

Automatic Coding at Scale: Design and Deployment of a Nationwide System for Normalizing Referrals in the Chilean Public Healthcare System

Fabián Villena

Center for Mathematical Modeling
& Department of Computer Sciences
University of Chile
fabian.villena@uchile.cl

Matías Rojas

Center for Mathematical Modeling
University of Chile
matias.rojas.g@ug.uchile.cl

Felipe Arias

Center for Mathematical Modeling
University of Chile
felipe.arias.t@ug.uchile.cl

Jorge Pacheco

Dept. of Statistics and Health Information
Chilean Ministry of Health
jorge.pacheco@minsal.cl

Paulina Vera

Dept. of Statistics and Health Information
Chilean Ministry of Health
paulina.vera@minsal.cl

Jocelyn Dunstan

Dept. Computer Science & IMC
Pontifical Catholic University of Chile
jdunstan@uc.cl

Abstract

The disease coding task involves assigning a unique identifier from a controlled vocabulary to each disease mentioned in a clinical document. This task is relevant since it allows information extraction from unstructured data to perform, for example, epidemiological studies about the incidence and prevalence of diseases in a determined context. However, the manual coding process is subject to errors as it requires medical personnel to be competent in coding rules and terminology. In addition, this process consumes a lot of time and energy, which could be allocated to more clinically relevant tasks. These difficulties can be addressed by developing computational systems that automatically assign codes to diseases. In this way, we propose a two-step system for automatically coding diseases in referrals from the Chilean public healthcare system. Specifically, our model uses a state-of-the-art NER model for recognizing disease mentions and a search engine system based on Elasticsearch for assigning the most relevant codes associated with these disease mentions. The system's performance was evaluated on referrals manually coded by clinical experts. Our system obtained a MAP score of 0.63 for the subcategory level and 0.83 for the category level, close to the best-performing models in the literature. This system could be a support tool for health professionals, optimizing the coding and management process. Finally, to guarantee reproducibility, we publicly release the code of our models and experiments.

1 Introduction

The clinical text represents a significant proportion of patient's health records, commonly found in a non-structured format. These texts have particular challenges due to the extensive use of abbreviations, the variability of clinical language across medical specialties, and its restricted availability for privacy reasons (Dalianis, 2018). Due to the complexity of its analysis, this data is commonly discarded in projects that seek to support clinical decision-making (Kong, 2019).

Clinical coding involves mapping medical texts into codes using a controlled vocabulary consistent across different departments, hospitals, or even countries (Dong et al., 2022). The World Health Organization maintains an open, controlled vocabulary called the International Classification of Diseases (ICD), which is used in almost every country. Currently, the most widely used revision is the tenth (ICD-10) (World Health Organization, 2015), and they are developing its eleventh revision, which will include not only diseases (World Health Organization, 2023).

Regarding the Chilean public health system, the ICD-10 terminology is used for coding hospital discharges (morbidity coding by each healthcare provider) and deaths (mortality coding by the Ministry of Health). Having patients' data normalized using these controlled vocabularies enables the ability to summarize information automatically and

not deal with the noisiness of free-text data. The already-digested information from the normalized data empowers data analysts who are not experts in NLP to add more complex information into their workflows.

The Waiting Time Management System (SIGTE, in Spanish) contains electronic records of referrals from the Chilean Waiting List, which is the system that manages the high demand existent for consultation by specialists (Ministerio de Salud de Chile, 2011). This data provided by 29 health services contain information about the medical diagnoses of patients but is not standardized (Báez et al., 2022).

As of November 2022, SIGTE recorded 25,374,491 waiting list referrals, of which 18,716,629 correspond to "new specialty referrals" and are associated with patient pathologies. Of these referrals, approximately 5,760,750 (30.7 %) have an ICD-10 code. This calculation was performed by searching for a regular expression formatted as an ICD-10 code in the free-text diagnosis fields.

Clinical experts perform the disease coding task manually, which is not optimal for several reasons. Firstly, since this process is subject to errors, medical personnel must have significant competence in coding rules and a thorough knowledge of specialized terminologies, such as ICD, which also get updated frequently. In other words, expert coding staff must be familiar with the clinical field, analytical and focused, and have fundamental skills for inspecting and analyzing highly specialized texts. In addition, manual coding is time-consuming (Yan et al., 2022), which could be optimized by a support system, and this time could be used for other tasks relevant to clinical decision-making.

