Enhancing Deep Learning Models with Automated Knowledge Graphs for Improved Classification Performance and Explainability

By

Mutahira Khalid

00000330108

MS IT-2K20

Supervisor

Dr. Hasan Ali Khattak

Department of Computing

A thesis submitted in partial fulfillment of the requirements for the degree of Masters of Science in Information Technology (MS IT)

In

School of Electrical Engineering & Computer Science (SEECS) ,

National University of Sciences and Technology (NUST),

Islamabad, Pakistan.

(Feburary 2023)

THESIS ACCEPTANCE CERTIFICATE

Certified that final copy of MS/MPhil thesis entitled "Enhancing Deep Learning Models with Automated Knowledge Graphs for Improved Classification Performance and Explainability" written by Mutahira Khalid, (Registration No 00000330108), of SEECS has been vetted by the undersigned, found complete in all respects as per NUST Statutes/Regulations, is free of plagiarism, errors and mistakes and is accepted as partial fulfillment for award of MS/M Phil degree. It is further certified that necessary amendments as pointed out by GEC members of the scholar have also been incorporated in the said thesis.

Signature: ____________________

Name of Advisor: Prof. Hasan Ali Khattak

Date: 16-Feb-2023

HoD/Associate Dean: ____________________

Date: ____________________

Signature (Dean/Principal): ____________________

Date: ____________________

Approval

It is certified that the contents and form of the thesis entitled "Enhancing Deep Learning Models with Automated Knowledge Graphs for Improved Classification Performance and Explainability" submitted by Mutahira Khalid have been found satisfactory for the requirement of the degree

Advisor : Prof. Hasan Ali Khattak

Signature: ______________________

Date: 16-Feb-2023

Committee Member 1:Dr. Mehdi Hussain

Signature:

16-Feb-2023

Committee Member 2:Dr Mehvish Rashid

Signature: ______________________

Date: 16-Feb-2023

Signature: ______________________

Date: ______________________

Dedication

To my devoted parents and brother, who supported me throughout my life and to all the deserving children who do not have access to quality education especially young girls.

Certificate of Originality

I hereby declare that this submission titled "Enhancing Deep Learning Models with Automated Knowledge Graphs for Improved Classification Performance and Explainability" is my own work. To the best of my knowledge it contains no materials previously published or written by another person, nor material which to a substantial extent has been accepted for the award of any degree or diploma at NUST SEECS or at any other educational institute, except where due acknowledgement has been made in the thesis. Any contribution made to the research by others, with whom I have worked at NUST SEECS or elsewhere, is explicitly acknowledged in the thesis. I also declare that the intellectual content of this thesis is the product of my own work, except for the assistance from others in the project's design and conception or in style, presentation and linguistics, which has been acknowledged. I also verified the originality of contents through plagiarism software.

Student Name:Mutahira Khalid

Student Signature: 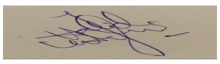

Acknowledgments

Glory be to Allah (S.W.A), the Creator, the Sustainer of the Universe. Who only has the power to honour whom He please, and to abase whom He please. Verily no one can do anything without His will. From the day, I came to NUST till the day of my departure, He was the only one Who blessed me and opened ways for me, and showed me the path of success. Their is nothing which can payback for His bounties throughout my research period to complete it successfully.

Mutahira Khalid

Contents

List of Figures

List of Tables

Abstract

Medical coding works by assigning standardized medical codes to clinical records' diagnoses, prognoses, and prescriptions. These codes are necessary for accurate medical billing and claims processing, both of which are vital for sustaining effective revenue cycles. Computer Assisted Coding (CAC) automates the process of assigning medical codes, with the aid of the Artificial Intelligence (AI) model. Despite the extraordinary results, there are certain limitations. These AI models rely on training data and collapse because they lack domain-specific knowledge, which results in false-positive predictions or just no predictions at all. Apart from this, the users' ability to trust these AI applications is also hampered by the black-box nature of deep learning models. Even the explainable attention mechanism is unable to explain its certain predictions. These limitations can be addressed with the consolidation of Symbolic AI with deep learning leading to explainable and trustable predictions with an overall increase in accuracy. The hybrid AI approach has a number of benefits, but creating knowledge graphs—the brain behind symbolic AI—is a laborious process. Thus, I have automated the construction of knowledge graphs using a few processes that include Data preprocessing, ontology mapping, concept enrichment, and Neo4j knowledge graph creation. Additionally, I have suggested two distinct NeuroSymbolic AI approaches to get around some of the deep learning's drawbacks. The first approach "Domain-specific knowledge infusion" enriches the medical terms leading to an overall increase in classification accuracy of nearly 81%. The second approach of "Explainable Deep Learning Predictions" explains the attention mechanism results by visualizing the word-to-word and word-to-code level connections with an accuracy of 64% and 53%. This research is novel as the knowledge graph creation in few and easy steps has not been done before. Additionally, it is the earliest study on knowledge graphs for explainability and domain-specific knowledge infusion to medical coding.

