

Assessing Data Availability for Answering Bedside Questions in the Critical Care Settings

(Research-in-progress, IQ Assessment, Measures and Models)

Abdul Ghani Mohammed, MS

University of Arkansas at Little Rock, College of Engineering and Information Technology,
Department of Information Science
ggmohammed@ualr.edu

Meredith Nahm Zozus, PhD

University of Arkansas for Medical Sciences, College of Medicine,
Department of Biomedical Informatics
mzozus@uams.edu

Anne Sofie Andreasen, MD, Ph.D

University of Arkansas for Medical Sciences, College of Medicine
Department of Anesthesiology,
Department of Anesthesiology and Intensive Care, Herlev University Hospital, Denmark
anne.sofie.andreasen@regionh.dk

Melody L Penning, PhD

University of Arkansas for Medical Sciences, College of Medicine,
Department of Biomedical Informatics
mlpenning@uams.edu

Abstract: With ubiquitous implementation of Electronic Health Records (EHRs), healthcare organizations hold large volumes of clinical data. Historic, observational data from routine care may provide useful information for bed-side decisions in patient care. We take the initial step toward this goal through analysis of 53 clinical questions relevant to clinical decision making in the Intensive Care Unit. Each question was decomposed into population, intervention and outcome statements then into data elements. Overall, 92.5 % of the questions were supported by data elements. However, algorithms were needed for population determination for 98% of the questions with available data elements. Thirty-one (63%) of the interventions required algorithms. Seven of the standard outcomes required algorithms. The work reported here is the initial step in evaluating the feasibility of observational data for use in clinical decision support. The results are encouraging enough to support further analysis.

Introduction

Electronic Health Records (EHRs) are almost ubiquitous in U.S. hospitals. We cannot help but wonder if these existing large stores of data could be used to support clinical decision-making in areas where gaps in evidence-derived knowledge exist and what might be needed to evaluate the safety of such use of observational data.

Today, an evidence gap exists in critical care as well as in many other medical specialties. It is estimated that currently, approximately 65 years after publication of the first randomized controlled trial, only 10-20% of clinical decisions are based upon evidence derived from this gold-standard of knowledge discovery.⁽¹⁾ Since 2010, an average of 21,000 clinical studies are registered annually in clinicaltrials.gov with a 13% increase in the number of studies reported from 2015 to 2016.⁽²⁾ Ideally, this increase in trials should expand the amount of evidence and thus strengthen the foundations on which we base patient management. However, the RCT as an investigation type has several inherent limitations. RCTs are expensive and cumbersome to perform. Findings from a RCT on a bigger, heterogeneous group of patients might not apply to the physician's individual patient, and vice versa, findings from RCTs with tight inclusion criteria, exclusion criteria, and experimental controls have limits in their generalizability to real-world clinical practice. Further, there will most often be a significant delay from achieving RCT-derived knowledge till the findings become of benefit for the patients; findings from studies must be assessed, synthesized, incorporated into clinical guidelines then implemented into clinical practice and thus until the knowledge.

As such, there is a request and a need for alternative ways of deriving knowledge. In recent years, a growing interest in how to use the vast amounts of digitalized medical information stored in the EHRs, as a supplement to or an independent way of achieving knowledge, has emerged in critical care medicine.⁽³⁾⁽⁴⁾⁽⁵⁾ These observational data consist of patient information, medication lists, lab results, vitals, wave forms from monitor output, procedures, diagnoses etc. The so-called Dynamic Clinical Data Mining or "The Green Button solution" have been suggested as processes by which data can be used in clinical decision support⁽³¹⁾, where evidence-based guidelines do not exist.⁽⁶⁾ By searching population databases, similar patients based upon entered information and characteristics could be identified. The treatments and outcomes of these historical individuals could then be compared and displayed as bedside clinical decision support.⁽⁷⁾

