
Research and Applications

Formal representation of patients' care context data: the path to improving the electronic health record

Tiago K. Colicchio , Pavithra I. Dissanayake, and James J. Cimino 

Informatics Institute, University of Alabama at Birmingham, USA

Corresponding Author: Tiago K. Colicchio, PhD, MBA, Informatics Institute, University of Alabama at Birmingham, 1900 University Boulevard, Suite 142, Birmingham, AL 35294-3412, USA (tcolicchio@uabmc.edu)

Received 20 February 2020; Revised 15 May 2020; Editorial Decision 2 June 2020; Accepted 10 June 2020

ABSTRACT

Objective: To develop a collection of concept-relationship-concept tuples to formally represent patients' care context data to inform electronic health record (EHR) development.

Materials and Methods: We reviewed semantic relationships reported in the literature and developed a manual annotation schema. We used the initial schema to annotate sentences extracted from narrative note sections of cardiology, urology, and ear, nose, and throat (ENT) notes. We audio recorded ENT visits and annotated their parsed transcripts. We combined the results of each annotation into a consolidated set of concept-relationship-concept tuples. We then compared the tuples used within and across the multiple data sources.

Results: We annotated a total of 626 sentences. Starting with 8 relationships from the literature, we annotated 182 sentences from 8 inpatient consult notes (initial set of tuples = 43). Next, we annotated 232 sentences from 10 outpatient visit notes (enhanced set of tuples = 75). Then, we annotated 212 sentences from transcripts of 5 outpatient visits (final set of tuples = 82). The tuples from the visit transcripts covered 103 (74%) concepts documented in the notes of their respective visits. There were 20 (24%) tuples used across all data sources, 10 (12%) used only in inpatient notes, 15 (18%) used only in visit notes, and 7 (9%) used only in the visit transcripts.

Conclusions: We produced a robust set of 82 tuples useful to represent patients' care context data. We propose several applications of our tuples to improve EHR navigation, data entry, learning health systems, and decision support.

Key words: electronic health records, knowledge representation, clinical decision support, clinical concepts, clinical documentation

INTRODUCTION

The US health system has recently achieved nationwide electronic health record (EHR) adoption,¹ with large-scale adoption of commercial systems.² These systems were adopted without a much-needed redesign^{3,4} to address critical limitations, such as suboptimal interfaces,^{5,6} overzealous alerts,⁷ and bloated notes;⁸ as a result, unintended consequences have emerged at different levels of the US health system.⁹ Recently proposed solutions to address EHRs' limitations include improvements to EHR design and interoperability,

clinical decision support (CDS), and transferring some data entry to patients.^{10,11} Others suggest that such improvements, if implemented in isolation, may not be sufficient because the EHR would still miss relevant information that is buried in clinical notes or the clinicians' minds.^{12,13} The missing information is what we define as the patient's care context data.

The patient's care context can refer to information about his/her health and care (eg, a patient reports that a symptom is getting worse) or the clinician's interpretations and decisions (eg, current

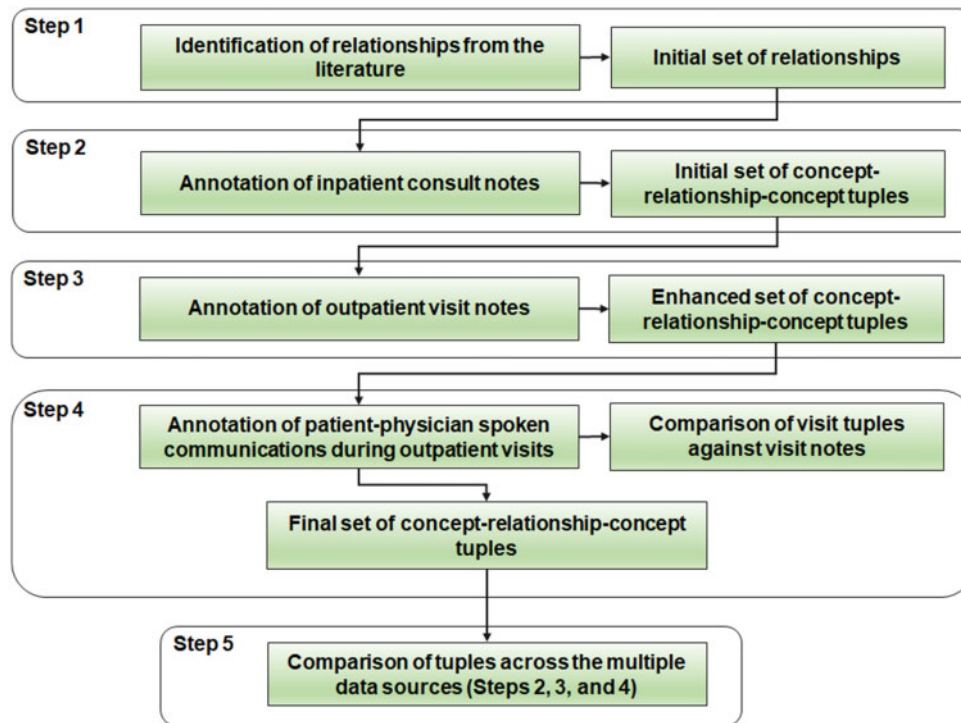


Figure 1. Flowchart illustrating the multi-method approach.

regimen was discontinued because the patient failed antibiotics). An EHR's inability to represent such information is due, in part, to the fact that the problem-oriented structure of most EHRs does not require explicit statements about care decisions, or what Cimino¹⁴ defines as the “why” in the EHR. For example, a treatment can be linked to a problem, indicating that it was chosen to treat the problem, but the underlying reason behind this decision (eg, “why” this particular treatment was chosen) is not captured. We believe that if such information was routinely captured (in computable form, as well as modeled, structured, and coded), it could be used to inform better navigation, CDS, and data entry.⁹ We hypothesize that a formalism with linkages among clinical concepts, or semantic relationships, would be useful to represent patients' care contexts as a semantic network of subject-predicate-object tuples.¹⁵

For instance, the inference “the patient's mastoiditis is from her recurrent ear infection” could be represented with a formalism containing a subject (ear infection), a predicate (suggest_finding), and an object (mastoiditis). By creating such a tuple and adding the corresponding classifications (eg, ear infection = condition), coding (eg, Systematized Nomenclature of Medicine-Clinical Terms [SNOMED CT] ID: 129127001), and modifiers (eg, frequency = recurrent), one could formally represent a clinician's interpretation about a patient's finding, allowing the EHR to use this information for multiple purposes.

