

Summaries of Knowledge Graph Entities: First Steps to Measure the Relevance of Symptoms to Infer Diseases

Miguel Ángel Rodríguez-García¹, Carlos Badenes-Olmedo² and Soto Montalvo Herranz¹

¹Department of Computer Science, Universidad Rey Juan Carlos, 28933 Móstoles, Madrid, Spain

²Department of Artificial Intelligence, Universidad Politécnica de Madrid, 28040 Madrid, Spain

Abstract

Knowledge Graphs (KG) are concerned as one of the most efficient and effective knowledge integration approaches. In health domain, they have proven to be valuable resources that link clinical concepts by meaningful relations. This graph-structured information is usually extensive, and the data density it generates may make it difficult to perform tasks that involve human judgement, where the complexity and amount of information provided must be reduced. Consequently, it is required to develop techniques to reduce that large amount of data to more concise forms that facilitate their usage, visualization and analysis. In this paper, we propose a method for distilling the information available in a knowledge graph by creating entity summaries in the form of bags-of-words (BoW). Specifically, we create summaries of symptoms and diseases to measure their presence in medical records of patients. Our evaluation is focused on a vital healthcare worldwide problem, the early diagnosis of HIV in medical records. The proposed method summarizes the KG entities that represent each sign and symptom of acute HIV infection as a BoW and measures its relevance in a set of medical records. A labelled dataset with clinical notes has been compiled to evaluate the method and the results, with a precision and recall close to 0.6, make us optimistic about its performance as only syntactic matching of terms has been considered.

Keywords

Knowledge Graph Summaries, Bag-of-Word Representation, Medical Records, VIH Diagnosis

1. Introduction

Knowledge Graphs (KG) define a general model for representing information using a graph structure made of interlinked concepts [1]. This structured representation has made KGs one of the most effective and efficient knowledge integration technique, capable of incorporating information easily from diverse data sources independently of their structure [2]. This integrability and versatility have boosted the usage of KGs in academics and industry for knowledge harvesting in any area. Several KGs have been published, such as YAGO, NELL, Freebase or

KGSum'22: International Workshop on Knowledge Graph Summarization, October 23–24, 2022, Hangzhou, China

✉ miguel.rodriguez@urjc.es (M. Á. Rodríguez-García); carlos.badenes@upm.es (C. Badenes-Olmedo); soto.montalvo@urjc.es (S. M. Herranz)

🌐 <https://dblp.org/pid/124/9748.html> (M. Á. Rodríguez-García); <https://dblp.org/pid/186/2838.html> (C. Badenes-Olmedo); <https://dblp.org/pid/44/1331.html> (S. M. Herranz)

🆔 0000-0001-6244-6532 (M. Á. Rodríguez-García); 0000-0002-2753-9917 (C. Badenes-Olmedo); 0000-0001-8158-7939 (S. M. Herranz)



© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

Wikipedia, which are made of billions of entities that represent different facts about the world and are related by semantically meaningful relations. Generally, KGs are not a static graph structure but are continually being augmented with new facts that increase their size and complexity [3]. This growth has given rise to many challenges that aim at helping users to access, visualize and consume information efficiently.

But the massive scale of Knowledge Graphs could make them difficult to use, especially for users who need limited, domain-focused information [4]. Significant progress has been made to address this challenge by proposing data condensing and summarizing techniques that follow the same principle, promoting efficient information consumption and knowledge acquisition. These techniques have been employed in Recommendation Systems [5] and several areas such as Biomedicine [6] or Geographic [7]. Our work is focused on Medicine, specifically in the field of diseases, and aims to create a system that helps to infer diseases from the symptoms and related diseases described in patients' clinical notes by leveraging the information provided by Knowledge Graphs. Medical practitioners create clinical notes, using natural language text, to describe the symptomatology of a patient who comes to the hospital or health centre with a particular health problem. Many of the symptoms described in these documents may be mentioned explicitly (e.g. myalgia) or by related terms that could be more (e.g. fascia) or less specific (e.g. muscle pain). In this sense, identifying the presence of symptoms and diseases in clinical notes is crucial, as they may reveal diagnostically relevant diseases.

