

# CREATE: Clinical Report Extraction and Annotation Technology

Yichao Zhou<sup>†</sup>, Wei-Ting Chen<sup>†</sup>, Bowen Zhang<sup>†</sup>, David Lee<sup>†</sup>, J. Harry Caufield<sup>‡</sup>,  
Kai-Wei Chang<sup>†</sup>, Yizhou Sun<sup>†</sup>, Peipei Ping<sup>‡</sup> and Wei Wang<sup>†</sup>

<sup>†</sup> Department of Computer Science, University of California, Los Angeles.

<sup>‡</sup> Departments of Physiology, Medicine and Bioinformatics, University of California, Los Angeles.

{yz, weitingtw, bzhang0527, dkmllee, kwchang, yzsun, weiwang}@cs.ucla.edu;  
{jcaufield, pping}@mednet.ucla.edu

**Abstract**—Clinical case reports are written descriptions of the unique aspects of a particular clinical case, playing an essential role in sharing clinical experiences about atypical disease phenotypes and new therapies. However, to our knowledge, there has been no attempt to develop an end-to-end system to annotate, index, or otherwise curate these reports. In this paper, we propose a novel computational resource platform, CREATE, for extracting, indexing, and querying the contents of clinical case reports. CREATE fosters an environment of sustainable resource support and discovery, enabling researchers to overcome the challenges of information science. An online video of the demonstration can be viewed at <https://youtu.be/Q8owBQYTjDc>.

## I. INTRODUCTION

Case reports are a time-honored means of sharing observations and insights about novel patient cases [1], [2]. As of 2020, at least 160 case report journals were in existence, with over 90% having open access policies and almost half indexed by PubMed [3]. The narrative of a case report details the symptoms, diagnosis, treatment, and outcome of an individual, describing observations made over the course of clinical care.

Often these case reports contain exceptionally valuable clinical data, addressing unusual disease situations. To our knowledge, there has been no attempt to annotate, index, or otherwise curate these reports. Unlike other types of medical literature, there are no organizational frameworks or methodical review articles for case reports, and no metadata standards exist for their curation. A significant challenge exist for this rapidly growing corpus. Focusing our efforts, we address the domain of cardiovascular disease, an important area for which a range of research and clinical questions occur frequently. As shown in Figure 1, cardiovascular disease accounts for 20% of all case reports.

In this paper, we demonstrate an end-to-end systems incorporating algorithms for extracting, indexing, and querying the contents of clinical case reports. Our proposed system, Clinical Report Extraction and Annotation Technology (CREATE), automates generation of metadata about case reports, unlocking this important data resource through a searchable portal.

To make case reports findable and accessible, our primary innovation is a graph-based representation of each case report, where nodes signify concepts described in the narrative (e.g. sign/symptoms, diagnosis, etc). To build a case report’s graph,

we start by employing named entity recognition techniques to identify concepts (nodes), which are then standardized against existing biomedical ontology [2] to make the metadata interoperable. Concepts are then connected (edges) by detecting described temporal relationships in the text, facilitating retrieval over the patient chronologies that are hallmarks of case report descriptions. Collectively, CREATE framework will be used to share metadata around published case reports, making it reusable by others via public APIs.

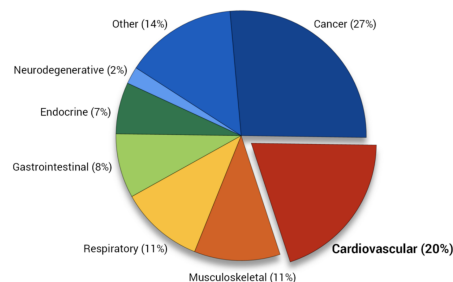


Fig. 1. Cardiovascular disease accounts for 20% of all case reports, and is the 2nd largest category of case reports after cancer.

In this demo, we present four key properties of our system.

- As of October 2020, it offers rich resources of over 10k reports for cardiovascular disease with depositions from a wide range of sources such as Scientific Literature and Authorized User Submissions, as shown in Figure 3.
- It provides a PDF submission service, based on Grobid [4], which is able to convert the publications in PDF format into well organized XML format. Metadata such as title, author, affiliation information can be automatically extracted for users by text mining technology.
- It is powered by CREATE-IR, a relation-based information retrieval system for clinical case reports, which outperforms solr [5]. Instead of simple keyword match, CREATE-IR embeds advanced deep learning algorithms to extract significant named entities and relations from the narrative. Relative case reports are retrieved from the database based on these structured knowledge. This also enables a temporal reasoning on the user queries and provides better search results.