We have conducted preliminary work to identify those relationships relevant for creating such tuples. We initially conducted a systematic review to explore how EHRs are used to support documentation of clinicians' reasoning; we found that they provide virtually no automation to support this task, and rely on clinicians' conscious decisions to document patients' care context data in their notes.¹⁶ We then identified an initial sample of relationships from clinical cases published in the literature.¹⁷ Next, we conducted another systematic review to identify ontology properties (eg, con-

cepts, relationships) from CDS systems that incorporate a formal ontology in their logic,¹⁸ and found that although most published ontologies contain relationships useful for CDS, they do not always provide relationships suitable for representing patients' care context. In this study, we expand our sample of relationships^{17,18} by manually annotating patients' care context data as represented in clinical notes and spoken communications during outpatient visits.

Objective

Our objective was to develop a collection of concept-relationship-concept tuples to formally represent patients' care context data to inform future research and EHR development.

MATERIALS AND METHODS

We conducted a multi-method study including the following steps: (1) identify an initial set of relationships from our prior work;^{17,18} (2) annotate inpatient consult notes; (3) annotate outpatient visit notes; (4) annotate spoken communications during outpatient visits and verify whether the tuples from this annotation contain information communicated in the respective notes of these visits; and (5) compare the tuples used across the multiple data sources. These steps are described in detail in the subsequent sections. [Figure 1](#) illustrates the multi-method approach. The University of Alabama-Birmingham (UAB) Institutional Review Board approved the study.

Step 1: identification of relationships from the literature

The 3 authors independently reviewed the 163 relationships from our previous studies,^{17,18} then met in person to reach consensus about an initial sample of relationships that seemed most likely, in our experience, to be useful for representing patients' care context data as reported in clinical notes. We also reached consensus on the

A	B	C	D	E	F	G
#	Sentence (provide context for tuples)	Subject	Subject class	Relationship	Object	Object class
13	He is not felt to be a candidate for any advanced therapies given his poor insight into disease and lack of social support.	he	Patient	(Patient)is_not_eligible	advanced therapies	Intervention
14	He is not felt to be a candidate for any advanced therapies given his poor insight into disease and lack of social support.	advanced therapies	Intervention	(Intervention)has_reason	poor insight into disease, lack of social support	reason
15	Since that time he has not had any more ICD shocks	he	Patient	(Patient)has_absence_of_finding/condition	ICD shocks	Finding/condition
16	Mr. PHI is medically treated for CAD, hypertension, paroxysmal atrial fibrillation	CAD, hypertension, paroxysmal atrial fibrillation	Finding/condition		medically treated	Intervention
17	Pt presented for abdominal pain with HF exacerbation and a renal mass of RCC (planned outpatient follow up) who diuresed but still had crackles on exam	Patient	Patient	(Finding/condition)has_course (Finding/condition)has_intention_of_intervention (Finding/condition)has_intervention (Finding/condition)suggest_finding/condition	abdominal pain with HF exacerbation	Finding/condition

Figure 2. Structured form used in the annotations. Note that the sentence in row 13 is repeated in row 14 because it contains multiple concept pairs. PHI: protected health information. CAD: Coronary Artery Disease; HF: Heart Failure; ICD: implantable cardioverter-defibrillator; RCC: renal cell carcinoma.

class (ie, the classification of a concept as finding, condition, intervention, etc.) of the subject concepts most likely to be associated with these relationships. The initial sample of concepts and relationships was used to start the annotation of inpatient consult notes (Step 2).

Step 2: annotation of inpatient consult notes

We extracted data from 2 note sections—history of present illness (HPI) and assessment and plan (AP)—from inpatient cardiology consult notes written between 1 August 2018 and 31 October 2018 at the UAB Hospital. We chose these sections because they contain the expressivity of patients' care context, which is not typically found in more structured sections, such as “medication list” or “laboratory results.” We decided to start the annotations with consult notes because they are more likely to contain explicit contextual statements, as they reflect specialists' interpretations and recommendations. We chose the specialty of cardiology because it is commonly referred to by attending physicians.

Annotations were conducted in rounds that included 35 sentences extracted from the HPI and AP sections of an initial random sample of 3 notes. We created a structured form for the annotations. First, 1 of the authors (TKC) removed all identifiable information from the content of HPI and AP sections of the selected notes and parsed the content into individual sentences. For each sentence, 2 authors (TKC, JJC) iteratively identified relevant text mentions (hereafter called concepts) to form concept pairs consisting of a subject concept (ie, the first concept in a tuple) and an object concept (ie, the second concept in tuple), along with their classes (eg, intervention, finding), with disagreements resolved via consensus. When a sentence contained more than 1 concept pair, it was repeated as many times as needed. [Figure 2](#) provides an example of the structured form.

All authors independently annotated the concept pairs by filling out the column “Relationship” in the structured form using the relationships from Step 1. The parsed sentences were used to provide context. Subject classes were used to restrict the relationships that can be selected by only allowing the use of relationships associated with the subject class of each row (see example in [Figure 1](#)). When the relationships from Step 1 were deemed inappropriate, the column Relationship was left blank. After completing the annotation of the first round, we calculated inter-rater agreement using a Fleiss Kappa. The authors then met in person to discuss disagreements and identify new relationships and classes not covered by the original sample (Step 1). We searched for equivalent relationships and classes in our previous studies^{17,18} to use them in the next round. For those

not found, we created new relationships/classes via consensus. Subsequent rounds were conducted with the updated list of relationships and classes from the previous round, with the inter-rater agreement recalculated. New notes were added to the annotations until we reached saturation by annotating all concept pairs of a new round using only relationships from the previous round. These annotations resulted in an initial set of concept-relationship-concept tuples.

Since a similar annotation schema is not available in the literature to serve as a model framework, we iteratively created our own guidelines for concept classification and relationship creation and use. Concepts more likely to have an equivalent term in standard vocabularies were generally classified as a finding, condition, intervention, or diagnostic process; when applicable, the Unified Medical Language System (UMLS) was consulted for this task. For other text mentions that provide contextual elements, we iteratively created new classes to classify them (eg, reason: “current regimen discontinued *because patient failed antibiotics*”). Relationships were created and used to indicate how 2 concepts are related (eg, patient-*has_intervention*-intervention). Patients and clinicians were treated as entities and had their own classes (ie, patient, clinician). Due to the conceptual, foundational stage of our schema, some restrictions were applied to facilitate the creation of a starter set of tuples: (1) normalization to standard vocabularies was out of the scope at this stage; (2) possible modifiers were directly added to the relationships (eg, negation: “*has_absence_of_intervention*”); and (3) concept classes were restricted to high-level classes (eg, pharmacotherapy, visit, and diet were all classified as “intervention”). The final version of the guidelines used in the annotation process can be found in the [Supplementary Material](#).