We address the identification of HIV/AIDS disease-related symptoms in clinical notes of undiagnosed patients due to it represents a severe problem in worldwide health. If all infected people were diagnosed, it would be easy to curb this illness pandemic and reach the target '95-95-95' of ONUSIDA, which means that 95% of people with HIV are diagnosed, from which 95% are undergoing treatment and at least 95% with undetectable viral load ¹.

The remainder of the work is organized as follows: Section 2 starts analyzing the applications of Knowledge Graphs in various fields related to research on healthcare and concludes by pointing out its needed in the context of HIV. Section 3 details the method developed, carrying out a modular decomposition that explains how it works in depth. Section 4 describe the dataset compiled and the strategy designed to evaluate the method. Finally, Section 5 itemizes the most relevant obtained findings and future research lines to explore.

2. Related Work

Medical knowledge graph have proven to be valuable resources in healthcare applications and medical research, with increasing use in recent years. [8] proposed a systematic approach to build medical KG from EMRs. [9] made a study to learn high quality knowledge bases linking diseases and symptoms directly from electronic medical records. They showed that direct and automated construction of high-quality health knowledge graphs from medical records using rudimentary concept extraction is feasible. [10] developed a rare disease classification algorithm that made effective use of a knowledge graph, even when the graph was imperfect. To assist in disease diagnosis it is usual to find works that use text classification from its medical records to find the diagnosis or identify relevant concepts about it [11][12][13]. Also, there are some works that

¹https://www.unaids.org/sites/default/files/media_asset/global-AIDS-strategy-2021-2026_en.pdf

use Knowledge Graphs to discover the diagnosis. In [14] is used a Knowledge Graph to connect trivial and scattered knowledge in various medical information systems. Experiments shown that the thyroid disease diagnosis method that combines Knowledge Graphs and deep learning has a better diagnostic effect, compared with traditional machine learning algorithms. Similarly, [15] proposed integrates machine learning algorithms and Knowledge Graph technology to help patients conduct online consultations.

In the specific case of HIV research with clinical notes or EMRs, Feller et al., in [16] examine whether pairing Natural Language Processing (NLP) with Machine Learning algorithms can improve HIV diagnosis. Under the same Machine Learning model, Random Forest, they conducted three experiments using input variables extracted from electronic health records (EHRs). In each experiment, they varied the inputs of the model, in the first experiment used only structured EHR data, in the second and third experiments employed two different Natural Language Processing strategies to extract features, automated keyword identification by utilising frequency analysis and automated topic modelling by using Latent Dirichlet allocation model. In their study, the results show that combination techniques perform better. In the same context, Oliwa et al., in [17] applied a similar combination of techniques, Machine Learning and Natural Language Processing monitoring the treatment plan of HIV-positive patients. In this work, they labelled clinical notes with several patient demographic parameters to identify the factors that push these patients to fall out of care when they are under treatment. This classification problem was modelled as a binary classification task in which they employed N-grams in conjunction with frequency analysis for extracting features and a supervised machine learning system with a linear model and elastic NET regularization for predicting this status. Closer to the domain of the work presented, Sun et al., in [18], proposed the construction of a visual knowledge graph with information about HIV-associated Neurocognitive Disorder (HAND). The KG contains entities and their relations extracted from text data, such as electronic medical records and the medical community. We have not found any other work in this research line, which encourages us to continue with the study we propose here.

3. Approach

Our proposal leverages KGs to identify the presence of symptoms and diseases in a medical record. The information in a KG is structured in triples containing facts about the value of a property of an entity. An entity is a real-world concept, for example a symptom or a disease, and the property may refer to an attribute of that entity and have a literal value (e.g. `<'myalgia[Q474959]','schema:description','muscle pain'>`), or it may be a relation to another entity (e.g. `<'myalgia[Q474959]','instance_of[P31]','physiological condition[Q7189713]'>`). Thanks to this structure of information, which contextualizes concepts through their relationships and describes them based on their attributes, we represent diseases and symptoms based on KG entities so that their presence in a medical record can be measured.

