

# Consensus Development of a Modern Ontology of Emergency Department Presenting Problems—The Hierarchical Presenting Problem Ontology (HaPPy)

Steven Horng<sup>1,2</sup> Nathaniel R. Greenbaum<sup>1,2</sup> Larry A. Nathanson<sup>1,2</sup> James C. McClay<sup>3</sup>  
Foster R. Goss<sup>4</sup> Jeffrey A. Nielson<sup>5</sup>

<sup>1</sup> Division of Clinical Informatics, Beth Israel Deaconess Medical Center, Harvard Medical School, Boston, Massachusetts, United States

<sup>2</sup> Department of Emergency Medicine, Beth Israel Deaconess Medical Center, Harvard Medical School, Boston, Massachusetts, United States

<sup>3</sup> Department of Emergency Medicine, College of Medicine, University of Nebraska Medical Center, Omaha, Nebraska, United States

<sup>4</sup> Department of Emergency Medicine, University of Colorado Hospital, University of Colorado School of Medicine, Aurora, Colorado, United States

<sup>5</sup> Northeastern Ohio Medical University, University Hospitals Samaritan Medical Center, Ashland, Ohio, United States

**Address for correspondence** Steven Horng, MD, MMSc, Division of Clinical Informatics, Beth Israel Deaconess Medical Center, Harvard Medical School, 330 Brookline Avenue, Boston, MA 02215, United States (e-mail: shorng@bidmc.harvard.edu).

Appl Clin Inform 2019;10:409–420.

## Abstract

**Objective** Numerous attempts have been made to create a standardized “presenting problem” or “chief complaint” list to characterize the nature of an emergency department visit. Previous attempts have failed to gain widespread adoption as they were not freely shareable or did not contain the right level of specificity, structure, and clinical relevance to gain acceptance by the larger emergency medicine community. Using real-world data, we constructed a presenting problem list that addresses these challenges.

**Materials and Methods** We prospectively captured the presenting problems for 180,424 consecutive emergency department patient visits at an urban, academic, Level I trauma center in the Boston metro area. No patients were excluded. We used a consensus process to iteratively derive our system using real-world data. We used the first 70% of consecutive visits to derive our ontology, followed by a 6-month washout period, and the remaining 30% for validation. All concepts were mapped to Systematized Nomenclature of Medicine—Clinical Terms (SNOMED CT).

**Results** Our system consists of a polyhierarchical ontology containing 692 unique concepts, 2,118 synonyms, and 30,613 nonvisible descriptions to correct misspellings and nonstandard terminology. Our ontology successfully captured structured data for 95.9% of visits in our validation data set.

**Discussion and Conclusion** We present the Hierarchical Presenting Problem ontology (HaPPy). This ontology was empirically derived and then iteratively validated by an expert consensus panel. HaPPy contains 692 presenting problem concepts, each concept being mapped to SNOMED CT. This freely sharable ontology can help to facilitate presenting problem-based quality metrics, research, and patient care.

## Keywords

- ▶ chief complaint
- ▶ presenting problem
- ▶ emergency department
- ▶ emergency medicine
- ▶ ontology

received  
February 27, 2019  
accepted after revision  
April 24, 2019

© 2019 Georg Thieme Verlag KG  
Stuttgart · New York

DOI <https://doi.org/10.1055/s-0039-1691842>.  
ISSN 1869-0327.

## Background and Significance

The precipitating reason for a visit is an important data element that is captured when a patient presents to an emergency department (ED) and is often one of the first questions providers ask a patient. This information is used to guide the patient's initial clinical care, and when aggregated, serves as a valuable tool for understanding patterns of patient visits for administrative and research purposes. Often referred to as a *chief complaint* (CC) or *reason for visit*, there is no formalized definition or required vocabulary for recording this information in ED information systems (EDISs).

The high degree of variability in recording this information both within and between hospitals greatly hampers reuse of this information for clinical, research, or quality purposes. For example, a specific set of symptoms can be recorded by different providers as "chest pain," "CP," or "cardiac pain," making evaluation of data difficult. This is further complicated by the prevalence of unstructured free-text entries, which commonly have misspellings, local abbreviations, and other errors that are unsuitable for use in computerized decision support, information exchange, and for secondary analysis.<sup>1-5</sup> The need for standardization has been well described in the past,<sup>6</sup> but no solution has received wide adoption. Any potential solution must also facilitate clinical workflows and support clinicians, rather than be optimized for secondary reuse. This focus on reducing the unintentional negative effects of electronic health records (EHRs) is important as EHRs are becoming increasingly linked to physician dissatisfaction and burnout.<sup>7-11</sup>

Over the past 15 years, attempts have been made to create a standardized method of recording CCs.<sup>12</sup> Prior attempts have failed to gain widespread acceptance for various reasons: they may not be freely sharable or may not have had the right level of specificity, structure, and clinical relevance to gain acceptance by the larger emergency medicine community.<sup>13</sup> EDIS vendors and other commercial vendors<sup>14</sup> offer vocabularies, but they cannot be used to compare data across EDs who do not have access to these proprietary vocabularies. Other efforts at nonproprietary vocabularies<sup>15,16</sup> were developed, but lacked the granularity for effective secondary use. Natural language processing has been successful for very specific use cases,<sup>17</sup> but has not yielded sufficient sensitivity or specificity when applied more generally.<sup>18</sup> Furthermore, most work in the field is in the secondary use of CCs to classify visits for syndromic

surveillance,<sup>19-28</sup> rather than characterizing a presenting problem (PP) for an ED visit.

To address these issues, the American College of Emergency Physicians Section of Emergency Medicine Informatics was tasked with developing a freely available, standardized vocabulary for the "PP" suitable for use in any ED EHR. We make a distinction between the "CC," defined as the patient's own words<sup>29</sup> and the PP which is a provider's clinical interpretation of the patient's concerns.

When patients present to the ED, they share a reason for their visit with the initial provider. These first spoken words are the CC. While classically taught in medical school that the CC should be in the patient's own words, there has been a shift toward recording the CC using a list of standardized terms recognized by the local EHR. The categorized version of the complaint requires a transformation from a patient's view of the problem to the provider's interpretation of the problem. Unfortunately, this transformed term is still often referred to as "CC," even though it is a distinct new entity. We choose to use the term CC as originally intended, and term the new entity the "PP." The PP is the provider interpretation of the patient's CC (→ **Table 1**). We collect the PP (provider's perspective) in a text field, separate from other free-text. We do not collect a separate CC (patient's perspective), although some EHRs might.

In 2006, a group of 40 stakeholders held a national consensus meeting to develop the framework for a standardized CC vocabulary<sup>13</sup>; a 10-point consensus plan was conceived. We executed portions of the plan to construct an ontology of PPs that can be used by a wide variety of users, leveraging existing standards, and ready for external validation studies.