Appreciating the opportunities in using observational data from electronic records, the critical care community has witnessed the development of a number of ICU databases. These large collections of observational datasets from patients admitted to the ICU may be commercial with data from multiple ICUs, eg. APACHE Outcome database or the Philips eICU (<http://www.usa.philips.com>), or they may be publicly accessible and from a single center.⁽⁸⁾⁽⁹⁾ They may also contain ICU data from an entire nation's ICUs in countries with a centralized health care system, such as Denmark.⁽¹⁰⁾ These databases have been used in a number of studies, searching for various associations, such as associations between fluid balance and survival,⁽¹¹⁾ or relationships between oliguria, creatinine, mortality and renal replacement therapy.⁽¹²⁾ However, in the ICU we are yet not able to perform bed-side, real-time searches to identify patients close in similarity, retrieve information on interventions and outcome of these similar patients, and use the information to impute the best choice of action for diagnosis, prognosis and treatment for future patients. The project reported here is an initial step towards enabling us to do so.

Background

The Institute of Medicine's landmark report *To Err is Human: Building a Safer Health System*, estimated that medical errors contributed to between 44,000 and 98,000 deaths annually in the United States.⁽¹³⁾ The subsequent report, *Crossing The Quality Chasm*, laid out a strategy to increase quality and safety in healthcare.⁽¹⁴⁾ Better synthesis and use of information through evidence based medicine and health information technology has aggressively been pursued as a solution for the last two decades.⁽¹⁶⁾⁽¹⁷⁾⁽¹⁸⁾⁽¹⁵⁾ arguably since the 1970's. However, regardless of technology, the effectiveness of evidence based medicine is limited by the available evidence and for the evidence that is available there is a need to perform quality checks to ascertain that it meets the appropriate standard.⁽¹⁹⁾

Computer-based clinical decision support (CDS) is the “use of the computer to bring relevant knowledge to bear on the health care and well-being of a patient.”⁽²⁰⁾ CDS applications range from simple alerts and reminders, e.g., flagging a high lab value, detecting a drug interaction or reminding that a screening test is due, to data displays that aid trend detection, to complex activities such as making a differential diagnosis or choosing an optimal plan of action.⁽²⁰⁾ The supporting methodologies include form-based templates and check lists, algorithm-based alerts and reminders using existing or external knowledge sources, assisted retrieval of relevant information from clinical guidelines or the medical literature, and rule-based or probabilistic decision aids. Probably the most important learning in the almost sixty year trajectory in the development of clinical decision support is that unlike the Greek oracle⁽²¹⁾, decision support in healthcare augments but does not replace the clinician’s assessment, reasoning and actions.⁽²²⁾⁽²³⁾

Believing that a person’s limited memory was one of the greatest weaknesses of clinical practice, Dr. Eugene Stead, then (1967) Chair of the Department of Medicine at Duke University, developed a database of patient characteristics, diagnoses, treatments, and outcomes. His vision was that the computer be used hospital-wide as a "computerized textbook of medicine," replacing a doctor's fallible memory of how to treat a condition or disease with a computer's infallible memory of each patient treated in the hospital.⁽²⁴⁾ Today, the Duke Databank for Cardiovascular Disease is the oldest and largest compilation of data on heart care outcomes, with comprehensive, long-term records on more than 200,000 people.

We are not the first to pursue the real-time use of observational data in clinical settings. In 2007, the developers of the behavioral health EMR, MindLinc created a tool that supported real-time comparison of the outcomes of similar patients receiving different pharmacological interventions.⁽²⁵⁾⁽²⁶⁾ The EMR system had a large multi-institutional database from which to draw - more than 500,000 patients collected during 7 million visits.⁽²⁷⁾ The Clinical Global Impression (CGI) score, collected at each encounter for each patient in the system was available as an outcome measure. Probably for these reasons, (1) a very large longitudinal dataset and (2) standardized, interventions limited to pharmacotherapy and standardized DSM based and ICD mapped diagnoses by which the population was easily determined in conjunction with a standardized outcome measure, the approach was successful in behavioral health supported by the described system. While institutional EHR systems at single or small hospital groups rarely have specialty populations in these numbers, multi-institutional networks federating data and integrated health systems do, and regional collaborations have the potential to make these numbers available to even single facilities. Thus, we desired to start testing this approach - using observational EHR-based data to support clinical decision-making at the point of care, in specialties and settings with more varied data such as critical care settings.

