

A Query Taxonomy Describes Performance of Patient-Level Retrieval from Electronic Health Record Data

Steven R. Chamberlin^{1*}, Steven D. Bedrick^{1,2}, Aaron M. Cohen¹, Yanshan Wang³, Andrew Wen³, Sijia Liu³, Hongfang Liu³, William R. Hersh¹

¹Department of Medical Informatics & Clinical Epidemiology, Oregon Health & Science University, Portland, OR

²Center for Spoken Language Understanding, Oregon Health & Science University, Portland, OR

³Division of Digital Health Sciences, Department of Health Sciences Research, Mayo Clinic, Rochester, MN

***Correspondence:**
Steven R. Chamberlin
chambest@ohsu.edu

Abstract

Performance of systems used for patient cohort identification with electronic health record (EHR) data is not well-characterized. The objective of this research was to evaluate factors that might affect information retrieval (IR) methods and to investigate the interplay between commonly used IR approaches and the characteristics of the cohort definition structure.

We used an IR test collection containing 56 patient cohort definitions, 100,000 patient records originating from an academic medical institution EHR data warehouse, and automated word-base query tasks, varying four parameters. Performance was measured using B-Pref. We then designed 59 taxonomy characteristics to classify the structure of the 56 topics. In addition, six topic complexity measures were derived from these characteristics for further evaluation using a beta regression simulation.

We did not find a strong association between the 59 taxonomy characteristics and patient retrieval performance, but we did find strong performance associations with the six topic complexity measures created from these characteristics, and interactions between these measures and the automated query parameter settings.

Some of the characteristics derived from a query taxonomy could lead to improved selection of approaches based on the structure of the topic of interest. Insights gained here will help guide future work to develop new methods for patient-level cohort discovery with EHR data.

Key Words

Information retrieval, patient cohort discovery, electronic health record, topic taxonomy

1. Background

The intent of this research is to define and test a query taxonomy, applied to patient cohort definitions, which can

explain the performance variations seen when retrieving these cohorts from electronic health record (EHR) data using automated methods. Also of interest is the possible relationship between a query taxonomy and different methods of retrieval and associated parameter settings. Patient cohort discovery in health records is an important task that is often used in academic institutions for research purposes, such as recruiting for clinical trials [1]. This can be a very labor-intensive task requiring time spent to design custom queries for each cohort definition, or topic. Automated methods could improve the efficiency of this task, but information retrieval methods have not been well-studied in this domain [2-5].

There has been promising research in medical record retrieval methods, some using publicly available EHR test collections [6, 7]. The Text Retrieval Conference (TREC) Medical Records Track, in 2011 and 2012, used one of these public EHR sources for methods development, but retrieval was only at the encounter level [8, 9]. Other patient-level cohort identification research has used only structured data [10], or focused on cohorts more broadly defined than that seen for research recruitment [11]. Methods have also been developed to locate clinical study inclusion criteria in EHR data, but not patient-level cohorts [12]. Methods using natural language processing, deep learning, and structured interfaces have been developed to optimize queries by the addition of context categories to EHR data, automate structuring of free text and classification, and to automatically convert criteria into structured queries [13-16]. Some methods focus on task definitions that differ from cohort identification, such as phenotyping [17-19]. For our purposes, cohort definitions not only contain disease diagnoses, but other complex features, such as lab tests, surgical procedures, medications, lab values, temporal relationships as well as combinations of structured and unstructured data.

Our previous research studied the performance of automated word-based queries used for complex patient-level cohort discovery with raw EHR data [5, 20], testing different parameter variations (n=48) for these queries

against 56 complex patient cohort definitions, or topics. Performance was generally poor for these queries, with 86% of the topics having a median B-Pref [21] under 0.25 (scale 0-1) across the 48 query parameter variations. These queries also underperformed when compared to custom designed Boolean queries. There were also large performance variations between and within the 56 topics. The range of median B-Pref across topics was 0 at a minimum and .895 at a maximum, and within topics ranges were seen as small as .03 points and as large as .60 points. And finally, there were also differences in median B-Pref between the 48 query parameter settings, although these differences were not as dramatic as that seen for the topics.

