DEVELOPMENT OF A BUSINESS INTELLIGENCE TOOL FOR THE EXTRACTION, ANALYSIS, DISPLAY, AND UTILIZATION OF EARLY WARNING SYSTEM ALERTS

By

James M Bailey M.D.

A CAPSTONE PROJECT

Presented to the Department of Medical Informatics and Clinical Epidemiology

and the Oregon Health & Science University School of Medicine

In partial fulfillment of the requirements for the degree of

Master of Biomedical Informatics

December, 2012

School of Medicine

Oregon Health & Science University

CERTIFICATE OF APPROVAL

This is to certify that the Master's Capstone Project of

James M Bailey

"DEVELOPMENT OF A BUSINESS INTELLIGENCE TOOL FOR THE EXTRACTION, ANALYSIS, DISPLAY, AND UTILIZATION OF EARLY WARNING SYSTEM ALERTS"

Has been approved

Capstone Advisor

ACKNOWLEDGEMENTS

I would like to thank Dr. Vishnu Mohan, Department of Department of Medical Informatics and Clinical Epidemiology, Oregon Health and Science University, for his wise and patient guidance in the conception and completion of this capstone project. I would also like to thank Professor Wassim Haddad, Department of Aerospace Engineering, Georgia Institute of Technology, for his expert review of and attestation to the mathematical exposition in the theoretical appendix to this manuscript. I appreciate the support for my foray into medical informatics that I have consistently received from Ms. Allana Cummings, CIO, Northeast Georgia Health System. This project would not have been possible without the expertise and patient tutelage of Mr. Dalton Ruer, Information and Technology Services, Northeast Georgia Health System. He imparted a small portion of his mastery of QlikView to me and I am grateful. Finally, I express my heartfelt gratitude to my fabulous wife Rachel, who saw me through this Masters program, including tolerating my obsessing over a Java program one Thanksgiving morning while she was preparing dinner for more than 20 people!

INTRODUCTION:

Interest in early warning systems or scores originated in the observation that catastrophic deterioration of patients in hospitals is frequently preceded by documented abnormalities of vital signs and other physiological parameters that were sometimes ignored by clinical staff (1-5).

One solution to the problem of unrecognized patient deterioration on the floor was the medical emergency team (MET) concept, introduced in 1990 at Liverpool Hospital, New South Wales, Australia (6).

The medical emergency or rapid response team is a multi-disciplinary team of intensive care unit nurses, respiratory therapists, and sometimes physicians who respond whenever floor nurses have a concern that their patient is deteriorating. While the basis for activation of the MET is quite broad, the original implementation used defined physiological parameters as calling criteria. (Also please note that pediatric patients will not be considered in this discussion).These included

Temperature	<35.3 degrees C	> 39.5 degrees C
Systolic blood pressure	< 100 mm Hg	> 200 mg Hg
Respiratory rate	< 10 bpm	> 30 bpm
Urine output (24 hrs)	< 500 ml	
Level of consciousness	Altered	

Most institutions that implemented a rapid response team have added heart rate and hemoglobin oxygen saturation (SpO2) as criteria. Currently the criteria for activation of the MET at Northeast Georgia Medical Center are

Heart Rate	< 40 beats per minute	> 130 beats per minute
Systolic Blood Pressure	< 90 mm Hg	
Respiratory Rate	< 8 breaths per minute	> 24 breaths per minute
Hemoglobin oxygen saturation (SpO2)	< 90%	
Temperature	> 100.4 degrees F	

These MET criteria were the first early warning or track and trigger scoring systems. Note that these criteria are single parameter systems, i.e., any one parameter outside the defined limits will trigger the alert. In contrast to alerts based on more than one measure, which require abnormality of more than one vital sign to be triggered, single parameter alerts tend to be less specific. It is important to note that the selection of parameters, and more significantly, the threshold values, for the original MET criteria were based on expert opinion. And that approach has continued in most implementations of rapid response teams.

Shortly after the introduction of the MET and the calling criteria, Morgan, Williams, and Wright developed a composite scoring system, the Early Warning Score, based on Systolic Blood Pressure (SBP), Heart Rate (HR), Respiratory Rate (RR), Temperature (Temp), and a simple mental status score called **AVPU** (***Alert***, responds to ***Voice***, responds to ***Pain*** only, ***Unresponsive***) (7). This was subsequently modified into the Modified Early Warning System or MEWS (8). This was a composite scoring system in which ranges for SBP, HR, RR, temp, and AVPU were given inter-parameter weights and a total score calculated. Once again, in this study, the selection of parameters, the threshold ranges, and the weights were selected by expert opinion. The authors of this publication demonstrate that a composite score of 5 or more (out of a maximum of 15) was associated with an increased risk of death or transfer to a higher level of care (9).

Since the publication of these pioneering manuscripts, there have been multiple modifications of both the single parameter (in the sense that a single parameter out of several can trigger the alert) MET system as well as the composite EWS or

MEWS system. These have included retrospective studies focusing on the sensitivity and specificity of the alert for the prediction of clinical deterioration, as well as prospective studies examining the effect of an alert on adverse events. In all cases, the parameters initially selected for use are based on expert opinion and in most cases the thresholds were selected by expert opinion also. The use of expert opinion could be a potential shortcoming of the scoring systems. Ideally, the parameters and their thresholds, either for single parameter or weighted multiple parameter systems, would be a function of outcome data. Two examples of this data-driven scoring system are the publication of the Worthington physiological score and the more recent paper by Bleyer et al (10, 11) where the authors use multivariate logistic regression to determine odds ratios for clinical deterioration and then use the odds ratios to establish scoring parameters and thresholds.

There have been multiple reviews (12- 18) of early warning systems as well as the publication of the conclusions of a consensus panel held in 2008 (19). These early warning systems are also referred to as “track and trigger” scoring, indicating that vital signs, as well as other clinical observations, are consistently monitored and tracked and when a threshold of abnormality is reached, a defined response by clinicians is instigated. Studies of whether track and trigger scoring systems actually improve outcomes such as reducing mortality or the incidence of cardiac arrest, have all been part of the larger question of whether a MET or rapid response team using a scoring system improve outcome. There have been a number of single centers, before-and –after studies of the MET/rapid response team outcomes (I have excluded studies at pediatric facilities) (20-34). The results of these studies are mixed. Bellomo et al (20), Buist et al (27), and Moon et al (32) each report a decrease in cardiac arrests as well as hospital mortality and Mitchell et al (31) reported a decrease in “unexpected” hospital deaths. Paterson et al (33) report decreased mortality based on before-and-after (post introduction of a composite scoring system) 11 day audits. Jones et al (21), Baxter et al (24), Dacey et al (25), and DeVita et al (28) reported a decrease in the incidence of cardiac arrest but no change in overall hospital mortality. The remainder of these references reported no change in either incidence of cardiac arrest or hospital mortality, although Kyriacos et al (18) reported a decrease in time to treatment.

A 2009 Cochran review of all randomized controlled trials, controlled clinical trials, controlled before-and-after trials, and interrupted time series of outreach programs (MET and rapid response teams) with track and trigger scoring systems and identified only two studies (out of 16) that met pre-determined criteria for study design and rigor (35,36).

One study was a prospective cluster randomized controlled trial of general patient wards in 23 hospitals over 12 months (35). This study found no difference in a composite score of death, cardiac arrest, or unplanned admission to the ICU. Mortality was a secondary outcome and this did not improve with the MET scoring system. In contrast, the second study was a prospective stepped-wedge randomized controlled trial of a phased introduction of critical care outreach (36). In this study there was a reduction in mortality, the primary outcome of the study.

More recently, a meta-analysis of rapid response teams was published (37). In this meta-analysis it was found that rapid response teams decreased the incidence of cardiac arrests but had no impact on mortality. This review concluded that "although RRTs (rapid response teams) have broad appeal, robust evidence to support their effectiveness in reducing hospital mortality is lacking". This conclusion initially seems counter-intuitive, however, there is a plausible explanation. Patients who suffer a cardiac arrest most commonly do so because of underlying disease, and often this disease is not reversible. Examples would be the patient with end-stage chronic obstructive pulmonary disease or heart failure. While the rapid response team may prevent an unexpected cardiac arrest, the activation of this team often results in a transfer to an intensive care unit. And in many of the studies, cardiac arrests in the intensive care unit were not included in the outcome measure. Furthermore, many patients with end-stage underlying disease who are transferred to the intensive care unit may ultimately be placed in a do-not-resuscitate status or palliative care. So fundamentally these studies may not have had adequate power to see a change in overall hospital mortality.

The above literature review indicates that while early warning systems have been in use for some time and have the potential to alert clinicians of clinical deterioration and possibly avert cardiopulmonary arrest, the impact of these systems is still not fully understood. The plethora of systems, the use of expert

opinion rather than outcomes-driven scoring tools, and the conflicting results are indicative of the need for further study.

Optimization of any early warning system will require the capability to load large amounts of pertinent data from source systems (vital signs and also outcomes such as mortality, cardiopulmonary resuscitation, ICU transfer) and then perform analytics, ranging from the simple tasks of identifying the number of patients with various vital sign abnormalities to more complex tasks, such as multi-variate logistic regression to determine the variables that are correlated with clinical deterioration. This will require a business intelligence platform for the extraction, display, and analysis of early warning system data. Business intelligence is defined as "the ability of an organization to collect, maintain, and organize knowledge" (http://en.wikipedia.org/wiki/Business_intelligence). And just as major corporations use business intelligence to accumulate and analyze large amounts of information that can help develop new opportunities and provide a competitive market advantage so too healthcare is beginning to exploit the analysis of large datasets to improve quality of care and patient safety.

In this capstone project, we will describe the development of a business intelligence platform for analysis of an early warning in use at Northeast Georgia Medical Center. The goal is to create a platform that accepts data in multiple formats (such as electronic nursing documentation, laboratory data, or even simple spreadsheets), that facilitates analysis along multiple dimensions so that questions such as which component of the alert or which combination of components is most often associated with clinical deterioration, what is the longitudinal (temporal) evolution of alert frequency or adverse events, or what is the relationship between interventions triggered by the alert and subsequent adverse events, can be readily answered, and that facilitate convenient visual display of the data.

METHODS:

DATABASE

Northeast Georgia Medical Center is a 557 bed tertiary care facility approximately 50 miles north of Atlanta. In the fall of 2011, the nursing service system requested the generation of an electronic alert that was transmitted via email to the nursing supervisor on duty. This alert is currently triggered by any of the following vital sign abnormalities

Heart Rate	< 40 beats per minute	> 130 beats per minute
Systolic Blood Pressure	< 90 mm Hg	
Respiratory Rate	< 8 breaths per minute	> 24 breaths per minute
Hemoglobin oxygen saturation (SpO2)	< 90%	
Temperature	> 100.4 degrees F	

An example of the actual alert is demonstrated below

MET TEAM ALERT - XXXXXXXXXXX Systolic BP - 83.0 Code Status - Full

Patient: XXXXXXXXXXXXX 5/19/2012 5:27:40 AM

Location: S5E - 5431 Patient Type: I

The Patient has the following Code Status: Full

The patient's systolic blood pressure is less than 90 mmHg.

Systolic BP: 83.0

Patient Vitals:

Level of Consciousness

Alert 05/19/12 01:18

BP

88/53 Ruparm 05/19/12 02:42

85/47 Ruparm 05/19/12 03:10

83/55 Ruparm 05/19/12 03:10

O2 Saturation

100% L2l 05/19/12 02:42

100% L2l 05/19/12 03:10

99% L2l 05/19/12 05:26

Temperature

97.8F Oral 05/19/12 02:42

98F Oral 05/19/12 03:10

97.8F Oral 05/19/12 05:26

Heart Rate

79 Brachl 05/19/12 02:42

76 Brachl 05/19/12 03:10

96 05/19/12 05:26

Respiration Rate

18 Visual 05/19/12 02:42

18 Visual 05/19/12 03:10

16 05/19/12 05:26

Note that the alert not only indicates the specific time-stamped vital sign abnormality but also includes the three time-stamped prior measurements for all other vital signs as well as an assessment of level of consciousness (expressed as alert, delayed, lethargic, obtunded, or stuporous). Note that the timing of vital signs measurements are specific to the nursing unit and to the patient and are recorded at least once a shift.

This alert has proven to be highly non-specific and, indeed, is triggered approximately 70-80 times a day a number significantly in excess of the number of patients who actually deteriorate (approximately 1-2 patients a day deteriorate to the point of activation of the medical emergency team, transfer to the intensive care unit, or cardiopulmonary arrest) raising concern that the low specificity will lead to the alert being ignored. This alert is distinct from the actual activation of the medical emergency team and is simply a warning to the nursing supervisors, sent to them as email messages. Typically, the nursing supervisor will contact the nursing unit for more information and then intervene as they deem appropriate. While the nursing service believes it has been beneficial, facilitating early identification of patients who are clinically deteriorating, it has clearly not been optimized in any formal sense. Currently the record of these alerts is stored in an Oracle database table. Patient outcomes, such as mortality or transfer to the ICU, are stored in the administrative database, and occurrences of cardiopulmonary resuscitation or true hands-on MET calls are recorded in simple Excel spreadsheets. The alert records as well as these clinical outcomes provide the data necessary for development of a business intelligence tool for analysis of the efficacy of the alert.

Business Intelligence Platform

The analysis of the early warning alert in use by the nursing service at Northeast Georgia Medical Center was implemented as a QlikView™ (QlikTech International AB, Sweden) application. QlikView™ is described as an "associative in-memory" business intelligence tool (38, 39). Qlikview applications hold all the analytic query data in memory (RAM versus a disk storage system) by a proprietary compression algorithm leading to minimal wait times for data query retrieval. The term associative is used to indicate that a Qlkview application will associate database fields that have the same name (precisely) facilitating associations over these fields. To get its user interface flexibility, QlikView implicitly assumes a "star/snowflake schema" (40), **t**hat is, there is no more and no less than one possible path between any pair of tables (41).

QlikView has an associative search engine and by searching on a particular field, all other associated fields are displayed. This is in contrast to traditional on-line processing (OLAP) tools query-based business intelligence tools that require hypercube development for query associations. With query-based tools creating associations among all available fields may require a high level of IT support or else result in delays in business intelligence development that compromise the enterprise. With QlikView there is immediate association of connected data. As an example cited in "QlikView Architectural Overview" is suppose a user wants to identify a sales representative but can't recall the name only that the sales representative sells fish to customers in Nordic regions. By searching the sales representative list box for "Nordic" and "fish" all sales representatives who meet those criteria are listed.

A good example of the table structure is shown below with five connected tables with three connected by PATIENT_ACCT, two connected by DIAGNOSIS_CONS, two connected by DIAG_LINK_DATE and another pair by ORDERING_ID.

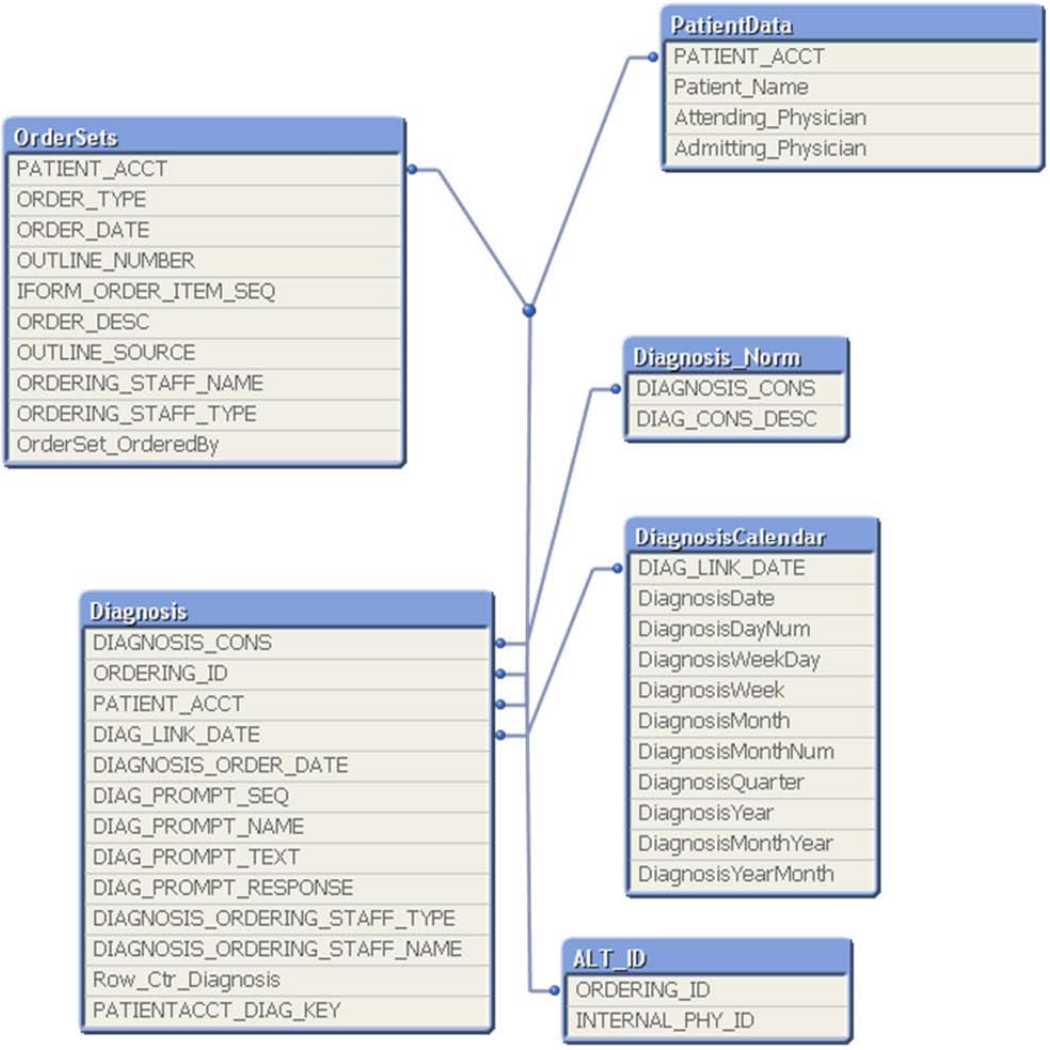
PatientData
PATIENT_ACCT
Patient_Name
Attending_Physician
Admitting_Physician
OrderSets
PATIENT_ACCT
ORDER_TYPE
ORDER_DATE
OUTLINE_NUMBER
IFORM_ORDER_ITEM_SEQ
ORDER_DESC
OUTLINE_SOURCE
ORDERING_STAFF_NAME
ORDERING_STAFF_TYPE
OrderSet_OrderedBy
Diagnosis_Norm
DIAGNOSIS_CONS
DIAG_CONS_DESC
DiagnosisCalendar
DIAG_LINK_DATE
DiagnosisDate
DiagnosisDayNum
DiagnosisWeekDay
DiagnosisWeek
DiagnosisMonth
DiagnosisMonthNum
DiagnosisQuarter
DiagnosisYear
DiagnosisMonthYear
DiagnosisYearMonth
Diagnosis
DIAGNOSIS_CONS
ORDERING_ID
PATIENT_ACCT
DIAG_LINK_DATE
DIAGNOSIS_ORDER_DATE
DIAG_PROMPT_SEQ
DIAG_PROMPT_NAME
DIAG_PROMPT_TEXT
DIAG_PROMPT_RESPONSE
DIAGNOSIS_ORDERING_STAFF_TYPE
DIAGNOSIS_ORDERING_STAFF_NAME
Row_Ctr_Diagnosis
PATIENTACCT_DIAG_KEY
ALT_ID
ORDERING_ID
INTERNAL_PHY_ID

