International Journal of Knowledge Discovery in Bioinformatics
Volume 8 • Issue 2 • July-December 2018

Augmenting Medical Decision Making With Text-Based Search of Teaching File Repositories and Medical Ontologies:

Text-Based Search of Radiology Teaching Files

Priya Deshpande, DePaul University, Chicago, USA
Alexander Rasin, DePaul University, Chicago, USA
Eli T Brown, DePaul University, Chicago, USA
Jacob Furst, DePaul University, Chicago, USA
Steven M. Montner, Department of Radiology, The University of Chicago, Chicago, USA
Samuel G. Armato III, Department of Radiology, University of Chicago, Chicago, USA
Daniela S Raicu, DePaul University, Chicago, USA

ABSTRACT

Teaching files are widely used by radiologists in the diagnostic process and for student education. Most hospitals maintain an active collection of teaching files for internal purposes, but many teaching files are also publicly available online, some linked to secondary sources. However, public sources offer very limited (and ad-hoc) search capabilities. Based on the previous work on data integration and text-based search, the authors extended their Integrated Radiology Image Search (IRIS 1.1) engine with a new medical ontology, SNOMED CT, and the ICD10 dictionary. IRIS 1.1 integrates public data sources and applies query expansion with exact and partial matches to find relevant teaching files. Using a set of 28 representative queries from multiple sources, the search engine finds more relevant teaching cases versus other publicly available search engines.

KEYWORDS

Data Integration, Image-Based Search, Integrated Radiological Image Search, Medical Ontologies, Natural Language Processing, Radiological Teaching Files

DOI: 10.4018/IJKDB.2018070102

1. INTRODUCTION

A radiology teaching file repository is a collection of important cases for teaching and clinical follow-up, and references to better understand the classification of diseases (Dashevsky et al., 2015). All teaching files share a similar general structure but significant variations exist, even within the same data source. Teaching files may include information such as patient history, findings, diagnosis, differential diagnosis, and images related to clinical reports. Teaching files can be categorized into three types: (1) personal teaching files that are meant for the general use of the teaching file owner, (2) shared in-house teaching files in which the owner makes the teaching file content available for viewing within their institution, and (3) public teaching files built on a shared model but with more comprehensive content that may undergo a formal review before publication (De-Arteaga et al., 2015).

A recent national survey assessing the role and desired features of radiology teaching files found that, among the 396 respondents from 115 institutions, 89% use some form of teaching file from which 76% keep a personal teaching file containing a variety of media and 67% use a shared in-house teaching file, while 83 institutions had paid subscriptions to a public teaching file repository (Dashevsky et al., 2015). Public teaching file solutions have become increasingly popular, providing users with instant access to thousands of cases (although of inconsistent data) (Seitz et al., 2003), sometimes for a fee. While all of these public and commercial solutions are available, most do not permit users to (1) easily submit personal cases to their libraries, (2) perform efficient querying, categorization and search for particular cases, (3) simulate basic PACS (picture archiving and communication system) functionality, or (4) enable self-directed and assessed learning – all important teaching file repository features as identified by at least 50% of the survey respondents (Dashevsky et al., 2015).

Therefore, as the first step to organize and extract medical knowledge from large teaching file repositories, we have 1) developed a database schema for teaching file integration and a framework for a radiology image search engine and 2) evaluated the framework on the Radiology Society of North America Medical Imaging Resource Community (RSNA MIRC) (2018) and MyPacs (Group, 2018) repositories indexed using the Radiology Lexicon (RadLex) (RSNA, 2018). We normalized all data sources and augmented the integration process with data cleaning and validation to account for different format representations. Many data sources include noisy entries – for example, different teaching files, even though stored in the same data repository, do not use the same text category names. For the teaching files that did not come indexed by RadLex (as is done for MIRC), we annotated all imported data with RadLex terms.

In this paper, we propose an extension of the data repository indexing by integrating Unified Medical Language System (UMLS) Systematized Nomenclature of Medicine – Clinical Terms (SNOMED CT) (2017) using UMLS Metamorphosis (SNOMED, 2017) and show that this extension improves search results, particularly for queries that originally retrieved few teaching files. To evaluate and quantify search results, we propose a new evaluation criterion in consultation with medical experts that measures the accuracy of search results based on the term appearance in different categories of teaching file text. Based on domain knowledge and surveys, we found that findings and diagnosis are the most relevant search categories in teaching files. Our original search engine is referred to as Integrated Radiology Image Search (IRIS) (Deshpande, 2017) IRIS 1.0, and the proposed improvement as IRIS 1.1.

The rest of the paper is organized as follows. In Section 2 we present related relevant data sources and other research that motivated our work. In Section 3 we present our database schema, architecture of IRIS 1.1, teaching file repositories we integrated, RadLex and SNOMED CT ontologies, Natural Language Processing (NLP) techniques we applied to perform the smart search, and IRIS 1.1 evaluation metric. In Section 4 we present a comparative analysis of our proposed work (IRIS 1.1) to the previous system (IRIS 1.0), the MIRC and MyPacs search engines as well as customized Google searches. In Section 5 we summarize the proposed approach and outline new avenues for expanding this work.

2. RELATED WORK

In this section we present a literature review of papers that discuss the need for data integration of radiological sources. We also summarize existing radiological sources along with their search engine capabilities, advantages, and limitations. While our current implementation integrates MIRC and MyPacs data sources, we intend to integrate the data sources described here into our system.

Several studies have highlighted the need to integrate clinical reports and images into an integrated database with advanced search capabilities. Gutmark et al. (2007) argued for building a system that reduces errors in radiological images' interpretation using teaching file databases. Easy-to-use computer-based teaching files are useful for training physicians and as a reference tool for experienced physicians with the long-term goal of improving diagnostic accuracy. Talanow et al. (2009) discussed how radiological images are critical for diagnosis, why teaching file are helpful in radiological diagnosis, and how they developed internet-based radiology teaching file systems. Dos-Santos et al. (2012) discussed how the availability of a large and diverse set of clinical cases drives the need for the integration of profiles published by Integrating the Healthcare Enterprise (IHE). Margolies et al. (2016) found that a repository of pathology-proven cases in a dashboard has the potential to enhance and encourage the formation of accurate teaching files, as well as educational publications in the form of case series or "case of the day" submissions. Hwang et al. (2016) discussed how the use of Positron Emission Tomography Computed Tomography (PET-CT) increased the need to retrieve relevant medical images that can assist in diagnosis. Hwang et al. designed a database using annotation and image markup as well as an image search engine. This paper further motivated our research work – to design a logical schema for a database with stored image features and to develop a search engine that supports both text and image-based search. Furthermore, Kansagra et al. (2016) presented the idea of having a global database that integrates multiple data sources (such as clinical data, patient history, physical exam findings, laboratory data, or imaging data) for more precise and accurate diagnosis.

In the rest of this section, we present a list of currently or previously available data repositories and medical search engines along with their advantages and limitations (summarized in Table 1). We marked some entries with "*," indicating that these search engines provide links to teaching file repositories from external data sources. For example, Google provides links to MyPacs or MIRC teaching files but does not show results in terms of teaching files.

Radiology Society of North America (RSNA) Medical Imaging Resource Community (MIRC) (2018) is a large repository (over 2,000 public teaching files with more than 12,000 images) with teaching files including patient history, diagnosis, differential diagnosis, findings, discussion, and external references to journal articles. Radiological terms are highlighted and linked to RadLex terms as described in section 3.1. MIRC links to the RadioGraphics and Radiology journals, which can provide additional data for medical knowledge extraction. Although a rich source of medical knowledge, the MIRC repository search engine does not support specialized search fields (such as anatomy, age, and imaging modality), does not recognize negation, and does not have the ability to perform query expansion through synonyms.

MyPacs data repository (with 17,000 publicly available teaching files) has over 33,000 cases in total. Using the search engine of MyPACS, users can search records based on anatomy, pathology, image modality, age, gender, etc. The data repository also links to MedScape (LLC, 2017) for additional supplemental information; however, the built-in search engine does not support searching through synonyms and negation, and, similar to the MIRC search engine, is not able to perform image-based search. Although beyond the scope of this paper, our database schema and engine design include facilities to support image-based search.

