Monitoring Food Safety by Detecting Patterns in Consumer Complaints

Artur Dubrawski[1], Kimberly Elenberg[2], Andrew Moore[1] and Maheshkumar Sabhnani[1]

(1) The Robotics Institute, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213
Artur.Dubrawski@cs.cmu.edu, Andrew.W.Moore@cs.cmu.edu, Maheshkumar.Sabhnani@cs.cmu.edu
(2) Food Safety Inspection Service, USDA, 1400 Independence Ave. S.W., Washington, D.C. 20250, Kimberly.Elenberg@fsis.usda.gov

Abstract

EPFC (Emerging Patterns in Food Complaints) is the analytical component of the Consumer Complaint Monitoring System, designed to help the food safety officials to efficiently and effectively monitor incoming reports of adverse effects of food on its consumers. These reports, collected in a passive surveillance mode, contain multi-dimensional, heterogeneous and sparse snippets of specific information about the consumers' demographics, the kinds, brands and sources of the food involved, symptoms of possible sickness, characteristics of foreign objects which could have been found in food, involved locations and times of occurrences, etc. Statistical data mining component of the system empowers its users, allowing for increased accuracy, specificity and timeliness of detection of naturally occurring problems as well as of potential acts of agro-terrorism. The system's main purpose is to enhance discovery and mitigation of food borne threats to public health in the USDA Food Safety Inspection Service regulated products. As such, it is being envisioned as one of the key components of the nationwide bio-security protection infrastructure. It has been accepted for use and it is currently going through the final stages of deployment. This paper explains the motivation, key design concepts and reports the system's utility and performance observed so far.

Introduction

Challenges of Food Safety. The United States agriculture generates more than one trillion dollars worth of economic activity and over 60 billion dollars in food exports (GAO 2005). Unfortunately, the agriculture and food systems are vulnerable to disease, pest, or poisonous agents that occur naturally, are unintentionally introduced, or are intentionally delivered by acts of terrorism. In the US, there are 76 million recorded cases of food borne illness occurring every year, and about 5,000 of them are terminal (Mead et al. 1999). Costs of food borne illnesses or injuries adversely impact the food industry, households, and health sector. The attacks on September 11th, 2001, brought special attention to vulnerability of the US food supply system. It is an extensive, open, interconnected, diverse, and complex structure providing potential targets for terrorist attacks which could have catastrophic health and economic effects.

Consumer Complaint Monitoring System. In order to provide the best possible protection against all those threats, the USDA Food Safety Inspection Service (FSIS) is responsible for numerous initiatives aimed at improving food supply safety and security, such as monitoring and surveillance activities, which require real time tracking and assessment of data in an effort to serve as early warning systems. A significant example is the implementation of an adverse event consumer based surveillance system for food safety and food security referred to as the Consumer Complaint Monitoring System (CCMS). Located within the FSIS Office of Public Health and Safety, the CCMS has evaluated nearly 4,000 consumer complaints from January 2001 through January 2005. Results of these evaluations have led to recalls of unwholesome adulterated products, improved quality assurance, and analyzes of hazards and critical control points in slaughter and processing plants.

The CCMS is an intranet electronic database used to record, evaluate, and track all adverse food events involving specifically meat, poultry, and egg products reported to FSIS. The two main portals of entry for consumer complaints are through phone calls to the field Compliance Officers stationed throughout the United States and through the 1-800 phone number of the Meat and Poultry Hotline. Additional data sources include consumer complaints reported by a state or local health department or another federal agency, such as the Food and Drug Administration, as well as complaints that involve imported products that have been re-inspected at the port of entry. Most of the consumer complaints involve: illnesses that occurred after eating a product; injuries that occurred while eating a product; foreign objects that were found in a product; allergic reactions that occurred after eating a product; suspected under-processing of ready-to-eat product; allegations of improper labeling of a product; and dissatisfaction with the quality of a product.