Implementation

The primary challenges implementing our QlikView application were that some patients had multiple adverse events (most typically medical emergency team activation) and that in many incidences there were at times no temporal relationship between the alert and the adverse event. For example, a patient may have had an alert on one day and an adverse event two days later. Or, as another example, the patient may have had an alert and then a medical emergency team activation 2 hours later but then a second medical emergency team activation the next day. This complicated simple “joins” on the patient ID. The solution was to load data from the MET activation, cardiopulmonary resuscitation (referred to simply as a code or code “blue”), and ICU transfer tables into the alert table rather than trying to “join” tables. This workaround was efficient for us as the MET activations and code blues were recorded in spreadsheets and loading this data into QlikView is quite simple. The resultant table structure for this view is

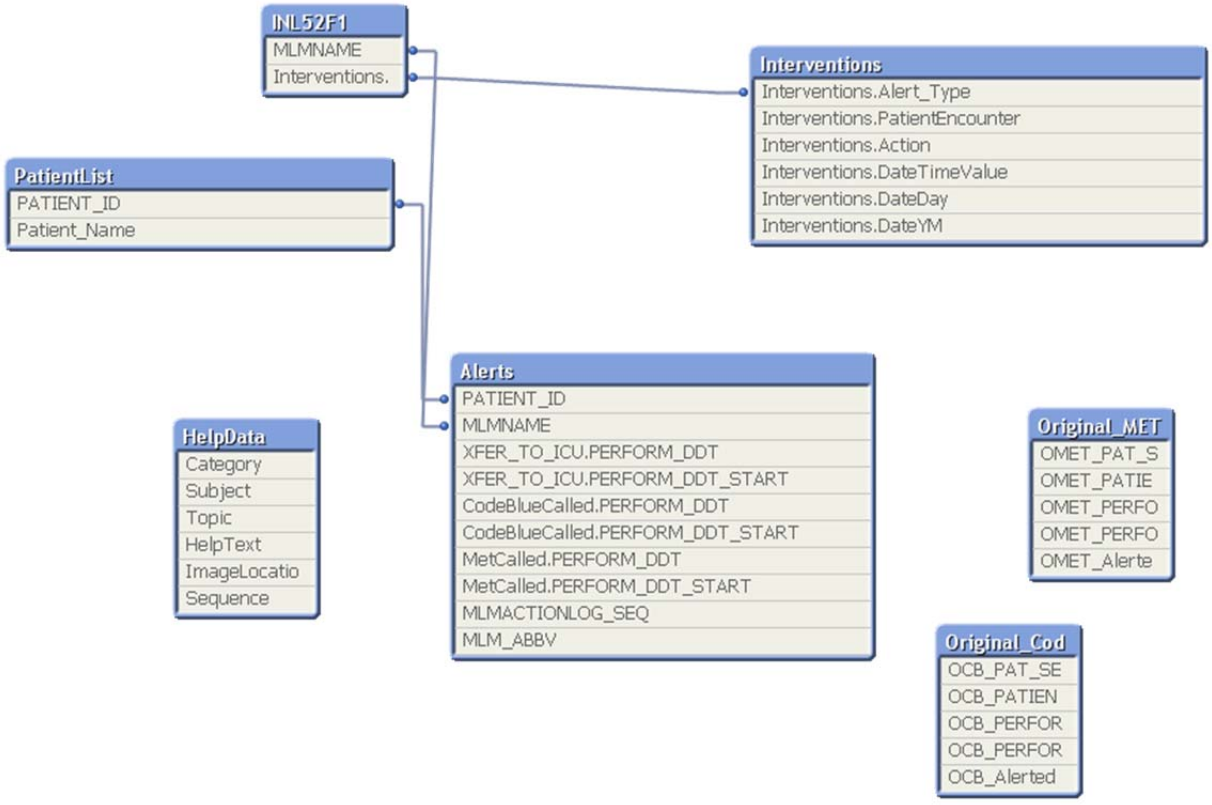

Note that we have included in the application several free-standing tables, including codes, MET activations, and a “help” table. We have joined the original alert table (INL52F1) with a separate table (stored on a “SharePoint” site) where the nursing staff records any interventions triggered by an alert to the nursing supervisor.

The most basic component of any QlikView document is a sheet and any application may have multiple sheets. Each sheet may contain multiple objects including

List boxes

Statistics boxes

Multi boxes

Table boxes

Charts (including pivot tables and straight tables)

Input boxes

Slider objects

Current Selection boxes

Bookmark objects

Buttons

Text objects

Line/Arrow objects

By right mouse clicking on the banner of a QlikView table one can access the properties tabs, which control the organization of the table. An example of one of the sheets in this application, “Dashboard-unique patient %” (a sheet that presents alert and adverse event data for unique patients), is shown below.

.

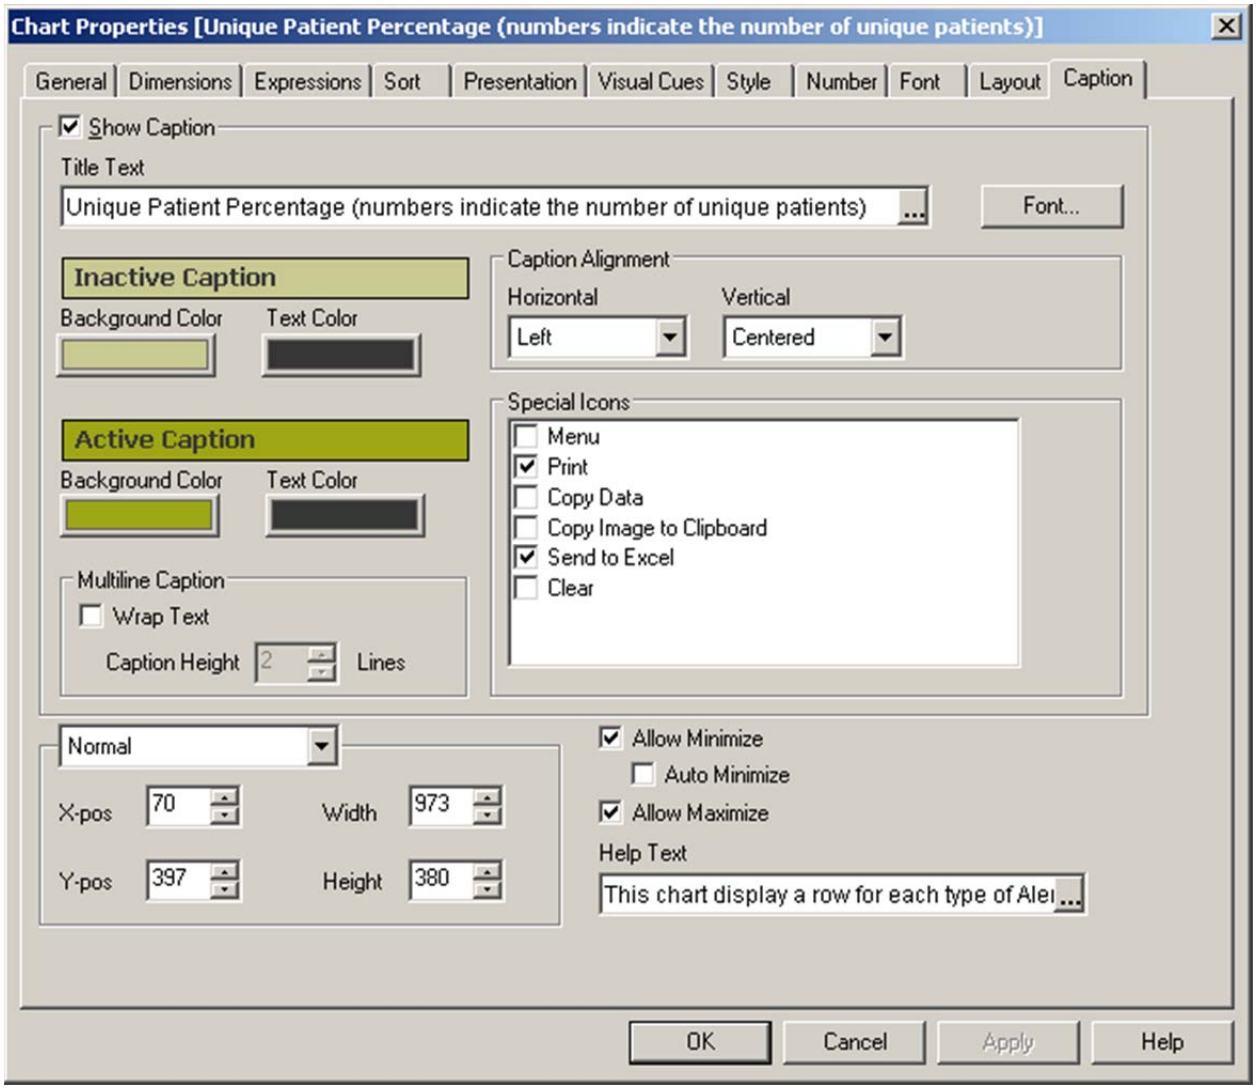

While the “General” tab focuses on the color scheme, font, caption, etc. of the table, more relevant tabs include the “Dimensions” and “Expressions” tabs. In QlikView “Dimension” refers to how data is grouped while “Expression” refers to what is calculated. For example, consider a company that wishes to know the “sales per region”. The dimension is what comes after the “per” (in this case, region), while the expression is what comes before the “per” (in this case, sales). Shown below is the “Dimensions” tab for the “Dashboard-unique patient %” sheet.

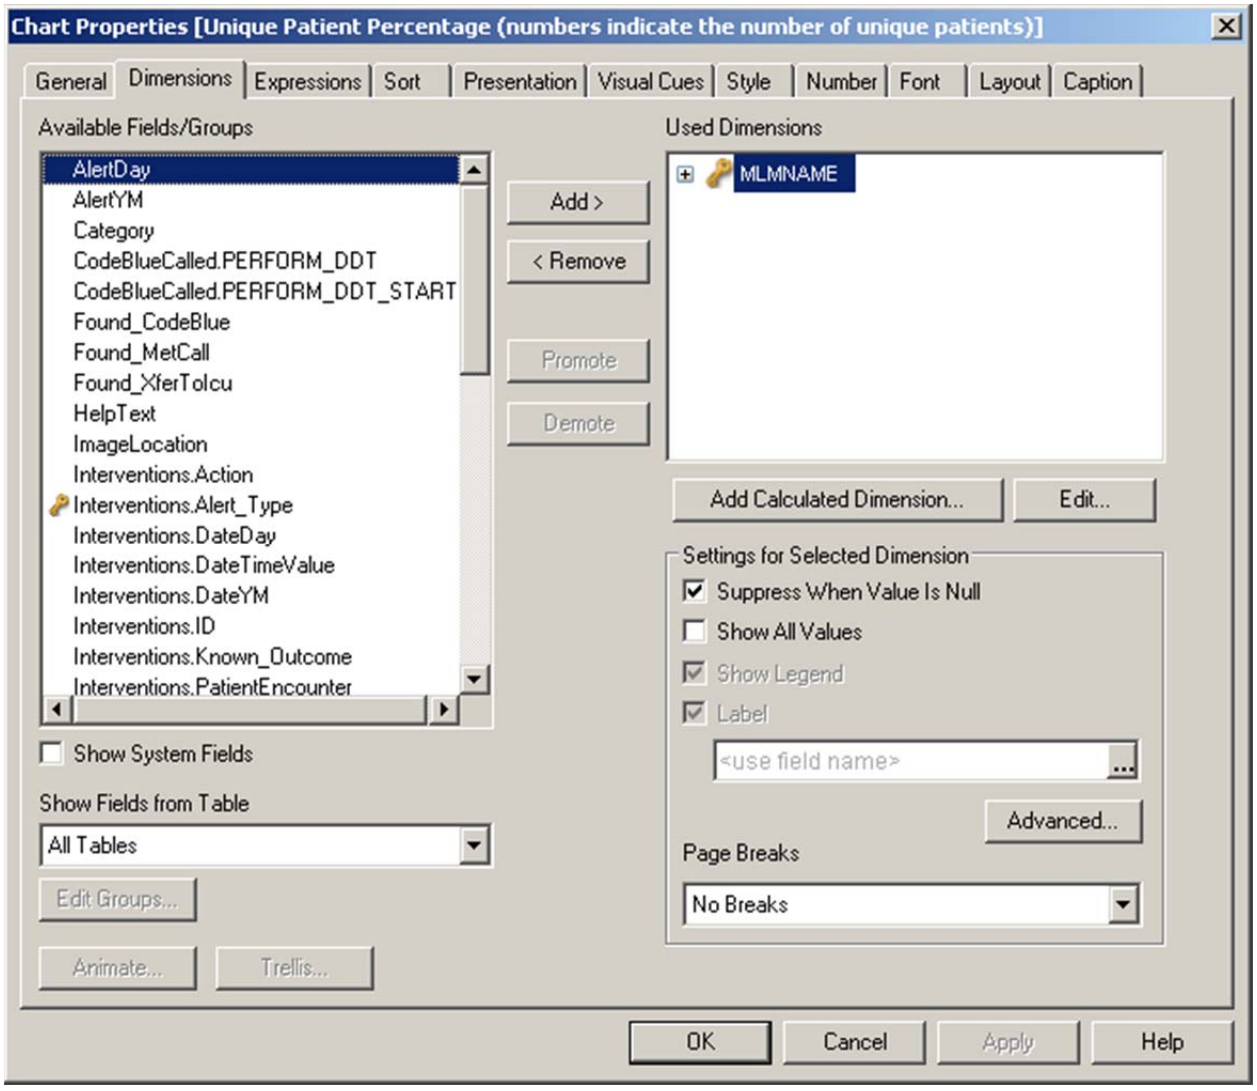

For this sheet, the dimension is the name of the alert from the original alert table (INL52F1) described above. Thus the data is grouped by the name of the alert. The “Expressions” tab is now shown.