While there has certainly been a national push for use of observational EHR data in research, for example, the Patient Centered Outcomes Research Network (<http://www.pcornet.org>), the Healthcare Systems Research Collaboratory (<https://www.nihcollaboratory.org>) for comparative effectiveness research and clinical trials alike, significant questions naturally arise about using raw observational data in decision support rather than high quality evidence, i.e., vetted clinical guidelines synthesized from multiple randomized clinical trials. However, two clear arguments support the application of the use of EHR data. First, the legal position codified in the New Jersey Product Liability Act with respect to drug prescribing that has been applied to EHR technology⁽²⁸⁾⁽²⁹⁾ has been that the technology itself does not directly act on the patient; a *learned intermediary*, i.e., a physician who makes and is responsible for decisions, mediates the interaction. Second, and more important, the output of such a system is one input to the clinical decision-making process.⁽²⁰⁾ Therefore, these data do not constitute a guideline, best practice, or rule. Instead, the output is a simple comparison of an outcome of interest for the indicated intervention and population – a piece of information that may be factored into decision-making and needs to be weighed with the often many other factors considered in clinical decision-making. Given these, and should such an approach turn out to be feasible, testing will be required to assess the frequency with which the output differs from

guideline-based recommendations and why. Testing will also be required with clinicians at the point of care to assure that with appropriate training, that clinicians use the information as one input into clinical decision-making rather than as a best practice of directive. We would also like to rest the approach with patients to assess whether patients would want their provider to factor such information into clinical decision-making about important aspects of their care.

We are, however, a long way from such testing. We must first assess whether (1) the types of clinical questions arising at the bedside can easily be formulated in terms of a population, intervention and an outcome of interest, (2) the required data are collected in the EHR at the average institution, (3) there are sufficient numbers of patients once a population is defined and how large of a dataset would be required to support population definitions for common questions, (4) the EHR data sufficiently standardized in terms of format and meaning to support such direct uses of the data, and (5) the data accurate enough, i.e., do the data support the same conclusions as error free data. In information quality parlance, factors two through five are the dimensions of data quality that are important for the desired use of the data, in this case (2) data element availability, (3) data completeness and volume, (4) data standardization, and (5) data accuracy. If the data quality is not adequate to support the desired use – answering common clinical questions at the bedside, then we need go no further. The preliminary results reported here resulted from the feasibility assessment described for steps (1) and (2).

Methods

We developed an initial list of 51 questions. These questions were designed to address five patterns of inquiry: 1) Areas of medicine with two competing outcomes (e.g. deep sedations favors early wound closure, but deep sedation is also an independent risk factor for worse outcome in the general ICU population); 2) Clinical practice is only guided by expert opinion; 3) to test whether findings from previous RCTs can be confirmed with the data captured and available to clinicians in our EHR; 4) to investigate whether the data in our EHR supports the use of continuous variables for optimized clinical decision making rather than the discreet variables often used in prior RCTs (e.g. transfusion trigger levels); and 5) to determine whether EHR data can be used to identify patient characteristics that may impact outcome from a specific condition or a specific intervention. We applied the decomposition methodology described by Sim et al. (2004) ⁽³⁰⁾ to partition each question into three parts (1) population of interest, (2) interventions of interest, i.e., the independent variable or any variable used as a foundation for decision making that independently could affect outcome, and (3) outcomes of interest, i.e., the dependent variable (Table 1). In the second stage of decomposition, each question part was further broken down into the data elements necessary to retrieve the needed information from EHR-based data sources. In order to simplify this pilot study, we decided to focus on a relatively limited number of common outcome parameters, including those of greatest importance to patient and clinician. Completion of the decomposition for each example question provided a comprehensive set of data elements needed to answer the question. The decomposition required several collaborative iterations between the ICU Physician and the informatics team.

