

A KNOWLEDGE-BASED APPROACH TO TEMPORAL ABSTRACTION OF CLINICAL DATA FOR DISEASE SURVEILLANCE

Mark A. Musen, M.D.,¹ Ph.D., Martin J. O'Connor, M.S.,¹
David L. Buckeridge, M.D., M.P.H.,¹ Justin Graham, M.D.,¹
Natalya Fridman Noy, Ph.D.,¹ Yuval Shahar, M.D., Ph.D.,¹ and Kurt A. Henry, M.D.²

¹Stanford Medical Informatics
251 Campus Drive, Room X-215
Stanford University School of Medicine
Stanford, California 94305-5479
USA

²Defense Sciences Office
Defense Advanced Research Projects Agency
3701 North Fairfax Drive
Arlington, Virginia 22203-1714
USA

Summary

Automated surveillance for disease detection requires the integration of many kinds of data. The clinical information systems in hospitals, clinics, and emergency rooms that record patient findings and conditions are a fruitful source of information that may suggest the presence of incipient epidemics in a timely manner. Analysis of data from electronic patient record systems requires methods that can integrate both numeric and nonnumeric information, and that can use the clinical context to aid in data interpretation. Techniques from artificial intelligence that can take advantage of clinical knowledge that has been encoded in machine-processable form are particularly well suited for disease surveillance. We describe the knowledge-based temporal abstraction (KBTA) method, which we believe will be particularly useful in performing automated surveillance of clinical data to detect occult acts of bioterrorism.

1. Introduction

Disease surveillance—including surveillance for nascent epidemics that could reflect occult acts of bioterrorism—requires the continuous analysis, interpretation, and feedback of systematically collected data. Surveillance can support many activities such as planning and research, but the most important reason for conducting surveillance is to identify changes in population health status that are amenable to control by intervention. The changes, or aberrations, must be detected from data sources that often have a highly variable baseline. Yet, because of the urgency of detecting incipient epidemics, methods for disease surveillance that are distinguished by their practicality, uniformity, and rapidity are preferred to those that may be most accurate and most complete.

Methods for disease surveillance generally have relied on traditional statistical models (see Stroup, 1994). Such approaches typically take as input disease reports from passive surveillance and generate as output notification of diseases or clinical conditions that may occur above certain thresholds within given geographic areas. Passive surveillance requires health-care personnel to be aware that a clinical situation is “reportable” and to initiate a report to the relevant department of health, and for that department of health to collate and analyze those reports as they come in. The reliability of passive surveillance systems is quite low, since many health-care workers may not even know which conditions are reportable, and may not see an immediate benefit to informing public-health officials of particular diseases or syndromes. Most important, however, passive surveillance is extraordinarily slow. Infected patients typically do not present for

treatment until they significantly ill, and reports of sentinel cases typically do not reach local authorities with any great urgency. By the time passive surveillance systems can detect incipient epidemics, many people will already have been infected and secondary spread of the contagion may already be underway.

The limitations of current public health methods are well recognized, particularly in light of increasing concern about the possibility of epidemics that may result from acts of biological warfare or bioterrorism. There is thus considerable activity to develop active surveillance systems that may be able to identify incipient epidemics rapidly from primary data available in electronic format in a manner that is not dependent on the acumen of health-care workers to recognize reportable conditions and on their good will to file such reports. Such work requires creative thought about sources of useful and timely data that may not be the traditional stuff of public-health decision making.

Retrospective analysis of natural disease outbreaks can identify important performance characteristics of potential data sources for detection of bioterrorism. For example, review of data sources related to the large outbreak of waterborne cryptosporidiosis in Milwaukee, Wisconsin, in 1993 showed that emergency-room visits for gastrointestinal symptoms peaked 5 days after the start of the epidemic; school absenteeism peaked at 9 days; laboratory identification of the pathogen, however, did not peak until 15 days after onset of the outbreak (Proctor et al., 1998). In the case of the Milwaukee epidemic, the challenge would have been to identify the increase in emergency-room visits for (nonreportable) gastrointestinal symptoms as quickly as possible.

Increasingly, health-care institutions store a wide range of patient information in electronic medical record systems. Such data typically include the results of laboratory tests (including the results of microbiological cultures), the results of other diagnostic studies, prescriptions and other clinician orders, clinical notes (generally in free text), and codes for diagnoses and procedures. Data available from clinical information systems have been suggested as a rich source of information for disease surveillance. The goal is to identify patterns in the complaints with which patients present to emergency departments, in the conditions of patients admitted to hospitals, and in the prescriptions written for both inpatients and outpatients that could suggest emerging epidemics in the general population as manifested in the subset of patients presenting themselves to health-care organizations. Rather than relying on humans to identify reportable situations, public-health authorities could monitor institutional databases continuously to identify the presence of public-health problems requiring immediate action or further investigation.