Step 3: annotation of outpatient visit notes

To expand the initial set of tuples, we annotated clinical notes written from outpatient urology and ear, nose, and throat (ENT) specialty visits between 1 March 2019 and 30 June 2019 at the UAB specialty care clinic. We chose secondary specialties (urology and ENT) to contrast with the inpatient specialty from Step 2 (cardiology), thus allowing for an assessment of the applicability of tuples across heterogeneous specialties. The specialties chosen provide a wide range of concepts ranging from common infections to complex surgeries, thus increasing their generalizability.

We used the same process from Step 2 to create a structured form. If we achieved a satisfactory agreement in Step 2 (agreement > 80% and K > .80), annotations were conducted by a single author (TKC) with the other authors consulted as needed; otherwise, they

Table 1. Initial sample of 8 relationships used to start the annotations

ID	Subject class	Semantic relationship	Description	Example
1	Patient	has_finding/ condition ^a	Statement indicating that the patient has present symptom, sign, or condition	The patient reported sore throat
2	Patient	has_history	Statement of fact about a patient that constitutes past medical history	The patient underwent 2 kidney transplants (2002, 2012)
3	Finding/condition	suggest_finding/ condition	Statement inferring that a finding/condition suggests another finding/condition	The pink sputum suggests hemoptysis
4	Finding/condition	has_intervention	Intervention is initiated/executed to treat a present/presumed condition	Levofloxacin was ordered for his pneumonia
5	Finding/condition	has_intention_of_ intervention	Description of the intention to treat a present/presumed condition	Encourage fluid intake
6	Finding/condition	has_course	The course of a finding/condition	Her shortness of breath is becoming worse
7	Intervention	has_reason	Justification/reason for ordering an intervention	She was placed on linezolid and Zosyn based on susceptibilities
8	Intervention	has_effect	Intervention has effect on patient's finding/condition	Patient reported that the prednisone is working to decrease her pain

Note: ^aWe were not able to find an unequivocal distinction between a finding and a condition and decided to merge them.

were conducted in rounds as in Step 2. The number of notes for this step was defined based on the number of notes needed to reach saturation in Step 2.

Step 4: annotation of patient-physician spoken communications

To expand our set of tuples and verify its usefulness to represent data from multiple data sources, we annotated transcripts of patient-physician spoken communications during 5 outpatient visits of 1 specialty from Step 3 (ENT). Upon physician and patient consent, we used a HIPAA-compliant audio recorder to record the entire visit. We transcribed the recordings verbatim, removing any identifiable information and labeling each speaker (eg, Speaker 1, Speaker 2). We used the same methods from previous steps to produce a structure form for annotations. Since spoken communications during clinical visits often involve informal conversations (eg, "hi, my name is ZZZ") and explanations generally not documented in the EHR (eg, clinician explaining to the patient the chemical structure of a cauterizer), we only annotated sentences containing clinical concepts (eg, orders planned, symptoms reported) and their context (eg, course of findings, reason for ordering a medication), because they are more likely to be reported in clinical notes. Annotations followed the same methods from Step 3 (independent annotation by all authors [if no satisfactory agreement was reached in Step 2] or annotation by a single author [TKC; if satisfactory agreement was reached]). The combination of tuples from Steps 2, 3, and 4 formed our final set of tuples.

After annotating the transcripts, we obtained the content of HPI and AP sections of the original visit notes. One of the authors (TKC) parsed the content into individual sentences and compared the concepts in each sentence against the content represented in the tuples used in the annotation of the transcripts. We verified whether these tuples would be useful to represent information documented in visit notes.

Step 5: comparison of the tuples used across the different data sources

Once we produced the combined list of tuples from Steps 2, 3, and 4, we compared the usage of tuples across the different data sources

(inpatient consult notes, outpatient visit notes, and visit transcripts) to identify the tuples more commonly used within and across these sources.

RESULTS

We annotated 626 sentences; of these, 414 were extracted from 18 clinical notes and 212 from 5 visit transcripts. Starting with an initial sample of 8 relationships (Step 1), we annotated 182 sentences from 8 cardiology consult notes (Step 2), producing an initial set of concept-relationship-concept tuples ($n=43$). We annotated 232 sentences from 10 outpatient visit notes (Step 3), producing an enhanced set of tuples ($n=75$). Next, we parsed the transcripts of 5 ENT visits into 212 sentences that were annotated to produce a final set of tuples ($n=82$; Step 4). We compared the tuples used to annotate the transcripts against the parsed sentences from their original visit notes (Step 4) and compared the usage of tuples across the different data sources (Step 5).

Step 1: identification of relationships from the literature

After independently reviewing the 163 relationships from our previous studies,^{17,18} the 3 authors iteratively formulated an initial sample of 8 relationships (Table 1).

Step 2: annotation of inpatient consult notes

During the annotation of inpatient consult notes 8 notes were used, which reached saturation in the fifth round after 182 sentences had been annotated with 43 unique tuples. In the first round, inter-rater agreement was moderately low (57% agreement, $K = .44$), and most disagreements were caused by unclear definitions and uses of the relationships from Step 1. These were resolved via consensus and the agreement of the subsequent rounds was consistently better (round 2: 88% agreement, $K = .85$; round 3: 83% agreement, $K = .81$; round 4: 82% agreement, $K = .80$; and round 5: 86% agreement, $K = .84$). The partial remaining disagreements were resolved via consensus and were mostly due to final adjustments applied to some classes and relationships. The complete list of sentences annotated in Step 2 can be found in [Supplementary Table S1](#).

Step 3: annotation of outpatient visit notes

With the satisfactory agreement achieved in Step 2, we proceeded with the annotations of outpatient visit notes conducted primarily by 1 author (TKC). Since saturation in Step 2 was achieved with 8 notes, we decided to use 10 visit notes (5 urology and 5 ENT notes). The content from the HPI and AP sections of the 10 notes produced 232 sentences that were annotated with 59 tuples; of these, 27 (46%) were identified in Step 2 and 32 (54%) were identified in Step 3. The combination of tuples from Steps 2 and 3 produced an enhanced set of tuples ($n=75$). The complete list of sentences annotated in Step 3 can be found in [Supplementary Table S2](#).