Ontologies provide a means to represent knowledge in a formal, reproducible, and useful way to capture, store, retrieve, aggregate, reuse, and reason from data.<sup>30</sup> An ontology is composed of (1) terms, which serve as human-understandable labels, (2) concepts, which are abstract representations of a unique entity, and (3) relationships, which link concepts to each other. Each concept has one or more terms. This multiplicity of terms helps capture the synonymy inherent in human language.<sup>31</sup> For example, the concept "dyspnea" could be described using the terms dyspnea, shortness of breath, SOB, difficulty breathing, etc. In our ontology, we utilize "is-a" relationships, which are parent-child relationships to link related concepts, which are essential for aggregating related concepts. For example, the author of a clinical decision support algorithm might want to arrange for dentistry follow-up on

**Table 1** Definition of terms

Term	Definition	Example
Chief complaint	The patient's reason for seeking care in their own words	I can't catch my breath
Reason for visit	The patient's motivation for seeking care	My doctor sent me for a chest radiograph
Presenting problem	A clinician's interpretation of the patient's symptoms	Dyspnea
Clinical syndrome	A constellation of patient's symptoms and demographics	65 years old + male + dyspnea + temp 102
Diagnosis	A condition reached at the end of a medical workup	Pneumonia
Problem list	A list of a patient's chronic and active diagnoses	1. Diabetes 2. Hypertension 3. Pneumonia

patients with dental complaints. Rather than enumerating all dental-related concepts such as “broken tooth,” “dental abscess,” etc., one could instead just specify the single concept “tooth disorder” and aggregate all of its children.

We convened a multicenter physician and nurse expert panel of emergency informatics experts with a goal of creating a PP ontology that fosters clinical decision support, standardized clinical pathways,<sup>32</sup> quality measures,<sup>33,34</sup> and research. It should also support information exchange from prehospital providers<sup>35</sup> as well as with the public health system for syndromic surveillance. Leveraging informatics techniques, the end product is more than just a simple list of words—it also includes the relationships between terms, creating a powerful way to easily group and categorize related visits.

## Objective

Our goal was to develop a standardized PP ontology that would be precise and easy to use while capturing at least 95% of all PPs at the primary institution site.

## Materials and Methods

### Overview

We conducted a quality improvement project over a 3-year period to develop and validate a standardized emergency medicine PP ontology. We call this the Hierarchical Presenting Problem ontology (HaPPy). We collected free-text PPs as well as patient demographics.

### Goals of the Investigation

The following goals motivated the development of this ontology:

- Goal 1—Standardize the collection of ED PPs.
- Goal 2—Develop an ontology that can be easily deployed to existing EHRs.
- Goal 3—Reliably cohort patients for quality measurement, specifically for development of electronic clinical quality measures.<sup>36</sup>
- Goal 4—Trigger PP-based decision support.
- Goal 5—Reliably cohort patients for research and quality improvement.

### Setting and Selection of Participants

The study was performed in a 55,000 visits/year Level I trauma center and tertiary, academic, adults-only, teaching hospital. All consecutive ED patient visits between March 10, 2013, and May 29, 2016, were included in the study. The first 106,695 consecutive visits (70%) were used to derive the ontology, followed by a 6-month washout period, and finally a validation period of 45,687 (30%) patient visits. No visits were excluded. The EHR used was custom developed at the institution.

### Iterative Presenting Problem Development

Our grounded theory approach builds upon Systematized Nomenclature of Medicine—Clinical Terms (SNOMED CT), an internationally developed and maintained hierarchical ontology, by using real-world data captured at the point of

care. It improves over previous list-based approaches and provides a foundation for future PP advances. We employed the four usability requirements for structured documentation described by Rosenbloom et al<sup>37</sup> to create a system that is easy to use and clinically meaningful.

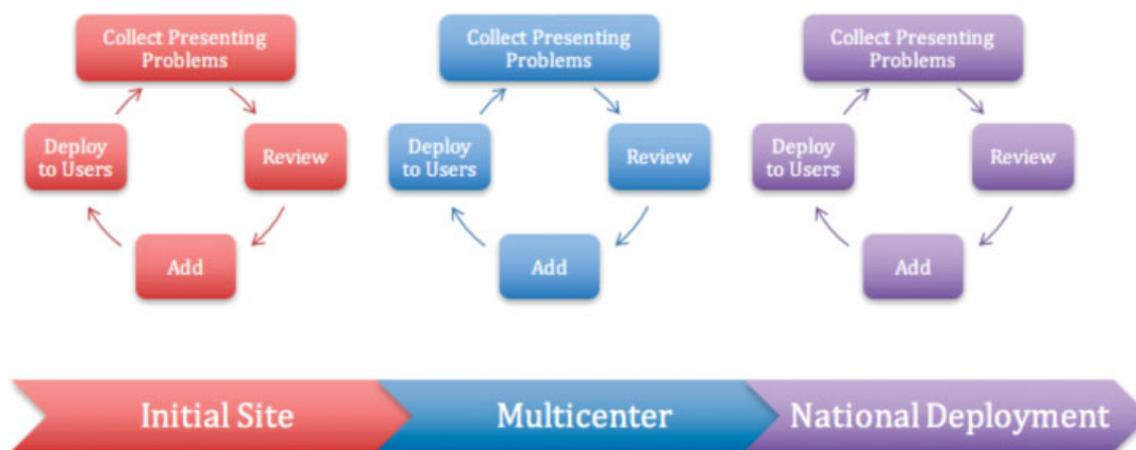
To create the PP vocabulary, we examined all consecutive PPs entered into our EHR from March 10, 2013, to February 11, 2015. After an initial period of data collection, unstructured (free-text) PPs were reviewed by the committee to create an initial PP data set. In an iterative process, free-text entries that were not already part of our ontology were reviewed to identify candidates for inclusion. During each review, free-text PPs not yet in the vocabulary were sorted by frequency with the most frequent PPs considered first for addition. Lexical variants, synonyms, and misspellings were also captured.

Free-text PP terms were normalized and then manually mapped to concepts in SNOMED CT (March 2013) and the U.S. Extension to SNOMED CT (September 2012). The level of granularity, as well as proper mapping to SNOMED CT, was performed using the consensus process defined below. The updated PP list, along with its synonyms and lexical variants, was then redeployed to users. After a brief collection period, this process was repeated until we reached the a priori termination point of 95% coverage, the same threshold chosen by the Clinical Observations Recordings and Encoding (CORE) Problem List Subset of SNOMED CT. This iterative process is illustrated in [Fig. 1](#).

SNOMED CT was used as the foundation vocabulary because it is licensed for use in the United States, freely available, internationally accepted, regularly maintained, and adheres to the principles of a formal ontology. SNOMED provides a structured set of concepts and synonyms, linked via well-defined relationships, allowing for complex computational queries based on these relationships. For example, a search for “abdominal pain” allows the user to query all its child terms (right-sided abdominal pain, left lower quadrant abdominal pain, epigastric pain, etc.) with a single search. SNOMED CT is also a designated terminology for Meaningful Use regulations and the terminology standard for encoding patient problems in EHRs, making it uniquely suited for the development of a PP ontology. Our PP ontology is composed of two parts: (1) a reference set (subset) of the SNOMED CT terminology that represents ED PPs and (2) an interface terminology that consists of all the lexical variants end users can enter into the EHR to express those terms.