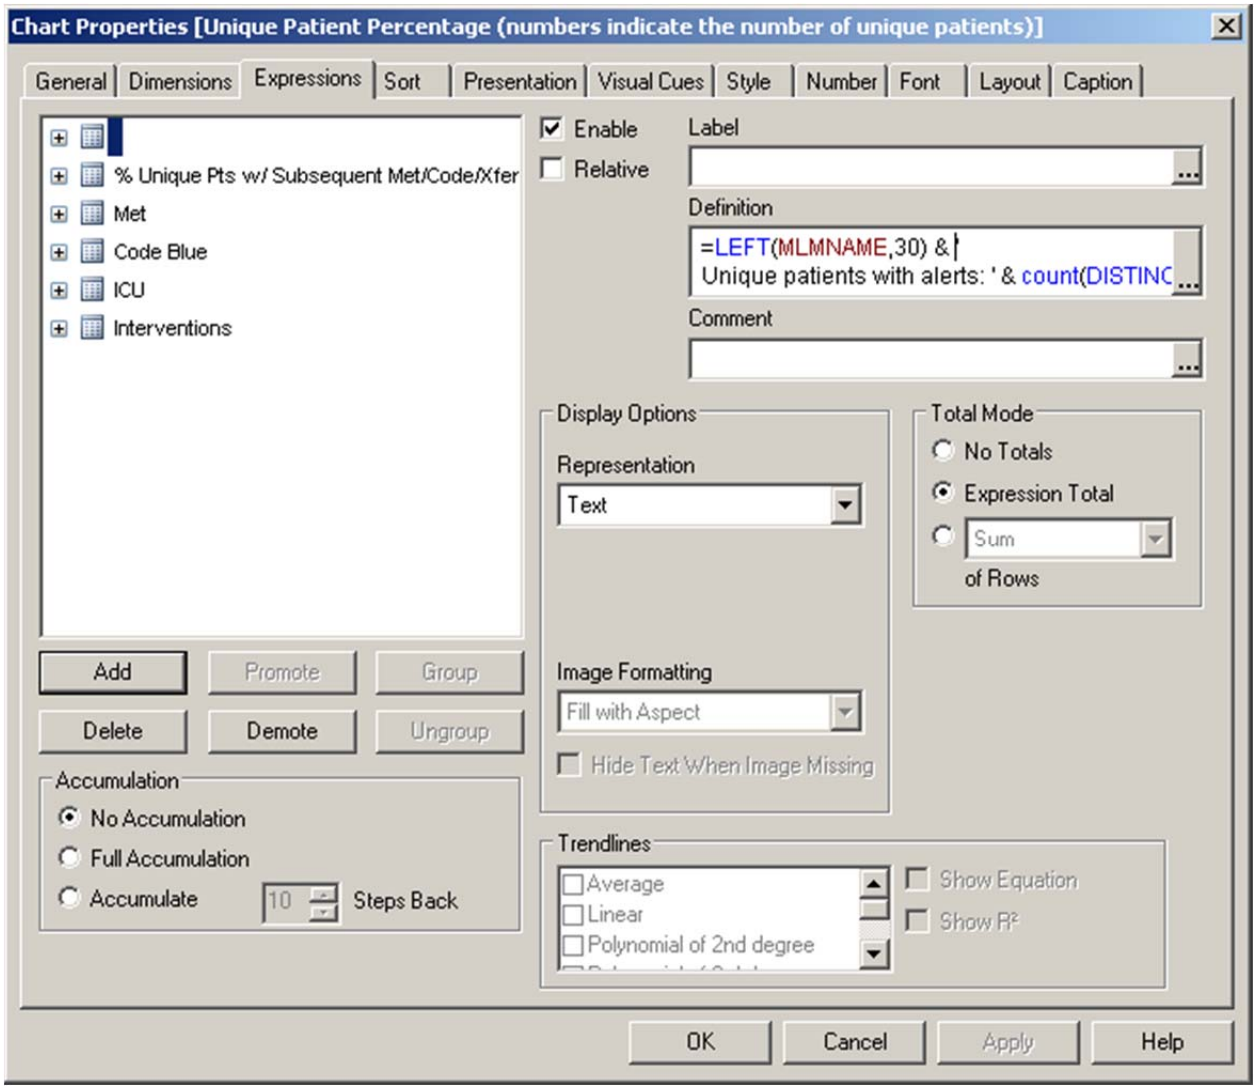

The definition of the expression, i.e., what is calculated, is shown in the "Definition" box. The full expression is

```
=LEFT(MLMNAME,30) & '
Unique patients with alerts: ' & count(DISTINCT PATIENT_ID)
```

and it should be clear that this expression calculates the number ("count") of unique ("distinct") patients who had an alert of the specified type (dimension).

We organized this “Early Warning Alert” application into 8 sheets, providing the following data

1. Home page: describes the purpose of the application and the most recent data refresh time, and offers buttons to quickly navigate to the three primary sheets (unique patient tab, number of alerts tab, and alert details tab)
2. Dashboard-unique patient %: this sheet contains a table that shows the percentage of unique patients who have a subsequent (within 12 hrs) MET activation, code, or transfer to the ICU and also the percentage of unique patients who had some intervention by the nursing supervisor triggered by the alert. This is classified by the type of alert (blood pressure, heart rate, respiratory rate, temperature and SpO2.
3. Dashboard-Number of alerts: This tab shows the total number of alerts, classified by type, as well as the number of patients with an alert who had a subsequent MET activation, code, transfer to the ICU, and intervention prompted by the alert.
4. Analysis: This tab presents the same data as found in the two prior tabs, however it presents all the data in a single table with more columns
5. Alert Details: This tab presents specific data about the alert including patient identification, the type of the alert, and the subsequent outcome.
6. Code Blue and MET details: This tab lists all patients who had a code in the time frame under analysis and indicates whether an alert was ever triggered for that patient.
7. Interventions: Enumerates details of the interventions that were prompted by the alert
8. Help: Provides access to QlikView help topics

For each of these tabs, there are selection boxes at the top of the tab to select the time frame of analysis. Note that the data for MET calls and codes is entered monthly.

Programming

The following programming code was used to create this application. Most of the following is QlikView scripting language although there are interspersed SQL statements. While this code is specific to the platform it is relatively transparent.

```
//Load interventions triggered by alert from SharePoint site
Qualify *;
Interventions:
LOAD PatientEncounter, Action, Alert_Type, ConvertToLocalTime([Date/Time],
'Eastern Time (US & Canada)') as DateTimeValue,
Date(Floor(ConvertToLocalTime([Date/Time], 'Eastern Time (US &
Canada)')),'MM/DD/YYYY') as DateDay,
Date(MakeDate(Year([Date/Time]),Month([Date/Time])),'YYYY-MM') as
DateYM,
ID, Known_Outcome, Room_Number;
SELECT a.nvarchar1 AS PatientEncounter, a.int2 AS Action, a.int1 AS
Alert_Type, a.datetime1 AS [Date/Time], a.tp_ID AS ID, a.nvarchar5 AS
Known_Outcome,
a.nvarchar3 AS Room_Number
FROM dbo.AllUserData AS a WHERE (tp_ListId= 'c8a43710-ca93-47cc-b83f-
5f3e0e32e43b') AND a.tp_ID not in (SELECT ListItemId FROM RecycleBin
WHERE ListId= 'c8a43710-ca93-47cc-b83f-5f3e0e32e43b');

UNQUALIFY *;

// Build a quick xref from the alert type number to the alert name so that they can
be joined
LOAD * INLINE [
   Interventions.Alert_Type, MLMNAME
   1, NGHS MET Team Alert HR
   2, NGHS MET Team Alert Resp
   3, NGHS MET Team Alert BP
   5, NGHS MET Team Alert O2 Sat
   6, NGHS MET Team Alert Temp
];


///$tab PatientsInTrouble
SET TimestampFormat='M/D/YYYY h:mm:ss[.fff] TT';
```

```
UNQUALIFY *;

// First thing we want to do is load up the patient names so we have them as a
master
TRACE ===== Load the Patient List so we have master of SEQ, ID and NAME
=====;
PatientList:
LOAD PAT_SEQ,
  LAST_NAME&', '&FIRST_NAME&' '&MIDDLE_NAME as Patient_Name,
     PATIENT_ID
FROM [..\HCI\DataFiles\HCI_Patients.qvd] (qvd);

// We need to qualify everything from now on so that only PATIENT_ID is the
link
Qualify *;
Unqualify PATIENT_ID;
UNQUALIFY PAT_SEQ;

// Now load the temporary Transfers, CodeBlue, MetCall tables
// Subtracting the .5 means take 12 hours from the value. This identifies transfers
within 12 hrs of alert
XFER_TO_ICU:
LOAD REPLACE(LTRIM(REPLACE("AcctNumber",'0',' ')),' ','0') as
PATIENT_ID,
  EffectDateTime as PERFORM_DDT,
  EffectDateTime-.5 as PERFORM_DDT_START
FROM ..\HBI_Extracts\DataFiles\XFER_TOICU.QVD (qvd);

CodeBlueCalled:
LOAD PAT_SEQ,
     PERFORM_DT as PERFORM_DDT,
   PERFORM_DT-.5 as PERFORM_DDT_START
FROM [..\HCI\DataFiles\HCI_ImporantPATResults.qvd] (qvd)
WHERE LABEL_NAME = 'Code Blue'
and RESULT_VALUE = 'Yes';

Left Join
LOAD PAT_SEQ,
```

```
PATIENT_ID
Resident PatientList;

MetCalled:
LOAD PAT_SEQ,
    PERFORM_DT as PERFORM_DDT,
  PERFORM_DT-.5 as PERFORM_DDT_START
FROM [..\HCI\DataFiles\HCI_ImporantPATResults.qvd] (qvd)
WHERE LABEL_NAME = 'Met Team Called'
and RESULT_VALUE = 'Yes';

Left Join
LOAD PAT_SEQ,
PATIENT_ID
Resident PatientList;
///$tab HCI - Alerts
OLEDB CONNECT TO [Provider=OraOLEDB.Oracle.1;Persist Security
Info=True;User ID=ccdev;Data Source=HCILIVE;Extended Properties=""]
(XPassword is SXULXRNMELZMHfA);

Unqualify *;
TRACE ===== Load the ALRTS =====;

//Create a temporary alerts table
Alerts_Temp:
LOAD MLMACTIONLOG_SEQ,
    MLMNAME,
    MID(MLMNAME, 20) as MLM_ABBV,
    URGENCY,
    MESSAGE,
    TRANSACTIONDTTM,
    Date(FLOOR(TRANSACTIONDTTM),'MM/DD/YYYY') as AlertDay,
  Date(MakeDate(Year(TRANSACTIONDTTM),Month(TRANSACTIONDTTM)
),'YYYY-MM') as AlertYM,
    PAT_SEQ
FROM [..\HCI\DataFiles\HCI_Alerts.qvd] (qvd)
WHERE LEFT(MLMNAME, 19) = 'NGHS MET Team Alert'
OR LEFT(MLMNAME, 15) = 'NGHS_VITAL_SIGN';
// Join in the field to the alert so that we can use Patient ID instead of pat seq
LEFT JOIN
```

```
LOAD PAT_SEQ,
PATIENT_ID
Resident PatientList;

// We need to use the interval match function to join the data with the actions
// We are keying on the transaction date and time of the Alert and seeing when it is within the range
// from the start time to the performed time.
TRACE ===== Utilize the Interval Match function to join the alerts with the actions =====;
Alert_CodeBlue_Facts:
IntervalMatch (TRANSACTIONDTTM, PATIENT_ID)
Left Join
LOAD CodeBlueCalled.PERFORM_DDT_START,
CodeBlueCalled.PERFORM_DDT, PATIENT_ID
Resident CodeBlueCalled;
Left Join Load *
Resident CodeBlueCalled;


Alert_MetTeam_Facts:
IntervalMatch (TRANSACTIONDTTM, PATIENT_ID)
Left Join LOAD MetCalled.PERFORM_DDT_START,
MetCalled.PERFORM_DDT, PATIENT_ID
Resident MetCalled;
Left Join Load *
Resident MetCalled;

Alert_XferToIcue_Facts:
IntervalMatch (TRANSACTIONDTTM, PATIENT_ID)
Left Join LOAD XFER_TO_ICU.PERFORM_DDT_START,
XFER_TO_ICU.PERFORM_DDT, PATIENT_ID
Resident XFER_TO_ICU;
Left Join Load *
Resident XFER_TO_ICU;

TRACE ===== Move the alerts from temp table to real table then drop alerts temp table =====;
Alerts:
LOAD *
```

```
,IF (Not IsNull( CodeBlueCalled.PERFORM_DDT), 1, 0) as Found_CodeBlue
,IF (Not IsNull( MetCalled.PERFORM_DDT), 1, 0) as Found_MetCall
,IF (Not IsNull( XFER_TO_ICU.PERFORM_DDT), 1, 0) as Found_XferToIcu
Resident Alerts_Temp;

DROP table Alerts_Temp;

TRACE ===== Drop the action based tables because we don't need them anymore
=====;
// Get rid of the underlying tables that fed our matches
Original_CodeBlue:
LOAD PAT_SEQ as OCB_PAT_SEQ,
PATIENT_ID as OCB_PATIENT_ID,
CodeBlueCalled.PERFORM_DDT as OCB_PERFORM_DDT,
Date(MakeDate(Year(CodeBlueCalled.PERFORM_DDT),Month(CodeBlueCalled
.PERFORM_DDT)),'YYYY-MM') as OCB_PERFORM_YM,
IF( ISNULL(Lookup('PATIENT_ID', 'PATIENT_ID', PATIENT_ID, 'Alerts')),
'Not Alerted','Alerted') as OCB_Alerted
Resident  CodeBlueCalled;

Original_MET:
LOAD PAT_SEQ as OMET_PAT_SEQ,
PATIENT_ID as OMET_PATIENT_ID,
MetCalled.PERFORM_DDT as OMET_PERFORM_DDT,
Date(MakeDate(Year(MetCalled.PERFORM_DDT),Month(MetCalled.PERFORM
_DDT)),'YYYY-MM') as OMET_PERFORM_YM,
IF( ISNULL(Lookup('PATIENT_ID', 'PATIENT_ID', PATIENT_ID, 'Alerts')),
'Not Alerted','Alerted') as OMET_Alerted
Resident  MetCalled;


// Now that all tables have patient id we need to get rid of pat_seq since it would
cause synthetic key
drop Field PAT_SEQ;

drop table CodeBlueCalled;
drop Table MetCalled;
drop Table XFER_TO_ICU;
```

RESULTS:

The home page of this business intelligence application is shown below

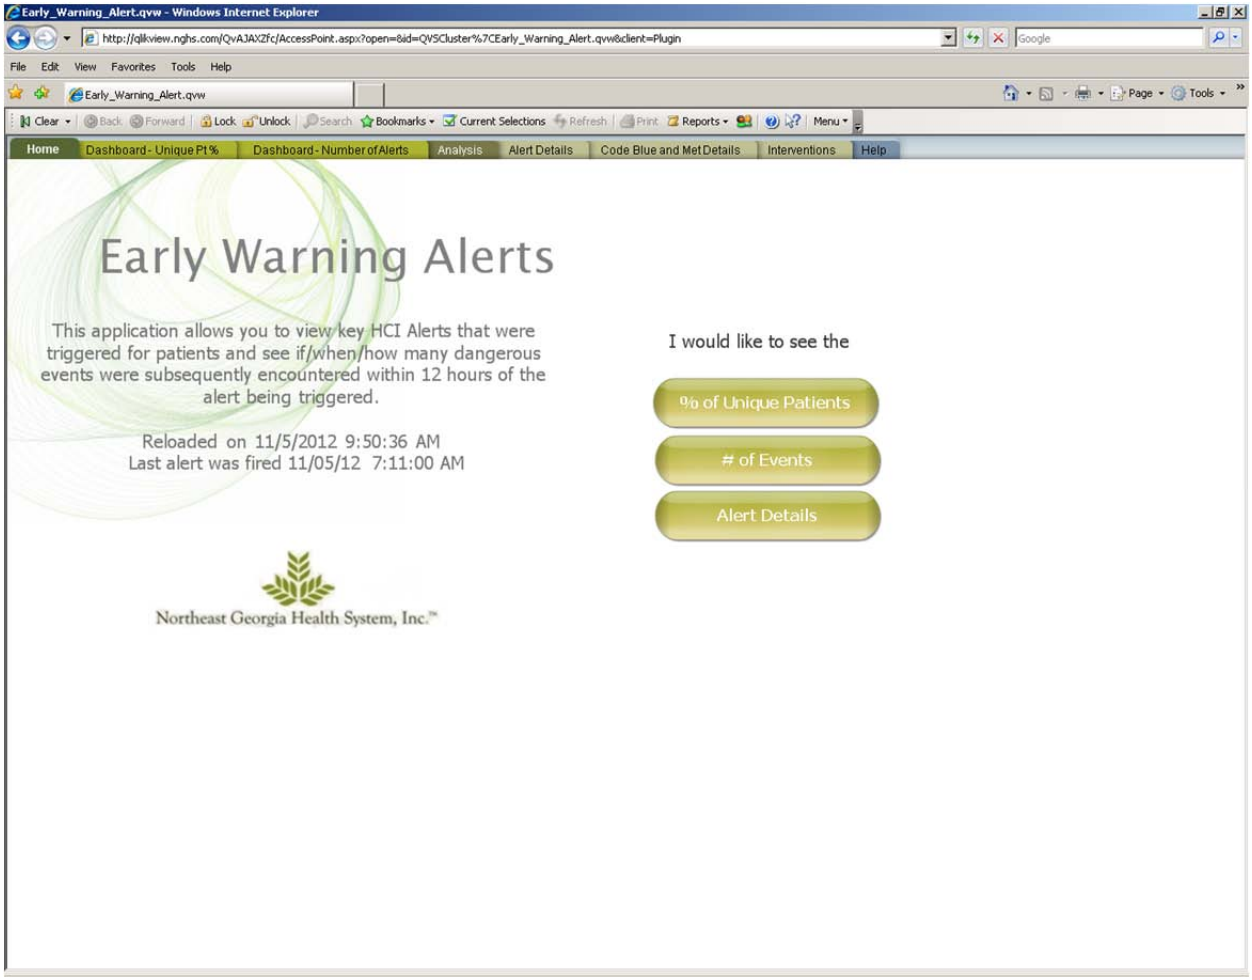

This sheet briefly explains the purpose of the application and provides buttons and tabs to access the sheets that summarize the data.

The next sheet is the “unique patients” sheet and is shown below.