Upon arrival of the new complaint case, the FSIS analysts search the CCMS database for similar complaints. Multiple cases bearing similar features reported independently on each other may indicate a possibility of a serious food safety problem. Analysts initiate an investigation and mitigating actions for all possible events detected that indicate potential health safety, or security issues. Investigation examples include: a laboratory confirmed food borne illness, an alleged allergic reaction due to a previously diagnosed food allergy to an unlabeled ingredient, or signs that a ready-to-eat product may be under-processed. If the CCMS database contains more than one apparently related complaint about a foreign material in a product produced at a particular establishment an investigation is also initiated.

Monitoring Emerging Patterns in Food Complaints with EPFC. The EPFC is the analytic core of the CCMS. It is the first practically applied instance of *Tip Monitor*, a broader concept of a statistical data mining approach to practical alerting from multivariate rare-event data (Sabhnani et al. 2005). The purpose of Tip Monitor is to identify small groups of significantly related records in an incoming stream of event-based data such as anonymous individual patient health events involving chief complaint strings, prescription orders, public safety hotlines or customer complaints. Very often, such data is sparse, noisy and it may contain little and spotty evidence of potentially crucial coincidences. That last feature makes it very hard to detect important events with more traditional approaches used by bio-surveillance analysts such as scan statistics or multivariate time series analysis, since they are designed to benefit from ample evidence (Buckeridge et al. 2003, Neill and Moore 2004, Wong et al. 2003).

Suppose that among the chief complaint strings of two unrelated patients in the same city on the same date there was mention of bloody stools in pediatric cases. The multiple mentions of “bloody stools” or “pediatric” might not be surprising, but the tying together of these two factors, given matching geographic locations and timings of reporting, is sufficiently rare that seeing only two such cases is of interest. This was precisely the evidence that was the first noticeable signal of the tragic Walkerton, Canada, waterborne bacterial gastroenteritis outbreak caused by contamination of tap water in May 2000 (Mackay 2002). That weak signal was spotted by an astute physician, not by a surveillance system. Reliable automated detection of such signals in multivariate data requires new analytic approaches.

Similar circumstances accompany the task of detecting systematic associations between individual food consumer complaints received by the USDA. FSIS analysts need to be alerted even if merely two independently collected complaints seem to be substantially related to each other. Potential relevance must be evaluated using spotty multivariate data subjectively reported by the complainants. Even seasoned analysts agree that it is a tedious task, prone to subjective judgment and human error. This difficulty is augmented by the underlying complexity of data interpretation. For instance, a food borne illness often presents itself with flu-like symptoms such as nausea, vomiting, diarrhea, or fever. That makes it difficult for analysts to determine if illness was caused by bacteria or other pathogens in food and if the pathogens occurred naturally or intentionally, however illness could also be caused by chemicals or heavy metals. The underlying complexity of the task at hand results in highly subjective judgments and it may lead to erroneous decisions.

There is a need for a statistical data mining system capable of automated and objective monitoring of the stream of incoming complaints for evidence of linkages with records from the recent past. It should help the analysts by alerting them early about emerging patterns in food complaints, and by focusing their attention to the most probable linkages. Effectively, it should improve their situation awareness, reliability of decisions and response times. The EPFC is aimed at fulfilling that need.

Approach

The EPFC is designed to screen sparse and noisy data for potential linkages between individual reports of adverse effects of food on its consumers. These reports contain multi-dimensional and heterogeneous snippets of specific information about the consumers’ demographics, the kinds, brands and sources of the food they ate, symptoms of sickness they may be experiencing, characteristics of foreign objects which could have been found in food, involved locations and times of occurrences, and so on.

In EPFC, the notion of two cases being similar is determined by a probabilistic model that is partially learned from historical data and partly obtained from experts. This is in contrast to simple pair-wise distance measures and string similarity scoring, which may indicate a correlation but not necessary a factual connection between the two events.

