

# Consensus Development of a Modern Ontology of Emergency Department Presenting Problems – the Hierarchical Presenting Problem ontology (HaPPy)

Steven Horng, MD, MMSc 1  
Nathaniel R. Greenbaum, MD 1  
Larry A. Nathanson, MD 1  
James C McClay, MD, MS 2  
Foster R. Goss, DO, MMSc 3  
Jeffrey A. Nielson, MD, MS 4

- 1: Beth Israel Deaconess Medical Center — Harvard Medical School — Department of Emergency Medicine Boston, MA, USA
- 2: University of Nebraska Medical Center College of Medicine — Department of Emergency Medicine, Omaha, NE, USA
- 3: University of Colorado Hospital — University of Colorado School of Medicine - Department of Emergency Medicine, Boulder and Aurora, CO, USA
- 4: Northeastern Ohio Medical University — Western Reserve Hospital, Cuyahoga Falls and Rootstown, OH, USA

## Corresponding Author:

Steven Horng, MD MMSc  
330 Brookline Ave.  
Boston, MA 02215  
shorng@bidmc.harvard.edu  
617.754.2527  
617.278.8164

**Keywords:** chief complaint, presenting problem, Emergency Department, Emergency Medicine, ontology

**Word count:** 3285

## 0. 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 none were freely sharable and contained 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 112,612 consecutive emergency department patient encounters at an urban, academic, Level I trauma center. No patients were excluded. We used a modified Delphi consensus process to iteratively derive our system using real-world data. We used the first 95% of encounters to derive our ontology; the remaining 5% for validation. All concepts were mapped to SNOMED-CT.

**Results:** Our system consists of a polyhierarchical ontology containing 690 unique concepts, 2,113 synonyms, and 30,605 non-visible descriptions to correct misspellings and non-standard terminology. Our ontology successfully captured structured data for 95.8% of visits in our validation dataset.

**Discussion and Conclusion:** We present the Hierarchical Presenting Problem ontology (HaPPy). This ontology was empirically derived then iteratively validated by an expert consensus panel. HaPPy contains 690 presenting problem concepts, each concept being mapped to SNOMED-CT. This freely sharable ontology should help to facilitate presenting problem based quality metrics, research, and patient care.

# 1. INTRODUCTION

## 1.1 Background and Significance

The precipitating reason for an emergency department (ED) visit is an important piece of information and 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 encounters for administrative and research purposes.

However, the high degree of variability both within and between hospitals greatly hampers this function. For example, a specific set of symptoms can be recorded by different providers as "chest pain", "CP" and "cardiac pain", making analysis considerably more difficult. This is further complicated by the prevalence of unstructured text entries, which commonly have misspellings and other errors that are unsuitable for decision support and secondary analysis.[1],[2] The need for standardization has been well described in the past[3], but no solution has received wide adoption.

Over the last 15 years attempts have been made to create a standardized method of recording chief complaints[4]. 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[5]. Emergency department information system (EDIS) vendors and other commercial vendors[6] offer vocabularies, but they cannot be used to compare data across EDs who do not have access to these proprietary vocabularies. Other efforts at non-proprietary vocabularies[7],[8] were developed, but lacked the granularity for effective secondary use. Attempts to solve this by computerized natural language processing did not yield sufficient sensitivity and specificity.[9]

To address these issues, we developed a freely available, standardized vocabulary for the "presenting problem" (PP) suitable for use in any ED electronic health record (EHR). We make a distinction between the "chief complaint" (CC), defined as the patient's own words[10] and the PP which is a provider's clinical interpretation of the patient's symptoms.

When patients present to the emergency department (ED), they share a reason for their visit with the initial provider. These first spoken words are the chief complaint (CC). While classically taught in medical school that the CC should be in the patient's own words, there has been a shift towards recording the CC using a list of standardized concepts recognized by the local electronic health record (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 "chief complaint", even

though it is a distinct new entity. We choose to use the term chief complaint as originally intended, and term the new entity the “presenting problem” (PP). The PP is the provider interpretation of the patient’s chief complaint.

