than summer. Additionally, the number of ED visits varies between weekends and weekdays. The definition of what is normal will change depending on these variations.

Rule-based Anomaly Pattern Detection

The basic question asked by all detection systems is whether anything strange has occurred in recent events. This question requires defining what it means to be recent and what it means to be strange. Our algorithm considers all patient records falling on the current day under evaluation to be recent events. Note that this definition of recent is not restrictive -- our approach is fully general and “recent” can be defined to include all events within some other time period. In order

to define an anomaly, we need to establish what is “normal.” Our algorithm is intended to be applied to a database of ED cases, and we need to account for environmental factors such as weekend versus weekday differences in the number of cases. Consequently, normal behavior is assumed to be captured by the events occurring on the days that are exactly five, six, seven, and eight weeks prior to the day under consideration. The definition of what is normal can be easily modified to another time period without major changes to our algorithm. We will refer to the events that fit a certain rule for the current day as C_{today} . Similarly, the number of cases matching the same rule from five to eight weeks ago will be called C_{other} .

From this point on, we will refer to our algorithm as WSARE, which is an abbreviation for “what’s strange about recent events.” WSARE operates on discrete data sets with the aim of finding rules that characterize significant patterns of anomalies. Due to computational issues, the number of components for these rules is two or less. Our description of the rule-finding algorithm will begin with an overview followed by a more detailed example.

Overview of WSARE

The best rule for a day is found by considering all possible one- and two-component rules governing events occurring on that day and returning the one with the best “score.” The score is determined by comparing the events on the current day with events in the past. The best scoring rule then has its p-value estimated by a randomization test. This p-value is the likelihood of finding a rule with as good a score under the hypothesis that the features of the case and the date are independent. The randomization-based p-value takes into account the effect of the multiple testing that went on during the rule search. If we were running the algorithm on a day-by-day basis we would end at this step. If we were looking at a history of several days, however, we would need the additional step of using the False Discovery Rate (FDR) method (10) to

determine which p-value is significant. The days with significant p-values are returned as anomalies.

One component rules

To illustrate this algorithm, suppose we have a large database of 1,000,000 ED records collected during a two-year span. This database contains approximately 1000 records a day, thereby yielding approximately 5000 records if we consider the cases for today plus those from five to eight weeks ago. We will refer to this record subset as DB_i , which corresponds to the recent event data set for day i . The algorithm proceeds as follows. For each day i , retrieve the records belonging to DB_i . We first consider all possible one-component rules. For every possible feature-value combination, obtain the counts C_{today} and C_{other} from the data set DB_i . As an example, suppose the feature under consideration is the Age Decile for the ED case. There are nine possible Age Decile values, ranging from zero to eight. We start with the rule Age Decile = 3 and count the number of cases for the current day i that have Age Decile = 3 and those that have Age Decile $\neq 3$. The cases from five to eight weeks ago are subsequently examined to obtain the counts for the cases matching the rule and those not matching the rule. The four values form a two-by-two contingency table such as the one shown in Table 1.

	C_{today}	C_{other}
Age Decile = 3	48	45
Age Decile $\neq 3$	86	220

Table 1: A Sample 2x2 Contingency Table

Scoring each one component rule

The next step is to evaluate the “score” of the rule using a test in which the null hypothesis is the independence of the row and column attributes of the two-by-two contingency table. In effect, this hypothesis test measures how different the distribution for C_{today} is compared with that of C_{other} . This test will generate a p-value, which we will call the *score* in order to differentiate it from the p-value obtained later from the randomization test. We use Fisher's Exact Test to find the score for each rule. Running Fisher's Exact Test (11) on Table 1 yields a score of 0.00005058, which indicates that the count C_{today} for cases matching the rule Age Decile = 3 is significantly different from the count C_{other} .

Two component rules

At this point, the best one-component rule for a particular day has been found. We will refer to the best one-component rule for day i as BR_i^1 . The algorithm then attempts to find the best two-component rule for the day by adding on the component to BR_i^1 that yields the best score for the resulting two-component rule, which we will refer to as BR_i^2 . BR_i^2 , however, may not be an improvement over BR_i^1 . We need to perform further hypothesis tests to determine if the presence of either component has a significant effect. For further details on the creation of two-component rules, consult (1).