In this ontology, we created a subset of SNOMED CT. Although SNOMED CT could have been used in its native form, clinical users would have had difficulty choosing from 399,117 unique concepts. Clinical users are not expert ontologists, and would be unlikely to reliably select from the correct hierarchy. Furthermore, an ontology of 399,117 unique concepts would be difficult for users to learn and use.

This HaPPy ontology subset is similar in approach to the CORE Problem List Subset of SNOMED CT. The CORE Problem list defines a subset of SNOMED CT that facilitates the use of SNOMED CT as the primary coding terminology for problem lists or other summary level clinical documentation. The HaPPy ontology is also a subset, with a different, more narrow



**Fig. 1** Iterative presenting problem development.

use case of capturing PPs in the ED. Though problem list and presenting problem sound similar, they are in fact two distinct entities. Generally speaking, an entry made in the problem list represents the conclusion of a medical workup. A PP, on the other hand, is the provider's interpretation of a patient's symptoms, before any medical workup is complete. This distinction between PP and diagnosis is critical to the motivation underlying this work. This also causes differences in mappings for related concepts. For example, we mapped the concept rash to "complaining of a rash (finding)," while the CORE Problem List maps it to "Pruritic rash (disorder)." Though related concepts, they are semantically different. One is a patient-reported symptom, the other is a disorder. There does exist overlap between the CORE Problem List Subset and the HaPPy subset, as a concept can both be a PP and problem on a problem list simultaneously. Note that 336 (49%) of our concepts overlap with the CORE Problem List. One explanation for this overlap is a common clinical practice to document problems on a problem list using symptoms, until a definitive diagnosis is made. This cognitive forcing strategy of using PPs as diagnoses mitigates the risk of anchoring bias.<sup>38,39</sup>

### Consensus Methodology

The content was developed by a group of eight emergency physicians and emergency nurses at an urban, academic medical centers who also are leaders in clinical informatics. Unanimous consensus was achieved after nine meetings over 2 years, with a combination of in-person meetings at national conferences and teleconference. The group developed a set of rules to determine which problems should be included in the vocabulary and mapped the PP terms to SNOMED CT. The group also developed a set of heuristics to determine whether a problem should be included in the vocabulary.

Heuristic 1—Would the addition of the PP change clinical decision making, workflow, or analysis?

Heuristic 2—Is the level of granularity appropriate for a triage nurse?

Heuristic 3—Is the PP supported by current user behavior?

Heuristic 4—Does the PP reduce ambiguity and improve communication?

The consensus process also determined how to appropriately map each PP to the SNOMED CT hierarchy as the same concept can be represented in multiple hierarchies within SNOMED. In this study, we represent signs and symptoms in the "Clinical Findings" hierarchy, diagnoses in the "Disorders" hierarchy, and events in the "Events" hierarchy. For example, the concept "abdominal pain" was mapped to "Abdominal Pain (finding)," which is in the clinical finding hierarchy with semantic type "Sign or Symptom." Alternatively, it could also have been mapped to "On examination - abdominal pain (finding)," also under the clinical findings hierarchy, but with semantic type finding. Since abdominal pain in this context is most likely a patient's reported symptom of abdominal pain, rather than a triage nurse elicited abdominal pain from a physical exam, it was more appropriate to map this to the former concept. Another alternative would be "abdominal pain characteristic (observable entity)." However, in this context, the abdominal pain would be most likely a patient's reported sign or symptom, rather than the triage nurse observing abdominal pain.

The group used the software package Gephi<sup>40</sup> to help visualize the "is-a" relationships between PPs. We suggest that users use the included .gexf graph file in the distribution to explore the ontology.

We chose to exclude some PPs that were frequently documented at triage based on the application of the above heuristics by the investigators. For example, the term "lethargy" is used colloquially to describe weakness, fatigue, or sleepiness. However, the clinical definition for a physician is a decreased level of consciousness requiring prompt intervention.<sup>41</sup> This discrepancy between the intent of how lethargy is used, and how it is interpreted, leads to much confusion. Similarly, the PP of "arrest" could be used to indicate a cardiac arrest, a respiratory arrest, or that a patient was detained by law enforcement. Omitting these terms reduces ambiguity and improves communication (Heuristic 4).

Additional omissions include items such as "patient referral for consultation," "minor complaint," and "imaging tests." Although these terms may describe why a patient presented to the ED, they do not meaningfully change medical decision making or workflow (Heuristic 1). We, however, did include

the concept “transfer,” to denote when a patient was transferred from an outside hospital, as it does meaningfully change workflow. We also built decision support so that transfer could not be used as the only PP.

### Interface Terminology

We started our interface terminology by manually reviewing all SNOMED CT descriptions for a particular concept to use as synonyms. We also automatically generated several lexical variants. For example, we generated common abbreviations such as UE for upper extremity, Fx for fracture, and Lac for laceration. By proactively generating these terms, we increase the likelihood that a user will find the term they are searching for, improving the amount of structured data captured.

We also automatically generated lateralizing prefixes for concepts that we indicated as having laterality. For example, for the concept “Arm pain” our system would automatically generate “Left arm pain,” “L arm pain,” “Lt arm pain,” “(L) arm pain,” as well as all the permutations with and without the flanking word, “-side,” “sided,” and “-sided.” Laterality was represented in the codes using postcoordination by adding -R, -L, or -B to the end of the concept. This allows developers to easily recover the base concept, as well as the laterality without having to perform a dictionary lookup (→Table 2). Though this may appear that we are creating a new pre-coordinated term, we mean this to be postcoordination only, and use a -R, -L, and -B for usability by EHRs only, as most EHRs are not capable of postcoordination.

To ensure that all anatomic variations for a PP were included, we developed a list of anatomic body parts from all existing PP. We then applied this list of anatomic body parts to existing PPs to uncover potential PPs for inclusion. We manually reviewed this potential list of PP for inclusion and included PP according to the heuristics described above. We repeated this process until no additional terms were discovered.

### Data Analysis

We derived the ontology using the first 106,695 consecutive visits (70%), followed by a 6-month washout period, and finally a validation set of 45,687 (30%) patient visits.

**Table 2** Automated synthesis of lexical variants

Left	L, Lt, (L)	side, -side, sided, -sided
Right	R, Rt, (R)	side, -side, sided, -sided
Bilateral	B, Bilat, Bil, (B)	
Upper extremity	UE	
Lower extremity	LE	
Fracture	Fx	
Injury	Inj	
Abnormal	Abnl, Abn	
Laceration	Lac	
Evaluation	Eval	

One or more PPs can be documented for each patient visit in our EHR. We defined the primary outcome measure as positive if all of the documented PPs listed for the patient were an exact match to a term in our interface terminology.