For this work, we convened a multi-year, multi-center physician and nurse expert panel of emergency informatics experts with a goal of creating a PP vocabulary that fosters clinical decision support, quality measures [11], syndromic surveillance, and research. We executed portions of a framework developed at a national consensus meeting in 2006[5]. Leveraging informatics techniques, the end product is more than just a simple list of words — it also include the interrelationships between terms. This structure is known as an “ontology” and means that concepts such as “broken tooth” and “dental abscess” are understood by the computer to be related to each other as a “tooth disorder” concept. This approach provides a very powerful way to easily group and categorize related visits.

## **1.2 Objective**

Our goal was to develop a consensus standardized presenting problem terminology that would be precise and easy to use while capturing at least 95% of all presenting problems at the primary institution site.

# **2. MATERIALS AND METHODS**

## **2.1 Overview**

We conducted a quality improvement project over a 2 year period to develop and validate a standardized emergency medicine presenting problem ontology. We call this the Hierarchical Presenting Problem ontology (HaPPy). We collected free text presenting problems as well as patient demographics.

## **2.2 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 3/10/2013 and 3/22/2015 were included in the study. No visits were excluded. The EHR used was custom developed at the institution.

## **2.3 Iterative Presenting Problem Development**

To create the presenting problem vocabulary, we examined all presenting problems entered into our EHR from 10 March 2013 to 22 March 2015. After an initial period of data collection, unstructured (free text) presenting problems were reviewed by the committee to create an initial PP dataset. In an iterative process, free text entries that were not part of our ontology were reviewed to identify candidates for inclusion. During each review, free text presenting problems not yet in the vocabulary were sorted by frequency with the most frequent presenting problems considered first for addition. Lexical variants, synonyms, and misspellings were also captured.

Free text presenting problem terms were normalized and then manually mapped to concepts in SNOMED CT (March 2013) and the US 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 presenting problem list, along with its synonyms and lexical variants, was then re-deployed to users. After a brief collection period, this process was repeated until we reached the a priori termination point of 95% coverage. This iterative process is illustrated in **Figure 1**.

SNOMED CT was used as the foundation vocabulary because it is licensed for use in the United States, freely available, internationally accepted, and regularly maintained. As a true ontology, terms are 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 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 presenting problems and 2) an interface terminology that consists of all the lexical variants end-users can enter into the EHR to express those terms.

## 2.4 Consensus Methodology

The content was developed by a group of 8 emergency physicians and nurses who also are leaders in clinical informatics. The group developed a set of rules to determine which problems should be included in the vocabulary and used a modified Delphi process to achieve consensus. This approach was also used to map the PP terms to SNOMED CT. The group developed a set of heuristics to determine whether a problem should be included in the vocabulary.

**Heuristic 1 - Would the addition of the presenting problem change clinical thinking, workflow, or analysis?**

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

### **Heuristic 3 - Is the presenting problem supported by current user behavior?**

### **Heuristic 4 - Does the presenting problem reduce ambiguity and improve communication?**

The consensus process also determined how to appropriately map each presenting problem 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. The group used the software package Gephi[12] to help visualize the relationships between presenting problems.

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[13]. This discrepancy between the intent of how lethargy is used, and how it is interpreted, leads to much confusion. Similarly, the presenting problem 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 emergency department, they do not meaningfully change medical decision making or workflow (Heuristic 1).

## **2.5 Interface Terminology**

To expedite interface terminology generation, we 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 post-coordination 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 (**Supplement Table 1**).

In order 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 PP's to uncover potential PP's 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.

## 2.6 Data Analysis

We divided our data into two sets of records: a derivation set of 106,695 (95%) and validation set of 5,919 (5%) patient visits.

One or more presenting problems can be documented for each patient visit in our EHR. We defined the primary outcome measure as positive if all of the documented presenting problems listed for the patient were able to be automatically encoded by our ontology.

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

## 3. RESULTS

### 3.1 Characteristics of study subjects

A total of 112,612 patient visits were included in the study. These patient characteristics are reported in **Table 1**.