Following the question decomposition, the operational data repository for our EHR (the Clarity tables associated with the Epic EHR) and when necessary the EHR screens were consulted to determine if the data elements required for a particular clinical question were collected. The data under investigation is composed of all clinical data elements, not record level data, observed and collected during the course of all past encounters limited by population filters. This data is typical of, and necessarily collected by, EHR systems. While different EHR systems can have very different front end data collection methods and back end data storage schemas the core elements that are needed for clinical work are common. Any system taking advantage of these data elements would of course need to be customized to the data environment provided but the same basic functionality should be achievable regardless of the differences in setting.

Common categories of data elements needed included clinical observations, measured physiologic parameters, diagnoses, procedures, medical orders, order fulfillment, and derived or otherwise composite data elements such as patient acuity scales or prognostic scores.

Question Component	Examples
<p><u>Population:</u> The characteristics of the patients of interest in the question. The population defines the data to be retrieved for a given clinical question. In the identified question set, population was often identified by service to which the patient was admitted, e.g., the Intensive Care Unit, medical conditions, procedures or prior or current treatment.</p>	<ul style="list-style-type: none"> • ICU patients with a diagnosis of delirium • ICU patients with a diagnosis of Acute Respiratory Distress Syndrome • ICU patients on mechanical ventilation • ICU patients with an open abdomen
<p><u>Intervention:</u> The treatment or action under consideration by the clinician. For example, drug A versus drug B, high dose or low dose of drug A, or timing of a procedure. Intervention in our decomposition was also used for exposures of interest.</p>	<ul style="list-style-type: none"> • The administration of antipsychotics • The tidal volumes applied during mechanical ventilation • An application of Positive end-expiratory pressure (PEEP) > 10 cmH₂O during mechanical ventilation • Maintaining a mean arterial pressure above or below 65 mmHg • Deep versus light sedation during mechanical ventilation as evaluated by Riker Sedation Agitation Score
<p><u>Outcome:</u> The outcomes of interest, on which the interventions or exposures will be judged. For example, is there a difference in length of stay for patients prescribed drug A versus drug B.</p>	<ul style="list-style-type: none"> • Discharge from ICU alive / Discharge from hospital alive • Discharge to home vs. Discharge to facility • Length of stay (in the ICU and in hospital) • 30 day Re-admission

Table 1: Decomposition of Clinical Questions into Components: Population, Intervention, and Outcome

Secondarily, any algorithms necessary to retrieve or use available data were described. The necessity of an algorithm (orchestrated combinations of *multiple* data elements) is an important consideration because such algorithms used in EHR data can get quite complex, can undermine generalizability of the method across questions, and require validation when used in a new data source. Having said this, data points built from algorithms are not at all inferior, they just require more time and effort to build and then validate afterwards. Data element availability and the necessity of an algorithm were documented in a matrix of the initial questions as the rows and the question components (population, intervention, and outcome) as the columns. If all the data elements required for a component of a question were present in the EHR (screens or data repository), the question component was scored affirmative for that question. Question components requiring an algorithm to retrieve or calculate the information needed for the component were marked as such.