There are significant difficulties with this approach, however. Apart from the obvious, surmountable problems of ensuring patient confidentiality, there is the need to translate numerous low-level laboratory values into meaningful abstractions that can drive epidemiological decision making. Public health officials need to be alerted to patterns of patients presenting to emergency departments with fever, not to the particular temperature measurements of specific patients. Furthermore, there is a need to identify complex patterns of findings (e.g., fever plus diarrhea) that may require integration of abstractions of detailed observations stored in the clinical record with additional qualitative patient attributes recorded as diagnosis codes or inferred from narrative text.

Even the simple data stored in electronic patient record systems rarely are in a form that is suitable for direct analysis by traditional statistical approaches. The results of individual microbial cultures, for example, typically need to be interpreted in the context of other cultures that have been taken from the same patient. Primary laboratory data, such as white-blood-cell

counts and serum enzyme concentrations, need to be understood in terms of relevant abstractions (e.g., “severe, worsening leukocytosis” and “sustained moderately elevated liver-function tests”) that occur over explicit temporal intervals. Standard statistical methods do not lend themselves to the generation of clinically meaningful temporal abstractions from the myriad point data available in electronic patient records—particularly when the clinical situation can be inferred only from patterns that occur in the data over time within specific clinical contexts.

In this paper, we describe a promising automated method for analysis of electronic patient-record data that uses medical knowledge to infer high-level patterns from primary data. The explicit encoding of knowledge for use by the computer allows for identification of associations among the data and of temporal trends that cannot be detected by standard statistical approaches. We describe our plans to apply our approach, known as *knowledge-based temporal abstraction*, to actual clinical data in an attempt to monitor for aberrant disease conditions.

2. Knowledge-Based Temporal Abstraction

Knowledge-based temporal abstraction (KBTA) is a problem-solving method based on artificial-intelligence techniques that has been developed and tested in our laboratory during the past decade (Shahar and Musen, 1993; Shahar, 1997). KBTA infers context-specific, temporally bound abstractions of point data. The method uses knowledge about the kind of primary data available as input (and the potential abstractions of those data) to inform a problem solver that takes as input a database of primary data and that generates as output all the relevant abstractions of those data and the time intervals during which those abstractions hold. Thus, KBTA can take as input specific body-temperature measurements and conclude, for example, that a patient has had an episode of increasing, high fever during the past 6 hours. Knowledge about the relationship between body temperature and fever, about what constitutes a *high* fever, and about the criteria for concluding that a fever might be *increasing*, is stored in a machine-processable knowledge base. KBTA uses knowledge of the clinical context to adjust its conclusions; for example, the method might use a different definition of what constitutes a “fever” for patients who are taking antipyretic drugs, such as aspirin.

KBTA can use the abstractions that it generates to create yet other abstractions. Thus, body-temperature measurements can be abstracted into points in time during which a patient has had *fever*. KBTA then can abstract the various points in time associated with fever into a *time interval* during which fever has occurred. A time interval of fever might be abstracted into an interval of *high, increasing fever*. An interval of fever during which there is concomitant chills or sweats might be abstracted into an interval of *constitutional symptoms*. The KBTA method thus has the ability to examine the myriad time-oriented point data that might be stored in an electronic patient record and to construct automatically abstractions of those data—qualitative abstractions that might suggest the presence of syndromes or diseases that might be of epidemiological significance.

Unlike standard statistical techniques such as time-series analysis (Quenel and Dab, 1998) or Kalman filtering (Stroup and Thacker, 1993), KBTA can easily mix a wide range of quantitative and qualitative data when performing the temporal-abstraction process. KBTA uses clinical knowledge encoded for the computer to recognize that concomitant fever and sweats (both qualitative abstractions) can combine to form another qualitative abstraction called *constitutional signs*. Other clinical knowledge informs the method when numeric values should be considered *low* or *high*. Still other knowledge elements allow the method to incorporate explicitly the clinical context when interpreting data stored in the electronic patient record.

There are three primary classes of abstractions that KBTA can make. These abstractions refer to the state of parameters (e.g., high or low), the gradient of change of parameters (i.e., increasing or decreasing), and the rate of change of parameters (e.g., fast or slow).