- It provides a user-friendly interface for entity and relation annotation. A graph visualization of temporal order of the clinical events is generated for each document.

## II. SYSTEM ARCHITECTURE AND DESIGN

CREATE is a cloud-based application and its service is mainly hosted on the Amazon ECS (Fargate) with a CI/CD pipeline as illustrated in Figure 2.

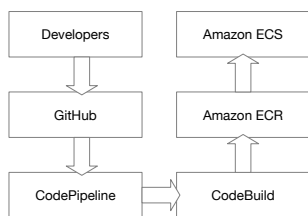


Fig. 2. The Development Pipeline of CREATE

As shown in Figure 3, the main feature of CREATE is to allow users perform CREATE-IR search, which will be elaborated in later sections. The two sources of data are case reports collected from PubMed and user-input reports. The frontend of CREATE is a single page application developed using React, which allows users to communicate with the RESTful API in the backend built with Express. The application is served by Nginx, a light-weight software for web serving. The majority of data for CREATE, is stored in the MongoDB server for persistency. The data in MongoDB server is queried via the Express backend to ensure security and consistency. This is the same for both Neo4j server and Elasticsearch server, which allow the application to perform complex search as explained in later section.

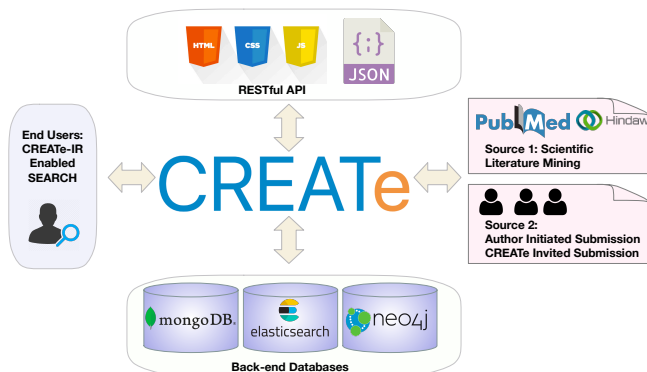


Fig. 3. The Architecture of CREATE

## III. CREATE-IR

### A. Data Source and Preprocessing

Identification of data source begins with a query to PubMed using the publication type and MeSH term filters to locate cardiovascular (CVD) case reports. A query of PubMed in six areas of CVD (cardiomyopathy, ischemic heart disease, cerebrovascular accidents, arrhythmias, congenital heart disease, and valve disease) returns around 118,000 case reports.

From these results, a web crawler (e.g. Apache Nutch) is used to locate the associated case reports and publication metadata. The contents can be captured in XML or online PDFs. The XML and PDF documents will be parsed into plain text to facilitate subsequent analysis, organized into case report sections and sentences.

### B. Data Annotations

To enable a supervised learning process for the automatic extraction of events/entities and relations from the clinical case reports, we build an embedded interface (Figure 4) for creating, editing, and visualizing document annotations based on the brat rapid annotation tool (BRAT) [6]. BRAT provides basic annotation primitives through the assignment of type labels to spans of text. These base annotations can be connected by directed or undirected edges using standard user interface gestures into  $n$ -ary associations, enabling the representation of sophisticated relational structures. In addition, we develop a comprehensive typing schema for information extraction from clinical narratives [2] and invite several medical experts to annotate hundreds of case reports based on this schema. Details are explained as following:

- We annotate EVENTS which are text elements representing situations or conditions and trigger a progression in a patient’s clinical course (e.g. *dyspnea* as *Sign/Symptom*).
- We annotate ENTITIES which are general non-trigger text elements which play a semantic role in the clinical narrative (e.g. *cotton farmer* as *Occupation*).
- We also annotate RELATIONS which define the associations between two EVENTS, or between EVENTS and ENTITIES. Relations include temporal relations, such as BEFORE, AFTER, and OVERLAP, which order events in time. Furthermore, they include semantic relations, such as IDENTICAL and MODIFY, which reflect the meaning between words.