Knowledge of the domain expert is modeled into the form of a list of causal scenarios. Each scenario, such as for instance a malicious contamination of raw food at processing plant, or a product-nonspecific illness occurring in some local community, will focus on a specific subset of features of multivariate records of events, which would be considered relevant to the selected cause. The experts assist in defining the structure and parameters of the probabilistic model of conditional dependencies between features of the pairs of events. In the current version of the system, these dependencies are modeled by a Bayesian network, given a predefined assumed cause. A separate model is constructed this way for each individual scenario of interest. These models, pre-constructed by hand, are combined into one system using Bayes’ rule.

Mathematically, the EPFC estimates how likely it is for a newly reported complaint case X_n to be a close copy of some other case in the past data, X_i, if both have been generated by the same specific underlying cause labeled

C_k. We can directly ask for probability Q_{ik} that a hidden event H of causal category C_k has occurred, and it resulted in the occurrence of cases X_i and X_n.

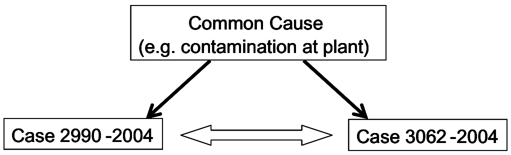

Fig. 1. In EPFC, similarity between two independent consumer complaints is judged under the assumption that they have been filed due to the same underlying problem (contamination of food at processing plant in this example).

These probabilities can be obtained for all considered causal categories C_k, and all old cases X_i to produce other directly useful measures such as the probability Q_k that X_n is a part of the consequence of a hidden event H of category C_k that has produced at least one other observed case, or the probability Q_i that X_n is related to X_i through at least one common consequence of a hidden event H of category C_k. Similarly, by marginalizing across all C_k's and X_i's we can obtain probability Q that X_n is causally related to any of the past cases with respect to any of the assumed causal scenarios. Very unusual new cases would have very low values of Q, which may be useful in screening for anomalous consumer reports.

The matches can be ordered according to the calculated probability scores (primarily Q_{ik}, but it also makes practical sense to consider Q_i and Q_k) and the top portion of the ranking list can then be reported to the human analysts. The users could focus their attention on investigating in more detail those of the past cases which were most likely to be relevant to the newly reported one, given specific common causes. The current implementation of EPFC reports top 10 unique pairing matches: 8 selected using Q_{ik} as the criterion of choice and one each maximizing Q_i and Q_k respectively.

Note that this reasoning also requires a model of normal, null-hypothesis cases, which can be learned from historical data. The full design of EPFC includes probabilistic models for the distribution of the observed features of reported cases, conditioned on the causal scenario category C_k. One way to formulate the complete model is expressed by the following formula which implements Bayes' rule:

$$Q_{ik} = \frac{I_{kn} \cdot R_{ik} \cdot K_{ikn}}{G_n\left(1 - \sum_k \sum_j R_{jk}\right) + \sum_k \left(I_{kn} \cdot \sum_j \left(R_{jk} \cdot K_{jkn}\right)\right)} \tag{1}$$

Here, G_n denotes the *full instantiation model*. It estimates the chance of seeing the particular new complaint X_n at all, based on past data, and it can be formulated as a density model over the existing collection of past complaints. We implement it as a Bayesian network. It learns its structure and parameters from more than 80 binary attributes derived from historical data. I_{kn} stands for *partial instantiation model* and it quantifies the chance of seeing the particular set of non-specific features of X_i, given its features specific to causal scenario C_k. Here we in fact use the same Bayesian network as above, but now it is being used in the inference mode. R_{ik} corresponds to the *relevance model* telling how likely it is for the historical case X_n to be caused by a reason matching causal scenario C_k. It is useful for pruning past cases which obviously cannot be due to the particular C_k. In turn, K_{ikn} denotes the *causal model*. It estimates the likelihood of the particular set of specific features of X_n to be linked with the corresponding features of case X_i if both cases were caused by scenario C_k. Note that the first component of the sum in the denominator quantifies how typical or strange is the current case given the complete collections of X_i's and C_k's.