We have created a method to summarize KG entities based on their property values. Figure 1 represents the pipeline that we have created in this approach. Firstly, the Knowledge Graph is queried with a list of HIV indicators created by experts. Given a list of HIV indicators from experts, it is first created a summary for each word. A summary is a bag-of-words (BoW), i.e., a

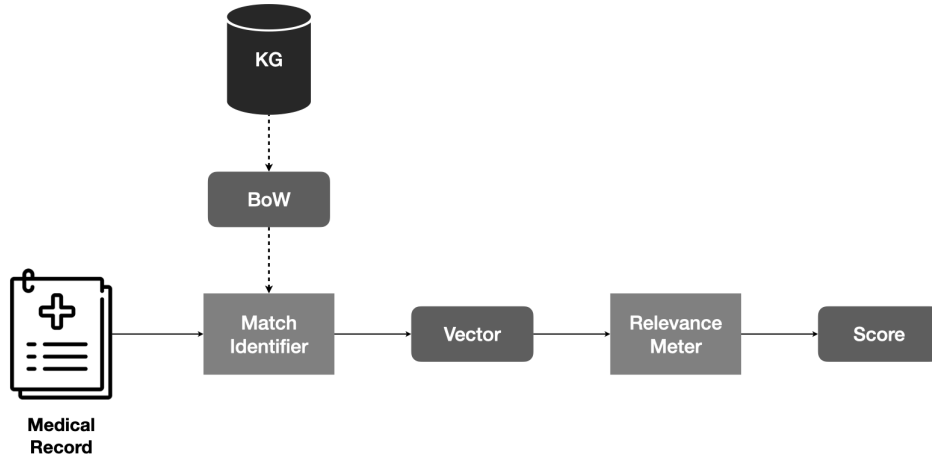


Figure 1: Tasks involved in the KG Entity Summarization algorithm

set of terms with weights that have no order or sequence, based on the property values of the entities. Given an index e from the HIV indicators, which has i properties p_i and i attributes a_i , the summarized BoW W_e consists of the attributes a_i . When a_i is another entity, the value associated with the 'alt:label' attribute of a_i is considered, otherwise a_i . Figure 2 depicts the summarization process of the Knowledge Graph for the myalgia indicator. It exemplifies the formal method described above.

Following the explanation of the pipeline design, the Match Identifier is responsible to measure the presence of KG entities in medical records using the BoW-based summaries. We propose a dictionary-oriented representation with the BoW retrieved from the KG that is used to tokenize the text of the medical records, changing their representations to vectors of occurrences found. How the entity is present depends on the terms identified in the medical record that belong to the dictionary created from the KG entities extracted. The measure of the presence of a concept represented by our BoW-based summary in a medical record depends on the number of shared terms and their relevance. The correspondence between terms can be *syntactic*, i.e. they match character to character, or *semantic*, i.e. they have related meanings. Our work considers syntactic matching only. Term weights, in turn, may or may not be influenced by the KG structure. They may all be equally important or vary according to the level of depth of the entity properties. In summary, we consider the following alternatives for defining the criteria for the relevance of terms when creating the BoW, taking into account syntactic matching, i.e. only the weight of terms from the BoW that match completely are considered:

- **Static Relevance:** all terms are equally important (e.g. weight=1 for each Bow term).
- **Dynamic Relevance:** each term has its own weight, which may or may not match the weight of other terms.
 - *based on Depth:* The main idea behind the strategy is to measure the relevance of related words taking into account the hierarchical distance between them. The strategy is mainly based on two parameters, initial weight and delta. The initial

weight stands for a pre-established value to express the relevance of each word in the BoW set. Initially, the same initial weight is assigned for all words belonging to the set. Delta represents a factor designed to penalize each word's weight considering its hierarchy depth relating to a word given. Thus, words closer to the hierarchy will be more relevant than those away from.