A *state abstraction* takes as input one or more values and generates as output the value of the corresponding condition. A state abstraction allows hematocrit measurements to be abstracted into the presence of low hematocrit, and the presence of low hematocrit in turn to be abstracted into the presence of *anemia*. State abstraction allows white-blood-cell counts to be abstracted into the presence of *low white-blood-cell counts*, and the presence of low white-blood-cell counts to be abstracted into the presence of *leukopenia*. In turn, state abstraction allows the concurrent presence of anemia, leukopenia, and thrombocytopenia to be abstracted into episodes of *pancytopenia*. Performing each of these state abstractions requires the KBTA method to have access to clinical knowledge that defines the relationships between the different descriptors (e.g., that low hematocrit is called anemia) as well as expectations for data value in different contexts (e.g., what constitutes a “low” hematocrit under ordinary circumstances; what constitutes a “low” hematocrit in the setting of chronic renal failure).

A gradient abstraction defines an interval during which the value of a parameter is changing. KBTA may therefore detect *decreasing hematocrit* or *constant leukopenia*. A rate abstraction would detect events such as *rapidly decreasing hematocrit* or *slowly rising fever*. Generating these kinds of abstractions requires domain-specific knowledge as well.

KBTA, like many current knowledge-based approaches (Musen, 2001), assumes that the method will have access to explicit knowledge representations, called *ontologies*, that define the concepts and the relationships among those concepts that are relevant to the domain of interest. For example, KBTA requires a parameter-properties ontology that defines the data available for processing (e.g., what data can be input from electronic patient records); the state, gradient, and rate abstractions of those data (i.e., what are the abstractions that the system should make and what are the relationships among them); and the criteria for making those abstractions (e.g., for each parameter, what makes a value “low”?; what makes a gradient “rapid”?). Other ontologies define the *events* that might happen to patients and the *contexts* that may influence the way in which patient data should be interpreted. It is this rich domain knowledge that gives the KBTA method its particular power. Experience in our laboratory demonstrates that it is feasible for clinical experts working alone to define and maintain the knowledge bases required by KBTA (Shahar et al., 1999).

The KBTA method uses several different problem-solving components that work together to derive abstractions from the primary data. The components are called problem-solving *mechanisms*. The mechanism known as *vertical temporal inference* takes input data such as hemoglobin values at different time points and generates as output abstractions such as the presence of anemia at the same time points. Another mechanism, called *temporal interpolation*, may join several time points of anemia into one long interval of anemia. A third mechanism, called *horizontal temporal inference*, may join several consecutive intervals of anemia into one long interval. The different problem-solving mechanisms work opportunistically on a data set to generate the complete set of abstractions that may be inferred from the data, given the available domain knowledge. The mechanisms continue to operate, taking as input all the previously generated abstractions and applying domain knowledge that provides the basis for making further abstractions of the input data. Because the mechanisms can run in parallel and may operate on different input data and abstractions, it is possible to execute KBTA on a supercomputer cluster that facilitates rapid, parallel computation (O’Connor et al., 2001).

The KBTA problem-solving mechanisms make use of additional domain knowledge as they perform the task of generating appropriate abstractions from the input data. For example, the temporal-interpolation mechanism requires probabilistic knowledge of the decay properties of data values to know when it is appropriate to join two different intervals. Thus, it may be appropriate to join two different intervals of anemia if the gap between them is less than a few days; two intervals of anemia separated by a month's time probably cannot be assumed to be representative of one long interval of anemia, however. Other semantic properties of the data and their abstractions are important when performing horizontal temporal inference: Two consecutive intervals of anemia easily can be combined into one longer interval of anemia; two consecutive intervals of pregnancy, however, generally cannot be combined into one longer interval of pregnancy.

In our laboratory, we have developed considerable experience in the application of KBTA to solve clinical problems. We have used the method to identify patterns in the blood sugar of patients who have diabetes (Shahar and Musen, 1996). We have used the method to detect aberrancies in the growth of children (Kuilboer et al., 1993). KBTA has been particularly helpful in the automation of protocol-based medical therapy (Musen et al., 1996), in which it often is necessary to predicate clinical actions based on the presence of complex patterns in the patient data that cannot be detected without the use of considerable domain knowledge (e.g., "Do not give the drug zidovudine if the patient experience a second episode of myelotoxicity of Grade III or more within six months). Because disease surveillance may require these kinds of complex abstractions of time-ordered data—abstractions that are impossible to make using standard numerical processing—there is good evidence that KBTA may be particularly germane to the task of automated disease surveillance to detect occult bioterrorism. We currently are performing research test the utility of KBTA for this purpose.