CTisus (Hospital, 2017) is a large repository of radiological images, quizzes, and CT protocols. Although there are 237,000 images available along with video files, the repository does not contain case diagnosis, patient medical history or differential diagnosis, and there is no support for image-based searches. Medscape (LLC, 2017) is the latest medical news and information source about drugs and diseases available for radiology students and physicians. Even though Medscape is rich in medical data, there is no search engine available nor do any teaching files contain images, differential diagnosis or other valuable case information.

Radiopaedia (Jones, 2017) is an open-edit radiology resource with 25,000 cases and 10,000 articles. As with all the other data sources, no image-based search is provided. Gamuts (Reeder, 2017) contains a comprehensive list of image differential diagnoses that are linked to symptoms, disease names, and causes. Although images are linked to the GoldMiner Radiology search engine, Gamuts does not offer a text-based search as it has no search engine. The Casimage database within the IRMA framework (Thies, et al., 2004) integrates multimedia teaching and reference data into the PACS environment. The database includes only 8,000 images and it does not feature concept-based image retrieval. ImageCLEFmed Teaching Files (Muller et al., 2010) is a collection of domain-specific photographs for the medical field, which was used in medical ad hoc retrieval tasks from 2004 to 2007. This medical archive comprises a total of 66,000 images and several composite medical sub-collections provided by independent medical institutions and hospitals.

Radiology Teacher (Talanow, 2009) is a web-based teaching file development and distribution program, which allows authors to create, edit, and delete cases and images with descriptions and annotations. A quiz mechanism and an image annotation feature integrate an interface to the Medical Illustrator software are also provided. Presently, the Radiology Teacher system contains less than 350 cases.

RADTF (Do et al., 2010) is a teaching file solution, which is compatible with RadLex. Differential diagnosis and quiz modes are available. RADTF uses RadLex anatomy concept terms and provides NLP features to process radiologic reports (Do et al., 2010), including stemming, ranking of results based on detected negation, hedge, and uncertainty expressions. Although RADTF has been described as open-source (Do et al., 2010), others have concluded that RADTF is not publically available (De-Artega et al., 2015) and therefore, we cannot evaluate the features it supports.

EURORAD (European Society of Radiology) (Neutorgasse, 2017) is a peer-reviewed educational tool based on teaching cases. It contains 6,000 teaching files with clinical history, image findings, discussion, final diagnosis and differential diagnosis. EURORAD currently supports three different languages (English, Spanish, and French). Users can search on anatomical body structure but, similar to other teaching file sources, there is no support for negations, synonyms, or image-based search. Another radiological teaching file system that can be integrated into a PACS environment is RadPix (Weadock, 2017). Complete radiological teaching files can be created by adding text, annotations and images. The current public user interface only has 11 teaching cases.

In addition to the work described so far, there were several efforts made to allow image search to use the text associated with the images through captions or text embedded in the images. Some of these systems proposed for the medical domain are summarized below.

The Biomedical Image Metadata Manager (BIMM) (Korenblum et al., 2011) system provides retrieval of similar images using semantic features of image metadata. Based on the imaging observation, 2D regions of interest are stored as metadata, and the system offers content-based image retrieval capabilities although this could not be tested as it is no longer available in the public domain (Korenblum et al., 2011).

Khresmoi for Everyone (Khresmoi, 2017) is a medical informatics and retrieval system that provides an access system for online biomedical information and documents. The result of a search is a web link to a discussion forum about diseases and quizzes. However, there is no unified solution with clinical reports with categories, such as patient history or differential diagnosis. Furthermore, the search capabilities are limited; for example, it does not support synonym- and negation-based search.

GoldMiner (Society, 2017) helps users search images and articles from peer-reviewed biomedical journals. It uses the National Library of Medicine to discover medical concepts in figure captions with the final goal of retrieving relevant images. GoldMiner recognizes abbreviations, synonyms, and types of diseases but search results depend on the explicit presence of specific words in figure's captions.

Yottalook (Solutions, 2017) is a radiologist-targeted search engine powered by Google Custom Search that searches a variety of sources such as radiopaedia.org, American Journal of Radiology, University of Michigan Medical School, and MyPacs. It provides users with the ability to choose the category of search (e.g., CT, ultrasound). Although Yottalook searches multiple sources it does not integrate them: users searching a category (e.g., X-ray) are then redirected to the original source in its specific format (e.g., an external webpage or a Power-Point file).

Open-i (NIH, 2017) (Open Access Biomedical Image Search Engine) is a service of the National Library of Medicine that enables search and retrieval of abstracts and images (including charts, graphs, clinical images, etc.) from the open source literature and biomedical image collections. Searching may be done using text queries as well as images. Open-i provides access to over 3.7 million images from about 1.2 million PubMed Central® articles; 7,000 chest x-rays with 3,000 radiology reports; 6,000 images from NLM History of Medicine collection; and 2,000 orthopedic illustrations. However, these articles are not teaching files and therefore do not include important categories such as patient history, diagnosis or findings.

3. MATERIALS AND METHODS

3.1. Creation of Logical Schema, Data Integration, and Indexing

We designed a generalized database (used to combine heterogeneous data sources) logical schema (Figure 1) and populated the database from public teaching file repositories. The base entry in the center of the schema is a combination of teaching file record and an image entry since teaching files are naturally built around visual examples (e.g., MRI, digital radiograph images). Each entry is then annotated with a variety of related information, both from the data source (e.g., differential diagnosis, patient history and discussion fields of the teaching file) and derived data (e.g., image properties, indices or image feature extracts). Each teaching file entry is further linked with references and diagnosis information as well as patient data, physician data and a pathology report (when available).

To implement the database schema, we considered many relational databases including MySQL, SQL Server, Oracle and PostgreSQL. Based on the radiology teaching files' types of data, we determined that the PostgreSQL database is better suited for heterogeneous database connectivity, user defined functions (UDFs), support for low-level image features, and data integration. After cleaning and loading data into a PostgreSQL database, based on the logical schema we developed, we enforced Health Insurance Portability and Accountability Act (HIPAA) constraints by censoring patients' protected health information. We searched for matching patterns such as date of birth, bank account number, or social security number to identify patient personal information and blanked it out in teaching files within our database. In some cases, the matching pattern approach has the limitation of not differentiating between personal information (e.g., patient date of birth) and non-personal information (e.g., date of publication of a teaching case); we currently eliminate all dates even if they are not part of patient personal records. In our next implementation, we plan to use the HIPAA Security Rule Toolkit by NIST (NIST HIPAA, 2017). The NIST HIPAA Security Rule Toolkit Application is intended to help organizations better understand the requirements of the HIPAA security rules, implement those requirements, and assess those implementations in their operational environment. Furthermore, we plan to detect potential HIPPA compliance violations not only in text, but also in images associated with the teaching files by modeling private data stamped or written within the image.

We initially integrated RSNA MIRC data source into IRIS 1.0. After the MIRC dataset integration, we added the content from MyPacs.net data repository. When querying large bodies of unstructured text (e.g., teaching file patient history and discussion categories), indexing is necessary to optimize query response time and to support custom analysis. Text indexing is particularly effective when applied to frequently used words in data repositories and guided by medical ontologies. An ontology is a way of representing the terms and their relationships in a domain; it

Table 1. A comparative study of available data sources and search engines NLP capabilities

Search Engine	Keyword Search	Synonyms	RBT	SEC	RB	Publicly Available	TF
RadTF	✓	✓	✓	X	X	X	X
GoldMiner	✓	X	X	✓	*	✓	X
Yottalook	✓	✓	X	✓	*	✓	X
Google	✓	X	X	✓	✓	✓	X
MIRC	✓	X	X	X	X	✓	✓
MyPacs	✓	X	X	X	X	✓	✓
Gamuts	X	X	X	X	X	✓	X
CTisus	✓	X	X	✓	X	✓	X
Casimage	X	X	X	X	X	X	X
RadICS	X	X	X	X	X	X	X
BIMM	✓	X	X	X	X	X	X
Radiology Teacher	✓	X	X	X	X	✓	X
Medscape	✓	X	X	✓	X	✓	X
ImageCLEFmed	X	X	X	X	X	X	X
Khresmoi	✓	✓	X	X	X	✓	X
Openi	✓	✓	X	✓	✓	✓	X
EURORAD	✓	X	X	X	X	✓	✓

Abbreviations: Relationships between terms (RBT): Hierarchical relation between terms – is a/has a, relevance feedback (RB): Based on results and feedback from user, retrieve relevant results, Spelling Error Correction (SEC): Prompt user with correct spelling, Teaching Files (TF) – data source in terms of categories: history, findings, diagnosis, and discussion.