- *based on Frequency*: repetitions between terms associated to attributes of different levels are penalized when they coincide, according to the TF/IDF strategy [19]. Less frequent terms with less depth will be more relevant than the others.

Considering variations in the matching and relevance criteria, we propose the following methods to summarize a KG entity in a BoW and measure its presence in a text: (1) *Syntactic Matching + Static Relevance*; (2) *Syntactic Matching + Depth-Based Dynamic Relevance*; (3) *Syntactic Matching + Frequency-based Dynamic Relevance*.

Finally, independently of the variation matching strategies utilized, once the vector representations are created, the Relevance Meter computes a score by utilising a measure of similarity to estimate the distance between both vectors, the indicators and the medical record. In this approach, we employed cosine similarity, since the documents to compare are represented as numerical vectors.

4. Evaluation

In this section, we first describe the dataset compiled and two types of experiments configured to analyze the accuracy of the methodology summarization proposed. The first consists of a generic evaluation in which we contemplate the HIV diagnosis as a binary classification problem, where the target is to differentiate between HIV diagnosed/non-diagnosed clinical notes. The remaining analysis is more specific, addressing the diagnosis problem from another perspective, the precociousness of being able to identify the disease in its early stages.

4.1. Dataset and Knowledge Graph

Compiling the dataset started with a review of existing ones created for challenges related. There were analyzed the followings: anonymization shared task [20], biomedical abbreviation challenge [21], PharmaCoNER shared task [22], Cantemist shared task [23], MEDDOPROF Shared Task [24], CodiEsp track for CLEF ehealth 2020 [25], SPACCC [26], a Spanish clinical corpus composed by 1,000 clinical cases from SciELO, among others. We used Apache Solr to carry out the first sieve and identify cases in which appears HIV illness. At the end of the sifting process, we collected 47 clinical notes, from which 28 notes were classified as HIV-diagnosed and 19 non-diagnosed. Within HIV-related clinical notes, we manually distinguish between Consolidated HIV (more than 5 years with the illness), Recent HIV (about 5 years or less with the illness) and Diagnosed HIV (the illness was diagnosed in the report).

Our experiments were conducted on Wikidata, a public knowledge graph aligned with Wikipedia, that is hugely popular as a crowdsourced collection of knowledge. It provides a new ways for accessing to information published in Wikipedia. In the experiments, we have utilised the last version, that contains specifically 602,538 items and 1,095,368 statements. We