When performing knowledge-based temporal abstraction for disease surveillance, it is necessary not only to perform abstraction of individual patient data, but also to perform abstractions across different patients. We therefore are developing new versions of the temporal-abstraction mechanisms that can generate summaries that hold for *groups* of patients whose data may be stored in an electronic record. The ultimate goal is to develop a surveillance system that can monitor a wide range of patient-related data sources, such that the application of computer-stored domain knowledge can drive the automated identification of aberrant patterns that could signal events requiring further investigation.

3. Discussion

Standard approaches to disease surveillance typically use passive methods that are not well suited to rapid detection of changes in disease patterns. Moreover, even when it is possible to monitor primary data sources in real time, traditional statistical techniques do not allow epidemiologists to evaluate rich data sources, such as electronic patient records, where detailed clinical knowledge is needed to determine how the data should be interpreted and viewed abstractly over time. The power of artificial-intelligence approaches, such as KBTA, is that the necessary clinical knowledge can be encoded directly in the computer and brought to bear in a principled way to detect automatically a wide range of high-level patterns. Although statistical techniques can "let the data speak for themselves," knowledge-based techniques have the unique ability to use qualitative data, contextual information, and explicit relationships among the data elements to make inferences about the data that simply cannot be performed when using standard approaches. Particularly when performing surveillance for bioterrorism, the low signal-to-noise ratio in most

available data sources requires the ability to use contextual information and the presence of concomitant conditions to adjust the thresholds used for identifying abnormal patterns. At the same time, the high covariance among many data elements would favor the use of automated monitoring systems that can use domain knowledge to form appropriate abstractions that avoid over-counting of interdependent data streams.

Data such as those related to the Milwaukee cryptosporidium epidemic suggest that, for some epidemics, there will be signals detectable from emergency-room visits and possible hospital admissions several days before clinical laboratories begin to report positive cultures (Proctor et al., 1988). The ability to use clinical data from emergency rooms and hospitals in an effective manner requires the ability to generate useful temporal abstractions of those data across patient groups. In the United States, the Defense Advanced Research Projects Agency (DARPA) has begun a substantial research program to identify how nontraditional data sources—many of them both sparse and qualitative in nature—can contribute to a comprehensive program of disease surveillance for bioterrorism (Burke, 2001). The vision is that data as diverse as records of school absenteeism, reports of animal deaths in zoos, and indicators of over-the-counter drug sales can contribute to an overall autonomous disease surveillance program. Because these data sources are extremely heterogeneous in nature and often nonnumeric in their form, they are excellent candidates for processing via knowledge-based techniques, such as KBTA.

Developers might consider the application of knowledge-based techniques other than KBTA to the interpretation of heterogeneous data streams. Bayesian networks, for example, provide a well understood modeling methodology and a principled inference mechanism for integrating the ways in which diverse inputs may affect a set of enumerable output states. Indeed, there has been promising work on the application of Bayesian networks to domains where data are time-varying in nature (Dagum and Galper, 1995). Traditional Bayesian systems, however, do not make explicit the ontology of input data and output conditions that in KBTA clarifies the data being monitored, all the abstractions of those data, and the relationships among the data and their abstractions. The ability to represent these ontologies declaratively is a considerable aid to the development and maintenance of electronic knowledge bases used by the KBTA method (Shahar et al., 1999). Although KBTA uses probabilistic reasoning to perform temporal interpolation (Shahar, 1999), most of the inferences required for the temporal-abstraction task actually are categorical in nature, and stem directly from the ontology of abstractions created at the time the knowledge base is designed.

The design of the knowledge bases required by KBTA is nontrivial, however. Developers need to be explicit about the relationships among all the data and abstractions of those data, and need to define distinct criteria for making each possible abstraction. Managing this complexity in the knowledge base requires good tool support. We have come to rely on the Protégé-2000 knowledge-acquisition system for building all our intelligent systems at Stanford University (Musen et al., 2000). Our experience suggests that the knowledge-base visualization capabilities of the Protégé system, coupled with its relatively simple user interface, make the tasks of knowledge-base development and maintenance practical, even for domain experts with minimal computer experience (Shahar et al., 1999).