These difficulties can be efficiently addressed using computational systems capable of automatically performing the coding task using NLP. Currently, most automatic coding systems are based on an end-to-end architecture based on deep learning techniques. Although these systems have boosted the performance of several coding tasks, they cannot incorporate context-specific rules, such as code priority, medical assumptions, code definition, and synonyms.

In this work, we developed an automated disease coding system, thus being able to code the entire historical waiting list in Chile, identifying a total of 18,716,629 referrals. Our system is based on two steps; first, the automatic extraction of diseases is

addressed using a state-of-the-art NER model, and then, using a search engine, the most probable code for each disease found is identified. Finally, we explored the potential applications derived from this system and studied in more depth the most frequent diseases in the country today.

2 Related Work

The disease coding task involves transforming clinical texts, commonly written by physicians in a non-structured format, into codes following medical terminologies. This is not an easy task since a medical ontology such as ICD in Spanish has 14,668 codes, an example of extreme multi-label classification (Barros et al., 2022). We have identified two major groups of computational methods proposed to solve this task; rule-based coding and neural network-based coding.

2.1 Rule-based Models

This approach involves designing hand-crafted rules to represent and simulate the flow that clinical experts follow when assigning codes. Most of the studies are based on using regular expressions and keywords to transform diseases found in the text into their respective codes. However, these methods are not feasible since manually capturing all the relations between texts and codes is time-consuming and complex.

Different approaches based on machine learning have been proposed to address this issue. In this way, features extracted from statistical models such as decision trees and support vector machines, among others, are incorporated into the manual rules (Stanfill et al., 2010; Teng et al., 2022; Farkas and Szarvas, 2008). Another method is to create a list of synonyms of the original text to calculate a word distance with respect to the code descriptions of the terminology. Despite their disadvantages, these methods have yielded high results in the literature, effectively supporting manual coding performed by humans (Zhou et al., 2020).

2.2 Models based on neural networks

Deep learning-based methods have significantly improved the disease coding task in recent years. The advantage of using these models is that the healthcare-specific domain knowledge is no longer needed for the manual development of complex rules. In contrast, these methods can automatically build features powerful enough to capture the rela-

tionships between clinical texts and their respective codes.

Most proposed systems are based on posing the problem as a multi-label text classification task (Karimi et al., 2017; Mullenbach et al., 2018; Yu et al., 2019; Cao et al., 2020). Thus, the algorithm's input is text, while the output can be one or more codes associated with diseases. Unlike traditional text classification problems, this problem is considered extreme since the number of possible labels increases to thousands (depending on the terminology).

The main disadvantage of this approach is that manual coding requires incorporating context-specific rules, such as code priority, medical assumptions, code definition, and synonyms, among other types of information, to improve system performance. In the case of deep learning, this is not considered since the systems are commonly created using an end-to-end approach, meaning that no human knowledge is involved when creating the features or making the predictions.

To solve the previous problem, we followed another approach used in the literature, which consists of mixing the previous ideas using two sequential steps; the first one uses deep learning algorithms, while the second allows us to incorporate medical knowledge into the computational system. Firstly, we used a Named Entity Recognition model for automatically recognizing sequences of words in the text which are associated with diseases. Then, each disease found is associated with its most likely ICD-10 code, a task better known as Entity Linking (Kraljevic et al., 2021; Wiegrefe et al., 2019). Nowadays, the most commonly used methods for solving the NER task are based on deep neural networks such as transformers-based models or recurrent neural networks, while a frequent technique for assigning codes is to use distance algorithms or search engines to compare the diseases found with the code descriptions of the terminology.

2.3 Commercial Systems

A handful of commercial products offer information extraction from clinical data, including automatic coding. These products usually are delivered as services and offered by leading cloud providers such as Amazon Web Services with Amazon Comprehend Medical¹, Google Cloud with

¹<https://aws.amazon.com/comprehend/medical/>

Google Cloud Healthcare Data Engine² and Microsoft Azure with Azure Cognitive Service for Language³. The problem with these services is that they do not offer automatic coding for languages other than English.

Data privacy concerns may arise from using this third-party software to extract patients' information. Some healthcare providers may prohibit sending data to systems outside the primary source due to potential cybersecurity issues.