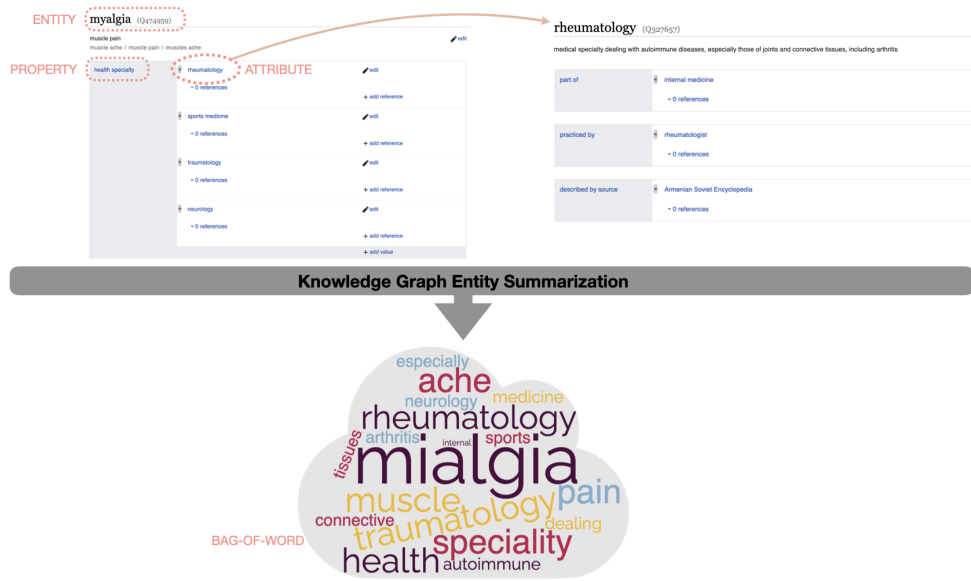


Figure 2: Creation of a Bag-of-Words representation of a Knowledge Graph Entity

extract disease-associated terms from entities in Wikidata, and create dictionaries that are used to analyse medical records.

4.2. Results

The evaluation consisted of preparing two experiments to analyse the accuracy of the summarization techniques proposed. Both experiments are focused on finding indicators of the text content of clinical notes to enable inferring a possible infection of HIV. The first experiment is a binary classification problem and stands for identifying clinical notes with HIV diagnosed from symptoms and diseases related. The second experiment is a more complex problem, concerning a multiclass classification problem, which includes a relevant parameter in the classification, the precociousness of the clinical diagnosis. Hence, we open the range of classification possibilities, including more labels and consequently making a more complex scenario. Below, we detail the experiments and their results gained.

The experiment was conducted as follows, considering the pre-selected symptoms and related diseases given by clinicians and the summary strategies described in Section 3. BoW distributions are created and utilised as a dictionary to find occurrences in the content text of clinical notes. Next, clinical notes are codified according to these dictionaries, considering the terms identified. These representations are compared by utilising a measure of similarity that quantifies the relativeness between both representations, clinical notes and symptoms/related diseases. Lastly, the clinical note is classified as HIV positive or negative, depending on a threshold defined. For this experiment, we set as such limit below 0.6. Table 1 shows the result obtained.

In this first experiment, the best precision was obtained by utilising symptoms as word

Table 1

Results obtained from the first experiment in terms of Precision, Recall and F-measure. The strategies considered are (1) *Syntactic Matching + Static Relevance*; (2) *Syntactic Matching + Depth-Based Dynamic Relevance*; (3) *Syntactic Matching + Frequency-based Dynamic Relevance*.

| | | Precision | Recall | F-measure |
|----------|-----|-----------|--------|-----------|
| Symptoms | (1) | 1 | 0.07 | 0.13 |
| | (2) | 0.5 | 0.36 | 0.42 |
| | (3) | 0.57 | 0.61 | 0.59 |
| Diseases | (1) | 0.83 | 0.36 | 0.5 |
| | (2) | 0.58 | 0.5 | 0.54 |
| | (3) | 0.42 | 1 | 0.59 |

distributions and the first strategy based on static relevance, in which the classifier has reached the max score of 1. On the other hand, the best recall was achieved by utilising diseases and the technique based on frequency, where the classifier scores the max value, reaching the highest F-measure score of 0.59.

The second experiment goes to a more specific field, with the target of the precocity diagnostic of HIV. The hypothesis behind the second experiment is concerning to the existence of illnesses that patients suffered during the first stages of contagion. For instance, “tuberculosis” or “pneumonia” are considered defining diseases and are related to advanced settings of the contagion, where if the patient suffers from them, there are a high probability of having AIDS [27]. Conversely, “mononeuritis” or “hepatitis” are classified as indicating diseases related to initial stages, where the early diagnosis can help to considerably increase the effectiveness of the treatments against the virus, increasing the patients’ life expectancy and even saving their life. Hence, the second experiment passed to a more complex problem, a multiclass classification task that addressed the prediction of stages of HIV disease. The defined labels were: i) diagnosed phase, which means the first stages; ii) recent phase to describe recent contagious, and iii) consolidated phase in which the patient has a high percentage of suffering from the disease. To carry out the experiment, it was utilised a set of related diseases was provided to clinicians. Then, we utilise these sets in the same way as the first experiment to codify the clinical notes and compute the results. An example of this experiment would be, having a clinical note, where it has been diagnosed as a fungal infection and an intestinal disorder. Then, the clinical note will be expressed considering the pre-defined related disease for the three stages of the HIV disease, consolidated, recent and diagnosed. Concretely, in this example, we have described the worst case, where the patient has a high probability of suffering from this disease since the diagnosed symptoms are related to the Cryptosporidiosis and Candidiasis diseases, classified as a defining indicators. Hence, considering the consolidated markers and, specifically, the candidiasis disease, a similar BoW representation will be summarized from Wikidata: ‘Candidiasis’, ‘Opportunistic mycosis’, ‘Systemic mycosis’, ‘Opportunistic infection’, ‘fungal infectious disease’, ‘fungal infection’, among others. Then, each clinical note will be represented by using this set of words structure. Depending on the relevance criteria selected, each match will be represented with a different number. As a result, a string of numbers is