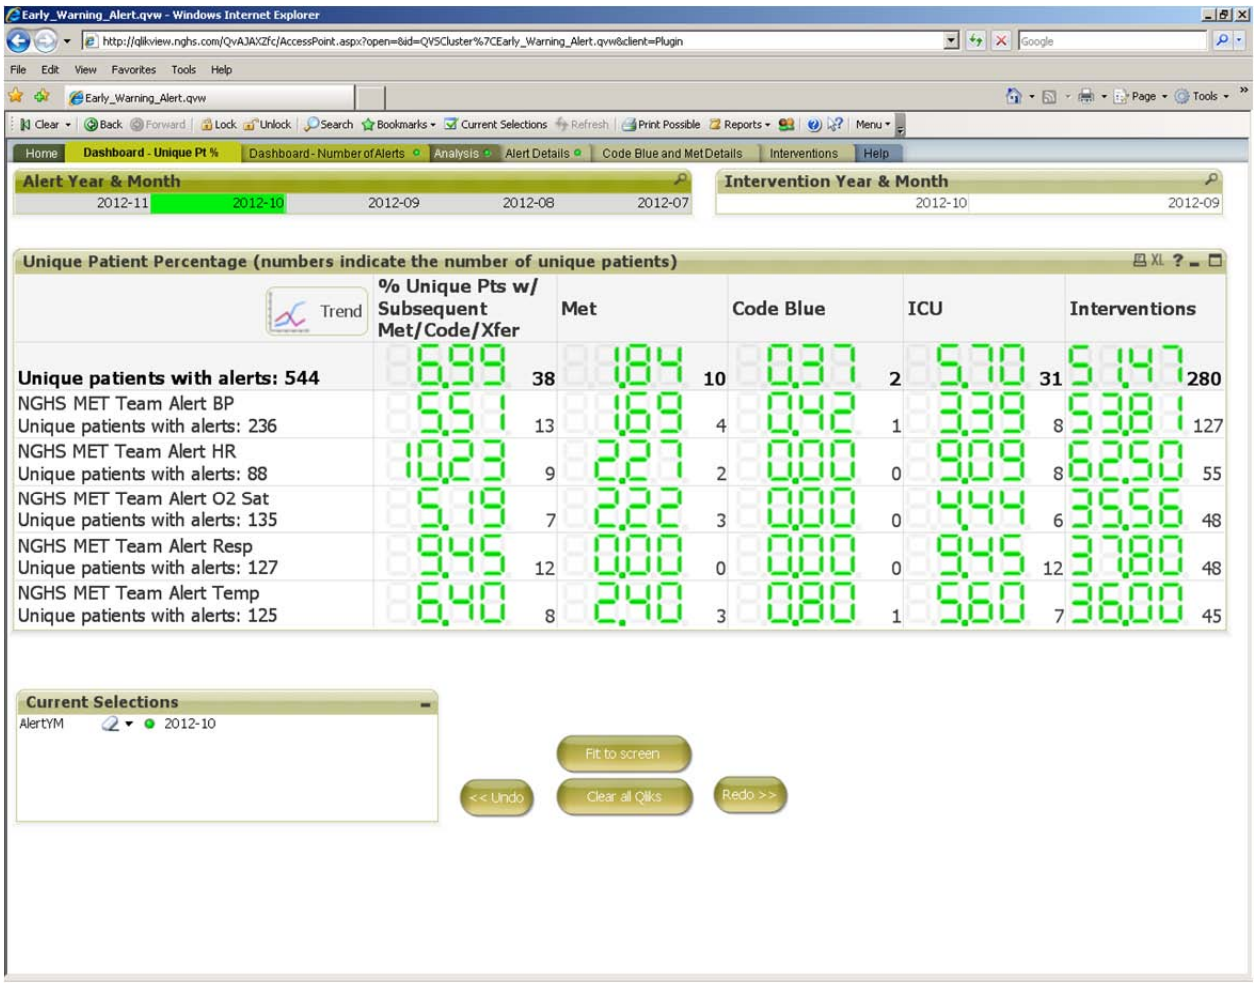

Note that we have selected data for October 2012 (highlighted in green at the top of the sheet). This sheet (1st row) shows that in October 2012 there were 544 unique patients for whom an alert was triggered (see column #1) and 6.99% (38, shown in smaller black number in column #2) of these had a subsequent MET activation, code, or transfer to the ICU. The breakdown into each of these categories is shown in columns #3-5. We also see that 51.47% of these patients (280) had an intervention by the nursing supervisor who received the alert. Subsequent rows show the statistics for the specific types of alerts (blood pressure, heart rate, SpO2, respiratory rate, and temperature). The most obvious conclusion one reaches upon examination of this sheet is that the alert is non-specific. Less than 10% of the unique patients for whom an alert was triggered go on to have an adverse event (defined as a MET call, code, or transfer to the ICU). However, note that over 50% of the patients had some sort of intervention. What we cannot yet analyze is whether interventions result in a lower incidence of adverse events. This

reflects the fact that currently the interventions are recorded on paper and then transferred to an Excel spreadsheet, which is loaded into the QlikView application. The nursing service has not yet used a unique patient identifier for this record. This is under correction.

We see that heart rate and blood pressure alerts led to more interventions than the other specific alerts. We also see that the heart rate alert and the respiratory rate alerts were most often associated with subsequent adverse events. Also note that when we total the number of unique patients for the various categories of alerts the sum (711) exceeds 544, implying that some patients have triggered multiple alerts.

By clicking on the “Trend” icon the viewer can see a line graph illustrating the frequencies of adverse events in unique patients as a time series. This function should prove useful in the future, although of limited interest at this time as the data for months prior to October 2012 was not consistently collected and verified.

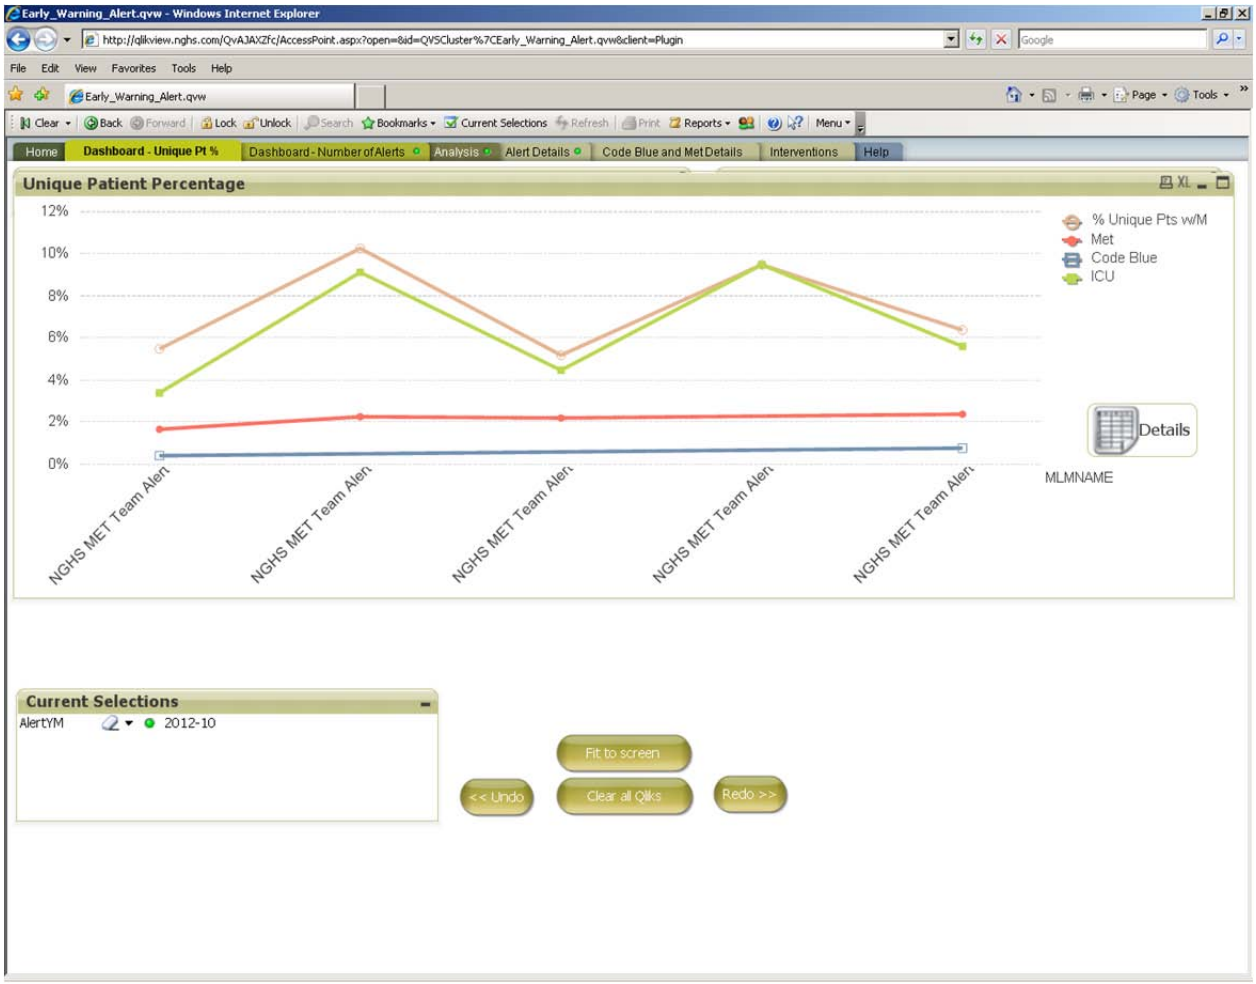

The sheet labeled “Dashboard-Number of Alerts” presents an alternative analysis of the early warning alert. This sheet presents the number of alerts rather than unique patients as shown below. In the first row, first and second columns we see that in October of 2012 there were 2013 total alerts and there were 112 alerts that had a subsequent adverse event, i.e., only 6% had a subsequent adverse event and this occurred in 38 unique patients (these latter two statistics are seen as smaller font black numbers surrounding the total number of alerts with a subsequent adverse event). The observation that the number of alerts with a subsequent adverse event was almost 3 times larger than the number of unique patients suggests that multiple alerts were triggered prior to the adverse event. The remainder of the table breaks these overall statistics into the various categories of alerts and also of adverse events. We also see that 560 interventions (6th column) in 280 unique patients were stimulated by these alerts.

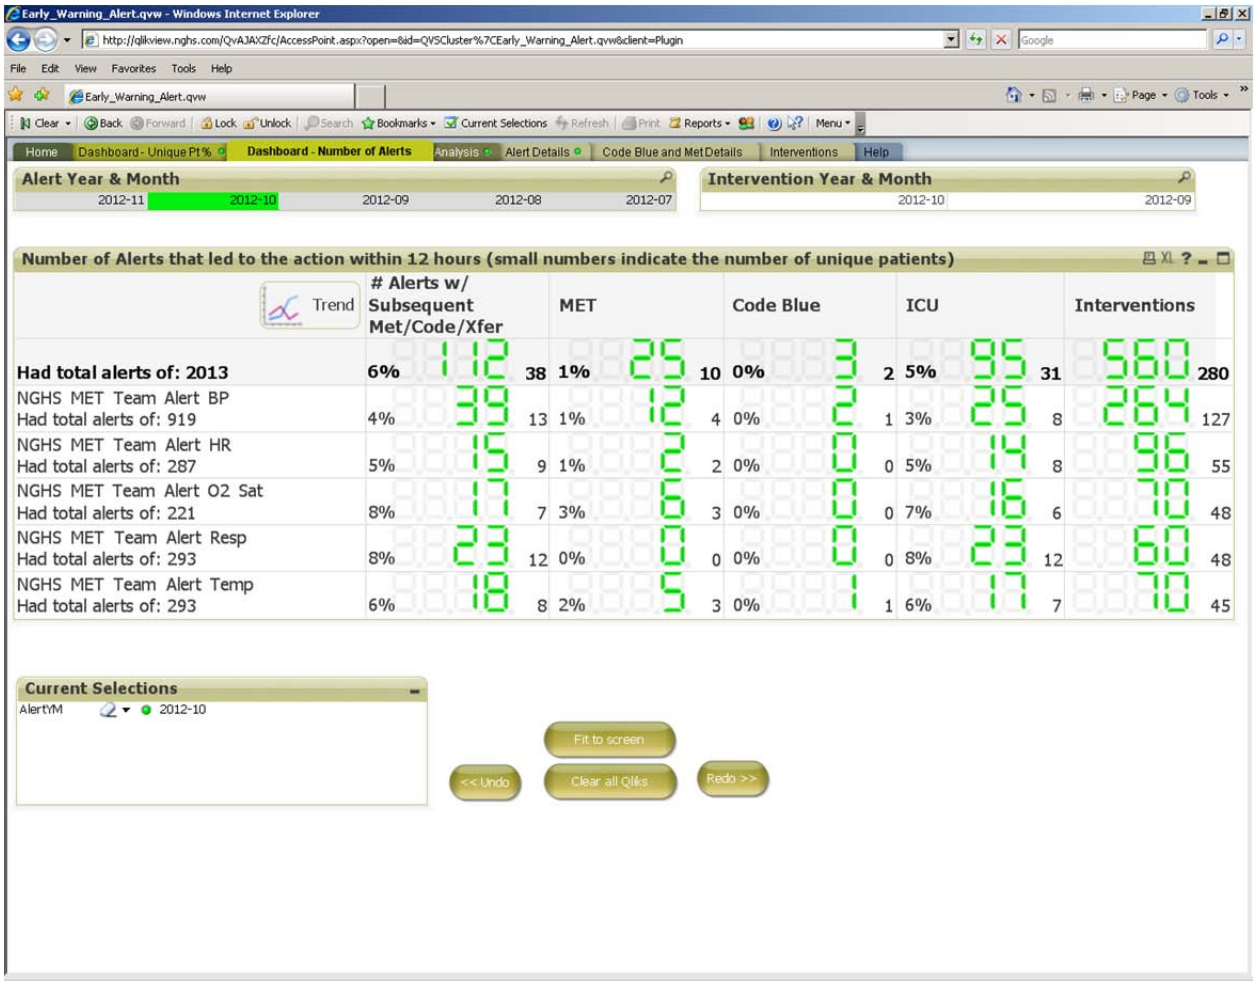

Both the “Dashboard-Unique Pts” and “Dashboard-Number of Alerts” show that the number of unique patients is much smaller than the number of alerts. There are multiple alerts that are triggered by some unique patients. We see that while roughly 50% of unique patients with one or more alert had some sort of intervention, in terms of raw numbers 27.8% (560/2013) of the alerts led to an intervention. The implication is that the nursing supervisor is intervening only when there are multiple alerts for a unique patient.

We can see in the SpO2 and respiratory rate alerts are more likely to be associated with a subsequent adverse event. The heart rate and blood pressure alerts seem to have triggered the greatest number of interventions.

One interesting question and concern is whether patient can experience a serious adverse event, defined as MET activation or code, and not have had a prior alert. The “Code Blue and Met Details” sheet has a simple table that indicates whether patients experiencing a code blue or MET activation had an alert at some point in the time frame under analysis. This is illustrated below.

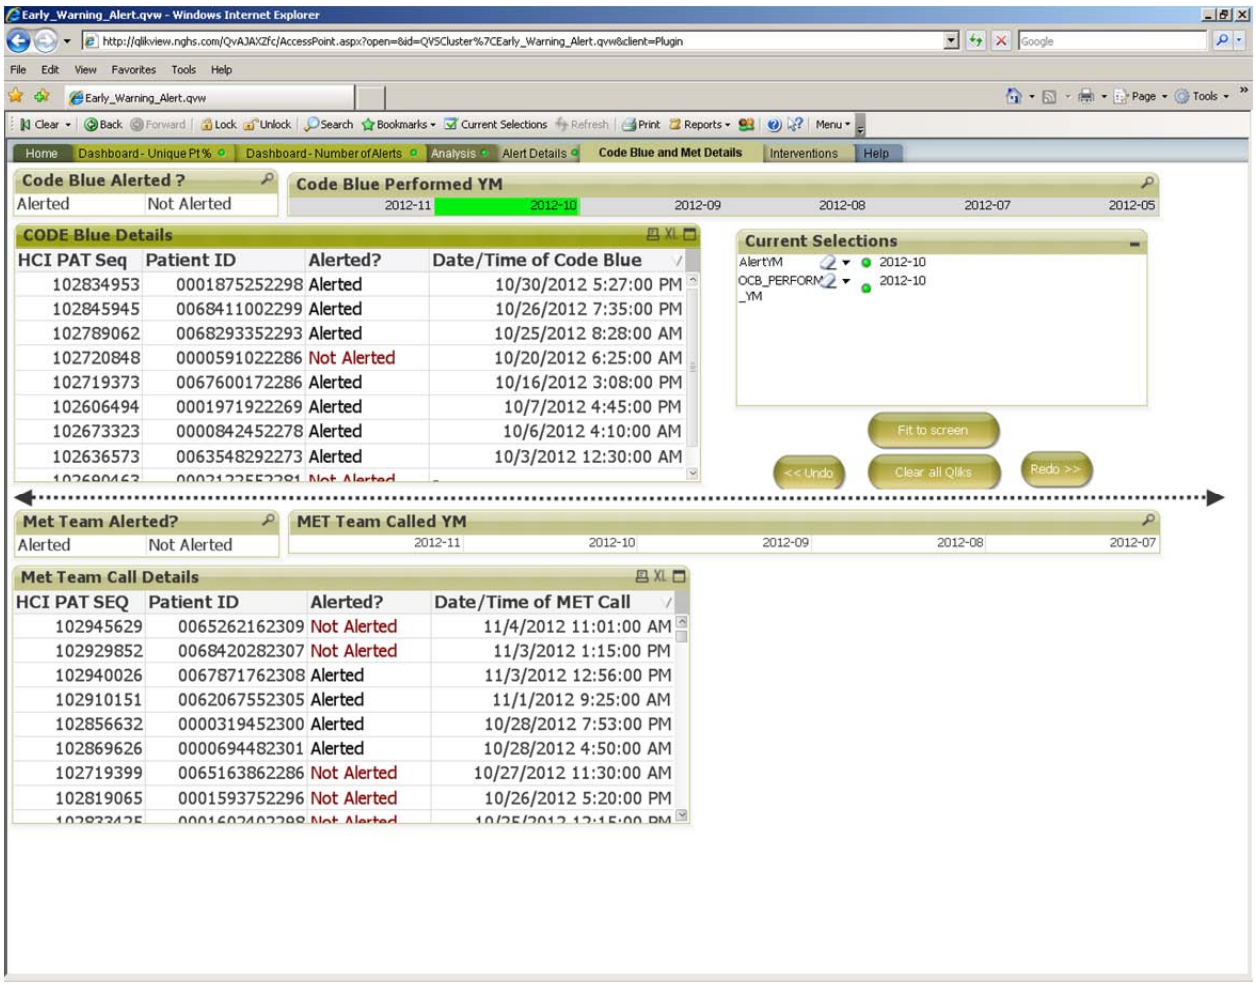

For the month of October 2012 there were 9 cardio-pulmonary arrests (codes) and 2 of this did not have an alert generated at some point in this time frame. There were 107 MET activations and of these 34 did not have an alert. At this point in the evolution of our application, there is insufficient data (particularly for codes) to draw conclusions. Furthermore, we currently aggregate data on a monthly basis and it will be more interesting to look at a shorter time interval since the number of patients who did not have an alert within a week (as an example) of an adverse event could be predicted to be greater than the number of patients who did not have an alert within a month period. In other words, the probability of an alert will, in general, increase with the time span of observation if only due to random variation in vital signs. A useful enhancement for the current application will be the option of aggregating data over a shorter time interval. By doing this we should learn more about patients who have an adverse event without a prior alert.