Step 4: annotation of patient-physician spoken communications

The audio recordings of 5 ENT visits were transcribed and parsed into 212 sentences. Note: in some cases, responses to clinicians' questions (eg, [Speaker 1]: Are you on blood thinners? [Speaker 2]: No.) were combined with their preceding question into a single sentence to facilitate the annotations. The 212 sentences were annotated with 50 tuples; of these, 43 (86%) were identified in the previous steps and 7 (14%) in this step.

The combination of tuples from Steps 2, 3, and 4 produced a final list of 82 unique tuples, combining 48 relationships and 14 concept classes. [Table 2](#) lists the top 15 most used relationships and their subject and object classes. [Figure 3](#) illustrates the connections produced by the top 15 most used relationships. The complete list of sentences annotated in Step 4 and the complete list of relationships and their concept classes can be found in [Supplementary Tables S3 and S4](#).

The visit transcripts had, on average, 1516 words, of which 176 (13%) were classified as informal conversations, 573 (39%) were classified as general explanations, and 740 (45%) formed the sentences included in the annotations. We also identified a small portion of questions that were answered with non-verbal communication (eg, nodding, gesture); these were classified as unclear communication and were excluded from the annotations. [Table 3](#) summarizes the audio recordings and visit transcripts.

We parsed the content of the HPI and AP sections of the original notes of the recorded visits into 140 sentences; of these, 103 (74%) contained concepts that had been included in the transcript tuples, 36 were not covered by the tuples, and 1 was inconsistent with what was represented in the tuples. We reviewed the 36 sentences not covered by the tuples and iteratively classified them as "reasoning not verbalized" (ie, clinician did not mention an interpretation added to the note) or "medical record information" (ie, information from the patient's record was not verbalized, but was documented in the note). There were 27 sentences classified as reasoning not verbalized; examples include reasons for performing an assessment ("based on patient's symptoms we performed a nasal endoscopy") and perceptions about a patient's comprehension ("barriers to learning: none evident"). There were 9 sentences classified as medical record information (eg, "complete blood count obtained and hemoglobin/hematocrit—8/26"). We found 2 tuples with information communicated during the visit (eg, "[Speaker 1]: Do you feel any pressure? [Speaker 2]: Yes, on my head") but not documented in the notes (eg, "patient denies pressure"). [Figure 4](#) provides examples of sentences with concepts covered by the tuples. The complete comparison of transcript tuples and note sentences can be found in [Supplementary Table S5](#).

Step 5: comparison of the tuples used across the different data sources

From the 82 tuples in our final set, 20 (24%) were used across all data sources, with the most common being a finding/condition reported by the patient ("I had a really bad nosebleed"), followed by the intention to start an intervention ("recommend beginning IV [intravenous] steroids"). There were 10 (12%) tuples used only in Step 2, which included mostly assessments and recommendations by a specialist. The most common were the absence of improvement of an intervention ("diuresis had no improvement in crackles") and eligibility for an intervention ("he is not a heart failure candidate for transplant"). There were 15 (18%) tuples used only in Step 3, with the most common being the reason for the visit ("patient presents today with urinary incontinence"), followed by a diagnostic process not suggesting a finding/condition ("cystogram showed no leakage"). There were 7 (9%) tuples used only in Step 4, with the most common being the frequency of an intervention ("nasal rinse 3 times a day"), followed by the date of a finding/condition ("last night I had another nosebleed"). [Table 4](#) lists the top 5 used tuples of each data source and across all data sources, and [Figure 5](#) illustrates the distribution of tuples used within and across all sources. The complete list of tuples used can be found in [Supplementary Table S6](#).

DISCUSSION

We report a collection of concept-relationship-concept tuples to represent patients' care context data. The formal representation of patients' care context has already been proposed with the use of relationships found in terminologies such as SNOMED CT.¹² Although SNOMED CT covers relationships of medical knowledge inferences (eg, sinusitis-is_a-disorder of nasal sinus), it does not provide relationships that represent the context surrounding a patient's care (eg, amoxicillin [*for a particular patient*]-has_reason-allergic to ciprofloxacin). Such a representation has also been proposed with the use of clinical element models¹⁹ but has not, to our knowledge, been implemented. The concepts and relationships identified can inform the development of formal knowledge representation schemas with normalization to the UMLS Concept Unique Identifiers, thus increasing their generalizability and application for EHR development.²⁰

Despite the variability of tuples, 459 (73%) sentences were annotated with only 25 (30%) tuples, indicating that with a relatively small combination of concepts and relationships, one can formally represent situations commonly found in clinical encounters, such as interventions ordered (81 instances), findings/conditions present (71 instances), or the intention to start an intervention (69 instances). Although the most used relationships came from the literature, our annotations still produced 31 (65%) new relationships, indicating that specific information reported in clinical notes or spoken communications has not been covered by previous studies.

We do not expect that clinicians or domain experts will have to manually create these tuples upfront. Rather, they are intended to be the foundation of a data infrastructure to improve the EHR. We envision a next-generation EHR capable of processing the patient's care context, represented as a collection of tuples added to what we defined as a patient-specific knowledge base (PSKB). In order to fulfill this vision, some practical, near-term steps for which our starter set of tuples is useful can be performed to build the foundation for more complex, long-term steps to improve EHR functions. The first