### C. Information Extraction from User Queries

Once we receive a user query like “A patient was admitted to the hospital because of fever and cough.”, we apply two machine learning modules to parse the query and extract important knowledge from them (i.e. hospital (*Non-biological Location*), fever (*Sign/Symptom*), cough (*Sign/Symptom*), and the temporal relation between cough and fever (OVERLAP)). **Named Entity Recognition.** We first develop a machine learning module, called named entity recognizer, using contextualized token representations to locate and classify clinical terminologies into predefined categories, such as diagnostic procedure, disease disorder, severity, medication, medication dosage, and sign symptom. In specific, we pre-train a deep contextualized model C-FLAIR<sup>1</sup> with millions of parameters, in order to provide rich token embeddings for knowledge extraction. The training process is developed on one NVIDIA Tesla V100 GPU (16GB) and takes one week to finish. We

<sup>1</sup>The pre-trained model can be download at [https://drive.google.com/drive/folders/1b8PQyzTc\\_HUa5NRDqI6tQXz1mFXpJbMw?usp=sharing](https://drive.google.com/drive/folders/1b8PQyzTc_HUa5NRDqI6tQXz1mFXpJbMw?usp=sharing).




Fig. 4. Example of the BRAT annotation interface, showing labeled textual elements linked by temporal and semantic relational structures.

conduct experiments with three public datasets to test C-FLAIR and find our approach outperforms the state-of-the-art methods by 1.5% on average F1 score in terms of the extraction performance.

**Temporal Relation Reasoning.** We continue to predict the temporal relations among the extracted named entities. There is a consensus with the clinical community regarding the difficulty of temporal information extraction, due to the high demand for domain knowledge and high complexity of clinical language representations. we notice the dependencies among

The patient is a 47-year-old woman with long-term use of **glucocorticoids** She was **confirmed with COVID-19** by tested **positive of antibody** and **admitted to the hospital**. Just **a day later**, she began to have a **mild cough** and **nasal congestion**.

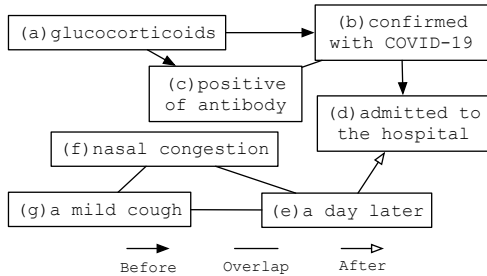


Fig. 5. An illustration of a clinical case report with its partial temporal graph where transitivity dependencies exist.

events in one clinical document are key enabler of classifying the temporal relations. For example in Figure 5, given that b happened before d, e happened after d and e happened simultaneously with f, we can infer according to the temporal transitivity rule that b was before f. However, existing methods either require expensive feature engineering or are incapable of modeling the global relational dependencies among the events. In our system, we build a temporal relation extraction module based on common dependencies such as transitivity and symmetry patterns. Specifically, as we predict a series of temporal relations from one clinical report, we compute a score to measure the satisfaction of all dependencies among these predicted relations, in order to generate an extra loss term to regularize the model training process. Extensive experiments on two public dataset, I2B2-2012 [7] and TB-Dense [8] show that our temporal relation extraction approach significantly outperforms baseline methods by 1.98% and

2.01% per F1 score.

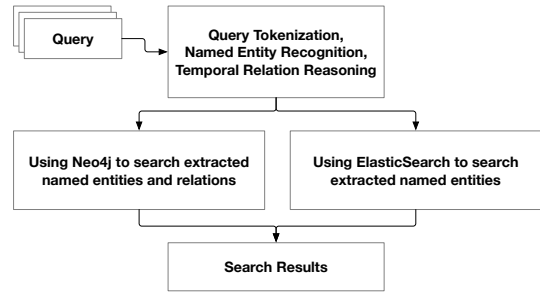


Fig. 6. The search workflow of CREATE-IR

#### D. Search Approach

We provides multiple search functions in our search engine, including search by keywords, search by entities and search by relations. ElasticSearch and Neo4j are utilized to build the search component, where ElasticSearch mainly handles keyword search and Neo4j handles entity&relation search. For independent resource occupation purposes, a collection of case reports are indexed separately on each search engine. Figure 6 illustrates the overall search flow. By default, Neo4j is the primary search engine in CREATE-IR system. The results returned by Neo4j will be placed on top, followed by results from ElasticSearch.

