Extended Abstract: How to Interpret Organizational Relations Mined from the Access Logs of Health Information Systems

You Chen[1], Nancy Lorenzi[1,2], Steve Nyemba[1], Jonathan S. Schildcrout[3,4], and Bradley Malin[1,5]

[1]Dept. of Biomedical Informatics, School of Medicine, Vanderbilt University, Nashville, TN
[2]School of Nursing, Vanderbilt University, Nashville, TN
[3]Dept. of Biostatistics, School of Medicine, Vanderbilt University, Nashville, TN
[4]Dept. of Anesthesiology, School of Medicine, Vanderbilt University, Nashville, TN
[5]Dept. of Electrical Engineering and Computer Science, Vanderbilt University, Nashville, TN

The healthcare community has made considerable strides in the development of information technology to support clinical operations in healthcare organizations (HCOs). These advances stem from a variety of factors, including commercialization of health information technology (HIT) and policy making that promotes the uptake of such technologies (e.g., the "meaningful use" incentives offered in the United States). While there is evidence that HIT can improve the safety and the efficiency of healthcare delivery, there remain considerable obstacles to adoption and realization of these benefits on a massive scale.

In particular, as HIT, and the healthcare workforce more generally, grows in diversity, so too do its complexity. This is a concern because, despite the aforementioned benefits, there is also evidence to suggest that HIT can contribute to (though is not necessarily the cause of) the interruption of care services, induce medical errors, and expose patient data to privacy breaches. Moreover, such events tend to be discovered only after they have transpired *en masse*, leading to undesirable popular media coverage, loss of patients' trust, and sanctions imposed by state and federal agencies.

It has been suggested that such problems can be mitigated through the integration of rules to recommend against or even prohibit certain actions (e.g., the prescription of two drugs in combination that are known to cause an adverse reaction). At the same time, it is recognized that no rules-based system is perfect and that exceptions need to be granted. These exceptions can, in turn, be audited to determine if the existing set of rules are in alignment with the expectations of the HCO or if they need to be revised to more accurately represent healthcare operations. [1,2] For instance, it has been shown that the exposure of patient records (and thus the violation of their privacy) can be lessened through access control [3,4], which allocates permission to patient information on a need-to-know basis. Yet, access control configurations are often rigidly defined and may be refined by evolving it to reflect how people actually utilize the system. [5-7]

Data-driven approaches to HIT improvement will only be acceptable to HCO administrators if the patterns of HIT utilization reflect the expected operations of healthcare environments. Our recent work begins to address this issue by investigating how a specific type of HIT utilization pattern, which has been suggested for use in audit and refinement of access control models [8], aligns with the expectations of employees in a large academic medical center. To do so, we designed a survey to capture the degree to which employees agree with relational patterns (i.e., the likelihood that certain HCO areas in a medical center coordinate to support a patient) as inferred by actual utilization of an electronic health record (EHR) system. [9] This survey was conducted with employees from four areas in the Vanderbilt University Medical Center VUMC. It was designed to determine if employees 1) agreed with the distinction between relationships of high and non-high likelihood and 2) were better at assessing relationships regarding their own area as opposed to others in the institution.

For our study, we selected two clinical areas and two operational areas with direct responsibilities that involve the electronic record, who were relatively evenly matched on the number of users and patient records they accessed over a single week. For the clinical areas, we selected *Anesthesiology* (~250 users' accessed ~2900 patient records) and *Psychiatry* (~160 user's accessed ~1900 patient records). For the operational areas, we selected *Medical Information Services* (~100 user's accessed ~7500 patient records) and *Coding & Charge Entry* (~75 user's accessed ~6000 patient records). A leader from each area was contacted and all agreed to participate and to identify 10 users as potential participants of the study.

Our first hypothesis is that employees in a certain HCO area are capable of distinguishing between HCO interactions of high and non-high likelihood when their own HCO area is involved in the interaction. To assess this hypothesis, we model each HCO area X as set of association rules of the form $Area\ X \Rightarrow Area\ Y_i$, where Y_i corresponds to any area in the HCO. This area may be the same, such that $X = Y_i$ or different, such that $X \neq Y_i$. The rule corresponds to the conditional probability that a user from area Y_i accessed a patient's record given that a user from area X accessed the patient's record, as defined in an earlier study.

To prevent fatigue in evaluation, for each area so we sampled 30 rules, stratified across groups of high, moderate, and low likelihood. Without observing the likelihood, employees were asked to score the expectation that a rule represents the hospital system on a 5-point Likert scale. We adopted a linear-mixed effects model to test our hypotheses since observations made within respondents are likely to exhibit correlation with one another. Statistical significance was assessed using (non-restricted) maximum likelihood ratio tests at the two-sided α=0.05 significance level.

Our preliminary findings confirmed that each group of respondents from a specific organizational area was able to distinguish between the high and non-high likelihood rules for their own organizational area. Moreover, respondents could distinguish between high and non-high likelihood rules regardless of the organizational area on which the rules are conditioned. However, the evidence further suggested respondents from a specific area were no better at distinguishing between classes of rules associated with their own HCO area than rules associated with other HCO areas.

The findings of this study are notable because it suggests that the application of automated learning strategies to EHR utilization patterns may be useful for modeling the workings of a large healthcare system. These models may allow for assessments of organizational efficiency, workflow, and provide insight into how to refine such models into more effective organizational structures. These findings indicate that EHR auditing models, such as those leveraged to redefine access control systems may be trustworthy and reflect actual organizational behavior.

References

1. Isaac T, Weissman JS, Davis RB, Massagli M, Cyrulik A, Sands DZ, Weingart SN. Overrides of medication alerts in ambulatory care. Archives of Internal Medicine. 2009; 169: 305-311.

2. Yeh ML, Chang YJ, Wang PY, Li YC, Hsu CY. Physicians' responses to computerized drug-drug interaction alerts for outpatients. Computers Methods and Programs in Biomedicine. 2013; 111: 17-25.

3. Anderson R. Clinical system security: interim guidelines. British Medical Journal. 1996; 312: 109-11.

4. Blobel B. Authorisation and access control for electronic health record systems. International Journal of Medical Informatics. 2004; 31: 251-257.

5. Bhatti R, Grandison T. Towards improved privacy policy coverage in healthcare using policy refinement. Proceedings of the VLDB Workshop on Secure Data Management. 2007: 158-73.

6. Gunter C, Liebovitz D, Malin B. Experience-based access management: a life-cycle framework for identity and access management systems. IEEE Security and Privacy Magazine. 2011; 9: 48-55.

7. Zhang W, Chen Y, Gunter C, Liebovitz D, Malin B. Evolving role definitions through permission invocation patterns. Proceedings of the ACM Symposium on Access Control Models and Technologies. 2013: 59-70.

8. Chen Y, Nyemba S, Malin B. Detecting anomalous insiders in collaborative information systems. IEEE Transactions on Dependable and Secure Computing. 2012; 9(3): 332-344.

9. Malin B, Nyemba S, Paulett J. Learning relational policies from electronic health record access logs. Journal of Biomedical Informatics. 2011; 44: 333-42.