

# Chapter 16

## Methods for Data Quality Studies

*Francis Lau*

---

### 16.1 Introduction

The proliferation of eHealth systems has led to a dramatic increase in the volume of electronic health data being collected. Such data are often collected as part of direct patient care delivery to document the patient's conditions and the care being provided. When the collected health data are used as intended, it is referred to as primary use. Once collected the health data can be used for other purposes such as clinical quality improvement, population health surveillance, health systems planning and research. These are referred to as secondary uses (Safran et al., 2007). In Canada, a further distinction is made where all secondary uses except for research are labelled as health system use (Canadian Institute for Health Information [CIHI], 2013).

The quality of routinely collected eHealth data is a major issue for healthcare organizations. To illustrate, a systematic review by Thiru, Hassey, and Sullivan (2003) on EHR data quality in primary care found a great deal of variability depending on the type of data collected. In 10 EHR studies on sensitivities they found data completeness ranged from 93% to 100% for prescriptions, 40% to 100% for diagnoses, 37% to 97% for lifestyle in alcohol use and smoking, to 25% for socio-economic data. A 2010 review by Chan, Fowles, and Weiner (2010) showed that the variability in the quality of EHR data is an ongoing issue especially with problem lists and medications. Iron and Manuel (2007) reported in their environmental scan of administrative health data quality assessment that: the concepts of accuracy and validity are often confused; there are no standard methods for measuring data quality; and the notion of data quality depends on the purpose for which the data are used. These findings suggest data quality can affect the performance of the eHealth systems and care delivery within organizations.

In this chapter we describe approaches to eHealth data quality assessment that are relevant to healthcare organizations. The approaches cover concepts, practice and implications. Concepts refer to dimensions, measures and methods of data quality assessment. Practice refers to how data quality assessment is done in different settings as illustrated through case examples. Implications refer to guidance and issues in eHealth data quality assessment to be considered by healthcare organizations.

## 16.2 Concepts of Data Quality

In this section we describe the key concepts in eHealth data quality. These are the conceptual quality dimensions, measures used to assess the quality dimensions, and methods of assessment. These concepts are described below.

### 16.2.1 Data Quality Dimensions

An overriding consideration when defining data quality concepts is “fitness for use”. This suggests data quality is a relative construct that is dependent on the intended use of the data collected. Different terms have been used to describe the conceptual dimensions of data quality, with no agreement on which should be the standard. Sometimes the meanings of these terms overlap or conflict with each other. Drawing on the studies by Weiskopf and Weng (2013) and Bowen and Lau (2012), we arrived at the following five commonly cited terms for this chapter:

- *Correctness* – Reflects the true state of a patient, also known as accuracy. An example is whether a high blood pressure value for a patient is true or not.
- *Completeness* – Covers all truths on a patient, also known as comprehensiveness. An example is the blood pressure measurement that contains the systolic and diastolic pressures, method of assessment, and date/time of assessment.
- *Concordance* – Agreement of the data with other elements or sources, also known as reliability, consistency and comparability. An example is the use of metformin as a diabetic medication in the presence of a diabetes diagnosis.
- *Plausibility* – Does the data make sense in what is being measured given what is known from other elements? This is also known as validity, believability and trustworthiness. An example is the presence of a hypertension diagnosis in the presence of recent abnormal blood pressure measurements.

- *Currency* – Reflects the true state of a patient at a given point in time, also known as timeliness. An example is the presence of a recent blood pressure measurement when considering a hypertensive condition.

In the literature review by Weiskopf and Weng (2013), completeness, correctness and concordance were the most common dimensions assessed. Other less common data quality dimensions described in the literature (Bowen & Lau, 2012) include comprehensibility, informative sufficiency, and consistency of capture and form. These terms are defined below.

- *Comprehensibility* – The extent to which the data can be understood by the intended user.
- *Informative sufficiency* – The extent to which the data support an inference on the true state of condition.
- *Consistency of capture* – The extent to which the data can be recorded reliably without variation by users.
- *Consistency of form* – The extent to which the data can be recorded reliably in the same medium by users.