Of the four components described above, the most elaborate is the causal model. It implements domain knowledge extracted from experts and it plays a key role in scoring possible matches. Currently, there are 42 specific causal scenarios implemented in EPFC (cf. Fig. 2), covering a range of possible assumed origins (such as food processing plant, point of purchase), complaint classifications (reported illness, foreign object found in food, allergic reaction) and product-specificity of the types of problems under consideration. Each individual causal scenario has a corresponding causal model. These models encapsulate experts' knowledge about what makes two reported complaint cases similar or dissimilar. They are based on a few criteria such as timeframe relevance of the two complaints, symptom lists match or foreign objects match (for illness and foreign object case types respectively), products match (source, timing, ingredients) and locations match (complainants, points of purchase).

Case Classification	Product specificity	Origin					
		Plant	POP	Community	Chain	Brand	Plant Co-loc
Illness	Yes	√	√	√	√	√	
	No	√	√	√	√	√	
	N/A						√
Foreign Object	Yes	√	√		√	√	
	No	√	√		√	√	
	N/A			√			√
Injury	Yes	√					
	No	√					
	N/A		√	√			
OT/OO/OC	Yes	√	√		√	√	
	No	√	√		√	√	
	N/A						
Allergy	Yes		√		√	√	
	No		√		√	√	
	N/A	√		√			√

Fig. 2. Matrix of causal scenarios implemented in the EPFC. The system consists of 42 predefined typical causal scenarios, and each of them is tested for relevance to the features of the analyzed pair of complaints.

Figure below depicts the processing flow leading to computation of K_{ikn}'s for two particular cases X_i and X_n and each possible scenario C_k. In this example, the most likely causal scenario that may explain similarity of features of the two cases under consideration is a community-related illness which is not attributable to a particular food product. It may be an effect of contamination of water supply, or have contagious etiology, rather than being related to a food supply problem. Note that EPFC is designed to also evaluate scenarios which are not directly related to food, but which may explain observations resembling characteristic effects of problems with food safety.

In the considered example, the particular causal scenario seems to be relevant to the pair X_n and X_i, because the particular subsets of features of the two cases, specific to this scenario, match closely. Community-related illness that is not product-specific requires that the illness symptoms reported in the two complaint cases significantly match in terms of their types and onset times, it is also required that the complainants locations do match, while the implicated food products do not match in terms of their origin.

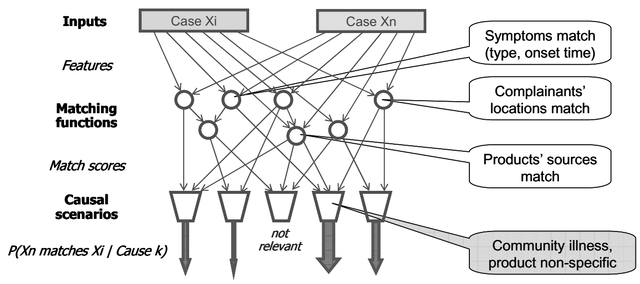

Fig. 3. Diagrammatic view of the processing flow leading to computation of likelihoods of the individual causal scenarios' responsibility for the apparent relationship between two consumer complaint cases.

Example contents of the EPFC output produced when analyzing a hypothetical case #2418:

- Total probability of #2418 to be causally related to any of the past records in the reference data is Q=0.85. #2418 does not look anomalous.
- The 1st highest ranking past record is #2215 with the score Qik=0.0561; #2418 and #2215 are most likely caused by an illness-causing problem originating at the food processing plant XYZ and involve beef products.
- The 2nd highest ranking past record is #2023 with the score Qik=0.0175; #2418 and #2023 may be linked through some local environment phenomenon (such as tap water contamination or an outbreak of infectious gastrointestinal ailment) happening in the proximity of the location of the complainant associated with the query record.
- ...