Figure 1. Logical schema used in IRIS 1.1

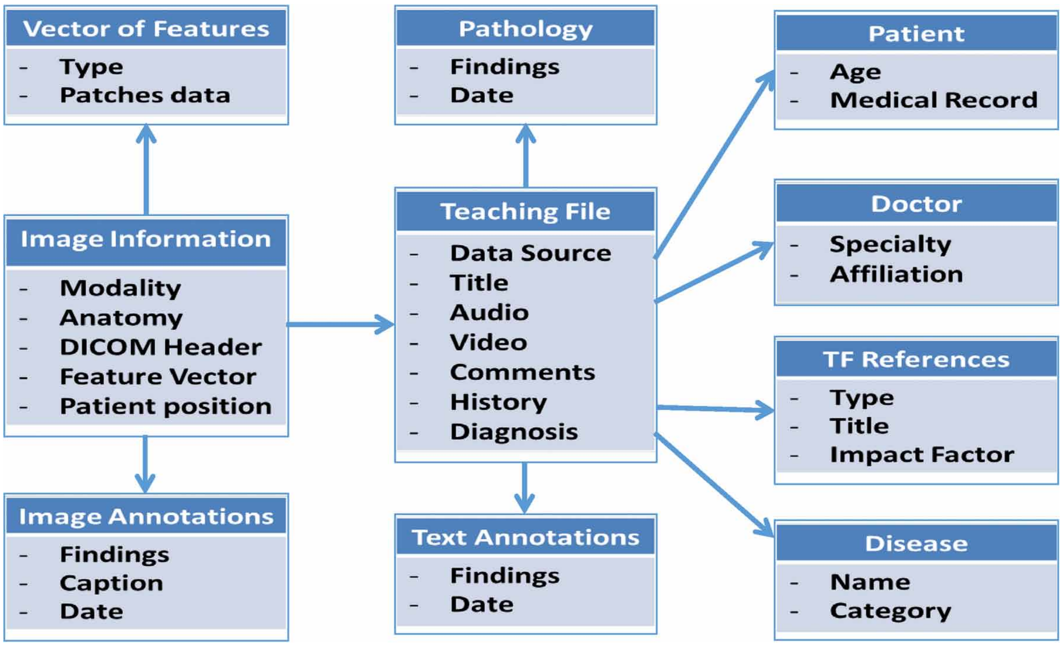

gives the description of term concepts and the relations that connect them. Our current indexing implementation integrates two popular ontologies, RadLex and SNOMED CT.

RadLex (RSNA, 2017) is an ontological system that provides a comprehensive lexicon vocabulary for radiologists. A radiology-specific lexicon was needed to make more efficient use of the growing amount of electronic information in the radiology environment, in particular in the creation of electronic teaching materials, and to more accurately search reports and perform data-mining. The RadLex browser was developed by the RSNA and links to articles from journals including the British Institute of Radiology and American Journal of Neuroradiology.

The Systematized Nomenclature of Medicine Clinical Terms (SNOMED CT) (2017) ontology provides a standardized, multilingual vocabulary of clinical terminology that is used by physicians and other healthcare providers for the electronic exchange of clinical health information. The SNOMED CT ontology follows the National Library of Medicine (NLM) Unified Medical Language System (UMLS) format (SNOMED, 2017); it has a hierarchical structure and includes clinical findings, anatomy, test findings, and morphological connections. This ontology covers more than 311,000 terms with preferred name, synonyms, definition, and semantic meaning.

We used UMLS (2017) Metathesaurus to build vocabulary dictionaries using ontologies such as SNOMED CT and International Classification of Diseases (ICD). The UMLS includes the Metathesaurus, the Semantic Network, and the specialist Lexicon and Lexical Tools. The Metathesaurus is a large biomedical thesaurus that is organized by concept, or meaning, and links similar names for the same concept from nearly 200 different vocabularies. The Metathesaurus also identifies useful relationships between concepts and defines the meanings (using the definition of terms), concept names, and relationships from each vocabulary.

The ICD is a widely recognized international system for recording diagnoses with standardized codes. It is developed, monitored, and copyrighted by the World Health Organization (WHO). We have expanded our medical dictionary by building a ICD10 (10^{th} revision) (2017) dictionary. We have used UMLS SNOMED CT, which provides us with concept ids that are mapped to the ICD10 dictionary; this coding standard can be used to analyze the text by disease name, for example, “Neoplasm” corresponds to the ICD10 code C00-D49. ICD10 codes can enrich search vocabulary because they group together diseases, procedures and entity relationship information into categories.

Figure 2 summarizes the IRIS 1.1 architecture: the two data sources (MIRC and MyPACS), the RadLex lexicon and SNOMED CT ontology (which integrated ICD10 standards through UMLS), the NLP modules, and the display of the query results including text categories and images associated with the teaching file. The IRIS 1.1 evaluation module stores user’s feedback in an accuracy measure table that includes user feedback ranking of the search results.

To keep IRIS 1.1 content up-to-date, we intend to update the database bi-monthly by looking at the “Date of Modification” attribute of each teaching file; if the data repository is more recent, IRIS 1.1 will update the corresponding stored teaching file and all of its associated indexes. Similarly, IRIS 1.1 content will be updated based on any changes in the SNOMED CT and RadLex ontologies. IRIS 1.1 expanded vocabulary includes a large number of defined terms (300,000), further improving our search results from previous work. For example, Angiosarcoma is a term which belongs to one type of disease; using our previous dictionary (in IRIS 1.0), we were not able to find related results. Using SNOMED CT, an IRIS 1.1 search for “Angiosarcoma” was considered along with synonyms such as “malignant Hemangioendothelioma,” “hemangio-sarcoma,” “hemangio-endothelial sarcoma,” and “haemangiosarcoma”; as a result, IRIS 1.1 found 30 teaching cases related to “angiosarcoma” which were missing from the IRIS 1.0 search. This IRIS 1.1 improvement is particularly significant in minority search cases (i.e., very few matches) which are most sensitive to missing search results.

3.2. Smart Search Through Synonyms and Negation Interpretation

3.2.1. Synonym Expansion

Rather than limiting searches to an exact match between the query and the data in the integrated database, our search engine performs automatic query expansion, augmenting the search with term synonyms found using medical ontologies. For example, if a user searches for “paranephric,” IRIS 1.1 will expand the query, augmenting the search with synonymous terms such as “adrenal gland,” “glandula suprarenalis,” “nebenniere”.