### 16.2.2 Data Quality Measures

The dimensions of correctness and completeness can be quantified through such measures as sensitivity, specificity, positive predictive value and negative predictive value. Quantifying these data quality measures requires some type of reference standard to compare the data under consideration. Using a health condition example such as diabetes, we can take a group of patients where the presence or absence of their condition is known, and compare with their charts to see if the condition is recorded as present or absent. For instance, if the patient is known to have diabetes and it is also recorded in his chart then the condition is true. The comparison can lead to different results as listed below.

- *Sensitivity* – The percentage of patients recorded as having the condition among those with the condition.
- *Specificity* – The percentage of patients recorded as not having the condition among those without the condition.
- *Positive predictive value* – The percentage of patients with the condition among those recorded as having the condition (i.e., condition present).

- *Negative predictive value* – The percentage of patients without the condition among those recorded as not having the condition (i.e., condition absent).
- *Correctness* – The percentage of patients with the condition among those recorded as having the condition. It can also be the percentage of patients without the condition among those recorded as not having the condition. These are also known as the positive predictive value and negative predictive value, respectively. Often only positive predict value is used to reflect correctness.
- *Completeness* – The percentage of patients recorded as having the condition among those with the condition. It can also be the percentage of patients recorded as not having the condition among those without the condition. These are also known as sensitivity and specificity, respectively. Often only sensitivity is used to reflect completeness.

The comparison of the patients' actual condition against the recorded condition can be enumerated in a 2x2 table (see Table 16.1). The actual condition represents the true state of the patient, and is also known as the reference standard.

**Table 16.1**

*Calculation of Completeness and Correctness Using Sensitivity and Positive Predictive Value*

Reference Standard Data					
		Condition is Present	Condition is Absent		
Data under evaluation	Condition Appears Present	A – True Positive	B – False Positive	Correctness Positive Predictive Value (PPV) = $A/(A+B)$ in %	
	Condition Appears Absent	C – False Negative	D – True Negative	Negative Predictive Value (NPV) = $D/(C+D)$ in %	
		Completeness Sensitivity = $A/(A+C)$ in %	Specificity = $D/(B+D)$ in %		

Note. From "Defining and evaluating electronic medical record data quality within the Canadian context," by M. Bowen and F. Lau, 2012, *ElectronicHealthcare*, 11(1), e5–e13.

### 16.2.3 Data Quality Methods

Different data quality assessment methods have been described in the literature. Some methods are focused on ways to measure different dimensions of data

quality such as correctness and completeness of the data in an eHealth system. Others are concerned with the means of carrying out and reporting data quality assessment studies. There are also methods that apply predefined criteria to identify and validate specific health conditions recorded in the eHealth system. The types of methods covered in this chapter are defined below and elaborated in the next section.

- Validation of data from single and multiple sources – The use of predefined knowledge and query rules to validate the integrity of the data in one or more eHealth systems and/or databases.
- Designing, conducting and reporting data quality studies – The use of a systematic process to carry out data quality assessment studies.
- Identification and validation of health conditions – The use of predefined criteria to identify and validate specific health conditions in an eHealth system or database. The process is also known as case definition or case finding, and the criteria may be from evidence-based guidelines or empirically derived with expert consensus.

### 16.3 Methods of Data Quality Assessment

This section examines the three types of data quality assessment methods defined in section 16.2.3. Most of the methods were developed as part of data quality assessment studies or as validation of previously developed methods. The analysis in these methods typically involves the use of frequency distributions, cross-tabulations, descriptive statistics and comparison with a reference source for anomalies. These methods are described below.

#### 16.3.1 Validation of Data from Single and Multiple Sources

Brown and Warmington (2002) introduced Data Quality Probe (DQP) as a method to assess the quality of encounter-driven clinical information systems. The principle behind DQP is that predefined queries can be created from clinical knowledge and guidelines to run against the system such as an EHR as measures of its quality. Typically the DQPs examine two or more data recordings that should or should not appear together in the patient record. The most common DQPs involve checking for the presence of a clinical measurement that either should always or never be associated with a diagnosis, and a therapy that either should always or never be accompanied by a diagnosis or clinical measurement. In an ideal system there should be no data inconsistencies detected when the queries are run. Examples are the absence of Hemoglobin A1c (HbA1c) test results on patients with diabetes and prescriptions of penicillin on patients with a penicillin allergy. Two types of errors can be detected: failure to record the

data or error of omission, and suboptimal clinical judgment or error of commission. Once detected, these errors should be reported, investigated and corrected in a timely fashion. To be effective, the DQPs should be run periodically with reports of any inconsistencies shared with providers at the individual and/or aggregate level for purposes of education or action.