Further insight into early warning of cardiopulmonary arrests and MET calls can be found in the “Alert Details” sheet. As seen below one can determine by appropriate filtration of this sheet which alerts occur most commonly prior to either a cardiac arrest or a MET call. This screen shot shows 5 examples of cardiac arrests which were preceded by a respiratory rate alert. To date, the respiratory rate alert seems to be the one most commonly preceding a cardiopulmonary arrest, although we have not yet collected enough data to have confidence in this conclusion.

Alert Type	Alert Seq#	Patient ID	Alert Date Time	Code Blue Date & Time	MET Date & Time	XFER Date & Time
NGHS MET Team Alert HR	1134620	0065272672272	9/29/2012 3:57:14 PM	-	-	-
NGHS MET Team Alert O2...	1134619	0000752582271	9/29/2012 3:52:15 PM	-	-	-
NGHS MET Team Alert Temp	1134577	0001411232272	9/29/2012 3:17:48 PM	-	-	-
NGHS MET Team Alert Resp	1134557	0068346332249	9/29/2012 3:14:16 PM	9/29/2012 3:58:00 PM	-	-
NGHS MET Team Alert Temp	1134478	0001411232272	9/29/2012 2:13:06 PM	-	-	-
NGHS MET Team Alert Temp	1134477	0001411232272	9/29/2012 2:10:43 PM	-	-	-
NGHS MET Team Alert Resp	1134441	0067481962264	9/29/2012 1:40:27 PM	-	-	-
NGHS MET Team Alert Resp	1134360	0068346332249	9/29/2012 12:56:07 PM	9/29/2012 3:58:00 PM	-	-
NGHS MET Team Alert Resp	1134339	0068346332249	9/29/2012 12:05:59 PM	9/29/2012 3:58:00 PM	-	-
NGHS MET Team Alert Temp	1134318	0068377002272	9/29/2012 11:53:00 AM	-	-	-
NGHS MET Team Alert BP	1134199	0068372762270	9/29/2012 10:28:31 AM	-	-	-
NGHS MET Team Alert BP	1134160	0068372762270	9/29/2012 10:17:10 AM	-	-	-
NGHS MET Team Alert O2...	1134178	0001610322268	9/29/2012 10:17:05 AM	-	-	-
NGHS MET Team Alert HR	1134057	0067481962264	9/29/2012 9:10:06 AM	-	-	-
NGHS MET Team Alert BP	1134037	0068372762270	9/29/2012 8:41:37 AM	-	-	-
NGHS MET Team Alert BP	1134017	0068376282272	9/29/2012 8:35:47 AM	-	-	-
NGHS MET Team Alert HR	1134001	0067481962264	9/29/2012 8:12:47 AM	-	-	-
NGHS MET Team Alert O2...	1133999	0000430282271	9/29/2012 8:05:15 AM	-	-	-
NGHS MET Team Alert BP	1133998	0068376282272	9/29/2012 8:05:03 AM	-	-	-
NGHS MET Team Alert BP	1133997	0068376282272	9/29/2012 8:05:02 AM	-	-	-
NGHS MET Team Alert Resp	1133978	0068346332249	9/29/2012 8:00:12 AM	9/29/2012 3:58:00 PM	-	-
NGHS MET Team Alert BP	1133957	0068376282272	9/29/2012 7:31:49 AM	-	-	-
NGHS MET Team Alert Temp	1133839	0060393282264	9/29/2012 6:45:20 AM	-	-	-
NGHS MET Team Alert O2...	1133717	0060119662268	9/29/2012 5:32:18 AM	-	-	-
NGHS MET Team Alert HR	1133698	0064458022272	9/29/2012 5:15:02 AM	-	-	-
NGHS MET Team Alert BP	1133620	0061983082272	9/29/2012 4:25:58 AM	-	-	-
NGHS MET Team Alert Resp	1133619	0068346332249	9/29/2012 4:25:33 AM	9/29/2012 3:58:00 PM	-	-

Examples of the types of interventions by the nursing supervisors can be found in the “Interventions” sheet, shown below. As implied earlier, the documentation of interventions by our nursing supervisors is a relatively recent development and the documentation is often incomplete as our nursing supervisors adjust to a new workflow. The potential for interesting analytics is obvious, although further analysis of this type of data will be contingent on the nursing service using a unique patient identifier. And it will be facilitated by the nursing service using

structured data entries in addition to free text. For example, simple 0/1 (no/yes) entries could be used to stipulate that the patient's physician was notified of a change in status or that supplemental oxygen was administered or that anti-hypertensive medications were withhold in the case of a low blood pressure alert. Anecdotally, the nursing supervisors report that on occasion the alert triggers a conversation with the family that leads to a change in resuscitation status. This, too, could be recorded as a simple binary variable. Creating some structured data entries in the intervention record will enhance the business intelligence capabilities of the application.

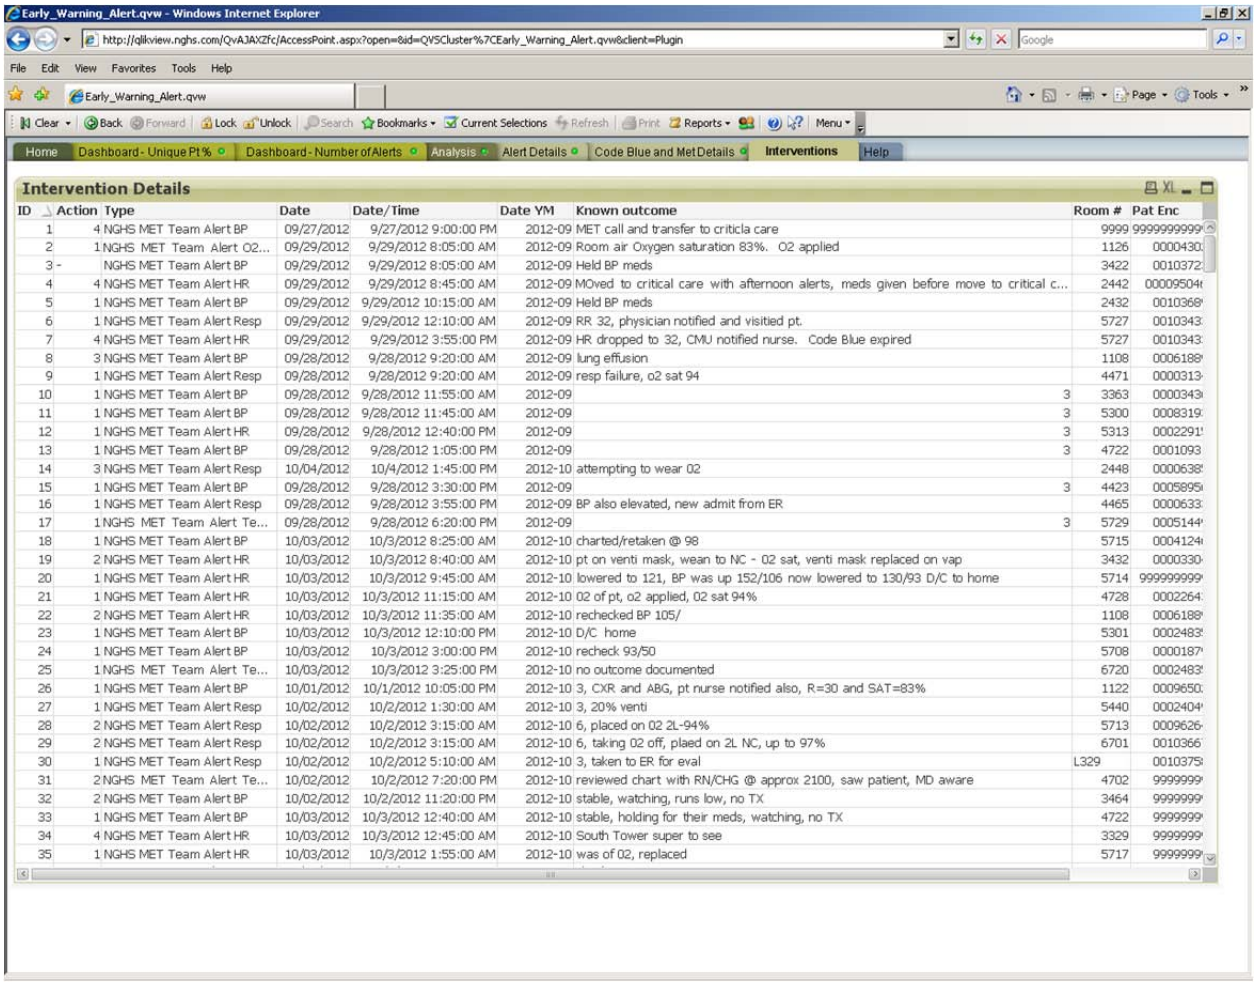

ID	Action	Type	Date	Date/Time	Date YM	Known outcome	Room #	Pat Enc
1	4	NGHS MET Team Alert BP	09/27/2012	9/27/2012 9:00:00 PM	2012-09	MET call and transfer to criticla care	9999	9999999999
2	1	NGHS MET Team Alert O2...	09/29/2012	9/29/2012 8:05:00 AM	2012-09	Room air Oxygen saturation 83%. O2 applied	1126	0000430
3	-	NGHS MET Team Alert BP	09/29/2012	9/29/2012 8:05:00 AM	2012-09	Held BP meds	3422	0010372
4	4	NGHS MET Team Alert HR	09/29/2012	9/29/2012 8:45:00 AM	2012-09	MOved to critical care with afternoon alerts, meds given before move to critical c...	2442	00009504
5	1	NGHS MET Team Alert BP	09/29/2012	9/29/2012 10:15:00 AM	2012-09	Held BP meds	2432	0010368
6	1	NGHS MET Team Alert Resp	09/29/2012	9/29/2012 12:10:00 AM	2012-09	RR 32, physician notified and visitied pt.	5727	0010343
7	4	NGHS MET Team Alert HR	09/29/2012	9/29/2012 3:55:00 PM	2012-09	HR dropped to 32, CMU notified nurse. Code Blue expired	5727	0010343
8	3	NGHS MET Team Alert BP	09/28/2012	9/28/2012 9:20:00 AM	2012-09	lung effusion	1108	0006188
9	1	NGHS MET Team Alert Resp	09/28/2012	9/28/2012 9:20:00 AM	2012-09	resp failure, o2 sat 94	4471	0000313
10	1	NGHS MET Team Alert BP	09/28/2012	9/28/2012 11:55:00 AM	2012-09	3	3363	0000343
11	1	NGHS MET Team Alert BP	09/28/2012	9/28/2012 11:45:00 AM	2012-09	3	5300	0008319
12	1	NGHS MET Team Alert HR	09/28/2012	9/28/2012 12:40:00 PM	2012-09	3	5313	0002291
13	1	NGHS MET Team Alert BP	09/28/2012	9/28/2012 1:05:00 PM	2012-09	3	4722	0001093
14	3	NGHS MET Team Alert Resp	10/04/2012	10/4/2012 1:45:00 PM	2012-10	attempting to wear 02	2448	0000638
15	1	NGHS MET Team Alert BP	09/28/2012	9/28/2012 3:30:00 PM	2012-09	3	4423	0005895
16	1	NGHS MET Team Alert Resp	09/28/2012	9/28/2012 3:55:00 PM	2012-09	BP also elevated, new admit from ER	4465	0000633
17	1	NGHS MET Team Alert Te...	09/28/2012	9/28/2012 6:20:00 PM	2012-09	3	5729	0005144
18	1	NGHS MET Team Alert BP	10/03/2012	10/3/2012 8:25:00 AM	2012-10	charted/retaken @ 98	5715	0004124
19	2	NGHS MET Team Alert HR	10/03/2012	10/3/2012 8:40:00 AM	2012-10	pt on venti mask, wean to NC - 02 sat, venti mask replaced on vap	3432	0000330
20	1	NGHS MET Team Alert HR	10/03/2012	10/3/2012 9:45:00 AM	2012-10	lowered to 121, BP was up 152/106 now lowered to 130/93 D/C to home	5714	999999999
21	1	NGHS MET Team Alert HR	10/03/2012	10/3/2012 11:15:00 AM	2012-10	02 of pt, o2 applied, 02 sat 94%	4728	0002264
22	2	NGHS MET Team Alert HR	10/03/2012	10/3/2012 11:35:00 AM	2012-10	rechecked BP 105/	1108	0006188
23	1	NGHS MET Team Alert BP	10/03/2012	10/3/2012 12:10:00 PM	2012-10	D/C home	5301	0002483
24	1	NGHS MET Team Alert BP	10/03/2012	10/3/2012 3:00:00 PM	2012-10	recheck 93/50	5708	0000187
25	1	NGHS MET Team Alert Te...	10/03/2012	10/3/2012 3:25:00 PM	2012-10	no outcome documented	6720	0002483
26	1	NGHS MET Team Alert BP	10/01/2012	10/1/2012 10:05:00 PM	2012-10	3, CXR and ABG, pt nurse notified also, R=30 and SAT=83%	1122	0009650
27	1	NGHS MET Team Alert Resp	10/02/2012	10/2/2012 1:30:00 AM	2012-10	3, 20% venti	5440	0002404
28	2	NGHS MET Team Alert Resp	10/02/2012	10/2/2012 3:15:00 AM	2012-10	6, placed on 02 2L-94%	5713	0009626
29	2	NGHS MET Team Alert Resp	10/02/2012	10/2/2012 3:15:00 AM	2012-10	6, taking 02 off, plaed on 2L NC, up to 97%	6701	0010366
30	1	NGHS MET Team Alert Resp	10/02/2012	10/2/2012 5:10:00 AM	2012-10	3, taken to ER for eval	L329	0010375
31	2	NGHS MET Team Alert Te...	10/02/2012	10/2/2012 7:20:00 PM	2012-10	reviewed chart with RN/CHG @ approx 2100, saw patient, MD aware	4702	9999999
32	2	NGHS MET Team Alert BP	10/02/2012	10/2/2012 11:20:00 PM	2012-10	stable, watching, runs low, no TX	3464	9999999
33	1	NGHS MET Team Alert BP	10/03/2012	10/3/2012 12:40:00 AM	2012-10	stable, holding for their meds, watching, no TX	4722	9999999
34	4	NGHS MET Team Alert HR	10/03/2012	10/3/2012 12:45:00 AM	2012-10	South Tower super to see	3329	9999999
35	1	NGHS MET Team Alert HR	10/03/2012	10/3/2012 1:55:00 AM	2012-10	was of 02, replaced	5717	9999999

DISCUSSION:

A relational database is a two –dimensional object with rows (“records”) and columns (“fields”) with some column or combination of columns that uniquely identifies each row. Each field will contain information relevant to that record. For example, a company may want to keep sales records in which each row is uniquely identified by an invoice number in the first field with subsequent fields containing the salesperson identification, the product identification, the units of product sold, the total price of the sale, and the store location identification. Relational databases have been and still are extensively used for storing data. However, enterprises often want the ability to aggregate and summarize the data in order to achieve the goals of the organization. This endeavor is referred to as “business intelligence”. Using the simple example described above, the company may want to analyze their sales to understand which products are popular, which are profitable, which sales people are most productive, what is the relationship between store location and items sold, etc. While the relational database contains the information, it is not an efficient platform for this type of analysis. This has led to the development of business intelligence tools and alternate models for storing data, particularly in so-called data warehouses.