**Table 1: Patient Demographics**

	<b>Derivation (n=106,695; 95%)</b>	<b>Validation (n=5,919; 5%)</b>
<b>Age – mean yrs. (95% CI)</b>	54.1 (54.0 - 54.2)	53.4 (52.8 - 53.9)
<b>Male gender – number (%)</b>	48,611 (45.6%)	2709 (45.8%)
<b>Severity – median ESI [IQR]</b>	3 [2-3]	3 [2-3]

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

## 3.2 Hierarchical Presenting Problem ontology (HaPPy)

A total of 690 unique presenting problems were included in our vocabulary (Appendix 1). In the validation phase, we found that our presenting problem ontology covered 95.8% of all visits. There were 2,113 synonymous terms (Appendix 2) that were shown to the user, and an additional 30,605 non-visible descriptions (Appendix 3) to correct misspellings and non-standard terminology that were not displayed to the user. **(Table 2)**

**Table 2: A Typical Presenting Problem, Synonyms, and Non-visible Descriptions**



Presenting Problem [1]	Displayed Synonyms [4]	Non-visible Descriptions [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

Relatively few concepts were required to capture the presenting problem for most visits. One-half of all encounters could be described with just 19 terms, and 75% of visits were described with 54 concepts. Only 106 concepts were required for 90% coverage. For 99.9% coverage, 355 concepts were required, roughly half of our vocabulary. A histogram of presenting problem usage frequency is presented in **Figure 2**. The Top 25 presenting problems are presented in **Table 3** and summary of our ontology by SNOMED Semantic Tag appears in **Supplement Table 2**.

**Table 3: Top 25 Presenting Problems**

Presenting Problem	Description	Count (%)
Abd pain	Abdominal Pain	8858 (7.8%)
Chest pain		7447 (6.6%)
s/p Fall	Status Post Fall	4998 (4.4%)
Dyspnea		4881 (4.3%)
Headache		3338 (3.0%)
Fever		3101 (2.7%)
Back pain		2614 (2.3%)
Weakness		2377 (2.1%)
ETOH	Alcohol Intoxication	2219 (2.0%)
Dizziness		2165 (1.9%)
N/V	Nausea & Vomiting	2019 (1.8%)
SI	Suicidal Ideation	1839 (1.6%)
Cough		1732 (1.5%)
Altered mental status		1585 (1.4%)
MVC	Motor Vehicle Collision	1510 (1.3%)
Flank pain		1491 (1.3%)
Lower back pain		1397 (1.2%)
Leg pain		1373 (1.2%)
Syncope		1343 (1.2%)
Transfer		1222 (1.1%)
Wound eval	Wound Evaluation	1211 (1.1%)
Knee pain		1120 (1.0%)
Sore throat		1110 (1.0%)
BRBPR	Bright Red Blood Per Rectum	1058 (0.9%)
Palpitations		1057 (0.9%)

### 3.3 Error analysis

A first pass error analysis of our ontology yielded 168 entries that were not mapped automatically to SNOMED CT. We manually reviewed each of these complaints and then normalized them 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[14], we then manually mapped each normalized term to their corresponding SNOMED CT concept if available (**Supplement Table 3**). Post-normalization, 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 4**).

**Table 4. Analysis of Non-matching Entries**

<b>Ambiguous Terms</b>	<b>n=55 (%)</b>	<b>Examples</b>
Lexical Variant	7 (13%)	30 foot fall, 30ft fall
Logical Operator	11 (20%)	S/p, r/o
Multi-complaint	16 (29%)	Cardiac tamponade, transfer
Abbreviation	12 (22%)	C2 7x, cz fx
Malformed	9 1(6%)	For eval/?SZ
Granularity	9 (16%)	Frontal Lobe mass

### **3.4 Proposed SNOMED Additions**

We discovered a set of 12 concepts that were not yet in SNOMED that we believe should be added to future revisions of SNOMED: arm redness, decreased PO intake, eye chemical exposure, finger redness, hand redness, jaw numbness, jejunostomy evaluation, leg redness, midline catheter evaluation, status post colonoscopy, safe bed, and thumb redness.