3 Data and Methods

The Chilean Waiting List is characteristic of the the public healthcare system. This list arises due to the high demand for medical care and the limited capacity of the public health system to meet it. Entry on the waiting list begins when a patient goes to primary care or secondary care physician to treat pathology. The patient has two possible paths: if the pathology is included in the "Garantías Explícitas en Salud" (GES) program, the patient enters a process where his or her health problem is assured a maximum waiting time for medical attention. If the GES program does not cover the pathology, the referral is classified in one of these five options: New Specialty Consultations (CNE), Follow-up Consultations (CCE), Diagnostic Procedures (Proc), Surgical Intervention (IQ) and Complex Surgical Intervention (IQC). In any of these alternatives, the patient is placed on a waiting list and must wait a variable amount of time to receive medical attention from a specialist.

The Chilean Waiting List comprises 25,374,491 referrals, divided into five categories: 18,716,629 correspond to CNE type referrals, 4,391,257 to Proc type referrals, 2,222,545 to IQ type referrals, 39,266 to CCE type referrals, and finally, 4,794 to IQC type referrals. In particular, this work will focus on CNE-type referrals.

Within the Chilean Waiting database, 73 attributes are separated into two main types of sets. The first set corresponds to the attributes associated with the person (date of birth, sex, national identifier). In contrast, the second set corresponds to the administrative information associated with the referral given to the person (date of admission, date of discharge, the benefit provided, specialty, diagnostic suspicion, and diagnostic confirmation).

²<https://cloud.google.com/healthcare>

³<https://azure.microsoft.com/en-in/products/cognitive-services/language-service>

For the analysis of the diagnoses present in the referrals, two free-text attributes representing medical diagnoses are considered: diagnostic suspicion and diagnostic confirmation. Table 1 shows the frequency of referrals according to medical specialty, while Table 2 shows corpus statistics of the texts analyzed.

Specialty	Referrals	Relative Freq. (%)
Ophthalmology	3,352,203	17.91
Otorhinolaryngology	1,270,563	6.79
Traumatology	1,066,814	5.70
Gynecology	991,166	5.30
General Surgery	982,500	0.05
Dermatology	762,758	4.08
Internal Medicine	703,844	3.76
Endodontics	662,607	3.54
Removable Prosthesis	652,604	3.49
Urology	605,425	3.23

Table 1: Top 10 specialties with the highest presence on the Chilean Waiting List for a medical appointment.

We used 10,000 referrals from the historical Chilean Waiting List to train the NER module for disease recognition. As detailed in (Báez et al., 2020; Báez et al., 2022), these referrals were previously consolidated by a team of clinical experts, thus constituting the so-called Chilean Waiting list corpus. In addition, we performed rounds of evaluation of the NER performance, identifying diseases that the model could not identify. Thus, these diseases were incorporated as new examples of the model training process.

4 Proposed System

To code the narratives, we first used a NER model to automatically recognize sequences of words in the text associated with diseases. Then, each disease found is associated with its most likely ICD-10 code through a search engine. Figure 1 shows an overview of our proposed system.

4.1 NER Model

As shown in Figure 1, the input of our system is the referral written by the physician in an unstructured format. These texts are used as input for the automatic disease recognition model. In particular, this NER model is based on the work proposed in (Rojas et al., 2022a), where a simple but highly effective architecture for medical entity recogni-

tion is introduced. This model, named Multiple LSTM-CRF (MLC), is a deep neural network system composed of three main modules, emphasizing the impact of using domain-specific contextualized embeddings.

The first layer of the MLC approach, the “stacked embedding layer”, transforms the texts associated with the diagnoses into a vector representation using character-level contextual embeddings and static word embeddings, both trained in the clinical domain. Then, in the encoding layer, a recurrent neural network is used to obtain long-distance dependencies between words in the sentence, thus obtaining a better context to improve the previous layer’s representations. Finally, the classification layer assigns the most probable label to each word in the diagnosis using the CRF algorithm, identifying which parts of the text correspond to the beginning and end of a disease.

Regarding the experimental setup, the disease model was trained to 150 epochs using an SGD optimizer with mini-batches of size 32 and a learning rate of 0.1. As mentioned, to encode sentences, we used two types of representations; a 300-dimensional word embedding model trained on the Chilean Waiting List corpus⁴ and character-level contextualized embeddings retrieved from the Clinical Flair model (Rojas et al., 2022b). To implement the model and perform our experiments, we used the Flair framework, widely used by the NLP research community (Akbik et al., 2019).