obtained and the cosine similarity is employed to check the similitude between both vectors. Then, if the obtained value is higher than a pre-established threshold, the note will be classified as a consolidated HIV. Table 2 shows the results obtained in the experiment.

Table 2

Results obtained from the second experiment in terms of Precision, Recall and F-measure.

| | | Precision | Recall | F-measure |
|------------------|-----|------------------|---------------|------------------|
| Consolidated HIV | (1) | 0.1 | 0.14 | 0.12 |
| | (2) | 0.14 | 0.29 | 0.19 |
| | (3) | 0.3 | 0.43 | 0.35 |
| Recent HIV | (1) | 0 | 0 | 0 |
| | (2) | 0.4 | 0.18 | 0.25 |
| | (3) | 0 | 0 | 0 |
| Diagnosed HIV | (1) | 0 | 0 | 0 |
| | (2) | 0.4 | 0.22 | 0.28 |
| | (3) | 0.31 | 0.89 | 0.46 |

In the second experiment, the precision scores are quite low compared to the first one. The highest results are reached with the strategy based on depth for the recent and diagnosed use cases. For recall, however, the highest value has been obtained by the matching method based on the frequency that achieved 0.89, making reaching the best F-measure through all experiments conducted. Concretely, the 0s obtained in Recent and Diagnosed HIV are due to, although the symptoms and diseases are correctly identified as we can see in Table 1, the classifier fails to assign the clinical note to the disease stage accurately.

5. Conclusions and Future Work

In this work, we proposed a preliminary study to assess the viability of knowledge graph summarization in the medical domain, explicitly assisting in diagnosing problems. We proposed various techniques to summarize knowledge graphs as a primary data structure for organizing the information. We validate the designed techniques for a vital problem in our society, early diagnosis ADS. In particular, we face a problem in evaluating from two different perspectives. First, we assess the word distributions generated by conducting a binary classification between HIV diagnosed/non-diagnosed clinical notes. The evaluation is performed by utilizing symptoms provoked by the virus and related diseases. Second, we accomplished a more concrete assessment, where we further tried to identify early HIV diagnoses in the clinical notes taking into account different phases of the virus contagion.

Given the results obtained in this preliminary study, we conclude that the word distributions obtained from the graph can be useful in this application domain. However, the results are lowly in terms of accuracy of HIV early diagnosis, and there is still scope for improvement.

In future work, we would like to explore more classifying techniques for the experiments and include more data resources in the study since we believe enriching the word distribution sets utilised could increase the system's accuracy. Besides, we would like to have a closer

collaboration with clinicians for making a more robust set of symptoms and diseases since, as the first drawback to attending carefully, we have considered that all of them are relevant for the diagnosis, but we believe that it is not entirely true.