<p><u>Bedside Question:</u> Is deep or light sedation better for intubated patients in the ICU with open abdomen</p>
<p><u>Population:</u> ICU Patient with an open abdomen</p> <p><u>Algorithm for Population:</u></p> <p>SELECT Patients IF (The Patient Accommodation is ('Cardiovascular Intensive Care' OR 'Medical Intensive Care' OR 'Neuro Intensive Care' OR 'Level IV-ICU-Bone Marrow' OR 'Surgical Intensive Care') AND the Instant of Admission was during the Specified Time AND the Record has not been canceled AND the Patient ID is not NULL AND the Location is not NULL AND the Department Name = an ICU department) with any of (Explorative laparotomy: SHX1544, UR051, SUR591)) AND one or more day later any of the following procedures (Abdominal closure: 13160) AND Procedure code for Mechanical Ventilation: 94002, 94003 during the same time period</p>
<p><u>Intervention:</u> Deep versus light sedation vs. agitation</p> <p><u>Algorithm for Intervention:</u></p> <p>Sedation level (signified by RIKER score) with or without simultaneous administration of non- depolarizing Neuromuscular Blockers</p> <p>Riker levels: 1-2: Deep sedation, Riker 3-4: light sedation, Riker 5-7: agitation</p>
<p><u>Outcome:</u> Discharge from ICU and hospital alive</p> <p><u>Algorithm for Outcome:</u></p> <p>PATIENT.DEATH_DATE is not Null / ZC_DISCH_DISP.NAME == 'Expired' / PAT_ENC_HSP.HOSP_DISCH_TIME is not null + CLARITY_ADT.EFFECTIVE_TIME + EVENT_TYPE_C</p>

Table 2: Example Decomposition of a Bedside Question into Question Components

The above figure illustrates our approach towards categorizing the questions into population, intervention and outcome. This particular question can be further broken down into data elements and be linked to the elements present in the database with the help of algorithms for all the three components.

Our goal was to retrieve data in the historical database from patients so that it would be possible to perform a clinically relevant comparison to the specific patient of interest. We wanted to be able to perform a “patient like this one” search among a large patient population, comparable to our one patient in a clinically meaningful way, but also to ensure that the population for comparison was not minimized too much by being overly specific in our search criteria.

In order to do so, we used of a limited number of relevant characteristics which would define the “patient like this one” population, when doing the bedside query might be necessary. Those characteristics were chosen based upon some traditionally used variables expected (independently or dependently) to influence outcome of a condition or intervention. Such variables included: age, gender, race, height and weight for BMI calculation, history of alcohol and tobacco, APACHE II score, SOFA score, admitting department,

whether the patient was re-admitted to the ICU and whether the patient had sepsis, acute renal failure, was on mechanical ventilation, or was on renal replacement therapy.

For outcomes, we also employed a simplifying approach. While a large variety of outcomes might certainly be of interest in a given clinical scenario, we decided to focus only on those outcomes with the greatest importance in clinical decision making and for the patient. Based on the clinical setting and the questions, we defined eight standard outcomes on which interventions could be compared. These include (1) Discharge from ICU or hospital alive, (2) Discharge to home versus discharge to facility (3) Length of stay in the ICU / in the hospital, (4) 30 day mortality, (5) 1 year mortality, (6) Need for dialysis, (7) Costs associated with hospital admission, (8) 30 day re-admission rate to the ICU.

Since both UAMS and the Danish Capital Region for Healthcare use Epic as the platform for the EHR, our pilot study was conducted using this system.

Results

Overall, data elements were available for 92.5% of the questions. Two questions were marked as unanswerable during the decomposition, as it became clear that these could not be broken into data elements with sufficient level of granularity to answer the question in a clinical meaningful way. An additional two questions were deemed unanswerable at the data element stage of decomposition. In addition, during decomposition, it was realized that as stated two of the initial questions were really two separate questions, increasing the number of questions analyzed to fifty-three. In total, four (7.6%) of the research questions were unanswerable due to data element-level incompleteness.

Algorithms were needed for all but one (2.0%) the 49 questions with available data elements. The same 48 questions required algorithms for the population determination. Thirty-one (63%) of the intervention data elements require algorithms. Seven of the eight standard outcomes required algorithms, however, offering only these eight means that the need for algorithms for the outcomes is much less impactful because the algorithms can be programmed, tested and offered for selection on the user interface.