Note that in general each complaint may be statistically explained by more than one hypothetical cause, and it can be found as potentially linked with more than one past case. The EPFC reports estimated probabilities of each such event and the final assessment is to be made by human analyst.

Evaluation

Observed Utility and User Perception. The EPFC is undergoing the final stages of security compliance tests (as a part of the new revision of CCMS labeled CCMS II) and it is due for full scale live deployment in the second half of FY 2006. It is receiving a very positive feedback from its evaluators and future users. When recently tested on historical CCMS data, it managed to instantaneously flag several related E. Coli O157:H7 cases. The relationship between those complaint cases, which arrived over a period of three weeks and were analyzed by two separate people, was not realized by the analysts until the relationship was signaled by EPFC. This has been possible due to the ability of EPFC to remain sensitive to small signals in multivariate data; even when data is spotty, noisy and even if it comes in a short supply (the CCMS system currently logs only about 20 complaints per week on average).

The EPFC resolves the problem of manually checking for all possible multivariate associations between all possible pairs of complaints over recent weeks, and it helps to efficiently allocate limited analytical and investigative resources. Its unique feature is the ability to reliably detect signals supported by very little data – significant alerts can be raised on the basis of a very few complaints from consumers, provided that the few complaints contain significantly similar and explicable root causes.

EPFC often brings the analysts' attention to pairs of cases that do not necessarily call for immediate concern within USDA regulation. For example, cases where the nature of the complaint reflects food borne illness may simply be an impersonation of an infirmity associated to the air or water in a given region. Analysts, identifying a collection of illness cases that appear to have no common relationship to FSIS regulated product but do demonstrate a geographic relationship, have the ability to share this information with state, local, or tribal public health departments. These public health departments play an extremely important role in investigating, mitigating, and responding to community-based health hazards. Cases that reflect a highly probable relationship are highlighted by EPFC and may be either acted upon by FSIS or relayed to the appropriate authorities for further investigation or processing.

Historically, such instances were likely to go undetected due to the complex nature of the data which presents difficulties for analysts. EPFC facilitates the gaining of intelligence and awareness from disparate data points. Verified related cases of consumer complaints and adverse event reports aid in the detection of pre-incident indicators. This relates to food safety, food security, or terrorist activity in FSIS regulated facilities or products, or the encountered community.

The users are pleased with the computational efficiency of EPFC. Measured on a single CPU Pentium IV 2.4GHz (512KB cache) with 2GB RAM, it needs less than 10 seconds to re-train models to incorporate new evidence, and about 0.5 second on average to make predictions about an individual target complaint, based on the current CCMS database. The near-real-time operation greatly enhances the user's perception of the overall usability of the system.

Quantitative Evaluation. Ideally, the impact of a surveillance system like EPFC should be quantified with an economic criterion such as Return On Investment. It should take into account the costs of development, installation and maintenance on one hand and the benefits stemming from improved timeliness, accuracy and specificity of detection of adverse events on the other. Unfortunately, there are no established standards or widely accepted and easy to implement ways of doing so. A recent preliminary study (Jacobs et al. 2005) reviews the current situation and hints on a generic approach to estimate ROI for bio-surveillance systems. At this point, we are not in a position to execute such an evaluation for EPFC.

Hence, in the following, we focus on empirical evaluation of the system's predictive accuracy, based on labeled training data acquired from the domain experts, who have not participated directly in the development of the system or its underlying assumptions. We consider independence of evaluators to be an important factor ensuring objectivity. On the other hand, this requirement has drastically limited the amount of attainable testing data. In the process, we faced two major obstacles: the small amount of labeled testing data; and a high variance of scores caused by subjectivity of the individual expert opinions.

We have managed to execute two feedback collection cycles so far. One resulted in a labeled set of 158 semi-randomly selected pairs of complaint cases, each annotated by the domain expert in a binary way as either a match or a mismatch. The other set consisted of only 14 randomly selected pairs of complaints, however they were scored independently by three FSIS analysts.