Kahn, Raebel, Glanz, Riedlinger, and Steiner (2012) proposed a two-stage data quality assessment approach for EHR-based clinical effectiveness research that involves single and multiple study sites. In stage-1, source datasets from each site are evaluated using five types of data quality rules adapted from Maydanchik (2007). In stage-2, datasets from multiple sites are combined, with additional data quality rules applied, to compare the individual datasets with each other. Such multisite comparisons can reveal anomalies that may not be apparent when examining the datasets from one site alone. The five types of stage-1 data quality rules are outlined below (for details, see Kahn et al., 2012, p. S26, Table 3).

- Attribute domain constraints – Rules that restrict allowable values in individual data elements using assessment methods of attribute profiling, optionality, format, precision and valid values to find out-of-range, incorrect format or precision, missing or unusual data values. An example is a birthdate that is missing, unlikely, or in the wrong format.
- Relational integrity rules – Rules that ensure correct values and relationships are maintained between data elements in the same table or across different tables. An example is the use of diagnostic codes that should exist in the master reference table.
- Historical data rules – Rules that ensure correct values and relationships are maintained with data that are collected over time. For example, the recording of HbA1c results over time should correspond to the requested dates/times, follow an expected pattern, and be in a consistent format and unit.
- State-dependent objects rules – Rules that ensure correct values are maintained on data that have expected life cycle transitions. An example is a hospital discharge event should always be preceded by an admission or transfer.
- Attribute dependency rules – Rules that ensure the consistency and plausibility of related data on an entity. An example is the birthdate of a patient should not change over time; neither should a test be ordered on a deceased patient.

For stage-2 data quality rules, the focus is on semantic consistency to ensure data from different sites have the same definitions so they can be aggregated and analyzed meaningfully. The rules typically compare frequency distributions, expected event rates, time trends, missing data and descriptive statistics (e.g., mean, median, standard deviation) of the respective datasets to detect patterns of anomalies between sites. An example is the need to distinguish random versus fasting glucose tests or dramatic differences in the prevalence of diabetes between sites.

In both stage-1 and stage-2 data quality assessment, documentation is needed to articulate the rationale, methods and results of the assessments done. Often there can be hundreds of data quality rules depending on the complexity of the eHealth system and databases involved. Therefore some type of prioritization is needed on the key data elements and assessments to be included. The outputs generated can be daunting especially if every error encountered on each record is reported. The detailed errors should be grouped into categories and summarized into key areas with select performance indicators to report on the overall quality of the system or database, such as the percentage of records that pass all the data quality tests.

### 16.3.2 Designing, Conducting and Reporting Data Quality Studies

Bowen and Lau (2012, p. e10) published a 10-step method for conducting a context-sensitive data quality evaluation. These steps provide a systematic means of planning, conducting and reporting a data quality evaluation study that takes into account the intended use of the data and the organizational context. The 10 steps are:

- Focus on an activity that requires the use of the data being evaluated.
- Determine the context in which the activity is carried out, including the intent, tasks, people and results.
- Identify the tools/resources needed to evaluate the quality of the data and their alignment with the activity.
- Determine the degree of fitness between the activity and data being evaluated and the acceptable level of fitness.
- Select an appropriate data quality measurement method for the chosen fitness dimension being evaluated.
- Adapt the measurement method depending on the context and use data quality probes to aid in the evaluation.

- Apply the tools/resources identified in step-3 to evaluate the quality of the data being measured.
- Document the output of the fitness evaluation for each data being measured.
- Describe the overall fitness of the important data with the activity in context.
- Present data quality evaluation findings and provide feedback on the quality/utility of the data and improvement.