The variation in performance seen across complex patient cohorts in our previous research has led us to this research to explain, and predict, this variation by decomposing a patient-level cohort definition into a standard taxonomy. To do this we propose to use query performance data from this previous research to test our taxonomy definitions. Research on query decomposition with the intent to predict performance has been done in other domains [22], but to our knowledge, this type of taxonomy has not been previously defined or tested for this type of cohort discovery task using EHR data.

2. Materials and Methods

2.1 Test Data

To test the taxonomy developed for this project, we used the query performance data generated in a previous medical IR study [5]. We applied the Cranfield IR evaluation methodology [23] using the `trec_eval` program [24] (Fig 1). We used the B-Pref statistic as the performance measure for our evaluation. This statistic measures how many relevant patients were retrieved before non-relevant patients in ranked lists, and is used when relevance judging is incomplete, which is the case with the data used in this research. Due to the large volume of patients returned from the queries, random samples had to be selected for judging.

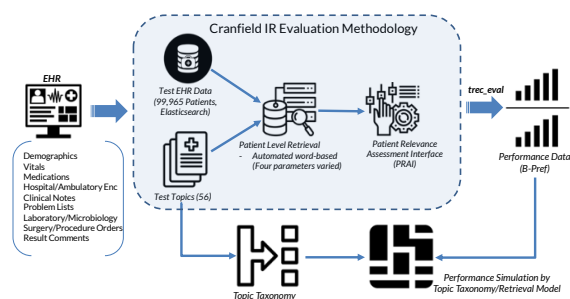


Figure 1. Overview of the generation of performance data used to test and model the query taxonomy definition.

Patient data originated from an Epic (Verona, WI) EHR system. A total of 99,965 unique patients with 6,273,137 associated encounters were stored in the Elasticsearch (v1.7.6) IR platform to evaluate the retrieval methods. There were a variety of document types associated with

each patient: demographics, vitals, medications (administered, current, ordered), hospital and ambulatory encounters with associated attributes and diagnoses, clinical notes, problem lists, laboratory and microbiology results, surgery and procedure orders, and result comments. These patients had to have at least three primary care encounters between 2009 and 2013.

The 56 topics were derived from actual patient cohort requests seen at two major medical research institutions, Oregon Health & Science University and The Mayo Clinic. A detailed example of a topic, in three representations, can be seen in **Table 1**. Examples of other topic summary descriptions include ‘Adults with IBD who haven’t had GI surgery’, ‘Adults with a Vitamin D lab result’, ‘Postherpetic neuralgia treated with topical and systemic medication’, ‘Children seen in ED with oral pain’, and ‘ACE inhibitor-induced cough’.

Table 1. Example of topic representations (A,B,C), Topic 15

- | |
|--|
| <p>A. Adults 18-64 years old with rheumatoid arthritis who have had a lab test for cyclic citrullinated peptide IgG antibody with a result greater than 40 units.</p> <p>B. Adults 18-64 years old with rheumatoid arthritis and lab result for positive anti-CCP IgG > 40 units.</p> <p> I. A 58-year-old female presents with morning stiffness and joint pain in her hands, especially her fingers, which improves after about 30 minutes, but doesn’t remit fully. On examination she is found to have ulnar deviation, decreased grip strength, and joint tenderness over the MCP and PIP joints. She has a positive rheumatoid factor and is positive for anti-CCP Ab at 45 units.</p> <p>C. Adults 18-64 years old with rheumatoid arthritis and lab result for positive anti-CCP IgG > 40 units.</p> <p> I. Demographics inclusion</p> <p> a. Age: 18-64 years</p> <p> II. Diagnosis inclusion</p> <p> a. Rheumatoid arthritis (ICD-9): 714.0</p> <p> III. Lab inclusion</p> <p> a. Cyclic citrullinated peptide IgG antibody (anti-CCP IgG): > 40 units</p> |
|--|

Four parameters for the word-based queries were varied to create 48 total iterations (queries) for each of the 56 topics. The four parameters are as follows:

1. Topic Representation (See **Table 1**) – A (summary statement), B (clinical case), or C (detailed criteria)
2. Text Subset – only clinical notes or all document types (including structured data reporting as text)