If any of the PPs were not coded (i.e., the triage nurse used free-text), we considered the outcome to be negative. For example, a PP entered as “Facial injury / Stab wound to the face” would be recorded as negative since “Facial injury” is in our ontology but “Stab wound to the face” is not. This all or none approach provides the most conservative estimate of the PP ontology’s performance.

## Results

### Characteristics of Study Subjects

A total of 180,424 patient visits were included in the study. These patient characteristics are reported in →Table 3.

### HierArchical Presenting Problem ontology

A total of 692 unique PPs were included in our vocabulary. In the validation phase, we found that our PP ontology covered 95.9% of all visits. There were 2,118 synonymous terms that were shown to the user, and an additional 30,614 nonvisible descriptions to correct misspellings and nonstandard terminology that were not displayed to the user (→Table 4).

Relatively few concepts were required to capture the PP for most visits. One-half of all patient visits could be described with just 21 terms, and 75% of visits were described with 58 concepts. Only 121 concepts were required for 90% coverage. For 99.9% coverage, 352 concepts were required, roughly half of our vocabulary. A histogram of PP usage frequency is presented in →Fig. 2. The top 25 PPs are presented in →Table 5 and summary of our ontology by SNOMED Semantic Tag appears in →Table 6. The top level concepts, ordered by number of children, appear in →Table 7.

### Error Analysis

A total of 1,086 entries failed to match. We randomly selected 121 entries (10%) for further analysis. We manually reviewed each of these entries and then normalized them

**Table 3** Patient demographics

	Derivation (n = 106,695; 70%)	Validation (n = 45,687; 30%)
Age – mean years (95% CI)	51.6 (51.5–51.7)	51.8 (51.6–52.0)
Male gender – number (%)	48,618 (45.6%)	20,859 (45.6%)
Severity – median ESI [IQR]	3 [2–3]	3 [2–3]

Abbreviations: CI, confidence interval; ESI, Emergency Severity Index; IQR, interquartile range.

Note: Means with 95% confidence intervals are reported for age and gender. Median and interquartile ranges are reported for Emergency Severity Index (ESI).

**Table 4** A typical presenting problem, synonyms, and nonvisible descriptions

Presenting problem (n = 1)	Displayed synonyms (n = 4)	Nonvisible descriptions (n = 516)
Headache	Headache B Headache L Headache R Headache	H/A HA HEAD ACHE HEADACHE HEADACHES HEAD PAIN MIGRAINE ...(more)... (L) HEAD ACHE (L) HEAD PAIN ...(more)... RT-SIDED HEADACHES RT-SIDED MIGRAINE

by removing ambiguous abbreviations (e.g., “BRADY” -> “Bradycardia”), punctuation (e.g., “\NUMBNESS” -> “Numbness”), or sentence structure (e.g., “FOR EVAL/? SZ” -> “Seizure”) that may have been present. Using a process similar to that described by Zhou et al,<sup>42</sup> we then manually mapped each normalized term to their corresponding SNOMED CT concept if available (→Table 8). Postnormalization, each complaint was scored as an exact match (e.g., aortic stenosis -> aortic stenosis), a partial match (e.g., frontal lobe mass -> brain mass), or, if no match could be found in SNOMED CT, as missing. We then examined why entries could not be matched (→Table 9).

### Proposed SNOMED Additions

We discovered a set of 18 concepts that were not yet in SNOMED CT that we believe should be added to future revisions of SNOMED CT (→Table 10). Three of these concepts have since been added to SNOMED in the 2018 release. We will request inclusion of these terms via the U.S. Edition SNOMED CT Content Request System.

## Discussion

### Pre-/Postcoordination

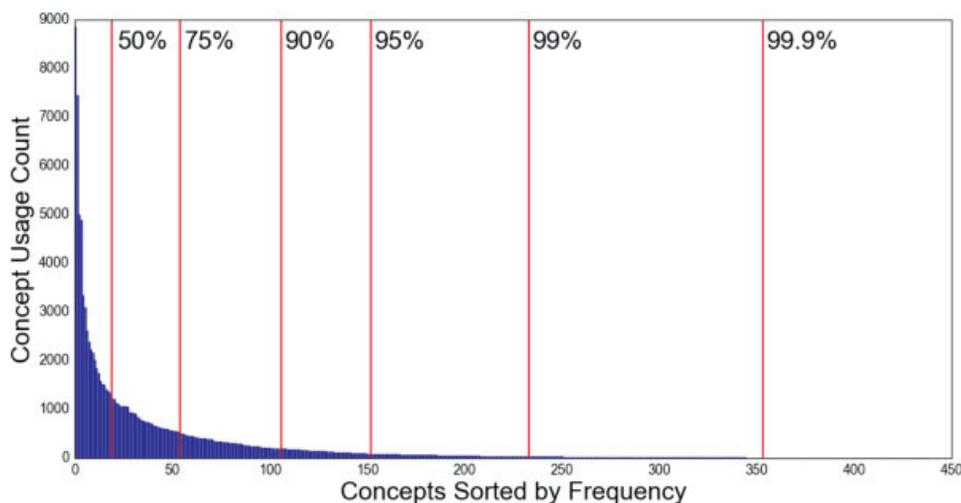
We examined the balance between precoordination and postcoordination and its implications at the point of entry.

With precoordination, a concept can be represented using one single concept identifier while postcoordination consists of two or more concepts that are used to represent a single complaint.<sup>43</sup> For example, entering the complaint of “left-sided chest pain” in a postcoordinated fashion entails the entry of three concepts starting with “pain,” adding the location modifier “chest,” and then a laterality modifier “left.” Alternatively, entry of this concept using precoordination would require selection of just one concept “left side chest pain.” We are not aware of any EDIS that supports postcoordination of PPs. Since our initial goal was to create an ontology that could be used in existing EHRs (Goal 2), postcoordination for anatomical location would be impractical. Furthermore, postcoordination could result in clinically nonsensical concepts, concept duplication, and inefficiency of concept composition.<sup>37</sup> Therefore, we chose to use precoordinated SNOMED CT terms, using postcoordination only to denote laterality. Though laterality may not be important to other users of SNOMED CT, laterality is extremely important in EDs. Wrong side errors is a well-recognized problem in medicine. For example, a patient with right wrist pain is more likely to get the correct X-ray than a patient with a PP of wrist pain.

### Polyhierarchy

We modeled PPs using a hierarchical approach to improve the usability of the terminology in identifying complaints and their relationships to one another. Building off SNOMED CT, we were able to exploit the hierarchy and relationships that have already been established to streamline our implementation.

Structural representation of concepts within the terminology is important both for finding concepts and understanding relationships between concepts. Terminologies are typically structured in either a flat (i.e., list) form or in a hierarchical form. Flat terminologies provide no relationship

**Fig. 2** Distribution of presenting problem concept usage frequency. Vertical lines denote the percentage of patient encounters covered by all concepts left of the line.