The data model most often used for business intelligence is a multi-dimensional representation rather than the two-dimensional representation of a relational database. Multi-dimensional representation simplifies data storage. For example, suppose a company records sales by product and location and that there are 5 store locations and 4 distinct products. This would require a two-dimensional table with 3 columns recording the product, location, and total sales, and 20 rows for the 20 permutations of location and products. Now suppose the company wished to follow sales over time by day. To do this in a two-dimensional table they would have to add a 4^{th} column (time) and then add 7300 (20 product x location permutations times 365 days) more rows. Data storage in this case is more readily done by simply adding a time dimension to the two-dimensional table (42). This multi-dimensional representation is often referred to as a “data cube” and is usually implemented as a star schema, with one central table, the fact table, at the center of the star linked to dimension tables. Using the same simple example, our company may want to store data along the dimensions of sales people, product, and store. An entity relationship diagram for this simple example is shown below

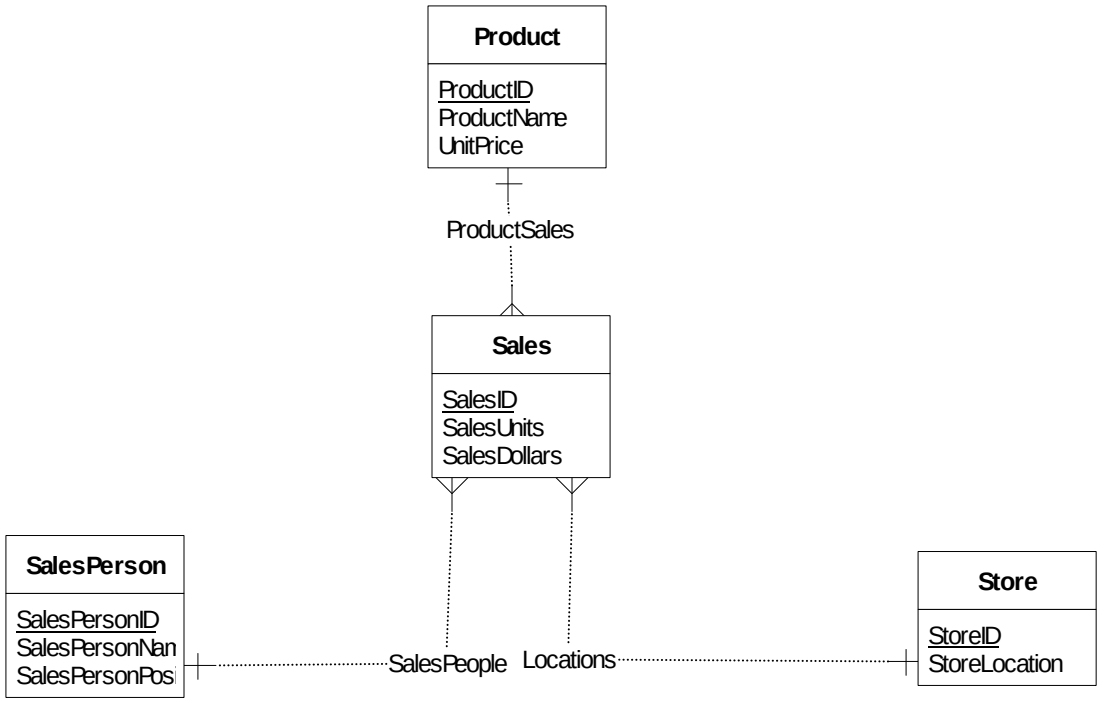

In this simple example, there are three dimensions and the data model could be viewed as a cube. This geometric concept may be extended; for example, in many cases businesses want to also analyze performance along a fourth (time) dimension. In the simple case above this would require 4 dimensions and we can view the data model as a hypercube. This type of multi-dimensional is a usable interface for the business analyst and there are multiple operations that can be performed (42) for data analysis. Furthermore, this representation can improve retrieval speed when data is stored on a disk. However, creation of the “fact” table requires discrete queries. For complex multi-dimensional data cubes, the effort to retrieve relevant facts requires a large time investment by information technology professionals. It becomes almost impossible to query for all possible associations between data elements and isolated queries lead to a loss of context between one query and another.

The business intelligence tool we used for this project, QlikView, purports to avoid these limitations. QlikView stores data in memory in a compressed form with associations defined between data items rather than joins as used in traditional databases. These associations are derived automatically by Qlikview during the

data load process based on matching column names across tables. Even though QlikView touts its "unique patented in-memory associative technology", the underlying data structures are not associative at all but rather the data is stored in a regular tabular format. QlikView has a proprietary compression algorithm that facilitates storage of this tabular format in memory, i.e., not on disk. Furthermore, the associative model speeds queries. For example, the zip codes for clients would be stored in a table of unique zip codes rather than as a zip code in each client table. This reduction in redundant information creates efficiency. In general, the use of associated tables in contrast to table joins preserves normalization and reduces redundancy, contributing to efficiency. The tabular format permits direct and indirect searches and great flexibility associating data fields. Different tables are connected by fields with the exact same name.

The construction of this particular QlikView application was complicated by the fact that a large number of individual patients had multiple alerts occurring at different times. This precluded our simply joining multiple tables in a star/snowflake schema. Consequently, we simply added data on adverse events (cardiopulmonary arrest, MET call, or ICU transfer) to the table comprised of the alert information, rather than attempting a join operation. This simplified the table structure for the application, but this table structure begs the question of how to organize the display to the user. We chose to include in our application one sheet that simply provided alert-number based statistics and another sheet that provided unique-patient statistics.

At this early date in the development of this business intelligence tool, the database is relatively sparse and drawing conclusions must be done with caution. Nevertheless, it does appear that for unique patients, the heart rate alert and the respiratory rate alert seem to be most often associated with subsequent adverse events. The more interesting question is which combination of alerts is most predictive of subsequent adverse event. As noted above we have not yet accumulated enough data to provide a definite answer at this point, but we have taken the data for October 2012 and have found that the adverse rate when patients have two or more different alerts (for example, temperature and blood pressure) is 13.3% (6/45) compared to 10.1% (49/483) for patients who only have one type of alert. This type of analysis can be extended to focus on specific pairwise alert sets

and then more extensive combinations. We eventually hope to identify an alert combination that is more specific, i.e., fewer false positives.

When we analyze our alerts from the unique patient perspective, we see that approximately 5-10% of the alerts had a subsequent adverse event (cardiopulmonary arrest, MET call, or ICU transfer). Thus, we might conclude that the alert is overly sensitive. However, approximately 50% of the unique patients had some sort of intervention by the nursing supervisor. We would like to conclude that the interventions had some effect on the frequency of subsequent adverse events. At this juncture we cannot do this as our nursing supervisors have been incomplete in recording a unique patient identifier on their intervention report. However, this is being rectified and soon we hope to be able to compare adverse event rates in patients who had an alert with an intervention to patients who had an event without an intervention.

When analyzed from the perspective of the number of alerts, our preliminary data indicates the criterion established at Northeast Georgia Medical Center for an early warning system alert (shown below) is quite non-specific.

Heart Rate	< 40 beats per minute	> 130 beats per minute
Systolic Blood Pressure	< 90 mm Hg	
Respiratory Rate	< 8 breaths per minute	> 24 breaths per minute
Hemoglobin oxygen saturation (SpO2)	< 90%	
Temperature	> 100.4 degrees F	

We observed multiple alerts for individual patients with the total number of alerts far greater than the number of unique patients. The incidence of adverse events (code blue, MET activation, or transfer to the ICU) subsequent to the alert is low, although this could reflect interventions by the nursing supervisor. It appears that adverse events are most common for the heart rate and respiratory rate alert, and the latter is in line with current literature (and somewhat ironic since respiratory rate documentation is often poorly done) (43).

Using the data from October 2012 (our most reliable data) we can see that the total number of alerts was approximately 4 times greater than the number of unique patients, i.e., there were multiple alerts for individual patients. This was also true in the subset of patients who had a subsequent adverse event. Approximately half of the unique patients had some sort of an intervention, although only 27.8% of the individual alerts led to some sort of an intervention. This leads to the question of whether interventions only occurred after multiple alerts, implying an element of alert fatigue. The primary conclusion that we have reached using this QlikView application is that the current alert is non-specific and hence, vulnerable to alert fatigue.

Alert fatigue is a well-known phenomenon. In the clinical informatics literature, it has most often been discussed in the context of clinical decision support within an electronic health record (44-49). However, anyone who has spent time in the intensive care unit can readily observe alert fatigue as clinicians routinely ignore alarms. The available literature suggests that alarms that are 90% accurate are seldom ignored, and, in contrast, alarms that are 10% accurate are usually ignored (49). If we can extrapolate from this general observation, then given that only 6% of our total alerts have a subsequent adverse event, the risk of alert fatigue for the current early warning system alert is substantial. The sheer volume of alerts (close to 70 a day) alone is a factor that could promote alert fatigue and in the appendix we present a theoretical analysis of how volume of alert may impact defining an optimal alert threshold. This theoretical analysis presupposes a model for the probability of clinical deterioration that we have not established in our institution, but the theoretical model raises **the key question of whether alert fatigue is a simple function of the alert volume. If so, the theoretical model suggests that there will be an optimal threshold for the alert, reflecting the balance between the efficacy of interventions and alert fatigue.**

Given our preliminary observations, the next steps in further development of this business intelligence application are the following.

1. Mandating a specific and unique patient identifier for documentation of interventions undertaken in response to the alert. This is essential to determination of whether interventions are preventing subsequent adverse events.

2. Analysis of the patients who have interventions to determine if there is a common theme to the type or volume of alert that characterizes the patients who have an intervention vs. those who do not. An interesting question is whether interventions typically occur after a patient has had multiple alerts.
3. Analysis of the different types of alerts to learn which are most predictive of specific adverse events, code blue, MET activation, or ICU transfer.
4. Analysis of combinations of alerts to determine if specific combinations have better specificity.
5. Preliminary analysis of alert fatigue by determining whether there is a correlation between the volume of alerts and the frequency of interventions. We will analyze the data on a 12 hour shift basis and determine whether the number of interventions is a linear function of the alert volume.
6. Inclusion of laboratory values (white blood cell count and lactic acid are likely candidates) in the alert. This will require definition of the time frame of the laboratory value vis-à-vis the timing of the alert.
7. Inclusion of structured assessment of level of consciousness (LOC) or changes in LOC in the alert using an ordinal scale. For example, some of the literature has scored LOC using a simple 4 point scale, where 3 is alert, 2 is responsive to voice, 1 is responsive to pain only, and 0 is unresponsive. This will require consistent documentation of LOC on the part of the nursing staff.

REFERENCES

1. Sax, FL, Charlson ME. Medical patients at high risk for catastrophic deterioration. Crit Care Med 1987; 15: 510515.
2. Smith AF, Wood, J. Can some in-hospital cardio-pulmonary arrests be prevented? A prospective survey. Resuscitation 1998; 37: 133-137.
3. Schein RMH, Hazday N, Pena, M, Ruben BH, Spring CL. Clinical antecedents to in-hospital cardiopulmonary arrest. Chest 1990; 98: 1388-1392.
4. Franklin C, Mathew J. Developing strategies to prevent inhospital cardiac arrest:analyzing responses of physicians and nurses in the hours before the event. Crit Care Med 1994; 22: 244-247.)
5. McQuillan P, Pilkington S, Allan A, Taylor B, Short A, Morgan G, Nielson M, Barrett D, Smith G. Confidential inquiry into quality of care before admission to intensive care. Br Med J 1998; 316: 1853-1858.
6. Lee A, Bishop G, Hillman KM, Daffurn K. The medical emergency team. Anaesth Intens Care 1995; 23: 183-186.
7. Morgan RJM, Williams F, Wright MM. An early warning scoring system for detecting critical illness. Clin Intens Care 1997; 8:100.
8. Stenhouse C, Coates M, Tivey M, Allsop P, Parker T. Prospective evaluation of a modified Early Warning System to aid earlier detection of patients developing critical illness on a general surgical ward. Br J Anaesth 2000; 84: 663p
9. Subbe C, Kruger M, Rutherford P, Gemmel L. Validation of a modified Early Warning System in medical admissions. Q J Med 2001; 94: 521-526.
10. Duckitt RW, Buxton-Thomas R, Walker J, Cheek E, Bewick V, Venn R, Forni LG. Worthing physiological scoring system:derivation and validation of a physiological early-warning system for medical admissions. An observational, population-based single-centre study. Br J Anaesth 2007; 98: 769-774.
11. Bleyer AJ, Vidya S, Russell GB, Jones CM, Sujata L, Daeihagh P, Hire D. Longitudinal analysis of one million vital signs in an academic medical center. Resuscitation 2011; 82: 1387-1392.

12. McArthur-Rouse F. Critical care outreach services and early warning systems: a review of the literature. J Advanced Nursing 2001; 36: 696-704.
13. Naeem N, Montenegro H. Beyond the intensive care unit: A review of interventions aimed at anticipating and preventing in-hospital cardiopulmonary arrest. Resuscitation 2005; 67: 13-23.
14. Gao H, McDonnell A, Harrison DA, Moore T, Adam S, Daly K, Esmonde L, Goldhill DR, Parry GJ, Rashidian A, Subbe CP, Harvey S. Systematic review and evaluation of physiological track and trigger warning systems for identifying at-risk patients on the ward. Intensive Care Med 2007; 33: 667-679.
15. Smith GB, Prytherch DR, Schmidt PE, Featherstone PI, Higgins, B. A review, and performance evaluation, of single –parameter "track and trigger" systems. Resuscitation 2008; 79: 11-21.
16. Smith GB, Prytherch DR, Schmidt PE, Featherstone PI. Review and performance evaluation of aggregate weighted 'track and trigger' systems. Resuscitation 2008; 77: 170-179.
17. Jansen JO, Cuthbertson BH. Detecting critical illness outside the ICU: the role of track and trigger systems. Curr Opin Crit Care 2010; 16: 184-190.
18. Kyriacos U, Jelsma J, Jordan S. Monitoring vital signs using early warning scoring systems: a review of the literature. J Nursing Management 2011; 19: 311-330.
19. DeVita MA, Smith GB, Adam SK, Adams-Pizarro I, Buist M, Bellomo R, Bonello R, Cerchiari E, Farlow B, Goldsmith D, Haskell H, Hillman K, Howell M, Hravnak M, Hunt EA, Hvarfner A, Kellett J, Lighthall GK, Lippert A, Lippert F, Mahroof R, Myers JS, Rosen M, Reynolds S, Rotondi A, Rubulotta F, Winters B. Identifying the hospitalized patient in crisis-A consensus conference on the afferent limb of Rapid Response Systems. Resuscitation 2010; 81: 375-382.
20. Bellomo R, Goldsmith D, Uchino S, Buckmaster J. Hart GK. Opdam H. Silvester W. Doolan L. Gutteridge G.. A prospective before-and-after trial of a medical emergency team. Med J Aust 2003; 179: 283-287.
21. Jones D, Bates S, Warrillow S, Goldsmith D. Hart G. Opdam H. Gutteridge G. Long term effects of a medical emergency team on cardiac arrests in a teaching hospital. Crit Care 2005; 9: R808-R815.

22. Sebat F, Musthafa AA, Johnson D, Kramer AA. Shoffner D. Eliason M. Henry K. Spurlock B.. Effect of a rapid response team for patients in shock on time to treatment and mortality during 5 years. Crit Care Med 2007; 35: 2568-2575.
23. Chan PS, Khalid A, Longmore LA, Berg RA, Kosiborod M, Spertus JA. Hospital-wide code rates and mortality and after implementation of a rapid response team. JAMA 2008; 300:2506-2513.
24. Baxter AD, Cardianl P, Hooper J, Patel R. Medical emergency teams at The Ottawa Hospital: the first two years. Can J Anaesth 2008; 55: 223-231.
25. Dacey MJ, Mirza ER, Wilcox V, Doherty M. Mello J. Boyer A. Gates J. Brothers T. Baute R.The effect of a rapid response team on major clinical outcome measures in a community hospital. Crit Care Med 2007; 35: 2076-2082.
26. Bristow PJ, Hillman KM, Chey T, Daffurn K. Jacques TC. Norman SL. Bishop GF. Simmons EG.. Rates of in-hospital arrest, deaths, and intensive care admissions: the effect of a medical emergency team. Med J Aust 2000; 173: 236-240.
27. Buist MD, Moore GE, Bernard SA, Waxman BP, Anderson JN, Nguyen TV. Effects of a medical emergency on reduction of incidence of and mortality from unexpected cardiac arrest in hospital: preliminary study. BMJ 2002; 324: 387-390.
28. DeVita MA, Braithwaite RS, Mahidhara R, Stuart S, Foraida RL, Simmons RL. Medical Emergency Response Improvement Team (MERIT). Use of medical emergency team responses to reduce hospital cardiopulmonary arrest. Qual Saf Health Care 2004; 13: 251-254.
29. Jones D, Opdam H, Egi M, Goldsmith D. Bates S. Gutteridge G. Kattula A. Bellomo R. Long-term effect of a Medical Emergency Team on mortality in a teaching hospital. Resuscitation 2007; 74: 235-241.
30. Kenward G, Castle N, Hodgetts T, Shaikh L. Evaluation of a medical emergency team one year after implementation. Resuscitation 2004; 61: 257-263.
31. Mitchell LA, McKay H, Van Leuvan C, Berry R, McCutcheon C, Avard B, Slater N, Neeman T, Lamberth P. A prospective controlled trial of the effect of a multi-faceted intervention on early recognition and intervention in deteriorating hospital patients. Resuscitation 2010; 81: 658-666.