Overview:

Questions	Algorithms	Total Count
Answerable	49	49
Unanswerable due to data element level incompleteness	-	2
Unanswerable due to algorithmic level completeness (could not be formulated because of missing data elements)	-	2
Total Count	49	53

Discussion

The high number of questions requiring algorithms, while not a surprise, is disheartening. Often called phenotypes by those who specialize in secondary use of clinical data, these algorithms have widely varying sensitivity and specificity and thus are tested prior to use. Use of eight standard outcomes circumvented this problem for the outcomes in that each of the eight algorithms can be defined, programmed and tested and provided for use. However, the population and intervention, at least in the sample set of questions analyzed here were specific to the clinical question. Algorithms are no problem in one-time analysis type

secondary data uses, but pose a significant obstacle in real-time interactions with the data, especially when the number of possibilities is too large for a form-based query. The extent of algorithms required is a significant difference between our application in the ICU setting versus the successful implementation in a behavioral health medical record where populations could be indicated based on diagnosis and a few additional parameters, there was one outcome measure and the interventions were restricted to pharmacotherapy. A broader set of ICU questions would need to be analyzed to explore what concepts would need to be included in standard sets of populations and interventions, if such a set is possible, and if the restrictions imposed by such a set significantly degrade the types of questions that can be asked at the bed-side.

The question, why such extensive use of algorithms is necessary warrants discussion. The need for such algorithms is pervasive in clinical applications. The industrial statistician George Box once said something similar to, “all models are wrong, some are useful”. The sentiment applies to all representations of reality. In healthcare, the data that we collect about individuals, their conditions and the treatment provided are a representation of the individual at the point in time at which the data are documented. The data collected are far from a complete description of the patient. When a clinician is treating a patient, the clinician is in reality, and the questions that arise may be about very subtle aspects of that reality. The likelihood that documented data cover the object of the clinician’s question, is slim. Thus, we must approximate the question (usually with an algorithm) and constrain the question by the available data. For our application here, our interest is to what extent do we need to constrain the questions, and when we do so, do we so severely limit the answerable questions that the approach is futile.

Other aspects of the need for algorithms are also important in our consideration. For example, once we understand the important bed-side questions, we may have success in changing what is documented to better support common questions. During the course of the work reported here, two prognostic scores (APACHE II and SOFA) were implemented in the EHR system. However, as with many such improvements, the scores will unfortunately not be retroactive once they are in place in the EHR. For Legacy data these will need to be determined using algorithms and some data points may need imputation with surrogate concepts. These concepts must of course be clinically acceptable. and the generated scores will need to be tested against the scores provided by physicians for accuracy.

Another consideration introduced with the need for extensive use of algorithms is use of data elements for which multiple values exist. In the algorithms, the needed value must be identified, for example, the first, last, min, max, or average. We may be able to standardize this sufficiently by calculating descriptive statistics for measures that are used to determine a patient’s condition, or by devising a standard method to flag important time points or durations. Similarly, for standardized scores such as RIKER, we used three categories as a surrogate for three levels of sedation. A similar approach is likely needed for use of medication dosages on algorithms where the clinical question and thus the important dosing differences will differ. Generalizing a strategy to “bin” or otherwise categorize the data on the fly at the user interface will be a particular challenge and will require further evaluation of the common cases and their amicability to such categorization in the context of a specific question.

Our initial question was answered by this small study – were the data elements available in the EHR data to answer common bed-side questions in the ICU. Most of them were, however, not in the exact combinations required to answer the questions easily. Answering most questions required one or more algorithm. Based on this, the next question becomes is it possible to standardize algorithms for populations and interventions, such as was done for the eight standard outcomes? Our sample of questions was too limited to answer this question and will need to be expanded for further exploration.

Limitations

The work reported here is subject to several limitations. As previously discussed, the basis of our study was 53 questions presented by ICU clinicians. These are a very small subset of the questions that may arise at the bedside. Further, the data that we considered were from just one institution. It is possible that at another institution, the results would vary.