**Table 5** Top 25 presenting problems

	Presenting problem	Description	Count (%)	Cumulative frequency
1	Abd pain	Abdominal pain	4,088 (6.3)	6.3%
2	Chest pain		3,738 (5.7)	12%
3	Transfer		2,966 (4.5)	16.5%
4	s/p Fall	Status post fall	2,509 (3.8)	20.4%
5	Dyspnea		2,413 (3.7)	24.0%
6	Headache		1,805 (2.8)	26.8%
7	Fever		1,576 (2.4)	29.2%
8	ETOH	Alcohol intoxication	1,422 (2.2)	31.4%
9	Weakness		1,276 (2.0)	33.3%
10	Dizziness		1,190 (1.8)	35.2%
11	N/V	Nausea and vomiting	1,184 (1.8)	37.0%
12	Back pain		1,128 (1.7)	38.7%
13	SI	Suicidal ideation	1,071 (1.6)	40.3%
14	Cough		1,059 (1.6)	42.0%
15	Wound eval	Wound evaluation	984 (1.5)	43.5%
16	Leg pain		783 (1.2)	44.7%
17	ILI	Influenza-like illness	772 (1.2)	45.9%
18	Flank pain		759 (1.2)	47.0%
19	Lower back pain		752 (1.2)	48.2%
20	MVC	Motor vehicle collision	751 (1.1)	49.3%
21	Altered mental status		714 (1.1)	50.4%
22	Abnormal laboratories		714 (1.1)	51.4%
23	Syncope		639 (1.0)	52.4%
24	Vomiting		605 (0.9)	53.3%
25	Nausea		587 (0.9)	54.2%

**Table 6** Distribution of concept types and their usage

SNOMED Semantic Tag	No. of concepts (n = 457)	No. of uses (n = 112,986)
Disorder	181 (40%)	20,527 (18%)
Event	12 (3%)	2,516 (2%)
Finding	228 (50%)	85,681 (76%)
Procedure	9 (2%)	1,604 (1%)
Other	27 (6%)	2,658 (2%)

Abbreviation: SNOMED, Systematized Nomenclature of Medicine.

information connecting-related concepts. As such, they are of limited utility when searching for concepts that are associated. Conversely, in a hierarchical structure, relationships between concepts are clearly defined and can be used computationally to create complex queries.

Most ontologies defined for biomedicine use a Basic Formal Ontology<sup>44</sup> approach where a monohierarchy is asserted. In a monohierarchy, each concept has exactly one parent, while in a polyhierarchy, each child may have more than one parent. In a monohierarchy, the concept "Right

**Table 7** Top level concepts, top 10 ordered by number of children

Concept ID	Number of children	Max depth	Fully specified name
127278005	73	5	Injury of upper extremity (disorder)
127279002	64	5	Injury of lower extremity (disorder)
312608009	42	5	Laceration - injury (disorder)
22253000	40	6	Pain (finding)
82271004	36	2	Injury of head (disorder)
125666000	23	2	Burn (disorder)
125665001	22	3	Crushing injury (disorder)
418975000	21	3	Bite of animal (event)
128045006	11	2	Cellulitis (disorder)

lower quadrant (RLQ) pain" would be required to have a parent of *either* "Right-sided abdominal pain" *or* "Lower abdominal pain," but not both. Coercing "RLQ pain" into one of these categories would force the ontology to diverge from

**Table 8** Error analysis

	No. of SNOMED CT ( <i>n</i> = 120)
Exact	57
Partial	58
Missing	5

Abbreviation: PP, presenting problem; SNOMED, Systematized Nomenclature of Medicine.

Note: An exact match denotes a PP that can be mapped to a single SNOMED concept that accurately captures the scope and granularity of the PP entered at triage. A partial match means that the PP entered is similar to, but distinct from, existing SNOMED phraseology.

**Table 9** Error analysis

Ambiguous terms	<i>n</i> = 120 (%)	Examples	Matching concept
Lexical variant	41 (34)	Febrile	Fever (finding)
Logical operator	4 (3)	S/P Bee Sting	Bee sting (disorder)
Multicomplaint	13 (11)	SI/SA	Suicidal (finding) Suicide attempt (event)
Abbreviation	11 (9)	FB	Foreign body (disorder)
Malformed	5 (4)	BICYCLIST ST5RUCK	Bicycle accident (event)
Granularity	29 (24)	T9 MASS	Mass of thoracic vertebrae (finding)
Exact match	17 (14)		

**Table 10** Terms that need to be added to SNOMED CT

Concept description	Temporary concept ID	Found in 2018 release
Arm redness	0552bee6c5e811e48aee525400010026	
Hand redness	098fe948c5e811e4899b525400010026	
Finger redness	f178d592c5ea11e4b987525400010026	
Thumb redness	e65a02c6c5ea11e4a4a5525400010026	
Leg redness	02192a6cc5e811e4a5e0525400010026	
Foot redness	d6f96bc0c73711e4b8b2525400010026	
Toe redness	e98210cec5ea11e4a9a0525400010026	
Arm weakness	6dc81860c5e611e4a6ab525400010026	C0751409
Midline eval	850176eeeb6211e48ada525400010026	
Jtube eval	41449a2ec5f111e4933b525400010026	
Positive blood cultures	e3d61a80c5ef11e4a050525400010026	C0740299
Decreased PO intake	943b48f4c5f111e491a2525400010026	
Safe Bed	9665c970c5f011e4b3d7525400010026	
Eye chemical exposure	9b606f0ecb6611e4855e525400010026	
Throat foreign body sensation	9bec4e7ac5e911e480fc525400010026	C3874344
Sickle cell crisis	b18adffcc72511e4a351525400010026	
s/p colonoscopy	c5f2efac5f011e4bda1525400010026	
Infected Fistula	ffcd89e8cb6311e4a958525400010026	

a clinicians' mental model of the PP. The prior probability of disease differs greatly given the location of pain in a patient's abdomen. Therefore, specifying the exact location of abdominal pain is critical. In fact, we noted that users would override the system and enter PPs not found in the ontology when an appropriately specific PP was not found. Although this could be represented as pain, with location RLQ, which is part of the abdomen through postcoordination, most EHRs do not support postcoordination. We discuss the tradeoffs of precoordination and postcoordination more extensively in the previous section. Because of the aforementioned reasons and current limitations of EHRs, we felt strongly that only a polyhierarchical ontology would suit our original goals.

Existing standards, such as International Statistical Classification of Diseases and Related Health Problems (ICD) and Current Procedural Terminology codes, are well suited for public health and billing applications but use monohierarchy; this limits the granularity of a concept as a term can only ever have one parent. Therefore, ICD10 was specifically rejected as the host terminology because it is not polyhierarchical.

Polyhierarchical structures mirror clinical reasoning and, as our ontology expands, provide a framework to add new concepts. In contrast, flat terminologies provide no relationship information connecting various concepts and invariably fail to provide the level of granularity that maximizes utility.

### Data Mapping

Terms were mapped to a reference terminology, SNOMED CT, where they can be described using formal relationships or descriptions (e.g., chest pain "is a" disorder that has a "finding site" of "chest").<sup>37</sup> SNOMED CT is particularly useful

for ED PPs, where the patient's complaint may be a symptom, the name of a disease, a physical finding, or an event.