The later set was used to quantify objectivity of the analysts' judgments. Each participant was supposed to rate each pair of complaints at a scale from 1 to 5; with 1 meaning the perfect match; 2: high similarity warrants further consideration; 3: some critical features match, finding is good to know though it does not constitute a potential immediate threat; 4: most features do not match; and 5: the pair of cases is completely unrelated. The three subsets of results, normalized to $N(0,1)$ individually for each analyst A, B and C, were then correlated against each other. The graph in Fig. 4. depicts the scatter plot of 3x14=52 data points (labeled as "Experts") obtained by concatenating 3 sets of pairs of scores: {A,B},{B,C} and {C,A}. Both axes of the graph correspond to normalized ratings obtained from analysts A, B and C, and a given diamond shaped point is located at the coordinates corresponding to a normalized score assigned to a pair of cases by one person (e.g. A - horizontal coordinate) vs. the respective score assigned to the same pair of cases by another person (e.g. B - vertical axis). The comparison is being performed for each of the possible couplings of available analysts: A vs. B, B vs. C and C vs. A, resulting in 52 separate data points in the scatter plot. It is evident from the graph that the ratings often substantially vary from person to person. The coefficient of linear correlation calculated for the "Expert" scatter plot is about 0.8. This observation indicates a potential fragility of the empirical evaluation of the system, if it is to be based on such a highly subjective reference.

Graph in Fig. 4. contains also a scatter plot of the normalized ranks assigned to the considered 14 pairs of cases by the EPFC (vertical axis) vs. the corresponding normalized scores assigned by independent experts A, B and C (horizontal axis). The ranks have been obtained while using the complete set of available real historical data to execute searches. The coefficient of correlation between EPFC and human expert scores was 0.63.

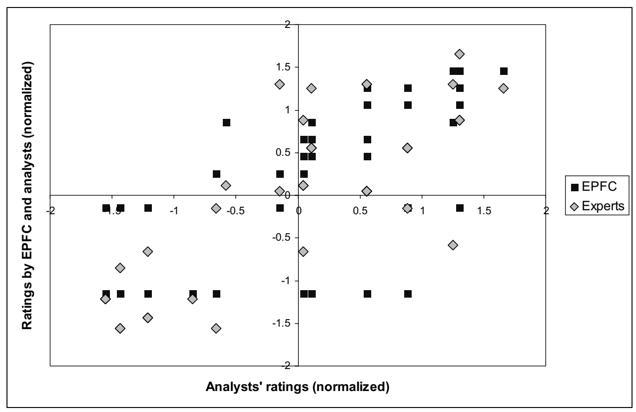

Fig. 4. Observed subjectivity of analysts' evaluations.

It is interesting to check accuracy of the system using human experts' scores as the baseline. Let us define a binary score for human experts such that it is equal 1 if the sum of ratings assigned by the three analysts to the evaluated pair of complaints equals 9 or less; and 0 otherwise. Let us also define a binary score for the EPFC as equal 1 if the analyzed pair of cases makes the top 10 on the ranking list of all Q_{ik} scores and its Q_{ik} is positive; or 0 otherwise. Then, 11 times out of 14 the EPFC correctly predicts the combined vote of the human experts. All three erroneous predictions are false positives. In fact, they turn out to be very close calls. One of them was related to community-level effects, which were not in the specific scope of responsibility of FSIS analysts, and thus it was assigned marginally interesting or not matching rates. The other two revealed enough similarities that the analysts admitted willing to know they existed, but they were not compelled to investigate at the first sight.