Table 2. The top 15 most used relationships and their subject and object classes

#	Subject class(es)	Relationship	Object class(es)	Description	Example	Source
81	Finding/condition Patient	has_intervention	Intervention Diagnostic process	Intervention is initiated/executed to treat a present/presumed finding/condition	The patient began to receive levofloxacin to treat his pneumonia	Prior work ^{17,18}
71	Finding site Patient	has_finding/condition	Finding/condition	Statement indicating that the patient or a specific finding site has a finding/condition	The patient reported sore throat and his physical indicated swollen lymph nodes	Prior work ¹⁸
69	Finding/condition Patient	has_intention_of_intervention	Intervention	Statement indicating the intention to perform an intervention	Avoid dehydration—encourage fluid intake	Prior work ¹⁷
49	Diagnostic process Intervention	has_reason	Finding/condition Intervention Reason	Justification/reason for ordering a procedure or intervention	She was placed on linezolid and Zosyn based on susceptibilities	Prior work ¹⁸
36	Diagnostic process Documentation Finding/condition Intervention	suggest_finding/condition	Finding/condition	Finding/condition suggested by a diagnostic process, documentation, finding/condition, or intervention	The pink sputum suggests hemoptysis	Prior work ^{17,18}
32	Diagnostic process Finding/condition Intervention	has_date	Time	Inference about the time when an event has happened or will happen	She has an appointment with Dr. PHI tomorrow	Prior work ¹⁸
27	Patient	has_history	Finding/condition Intervention	Statement about a patient that constitutes past medical history	The patient underwent 2 kidney transplants (2002, 2012)	Prior work ^{17,18}
20	Intervention	has_effect	Finding/condition	Statement about the effect of an intervention	Patient reported that the prednisone is working to decrease her pain	Prior work ¹⁸
17	Intervention Task	has_trigger	Finding/condition Reason	Statement indicating that an intervention or task will be triggered by an event	Please call if we can be of further assistance	Step 2
16	Patient	has_absence_of_finding/condition	Finding/condition	Statement indicating absence of finding/condition	She is not symptomatic	Step 2
16	Patient	has_visit_reason	Finding/condition Intervention	Reason for an outpatient visit	Patient returns for packing and splint removal	Step 3
12	Finding/condition Intervention	has_frequency	Time	Statement indicating the frequency of a finding/condition or intervention	The nosebleeds have become daily	Step 4
12	Patient	has_intention_of_diagnostic_procedure	Diagnostic process	Statement indicating the intention to perform a diagnostic procedure	Obtain fluid cultures	Step 2
11	Finding/condition	has_course	Course	The course of a finding/condition	Her shortness of breath is becoming worse	Prior work ¹⁸
8	Finding/condition	can_cause_finding/condition	Finding/condition	Statement of medical knowledge about the potential cause of a finding/condition	If systemic release of toxins occurs, a diffuse pustular variant of the toxic shock syndrome can also occur	Prior work ^{17,18}

Note: Relationships are sorted by descending order of tuples in which they were used (column #). The Source column denotes the source of the relationship used during the annotations (ie, if it was found in our previous studies or had to be created via consensus).

PHI: protected health information.

step involves the creation of a formal ontology representing our classes and relationships, which could be done by applying a widely implemented formalism, such as the Ontology Web Language Description Logic;^{21,22} however, the most adequate formalism will likely be defined through the evolution of our starter set of tuples in future studies. As part of this process, normalization to UMLS Concept Unique Identifiers can be performed and sub-classes and modifiers that have not been defined but are implicitly represented in our tuples can be created. Also during this process, our current guidelines can be expanded to cover these additional items, allowing expansion of the ontology to cover other potential contextual elements

and relationships to increase the dimensionality of our schema, which needs to be tested in future empirical studies. Next, the application of the ontology can be tested with a practical use case, such as the creation of a semantic record navigation, with a PSKB populated with a combination of natural language processing and manual chart review. The hypothesis to be tested is that such a solution would be preferred by clinicians when compared to current navigation patterns that require clinicians to access data through isolated EHR components, or to manually find and read bloated notes that frequently contain redundant information and errors.^{23,24} For example, to verify why a medication was ordered for a patient,

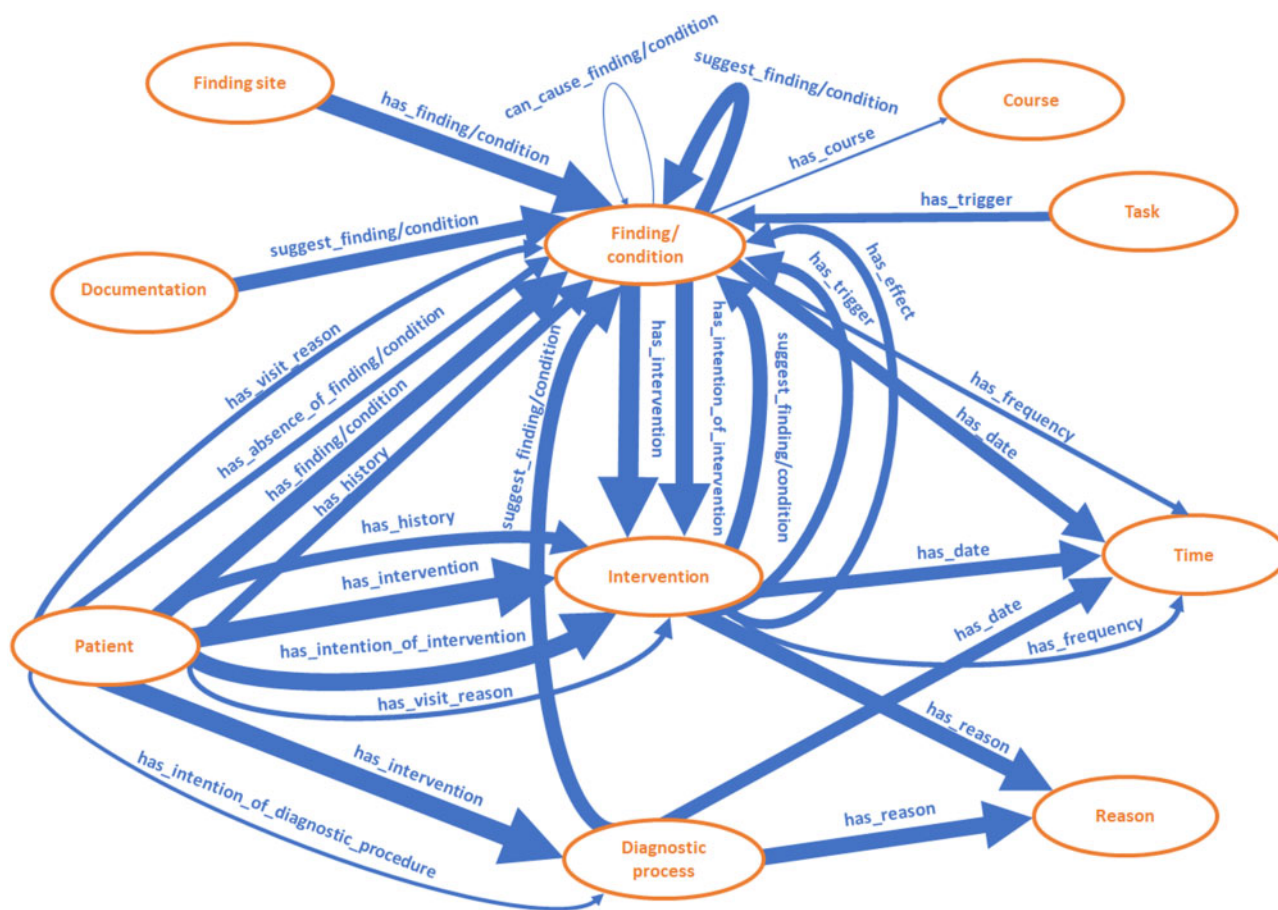


Figure 3. Connections produced by the top 15 most used relationships. The line thickness indicates frequency of use, ranging from 8 to 81.