Owing to its widespread use and robust nomenclature, SNOMED has become the de facto standard for clinical terminology. Mapping to SNOMED enhances data reuse and facilitates translation into non-English languages.

### Granularity and Data Binding

There are two major approaches to data binding when developing a data model. In early binding data models, how data are captured is highly coupled to a predefined set of use cases. In late binding data models, how data are captured is only loosely coupled to how it might be used in the future. Given the rapidly changing requirements of health care analytics, it is inevitable that new analytic use cases will be developed in the future, favoring a late binding data modeling approach. Aggregation can be used to transform a concept from a highly granular concept to a less granular one. It is not possible, however, to accurately go from a less granular concept to a more granular concept. For example, one can infer a patient has the more general concept "pain" if they have the more specific concept "abdominal pain." One cannot, however, infer a patient has the more specific concept "abdominal pain" if they have the more general concept "pain." Since different use cases will require different levels of granularity, and we anticipate novel use cases in the future, we represented PPs as concepts with the highest granularity that is clinically meaningful per our heuristics. Using SNOMED CT relationships, users of the ontology can then aggregate PPs into less granular value sets that are meaningful for a specific use case, whether it be for real-time clinical decision support, clinical research, population health management, or quality improvement.

### Comparative Analysis

There have been various attempts to create coded lists of CCs over the last 4 decades. →Table 11 shows comparison infor-

mation of openly available and published PP lists as well as the PP list described here. Most notable differences are the polyhierarchical structure and mapping to SNOMED CT.

### Limitations

Although we had a multicenter informatics team, our interface terminology was generated from PPs obtained from one tertiary academic medical center. This terminology may not be generalizable to other EDs practicing in other geographical areas or with a different patient population. Extensive external validation across different care settings (urban, rural, etc.) will be needed to further refine this interface terminology.

A major limitation of this work is that it only represents adult PPs, not pediatric PPs in the ED. Pediatric PPs were deliberately excluded as our data did not include pediatric patients.

Terms not present in SNOMED CT will require new concepts to be added and the appropriate mappings to be created. Terminologists at SNOMED CT may differ in their mapping of concepts. Although our terminology adheres to Rosenbloom et al's criteria for an interface terminology, the PPs were generated in an empirical fashion and the balance between pre- and postcoordination was determined by content experts.

Lastly, we recognize that all classification schemes are a product of expert opinion. Like any classification system, ours could unwittingly be manipulated to over- or underrepresent certain situations, events, and conditions. For example, both "snake bite" and "animal bite" appear in our ontology. Had "snake bite" been omitted, it could have falsely underestimated the incidence of these events, which could inadvertently lead to reduced funding for research and antidote development. To help mitigate these challenges, we intend to regularly update and refine our ontology based on community feedback.

### Implementation Suggestions

Whereas our ontology will be updated on an ongoing basis, we recommend that developers store PPs as the text entered by the user during the patient visit, as opposed to translating

**Table 11** Comparative analysis of presenting problem approaches

	CEDIS presenting complaint categories (Canadian standards)	Comprehensive	Emergency care data set (chief complaint)	EMT-P	HaPPy	Structured classification for ED presenting complaints
Initial paper	2008	2001	2015	2004	2016	2012
Latest paper (Update)	2012 (last update)		2016	2007	2016	
Structure	Flat	Flat	Flat	Flat	Polyhierarchical	Flat
Language	EN/FR	EN	EN	EN	EN w/ SNOMED mapping to others	EN
Mapped to/ Taken from	N/A	N/A	SNOMED CT	UMLS	SNOMED CT	ICPC-2
No. of concepts	163	57	133	228	692	89

Abbreviations: CEDIS, Canadian Emergency Department Information System; ED, emergency department; EMT-P, emergency medical technician-paramedic; EN, English; FR, French; HaPPy, Hierarchical Presenting Problem ontology; ICPC, International Classification of Primary Care; SNOMED CT, Systematized Nomenclature of Medicine–Clinical Terms.

the user's input into a SNOMED code. This late binding will allow for retroactive reclassification of complaints as our ontology is refined and matures.

Similarly, we recommend a user interface design that permits concurrent auto-completed items from our ontology to be used alone or in combination with free-text entry from the provider.<sup>45</sup> This hybrid approach will increase the generation of structured data while still enabling providers to enter free-text information that reflects their clinical judgment. These free-text additions should be captured and submitted as candidates for future inclusion in our ontology.

### Analysis Toolkit

It may be challenging for non-informaticians to aggregate PPs into useful groups using SNOMED CT is-a relationships. For example, a common task might be to create a patient demographics table, describing the distribution of PPs for a clinical study. For that use case, one would want to aggregate abdominal pain and all of its children concepts. To facilitate the use of this ontology, we have built an analysis toolkit using Microsoft Excel, analytics software familiar to a wide range of users, and included it in our ontology release.

### Future Directions

The next steps are to deploy this ontology to additional institutions, both retrospectively and prospectively. In the retrospective arm, we would see how well our interface terminology maps previously documented PPs. In the prospective arm, we would deploy the ontology, and after a washout period, we would analyze how well the ontology captures PPs. We would then analyze usage for any terms that did not match, and consider adding them to the ontology, using the same heuristics developed earlier. This could mean either adding new interface terminology, or new concepts. We will continue this iterative development cycle until we reach saturation of concepts.

An important future direction for any ontology is securing organizational and financial support for its ongoing maintenance and curation. Part of this maintenance would be adding the ontology to the National Library of Medicine's Value Set Authority Center.<sup>46</sup> As part of our postpublication evaluation process, we will require users to agree to an annual survey, so that we can better understand how the ontology is being used as well as penetrance.

Another exciting future direction would be to use machine learning to help curate the ontology by suggesting new concepts and interface terminology. We have previously developed machine learning models to predict PPs based on triage data. We already use this to help users input PPs.<sup>45</sup> We could similarly use machine learning to help curate the ontology.

We also plan to work with the relevant national disease registries, quality registries,<sup>47</sup> research networks,<sup>48</sup> EHR vendors, and standards organizations to further refine and develop this ontology.

As existing ontologies for PPs (CCs) already exist, it would be interesting to develop a cross-walk to map this ontology to prior ontologies, as not all ontologies map to a standard ontology such as SNOMED CT.

Lastly, we believe this reference set could be adapted to any language. A possible future direction would be to create an interface terminology for other English language regions such as the United Kingdom or Australia, followed by other languages such as Spanish or French.

### Submitting New Terms and Concepts

The HaPPy ontology contains 692 PP concepts of what is probably many more. In addition to missing PPs in the pediatric population, this ontology could also miss PPs that occur regionally. Users are encouraged to use our Web site <http://bit.ly/2KVZJPP> to submit new interface terminology and concepts for inclusion in future versions of the ontology to make it a living entity.

### Conclusion

We present the Hierarchical Presenting Problem ontology (HaPPy). This ontology was empirically derived and then iteratively validated by an expert consensus panel. HaPPy contains 692 PP concepts, each concept being mapped to SNOMED CT. This freely sharable ontology should help to facilitate PP-based quality metrics, research, and patient care.