**Neo4j.** In Neo4j, data is saved as a graph of nodes and edges. Therefore, indexing each case report into Neo4j requires a transformation from texts to nodes and edges. A particular node will contain a *nodeId*, a *label* and a *entityType*. Property *label* keeps a natural language description for the node and property *entityType* represents the classification of this node. An edge will contain a *source*, a *target* and a *label*, which records the source nodeId and destination nodeId, as well as the relation type. Then, all nodes and edges are put into Neo4j via cypher query.

**ElasticSearch.** To better cater to the keyword search demand in CREATE-IR system, we build the document index with customized analyzer. In ElasticSearch, an analyzer can be divided into three sub-components: token filters, tokenizers and character filters. For token filter, we choose *asciifolding*, *lowercase*, *snowball*, *stop* and *stemmer*. For tokenizer, considering that some of the symptoms or medications may have longer names, we select N-gram tokenizer and customize it with *min\_gram=3* and *max\_gram=25*.




Fig. 7. Example network graph visualization representing a clinical case matching the query: “A patient was admitted to the hospital because of fever and cough”. The start of the graph reflects that the patient presented overlapping symptoms of fever and cough, thus aligning with the temporal and semantic structure of the query. The nodes and edges that follow reflect the evolution of the patient’s case, which ultimately ended in death due to respiratory failure.

### E. Event Visualization in Temporal Order

Annotations break the text structure down to event anchors and the temporal relations between them. These representations emphasize the atomic units and relational structures that are crucial for an understanding of the progression of a clinical case. CREATE-IR represents these underlying structures through network graph visualizations, which show the interconnections between sets of entities and events, giving focus to the semantic roles played by these fine-grained elements over the course of a clinical narrative. An example of such a visualization is shown in Figure 7. This visualization component is rendered using scalable vector graphics under a force-directed algorithm, which distributes nodes and clusters in space to minimize their repulsive energies and crossing edges. This component also possesses functionality for recognizing standard user interface gestures. The layout of nodes can be reconfigured by selecting a node with the mouse and dragging it to a different location in space. Similarly, the visualization window can be adjusted by zooming and panning using mouse wheel and drag gestures, respectively.

## IV. CONCLUSION

In this paper, we demonstrate CREATE, the first end-to-end system for annotating, indexing and curating the clinical case reports. CREATE is powered by CREATE-IR, which includes two state-of-the-art information extraction techniques to parse the important clinical events and temporal relationships for retrieving related clinical documents. We also embed a comprehensive annotation system and PDF parsing module to our system in order to facilitate further data collection.

## REFERENCES

- [1] A. J. Cabán-Martínez and W. F. García-Beltrán, “Advancing medicine one research note at a time: the educational value in clinical case reports,” *BMC Research Notes*, vol. 5, no. 1, p. 293, 2012.
- [2] J. H. Cauffield, Y. Zhou, Y. Bai, D. A. Liem, A. O. Garlid, K.-W. Chang, Y. Sun, P. Ping, and W. Wang, “A comprehensive typing system for information extraction from clinical narratives,” *medRxiv*, p. 19009118, 2019.
- [3] J. McEntyre and D. Lipman, “Pubmed: bridging the information gap,” *Cmaj*, vol. 164, no. 9, pp. 1317–1319, 2001.
- [4] P. Lopez, “Grobid: Combining automatic bibliographic data recognition and term extraction for scholarship publications,” in *International conference on theory and practice of digital libraries*. Springer, 2009, pp. 473–474.
- [5] D. Smiley, E. Pugh, K. Parisa, and M. Mitchell, *Apache Solr enterprise search server*. Packt Publishing Ltd, 2015.
- [6] P. Stenertorp, S. Pyysalo, G. Topić, T. Ohta, S. Ananiadou, and J. Tsujii, “Brat: a web-based tool for nlp-assisted text annotation,” in *Proceedings of the Demonstrations at the 13th Conference of the European Chapter of the Association for Computational Linguistics*, 2012, pp. 102–107.
- [7] W. Sun, A. Rumshisky, and O. Uzuner, “Evaluating temporal relations in clinical text: 2012 i2b2 challenge,” *Journal of the American Medical Informatics Association*, vol. 20, no. 5, pp. 806–813, 2013.
- [8] T. Cassidy, B. McDowell, N. Chambers, and S. Bethard, “An annotation framework for dense event ordering,” in *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Baltimore, Maryland: Association for Computational Linguistics, Jun. 2014, pp. 501–506. [Online]. Available: <https://www.aclweb.org/anthology/P14-2082>