### 16.3.3 Identification and Validation of Health Conditions

Wright and colleagues (2011, pp. 2–6) developed and validated a set of rules for inferring 17 patient problems from medications, laboratory results, billing codes and vital signs found in the EHR system. A six-step process was used to develop and validate the rules. Additional analyses were done to adjust rule performance based on different rule options. These steps are listed below:

- Automated identification of problem associations – based on a previous study that identified lab-problem and medication-problem associations against gold standard clinical references and significant co-occurring statistics.
- Selection of problems of interest – an initial list of problems ranked according to three criteria: related pay-for-performance initiatives; existing decision support rules in EHR; strength of identified associations.
- Development of initial rules – confirmed relevant lab tests, medications and billing codes with medical references, added relevant free-text entries, then drafted initial rules which were reviewed by expert clinicians.
- Characterization of initial rules and alternatives – focused on patients with at least one relevant medication, lab test, billing code and vital sign but without the problem recorded in EHR; applied initial rules to identify rule-positive and rule-negative patients, then conducted chart review on a sample of patients to see if they have the problem (i.e., chart-positive, chart-negative); derived sensitivities, specificities, Positive Predictive Value (PPV) and Negative Predictive Value (NPV) of initial rule options by varying their thresholds such as lab values, drugs and counts.

- Selection of the final rule – had expert clinicians review different rule options for each problem with varying sensitivities, specificities, PPV and NPV; selected final rules with high PPV over specificity.
- Validation of the final rule – repeated above steps using an independent patient dataset from the same population.
- Additional analyses – derived sensitivity, specificity, PPV and NPV with coded problems then billing data only to adjust final set of rules based on F-measure for higher PPV over sensitivity (false negatives versus false positives).

An example of the final rules for diabetes is shown below (Wright et al., 2011, supplementary data):

- Rule 0: code or free-text problem on problem list for diabetes.
- Rule 1: any HbA1c result greater than or equal to 7.
- Rule 2: 2 or more ICD-9 billing codes in diabetes (250, 250.0, 250.00, 250.01, 250.02, 250.03, 250.1, 250.10, 250.11, 250.12, 250.13, 250.2, 250.20, 250.21, 250.22, 250.23, 250.3, 250.30, 250.31, 250.32, 250.33, 250.4, 250.41, 250.42, 250.43, 250.5, 250.50, 250.51, 250.52, 250.53, 250.6, 250.60, 250.61, 250.62, 250.63, 250.7, 250.71, 250.72, 250.73, 250.8, 250.80, 250.81, 250.82, 250.83, 250.9, 250.91, 250.92, 250.93).
- Rule 3: at least one medication in injectable anti-diabetic agents or oral anti-diabetic agents.

## 16.4 Case Examples

This section includes two published examples of eHealth data quality studies: one is on multisite data quality assessment while the other is on primary care EMRS.

### 16.4.1 Multisite Data Quality Assessment

Brown, Kahn, and Toh (2013) reviewed multisite data quality checking approaches that have been field-tested in distributed networks for comparative effectiveness research such as the Observational Medical Outcomes Partnership in the United States. Typically these networks employ a common data model and different types or levels of data quality checks for cross-site analysis as described in Kahn's 2-stage data quality assessment approach (Kahn et al., 2012). These data quality-checking approaches are:

- Common data model adherence – These are checks on extracted data against the common data model dictionary for consistency and adherence to the model. They are: (a) syntactic correctness on transformed variable names, values, lengths and format meeting data model specifications; (b) table structure and row definition correctness; and (c) cross-table variable relationships for consistency. Examples are valid codes for sex and diagnosis, linkable tables by person or encounter identifier, and presence of valid enrolment for all prescription records.
- Data domain review – These are checks on the frequency and proportion of categorical variables, distribution and extreme values for continuous variables, missing and out-of-range values, expected relationships between variables, normalized rates and temporal trends. The domains may cover enrolment, demographics, medication dispensing, prescribing, medication utilization, laboratory results and vital signs. Examples of checks are enrolment periods per member, age/sex distribution, dispensing/prescriptions per user per month, diagnoses/procedures per encounter, weight differences between men and women, and number of tests conducted per month.
- Review of expected clinical relationships with respect to anomalies, errors and plausibility – Within and cross-site co-occurrence of specific clinical variables should be assessed, such as the rate of hip fractures in 60- to 65-year-old females, male pregnancy and female prostate cancer.
- Member/study-specific checks – These are checks to ensure proprietary and privacy-related policies and regulations for specific members are protected, such as the inclusion of unique product formulary status, clinical/procedure codes, patient level information and tables with low frequency counts; and to detect data variability across study sites such as the exposure, outcome and covariates under investigation.