### Clinical Relevance Statement

Accurately capturing presenting problems is a vital tool for understanding patterns of patient visits, clinical decision support, quality measures, syndromic surveillance, and research. Our empirically derived, expert-validated system provides the first freely sharable, robust ontology capable of accomplishing this task.

### Multiple Choice Questions

1. ICD-10 and SNOMED CT are two potential base ontologies when creating a new ontology for a specific indication. Which ontology serves as the best base ontology for emergency department presenting problems?
  - a. ICD-10 because it is the standard already used by hospitals to encode diagnoses for billing purposes, while SNOMED CT is not commonly used in hospitals.
  - b. ICD-10 because it is already routinely used in epidemiology and would aid in biosurveillance, while SNOMED CT is not used in biosurveillance.
  - c. ICD-10 because it is hierarchical, so more granular concepts can be grouped into more general concepts, while SNOMED CT is not hierarchical.
  - d. SNOMED CT because it is polyhierarchical, which allows for highly granular terms, while ICD-10 is monohierarchical.
  - e. SNOMED CT because it is international, with interface terminologies in multiple languages, while ICD-10 is not.

**Correct Answer:** The correct answer is option d. SNOMED CT serves as the best base ontology for emergency department presenting problems. It is polyhierarchical, which allows for highly granular terms, unlike ICD-10 which is

monohierarchical which limits granularity. SNOMED CT also describes rich relationships between concepts that ICD-10 lacks. Therefore, options a, b, and c are not the best answer. Answer e is not true, as both SNOMED CT and ICD-10 are international.

2. Why should emergency department presenting problems be standardized?
  - a. Standardization would allow users to enter data more quickly and efficiently.
  - b. Standardization would allow the systematic creation of patient cohorts with the same presenting problems for research, quality measurement, or decision support.
  - c. Standardization would decrease training times for users.
  - d. Standardization would increase the amount of variability in clinical care.

**Correct Answer:** The correct answer is option b. The systematic creation of patient cohorts by presenting problems has been identified as the major motivator for a national ontology of presenting problems. Answer A is incorrect as using standardized ontologies could be both faster or slower for users, depending on the implementation. Answer C is incorrect as more training is required to learn a standardized ontology over no ontology at all. Answer D is incorrect as standardization would decrease the amount of variability in clinical care.

3. What is an interface terminology?
  - a. A terminology that interfaces different systems in an EHR such as between an emergency department system and radiology.
  - b. A set of user-friendly phrases that supports entry of information into a computer.
  - c. A set of relationships between different concepts.
  - d. A semantic triple of subject-predicate object.

**Correct Answer:** The correct answer is option b. This is the definition of an interface terminology.

4. Which of the following is a reason to use precoordination instead of postcoordination for a presenting problem ontology?
  - a. Precoordination allows more granular terms than postcoordination would allow.
  - b. Precoordination allows anatomical location to be prepended to a concept.
  - c. Precoordination would allow multiple synonyms to facilitate natural language.
  - d. Precoordination will not have the potential clinically nonsensical concepts that postcoordination might have.

**Correct Answer:** The correct answer is option d. Postcoordination can potentially create clinically nonsensical concepts. Precoordination would prevent this. Option A is incorrect as pre-/postcoordination does not affect granularity. Option B describes postcoordination not precoordination. Option C describes an interface terminology, not precoordination.

#### Data Availability

As a derivative work of SNOMED CT, the Hierarchical Presenting Problem ontology (HaPPy) is released freely to anyone with a valid SNOMED CT license. The ontology as well as analysis toolkit can be downloaded via our Web site <http://bit.ly/2KVZJPP>.

#### Protection of Human and Animal Subjects

This project was reviewed by the Committee on Clinical Investigations at Beth Israel Deaconess Medical Center and a determination (#2019D000313) was made that this activity did not constitute Human Subjects Research and no further review was required.

#### Funding

Administrative support was partially funded by an American College of Emergency Physicians Section Grant.

#### Conflict of Interest

F.R.G. reported grants from the Agency for Healthcare Research and Quality during the conduct of the study. Dr. Goss provides consulting for RxREVIEW, which develops web-based decision support for prescribing of medications and he receives cash compensation. The other authors report no conflict of interest.

#### Acknowledgments

We would like to acknowledge Stacie Jones for administrative support, as well as Laura Heermann Langford, Kevin Coonan, and Adam Landman for their participation in the consensus process. We would also like to acknowledge Mark Sutherland for the design and development of the Analysis Toolkit built in Microsoft Excel.

#### References

- 1 Macfarlane D. The lexeme hypotheses: their use to generate highly grammatical and completely computerized medical records. *Med Hypotheses* 2016;92:75–79
- 2 Bronnert J, Masarie C, Naeymi-Rad F, Rose E, Aldin G. Problem-centered care delivery: how interface terminology makes standardized health information possible. *J AHIMA* 2012;83(07):30–35
- 3 Boyce RD, Jao J, Miller T, Kane-Gill SL. Automated screening of emergency department notes for drug-associated bleeding adverse events occurring in older adults. *Appl Clin Inform* 2017;8(04):1022–1030
- 4 Kummer BR, Lerario MP, Navi BB, et al. Clinical information systems integration in New York City's first mobile stroke unit. *Appl Clin Inform* 2018;9(01):89–98
- 5 Sundaresan AS, Schneider G, Reynolds J, Kirchner HL. Identifying asthma exacerbation-related emergency department visit using electronic medical record and claims data. *Appl Clin Inform* 2018; 9(03):528–540
- 6 Data Elements for Emergency Department Systems, Release 1.0 (DEEDS): a summary report. *Acad Emerg Med* 1998;5(02): 185–193
- 7 Collier R. Rethinking EHR interfaces to reduce click fatigue and physician burnout. *CMAJ* 2018;190(33):E994–E995
- 8 Factors Affecting Physician Professional Satisfaction and Their Implications for Patient Care, Health Systems, and Health Policy | RAND. Available at: [https://www.rand.org/pubs/research\\_reports/RR439.html](https://www.rand.org/pubs/research_reports/RR439.html). Accessed April 14, 2019

