Avoiding Alert Fatigue

The term "alert fatigue" describes how clinicians in busy hospital settings can become desensitized to safety alerts, and as a result ignore or fail to respond appropriately to such warnings. This phenomenon occurs because of the sheer number of clinically inconsequential alerts, many emanating from physiological monitors or decision support rule engines processing medical record data in real time. The resultant problem is that "real" alerts that warn of impending serious risk to patients get lost in this noise of useless alarms. In essence, alerting tools that are intended to improve patient safety actually result in a paradoxical increase in the chance patients will be harmed due to alert fatigue… now recognized as a major unintended consequence of electronic medical records and digital monitoring systems.

Recent studies of alerting tools (e.g. AHRQ) consistently show that alert fatigue is common, increases with exposure to meaningless alerts and are consequentially only modestly clinically effective at best. Although there is intense interest in developing specific methods to combat alert fatigue, as yet there is no consensus on an optimal approach. It is generally recognized that the principles of human factors "usability" workflow engineering as well as those of precision informatics will be essential elements of future solutions. One informatics approach being used is the use of algorithms that increase alert specificity by reducing or eliminating clinically inconsequential alerts. Although potentially useful, such "high specificity" approaches have limitations.

For example, many hospitals are trying to harness their EHRs to flag patients becoming septic, a continuing battle with this deadly disease that is more common than heart attacks and claims more lives than cancer. A great recent example is the Yale-New Haven's EHR-based protocol that is designed to trigger an alert when a medical inpatient registers at least four of six sepsis criteria, extending the traditional 2-"SIRS" criteria proven to be of low specificity. In contrast the new 4-criteria protocol is highly specific (97%), but unfortunately operates at a sensitivity of 16%.

A published study on the use of this protocol in a pilot study concluded that in actual use clinicians believed that most flagged patients based on the 4-criteria alerting trigger were already "stable" before and after the alert. In effect the algorithm only flagged the most critically ill already recognized as septic, leading to a "high specificity/low sensitivity" type of alert fatigue. In this study only 30% of users deemed the alerting tool useful, and most importantly, no significant change in patient outcomes (e.g. lower mortality) was achieved. We believe there needs to be a significant improvement in usability and a combined specificity and sensitivity for such tools to make a difference in clinician perception and patient outcomes.

Our view is that such improvement is possible by alerting tools that leverage the combined analytic power of computerized semantic models (rules, natural language processing) and machine learning, and employ human factors/usability engineering to achieve highly synchronicity with clinical workflows. Late last year we were awarded a NIH research grant focused on the use of advanced ontological models combined with "big data" machine learning (ML) techniques as a foundation for a Clinical Decision Support (CDS) system. Our research goal is an exceptionally sensitive and specific tool that monitors hospitalized patient electronic medical record (EMR) data to accurately predict patients at risk for impending acute clinical deterioration due to sepsis. Our goal is also to deploy highly useable technology that achieves high levels of clinician acceptance and demonstrably influences timely treatment and patient outcomes.

Recent studies indicate currently available early warning tools do not reduce risk of sepsis death in hospitalized patients. We believe this may be because diseases such as sepsis are time-sensitive, complex syndromes and also due to the challenges of computerized reuse of unstructured EMR data. Our sepsis ontology mirrors this complexity with a model consisting of over 300,000 concepts and 2M axioms. Using de-identified ICU data, our initial study results demonstrate a “native” SCT classification sensitivity/specificity detection performance of .95/.84 and, when used in conjunction a machine learning algorithm over a large sample of ICU patients (N = 15,811), achieves an AUC in excess of .98. Work to be completed includes the human factors component that can only be achieved via actual in-use pilot-type studies.

VFusion™ is an early sepsis detection CDS with actionable accuracy and usability, effective in reducing sepsis mortality in both critical and non-critical care settings.