Figure 2. IRIS 1.1 architecture

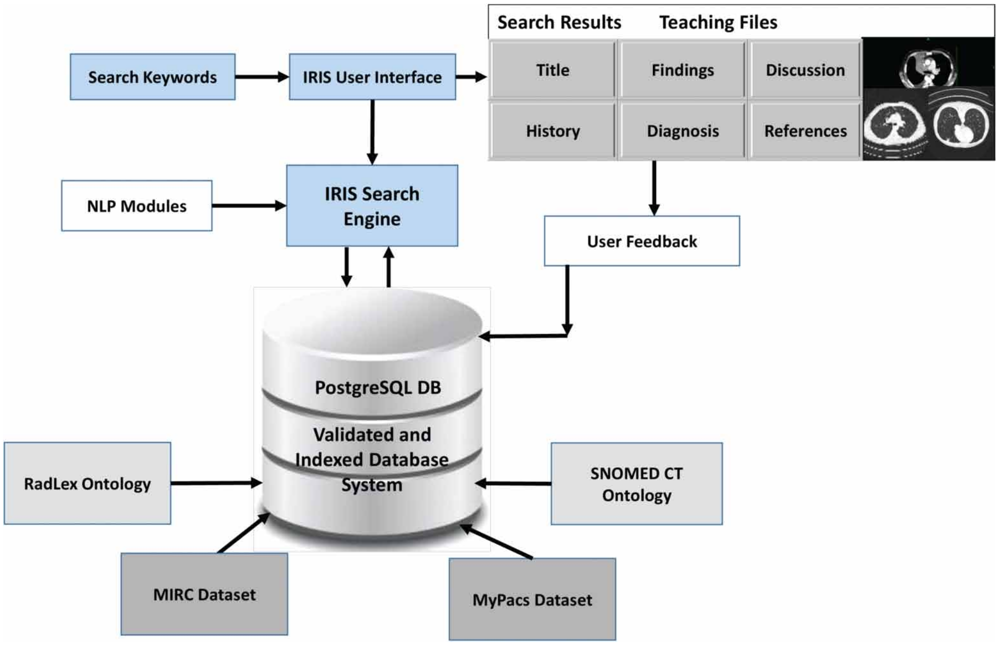

3.2.2. Negation

We currently handle negation by expanding the search from exact matches to queries that contain synonyms for that negation; for example, a typical search for "no X" is expanded with "missing X," "lacking X," "absent X," and "without X." We also substitute negation expressions with antonyms: for example, query such as "no abnormal renin secretion" will be substituted with "normal renin secretion".

We surveyed Clinical Text Analysis Knowledge Extraction Terms (CTAKES) (Apache, 2017), a well-known Natural Language Processing (NLP) tool. CTAKES can be used for extraction of information from electronic medical record clinical free-text by identifying the clinical entities, such as diagnosis, procedure, diseases, and anatomical structures. Each entity has mapping code context and related information about that entity. CTAKES with negEx library (Apache, 2017) can identify negation in search query terms. However, this tool does not assist with replacing negation with antonyms and therefore we implemented our own algorithm to interpret negation in search terms.

Currently, we are using the nltk library to pre-process search queries, apply stemming, and remove stop words. We removed negation stop words from the standard stop word list (e.g., with, without, not, between, below) to keep them in our search, because these search terms are medically relevant. IRIS 1.1 identifies negation terms, finds concepts associated with each term, and introduces antonyms for negated terms, implemented through our own algorithm and an integrated antonym dictionary. For example, for "no cardiomegaly" IRIS 1.1 will search for cases with "normal heart," "no enlarged heart," and "no cardiomegaly". Our results show that by adding antonyms based on ontology' definitions to recognize negation and handling negation queries, the recall for the search is improved. For example, Table 5 shows that for "no cardiomegaly" query, there are seven relevant retrieved teaching files without the use of antonyms. When using the antonym for "cardiomegaly" ("normal structure of heart" based on the SNOMED CT ontology definitions), there are eleven relevant teaching files retrieved. Since negation, just as any other NL search query construct, is inherently ambiguous (Korenblum et al., 2011) – the interpretation of the negation, in particular of the antonym, can be ambiguous as well. To reduce this ambiguity, as part of our future work we will implement a context-aware search feature to select an appropriate meaning based on context.

Figure 3 shows the flow of IRIS 1.1 with query expansion: once user inputs a query, IRIS 1.1 will apply basic stemming and pass the query to the next module. Next, a negation module will identify if negation is present; if there is negation IRIS 1.1 will search for antonyms in the medical dictionary, otherwise IRIS 1.1 will search for term synonyms. Next, the synonym/negation module will look for synonym/antonym options for the search query, using integrated medical ontologies (RadLex and SNOMED CT). If synonyms or antonyms are present, IRIS 1.1 performs query expansion and forwards the results to the search module. Based on both original and expanded query terms, IRIS 1.1 retrieves and displays teaching files images along with the associated text. In Figure 3, we illustrated the synonym substitution process using "Angiosarcoma" query term.

Figure 3. IRIS 1.1 Search Flow: IRIS 1.1 identifies that there is no negation term in query search term "Angiosarcoma" and returns teaching cases with "Angiosarcoma" along with teaching files that contain term synonyms such as "Malignant Hemangioendothelioma" and "Hemangio-sarcoma"

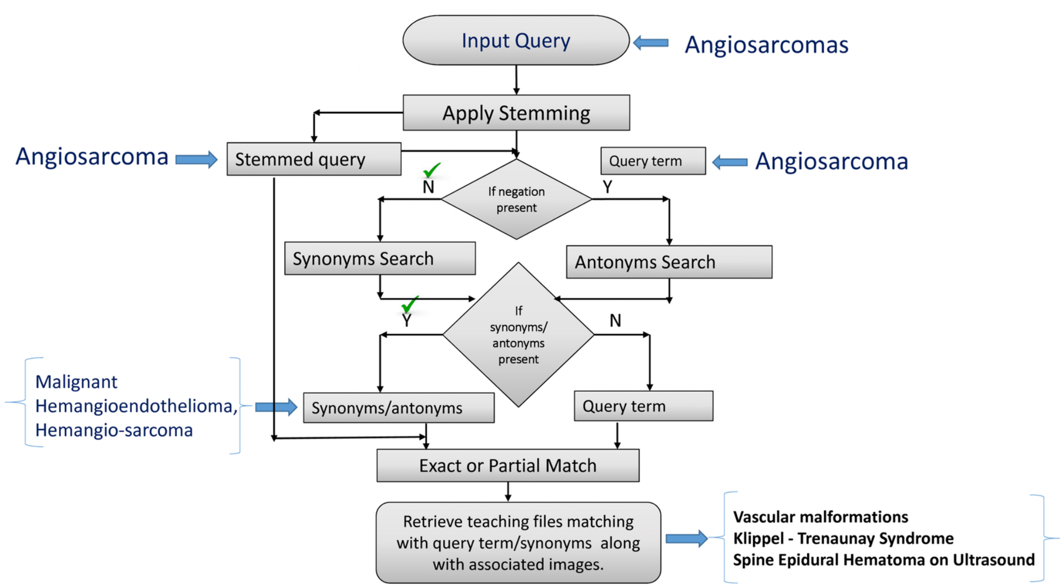

3.3. Search Methodology

We extended our previous study by improving search results through query expansion and by considering partial matches and search term frequency in target text. In our evaluation, we used queries received from our radiologist collaborators at a major medical university and from a literature survey. When medically relevant phrases do not generate exact match results, our search engine splits the search text and looks for individual words from the query. Partial text matching is likely to produce less relevant results, but is nevertheless greatly preferable to having no results. IRIS 1.1 ranks and displays final search results based on the frequency of query term occurrence in found teaching files. We use the date of modification in teaching files to break ties in relevance ranking; the most recent teaching file will get a higher rank.

3.4. User Study for IRIS 1.1 Evaluation

In order to further validate the IRIS 1.1 search, we designed a user study to create reference truth for quantifying search result accuracy. Four human annotators provided feedback on search results by annotating the retrieved teaching files as irrelevant, relevant, more relevant, most relevant, or best result (as summarized in Table 2).