Table 3. Summarization of the audio recordings and visit transcripts

Visit	Time	Words, Total	Informal Conversation, n (%)	General Explanations, n (%)	Unclear Communication, n (%)	Annotated, n (%)
Visit 1	16	1481	110 (7%)	652 (44%)	82 (6%)	637 (43%)
Visit 2	5	646	5 (0.8%)	403 (62%)	0 (0%)	238 (37%)
Visit 3	23	2115	140 (7%)	1023 (48%)	40 (2%)	912 (43%)
Visit 4	14	1149	551 (48%)	187 (16%)	0 (0%)	411 (36%)
Visit 5	24	2188	74 (3%)	601 (27%)	11 (1%)	1502 (69%)
Average	16.4	1515.8	176 (13%)	573 (39%)	27 (2%)	740 (45%)

Note: The Time column denotes the time in minutes of an active conversation between a clinician and the patient or a family member.

a clinician could access a network of concepts linked to the medication (ie, the medication knowledge graph) and navigate to other parts of the record (eg, to notes mentioning the reason for ordering it). Similar knowledge representation has been successfully applied to other informatics solutions²⁵ and is commonly used in other industries.²⁶ Although some assertions of the knowledge graph might change over time (eg, patient has pain in 1 visit which has resolved in another), each assertion will have a time stamp in the knowledge base, which can potentially be used to verify when a new assertion negates a previous assertion; however, this hypothesis needs to be tested in future studies.

Once the ontology has been created, expanded, and validated, another barrier to implementing this vision needs to be overcome: the need to continuously populate a PSKB in real clinical scenarios. This is challenging because clinicians are the only ones who can tell

what they are thinking about the patients’ problems, and they are already burdened with EHR-associated documentation.²⁷ However, the current challenge may become a future opportunity, which leads us to the potential long-term applications of our schema to improve EHR functions. As demonstrated by our transcript annotations, most concepts documented in visit notes tend to be verbally communicated during a visit; although sometimes this communication uses informal language (eg, “nosebleed” as opposed to “epistaxis”), with the use of a formal ontology these terms could be converted into structured, coded concepts. Based on recent breakthroughs of conversational speech recognition,²⁸ the accurate transcription of spoken communications will likely be accessible in the near future, allowing the application of a PSKB to improve data entry at the point of care. With proper translation methods,^{22,29} an ontology of patients’ care context could be parsed by natural language process-

B	C	D	E	F	G	H	I	J
Visit #		Transcript sentence	Subject	Subject class	Relationship	Object	Object class	Note Sentence
3	1	[Speaker 1]: All right so tell what you had going on.... [Speaker 2]: Sunday night, I had a really bad nose bleed, I mean	Patient	Patient	(Patient)has_visit_reason	really bad nose bleed	Finding/condition	PHI presents following a severe epistaxis episode 01/12/2020 that lasted around 30 minutes with recurrent mild episodes on 01/15/2020 and 01/16/2020
3	2	[Speaker 1]: All right so tell what you had going on.... [Speaker 2]: Sunday night, I had a really bad nose bleed, I mean	really bad nose bleed	Finding/condition	(Finding/condition)has_onset	Sunday night	Time	PHI presents following a severe epistaxis episode 01/12/2020 that lasted around 30 minutes with recurrent mild episodes on 01/15/2020 and 01/16/2020
3	3	[Speaker 2]: How long did you bleed Sunday night? [Speaker 1]: About that, 30 minutes, and it was pouring....	really bad nose bleed	Finding/condition	(Finding/condition)has_duration	30 minutes	Time	PHI presents following a severe epistaxis episode 01/12/2020 that lasted around 30 minutes with recurrent mild episodes on 01/15/2020 and 01/16/2020
3	4	[Speaker 2]: I had another nosebleed, last night, it was the night before, I then I just had a little bit more last night.	Patient	Patient	(Patient)has_finding/condition	I just had a little bit more	Finding/condition	PHI presents following a severe epistaxis episode 01/12/2020 that lasted around 30 minutes with recurrent mild episodes on 01/15/2020 and 01/16/2020
3	5	[Speaker 2]: I had another nosebleed, last night, it was the night before, I then I just had a little bit more last night.	Patient	Patient	(Patient)has_finding/condition	another nosebleed	Finding/condition	Initial bleed was followed by an episode of emesis

Figure 4. Examples of note sentences covered by the transcript tuples. PHI: protected health information.

Table 4. Top 5 tuples in each data source and top 5 tuples used in all data sources

Subject class	Relationship	Object class	Step 2	Step 3	Step 4
Tuples used only in inpatient consult notes					
Intervention	has_absence_of_improvement	Finding/condition	6	–	–
Patient	is_not_eligible	Intervention	5	–	–
Patient	is_admitted	Setting/location	3	–	–
Intervention	has_abandonment_reason	Finding/condition	2	–	–
Clinician	has_action	Task	2	–	–
Tuples used only in outpatient visit notes					
Patient	has_visit_reason	Intervention	–	5	–
Diagnostic process	does_not_suggest_finding/condition	Finding/condition	–	4	–
Diagnostic process	has_date	Time	–	3	–
Patient	is_referred_by	Clinician	–	2	–
Diagnostic process	has_abandonment_reason	Reason	–	1	–
Tuples used only in the visit transcripts					
Intervention	has_frequency	Time	–	–	8
Finding/condition	has_date	Time	–	–	5
Finding/condition	has_frequency	Time	–	–	4
Intervention	has_value	Value	–	–	3
Patient	has_contraindication	Intervention	–	–	2
Tuples used in all data sources					
Patient	has_finding/condition	Finding/condition	20	15	28
Patient	has_intention_of_intervention	Intervention	14	22	15
Patient	has_intervention	Intervention	12	25	9
Finding/condition	has_intervention	Intervention	9	6	10
Diagnostic process	suggest_finding/condition	Finding/condition	14	6	4