32. Moon A, Cosgrove JF, Lea D, Fairs A, Cressey DM. An eight year audit before and after the introduction of modified early warning score (MEWS) charts, of patients admitted to a tertiary referral intensive care unit after CPR. Resuscitation 2011; 82: 150-154.
33. Paterson R, MacLeoad DC, Thetford D, Beattie A, Graham C, Lam S, Bell D. Prediction of in-hospital mortality and length of stay using an early warning scoring system:clinical audit. Clinical Medicine 2006; 6: 281-284.
34. Green AL, Williams A. An evaluation of an early warning clinical marker referral tool. Intensive and Critical Care Nursing 2006; 22: 274-282.
35. MERIT study investigators. Introduction of the medical emergency team (MET) system: a cluster randomized controlled trial. The Lancet 2005; 365: 2091-2097
36. Priestly G, Watson W, Rashidian A, Mozley C, Russell D, Wilson J, Cope J, Hart D, Kay D, Cowley K, Pateraki J. Introducing Critical Care Outreach: a ward-randomized trial of phased introduction in a general hospital. Intensive Care Medicine 2004; 30: 1398-1404.
37. Chan PS, Jain R, Nallmothu B, Berg RA, Sasson C. Rapid response teams. A systematic review and meta-analysis. Arch Int Med 2010; 170: 18-26.
38. http://www.qlikview.com/us/explore/resources/whitepapers/what-makes-qlikview-unique
39. http://www.qlikview.com/us/explore/resources/whitepapers/qlikview-architectural-overview
40. Mannino MV. Database Design, Application, Development, and Administration, 3rd Ed. Boston: McGraw-Hill, 2007, pp. 567-570.
41. Mannino MV. Database Design, Application, Development, and Administration, 3rd Ed. Boston: McGraw-Hill, 2007, pp. 59-61.
42. Mannino MV. Database Design, Application, Development, and Administration, 3rd Ed. Boston: McGraw-Hill, 2007, pp. 560-566.
43. Hogan J. Why don't nurses monitor the respiratory rates of patients? Br J Nursing 2006; 15: 489-492.
44. Isaac, T., Weissman, J.S., 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.
45. Baker, DE. Medication alert fatigue: The potential for compromised patient safety. Hospital Pharmacy 2009; 44: 460-462.

46. Weingart, S, Toth M., Sands DZ, Aronson MD, Davis RB, Phillips RS. Physicians' decisions to override computerized drug alerts in primary care. Archives of Internal Medicine 2003; 163: 2625-2631.
47. Van der Sijs, H., Aarts, J. Vulto A, Berg M. Overriding of drug safety alerts in computerized physician order entry. Journal of the American Medical Informatics Association 2006; 13: 138-147.
48. Kuperman, GJ, Reichley, RM and Bailey, TC. Using commercial knowledge bases for clinical decision support: Opportunities, hurdles, and recommendations. Journal of the American Medical Informatics Association 2006; 13: 369-371.
49. Paterno, M.D., Maviglia, S.M., Gorman PN, Seger DL, Yoshida E, Seger AC, Dates DW, Gandhi TK. Tiering drug–drug interaction alerts by severity increases compliance rates. Journal of the American Medical Informatics Association 2009; 16: 40-46.
50. Bliss JP, Dunn MC. Behavorial implications of alarm mistrust as a function of task workload. Ergonomics 2000; 43: 1283-1300.

APPENDIX

BACKGROUND

As noted earlier, interest in early warning for clinical deterioration was the impetus for the medical emergency team (MET) concept, introduced in 1990 at Liverpool Hospital (6). The original implementation used defined physiological parameters as calling criteria. These included

Temp < 35.3 or > 39.5

SBP <100 or > 200

RR<10 or >30

Urine output over 24 hrs < 500 ml

Altered LOC

These MET criteria were the first early warning or track and trigger scoring systems. Since then there have been multiple modifications of the criterion for the early alert of potential clinical deterioration. However, with two exceptions, the criteria for activation of an alert have been based on expert opinion and validation of the alert has relied on the demonstration that the alert discriminates between patients who subsequently deteriorate (typically as defined by cardio-pulmonary arrest or the need for transfer to the intensive care unit) and those who do not deteriorate. The exceptions to this methodology are the studies by Duckitt et al (10) and Bleyer et al (11) who used multi-variate logistic regression to identity predictors of clinical deterioration. The advantage of this approach is not only that it uses patient data, in contrast to expert opinion, to define the alert, but also it facilitates a probabilistic model for clinical deterioration.

THEORY

Assume that we have an alert that predicts the probability of clinical deterioration, denoted P where P is a function of the observed variables (V), for example vital signs, select laboratory results, etc. Further assume that we trigger the alert when

$P > T$

where T is the alert threshold (as T increases the alert is more and more selective). We also assume that if the alert is triggered an intervention by clinicians may occur that changes the probability of deterioration to P_I where P_I is the conditional probability of deterioration given that an alert has been triggered [P(deterioration|alert)] and we assume it is a function of the probability of deterioration without intervention, i.e., $P_I = P_I(P)$. We will denote the number of patients who at any time have a probability of deterioration between P and P+dP as N(P). Thus

$\int_0^1 N(P)dP = C$ where C is the hospital census at this time

At any one time the census will be comprised of patients with varying probabilities of deterioration. The <u>mean</u> probability of deterioration across all patients assuming that we have not implemented an alert will be denoted as $\langle P \rangle_{\text{no alert}}$ and is given by

$$\langle P \rangle_{\text{no alert}} = \left(\frac{1}{C}\right) \int_0^1 PN(P)dP = \left(\frac{1}{C}\right) \int_0^T PN(P)dP + \left(\frac{1}{C}\right) \int_T^1 PN(P)dP$$

If we now implement an alert with alert threshold T, the <u>mean</u> probability of deterioration across all patients (denoted $\langle P \rangle_{\text{alert}}$) will be given by

$$\langle P \rangle_{\text{alert}} = \left(\frac{1}{C}\right) \int_0^T PN(P)dP + \left(\frac{1}{C}\right) \int_T^1 P_I N(P)dP$$

Thus with the alert implemented the change in the mean probability of deterioration is

$\Delta\langle P\rangle = \langle P\rangle_{\text{no alert}} - \langle P\rangle_{\text{alert}} = (\frac{1}{C}) \int_T^1 \Delta P(P) N(P) dP$

We will refer to $C\Delta\langle P\rangle$ somewhat loosely as the number of patients who will be "rescued" by the alert, N_R and note that

$N_R = \int_T^1 \Delta P(P) N(P) dP \qquad (1)$

where again N(P) is the number of patients with a probability of deterioration equal to P and $\Delta P = P - P_I(P)$. Note that the maximum value of $\Delta P(P) = P$. Since we assume that $P \geq P_I$ (P) ΔP is always positive and equation (1) implies that in order to minimize the number of patients who deteriorate (or more correctly maximize the change in the mean probability of deterioration with alert implementation) we should set the threshold at 0, i.e., trigger the alert for all patients.

However, alert fatigue does occur and if the alert were triggered for all patients a certain proportion of the alerts would be ignored. We will denote the alert success rate as f (by success we mean the frequency at which the alert generates an intervention rather than being ignored) and consider two situations, the first where f is a function of P only and the second where f is a function of T only.

Considering the case where f = f(P)

$N_R = \int_T^1 f(P) \Delta P(P)\, N(P)\, dP \qquad (2)$

But again this equation implies (since $\Delta P \geq 0$) that the integrand is always positive and to maximize N_R we would set T equal to zero.

On the other hand, if f = f(T), i.e., the rescue rate is a function of the threshold only

$N_R = f(T) \int_T^1 \Delta P(P) N(P) dP \qquad (4)$

To identify the threshold that will maximize N_R a necessary (but not sufficient condition, which requires that the solution to equation (5) is unique which shall be assumed) is that the derivative of the above expression be equal to zero, resulting in the following identity.

$$\frac{df}{dT}\int_T^1 \Delta P(P)N(P)dP \; - f(T)\Delta P(P)_{P=T}N(T) = 0 \qquad (5)$$

Or

$$\frac{df}{dT}\int_T^1 \Delta P(P)N(P)dP = f(T)[T-P_I(T)]N(T) \qquad (6)$$

With further manipulation

$$T = (\frac{dlnf}{dT})\int_T^1 \Delta P(P)[\frac{N(P)}{N(T)}]dP + P_I(T) \qquad (7)$$

INFERENCES

The optimal alert threshold can be determined by numerical solution of equation (7) although qualitative solutions can be identified by noting that the optimal threshold is the intersection of the graphs of the line of identity Z =T and

$Z = (\frac{dlnf}{dT})\int_T^1 \Delta P(P)[\frac{N(P)}{N(T)}]dP + P_I(T)$ as a function of T.

We anticipate that $(\frac{dlnf}{dT})$ is a positive number, i.e., the alert is less likely to be ignored as the threshold increases and the alert is more selective. We infer that the optimal alert threshold will increase as $(\frac{dlnf}{dT})$ increases since this will lead to larger values of the function Z(T) with a increased intersection with the line of identity

Z =T.

We also observe that **<u>if</u>** $P_I(P) = \alpha P$ ($\alpha < 1$) so that the probability of deterioration with alert is simply a fraction of the probability of deterioration without the alert it is easy to show that

$$T = (\frac{dlnf}{dT})\int_T^1 P[\frac{N(P)}{N(T)}]dP \quad (8)$$

i.e., the optimal alert threshold is <u>independent</u> of α, i.e., the optimal alert threshold is the same whether the alert results in a large reduction of probability of deterioration (small α) or a small reduction in the probability of deterioration (large α).

We can carry this further by expanding ΔP(P) in a Taylor's series. To second order

$$\Delta P(P) = [K_0 + K_1P + K_2P^2](\Theta(P-K_0 - K_1P - K_2P^2)) + P(1- \Theta(P-K_0 - K_1P - K_2P^2))$$

where Θ is the Heaviside functions that serves to keep $\Delta P(P) \leq P$. Note also that

$K_0 + K_1P + K_2P^2$ must be greater than or equal to zero since $\Delta P(P)$ is assumed to be positive.

Note that this approximation permits a non-linear relationship between P and ΔP.

Substituting into equation (7) yields

$$T = \left(\frac{dlnf}{dT}\right)\int_T^1 [K_0 + K_1P + K_2P^2]\left[\frac{N(P)}{N(T)}\right](\Theta(P - K_0 - K_1P - K_2P^2))\,dP +$$

$$\left(\frac{dlnf}{dT}\right)\int_T^1 [1 - \Theta(P - K_0 - K_1P - K_2P^2)]P\left[\frac{N(P)}{N(T)}\right]\,dP + (T - K_0 - K_1T - K_2T^2)\,\Theta(T - K_0 - K_1T - K_2T^2) \quad (9)$$

In this case the optimal alert threshold depends on the efficacy of the intervention via K_0, K_1, and K_2, although a solution to equation (9) is not guaranteed.

Further inferences are realized by focusing on the term denoted above as Z

$$Z = \left(\frac{dlnf}{dT}\right)\int_T^1 \Delta P(P)\left[\frac{N(P)}{N(T)}\right]dP + P_I(T)$$

To make any inferences about this term we have to make assumptions about the quasi-distribution N(P). If N(P) were proportional to a normal distribution, there are a number of possible inferences. However, this is a strong assumption. A weaker assumption is to assume that <u>some function of P has a normal distribution</u>. For example, $\ln\left(\frac{P}{1-P}\right)$ might be normally distributed. To be more general, we will assume that the existence of a normally distributed function g with mean $\bar{g}$ and variance σ_g^2 which has an invertible transformation involving P, that is, g(P). It is also assumed that after substituting the new variable g, $\Delta P(P)N(P)$ is transformed

to a normal distribution function $h(g^{-1})$ NORM($\bar{g}$, σ_g^2)dg where $h(g^{-1})$ = h(P) = $[\frac{\Delta P(P)}{(\frac{dg}{dP})}]$. Thus,

$$Z = (\frac{dlnf}{dT})[\frac{1}{N(T)}]\int_{g(T)}^{g(1)} h(g^{-1})\text{NORM}(\bar{g}, \sigma_g^2)\, dg + P_I(T) \quad (10)$$

Where NORM denotes a normal distribution with mean $\bar{g}$ and variance σ_g^2 over the distribution of values of P. In general this type of transformation is useful only if g(1) = ∞ so we assume this and note

$$Z = (\frac{dlnf}{dT})[\frac{1}{N(T)}]\int_{f(T)}^{\infty} h(f)\text{NORM}(\bar{g}, \sigma_g^2)\, df + P_I(T)$$

We can express the above integral in terms of the standardized normal distribution with the change of variables u = $\frac{g-\bar{g}}{\sigma}$

$$Z = (\frac{dlnf}{dT})[\frac{\sigma}{N(T)}]\int_{\frac{g(T)-\bar{g}}{\sigma}}^{\infty} h(\sigma u + \bar{g})\text{NORM}_S(u)du + P_I(T)$$

There are certain inferences that one can reach from this formulation. If $h(\sigma u + \bar{g})$ is an increasing function of its argument and σ, the standard deviation of g increases as the standard deviation of N(P) increases [as would be the case if g(p)= $\ln(\frac{P}{1-P})$] then Z will increase with increasing standard deviation of the distribution of probabilities, as there is a direct factor of σ and also because the integral term, exclusive of the factor of $h(\sigma u + \bar{g})$, will increase with increasing σ. This implies that the optimal alert threshold will increase (the alert will be more selective). However, if $h(\sigma u + \bar{g})$ is not an increasing function of its argument, one cannot reach this conclusion.

SIMULATIONS

In order to confirm these qualitative predictions, we must make assumptions about the functions that describe N(P), $P_I(P)$, and f(T). While there is no literature to establish a quantitative model our goal will be to create tractable models that are clinically plausible.

We first consider N(P). Given that a rigorous approach for defining the criteria activating an alert is logistic regression, a plausible model for N(P) is to assume that $\ln\left(\frac{P}{1-P}\right)$ has a normal distribution. This is consistent with the use of multi-variate logistic regression used in the studies by Duckitt et al (10) and Bleyer et al (11) in which the model is

$$\ln\left(\frac{P}{1-P}\right) = A_0 + A_1V_1 + A_2V_{2+} \ldots$$

where A_i are coefficients and V_i are variables (systolic blood pressure, respiratory rate, etc) which we can typically assume are approximately normally distributed.

If we assume that $\ln\left(\frac{P}{1-P}\right)$ is normally distributed, then the distribution of P is given by

$$\left(\frac{1}{\sqrt{2\pi\sigma^2}}\right) \exp\left[-\frac{(g-\bar{g})^2}{2\sigma^2}\right]dg$$

where $g= \ln\left(\frac{P}{1-P}\right)$, $\bar{g}$ is the mean value of $\ln\left(\frac{P}{1-P}\right)$, and σ is the standard deviation. Using a standard change of variable

$$dg = \frac{dg}{dP}\, dP = \frac{dP}{P(1-P)} \text{ and}$$

$$N(P) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{(g(P)-\bar{g})^2}{2\sigma^2}\right] \frac{dP}{P(1-P)} \quad (10)$$

An example of this distribution is shown below as a graph of N(P) vs. P (Figure 1) using as parameters $\bar{g}$ = -1 and σ^2 = 1.

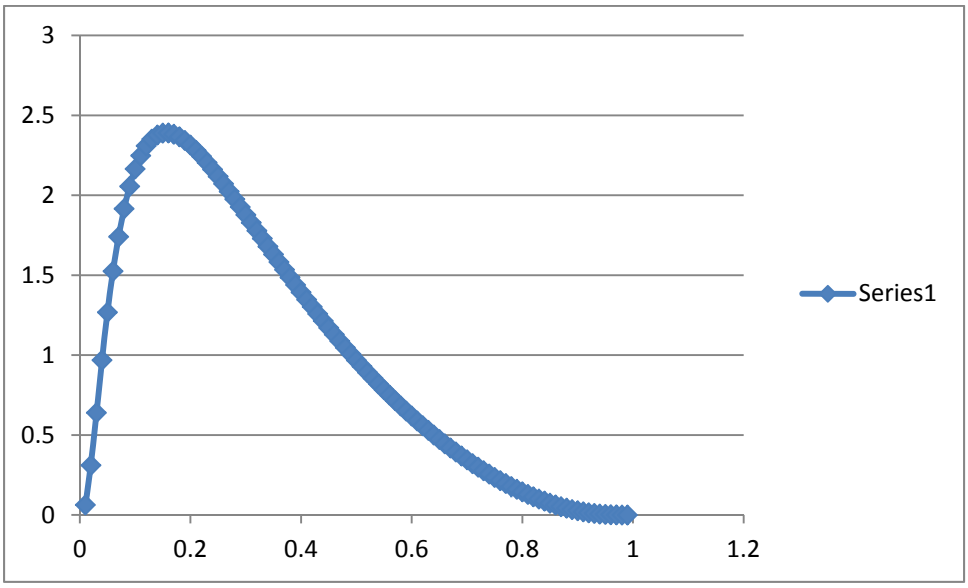

The distribution is clinically plausible, with a mode of P = 0.14 for a distribution that is skewed to the right but nevertheless indicating that patients with either a very low or very high probability of deterioration are uncommon.

To proceed further we need to postulate a functional form for ΔP(P). As indicated earlier if we assume that $P_I(P) = \alpha P$, a particularly simple equation (8) is obtained. To allow for a more complex function we can expand ΔP(P) in a Taylor's series, as illustrated to 2nd order. We note that if the probability of deterioration is low, ΔP(P) is limited by the fact that ΔP(P) ≤ P. Thus at low values of P, ΔP(P) is a linear function of P. Clinically, this is plausible since at low values of P we would expect to be able to "rescue" almost all patients. However, it is also clinically plausible to postulate that at very high values of P, i.e., when the patient has a very high probability of deterioration, our ability to rescue the patient may be limited. This

is illustrated in Figure 2, a plot of $\Delta P(P)$ vs. P, for a second order expansion of $\Delta P(P)$ with the following parameters, $K_0 = 0.2$, $K_1 = 0.5$, and $K_2 = -0.5$.