Finding the p-value for a rule

The algorithm for determining scores shown above is extremely prone to overfitting. Even if data were generated randomly, most single rules would have insignificant p-values but the best rule would be significant if we had searched more than 1000 possible rules. In order to illustrate this point, suppose we follow the standard practice of rejecting the null hypothesis when the p-

value is $< \alpha$, where $\alpha = 0.05$. In the case of a single hypothesis test, the probability of making a false discovery under the null hypothesis would be α , which equals 0.05. On the other hand, if we perform 1000 hypothesis tests, one for each possible rule under consideration, then the probability of making a false discovery could be as bad as $1 - (1 - 0.05)^{1000} \approx 1$, which is much greater than 0.05 (12). Thus, if our algorithm returns a significant p-value, we cannot accept it at face value without adding an adjustment for the multiple hypothesis tests we performed. We address this problem by using a randomization test in which the date and each ED case features are assumed to be independent. In this test, the case features in the data set DB_i remain the same for each record, but the date field is shuffled among records from the current day and records from five to eight weeks ago. The best rule is obtained on this randomized dataset. We repeat this procedure 1000 times and obtain a compensated p-value for BR_i , which we will refer to as CPV_i . The full description for the randomization test is given in (1).

Using FDR to determine which p-values are significant

WSARE can be used on a day-to-day basis similar to an online algorithm or it can be used to review a history of several days to report all significantly anomalous patterns. When using our algorithm on a day-to-day basis, the compensated p-value CPV_i obtained for the current day through the randomization tests can be interpreted at face value. When analyzing historical data, however, we need to compare the CPV values for each day in the history, thereby creating a second overfitting opportunity due to yet another multiple hypothesis testing problem. We address this problem through the use of the False Discover Rate (FDR) method (10, 12). For an in-depth discussion of the use of FDR in WSARE, see the related section in (1).

Results

We evaluated our algorithm using data from a simulator that simulated (to a first approximation) the effects of an epidemic on a grid world populated by people of varying characteristics. Thus, whenever an infected person exhibited the monitored symptom, an entry would be added to the log file to simulate an ED record.

Our results were obtained by running the simulator for 180 simulated days with the epidemic, named Epidemic0, introduced to the environment on the 90th day. Epidemic0 had a target demographic group of males 50-59 years old. Additionally, there were nine non-epidemic background diseases that spontaneously appeared at random points in the simulation. At certain stages, these background diseases caused infected people to display the monitored symptom. These background diseases had low infection probabilities as they were intended to provide a baseline for the number of ED cases. Our initial publication on this subject (1) contains a detailed description of the simulator and the experimental settings.

Evaluation of performance

We treated our algorithm as if it ran on a day-by-day basis. Thus, for each day in the simulation, WSARE was asked to determine if the events on the current day were anomalous. We evaluated the performance of WSARE against a standard anomaly detection algorithm that treated a day as anomalous when the daily count of ED cases for the monitor symptom exceeded a threshold.

The standard detector was applied to the ED case data from day 30 to day 89 in the simulation to obtain the mean μ and variance σ^2 . The threshold was calculated by the formula below, in which Φ^{-1} is the inverse to the cumulative distribution function of a standard normal.

$$\text{threshold} = \mu + \sigma * \Phi^{-1}\left(1 - \frac{\text{p-value}}{2}\right)$$

Both the standard algorithm and WSARE were tested using five levels of p-values (0.1, 0.05, 0.01, 0.005, and 0.001). In order to evaluate the performance of the algorithms, we measured the number of false positives and the number of days until the epidemic was detected. The specific rules for determining detection time and counting false positives are stated in the full version of this paper (1).

Figures 1 and 2 plot the detection time in days versus the number of false positives for five different p-value thresholds used in both the standard algorithm and WSARE. In Figure 1, the error bars for detection time and false positives are shown. Figure 2 fills in the lines to illustrate the asymptotic behavior of the curves. These values were generated by taking the average of 100 runs of the simulation.