#### **16.4.2 Improving Data Quality in Primary Care EMRs**

The Canadian Primary Care Sentinel Surveillance Network (CPCSSN) is a pan-Canadian practice-based research network made up of 10 regional networks involving more than 1,000 primary health care providers in eight provinces and territories (CPCSSN, n.d.). Its mission is to improve primary health care delivery and outcomes, epidemiological surveillance, research excellence, and knowledge translation. The effort involves the extraction and use of EMR data from community-based primary health care practices to inform and improve the manage-

ment of the most common chronic diseases in Canada. Here we describe the CPCSNN Data Presentation Tool (DPT) that has been developed to improve the management of individual and groups of patients within and across practices (Moeinedin & Greiver, 2013). In particular, we emphasize the effort undertaken to improve the quality of the EMR data and its impact in the DPT initiative.

- DPT purpose and features – The DPT is an interactive software developed as a quality dashboard to generate custom reports at the provider, office and organizational levels. It uses EMR data that have been de-identified, cleaned and standardized through a systematic process which are then returned to the providers for use in quality improvement purposes. These include the ability to improve data quality at the practice level, re-identify at-risk patients for tracking and follow-up, and produce custom reports such as prescribing patterns and comorbidities in specific chronic diseases (Williamson, Natarajan, Barber, Jackson, & Greiver, 2013; Moeinedin & Greiver, 2013).
- DPT study design – The DPT was implemented and evaluated as a quality improvement study in a family health team in Ontario. The study used mixed methods to examine practice change before and after DPT implementation from May to August 2013. Sixty-one primary care providers took part in the study. The qualitative component included field notes, observations, key informant interviews and a survey. The quantitative component measured the change in data quality during that period (Moeinedin & Greiver, 2013; Greiver et al., 2015).
- Data quality tasks – CPCSNN has developed an automated approach to cleaning EMR data. The data cleaning algorithms are used to identify missing data, correct erroneous entries, de-identify patients, and standardize terms. In particular, the standardization process can reduce the various ways of describing the same item into one term only. Examples are the use of kilograms for weights, one term only for HA<sub>1c</sub>, and three terms only for smoking status (i.e., current smoker, ex-smoker, never smoked). The cleaned data are then returned to the providers, allowing them to assess the data cleaning needed at the local level within their EMRS. To ensure transparency, CPCSNN has published its data cleaning algorithms in peer-reviewed journals and on its website (Greiver et al., 2012; Keshavjee et al., 2014).
- Key findings – The family health team in the DPT study was able to identify missing and non-standardized data in its EMRS. The

DPT was used to produce quality reports such as the prevalence of hypertension and dementia in the region, re-identification of high-risk patients for follow-up, and specific medication recall. Overall, the updating and standardization of the EMR data led to a 22% improvement in the coding of five chronic conditions and the creation of registries for these conditions (Moeinedin & Greiver, 2013; Greiver et al., 2015).

## 16.5 Implications

As healthcare organizations become more dependent on eHealth systems for their day-to-day operations, the issue of eHealth data quality becomes even more prominent for the providers, administrators and patients involved. The consequence of poor-quality data can be catastrophic especially if the care provided is based on incomplete, inaccurate, inaccessible or outdated information from the eHealth systems. The data quality assessment approaches described in this chapter are empirically-derived pragmatic ways for organizations to improve the quality and performance of their eHealth systems. To do so, there are a number of policy and practice implications to be considered.

For policy implications, healthcare organizations need to be aware of the task-dependent nature of data quality, or fitness for use, in order to embark on data quality policies that are most appropriate for their needs. An important first step is to adopt a consistent set of eHealth data quality concepts with clearly defined evaluation dimensions, measures and methods. More importantly, it should be recognized that data quality evaluation is only a means to an end. Once the state of eHealth data quality has been identified, there must be remedial actions with engaged owners and users of the data to rectify the situation. Last, organizational leaders should foster a data quality culture that is based on established best practices.

For practice implications, healthcare organizations need to dedicate sufficient resources with the right expertise to tackle data quality as a routine practice. Data quality evaluation is a tedious endeavour requiring attention to detail that includes meticulous investigation into the root causes of the data quality issues identified. There should be detailed documentation on all of the data quality issues found and remedial actions taken to provide a clear audit trail for references. Last, since providers are responsible for a substantial portion of the routine clinical data being collected, they need to be convinced of the value in having high-quality data as part of patient care delivery.