The other set of human-labeled data consisted of 158 pairs of cases, 68 of which were known as positive matches. We call this selection semi-random because the actual fraction of known matching pairs of complaints in the total number of possible pairs is extremely low. According to the domain experts, a new complaint case stands about 8-10% chance of having a close match in the recent past data at all, and about half of such cases would relate to more than one older complaint. In this experiment, the expert had to make a binary decision whether a given pair was sufficiently interesting to be labeled as a match. The EPFC (with binary scoring rule defined as above, and running searches through all available historical data) managed to predict correctly 80% of the time, raising false alerts in 11 instances and missing human-labeled true matches 21 times. That amounts to the specificity score of 87.8% and the sensitivity score of 81%. The ROC diagrams (labeled in Fig. 5 as real data curves) constructed using Q_{ik} scores and ranks (indices of the evaluated pairs of complaints on the ranking lists of Q_{ik}'s calculated for the given search cases), revealed AUC scores of 72.5% and 78.6% respectively.

The end users consider those empirical results obtained with real data perfectly acceptable. However, they seem to be slightly hard to judge given known subjectivity and fuzziness of the reference. Therefore, we executed another experiment: in controlled conditions, using synthetic data.

Synthetic data generation model is based on the real CCMS database entries. Using this model we can generate independent random complaints. We can also inject small clusters of noisy complaints linked through a known common cause. This generative model closely follows the actual data distributions for symptoms, foreign object types, products, and food processing plants, in order to produce synthetic data for illness and foreign object related causes. It also involves random insertion of typos in textual fields (such as product names, brands, etc.), as well as inducing missing entries, such that the end effect very closely resembles the real food complaint reports.

Fig. 5 shows the results of evaluating EPFC on a set of synthetic 4,600 complaints, spanning a period of 5 years, with about 10% of cases having one or two related complaints in the past. From this set, we randomly selected about 8,200 pairs of complaints, 10% of them were known positive matches. For each selected pair, we used EPFC to calculate the corresponding Q_{ik} and its rank in the similarity matrix Q. The resulting ROC curves are shown in Fig. 5, and the obtained AUC scores are 91.8% and 91.7% respectively. These results are substantially better than those obtained on the real data. It can be partially explained by the attained sharper distinction between known matches and random background in synthetic data relative to real data, and consequently a narrower margin for subjectivity in reference scoring.

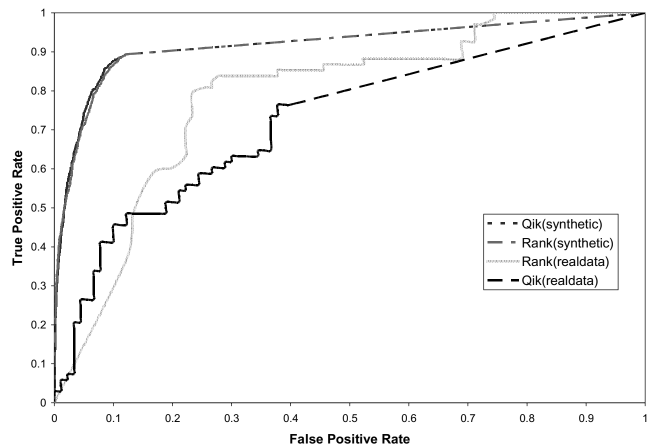

Fig. 5. ROC curves obtained during experiments with the EPFC.

Lessons Learned and Next Steps

Empirically, the most time consuming stage of system development was derivation of causal models. It involved a significant initial investment in terms of domain experts' time followed by multiple model refinement cycles. Our ongoing research leads towards active exploitation of the user feedback and self-, semi-supervised learning. This way we hope to substantially mitigate costs of estimation of causal models, and to make them adaptive to new trends in the collected data which would emerge during the regular use of the system.

The second important issue faced at this stage of the project was the limited availability of labeled training data. That made the quantitative evaluation of the system very hard. We plan to structure a runtime feedback acquisition process so that it would not be perceived as an additional burden upon already extremely busy food safety analysts. Instead it would emphasize the ultimate positive effect of frequently devoting a few minutes of spare time on the performance of the system. We hope we would collect much larger amount of labeled reference data that way. The next revision of the EPFC will also have to take into deeper consideration the fact that the users' feedback carries a substantial variance due to the natural subjectivity of the individual analysts' opinions.