## **4. DISCUSSION**

### **Historical Context**

In 2006, a group of 40 stakeholders held a national consensus meeting to develop the framework for a standardized CP vocabulary[5]; a ten-point consensus plan was conceived. We executed portions of the plan set forth by this work to construct a controlled vocabulary of presenting problems that can be used by a wide variety of customers, leverages existing standards, and is ready for external validation studies.

## Usability

Our grounded theory approach builds upon 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.[15] to create a system that is easy to use and clinically meaningful.

## Pre/Post coordination

We examined the balance between pre-coordination and post-coordination and its implications at the point of entry.

With pre-coordination, a concept can be represented using one single concept identifier while post-coordination consists of two or more concepts that are used to represent a single complaint[16]. For example, entering the complaint of “left sided chest pain” in a post-coordinated 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 pre-coordination would require selection of just one concept “left side chest pain”. Ideally, an interface terminology has pre-coordinated terms for the commonly entered concepts and is able to support post-coordination for finer grained concepts. In general, we preferred using pre-coordination if a suitable SNOMED CT term already existed, and used post-coordination to denote laterality.

## Hierarchy

We modeled presenting problems 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 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.

Specifically, we selected a polyhierarchical approach. Unlike a tree structure, where each concept has a single parent, in a polyhierarchy each child may have more than one parent. In a tree structure the concept ‘Right lower quadrant 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 clinical reasoning.

Existing standards, such as International Statistical Classification of Diseases and Related Health Problems (ICD) and Current Procedural Terminology (CPT) codes, are well suited for public health and billing applications but use a tree structure hierarchy; 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”).[15] SNOMED CT is particularly useful for emergency department presenting problems, 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 re-use and facilitates translation into non-english languages.

## Comparative Analysis

There have been various attempts to create coded lists of chief complaints over the last 4 decades. **Supplement Table 4** shows comparison information 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.

## 4.1 Limitations

Although we had a multi-center informatics team, our interface terminology was generated from presenting problems obtained from one tertiary academic medical center. This terminology may not be generalizable to other emergency departments practicing in another geographical areas or with a different patient population. Pediatric complaints 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 presenting problems were generated in an empirical fashion and the balance between pre- and post-coordination 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 under represent 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.

## 4.2 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 encounter, as opposed to translating 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 autocompleted-items from our ontology to be used alone or in combination with free text entry from the provider. [17] This hybrid approach will increase the generation of structured data while still enabling providers to enter free text information that reflects their clinical judgement. These free text additions should be captured and submitted as candidates for future inclusion in our ontology

## 4.3 Future Directions

The next steps are to deploy this ontology to additional institutions. We will then analyze usage for any terms that did not match, and consider adding them to the ontology. We will continue this iterative development cycle until we reach saturation of concepts.

We also plan to work with the relevant national disease registries, quality registries, EHR vendors, and standards organizations to adopt this as a national standard.

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.

# 5. CONCLUSION

We present the Hierarchical Presenting Problem ontology (HaPPy). This ontology was

empirically derived then iteratively validated by an expert consensus panel. HaPPy contains 690 presenting problem concepts, each concept being mapped to SNOMED CT. This freely sharable ontology should help to facilitate presenting problem based quality metrics, research, and patient care.

## Acknowledgements

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

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.