Results from Simulated Data

These results indicate that for p-value thresholds above 0.01, the detection time for WSARE is significantly smaller than that of the standard algorithm. On the other hand, as the p-value threshold decreases, the detection time for WSARE is somewhat worse than that of the standard algorithm. Choosing an extremely low threshold would be unprofitable, however, since all anomalies except those at an unusually high significance level would be ignored. For example, using a threshold of 0.01 corresponds to a 99% significance level.

The results also demonstrate that WSARE signals more false positives than the standard algorithm for higher p-value thresholds. Although this behavior is not desirable, it is tolerable since the number of false positives produced by WSARE differs by a small amount from the count generated by the standard algorithm. In Figure 1, there are at most three more false positives identified by WSARE that were not identified by the standard algorithm.

We now show some of the rules learned by WSARE. The rules below were obtained from one of the result-generating simulations.

Rule 1: Sat Day97 (daynum 97, dayindex 97)

SCORE = -0.00000011 PVALUE = 0.00249875

33.33% (16/ 48) of today's cases have Age Decile = 5 and Gender = Male

3.85% (7/182) of other cases have Age Decile = 5 and Gender = Male

Rule 2: Tue Day100 (daynum 100, dayindex 100)

SCORE = -0.00001093 PVALUE = 0.02698651

30.19% (16/ 53) of today's cases have Age Decile = 5 and Column less than 25

6.19% (12/194) of other cases have Age Decile = 5 and Column less than 25

In Rule 1, WSARE demonstrates that it is capable of finding the target demographic group that Epidemic0 infects. This rule proves to be significant above the 99% level. On the other hand, Rule 2 discovers something that was not deliberately hardcoded into Epidemic0. Rule 2 states that on Day 100, an unusually large number of cases involves people in their fifties in the left half of the grid. Since we designed the people in the simulation to interact with places that are in geographic proximity to their homes, we suspected that the locality of interaction of infected individuals would form some spatial clusters of ED cases. Upon further inspection of the log files, we discovered that 12 of the 16 cases from the current day that satisfied this rule were, in fact, caused by Epidemic0. This example illustrates the capability of WSARE to detect significant anomalous patterns that are completely unexpected.

Results from Real ED data

We also ran WSARE on an actual ED data collected from hospitals in a major US city. This database contained approximately 70000 records collected during a period of 505 days. Since we are looking at historical data, we need to use FDR to determine which of the p-values is significant. The results are shown below with α for FDR equal to 0.1.

Rule 1: Tue 05-16-2000 (daynum 36661, dayindex 18)

SCORE = -0.00000000 PVALUE = 0.00000000

32.84% (44/134) of today's cases have Time Of Day after 6:00 pm

90.00% (27/ 30) of other cases have Time Of Day after 6:00 pm

Rule 2: Fri 06-30-2000 (daynum 36706, dayindex 63)

SCORE = -0.00000000 PVALUE = 0.00000000

19.40% (26/134) of today's cases have Place = NE and Lat = d

5.71% (16/280) of other cases have Place = NE and Lat = d

Rule 3: Wed 09-06-2000 (daynum 36774, dayindex 131)

SCORE = -0.00000000 PVALUE = 0.00000000

17.16% (23/134) of today's cases have Prodrome = Respiratory and Age less than 40

4.53% (12/265) of other cases have Prodrome = Respiratory and Age less than 40

Rule 4: Fri 12-01-2000 (daynum 36860, dayindex 217)

SCORE = -0.00000000 PVALUE = 0.00000000

22.88% (27/118) of today's cases have Time Of Day after 6:00 pm and Lat = s

8.10% (20/247) of other cases have Time Of Day after 6:00 pm and Lat = s

Rule 5: Sat 12-23-2000 (daynum 36882, dayindex 239)

SCORE = -0.00000000 PVALUE = 0.00000000

18.25% (25/137) of today's cases have ICD9 = shortness of breath and Time Of Day before 3:00 pm

5.12% (15/293) of other cases have ICD9 = shortness of breath and Time Of Day before 3:00 pm

Rule 6: Fri 09-14-2001 (daynum 37147, dayindex 504)

SCORE = -0.00000000 PVALUE = 0.00000000

66.67% (30/ 45) of today's cases have Time Of Day before 10:00 am

18.42% (42/228) of other cases have Time Of Day before 10:00 am

Rule 1 notices that there are fewer cases after 6:00 pm, possibly due a lack of reporting by some hospitals. Rule 6 correctly identifies a larger volume of data being collected before 10:00 am on Day 504. Since Day 504 was the last day of this database, this irregularity was the result of the database being given to us in the morning.

