Performance of an Adaptive Anomaly Detection Algorithm for a Low Incidence Syndrome Before and After a Major Outbreak

Sylvia Halasz[1], Ph.D., Colin R. Goodall[1], Ph.D., John Allegra[2], M.D., Ph.D., Dennis Cochrane[2], M.D.

[1] *AT&T Labs – Research.* [2] *Emergency Medical Associates of New Jersey Research Foundation.*

Background

Ideal anomaly detection algorithms should detect both sudden and gradual changes, while keeping the background false positive alert rate at a tolerable level. Further, the algorithm needs to perform well when the need is to detect small outbreaks in low-incidence diseases. For example, when surveillance is done based on the specific ICD9 diagnosis of flu rather than a larger syndromic grouping, the baseline counts will generally be low, in the range of 0 or 1 per day even in a large sample of EDs.

Objective

Our goal was to determine the sensitivity of detection of various inserted outbreak sizes and shapes using a modified Holt-Winters (HWR) detection algorithm applied to daily flu count data before the flu season and after its peak. We compare our results to C3 of EARS [3].

Methods

We tested the algorithm on data based on ICD9 diagnosis of flu from 18 emergency departments (EDs) with simulated outbreaks of varying maxima and shapes injected at different times using a method described in [1]. Outbreaks were injected at 30 different starting dates in October, with all injected outbreaks ending before the flu season. We injected the same outbreaks with 60 consecutive starting dates beginning 1/2/04, eight days after the peak of the major outbreak. These sensitivities are reported for three 20-day periods separately in red: 1-20 days, 21-40 days, and 41-60 days. In the calculation of specificities all alerts outside of the simulated outbreak and the flu season were considered to be false. Sensitivity and average number of false alerts were calculated as an average over the different starting dates. The parameters of HWR were kept constant throughout this study.

Results

The time series for daily counts based on ICD9 diagnosis of flu is shown in the figure. The sensitivity and the average number of false alerts per year are shown in the table for five different outbreak profiles (counts per day).

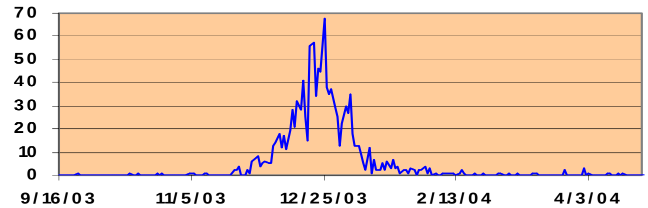

	HWR		C3	
Outbreak profile	**Sensitivity**	**Avg. false alerts/ yr**	**Sensitivity**	**Avg. false alerts/ yr**
1,4,2	1.00	0.77	0.67	0
	0, 0.25, 0.95 *		0, 0.15, 1 *	
2,9,3	1.00	0.65	0.60	0
	0, 0.85, 1.00 *		0, 0.65, 1 *	
2,3,2,1	1.00	0.85	0.67	0
	0, 0.35, 1.00 *		0, 0.4, 0.8 *	
4,7,4,1	1.00	0.63	0.67	0
	0, 0.70, 1.00 *		0, 0.75, 1 *	
1,1,2,2,2,1,1	1.00	0.73	1.00	0
	0, 0.20, 0.90 *		0, 0.35, 0.85 *	

* values for the three consecutive 20-day periods after the peak of the outbreak

Conclusion

For low daily background counts the HWR algorithm had greater sensitivity for small injected outbreaks with a small increase in the number of false alerts per year compared to the C3 algorithm. All sensitivities decreased markedly after the peak of the major outbreak and only recovered more than a month after the peak. As expected, recovery was faster for larger outbreaks. Further work is needed to detect smaller outbreaks sooner after a major peak.

References

[1] H. S. Burkom: Accessible Alerting Algorithms for Biosurveillance. 2005 National Syndromic Surveillance Conference

[2] S. Halász, P. Brown, C.R. Goodall, A. Lent, D. Cochrane, J. R. Allegra: Detecting the Start of the Flu Season. Biosurveillance Workshop 2007, LNCS 4506, Springer.

[3] Y. Zhu, W. Wang, D. Arrubin, P.Yu in MMWR vol. 54. (2005).