Acknowledgments

This work has been partially supported by the projects *DOTT-HEALTH* (PID2019-106942RB-C32, MCI/AEI/FEDER, UE) and ‘*DRUGS4COVID++*’ through grants “Ayudas Fundación BBVA a equipos de investigación científica SARS-CoV-2 y COVID-19”, ISCIII (PI20/00715, co-funded by ERDF/ESF, “A way to make Europe”/“Investing in your future”), “Programa para la Recualificación del Sistema Universitario Español 2021-2023”, and the Community of Madrid, through the Young Researchers R+D Project. Ref. M2173 – SGTRS (co-funded by Rey Juan Carlos University).

References

- [1] Z. Chen, Y. Wang, B. Zhao, J. Cheng, X. Zhao, Z. Duan, Knowledge graph completion: A review, *Ieee Access* 8 (2020) 192435–192456.
- [2] S. Tiwari, F. N. Al-Aswadi, D. Gaurav, Recent trends in knowledge graphs: theory and practice, *Soft Computing* 25 (2021) 8337–8355.
- [3] T. Safavi, C. Belth, L. Faber, D. Mottin, E. Müller, D. Koutra, Personalized knowledge graph summarization: From the cloud to your pocket, in: 2019 IEEE International Conference on Data Mining (ICDM), IEEE, 2019, pp. 528–537.
- [4] R. Marois, J. Ivanoff, Capacity limits of information processing in the brain, *Trends in cognitive sciences* 9 (2005) 296–305.
- [5] J. A. Sacenti, R. Fileto, R. Willrich, Knowledge graph summarization impacts on movie recommendations, *Journal of Intelligent Information Systems* 58 (2022) 43–66.
- [6] Y. Yu, K. Huang, C. Zhang, L. M. Glass, J. Sun, C. Xiao, Sumgnn: multi-typed drug interaction prediction via efficient knowledge graph summarization, *Bioinformatics* 37 (2021) 2988–2995.
- [7] B. Yan, K. Janowicz, G. Mai, R. Zhu, A spatially explicit reinforcement learning model for geographic knowledge graph summarization, *Transactions in GIS* 23 (2019) 620–640.
- [8] L. Li, P. Wang, J. Yan, Y. Wang, S. Li, J. Jiang, Z. Sun, B. Tang, T.-H. Chang, S. Wang, Y. Liu, Real-world data medical knowledge graph: construction and applications, *Artificial Intelligence in Medicine* 103 (2020).
- [9] M. Rotmensch, Y. Halpern, A. Tlimat, S. Horng, D. Sontag, Learning a health knowledge graph from electronic medical records, *Scientific Reports* 7 (2017).
- [10] X. Li, Y. Wang, D. Wang, W. Yuan, D. Peng, Q. Mei, Improving rare disease classification using imperfect knowledge graph, *BMC Medical Informatics and Decision Making* 19 (2019).
- [11] M. Almagro, R. Martínez, V. Fresno, S. Montalvo, ICD-10 coding of spanish electronic discharge summaries: An extreme classification problem, *IEEE Access* 8 (2020) 100073–100083. URL: <https://doi.org/10.1109/ACCESS.2020.2997241>. doi:10.1109/ACCESS.2020.2997241.
- [12] M. Jamaluddin, A. D. Wibawa, Patient diagnosis classification based on electronic medical