## 16.6 Summary

This chapter described eHealth data quality assessment approaches in terms of the key concepts involved, which are the data quality assessment dimensions, measures and methods used and reported. Also included in this chapter are two

examples of data quality assessment studies in different settings, and related implications for healthcare organizations.

## References

- Bowen, M., & Lau, F. (2012). Defining and evaluating electronic medical record data quality within the Canadian context. *ElectronicHealthcare*, 11(1), e5–e13.
- Brown, J., & Warmington, V. (2002). Data quality probes – exploiting and improving the quality of electronic patient record data and patient care. *International Journal of Medical Informatics*, 68(1–3), 91–98.
- Brown, J., Kahn, M., & Toh, S. (2013). Data quality assessment for comparative effectiveness research in distributed data networks. *Medical Care*, 51(8 suppl 3), S22–S29.
- Canadian Institute for Health Information (CIHI). (2013). *Better information for improved health: A vision for health system use of data in Canada*. Ottawa, ON: Author.
- Chan, K. S., Fowles, J. B., & Weiner, J. P. (2010). Electronic health records and the reliability and validity of quality measures: A review of the literature. *Medical Care Research and Review*, 67(5), 503–527.
- Canadian Primary Care Sentinel Surveillance Network (CPCSSN). (n.d.). *Canadian primary care sentinel surveillance network* (website). Retrieved from <http://cpcssn.ca/>
- Greiver, M., Drummond, N., Birwhistle, R., Queenan, J., Lambert-Lanning, A., & Jackson, D. (2015). Using EMRS to fuel quality improvement. *Canadian Family Physician*, 61(1), 92.
- Greiver, M., Keshavjee, K., Jackson, D., Forst, B., Martin, K., & Aliarzadeh, B. (2012). Sentinel feedback: Path to meaningful use of EMRS. *Canadian Family Physician*, 58(10), 1168.
- Kahn, M. G., Raebel, M. A., Glanz, J. M., Riedlinger, K., & Steiner, J. F. (2012). A pragmatic framework for single-site and multisite data quality assessment in electronic health record-based clinical research. *Medical Care*, 50(7), S21–S28.

- Keshavjee, K., Williamson, T., Martin, K., Truant, R., Aliarzadeh, B., Ghany, A., & Greiver, M. (2014). Getting to usable EMR data. *Canadian Family Physician, 60*(4), 392.
- Iron, K., & Manuel, D. G. (2007, July). *Quality assessment of administrative data (Quaad): An opportunity for enhancing Ontario's health data* (ICES Investigative Report). Toronto: Institute for Clinical Evaluative Sciences.
- Maydanchik, A. (2007). On hunting mammoths and measuring data quality. *Data Management Review, 6*, 14–15, 41.
- Moeinedin, M., & Greiver, M. (2013). *Implementation of the data presentation tool in a primary care organization*. Presentation to Trillium Research Day, Toronto, June 19, 2013. Retrieved from [http://www.trilliumresearchday.com/documents/2013\\_Moeinedin\\_DPT\\_implementation\\_Final\\_June192013.pdf](http://www.trilliumresearchday.com/documents/2013_Moeinedin_DPT_implementation_Final_June192013.pdf)
- Safran, C., Bloomrosen, M., Hammond, W. E., Labkoff, S., Markel-Fox, S., Tang, P. C., & Detmer, D. (2007). Toward a national framework for the secondary use of health data: An American Medical Informatics Association white paper. *Journal of American Medical Informatics Association, 14*(1), 1–9.
- Thiru, K., Hassey, A., & Sullivan, F. (2003). Systematic review of scope and quality of electronic patient record data in primary care. *British Medical Journal, 326*(7398), 1070–1072.
- Weiskopf, N. G., & Weng, C. (2013). Methods and dimensions of electronic health record data quality assessment: Enabling reuse for clinical research. *Journal of American Medical Informatics Association, 20*(1), 144–151.
- Williamson, T., Natarajan, N., Barber, D., Jackson, D., & Greiver, M. (2013). Caring for the whole practice — the future of primary care. *Canadian Family Physician, 59*(7), 800.
- Wright, A., Pang, J., Feblowitz, J. C., Maloney, F. L., Wilcox, A. R., Ramelson, H. Z., Schneider, L. I., & Bates, D. W. (2011). A method and knowledge base for automated inference of patient problems from structured data in an electronic medical record. *Journal of American Medical Informatics Association, 18*(6), 859–867.