Conclusion

This study showed positive results in terms of data element availability with data elements needed for 92.5% of the questions present in the EHR. The analysis demonstrates that data elements are present in the systems, and that further investigation into the use of observational data from the EHR could be fruitful. Our results also indicate that additional research is needed to more fully characterize common clinical questions arising at the bedside and whether standard population characteristics and interventions are feasible and useful. After this, directions for further research are many and include questions about other aspects of data quality such as completeness and accuracy. Further, the validity of the approach to using observational data for clinical decision support may be dependent on high data quality and improved data capture. Future studies are needed to investigate whether there are sufficient numbers of patients once a population is defined and how large of a dataset would be required to support population definitions for common questions, whether the EHR data are sufficiently standardized in terms of format and meaning to support such direct uses of the data, and whether the data are accurate enough, i.e., do the data support the same conclusions as error free data. The question that remains to be answered is whether practicing clinicians and health care facilities will see enough benefit from the use of observational EHR data to drive improvements in data quality.

References

1. Farley, H. L., Baumlin, K. M., Hamedani, A. G., Cheung, D. S., Edwards, M. R., Fuller, D. C., ... & Nielson, J. (2013). Quality and safety implications of emergency department information systems. *Annals of emergency medicine*, 62(4), 399-407.
2. McGinnis, J. M., Stuckhardt, L., Saunders, R., & Smith, M. (Eds.). (2013). *Best care at lower cost: the path to continuously learning health care in America*. National Academies Press.
3. Angus, D. C. (2015). Fusing randomized trials with big data: the key to self-learning health care systems?. *Jama*, 314(8), 767-768.
4. Anthony Celi, L., Mark, R. G., Stone, D. J., & Montgomery, R. A. (2013). "Big data" in the intensive care unit. Closing the data loop.
5. Ghassemi, M., Celi, L. A., & Stone, D. J. (2015). State of the art review: the data revolution in critical care. *Critical Care*, 19(1), 118.
6. ClinicalTrials.gov, accessed (April 23, 2016), *Trends, Charts and Maps*. Retrieved from <https://clinicaltrials.gov/ct2/resources/trends#RegisteredStudiesOverTime>
7. Celi, L. A. G., Tang, R. J., Villarroel, M. C., Davidzon, G. A., Lester, W. T., & Chueh, H. C. (2011). A clinical database-driven approach to decision support: Predicting mortality among patients with acute kidney injury. *Journal of healthcare engineering*, 2(1), 97-110.
8. Johnson, A. E., Pollard, T. J., Shen, L., Lehman, L. W. H., Feng, M., Ghassemi, M., ... & Mark, R. G. (2016). MIMIC-III, a freely accessible critical care database. *Scientific data*, 3.
9. Saeed, M., Villarroel, M., Reisner, A. T., Clifford, G., Lehman, L. W., Moody, G., ... & Mark, R. G. (2011). Multiparameter Intelligent Monitoring in Intensive Care II (MIMIC-II): a public-access intensive care unit database. *Critical care medicine*, 39(5), 952.