- 9 Babbott S, Manwell LB, Brown R, et al. Electronic medical records and physician stress in primary care: results from the MEMO Study. *J Am Med Inform Assoc* 2014;21(Suppl 1):e100–e106
- 10 Roberts DL, Shanafelt TD, Dyrbye LN, West CP. A national comparison of burnout and work-life balance among internal medicine hospitalists and outpatient general internists. *J Hosp Med* 2014;9(03):176–181
- 11 Hose B-Z, Hoonakker PLT, Wooldridge AR, et al. Physician perceptions of the electronic problem list in pediatric trauma care. *Appl Clin Inform* 2019;10(01):113–122
- 12 Travers DA, Haas SW. Unified medical language system coverage of emergency-medicine chief complaints. *Acad Emerg Med* 2006;13(12):1319–1323
- 13 Haas SW, Travers D, Tintinalli JE, et al. Toward vocabulary control for chief complaint. *Acad Emerg Med* 2008;15(05):476–482
- 14 Medical Terminology Software - IMO Intelligent Medical Objects. Available at: <https://www.e-imo.com/>. Accessed March 27, 2019
- 15 Aronsky D, Kendall D, Merkley K, James BC, Haug PJ. A comprehensive set of coded chief complaints for the emergency department. *Acad Emerg Med* 2001;8(10):980–989
- 16 Thompson DA, Eitel D, Fernandes CMB, Pines JM, Amsterdam J, Davidson SJ. Coded chief complaints—automated analysis of free-text complaints. *Acad Emerg Med* 2006;13(07):774–782
- 17 Jones BE, South BR, Shao Y, et al. Development and validation of a natural language processing tool to identify patients treated for pneumonia across VA emergency departments. *Appl Clin Inform* 2018;9(01):122–128
- 18 Handly N, Thompson DA, Li J, Chuirazzi DM, Venkat A. Evaluation of a hospital admission prediction model adding coded chief complaint data using neural network methodology. *Eur J Emerg Med* 2015;22(02):87–91
- 19 Silva JC, Shah SC, Rumoro DP, et al. Comparing the accuracy of syndrome surveillance systems in detecting influenza-like illness: GUARDIAN vs. RODS vs. electronic medical record reports. *Artif Intell Med* 2013;59(03):169–174
- 20 Wagner MM, Espino J, Tsui FC, et al. Syndrome and outbreak detection using chief-complaint data—experience of the Real-Time Outbreak and Disease Surveillance project. *MMWR Suppl* 2004;53:28–31
- 21 Tsui F-C, Espino JU, Dato VM, Gesteland PH, Hutman J, Wagner MM. Technical description of RODS: a real-time public health surveillance system. *J Am Med Inform Assoc* 2003;10(05):399–408
- 22 Rha B, Burrer S, Park S, Trivedi T, Parashar UD, Lopman BA. Emergency department visit data for rapid detection and monitoring of norovirus activity, United States. *Emerg Infect Dis* 2013;19(08):1214–1221
- 23 Centers for Disease Control and Prevention (CDC). Monitoring health effects of wildfires using the biosense system—San Diego County, California, October 2007. *Morb Mortal Wkly Rep* 2008;57(27):741–744
- 24 Dara J, Dowling JN, Travers D, Cooper GF, Chapman WW. Evaluation of preprocessing techniques for chief complaint classification. *J Biomed Inform* 2008;41(04):613–623
- 25 Chapman WW, Dowling JN, Wagner MM. Classification of emergency department chief complaints into 7 syndromes: a retrospective analysis of 527,228 patients. *Ann Emerg Med* 2005;46(05):445–455
- 26 Mikosz CA, Silva J, Black S, Gibbs G, Cardenas I. Comparison of two major emergency department-based free-text chief-complaint coding systems. *MMWR Suppl* 2004;53:101–105
- 27 Niiranen ST, Yli-Hietanen JM, Nathanson LA. Toward reflective management of emergency department chief complaint information. *IEEE Trans Inf Technol Biomed* 2008;12(06):763–767
- 28 Ferraro JP, Ye Y, Gesteland PH, et al. The effects of natural language processing on cross-institutional portability of influenza case detection for disease surveillance. *Appl Clin Inform* 2017;8(02):560–580
- 29 Bickley LS, Szilagyi PG, Hoffman RM. Bates' Guide to Physical Examination and History Taking. 12th ed. Philadelphia:Wolters Kluwer
- 30 Cimino JJ. In defense of the desiderata. *J Biomed Inform* 2006;39(03):299–306
- 31 Liu K, Hogan WR, Crowley RS. Natural Language Processing methods and systems for biomedical ontology learning. *J Biomed Inform* 2011;44(01):163–179
- 32 Lavelle J, Schast A, Keren R. Standardizing care processes and improving quality using pathways and continuous quality improvement. *Curr Treat Options Pediatr* 2015;1(04):347–358
- 33 Griffey RT, Pines JM, Farley HL, et al. Chief complaint-based performance measures: a new focus for acute care quality measurement. *Ann Emerg Med* 2015;65(04):387–395
- 34 Schuur JD, Hsia RY, Burstin H, Schull MJ, Pines JM. Quality measurement in the emergency department: past and future. *Health Aff (Millwood)* 2013;32(12):2129–2138
- 35 Martin TJ, Ranney ML, Dorroh J, Asselin N, Sarkar IN. Health information exchange in emergency medical services. *Appl Clin Inform* 2018;9(04):884–891
- 36 eCQMs | eCQI Resource Center. Available at: <https://ecqi.healthit.gov/ecqms>. Accessed March 27, 2019
- 37 Rosenbloom ST, Miller RA, Johnson KB, Elkin PL, Brown SH. Interface terminologies: facilitating direct entry of clinical data into electronic health record systems. *J Am Med Inform Assoc* 2006;13(03):277–288
- 38 Croskerry P. Cognitive forcing strategies in clinical decision making. *Ann Emerg Med* 2003;41(01):110–120
- 39 Daniel M, Khandelwal S, Santen SA, Malone M, Croskerry P. Cognitive debiasing strategies for the emergency department. *AEM Educ Train* 2017;1(01):41–42
- 40 Gephi - The Open Graph Viz Platform. Available at: <https://gephi.org/>. Accessed March 27, 2019
- 41 Walker HK, Hall WD. *Clinical Methods: The History, Physical and Laboratory Examinations*. Oxford: Butterworth-Heinemann; 1990:1087
- 42 Zhou L, Plasek JM, Mahoney LM, Chang FY, DiMaggio D, Rocha RA. Mapping Partners Master Drug Dictionary to RxNorm using an NLP-based approach. *J Biomed Inform* 2012;45(04):626–633
- 43 Goss FR, Zhou L, Plasek JM, et al. Evaluating standard terminologies for encoding allergy information. *J Am Med Inform Assoc* 2013;20(05):969–979
- 44 Arp R, Smith B, Spear AD. *Building Ontologies with Basic Formal Ontology*. Cambridge, MA: MIT Press; 2015:244
- 45 Greenbaum NR, Jernite Y, Halpern Y, et al. Contextual autocomplete: a novel user interface using machine learning to improve ontology usage and structured data capture for presenting problems in the emergency department. *bioRxiv* 2017. Doi: 10.1101/127092
- 46 Value Set Authority Center. Available at: <https://vsac.nlm.nih.gov/>. Accessed April 3, 2019
- 47 ACEP // CEDR Home. Available at: <https://www.acep.org/cedr/>. Accessed March 27, 2019
- 48 Deakynne Davies SJ, Grundmeier RW, Campos DA, et al; Pediatric Emergency Care Applied Research Network. The pediatric emergency care applied research network registry: a multicenter electronic health record registry of pediatric emergency care. *Appl Clin Inform* 2018;9(02):366–376