3. Aggregation Method – patient relevance score calculated by summation (sum) of all documents or by maximum (max) value
4. Retrieval Model – BM25, also known as Okapi [25], Divergence from randomness (DFR) [26], Language modeling with Dirichlet smoothing (LMDir) [27], Default Lucene scoring, based on the term frequency-inverse document frequency (TF*IDF) model [28]

Stratified random samples of 45 patients were selected from the top ranked 1000 patients retrieved from each run parameter iteration for each topic. The samples from all 48 iterations were combined for each topic and duplicates removed. The final judgment pools ranged in size from 450 to 780 patients for the 56 topics. Manual relevance judgment by clinically trained reviewers was performed on these pools. After the relevance assessment, final performance statistics were generated with the trec_eval program.

The final dataset contained the B-Pref performance statistic for all combinations of topics and query run parameters (56 topics x 48 run iterations = 2,688 unique queries).

2.2 Topic Taxonomy Characteristics

Our first step, to explain and predict the performance of the word-based queries, was to create a topic taxonomy composed of 59 features (**Table 2**). Three of the authors, who were trained clinically, iteratively developed a list of features that covered cohort inclusion or exclusion criteria of medical diagnoses and classifications, medications, procedures, lab tests, clinician information, patient demographics, information about the clinical setting, temporal measures and other aspects. Each of the 56 topics were then classified by these 59 features by the same three individuals. Fleiss Kappa was used to test interrater reliability [29].

Table 2. Taxonomy characteristics

Medical condition explicitly present?	Excluded medical diagnoses? ICD or Text
Systems: Cardiovascular	Included medical diagnosis? ICD or Text
Systems: Gastrointestinal	Included medications
Systems: Neuro/Psych	Excluded medications
Systems: Dental	Labs?
Systems: Accident	Labs? (if yes) Values also?
Systems: Ophthalmic	Imaging?
Systems: Metabolic	Imaging? (if yes) Interpretation also?
Systems: Infectious	Physical exam?
Systems: Genitourinary	Physical exam? (if yes) Values also?
Systems: Pulmonary	Included procedures/surgeries?
Systems: Hematology	Excluded procedures/surgeries?
Systems: Reproductive	Location? (none)
Systems: Oral	Location? (inpatient)
Systems: Skin	Location? (outpatient)
Systems: Cancer	Location? (emergency)
Systems: Breast	Location? (ICU)
Systems: Immunology	Provider Type
System: Endocrine	Visit Type
System: Musculoskeletal	Visit Specialty
Age: none	Status? (deceased)
Age: adult	Status? (living)
Age: child	Temporal: sequence of events
Age: senior	Temporal: age at first diagnosis
Sex: none	Temporal: time with diagnosis
Sex: female	Temporal: time on medication
Sex: male	Temporal: date of procedure
Free text descriptions required, simple term?	Temporal: age at encounter
Free text descriptions required, complex concept?	Temporal: date of procedure/lab
Number of encounters	

We wanted to examine any possible association between the query performance, as measured by B-Pref, and the 59 taxonomy characteristic classifications of the 56 topics. To do this we performed an exploratory data analysis by first creating a heatmap of run parameter settings by topics, with B-Pref as the performance metric. Using this heatmap, we clustered the 56 topics by query performance. Next, using this performance-based topic cluster order, we created a second heatmap of taxonomy characteristic assignment by topics, using the level of interrater agreement (0-3) as the performance metric. These heatmaps were compared for pattern similarities between performance clustering and taxonomy assignments to see if the B-Pref clustering patterns for topics were also seen with topic clusters associated with taxonomy characteristics.

2.3 Topic Taxonomy Structural Binary Features

To simplify the taxonomy definitions, focus more strictly on structural complexity, rather than content, representation, and to create features for model development and statistical testing, we defined the following six binary features by grouping some of the 59 taxonomy characteristics into categories (**Table 3**). We hypothesize that these six features capture the subset of taxonomy characteristics, and topic structure, and would more strongly correlated with performance. One investigator (SC) identified the taxonomy characteristics used to define these features as relevant to the topic, based on our experience designing and executing manual Boolean queries associated with each topic. These assignments were reviewed by the other investigators.