ing systems³⁰ for concept extraction and other modern language understanding methods^{31–33} for relation classification, forming the data infrastructure for capturing tuples from real-time transcripts of clinical visits. Contextual concepts not represented in standard terminologies (eg, a finding described as “room spinning sensation”) could be classified into an appropriate concept class (eg, value) and added to the PSKB. Coded concepts communicated during a visit could be used to facilitate structured data entry with either a single click (for items explicitly defined: Azithromycin, by mouth, once a day for 5 days) or 2 clicks (for items not defined: Azithromycin); the latter will demand an additional click to select the exact order. Furthermore, this infrastructure could be used to create content for narrative note sections. This could be done by combining standard sentences describing information commonly communicated in clinical notes with information represented in the PSKB or extracted from the patient’s record. For example, the sentence “Patient is a [Age] [gender] who presents for [type of visit] with [reason for

visit]” could be converted into “Patient is a 45-year-old female who presents for new patient visit with acute bronchitis.” Likewise, other sentences needed to populate HPI and AP sections could be automatically generated as new information is added to the PSKB. Information relevant for the note but not verbalized during a visit could be extracted from the patient’s record (eg, visit type, referring physician) or prepopulated automatically (eg, standard patient education sentences). As the clinicians realize the benefits of this infrastructure, they can be trained to “think aloud” and verbalize as many reasoning elements as possible. In addition to facilitating structured and narrative data entry at the point of care, this infrastructure would automatically populate a PSKB and, as more context is added to the knowledge base through new visits and other sources (eg, notes, problem lists, lab results), the PSKB would continuously capture the patient’s care context to empower several EHR functions in a continuous improvement cycle. Examples of long-term improvements that can be obtained from this data infrastructure in-

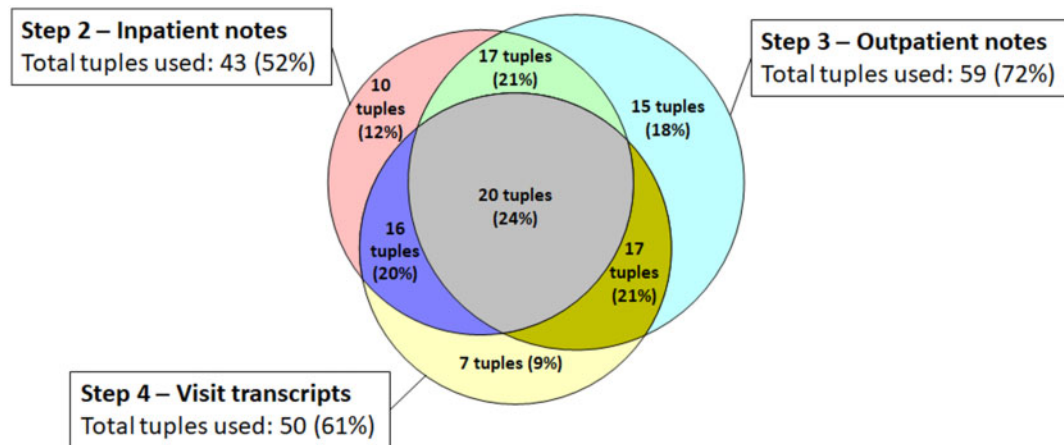


Figure 5. Venn diagram illustrating the number of tuples used within and across all data sources.

clude learning health systems capable of processing contextual elements in analyses of large-scale data sets³⁴ and more accurate CDS systems to potentially mitigate the infamous alert fatigue.³⁵ Imagine CDS systems capable of processing a patient's care context (eg, patient's symptoms are getting worse), and when a clinician orders a new chest x-ray (to compare with a previous test from 3 days ago), the system identifies that the frequently overridden alert asking the clinician to consider not ordering the same test within 3 days does not need to be fired (because now its logic can process the reason behind the clinician's decision: patient's symptoms worse; no alert needed!).

Limitations

We annotated information from only 3 specialties; the annotation of a greater, diverse group of specialties will likely reveal other tuples. However, we found a large number of tuples used across the different specialties and data sources, and they seem to be useful for representing what one would commonly find in a clinical encounter, thus increasing their generalizability. Patients' care context data were extracted and represented manually, which may be insufficient to improve EHR functionality; a formal knowledge representation needs to be developed and its validation for the automatic collection of our proposed tuples needs to be tested in simulated and real clinical scenarios. Our proposed solutions for improving the EHR will depend on improvements to and the application of computational methods that are in constant evolution and may still take several years to mature.

CONCLUSION

Based on data reported in previous studies and the annotation of 626 sentences from inpatient and outpatient notes and spoken communications during outpatient visits, we have identified 82 concept-relationship-concept tuples to represent patients' care context data, with multiple applications for EHR improvements. The tuples include 48 semantic relationships and concepts of 14 distinct classes. We propose several applications of our tuples to improve EHR navigation, data entry, learning health systems, and CDS.

FUNDING

This work was supported by research funds from the Informatics Institute of the University of Alabama at Birmingham.

AUTHOR CONTRIBUTIONS

TKC and JJC conceived this study. TKC collected the data from all sources and parsed the content for annotations. TKC, PID, and JJC conducted the annotations. TKC conducted the data analysis and led the writing of this manuscript, with PID and JJC commenting on subsequent drafts. All authors gave their approval for the final version to be submitted and published.

SUPPLEMENTARY MATERIAL

Supplementary material is available at *Journal of the American Medical Informatics Association* online.

ACKNOWLEDGEMENTS

The authors thank Ayme D. Miles and John D. Osborn for their assistance with data extraction from the clinical data repository and Heather D. Martin for her assistance with ear, nose, and throat visit recordings.

CONFLICT OF INTEREST STATEMENT

None declared.

REFERENCES

1. The Office of the National Coordinator for Health Information Technology. Health IT dashboard. <https://dashboard.healthit.gov/quickstats/quickstats.php> Accessed 27 July 2018.
2. KLAS Research. Hospital EMR market share. <https://klasresearch.com/report/us-hospital-emr-market-share-2018/1279> Accessed 10 May 2019.
3. Wachter R, Goldsmith J. To combat physician burnout and improve care, fix the electronic health record. *J Harv Bus Rev* 2018. <https://hbr.org/2018/03/to-combat-physician-burnout-and-improve-care-fix-the-electronic-health-record>
4. Colicchio TK, Del Fiore G, Cimino JJ. Health information technology as a learning health system: call for a national monitoring system. *Learn Health Sys* 2019; 4 (1): e10207.