- record using text mining and support vector machine, in: 2021 International Seminar on Application for Technology of Information and Communication (iSemantic), 2021, pp. 243–248. doi:10.1109/iSemantic52711.2021.9573178.
- [13] X. Li, H. Wang, H. He, J. Du, J. Chen, J. Wu, Intelligent diagnosis with chinese electronic medical records based on convolutional neural networks, *BMC bioinformatics* 20 (2019) 1–12.
 - [14] X. Chai, Diagnosis method of thyroid disease combining knowledge graph and deep learning, *IEEE Access* 8 (2020).
 - [15] F. Hao, K. Zheng, Online disease identification and diagnosis and treatment based on machine learning technology, *Knowledge-Based Systems Access* 235 (2022).
 - [16] D. J. Feller, J. Zucker, M. T. Yin, P. Gordon, N. Elhadad, Using clinical notes and natural language processing for automated hiv risk assessment, *Journal of acquired immune deficiency syndromes* (1999) 77 (2018) 160.
 - [17] T. Oliwa, B. Furner, J. Schmitt, J. Schneider, J. P. Ridgway, Development of a predictive model for retention in hiv care using natural language processing of clinical notes, *Journal of the American Medical Informatics Association* 28 (2021) 104–112.
 - [18] D. Sun, Y. Peng, H. Li, Construction of knowledge graph of hiv-associated neurocognitive disorders syndrome based on deep learning, in: *Proceedings of the 2020 ICAICE*, 2020, pp. 134–141.
 - [19] G. Salton, M. McGill, *Introduction to modern information retrieval* (pp. paginas 400), Google Scholar Google Scholar Digital Library Digital Library (1986).
 - [20] M. Marimon, A. Gonzalez-Agirre, A. Intxaurre, H. Rodriguez, J. L. Martin, M. Villegas, M. Krallinger, Automatic de-identification of medical texts in spanish: the meddocan track, corpus, guidelines, methods and evaluation of results., in: *IberLEF@ SEPLN*, 2019, pp. 618–638.
 - [21] A. Intxaurre, J. C. de la Torre, H. Rodríguez Betanco, M. Marimon, J. A. Lopez-Martin, A. Gonzalez-Agirre, J. Santamaria, M. Villegas, M. Krallinger, Resources, guidelines and annotations for the recognition, definition resolution and concept normalization of spanish clinical abbreviations: the barr2 corpus, in: *Proceedings of SEPLN*, 2018, pp. 1–9.
 - [22] A. Gonzalez-Agirre, M. Marimon, A. Intxaurre, O. Rabal, M. Villegas, M. Krallinger, Pharmaconer: Pharmacological substances, compounds and proteins named entity recognition track, in: *Proceedings of The 5th Workshop on BioNLP Open Shared Tasks*, 2019, pp. 1–10.
 - [23] A. Miranda-Escalada, E. Farré, M. Krallinger, Named entity recognition, concept normalization and clinical coding: Overview of the cantemist track for cancer text mining in spanish, corpus, guidelines, methods and results., in: *Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2020)*, CEUR Workshop Proceedings, 2020, pp. 303–323.
 - [24] S. Lima-López, E. Farré-Maduell, A. Miranda-Escalada, V. Brivá-Iglesias, M. Krallinger, Nlp applied to occupational health: Meddopprof shared task at iberlef 2021 on automatic recognition, classification and normalization of professions and occupations from medical texts, *Procesamiento del Lenguaje Natural* 67 (2021) 243–256. URL: <http://journal.sepln.org/sepln/ojs/ojs/index.php/pln/article/view/6393>.
 - [25] A. Miranda-Escalada, A. Gonzalez-Agirre, J. Armengol-Estapé, M. Krallinger, Overview of automatic clinical coding: Annotations, guidelines, and solutions for non-english clinical

- cases at codiesp track of CLEF ehealth 2020, in: L. Cappellato, C. Eickhoff, N. Ferro, A. Névél (Eds.), Working Notes of CLEF 2020 - Conference and Labs of the Evaluation Forum, Thessaloniki, Greece, September 22-25, 2020, volume 2696 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2020. URL: http://ceur-ws.org/Vol-2696/paper_263.pdf.
- [26] A. Intxaurreondo, M. Marimon, A. Gonzalez-Agirre, J. A. López-Martín, H. Rodriguez, J. Santamaría, M. Villegas, M. Krallinger, Finding mentions of abbreviations and their definitions in spanish clinical cases: The barr2 shared task evaluation results, in: *Proceedings of the IberEval 2018*, 2018, pp. 280–289.
- [27] R. Griesel, A. Stewart, H. van der Plas, W. Sikhondze, M. Mendelson, G. Maartens, Prognostic indicators in the world health organization’s algorithm for seriously ill hiv-infected inpatients with suspected tuberculosis, *AIDS research and therapy* 15 (2018) 1–9.