The first binary feature was positive if there was a temporal component in the topic ('Temporal', y/n). The 56 topics contain a variety of temporal conditions, including age at first diagnosis, time with diagnosis, chronological order of disease onset for several diagnoses, and medication use before or after first diagnosis. The second binary feature was positive if the topic could not be defined exclusively with the structured data present in the data set (ICD, CPT, disease and drug names) and required some free text ('Text', y/n). An example for this would be a topic that checked for the presence of a side effect, only included in clinical notes in the data set, associated with a medication. The third binary feature was positive if the topic required a medication list check, either exclusions or inclusions or both ('Medication', y/n). The fourth binary feature was positive if there was a procedure in the topic. This includes any surgical or non-surgical procedure ('Procedure', y/n). The fifth binary feature was positive if additional value criteria were required from lab tests, imaging, or physical exams beyond just having these tests in the record. ('Additional', y/n). And finally, the sixth binary feature was positive if the topic required a specific disease diagnosis or diagnoses. Some topics were defined for cohorts who only received certain screening tests without an explicit disease requirement ('Condition', y/n). Using the example for Topic 15 in **Table 1**, there is a medical condition explicitly required (rheumatoid arthritis, Condition=Y), an included lab test (anti-CCP, Procedure=Y), and an additional value criteria required for

the lab test (IgG>40 units, Additional=Y). The other three binary features would be 'N' for this topic

Table 3. Taxonomy binary features.

Temporal	The topic included a temporal component
Text	The structured query version of the topic also required unstructured data
Medication	The topic included or excluded a medication list.
Procedure	The topic included or excluded surgical procedures.
Additional	The topic included additional information about labs, imaging or physical exams.
Condition	The topic explicitly included information about a specific medical condition.

The relationship between these six taxonomy features and query performance was investigated, by testing for any performance related interactions between these features and the four word-based query parameters (topic representation, text subset, aggregation method, and retrieval model). These interactions capture the relationship between word-based query parameters and inherent topic structure related to complexity (binary taxonomy features).

We used a beta regression model for this investigation. This model was trained on the B-Pref performance data as the dependent variable, with the four word-based query parameters, the six binary taxonomy features and all first-order interactions between the parameters as the independent variables. Due to data limitations we felt that model coefficients and tests of significant might not be generalizable beyond this data set. We instead used this model to predict B-Pref on all possible permutations of values of the parameters and features, and to investigate the patterns of the predicted B-Pref in this predicted and simulated parameter/feature space. Since this simulated data contained all possible combinations of values of the four word-based parameters and the six binary taxonomy features, there were a total of 3,072 entries. Using the simulated data, we estimated the effect of the six binary taxonomy features individually, and the effect of the Topic Representations. We also used this simulated data for an exploratory data analysis, using a heatmap, to assess more complex interactions between the parameter space (interventions) and the binary feature space (inherent topic structure).

A beta regression mean model was selected because the response variable, B-Pref, is continuous, restricted to the unit interval [0,1], and asymmetrically distributed. The logit link function was used for these analyses. The regression was done with R (v3.3.1) using the package betareg (v3.1-2).

3. Results

3.1 Taxonomy Analysis – 59 Characteristics

We found moderate, substantial or almost perfect agreement by Fleiss kappa on 50 of the 56 topics, rated by the three clinically trained raters for the 59 query taxonomy characteristics (Fig 2). Topic distribution is in Fig 3.

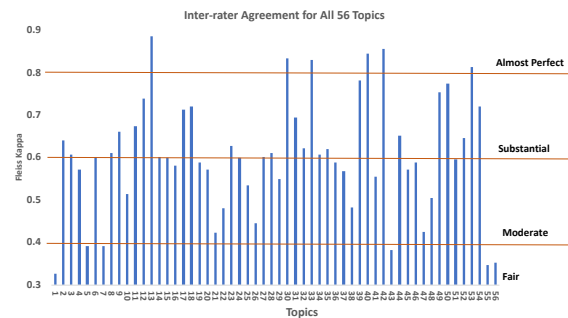


Figure 2. Interrater agreement for the 59 taxonomy characteristics applied to 56 topics. Fleiss Kappa was calculated for each topic based on agreement between three clinical trained raters on the 59 taxonomy characteristic assignments.