We are beginning the process of using input from public health officials of the target city to help us validate and measure WSARE's performance.

Conclusion

WSARE has successfully identified anomalous patterns in the data. Our simulation results indicate that WSARE has significantly lower detection times than a standard detection algorithm, provided the p-value threshold is not extremely low. This should not be a problem, since most anomalies are reported at a significance level of 95 or 99%, corresponding, respectively, to p-values of 0.05 and 0.01. WSARE also has a slightly higher false positive rate than the standard algorithm. This difference, however, was approximately three more false positives in the worst case for our particular simulation.

We believe the three main innovations in this paper are:

1. Turning the problem of “detect the emergence of new patterns in recent data” into the question “is it possible to learn a propositional rule that can significantly distinguish whether records are most likely to have come from the recent past or from the more distant past?”
2. Incorporating several levels of significance tests into rule learning in order to avoid several levels of overfitting caused by intensive multiple testing
3. Examining the domain of early outbreak detection by means of machine-learning tools

Acknowledgements

Thanks to Howard Burkom, Rich Tsui, Bob Olszewski, Jeremy Espino and Jeff Schneider for their help and suggestions. The work was sponsored by a grant from the DARPA IAO Biosurveillance Program and by the State of Pennsylvania.

References

- (1) Wong W, Moore AW, Cooper G, Wagner M. Rule-based Anomaly Pattern Detection for Detecting Disease Outbreaks. In *Proceedings of the 18th National Conference on Artificial Intelligence (AAAI-02)* 2002. MIT Press. Also available online from <http://www.autonlab.org>
- (2) Lane T, Brodley CE. Temporal sequence learning and data reduction for anomaly detection. *ACM Transactions on Information and System Security* 1999; 2:295--331.
- (3) Eskin E. Anomaly detection over noisy data using learned probability distributions. In *Proceedings of the 2000 International Conference on Machine Learning (ICML-2000)* 2000.
- (4) Maxion RA, Tan K. MC. Anomaly detection in embedded systems. Technical Report CMU-CS-01-157, 2001; Carnegie Mellon University.
- (5) Wagner MM, Tsui FC, Espino JU, Dato VM, Sittig DF, Caruana RA, McGinnis LF, Deerfield DW, Druzdel MJ, Fridsma DB. The emerging science of very early detection of disease outbreaks. *Journal of Public Health Management Practice* 2001; 7(6):51--59.
- (6) Bishop CM. Novelty detection and neural network validation. *IEEE Proceedings - Vision, Image and Signal Processing* 1994; 141(4):217--222.
- (7) Hamerly G, Elkan C. Bayesian approaches to failure prediction for disk drives. In *Proceedings of the eighteenth international conference on machine learning*, 2001; 202--209. Morgan Kaufmann, San Francisco, CA.

- (8) Bay SD, Pazzani MJ. Detecting change in categorical data: Mining contrast sets. In *Knowledge Discovery and Data Mining*, 1999; 302--306.
- (9) Brin S, Motwani R, Ullman JD, Tsur S. Dynamic itemset counting and implication rules for market basket data. In Peckham, J., ed., *SIGMOD 1997, Proceedings ACM SIGMOD International Conference on Management of Data, May 13-15, 1997, Tucson, Arizona, USA*, 1997; 255--264. ACM Press.
- (10) Benjamini Y, Hochberg Y. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society, Series B*. 1995; 57:289--300.
- (11) Good P. *Permutation Tests - A Practical Guide to Resampling Methods for Testing Hypotheses*. New York: Springer-Verlag, 2nd edition, 2000.
- (12) Miller CJ, Genovese C, Nichol RC, Wasserman L, Connolly A, Reichart D, Hopkins A, Schneider J, Moore A. Controlling the false discovery rate in astrophysical data analysis. Technical report, 2001; Carnegie Mellon University.

Figures

Figure 1: Scatterplot of Detection Time Versus False Positives with Error Bars for Detection Time and False Positives




Figure 2: Plot of Detection Time Versus False Positives