5. Howe JL, Adams KT, Hettinger AZ, Ratwani RM. Electronic health record usability issues and potential contribution to patient harm. *JAMA* 2018; 319 (12): 1276–8.
6. Ratwani RM, Savage E, Will A, et al. Identifying electronic health record usability and safety challenges in pediatric settings. *Health Aff (Millwood)* 2018; 37 (11): 1752–9.
7. Ash JS, Sittig DF, Campbell EM, Guappone KP, Dykstra RH. Some unintended consequences of clinical decision support systems. In: proceedings of the AMIA Annual Symposium; November 10–14, 2007; Chicago.
8. Siegler EL, Adelman R. Copy and paste: a remediable hazard of electronic health records. *Am J Med* 2009; 122 (6): 495–6.
9. Colicchio TK, Cimino JJ, Del Fiol G. Unintended consequences of nationwide electronic health record adoption: challenges and opportunities in the post-meaningful use era. *J Med Internet Res* 2019; 21 (6): e13313.
10. Sittig DF, Wright A, Ash J, Singh H. New unintended adverse consequences of electronic health records. *Yearb Med Inform* 2016; 25 (1): 7–12.
11. Payne TH, Corley S, Cullen TA, et al. Report of the AMIA EHR 2020 Task Force on the status and future direction of EHRs. *J Am Med Inform Assoc* 2015; 22: 1102–10.
12. Hsu W, Taira RK, El-Saden S, Kangarloo H, Bui A. Context-based electronic health record: towards patient specific healthcare. *IEEE Trans Inform Technol Biomed* 2012; 16 (2): 228–34.
13. Evans RS. Electronic health records: then, now, and in the future. *Yearb Med Inform* 2016; 25 (S01): S48–61.
14. Cimino JJ. Putting the “why” in “EHR”: capturing and coding clinical cognition. *J Am Med Inform Assoc* 2019; 26 (11): 1379–84.
15. W3C. RDF 1.1 concepts and abstract syntax. World Wide Web Consortium; 2014.
16. Colicchio TK, Cimino J. Clinicians’ reasoning as reflected in electronic clinical note-entry and reading/retrieval: a systematic review and qualitative synthesis. *J Am Med Inform Assoc* 2019; 26 (2): 172–84.
17. Cimino JJ, Li Z, Weng C. An exploration of the terminology of clinical cognition and reasoning. In: proceedings of the AMIA Annual Symposium; November 3–7, 2018; San Francisco.
18. Dissanayake PI, Colicchio TK, Cimino JJ. Using clinical reasoning ontologies to make smarter clinical decision support systems: a systematic review and data synthesis. *J Am Med Inform Assoc* 2020; 27 (1): 159–74.
19. Coyle JF, Mori AR, Huff SM. Standards for detailed clinical models as the basis for medical data exchange and decision support. *Int J Med Inform* 2003; 69 (2-3): 157–74.
20. Scuba W, Tharp M, Mowery D, et al. Knowledge author: facilitating user-driven, domain content development to support clinical information extraction. *J Biomed Semant* 2016; 7 (1): 42.
21. Smith B, Brochhausen M. Putting biomedical ontologies to work. *Methods Inf Med* 2010; 49 (2): 135–40.
22. Hochheiser H, Castine M, Harris D, Savova G, Jacobson RS. An information model for computable cancer phenotypes. *BMC Med Inform Decis Mak* 2016; 16 (1): 121.
23. Hammond KW, Helbig ST, Benson CC, Brathwaite-Sketoe BM. Are electronic medical records trustworthy? Observations on copying, pasting and duplication. In: proceedings of the AMIA Annual Symposium; November 8–12, 2003; Washington, DC.
24. Wrenn JO, Stein DM, Bakken S, Stetson PD. Quantifying clinical narrative redundancy in an electronic health record. *J Am Med Inform Assoc* 2010; 17 (1): 49–53.
25. Hulse NC, Long J, Tao C. Combining infobuttons and semantic web rules for identifying patterns and delivering highly-personalized education materials. In: proceedings of the AMIA Annual Symposium; November 16–20, 2013; Washington, DC.
26. Chah N. OK Google, what is your ontology? Or: exploring freebase classification to understand Google’s Knowledge Graph; 2018. <http://arxiv.org/abs/1805.03885> Accessed 20 February 2019.
27. Downing NL, Bates DW, Longhurst CA. Physician burnout in the electronic health record era: are we ignoring the real cause? *Ann Intern Med* 2019; 170 (3): 216–7.
28. Xiong W, Droppo J, Huang X, et al. Achieving human parity in conversational speech recognition. <http://arxiv.org/abs/1610.05256> Accessed 25 April 2018.
29. Tseytlin E, Mitchell K, Legowski E, Corrigan J, Chavan G, Jacobson RS. NOBLE—flexible concept recognition for large-scale biomedical natural language processing. *BMC Bioinformatics* 2016; 17 (1): 32.
30. Nadkarni PM, Ohno-Machado L, Chapman WW. Natural language processing: an introduction. *J Am Med Inform Assoc* 2011; 18 (5): 544–51.
31. Peng Y, Yan S, Lu Z. Transfer learning in biomedical natural language processing: an evaluation of BERT and ELMo on ten benchmarking datasets. <http://arxiv.org/abs/1906.05474> Accessed 31 January 2020.
32. Alsentzer E, Murphy JR, Boag W, et al. Publicly available clinical BERT embeddings. <http://arxiv.org/abs/1904.03323> Accessed 31 January 2020.
33. Li F, Jin Y, Liu W, Rawat BPS, Cai P, Yu H. Fine-tuning bidirectional encoder representations from transformers (BERT)-based models on large-scale electronic health record notes: an empirical study. *JMIR Med Inform* 2019; 7 (3): e14830.
34. Cabitza F, Rasoini R, Gensini GF. Unintended consequences of machine learning in medicine. *JAMA* 2017; 318 (6): 517–8.
35. Baysari MT, Tariq A, Day RO, Westbrook JI. Alert override as a habitual behavior—a new perspective on a persistent problem. *J Am Med Inform Assoc* 2017; 24 (2): 409–12.