Figure 3. Topic counts by taxonomy characteristics.

We next used heatmaps to investigate the relationship between word-based query performance and assignment to the 59 taxonomy characteristics (Fig 4). Topic clusters, based on B-Pref performance (left heatmap), were maintained for taxonomy characteristics (right heatmap), but column clustering was allowed for this heatmap. Performance-based clustering of topics can be seen for the B-Pref heatmap, but there do not appear to be similar patterns found in the taxonomy heatmap, while maintaining the same topic order. There does not appear to be an association between performance and the 59 taxonomy characteristics.




Figure 4. Comparison B-Pref Performance and Taxonomy Characteristic Clustering. The heatmap on the left are the topics by word-based query parameter settings clustered for B-Pref performance. Darker colors represent higher B-Pref values. The heatmap on the right maintains topic order from B-Pref clustering on the left (rows), but the columns, which represent the 59 taxonomy characteristic assignments, are clustered. Darker colors on the right heatmap represent the strength of interrater agreement. We do not see patterns here similar to the performance patterns on the right

3.2 Taxonomy Analysis – 6 Binary Features

For our data, the beta regression model output did show that five of the six binary taxonomy features were associated with poorer performance, as measured by B-Pref. One feature, ‘text’, was associated with better performance. Features associated with poorer performance were designed to capture increased topic complexity in various ways, so this result is not surprising. The feature ‘text’ captures the ability of purely structured data to describe a medical topic, with or without added free text. Our result indicates that topics that require text, in addition to structured data, might perform better. And there were notable interactions between the taxonomy features and the run parameters, particularly between the feature ‘temporal’ and Topic Representation. Interestingly this analysis did not point to any notable interactions between the four word-based parameters. But it is not clear if these results are generalizable due to the specific nature of our 56 topic descriptions.

We then used the beta regression model, containing the four word-based parameters, six binary taxonomy features and the interactions between the parameters and features, to predict B-Pref with a simulated dataset. This dataset contained all possible permutations of the ten predictors. We varied each of the six binary taxonomy flags independently, while holding all other values constant, to estimate the impact of these flags. We also did this for Topic Representation (**Fig 5**). We again saw that five of the six binary taxonomy features were associated with poorer performance, and the feature ‘text’ was associated with

improved performance. We also saw Topic Representation B associated with improved performance.




Figure 5. B-Pref prediction changes on simulated data for selected characteristics. Values represent change in predicted B-Pref using a simulated dataset with all permutations of the four word-based query parameters and the six binary taxonomy flags. All other values are held constant while varying the characteristic listed.

We also created a heatmap of the predicted B-Pref values generated from the simulated data (**Fig 6**). The x axis contained all possible permutations of the four word-based query parameters and the y axis contained all possible permutations of the six binary taxonomy features, and hierarchical clustering was done in both dimensions. Clear patterns of performance clustering can be seen, particularly around the combinations of three of the binary taxonomy features, temporal, text and condition. These three features are conceptually different from the other three (medication, procedure, additional) in that the latter are simple additions of information but the former represent more complex topic structural aspects. In addition, within specific combinations of these flags there are also clear variations in performance across different word-based parameter settings. In the bottom horizontal cluster, the best performance for topics without a temporal, text and condition component (blue rectangle) is seen with a completely different set of parameter settings than for topics with all three of these structural components (red rectangle). This performance pattern is an example of a possible interaction between the parameters and features, which could help guide the selection of parameters to optimize retrieval results based on the taxonomic attributes of the topics.