4.2 Search Engine

The output of the NER step is a list containing all the diseases mentioned in the referral. This second module aims to assign an ICD-10 code to each disease found, which can be used later for clinical decisions or management. The assignment of the ICD-10 code is done through a search engine tool based on Elasticsearch⁵, an open-source search and analytics engine. This system can assign similarities between the mention of the disease and each of the codes of the ICD-10 tabular list.

Unlike the algorithms of distance comparison between words, this search engine has an index that contains each of the ICD-10 diseases represented through a series of synonymous sentences extracted from different sources of information, simulating

⁴<https://zenodo.org/record/3924799>

⁵Registered trademark of Elasticsearch B.V. Available at <https://www.elastic.co/elasticsearch/>

Specialty	Number of tokens (std)	Number of sentences (std)	Tokens per sentence
Infectology	28.59 (48.04)	1.50 (1.61)	18.94
Medical Oncology	20.09 (42.12)	1.15 (0.66)	17.40
Diabetology	19.03 (33.07)	1.33 (1.31)	14.22
Pediatric Rheumatology	15.05 (30.95)	1.19 (0.75)	12.61
Oral Pathology	14.54 (24.51)	1.17 (0.69)	12.34
Neonatology	12.92 (31.31)	1.10 (0.55)	11.64
Pediatric Hemato-Oncology	12.70 (26.74)	1.16 (0.75)	10.94
Pediatric Plastic Surgery	17.51 (16.33)	1.25 (0.51)	10.88
Pediatric Gynecology	13.01 (22.61)	1.22 (0.82)	10.61
Pediatric Traumatology	12.27 (18.39)	1.16 (0.74)	10.55

Table 2: Top 10 specialties with the highest number of tokens per sentence on average in Chilean Waiting List.




Figure 1: Overview of the proposed disease coding system.

TRANSLATION Clinical report: A 60-year-old patient presented with an incisional hernia with a history of dm 2 and morbid obesity. Entities: 1. incisional hernia, 2. dm 2, 3. morbid obesity.

in a better way the process followed by clinical experts to determine the code of a disease.

For example, in the index, the code “K02.2” contains the canonical code description “Caries of cementum” and multiple synonymous definitions, such as “Cement caries” and “Root caries”. This is important as disease mentions found in unstructured diagnoses are rarely equivalent to the exact definition.

The sources of information used for the extraction of synonymous disease definitions were as follows:

Tabular list of ICD-10 terminology: This is the

basis of the index, which tells us which codes we will assign to the disease mentions.

Alphabetical index of ICD-10 terminology:

The guide for the manual assignment of codes to diseases and was obtained using the “web scraping” technique from the website of the Spanish Ministry of Health ⁶.

IRIS dictionary: It maps natural language sentences to an ICD-10 code. This dictionary was built from the mortality coding rounds con-

⁶https://eciemaps.mscbs.gob.es/ecieMaps/browser/index_10_2008.html

ducted in the Chilean Department of Statistics and Health Information.

UMLS: Spanish definitions from multiple vocabularies were extracted from the metatresaurus database.

DEIS abbreviations: Manually constructed list of abbreviations and their expansions.

4.3 Experiments

In our experiments, we measure how well the predictions made by the model fit compared to the decisions made by clinical experts. In this way, a subset of the referrals described in Section 3 was selected to be manually coded by a team of two clinical coders. The manual annotation process and system validation steps are provided below.

4.3.1 Manual coding

The clinical experts carried out the annotation process using Excel software. For this purpose, a file containing a unique identifier for each referral, the associated diagnostic suspicion, and a blank column for the actual coding was provided to the coders. This way, the expert coders identified disease codes in 1,188 clinical narratives from the Chilean Waiting List for a new specialty.

It is important to mention that in this process, codes were identified at the referral level, not at the entity level; therefore, it is not possible to determine the performance of the NER model in this experiment. In future work, specialized software such as INCEPTION, could be used, as proposed in the work of (Báez et al., 2020). This software would make it possible to identify which parts of the text refer to diseases. On the other hand, only diseases were coded, but future research could extend it to new entity types, such as clinical procedures or clinical findings.