Original IRIS 1.0 results were tested based on Relevant, Not relevant, or Not sure ratings, which were sufficient as IRIS 1.0 results were not ranked based on search relevance. We selected a scale from 0 to 4 to rate the relevance of the results based on previous work on evaluating search engines and gathering reference truth for diagnostics interpretation in the medical domain. In particular, a study by Kushniruk et al. (2002) presented a comparative evaluation of an experimental automated text summarization system and three conventional search engines – Google, Yahoo and About.com, and showed that, even when a rating system based on ratings from 1 to 10 was used, only a few ratings were used. Furthermore, in another study by Armato et al. (2011), when asked to interpret the level of malignancy for lung nodules, the ratings used by the NIH/NCI Lung Image Database Consortium (LIDC) were on a scale from 1 to 5 (1 = "most likely benign," 2 = "somehow benign," 3 = "indeterminate," 4 = "somehow malignant," and 5 = "most likely malignant"), showing the need for including different levels of uncertainty in the ratings. Since there is an inherent uncertainty when evaluating the relevance and ranking retrieval results, our proposed evaluation metric incorporates this uncertainty by choosing the scale of 0= "irrelevant" to 4= "best result." In IRIS 1.1, we extend our search module and evaluate the search relevance with the help of four computing experts with experience in medical imaging retrieval (but without medical training). According to surveyed medical domain experts, when a query term appears in findings or diagnosis category, a teaching file is more relevant to the search than the teaching file with a query match in the discussion category. Experts also provided exemplar results for sample queries, which were used as the baseline for the evaluation of our results. IRIS 1.1 search result evaluation scores are summarized in Table 7. Evaluators were asked to mark the returned teaching files as irrelevant, relevant, more relevant, most relevant, or best result based on a query term, synonym or the definition of term, and which (term/synonyms) appear in any category text. If the query term was not in any category of the teaching file, evaluators were asked to check for synonyms of the term – including both the synonyms and the definition of terms based on the integrated medical ontologies. For evaluation we used a subset of 28 queries, as explained in Section 4.4. We focused on 5 queries (ACL Tear, Bronchus intermedius, Mega cistern magna, Angiosarcoma, and No Angiosarcoma) for a more detailed evaluation. We chose queries that return few results because evaluating the relevance of dozens or more results is harder (for the evaluators) to quantify. All chosen queries resulted in even fewer teaching files in the original IRIS 1.0 search, further emphasizing the benefit of integrating an additional ontology. We provided detailed documentation for the evaluators including relevant synonyms/antonyms, parent term information, term concept, and definitions of query terms drawn from ontologies.

4. RESULTS

We compared the built-in search engines that are associated with MIRC and MyPacs, Google search engine, and IRIS 1.1 search engine based on six queries formulated by medical experts at a major medical university and twenty-two sample queries from related work by De-Arteaga et al. (2015). As Google does not offer integrated search for teaching file repositories, we performed custom searches (e.g., Google search for “renal site: www.mypacs.net” to find content related to “renal” in MyPacs.net teaching file repository). This search was chosen because a regular Google search returns a wide variety of irrelevant (for our purposes) results such as PowerPoint presentations, videos, and PDF documents mixed in with the teaching files. Table 3 summarizes the significant differences between IRIS 1.0 and IRIS 1.1. As discussed above, the integration of SNOMED CT ontology improved the search results.

4.1. Comparison Results Between IRIS 1.1 With and Without SNOMED CT Integration on MIRC

Our current search engine results are shown in Table 4. The results show that out of 28 search queries, IRIS 1.1 search engine improved results for 11 queries using query expansion with exact match; 5 more queries were improved by using partial query match. For example, for exact match if we consider the example of “tracheal dilation,” IRIS 1.1 will search for “tracheal dilation” and also search for “Bronchoscopy with tracheal dilation,” returning 69 and 159 additional teaching files for MIRC (2k) and MyPacs (17k) datasets respectively.

IRIS 1.0 search retrieved 59 cases while IRIS 1.1 retrieved 63 cases, as it also searched for “enlargement of heart” and “enlarged heart” which are synonymous with “cardiomegaly”. Using the SNOMED CT ontology, IRIS 1.1 engine also searched for “cardiac dilation” and “congenital cardiomegaly,” “cor bovinum,” and “megalocardia”. The exact search count depends on the dataset, but our enhanced dictionary enabled IRIS 1.1 to identify new teaching files that were closely related to the search terms. For example, the “cardiomegaly” search matched a teaching file with the title “Ebstein Anomaly” –where “cardiomegaly” does not appear in any of the categories. In this teaching file only synonyms of “cardiomeglay” are used, such as “cardiac enlargement” which appears once in findings, "enlarged heart" in differential diagnosis, and "cardiac enlargement" in the discussion category. This example shows how IRIS 1.1 searches were improved through the integration of medical ontologies and synonym expansion.

4.2. Comparison of IRIS 1.1 With Other Search Engines

Our in-depth evaluation of IRIS 1.1 engine compares query results, using 6 queries with the ones produced by the MIRC and MyPacs search engines as well as those produced by the Google search.

We chose 6 queries (related to diagnosis) from different sources including (De-Artega et al., 2015) (queries are summarized in Table 5) and added a negation term to those queries to measure how our search engine and other engines handle negation-based search. The goal in choosing these 6 queries was to design a workload that is easy to evaluate by non-medical experts. For example, without the use of negation in the search, results should be related to abnormal structure of heart, enlargement of heart, and cardiac enlargement. With negation-based search, e.g. "no cardiomegaly," results should be related to normal structure of heart, no symptoms of heart failure, and no enlargement of heart.

We used "cardiomegaly" and "no cardiomegaly" from Table 4 for a detailed illustration of IRIS 1.1 in comparison with other search engines (shown in Figure 4). Query search results were significantly improved through applying synonyms and negation interpretation when compared with the other search engines. For example,

Table 2. IRIS 1.1 evaluation metric

Relevance Score	Score Term	Definition
0	Irrelevant	Not relevant
1	Relevant	The term/synonyms appear in any category
2	More relevant	The term/synonyms appear in discussion category
3	Most relevant	The term/synonyms appear in differential diagnosis or history category
4	Best result	The term/synonyms appear in title, findings, or diagnosis category

Table 3. Comparison between IRIS 1.0 and IRIS 1.1

Categories	IRIS 1.0	IRIS 1.1
Ontology	RadLex ontology	RadLex and SNOMED CT ontologies
Type of search	An exact match for query keywords/synonyms	Both exact matches and partial matches for query keywords/synonyms
Relevance rank	No	Relevance rank based on keywords/synonyms appearance in teaching file text categories
# of queries	10	22
Evaluation metric scale	1-3 (relevant/not relevant/ not sure)	0-4 (irrelevant/ relevant/more relevant/most relevant/ best)

Table 4. A comparative study of IRIS 1.1 search with previous IRIS 1.0 search on MIRC (2k) and MyPacs (17k) dataset with results improved (highlighted in bold) through the use of (4A) enhanced synonym dictionary and (4B) partial match for MIRC dataset queries and search term "Double duct sign" MyPacs (17k); (4C) shows no improvement in results

A. Query (Enhanced Synonym Dictionary)	IRIS 1.1 MIRC (2k)	IRIS 1.0 MIRC (2k)	IRIS 1.1 MyPacs (17k)	IRIS 1.0 MyPacs (17k)
Tracheal dilation	**200**	131	**786**	627
Cardiomegaly	**63**	59	**106**	99
Bronchus intermedius	**3**	1	**2**	2
Chiari	**38**	19	**154**	134
Angiosarcoma	**30**	1	**96**	26
Cystitis cystica	**3**	0	**2**	0
Cystitis	**10**	7	**96**	95
Cystitis glandularis	**5**	2	**2**	0
Innominate vein	39	39	**95**	68
Innominate artery	238	238	**866**	855
Varicocele	**4**	2	28	24
B. Query (Partial Match)	**IRIS 1.1 MIRC (2k)**	**IRIS 1.0 MIRC (2k)**	**IRIS 1.1 MyPacs (17k)**	**IRIS 1.0 MyPacs (17k)**
Baastrup disease	**868**	0	1	1
Limbus vertebra	**243**	0	5	5
Thornwaldt cyst	**577**	0	6	6
Splenic hemangioma	**89**	0	2	2
Double duct sign	**1003**	0	**4355**	0
C. Query (No Improvement)	**IRIS 1.1 MIRC (2k)**	**IRIS 1.0 MIRC (2k)**	**IRIS 1.1 MyPacs (17k)**	**IRIS 1.0 MyPacs (17k)**
Irregularly shaped	11	11	20	20
Acl tear	9	9	145	145
Study	117	117	776	776
Appendicitis	40	40	176	176
ACL graft tear	7	7	85	85
Hepatic adenoma	74	74	360	360
Annular pancreas	14	14	36	36
Perthe	20	20	63	63
Mega cisterna magna	3	3	9	9
Vertebra	243	243	753	753
Toxic	48	48	165	165
Buford complex	43	43	178	178