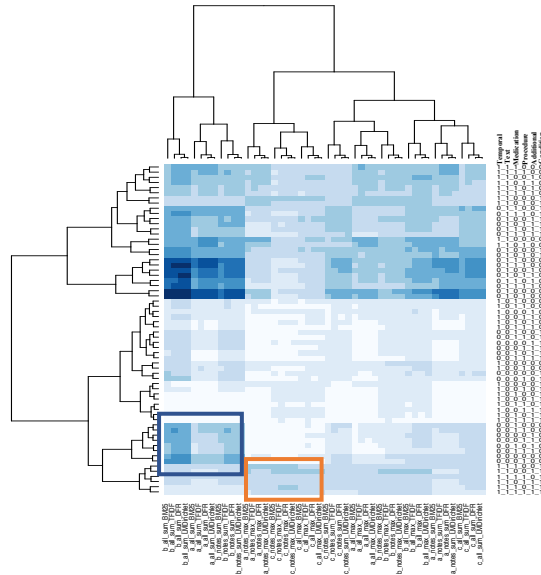


Figure 6. B-Pref predictions on simulated data. Darker values represent higher predicted B-Pref. Clustering was done on both axes. The x axis represents all permutations of the four word-based query parameters and the y axis represents all permutations of the six binary taxonomy features.

4. Discussion

The findings in our previous research, and the performance variation within and across topics, led us to pursue two further methods to understand and improve our results. In an attempt to understand our results, we developed a taxonomy for the topics that we hoped would identify characteristics associated with the differences in results. We first developed an exhaustive 59 parameter taxonomy that did not reveal any associations. However, when we reduced the taxonomy to six binary variables, we did find association with performance. As also shown by comparable work at Mayo Clinic [30], it may be possible with further prospective analysis that query taxonomy might lead to selection of different query approaches based on characteristics of the topic.

This work provides some evidence that applying a query taxonomy might improve performance. Further work with methods such as machine learning might yield improvements, although it is not clear what features will lead to performance improvement across varying topical criteria for different queries.

There were a number of limitations to this work. Our records were limited to a single academic medical center. There are many additional retrieval methods we could have assessed, but we would not have the resources to carry out the additional relevance judgments required as those additional methods would add new patients to be judged. It is also difficult to generalize our results due to the specificity of the topics. This will always be a limitation for this type of work since it would be extremely difficult to represent all possible cohort requests that could be seen for all forms of medical research. Finally, there is a global

limitation to work with EHR data for these sorts of use cases in that raw, identifiable patient data is not easily sharable such that other researchers could compare their systems and algorithms with ours using our data [31].

ACKNOWLEDGMENTS

This work was supported by NIH Grant 1R01LM011934.

References

1. Obeid, J., et al., *A survey of practices for the use of electronic health records to support research recruitment*. Journal of Clinical and Translational Science, 2017. **1**: p. 246-252.
2. Ni, Y., et al., *Increasing the efficiency of trial-patient matching: automated clinical trial eligibility pre-screening for pediatric oncology patients*. BMC Medical Informatics & Decision Making, 2015. **15**: p. 28.
3. Ni, Y., et al., *Automated clinical trial eligibility prescreening: increasing the efficiency of patient identification for clinical trials in the emergency department*. Journal of the American Medical Informatics Association, 2015. **22**: p. 166-178.
4. Ni, Y., et al., *A real-time automated patient screening system for clinical trials eligibility in an emergency department: design and evaluation*. JMIR Medical Informatics, 2019. **7**(3): p. e14185.
5. Chamberlin SR, B.S., Cohen AM, Wang Y, Wen A, Liu S, Liu H, Hersh WR *Electronic Health Record Data for a Cohort Discovery Task*. medRxiv, 2019.
6. Chapman, W., et al. *Creation of a repository of automatically de-identified clinical reports: processes, people, and permission*. in *Proceedings of the American Medical Informatics Association Clinical Research Informatics*. 2011. San Francisco, CA.
7. Johnson, A.E., et al., *MIMIC-III, a freely accessible critical care database*. Sci Data, 2016. **3**: p. 160035.
8. Voorhees, E.a.R.T. *Overview of the TREC 2011 Medical Records Track*. . in *The Twentieth Text REtrieval Conference Proceedings (TREC 2011)*. 2011. Gaithersburg, MD: National Institute of Standards and Technology.
9. Voorhees, E.a.W.H. *Overview of the TREC 2012 Medical Records Track*. in *The Twenty-First Text REtrieval Conference Proceedings (TREC 2012)*. 2012. Gaithersburg, MD: National Institute of Standards and Technology.
10. Glicksberg, B.S., et al., *Automated disease cohort selection using word embeddings from Electronic Health Records*. Pac Symp Biocomput, 2018. **23**: p. 145-156.
11. Sarmiento, R.F. and F. Dernoncourt, *Improving Patient Cohort Identification Using Natural Language Processing*, in *Secondary Analysis of Electronic Health Records*. 2016, Springer Copyright 2016, The Author(s). Cham (CH). p. 405-417.