4.3.2 Metric

The Mean Average Precision (MAP) metric is used to evaluate the performance of our coding system. This metric is widely used in works that address the same automatic coding task. This metric is defined as follows:

$$AveP = \frac{\sum(P(k) \cdot rel(k))}{\text{number of relevant documents}}, \quad (1)$$

where $P(k)$ represents the precision at position k , and $rel(k)$ is an indicator function equal to 1 if the

element in rank k is a relevant document and 0 otherwise.

The MAP is computed using the Python implementation of the TREC evaluation tool, `trectools`, by (Palotti et al., 2019), where an adaptation was applied, in which the coded diagnoses have to be ordered based on a ranking, which for this work is considered the order in which the mention was found and subsequently the code was assigned.

5 Results

Orthodontics (3.28)	Obstetrics (1.01)
Endodontics (2.52)	Plastic Surgery (0.93)
Oral Rehab. (1.26)	Pediatrics (0.84)
Nephrology (1.26)	P. Physical Med. (0.84)
P. Dentistry (1.26)	Dental Operatory (0.34)
Psychiatry (1.18)	Maxillo. Radiol. (0.34)
P. Urology (1.01)	STI (0.084)

P. = Pediatric, Maxillo. Radiol. = Maxillofacial Radiology, STI = Sexually transmitted infections.

Table 3: Specialties with a perfect MAP score. The relative frequency (in percentage) of referrals in the dataset is in parentheses.

5.1 Coding Performance

The ICD-10 consists of a solitary coded catalog composed of categories with three characters, each of which can be additionally subdivided into as many as ten subcategories of four characters.

We computed the MAP metric over the test set at the category (e.g. K02) and subcategory (e.g. K02.2) levels. We achieved a MAP of 0.83 for the category and 0.63 for the subcategory level.

To underline the difficulty of achieving outstanding results in coding, we analyzed the results obtained only by clinical experts. The expert coders achieved an agreement MAP of 0.75 for subcategory and 0.83 for category level. Several reasons, such as the subjectivity in clinical judgment, the complexity of coding guidelines, the evolving nature of medicine, the time pressure and workload, personal bias, and lack of standardization, could explain the low agreement score.

6 Error Analysis

To better understand the errors made by our coding system, we performed a granular analysis of the scores obtained among the different specialties in the corpus. Tables 3 and 4 show the top 14 best and

Specialty	MAP at Category Level	Relative freq. in %
Neurology	0.68	2.10
Immunology	0.67	0.84
Geriatrics	0.60	0.93
Pediatric Gastroenterology	0.58	0.84
Cardiothoracic surgery	0.53	0.84
Radiation therapy	0.50	0.84
Pediatric Family Medicine	0.45	0.84
Hematology	0.42	1.18
Diabetology	0.35	1.01
Pediatric Traumatology	0.29	0.93

Table 4: Top 10 worst scores according to the specialties.

10 worst scores according to the specialties. We noted that in the top 14 best specialties the diagnostic suspicions registered in the referral were written straightforwardly and were specific diagnoses, such as “lipoma”, “caries”, and “nephrolithiasis”, avoiding other clinical information like comorbidity, medication intake, or some other medical history. Furthermore, it can be noted that half of these referrals are related to dental diagnosis.

On the other hand, the top 10 worst specialties share in common that most of the diagnoses are very unspecific, with the incorporation of non-medical information such as the patient’s phone number, patient’s address, physician’s name, the specialty the patient is referred to and information about comorbidity. Besides, several referrals are without a diagnosis but with the text “unspecific consultation” or “other”.

7 Model Deployment and Use Cases

Due to internal regulations, we could not send patients’ data to third-party systems such as cloud providers or academic supercomputing clusters (Ministerio Secretaría Regional de la Presidencia, 1999). For this reason, we deployed the whole coding system on-premise on a bare metal machine with a GPU compute module (NVIDIA RTX A4000⁷) to process the coding requests from the whole department efficiently.

The complete automatic coding system was deployed as a pair of microservices running inside containers to ease portability. One container hosts

⁷The compute module has 16 GB of GPU memory and 6.144 CUDA cores. More information at <https://www.nvidia.com/en-us/design-visualization/rtx-a4000/>

the NER module and exposes an API as a web service listening to disease-mention detection requests. The other container consists of the recommended implementation of the Elasticsearch software, which also exposes its API as a web service listening to mention-coding requests.

To code the waiting list and schedule recurrent coding when new data arrives, we used the KNIME⁸ software, a visual-programming data mining platform. We chose this software because of its ease of use for non-expert developers. The workflow starts with the raw waiting list, which is first passed through the NER module to detect disease mentions, and then each mention is sent to the coding module to assign the most relevant code.