Table 5. Examples of queries based on our own queries and De-Arteaga (2011) (improvements in number of results highlighted in bold, DS: an IRIS integrated dataset)

Query	IRIS 1.1 (MIRC DS 2k)	MIRC (2k)	IRIS 1.1 (MyPacs DS 17 k)	MyPacs (33k)	Google Site: MIRC	Google Site: MyPacs
Cardiomegaly	**63**	56	**106**	110	21	64
No Cardiomegaly	**11**	56	**72**	110	7	7
ACL Tear	**9**	3	**71**	96	11	102
No ACL Tear	0	3	14	96	0	95
Appendicitis	40	42	179	162	10	99
No Appendicitis	3	42	1	162	4	100
Hepatic adenoma	**74**	8	**15**	20	0	14
No Hepatic adenoma	2	7	8	20	0	20
Annular pancreas	14	16	35	28	4	40
No Annular Pancreas	5	16	11	28	1	30
Toxic	48	53	166	99	9	90
No Toxic	12	53	2	99	4	60

while the built-in search engines for MIRC and MyPacs retrieved the same results for “cardiomegaly” and “no cardiomegaly” (effectively ignoring negation) showing 56 teaching files and 110 teaching files respectively, our search engine differentiated between these two searches by recognizing negation and returning different answers in the negation-based query.

For the MyPacs data repository, IRIS 1.1 retrieved 106 cases that were a subset of 110 MyPacs results; note that the MyPacs search has access to all 33k teaching files versus about 17k teaching files that are freely accessible (without fees) that were included in IRIS 1.1 search (i.e., we integrated roughly half of the MyPacs teaching files). We also used Google site search as one of the alternatives which found 21 and 72 teaching files for MIRC and MyPacs, respectively.

Figure 4a summarizes the comparative performance of different searches with “cardiomegaly”. IRIS 1.1 has 50 results that overlap with MIRC results, 106 results that overlap with MyPacs results, and 81 overlapping results with the Google search for the same query. There is an overlap of 21 cases between MIRC and Google and 81 cases between MyPacs and Google search. In general, for a simple single-term search query, our results are similar to other tested search engines, with the notable advantage of being returned through a single integrated search. Our next search in Figure 4b summarizes evaluation of negation showing results for “no cardiomegaly”. Surprisingly, no search engine of those compared here applies negation to the search.

Figure 4. Comparison of negation for "cardiomegaly" (the circles represent the sets of retrieved results and their overlapping sections show the counts of teaching files in the intersection, and the star indicates improvement in number of results)

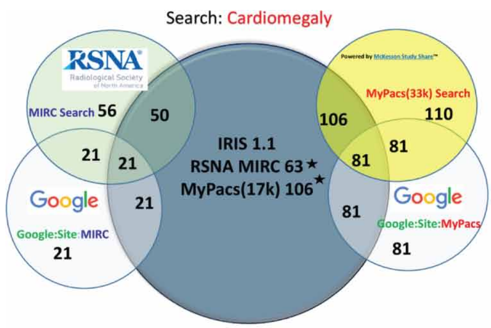

(4a) Search results for "cardiomegaly"

(4b) Search results for "No cardiomegaly"

MIRC and MyPacs return the exact same teaching files as with "cardiomegaly," meaning that the negation term was not considered in the query. Using the SNOMED CT ontology, the IRIS 1.1 search replaced "no cardiomegaly" with "normal heart" as "cardiomegaly" is defined as "morphologically abnormal structure of the heart". IRIS 1.1 also extended the search with "heart size normal" and retrieved 11 results from MIRC and 72 from MyPacs datasets. These results include a different set of cases and must not return "cardiomegaly" cases (which we manually verified). Only two cases from our negation-based search matched with the original "cardiomegaly" search; the overlap occurred because discussion referred to "usually normal heart" (i.e., accidental overlap that can be eliminated as a false-positive with additional analysis). Even if other search engines were to consider negation, searching for teaching files that do not include "cardiomegaly" is not the right strategy –an ontology is necessary to find correct results by considering antonyms. Google search showed 7 and 64 results for MIRC and MyPacs site search, respectively.

There is an overlap between the IRIS 1.1 and MIRC search with one teaching file and two teaching files with MyPacs. All these teaching files mention "normal heart" in a different context (in addition to "cardiomegaly"), thus we consider this a false positive. For example, in MIRC dataset we had one overlapping teaching file "D-Transposition of the great arteries"; on closer inspection we observed that "normal heart" and "cardiomegaly" were part of the discussion category in which radiologists wrote "the CXR may appear normal, with usually a normal heart size … the CXR may demonstrate mild cardiomegaly…"

We manually inspected all search engine results, including IRIS 1.1 results, and found that the teaching files retrieved by IRIS 1.1 were more relevant to the "no cardiomegaly" query as compared with the other search engines. IRIS 1.1 furthermore showed an improvement in number of teaching file matches over the previous IRIS 1.0 (IRIS 1.0 returned 54 cases, while IRIS 1.1 retuned 83 teaching cases) because of the integrated ontologies. We note that MyPacs returns more results than any other engine because it will always search for individual terms even if the search phrase is given in quotation marks.

From this comparative study we concluded that the MIRC, MyPacs, and even Google do not accurately distinguish between search terms with and without negation. Our search engine recognizes the presence of negation in search terms and retrieves teaching files that are different from matches without negation. Not all query results from IRIS 1.0 could be improved; for example, even though IRIS 1.1 augmented "No ACL Tear" with "No tear of ACL" and "No anterior cruciate ligament tear," even after applying stemming and considering antonyms, no new cases matched the search.

4.3. Synonym Coverage in RadLex and SNOMED CT Ontologies

To further validate query expansion and qualify the importance of considering synonyms in query search, we performed a study to determine synonym coverage in different ontologies (summarized in Table 6). We used a total of 28 queries from a major medical university and our own related literature survey. For these 28 queries, RadLex has synonyms for 9 queries (covering 32% with synonyms), while SNOMED CT has synonyms for 18 queries (covering 64% of our search terms). A union of the RadLex and SNOMED CT ontology results provides synonym coverage for 75% of our query dataset. SNOMED CT also has more than one synonym for most of the terms compared to RadLex ontology. RadLex has 3 overlapping terms (Brachiocephalic vein, Truncus brachiocephalicus, Malignant hemangioendothelioma for terms Innominate vein, Innominate artery, Malignant hemangioendothelioma, respectively) with SNOMED CT. Out of 18 synonyms from SNOMED CT only 3 synonyms overlap, while the other 4 terms have different synonyms (Bronchus intermedius, Appendicitis, Chiari, Cystitis). This demonstrates that both ontologies are important, and that a search engine is less effective when relying on just one integrated ontology.

RadLex encompasses approximately 50% of SNOMED CT synonyms terms for these queries, which suggests opportunities for expanding RadLex content. This lack of coverage can be overcome by combining RadLex with other medical lexicons such as SNOMED CT. Coverage with SNOMED CT ontology is 50% higher than that of RadLex for these 28 queries.

4.4. User Study for IRIS 1.1 Evaluation

We performed a user study evaluation of IRIS 1.1 results with 5 queries (ACL Tear, Bronchus intermedius, Mega cistern magna, Angiosarcoma, and No Angiosarcoma). We chose queries that return few results (as evaluating relevance of a large result set is less informative); same queries also resulted in very few teaching files in the original IRIS 1.0 search. We collected scores for all 5 queries from our 4 evaluators and averaged these results (Table 7).