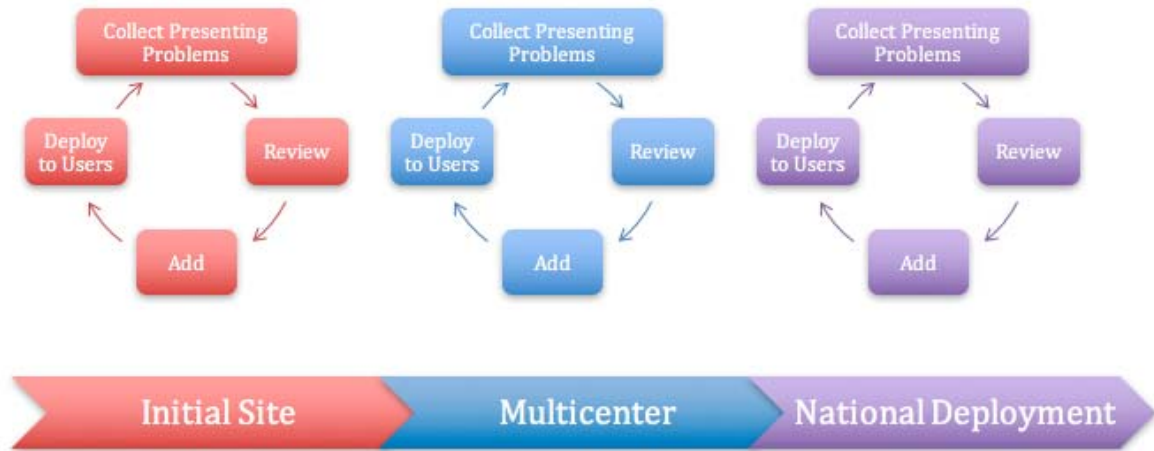
## Bibliography

- 1 Macfarlane D. The lexeme hypotheses: Their use to generate highly grammatical and completely computerized medical records. *Med Hypotheses* 2016;**92**:75–9. doi:10.1016/j.mehy.2016.04.031
- 2 Bronnert J, Masarie C, Naeymi-Rad F, *et al.* Problem-centered care delivery: how interface terminology makes standardized health information possible. *J AHIMA* 2012;**83**:30–5; quiz 36.
- 3 Data Elements for Emergency Department Systems, Release 1.0 (DEEDS): a summary report. *Acad Emerg Med* 1998;**5**:185–93. doi:10.1111/j.1553-2712.1998.tb02607.x
- 4 Travers DA, Haas SW. Unified medical language system coverage of emergency-medicine chief complaints. *Acad Emerg Med* 2006;**13**:1319–23. doi:10.1197/j.aem.2006.06.054
- 5 Haas SW, Travers D, Tintinalli JE, *et al.* Toward vocabulary control for chief complaint. *Acad Emerg Med* 2008;**15**:476–82. doi:10.1111/j.1553-2712.2008.00104.x
- 6 Intelligent Medical Objects, Inc.: Private Company Information - Bloomberg. Intelligent Medical Objects, Inc.: Private Company Information - Bloomberg. <http://www.bloomberg.com/research/stocks/private/snapshot.asp?privcapId=7984795> archived at <https://perma.cc/ZN4E-BGMW> (accessed 31 Mar2017).
- 7 Thompson DA, Eitel D, Fernandes CMB, *et al.* Coded Chief Complaints--automated analysis of free-text complaints. *Acad Emerg Med* 2006;**13**:774–82. doi:10.1197/j.aem.2006.02.013
- 8 Aronsky D, Kendall D, Merkley K, *et al.* A comprehensive set of coded chief complaints for the emergency department. *Acad Emerg Med* 2001;**8**:980–9.