12. Stubbs, A., et al., *Cohort selection for clinical trials: n2c2 2018 shared task track 1*. J Am Med Inform Assoc, 2019. **26**(11): p. 1163-1171.
13. Ateya, M.B., B.C. Delaney, and S.M. Speedie, *The value of structured data elements from electronic health records for identifying subjects for primary care clinical trials*. BMC Med Inform Decis Mak, 2016. **16**: p. 1.
14. Kang, T., et al., *EliIE: An open-source information extraction system for clinical trial eligibility criteria*. J Am Med Inform Assoc, 2017. **24**(6): p. 1062-1071.
15. Zhang, K. and D. Demner-Fushman, *Automated classification of eligibility criteria in clinical trials to facilitate patient-trial matching for specific patient populations*. J Am Med Inform Assoc, 2017. **24**(4): p. 781-787.
16. Yuan, C., et al., *Criteria2Query: a natural language interface to clinical databases for cohort definition*. J Am Med Inform Assoc, 2019. **26**(4): p. 294-305.
17. Denny, J., L. Bastarache, and D. Roden, *Phenome-wide association studies as a tool to advance precision medicine*. Annual Review of Genomics and Human Genetics, 2016. **17**: p. 353-373.
18. Richesson, R., et al., *Clinical phenotyping in selected national networks: demonstrating the need for high-throughput, portable, and computational methods*. Artificial Intelligence in Medicine, 2016. **71**: p. 57-61.
19. Robinson, J., et al., *Defining phenotypes from clinical data to drive genomic research*. Annual Review of Biomedical Data Science, 2018. **1**: p. 69-92.
20. Wu, S., et al., *Intra-institutional EHR collections for patient-level information retrieval*. Journal of the American Society for Information Science & Technology, 2017. **68**: p. 2636-2648.
21. Buckley, C. and E. Voorhees. *Retrieval evaluation with incomplete information*. in *Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2004. Sheffield, England: ACM Press.
22. al., F.e., *From Evaluating to Forecasting Performance: How to Turn Information Retrieval, Natural Language Processing, Recommender Systems into Predictive Sciences* Dagstuhl Manifestos, 2018. **7**(1): p. 96-139.
23. Cleverdon, C. and E. Keen, *Factors determining the performance of indexing systems (Vol. 1: Design, Vol. 2: Results)*. 1966, Aslib Cranfield Research Project: Cranfield, England.
24. Harman, D., *Information Retrieval Evaluation*. 2011, San Rafael, CA: Morgan & Claypool.
25. Robertson, S. and S. Walker. *Some simple effective approximations to the 2-Poisson model for probabilistic weighted retrieval*. in *Proceedings of the 17th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1994. Dublin, Ireland: Springer-Verlag.
26. Amati, G. and C. VanRijsbergen, *Probabilistic models of information retrieval based on measuring the divergence from randomness*. ACM Transactions on Information Systems, 2002. **20**: p. 357-389.
27. Zhai, C. and J. Lafferty, *A study of smoothing methods for language models applied to information retrieval*. ACM Transactions on Information Systems, 2004. **22**: p. 179-214.
28. Salton, G. and C. Buckley, *Term-weighting approaches in automatic text retrieval*. Information Processing and Management, 1988. **24**: p. 513-523.
29. Fleiss, J., B. Levin, and M. Paik, *The Measurement of Interrater Agreement, in Statistical Methods for Rates and Proportions, Third Edition*. 2003, John Wiley & Sons: Hoboken, NJ. p. 598-626.
30. Wang, Y., et al., *Test collections for electronic health record-based clinical information retrieval*. JAMIA Open, 2019: p. Epub ahead pf print.
31. Eickhoff, C.a.K., Yubin and White, Ryen. *Overview of the Health Search and Data Mining (HSDM 2020) Workshop*. 2020. ACM.