Monitoring for occult acts of bioterrorism requires more than reasoning about diverse data sources over time. It is necessary to reason about these data the spatial dimension as well. We already have experience extending the inference mechanisms in KBTA to reason about spatial events; Shahar and Molina (1996) readily adapted KBTA to identify patterns in vehicular traffic detected by sensors in the Madrid freeway system. This work involved reasoning to detect linear

spatial relationships. Future work will require extensions to KBTA that can detect patterns in two or three spatial dimensions.

Knowledge-based techniques offer considerable promise for monitoring multidimensional, heterogeneous data sources. KBTA in particular can perform considerable data reduction while leading to identification of salient events and intervals that take place over time within a given patient population.

Acknowledgments

This work has been supported in part by a grant from the United States Space and Naval Warfare Systems Command supported by the Defense Advanced Research Projects Agency.

References

- Burke, M. (2001). Bio-surveillance briefing. <http://www.darpa.mil/iso/BIOS/index.htm>. Accessed March 16, 2001.
- Dagum, P. and Galper, A. (1995). Time series prediction using belief network models. *International Journal of Human-Computer Studies*. **42**(6):573-712.
- Kuilboer, M.M., Shahar, Y., Wilson, D.M., and Musen, M.A. (1993). Knowledge reuse: Temporal-abstraction mechanisms for the assessment of children's growth. *Proceedings of the Sixteenth Annual Symposium on Computer Applications in Medical Care*, Washington, DC, November, pp. 449-453.
- Musen, M.A. (2001). Ontology-oriented design and programming. In: Cuenca, J., Demazeau, Y., Garcia, A., and Treur, J., eds. *Knowledge Engineering and Agent Technology*. Amsterdam: IOS Press.
- Musen, M.A., Ferguson, R.W., Grosso, W.E., Noy, N.F., Crubézy, M., and Gennari, J.H. (2000). Component-based support for building knowledge-acquisition systems. *Proceedings of the Conference on Intelligent Information Processing (IIP 2000) of the International Federation for Information Processing Sixteenth World Computer Congress (WCC 2000)*, Beijing, China, August, pp. 18-22. See also <http://protege.stanford.edu>.
- Musen, M.A., Tu, S.W., Das, A.K., and Shahar, Y. (1996). EON: A component-based approach to automation of protocol-directed therapy. *Journal of the American Medical Informatics Association*, **3**:367-388.
- O'Connor, M.J., Grosso, W.E., Tu, S.W., and Musen, M.A. (2001). RASTA: A distributed temporal abstraction system to facilitate knowledge-driven monitoring of clinical databases. *Proceedings of MedInfo 2001*. The Tenth World Congress on Medical Informatics, London, September.
- Proctor, M. E., K. A. Blair, et al. (1998). Surveillance data for waterborne illness detection: an assessment following a massive waterborne outbreak of *Cryptosporidium* infection. *Epidemiol Infect* **120**(1):43-54.

- Quenel, P. and Dab, W. (1998). Influenza A and B epidemic criteria based on time-series analysis of health services surveillance data. *European Journal of Epidemiology* **14**(3): 275–85.
- Shahar, Y. (1997). A framework for knowledge-based temporal abstraction. *Artificial Intelligence* **90**(1–2):79–133.
- Shahar, Y. (1999). Knowledge-based temporal interpolation. *Journal of Experimental and Theoretical Artificial Intelligence*. **11**:123–144, 1999.
- Shahar, Y. and Molina, M. (1996). Knowledge-based spatiotemporal abstraction. Proceedings of the AAAI-96 Workshop on Spatial and Temporal Reasoning, Portland, OR, July, pp. 21–30.
- Shahar, Y. and Musen, M.A. (1993). RÉSUMÉ: A temporal-abstraction system for patient monitoring. *Computers and Biomedical Research*. **26**:255–273.
- Shahar, Y. and Musen, M.A. (1996). Knowledge-based temporal abstraction in clinical domains. *Artificial Intelligence in Medicine*. **8**:267–298.
- Shahar, Y., Chen, H., Stites, D.P., Basso, L., Kaizer, H.J., Wilson, D.M., AND Musen, M.A. (1999). Semi-automated Entry of Clinical Temporal-abstraction knowledge. *Journal of the American Medical Informatics Association* **6**(6):494–511.
- Stroup, D.F. (1994). Special analytic issues. In: Teutsch, S.M. and Churchill, R.E., editors. *Principles and Practice of Public Health Surveillance*. New York: Oxford University Press. pp. 136–149.
- Stroup, D.F. and Thacker, S.B. (1993). A Bayesian approach to the detection of aberrations in public health surveillance data. *Epidemiology* **4**(5):435–43.