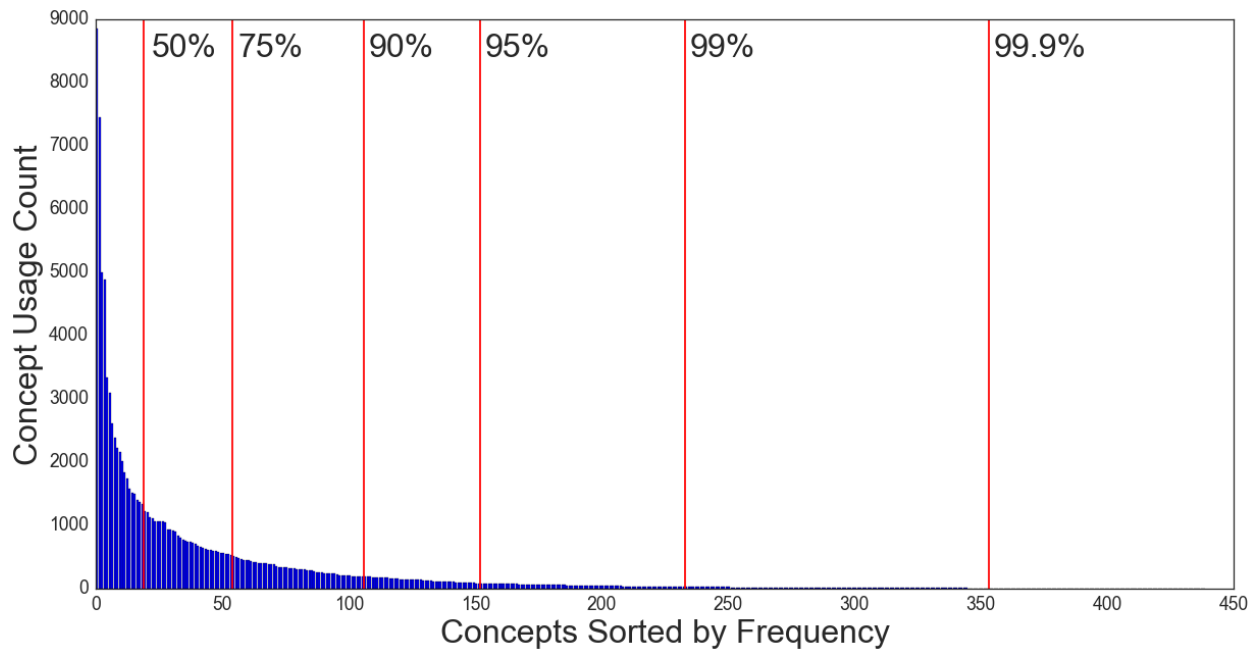
- 9 Handly N, Thompson DA, Li J, *et al.* Evaluation of a hospital admission prediction model adding coded chief complaint data using neural network methodology. *Eur J Emerg Med* 2015;**22**:87–91. doi:10.1097/MEJ.000000000000126
- 10 Barbara BM, Lynn S., M.D. Bickley, Robert A., M.D. Hoekelman. *A Guide to Physical Examination and History Taking/a Guide to Clinical Thinking*. 6th ed. Lippincott Williams 1995.
- 11 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**:387–95. doi:10.1016/j.annemergmed.2014.07.453
- 12 The Gephi Consortium. Gephi - The Open Graph Viz Platform. Gephi - The Open Graph Viz Platform. <http://gephi.org/> archived at <https://perma.cc/2FKZ-CSVL> (accessed 7 Apr2017).
- 13 Tindall SC. Level of Consciousness. In: Walker HK, Hall WD, Hurst JW, eds. *Clinical methods: the history, physical, and laboratory examinations*. Boston: : Butterworths 1990.
- 14 Zhou L, Plasek JM, Mahoney LM, *et al.* Mapping Partners Master Drug Dictionary to RxNorm using an NLP-based approach. *J Biomed Inform* 2012;**45**:626–33. doi:10.1016/j.jbi.2011.11.006
- 15 Rosenbloom ST, Miller RA, Johnson KB, *et al.* Interface terminologies: facilitating direct entry of clinical data into electronic health record systems. *J Am Med Inform Assoc* 2006;**13**:277–88. doi:10.1197/jamia.M1957
- 16 Goss FR, Zhou L, Plasek JM, *et al.* Evaluating standard terminologies for encoding allergy information. *J Am Med Inform Assoc* 2013;**20**:969–79. doi:10.1136/amiajnl-2012-000816
- 17 Jernite Y, Halpern Y, Horng S, *et al.* Predicting Chief Complaints at Triage Time in the Emergency Department. 2013. [http://cs.nyu.edu/~dsontag/papers/JerniteEtAl\\_nips13health.pdf](http://cs.nyu.edu/~dsontag/papers/JerniteEtAl_nips13health.pdf) archived at <https://perma.cc/3GR5-ETP4> (accessed 31 Mar2017).



**Figure 1: Iterative Presenting Problem Development**



**Figure 2: Distribution of Presenting Problem Concept Usage Frequency**



\*Vertical lines denote the percentage of patient encounters covered by all concepts left of the line.