Data-driven identification of unusual clinical actions in the ICU

Milos Hauskrecht[1], Shyam Visweswaran[2], Gregory F. Cooper[2], Gilles Clermont[3]

[1]Department of Computer Science, [2]Department of Biomedical Informatics, [3]CRISMA Center, Department of Critical Care Medicine, University of Pittsburgh, PA, USA

Abstract

Developing methods to identify unusual clinical actions may be useful in the development of automated clinical alerting systems. We developed and evaluated a data-driven approach for identifying clinical actions such as omissions of medication orders or laboratory orders in the intensive care unit (ICU) that are unusual with respect to past patient care. We generated 250 medication-omission alerts and 150 laboratory-omission alerts using a database of 24,658 ICU patient admissions. These alerts were evaluated by a group of intensive care physicians. Overall, the true positive alert rate was 0.52, which we view as quite promising.

Introduction

Medical errors remain a significant problem in healthcare. We developed an outlier-based monitoring and alerting method that relies on past patient data and on statistical methods for the identification of unusual clinical actions [1,2,3]. We evaluated our approach on EMR data for ICU patients. Our conjecture is that the detection of unusual clinical actions (e.g., medication orders and lab orders) will help identify medical errors. We believe outlier-based monitoring and alerting can complement the use of knowledge-based alerting systems that are currently deployed, thereby improving overall clinical coverage of current alerting systems.

Methods

We used electronic medical record (EMR) data from 24,658 ICU patient stays that included time-stamped data on medications, laboratory and physiological measurements. Each patient stay record was segmented into several time-series of increasing lengths of time (in 24 hour increments), and each time-series was summarized by a vector of over 12,500 temporal features [3,4]; this vector represented a patient instance. The data was split into a training set of 225,894 patient instances and a test set of 104,698 patient instances. We built a Support Vector Machine (SVM) model for predicting each type of action (medication orders and laboratory orders) from the training set and applied the models to the test set for identifying omissions of medication orders and laboratory orders.

We developed an alert score [2,3] computed from the SVM model predictions which measures how unusual the actions for the patients are. The alert score allows rank-ordering of alerts that are generated for each action and also allows the control of the rate at which alerts are raised for each clinical action by selecting an alert score threshold. We control the alert threshold as follows. Let A be a specific clinical action (e.g., aspirin order) and PosRate(A) be the average rate at which that clinical action is done (e.g., rate at which aspirin is prescribed in the ICU population). We limit the clinical action alert rate to AlertRate(A) = α x PosRate(A) where PosRate(A) is estimated from the training set. With this method, the alert rate can be increased or decreased by changing the multiplier α.

Results

We built SVM models for 1075 different medication orders and 222 types of laboratory orders. Out of all the alert models, 99 laboratory and 156 medication omission models were strong alert models based on $\geq$ 0.75 AUC ROC and $\geq$ 0.3 positive predictive value thresholds. From the strong models we randomly selected 85 medications and 45 labs models and used them to generate 400 alerts that included 250 medication-omission alerts and 150 laboratory-omission alerts. The alerts were evaluated by 16 physicians from the Departments of Critical Care Medicine and Surgery. Based on the evaluations, the true positive alert rate (TPAR) using α=0.05 is 0.52, which we view as quite promising. The TPAR for medication-order alerts and the same threshold was 0.44 and the TPAR for laboratory-order alerts was 0.66. Given α=0.05 and the strong alert models identified in the study, we estimate the outlier-based alerting system (if deployed) would raise in between 0.55 to 0.95 alerts per patient per day.

Acknowledgements: This research was supported by NIH grants R01-LM010019, and R01-GM088224.

References

[1] M. Hauskrecht, M. Valko, B. Kveton, S. Visweswaram, G. Cooper. Evidence-based anomaly detection in clinical domains. AMIA Fall Symposium; 2007. p. 319-23.

[2] Hauskrecht M, Valko M, Batal I, Clermont G, Visweswaran S, Cooper GF. Conditional outlier detection for clinical alerting. AMIA Fall Symposium; 2010. p. 286–90.

[3] Hauskrecht M, Batal I, Valko M, Visweswaran S, Cooper GF, Clermont G. Outlier detection for patient monitoring and alerting. J Biomed Inform. 2013 Feb; 46(1):47-55.

[4] Valko M, Hauskrecht M. Feature importance analysis for patient management decisions. Medinfo 2010, p 861-5.