The automatic coding result from the workflow mentioned above is persisted on a table inside a database that stores each disease mention for each referral along with the predicted code from the system.

8 Conclusions

In this work, we created a nationwide system to improve the management of the Chilean public healthcare system. Specifically, we addressed the challenge of creating an automated system to code the diseases present in the Chilean Waiting List referrals. We developed and validated a model based on two steps: a NER model to recognize disease mentions and a search engine based on Elasticsearch to assign the codes to each disease. This mapping system was enriched with several terminology resources used in real life by manual coders to assign codes, thus partially simulating

⁸Registered trademark of KNIME GmbH. Available at <https://www.knime.com/>

the pipeline followed by these professionals when solving this task.

The system allowed us to assign codes to 18,716,629 referrals, thus demonstrating its efficiency and effectiveness. The performance obtained in our experiments was 0.83 according to the MAP score, which is close to the most advanced systems currently in the coding task. The model was deployed into production in the Department of Health Statistics and Information Systems of the Ministry of Health of Chile.

The use of this system could be an important support for the management of waiting lists. In addition, since 75% of the Chilean population is in the public healthcare system, the analysis of the new specialty consultations can be used for epidemiological studies, such as the one done on the incidence of psoriasis (Lecaros et al., 2021).

Acknowledgements

This work was funded by ANID Chile: Basal Funds for Center of Excellence FB210005 (CMM); Millennium Science Initiative Program ICN17_002 (IMFD) and ICN2021_004 (iHealth), Fondecyt grant 11201250, and National Doctoral Scholarship 21220200. We also acknowledge Daily Piedra and Marcela Carmona for their work on annotating and coding the test dataset.

References

- Alan Akbik, Tanja Bergmann, Duncan Blythe, Kashif Rasul, Stefan Schweter, and Roland Vollgraf. 2019. [FLAIR: An easy-to-use framework for state-of-the-art NLP](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 54–59, Minneapolis, Minnesota. Association for Computational Linguistics.
- Pablo Báez, Antonia Paz Arancibia, Matías Ignacio Chaparro, Tomás Bucarey, Fredy Núñez, and Jocelyn Dunstan. 2022. Procesamiento de lenguaje natural para texto clínico en español: el caso de las listas de espera en Chile. *Revista Médica Clínica Las Condes*, 33(6):576–582.
- Pablo Báez, Felipe Bravo-Marquez, Jocelyn Dunstan, Matías Rojas, and Fabián Villena. 2022. [Automatic extraction of nested entities in clinical referrals in Spanish](#). *ACM Trans. Comput. Healthcare*, 3(3).
- Pablo Báez, Fabián Villena, Matías Rojas, Manuel Durán, and Jocelyn Dunstan. 2020. [The Chilean waiting list corpus: a new resource for clinical named entity recognition in Spanish](#). In *Proceedings of the 3rd Clinical Natural Language Processing Workshop*, pages 291–300, Online. Association for Computational Linguistics.
- Jose Barros, Matías Rojas, Jocelyn Dunstan, and Andres Abeliuk. 2022. [Divide and conquer: An extreme multi-label classification approach for coding diseases and procedures in Spanish](#). In *Proceedings of the 13th International Workshop on Health Text Mining and Information Analysis (LOUHI)*, pages 138–147.
- Pengfei Cao, Yubo Chen, Kang Liu, Jun Zhao, Shengping Liu, and Weifeng Chong. 2020. [HyperCore: Hyperbolic and co-graph representation for automatic ICD coding](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3105–3114, Online. Association for Computational Linguistics.
- Hercules Dalianis. 2018. *Clinical Text Mining*, first edition. Springer International Publishing.
- Hang Dong, Matúš Falis, William Whiteley, Beatrice Alex, Joshua Matterson, Shaoxiong Ji, Jiaoyan Chen, and Honghan Wu. 2022. [Automated clinical coding: what, why, and where we are?](#) *npj Digital Medicine*, 5(1).
- Richárd Farkas and György Szarvas. 2008. [Automatic construction of rule-based ICD-9-CM coding systems](#). *BMC Bioinformatics*, 9(S3).
- Sarvnaz Karimi, Xiang Dai, Hamed Hassanzadeh, and Anthony Nguyen. 2017. [Automatic diagnosis coding of radiology reports: A comparison of deep learning and conventional classification methods](#). In *BioNLP 2017*, pages 328–332, Vancouver, Canada., Association for Computational Linguistics.
- Hyoun-Joong Kong. 2019. [Managing unstructured big data in healthcare system](#). *Healthcare Informatics Research*, 25:1.
- Zeljko Kraljevic, Thomas Searle, Anthony Shek, Lukasz Roguski, Kawsar Noor, Daniel Bean, Aurelie Mascio, Leilei Zhu, Amos A. Folarin, Angus Roberts, Rebecca Bendayan, Mark P. Richardson, Robert Stewart, Anoop D. Shah, Wai Keong Wong, Zina Ibrahim, James T. Teo, and Richard J.B. Dobson. 2021. [Multi-domain clinical natural language processing with medcat: The medical concept annotation toolkit](#). *Artificial Intelligence in Medicine*, 117:102083.
- C Lecaros, J Dunstan, F Villena, DM Ashcroft, R Parisi, CEM Griffiths, S Härtel, JT Maul, and C De la Cruz. 2021. The incidence of psoriasis in Chile: an analysis of the national waiting list repository. *Clinical and Experimental Dermatology*, 46(7):1262–1269.
- Ministerio de Salud de Chile. 2011. [Norma técnica para el registro de las listas de espera](#).
- Ministerio Secretaría Regional de la Presidencia. 1999. [Ley 19628 sobre protección de la vida privada](#).