Average relevance score of these 5 queries (considering the top 3 results from each query) was 2.6. For "Bronchus intermedius" query, none of the teaching files had "Bronchus intermedius" in findings or diagnosis category, limiting achievable search relevance results to at most 2. This term is an anatomical structure, so the

Table 6. Synonyms in RadLex and SNOMED CT ontologies

Term	Radlex Synonyms	SNOMED CT Synonyms
Tracheal dilation	-	Bronchoscopy with tracheal dilation
Cardiomegaly	-	Enlarged heart, cardiac enlargement
Bronchus intermedius	Interlobar bronchus	Truncus intermedius of right main bronchus
Innominate vein	Brachiocephalic vein, vena brachiocephalica	Injury of innominate vein, brachiocephalic vein
Innominate artery	Truncus brachiocephalicus, brachiocephalic artery	Brachiocephalic artery, brachiocephalic trunk, truncus brachiocephalicus
Acl tear	Tear of acl	-
Study	-	Study
Appendicitis	Appendizitis	Ecphyaditis
Hepatic adenoma	-	Liver cell adenoma, hepatocellular adenoma
Annular pancreas	-	Annular pancreas
Varicocele	-	Venous varices, pampinocele
Perthe	-	Pseudocoxalgia
Chiari	Hindbrain hernia	Congenital abnormality
Angiosarcoma	Malignant hemangioendothelioma, angiosarkom	Malignant hemangioendothelioma, hemangio-sarcoma, hemangio-endothelial sarcoma, haemangiosarcoma
Mega cisterna magna	-	Mega cisterna magna
Baastrup disease	Lumbar interspinous bursitis	-
Vertebra	-	Structure of vertebra
Cystitis cystica	-	Cystitis cystic
Cystitis	Zystitis	Bladder infection
Cystitis glandularis	-	Cystitis glandularis

Table 7. Evaluation results ($\bar{X} \pm \sigma$: $\bar{X}$ = average of ratings, σ = standard deviation between ratings)

Query	$\bar{X} \pm \sigma$	Maximum Variation Between Ratings
ACL Tear	4 ± 0.5	1
Bronchus intermedius	3 ± 0.5	1
Mega cistern magna	2 ± 0.7	2
Angiosarcoma	2 ± 0.5	1
No Angiosarcoma	2 ± 0.8	2

results retrieved by IRIS 1.1 were about the diseases related to this anatomical structure. The rest of the queries scored in the range of 2 to 4. From our analysis we concluded two important things. First, a teaching file database is a supplemental learning resource and not a comprehensive decision support system which can offer an exact diagnosis. Teaching file databases can serve as reference resources to augment the diagnostic interpretation process. Moreover, we observed that query terms often appeared in the discussion category (rather than diagnosis category) because this is where the radiologists provided their opinion about the diagnosis. For example, for "angiosarcoma" query retrieved results contained the "Pulmonary zygomycosis" diagnosis; however, the discussion stated "…a broad spectrum of disease processes have been associated with the halo CT sign including vasculitic entities, angiosarcomas…" Search query term matching diagnosis category is the most relevant teaching file, yet discussion or history of the patient matching the search is also useful for radiologists to learn more about similar cases. Our second important observation was that we still have a limited set of data integrated into IRIS 1.1. Our aim is to provide a larger integrated repository with many teaching cases, so for future IRIS 1.1 versions we are going to continue integrating other publicly available data sources such as EURORAD and GoldMiner.

5. CONCLUSION AND FUTURE WORK

Currently, radiologists at best have access to an ad-hoc internal search engine that helps them find the internal teaching files available at their hospital. The reference coverage of internal engines is limited to the in-house teaching files and lacks the desired advanced analytical search capabilities. To our knowledge, none of the available engines provide radiologists with the ability to integrate in-house data with other data sources, even if these sources are publicly available and rich in medical content.

In this paper, we described a database and search framework for heterogeneous data integration to facilitate medical knowledge extraction from publicly available teaching file repositories. IRIS 1.1 supports negation and query expansion through synonyms; the integration of SNOMED CT ontology further improved results by enhancing synonyms dictionary. We are currently improving our search by incorporating the context of search into the query and ranking. Furthermore, we are improving IRIS 1.1 results by (increasingly) weighting ontology terms and considering co-occurrence of terms that appear frequently with query terms which will also lead to context-aware search. For future implementation versions, we are working on improving results by considering hierarchical structure of terms (e.g., "cardiomegaly" is a case of "enlarged heart") and proximity of words in search augmentation, that will further expand "cardiomegaly" with "heart problem" and "cardiac failure".

Based on our evaluation, we found that the IRIS 1.1 improves results in two ways: 1) search queries produce more relevant results compared with the existing search tools, and 2) results from multiple data sources can be merged into a single easy-to-query and interpret data source. Furthermore, our search engine can be hosted on a cloud to improve processing performance in a distributed computing environment. Our current implementation can be accessed through a prototype website and allows integration with additional ontologies and data sources for in-house use by other institutions.

Although our current implementation uses text-based search for images and teaching files, the existing database schema makes it easy to incorporate image-based search using image features derived from pixel content. We also plan to add functionality to the IRIS 1.1 enabling radiologists to mark regions of interest and create image annotations. In the long run, our search engine will provide capabilities for both text and image search that will be tailored for a variety of users, including radiologists, radiology residents, clinicians, and patients.

REFERENCES

Apache. (2017, November 11). Ctakes. Retrieved from http://ctakes.apache.org/

Armato, S. G. III, McLennan, G., Bidaut, L., McNitt-Gray, M. F., Meyer, C. R., Reeves, A. P., & Clarke, L. P. et al. (2011). The lung image database consortium (LIDC) and image database resource initiative (idri): A completed reference database of lung nodules on CT scans. *Medical Physics*, *38*(2), 915–931. doi:10.1118/1.3528204 PMID:21452728

Bhargava, P., Dhand, S., Lackey, A. E., Pandey, T., Moshiri, M., & Jambhekar, K. (2013). Radiology education 2.0—on the cusp of change: Part 2. ebooks; file sharing and synchronization tools; websites/teaching files; reference management tools and note taking applications. *Academic Radiology*, *20*(3), 373–381. doi:10.1016/j.acra.2012.11.001 PMID:23452484

Dashevsky, B., Gorovoy, M., Weadock, W. J., & Juluru, K. (2015). Radiology teaching files: An assessment of their role and desired features based on a national survey. *Journal of Digital Imaging*, *28*(4), 389–398. doi:10.1007/s10278-014-9755-3 PMID:25582529

De-Arteaga, M., Eggel, I., Do, B., Rubin, D., Kahn, C. E. Jr, & Müller, H. (2015). Comparing image search behaviour in the arrs goldminer search engine and a clinical pacs/ris. *Journal of Biomedical Informatics*, *56*, 57–64. doi:10.1016/j.jbi.2015.04.013 PMID:26002820

Deshpande, P., Rasin, A., Brown, E., Furst, J., Raicu, D., & Montner, S. & Armato III, S. (2017). An Integrated Database and Smart Search Tool for Medical Knowledge Extraction from Radiology Teaching Files. In *Proceedings of The First Workshop Medical Informatics and Healthcare held with the 23rd SIGKDD Conference on Knowledge Discovery and Data Mining, in PMLR 69* (pp. 10-18).

Do, B. H., Wu, A., Biswal, S., Kamaya, A., & Rubin, D. L. (2010). Informatics in radiology: Radtf: A semantic search–enabled, natural language processor–generated radiology teaching file 1. *Radiographics*, *30*(7), 2039–2048. doi:10.1148/rg.307105083 PMID:20801868

Dos-Santos, M., & Fujino, A. (2012). Interactive radiology teaching file system: the development of a mirc-compliant and user-centered e-learning resource. In *2012 Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)* (pp. 5871–5874). IEEE.