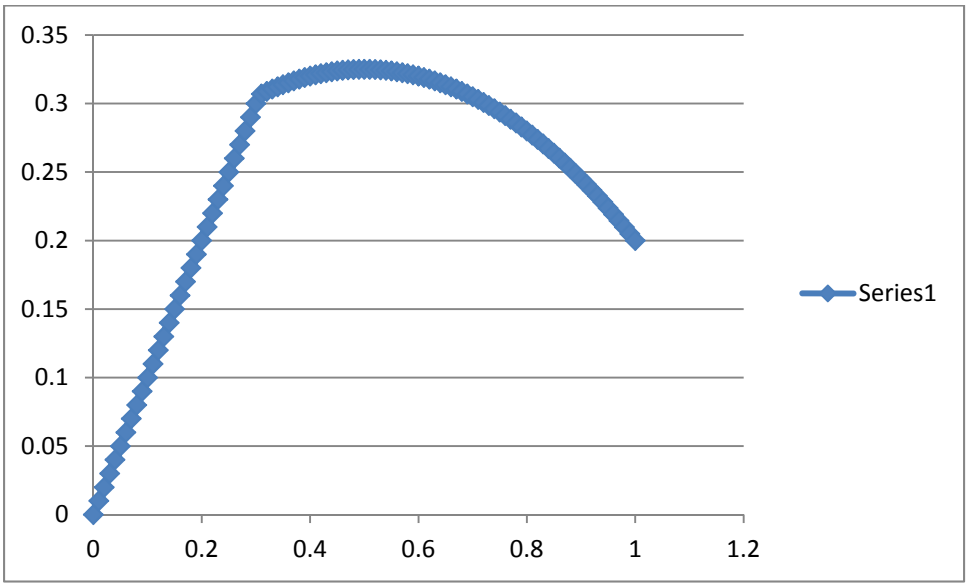

Finally we consider f(T). The functional form of f(T) is truly unknown and to move forward with simulations we must assume a function that is biologically plausible. We anticipate that the success of the alert success rate will depend on the threshold indirectly and that alert success is directly determined by the number of alerts that are triggered, which in turn is a function of the threshold. We anticipate that as the alert load increases the success rate will decrease. This could be captured in a simple linear relationship

$$F(T) = 1 - \beta C(T)$$

Where C(T) is the alert count, i.e.,

$$C(T) = \int_T^1 N(P)dP$$

where we must constrain β such that $f(T) \geq 0$. A more flexible formulation is to assume that

$f(T) = A \exp(-\beta C(T))$ (11)

Note that this implies a steady degradation in success as C(T) increases. An alternative model is to assume that the clinician has a finite capacity such that f(T) = 1 if C(T) is less than some critical value C_c and when the number of alerts exceeds this critical value the success rate is inversely related to C(T)

$$f(T) = \frac{C_c}{C(T)} \qquad (12)$$

With equation 11(and the assumption that A does not depend on T)

$$\frac{d\ lnf}{dT} = -\beta \frac{dC(T)}{dT} = \beta\ N(T)$$

And for the model embodied in equation 12

$$\frac{d\ lnf}{dT} = 0 \quad C(T) < C_c$$

$$\frac{d\ lnf}{dT} = \frac{N(T)}{C(T)} \quad C(T) \geq C_c$$

In order to confirm the inferences made earlier, simulations were performed assuming equation (10) as the model for N(P), both the simple linear model (equation 8) as well as the more complex second model (equation 9) to describe

$P_I(P)$ and $\Delta P(P)$, and both the monotonic and threshold models (equations 11 and 12) for f(T). Note that this leads to models that are parameterized by $\bar{g}$, σ^2 (describing the quasi-distribution N(P)), K_0, K_1, and K_2 (for the model for $\Delta P(P)$ and recalling that the linear model proves to be parameter-free) and β and C_c from the two models for f(T).

Integrals were evaluated using the trapezoidal method of numerical integration with a step size of 0.01with an Excel spreadsheet. The solutions to equation 8 or 9 were first approximated from the intersection of the graphs of these equations with the line of identity T=T and then confirmed using the Excel "Solver" function.

Initial simulations were begun using equation 11 for the function f(T). In the figure below we demonstrate the effect of changing $\bar{f}$ from -4 to 0 on the optimal threshold while assuming that $\sigma^2 = 1$, $K_0 = 0.2$, $K_1 = 0.5$, $K_2 = -0.5$, and $\beta = 2$. Series 1 refers to the results for equation 9 and series 2 to equation 8 (these are plots of optimal threshold as a function of $\bar{f}$). These simulations confirm the prediction that as $\bar{f}$ increases, the optimal threshold also increases (the alert becomes more selective).

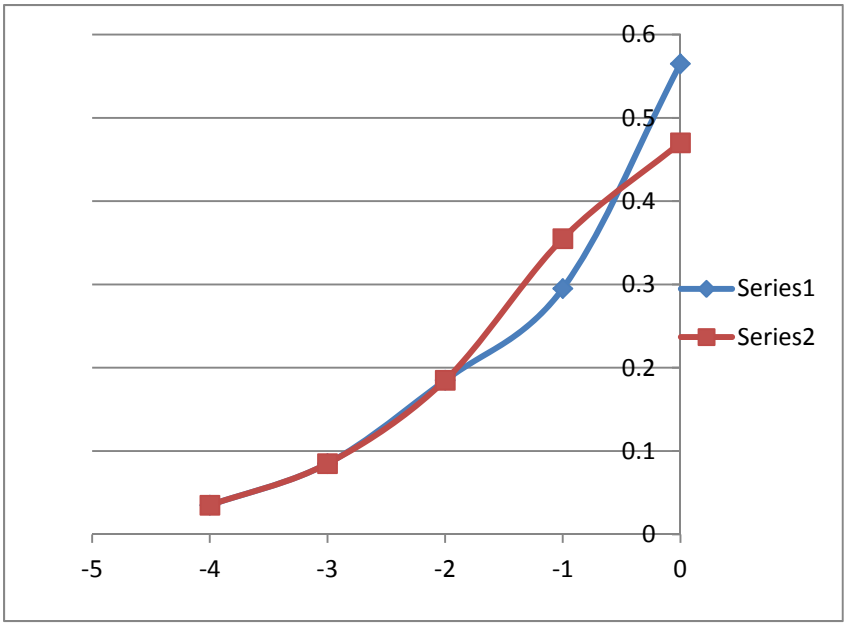

Next simulations were performed for verifying values of σ2 assuming that $\bar{f}$ = -2 and that K_0 = 0.2, K_1 = 0.5, K_2 = -0.5, and β = 2. Again Series 1 refers to the results for equation 9 (second order model for ΔP(P)) and series 2 to equation 8 (linear model for ΔP(P)).

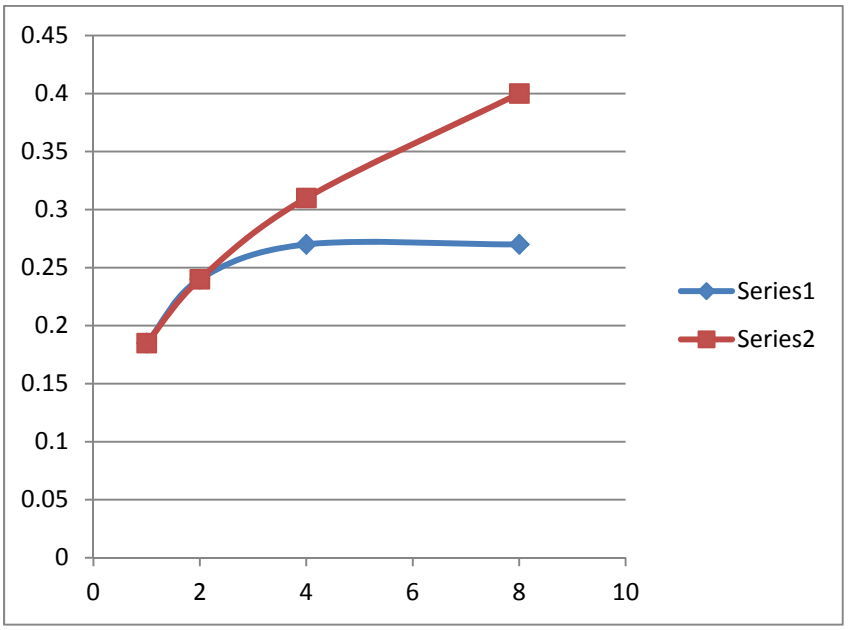

The figure confirms that as the standard deviation of N(P) increases, the optimal alert threshold increases. It is noteworthy that this effect is not as pronounced for the second order model for ΔP(P).

We next consider the effect of varying the success rate of the alert using the model embodied in equation 11. We present below a graph of optimal threshold vs β assuming that $\bar{f}$ = -2, σ^2 = 2, and K_0 = 0.2, K_1 = 0.5, K_2 = -0.5.

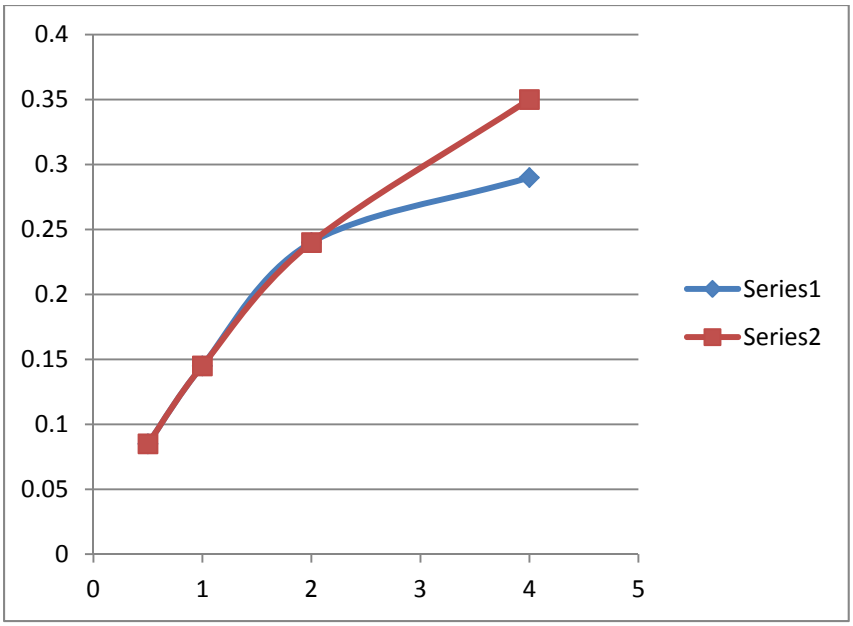

We see that as the relationship between alert success rate and threshold becomes steeper (as β increases) the optimal alert threshold also increases. Alert fatigue can be encapsulated by β and this just reflects the obvious fact that as alert fatigue increases, the optimal threshold increases (the alert needs to be more selective).

We repeated the same set of simulations using equation 12 as the model for f(T). Interestingly, this model, which has a discontinuity, implies that the optimal alert threshold is simply the critical value C_c. In the following two figures we illustrate the graphical determination of the optimal alert threshold as the intersection of equation 8 (series 1) or equation 9 (series 3) with the line of identity (series 2) for $C_c = 0.3$ (the critical value is 30% of total patients within the alert time frame).

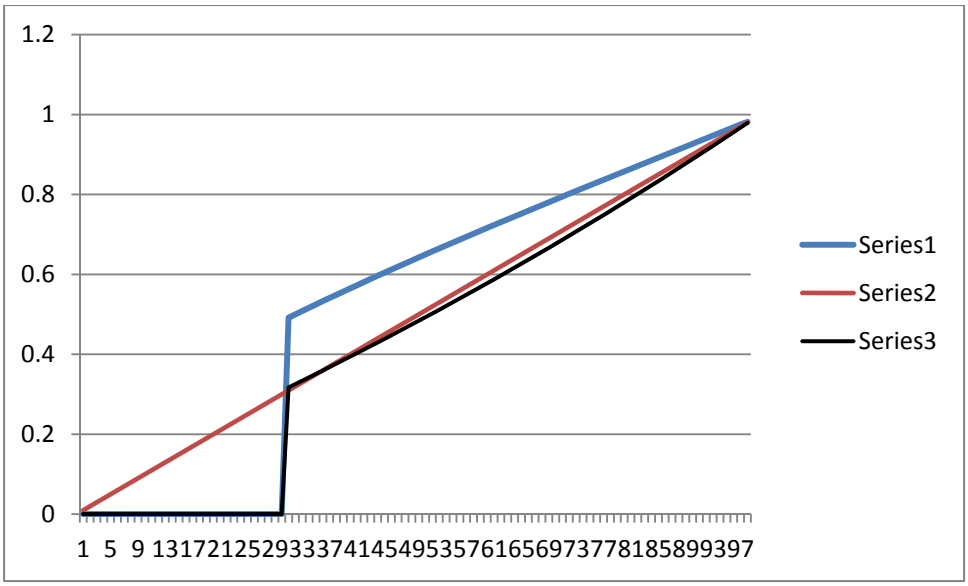

It can be seen that with the second order model for ΔP(P) the line defined by equation 10 appears to be nearly coincidental with the line of identity in contrast to the line defined by the simple linear model (equation 8). When the same simulation is done with $C_c = 0.5$ we find the following

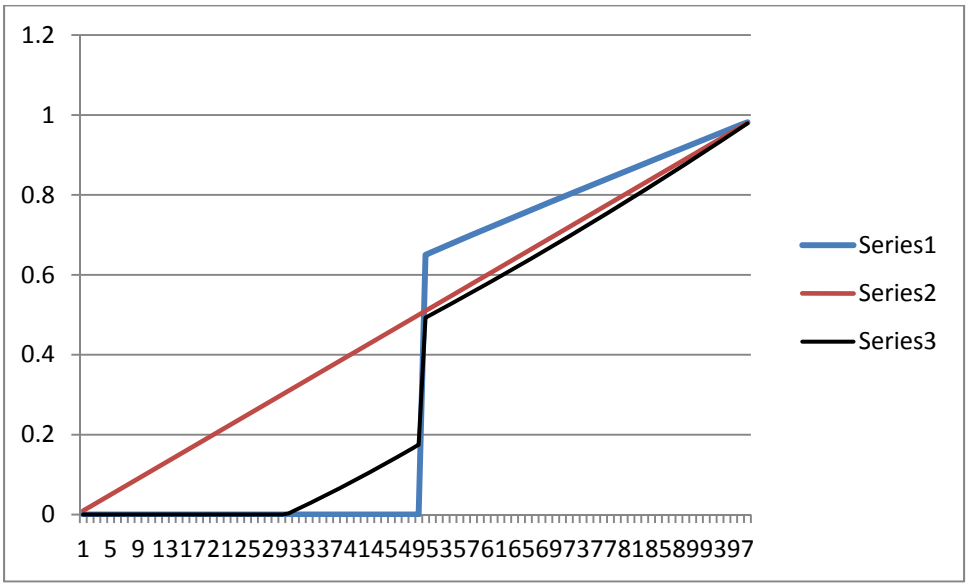

In this case the line for the second order $\Delta P(P)$ (series 3) does not appear to fully intersect the line of identity (series 2). We are unclear if this is simply a reflection of the approximate nature of numerical integration.

CONCLUSIONS:

1. If the alert success rate (the frequency at which the alert generates an intervention rather than being ignored) is a function of the alert threshold (as would be the case if the success rate depends on the alert load), there is an optimal alert threshold. This reflects the balance between the improvement in outcome with the alert (every time the alert fires it can improve outcome) and alert fatigue (but the more the alert fires the more likely the alert will be ignored).
2. If the intervention triggered by the alert simply reduces the probability of deterioration linearly, the optimal threshold for the alert is independent of the efficacy of the intervention. This is counter-intuitive.
3. As alert fatigue intensifies, the optimal alert threshold increases, i.e., the alert is more selective.
4. As the mean of the distribution of the probability of deterioration increases, the optimal alert threshold increases.
5. As the variance of the distribution of the probability of deterioration increases, the optimal alert threshold increases.
6. If alert fatigue occurs only when the clinician is exposed to some level of alerts in excess of a critical value, the optimal alert threshold is equal to the probability of deterioration associated with that critical value.

Georgia Institute of Technology

School of Aerospace Engineering
Atlanta, Georgia 30332-0150 U.S.A.
PHONE 404.894.3000
FAX 404.894.2760

December 1, 2012

Vishnu Mohan, MD, MBI, MBCS, FACP
Department of Medical Informatics
and Clinical Epidemiology
Oregon Health and Science University
Portland, Oregon

Dear Dr. Mohan:

I have reviewed the theoretical appendix of the capstone project, entitled "Development of a Business Intelligence Tool for the Extraction, Analysis, Display, and Utilization of Early Warning System Alerts" submitted by your advisee Dr. James M. Bailey. I have examined the mathematical exposition in detail and attest that the formulation of the problem of the optimal early warning system alert threshold in terms of the mean probability of deterioration, establishing Equations (1) and (2) as the foundation for further development is mathematically sound and well-grounded. Furthermore, there are no mathematical errors in the subsequent derivations and the simulations are consistent with the underlying theory. In general, the problem is appropriately formulated and there are no mathematical errors within the constraints of the stated model.

Sincerely,

Wassim M. Haddad
Professor