- James Mullenbach, Sarah Wiegrefe, Jon Duke, Jimeng Sun, and Jacob Eisenstein. 2018. [Explainable prediction of medical codes from clinical text](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1101–1111, New Orleans, Louisiana. Association for Computational Linguistics.
- João Palotti, Harris Scells, and Guido Zuccon. 2019. [Trectools: An open-source python library for information retrieval practitioners involved in trec-like campaigns](#). In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR'19*, page 1325–1328, New York, NY, USA. Association for Computing Machinery.
- Matias Rojas, Felipe Bravo-Marquez, and Jocelyn Dunstan. 2022a. [Simple yet powerful: An overlooked architecture for nested named entity recognition](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 2108–2117, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Matías Rojas, Jocelyn Dunstan, and Fabián Villena. 2022b. [Clinical flair: A pre-trained language model for Spanish clinical natural language processing](#). In *Proceedings of the 4th Clinical Natural Language Processing Workshop*, pages 87–92, Seattle, WA. Association for Computational Linguistics.
- Mary Stanfill, Margaret Williams, Susan Fenton, Robert Jenders, and William Hersh. 2010. [A systematic literature review of automated clinical coding and classification systems](#). *Journal of the American Medical Informatics Association : JAMIA*, 17:646–51.
- Fei Teng, Yiming Liu, Tianrui Li, Yi Zhang, Shuangqing Li, and Yue Zhao. 2022. [A review on deep neural networks for icd coding](#). *IEEE Transactions on Knowledge and Data Engineering*, pages 1–1.
- Sarah Wiegrefe, Edward Choi, Sherry Yan, Jimeng Sun, and Jacob Eisenstein. 2019. [Clinical concept extraction for document-level coding](#). In *Proceedings of the 18th BioNLP Workshop and Shared Task*, pages 261–272, Florence, Italy. Association for Computational Linguistics.
- World Health Organization. 2015. *International statistical classification of diseases and related health problems*, 10th revision, fifth edition, 2016 edition. World Health Organization.
- World Health Organization. 2023. *International statistical classification of diseases and related health problems*, 11th revision edition. World Health Organization.
- Chenwei Yan, Xiangling Fu, Xien Liu, Yuanqiu Zhang, Yue Gao, Ji Wu, and Qiang Li. 2022. [A survey of automated international classification of diseases coding: development, challenges, and applications](#). *Intelligent Medicine*, 2(3):161–173.
- Ying Yu, Min Li, Liangliang Liu, Zhihui Fei, Fang-Xiang Wu, and Jianxin Wang. 2019. [Automatic icd code assignment of chinese clinical notes based on multilayer attention birnn](#). *Journal of Biomedical Informatics*, 91:103114.
- Lingling Zhou, Cheng Cheng, Dong Ou, and Hao Huang. 2020. [Construction of a semi-automatic icd-10 coding system](#). *BMC Medical Informatics and Decision Making*, 20.