Group, M. M. I. (2018, January 30). Mypacs tfs. Retrieved from https://www.mypacs.net/

Gutmark, R., Halsted, M. J., Perry, L., & Gold, G. (2007). Use of computer databases to reduce radiograph reading errors. *Journal of the American College of Radiology*, *4*(1), 65–68. doi:10.1016/j.jacr.2006.08.016 PMID:17412226

HospitalJ. H. (2017, May 31). Ctisus. Retrieved from http://www.ctisus.com/

Hwang, K. H., Lee, H., Koh, G., Willrett, D., & Rubin, D. L. (2016). Building and querying rdf/owl database of semantically annotated nuclear medicine images. *Journal of Digital Imaging*, 1–7. PMID:27785632

ICD10. (2017, November 11). Icd10. Retrieved from http://www.icd10data.com/

Jones, D. J. (2017, May 31). Radiopaedia. Retrieved from https://radiopaedia.org/

Kansagra, A. P., John-Paul, J. Y., Chatterjee, A. R., Lenchik, L., Chow, D. S., & Prater, A. B. (2016). Big data and the future of radiology informatics. *Academic Radiology*, *23*(1), 30–42. doi:10.1016/j.acra.2015.10.004 PMID:26683510

khresmoi (2017, June 28,). khresmoi. http://everyone.khresmoi.eu/hon-search/

Korenblum, D., Rubin, D., Napel, S., Rodriguez, C., & Beaulieu, C. (2011). Managing biomedical image metadata for search and retrieval of similar images. *Journal of Digital Imaging*, *24*(4), 739–748. doi:10.1007/s10278-010-9328-z PMID:20844917

Kushniruk, A. W., Kan, M.-Y., McKeown, K., Klavans, J., Jordan, D., LaFlamme, M., & Patel, V. L. (2002). Usability evaluation of an experimental text summarization system and three search engines: implications for the reengineering of health care interfaces. In *Proceedings of the AMIA Symposium* (pp. 420–424).

WebMDLLC. (2017, May 31). medscape. Retrieved from http://www.medscape.com

Margolies, L. R., Pandey, G., Horowitz, E. R., & Mendelson, D. S. (2016). Breast imaging in the era of big data: Structured reporting and data mining. *AJR. American Journal of Roentgenology*, *206*(2), 259–264. doi:10.2214/AJR.15.15396 PMID:26587797

Müller, H., Clough, P., Deselaers, T., Caputo, B., and CLEF, I. (2010). Experimental evaluation in visual information retrieval. *The Information Retrieval Series*, *32*.

NeutorgasseE. (2017, May 31). Eurorad. Retrieved from http://www.eurorad.org/

NIH. (January 30, 2017). Openi. https://openi.nlm.nih.gov/

NIST HIPAA Security Rule Toolkit. (May 14, 2017). NIST. https://csrc.nist.gov/projects/security-content-automation-protocol/hipaa

S.I.I.H.T.S.D. Organization. (2017, May 31). Snomedct ontology. Retrieved from http://www.snomed.org/

ReederM. M. (2017, May 31). Gamuts. Retrieved from http://gamuts.isradiology.org/Gamuts.htm

RSNA. (2017a, May 31). Radlex ontology. Retrieved from http://www.radlex.org/

RSNA. (2018b, January 30). Rsna tfs. Retrieved from http://mirc.rsna.org/query

Seitz, J., Schubert, S., Völk, M., Scheibl, K., Paetzel, C., Schreyer, A., & Strotzer, M. et al. (2003). Evaluation radiologischer lernprogramme im internet. *Der Radiologe*, *43*(1), 66–76. doi:10.1007/s00117-002-0836-9 PMID:12552377

SNOMED. (2017, July 24). Snomednlm. Retrieved from https://www.nlm.nih.gov/healthit/ snomedct/index.html

The A.R.R. Society. (2017, May 31). Goldminer. Retrieved from http://goldminer.arrs.org/

Solutions, M. H. (2017, May 31). Yottalook. Retrieved from http://www.yottalook.com/

Talanow, R. (2009). Radiology teacher: A free, internet-based radiology teaching file server. *Journal of the American College of Radiology*, *6*(12), 871–875. doi:10.1016/j.jacr.2009.08.001 PMID:19945043

Thies, C., Güld, M. O., Fischer, B., & Lehmann, T. M. (2004). Content-based queries on the casimage database within the IRMA framework. In *Workshop of the Cross-Language Evaluation Forum for European Languages* (pp. 781–792). Springer.

UMLS. (2017, January 30). UMLS. Retrieved from https://www.nlm.nih.gov

Weadock Software. L. (2017, May 31). Radpix. Retrieved from http://radpix.com/

Priya Deshpande is a Ph.D. candidate at DePaul University. Her major is Computer Science with specialization in database systems. Her research interest encompasses Databases, Big data, and Computer-aided diagnosis. Priya's previous research focus on Big data analytics and cloud computing. Her current research focus on Medical informatics radiology teaching file systems, where natural language processing and machine learning techniques are applied to improve diagnosis accuracy and help radiologists to take decision for diagnosis.

Alexander Rasin is an Assistant Professor in the College of Computing and Digital Media (CDM) at DePaul University. He received his Ph.D. and M.Sc. in Computer Science from Brown University, Providence. He is a co-Director of Data Systems and Optimization Lab at CDM and his primary research interest is in database forensics and cybersecurity applications of forensic analysis. Dr. Rasin's other research projects focus on building and tuning performance of domain-specific data management systems -- currently in the areas of computer-aided diagnosis and software analytics. Several of his research projects are supported by NSF.

Eli T. Brown earned his B.A. from Cornell University in Math and Computer Science and gained five years of experience in software engineering before pursuing a PhD at Tufts University. He is an Assistant Professor at DePaul University, where his work continues to blend machine learning and data visualization to enable people to get the best of both worlds in data analysis.

Jacob D Furst is an Professor in the College of Computing and Digital Media (CDM) at DePaul University. His research interests are in medical informatics with applications of machine learning and data mining to medical image processing and computer vision. His current work concentrates on being able to generate semantically meaningful information about lung nodules in computed tomography images of the human torso. Dr. Furst also has a strong interest in computer security and is the director of the DePaul Information Assurance Center. He has helped design two majors and three courses in the CDM security curriculum. He has taught Secure Electronic Commerce, Social Aspects of Information Security, Information Systems Security, Host Based Security, and Introduction to Networking and Security. Dr. Furst earned his PhD in computer science from UNC Chapel Hill; he has a master's degree in education and a bachelor's degree in English literature.

Steven M. Montner is a Professor of Radiology, the University of Chicago, Section of Thoracic Imaging, specializing in interstitial lung disease.

Samuel Armato received a B.A. in physics from The University of Chicago in 1987. After spending several years outside of academics, he returned to the University of Chicago and entered the Graduate Program in Medical Physics to pursue research in computer-aided diagnosis. He earned a Ph.D. from the program in 1997. He is currently an Associate Professor of Radiology and the Committee on Medical Physics at The University of Chicago, where he continues his work on computer-aided diagnosis, specifically for thoracic imaging. His main research interests include the automated detection of lung nodules in computed tomography (CT) scans, enhanced visualization techniques for chest radiography, and imaging-based tumor response assessment for lung cancer and mesothelioma.

Daniela Stan Raicu is a Professor of School of Computing, College of Computing and Digital Media at DePaul University, Chicago. She is the co-director of the Medical Informatics and the Intelligent Multimedia Processing Laboratories. Daniela is the founding Director of the Data Mining and Predictive Analytics Center at DePaul. Her research interests include medical imaging, multimedia retrieval, pattern recognition and data mining. Daniela's projects have been funded by the National Science Foundation (NSF), Argonne National Laboratory, Department of Education, and McArthur Foundation. She is the recipient of the DePaul Excellence in Teaching Award 2008, the DePaul Spirit of Inquiry Award in 2010, and the IBM Faculty Innovation Award in 2010. She also serves on the National Council of the Upsilon Pi Epsilon (UPE) honor society in Computer and Information Sciences since 2008. Daniela holds a Ph.D. in Computer Science from Oakland University, Michigan, a M.A. in Computer Science from Wayne State University, Michigan, and a B.S. in Mathematics from University of Bucharest, Romania.