10. Lassen, A. T., Jørgensen, H., Jørsboe, H. B., Odby, A., Brabrand, M., Steinmetz, J., ... & Christiansen, C. F. (2016). The Danish database for acute and emergency hospital contacts. *Clinical epidemiology*, 8, 469.
11. Lassen, A. T., Jørgensen, H., Jørsboe, H. B., Odby, A., Brabrand, M., Steinmetz, J., ... & Christiansen, C. F. (2016). The Danish database for acute and emergency hospital contacts. *Clinical epidemiology*, 8, 469.
12. Mandelbaum, T., Lee, J., Scott, D. J., Mark, R. G., Malhotra, A., Howell, M. D., & Talmor, D. (2013). Empirical relationships among oliguria, creatinine, mortality, and renal replacement therapy in the critically ill. *Intensive care medicine*, 39(3), 414-419.
13. Donaldson, M. S., Corrigan, J. M., & Kohn, L. T. (Eds.). (2000). *To err is human: building a safer health system* (Vol. 6). National Academies Press.
14. Richardson, W. C., Berwick, D. M., Bisgard, J. C., Bristow, L. R., Buck, C. R., & Cassel, C. K. (2001). *Crossing the quality chasm: a new health system for the 21st century*.
15. U.S. Department of Health and Human Services, Office of the National Coordinator American Recovery and Reinvestment Act of 2009, Information Technology Professionals in Health Care: Program of Assistance for University-Based Training FOA 12172009 508. FOA release date, December 17, 2009.
16. Han, Y. Y., Carcillo, J. A., Venkataraman, S. T., Clark, R. S., Watson, R. S., Nguyen, T. C., ... & Orr, R. A. (2005). Unexpected increased mortality after implementation of a commercially sold computerized physician order entry system. *Pediatrics*, 116(6), 1506-1512.
17. US Department of Health and Human Services. Agency for Healthcare Research and Quality (AHRQ), National Guideline Clearinghouse Guideline Index.
18. Cronenwett, L. R., Bootman, J. L., Wolcott, J., & Aspden, P. (Eds.). (2007). *Preventing medication errors: quality chasm series*. National Academies Press.
19. National Research Council. (2009). *Computational technology for effective health care: immediate steps and strategic directions*. National Academies Press.
20. Zozus MN, Hammond WE, Green BG, Kahn MG, Richesson RL, Rusincovitch SA; Simon GE, Smerek MM. Assessing Data Quality for Healthcare Systems Data Used in Clinical Research (Version 1.0): An NIH Health Care Systems Research Collaboratory Phenotypes, Data Standards, and Data Quality Core White Paper. Available from www.nihcollaboratory.org
21. Greenes, R. A. (2007). A brief history of clinical decision support: technical, social, cultural, economic, and governmental perspectives. *Clinical Decision Support: The Road Ahead*, 31-91.
22. Miller, R. A. (1994). Medical diagnostic decision support systems—past, present, and future. *Journal of the American Medical Informatics Association*, 1(1), 8-27.
23. Miller, R. A. (1996). Evaluating evaluations of medical diagnostic systems. *Journal of the American Medical Informatics Association*, 3(6), 429.
24. Friedman, C. P. (2009). A “fundamental theorem” of biomedical informatics. *Journal of the American Medical Informatics Association*, 16(2), 169-170.
25. Duke University, accessed (March 17, 2017), *Duke Databank for Cardiovascular Disease*. Retrieved from <http://digitaldukemed.mc.duke.edu/databank/index.html>
26. Wu, L. T., Gersing, K. R., Swartz, M. S., Burchett, B., Li, T. K., & Blazer, D. G. (2013). Using electronic health records data to assess comorbidities of substance use and psychiatric diagnoses and treatment settings among adults. *Journal of psychiatric research*, 47(4), 555-563.
27. Gersing KR, Krishnan KR. Practical evidence-based medicine. *Psychopharmacology Bulletin*, 36, 20–26.
28. Mindlinc, accessed (March 18, 2017), *Mindlinc*, Retrieved from <http://www.mindlinc.com/>
29. Koppel, R., Metlay, J. P., Cohen, A., Abaluck, B., Localio, A. R., Kimmel, S. E., & Strom, B. L. (2005). Role of computerized physician order entry systems in facilitating medication errors. *Jama*, 293(10), 1197-1203.

30. Sim, I., Olasov, B., & Carini, S. (2004). An ontology of randomized controlled trials for evidence-based practice: content specification and evaluation using the competency decomposition method. *Journal of biomedical informatics*, 37(2), 108-119.
31. Longhurst, C. A., Harrington, R. A., & Shah, N. H. (2014). A 'green button' for using aggregate patient data at the point of care. *Health Affairs*, 33(7), 1229-1235.