The ongoing development of the CCMS focuses on incorporating new sources of consumer complaint data and on interoperability with other databases and surveillance systems run independently by other federal and state agencies, as well as with the National Bio-surveillance Integration System (NBIS). Currently, the available record of historical data is of limited size, and frequency of reporting new cases is relatively low as they come from federal-level sources only. Several state public health agencies have comparable or richer collections of data as well as means to collect new entries at comparable frequencies. Integration of those external sources of data should lead to improved reliability of predictions, though it may require additional work into scalability of the EPFC.

Conclusion

The EPFC at its current stage of development already lives up to its promise, as evidenced by testing on a collection of historical food complaints. It is receiving a very positive feedback from its evaluators and future users due to its high utility, sensitivity, and satisfactory accuracy.

It is able to overcome the limitations of typical health safety related event-based data by employing Bayesian techniques to model potential common causes obtained from domain experts' knowledge of scenarios of events triggered by specific causes of interest. In EPFC useful alerts can be generated using more specific information and on fewer cases than typically attainable in syndromic surveillance. In addition, it is sensitive to new emerging patterns of previously unknown or unanticipated adverse events.

There is a potential for secondary benefits from EPFC, going beyond its primary purpose of supporting food safety. They include giving food manufacturers a timely feedback on the safety of their products, which could positively impact stability and sustainable development of local economies which often heavily rely on food industry.

The EPFC is an instance of the more general Tip Monitor concept. As such, it illustrates the ability of this approach to become useful in other domains, where multivariate heterogeneous data comes in a relatively short supply and where early detection of relatively low amplitude signals is required. The natural areas of potential future applications of Tip Monitor include analyzing hospital records for signals of disease outbreaks, analyzing maintenance records for early evidence of systematic patterns of equipment failures, analyzing law enforcement reports, and so on.

Acknowledgement

This work has been performed under the US Government contract (Award #GS06K97BND0710 USDA/FSIS Consumer Complaint Monitoring System II).

References

Buckeridge, D.L.; Burkom, H.; Moore, A.; Pavlin, J.; Cutchis, P.; and Hogan, W. 2003. Evaluation of syndromic surveillance systems – design of an epidemic simulation model, *Morbidity and Mortality Weekly Report*, 2003. 53 (Supplement): 137-143.

GAO 2005. Protecting Against Agroterrorism. Report GAO-05-214, Government Accountability Office, March 2005.

Jacobs, B.; Greenwald, J..; Jain, A.; Park, A.; Wong D.; and Yang, M. 2005. Proposed Cost-Benefit Framework for Bio-surveillance Systems, Auton Lab Technical Report, Carnegie Mellon University, December 2005.

Mackay, B. 2002. Walkerton, 2 years later: Memory fades very quickly", *Canadian Medical Association Journal*, May 14, 2002. 166 (10).

Mead, P.S.; Slutsker, L.; Dietz, V.; McCaig, L.F.; Bresee, J.S.; Shapiro, C.; Griffin, P.M.; and Tauxe, R.V. 1999. Food-Related Illness and Death in the US, *Emerging Infectious Diseases*, 5(5), 1999. (Available from http://www.cdc.gov/ncidod/eid/vol5no5/mead.htm).

Neill D. and Moore A. 2004. Rapid detection of significant spatial clusters, *In Proceedings of the 10th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 256-265.

Sabhnani, M.; Neill, D.; Moore, A.; Dubrawski A.; and W.Wong, W.-K. 2005. Efficient Analytics for Effective Monitoring of Biomedical Security, *In Proceedings of the International Conference on Information and Automation*, Colombo, Sri Lanka, December 2005.

Wong, W.-K.; Moore, A.; Cooper, G.; and Wagner, M. 2003. What's Strange About Recent Events, *Journal of Urban Health*, 80: 66-75.