Chapter 1

Introduction

Medical Coding is the process that aids in billing and claim generation. Diagnosis, prognosis, prescriptions, and medical notes are all converted into alphanumeric codes through medical coding. This coding approach safeguards patient data while liberating insurance companies, hospitals, and governmental organisations from struggling to analyse lengthy discharge summaries [26]. These codes are from the Healthcare Standard Procedure Coding System (HSPCS), Current Procedural Therapy (CPT), or the International Classification of Diseases (ICD) [8].

Medical code assignment is a labour-intensive, time-consuming, and manual task. It requires the expertise of medical coders to assign standardized medical codes. Also, erroneous and flawed codes are a significant cause of lost revenue. An estimate states that four out of every five codes are incorrect [16], which is a huge loss for all parties involved [25]. Researchers have been particularly interested in Automated Medical Coding (AMC), also known as Computer-Assisted Coding (CAC), in which the human assignment of codes has been made automatic with the use of Artificial Intelligence (AI) and its subfields Deep Learning (DL) and Machine Learning (ML) [26]. These AI techniques commence the production of codes without human intervention after training models using manually annotated data.

Due to their superior and reliable performance and pattern recognition nature, Deep Learning models are the focus of interest for these CAC problems [35]. Hierarchical Attention Networks (HAN) [14], Convolutional Neural Networks (CNN), Recurrent Neural Networks (RNN), Long Short-Term Memory (LSTM) [2], Bidirectional Encoder Representations from Transformer (BERT) [18], etc [48]. are some models that have performed incredibly well. Discharge summaries are used as input by these trained models, producing medical codes as suggestions for medical coders. Even if these codes are incorrect, they can simplify the process of assigning codes by pointing to the appropriate blocks and portions of the coding structure [17].

CAC models perform precisely and accurately. However, there are several restrictions that deep learning cannot overcome. These problems stem from DL models' reliance on training data. Automated Medical Coding is a multi-label classification problem, and the models it uses are trained for specific medical conditions such as diabetes, COVID, hypertension, stroke, etc. These models don't work or give out misleading positive findings if any notion is from another domain.

The Deep Learning Black-box models' inability to explain predictions is another obstacle [46]. The Deep Learning models are complex and challenging to understand. The inclusion of an attention mechanism is a recent innovation. This mechanism highlights the keywords during decoding that were saved in the encoding process. The relationship between these elevated words and the outcomes is still unclear. As a result, the proposed explainable system cannot be explained. Models' "black-box" nature makes it difficult to explain how they operate within. By combining two distinct artificial intelligence approaches—one pattern-based and the other rule-based—these limitations can be overcome. The combination of Symbolic AI and Deep Learning may be used to improve classification accuracy and produce results that are comprehensible and reliable.

Symbolic AI was the first wave of AI where humans' nature of learning symbols and eventually making decisions on the rules was utilized. Knowledge Graphs (KG) are called the brains behind symbolic AI which contains hundreds to thousands of triplets. The nature of Knowledge Graphs being self-explainable can be used for the creation of applications which will be accurate and will be explainable [57].

1.0.1 Motivation

These KG are beneficial for many purposes but their creation is an extremely difficult and burdensome task which requires the human intervention of both medical and computer experts. It is highly time-consuming and massive efforts are needed. I had thus automated the cumbersome task of KG creation in four steps that are: Data Preprocessing, ontology-based information retravel, concept enrichment, and Neo4j KG creation. I had also purposed two different Neuro-Symbolic approaches by utilizing our Automated KG creation approach, one for the sake of increasing the overall accuracy of the multi-label classification and the second for the sake of explainability. I employed Bioportal [5] expert-created ontologies and Mimic-III data [13] for our proposed methodologies.

Our proposed methodologies for KG creation and Neuro-Symbolic AI will pave the paths towards the new era of AI, where the built applications will not only be accurate but also explainable in nature. This will lead to the third wave of AI also known as Neuro-Symbolic AI or eXplainable AI (xAI).

1.1 Problem Statement and Contribution

1.1.1 Problem Statement

- Medical coding is a manual process that has been automated in recent years, although few models have been trained for the 10th version of the ICD, which is now in use. ICD-10 codes are 13 times more numerous than ICD-9 codes, which reduces the precision of AI models developed to solve CAC problems.
- The CAC model's reliance on training data results in false positives or simply no results. The scarcity of domain-specific knowledge reduces the accuracy of AMC models.
- The black-box nature of Deep Learning models is causing AI applications to be rejected. This aspect makes it difficult for consumers to develop trust in these applications.
- The proposed Attention Mechanism does not explain the relationship and connection between highlighted words and predicted codes.
- Symbolic AI could be a solution to the challenges described above, but its construction is highly difficult. Both computer and medical knowledge are required. Manually creating the hub of hundreds to thousands of triplets 'Knowledge Graphs' can take months to years.

1.1.2 Research Contributions and Novelties

- A pretrained HLAN model is fine-tuned and upgraded from ICD-9 to ICD-10 to cover a broader problem domain.
- Automating the process of Knowledge Graph Creation by processing the discharge summaries and mapping the concepts to expert-created BioPortal Ontologies.
- Automating the creation of queries for knowledge graph creation and visualisation.
- In order to get over some of the constraints of the multi-label classification, the domain-specific knowledge is combined with the pretrained CAC model. Paper presented at IBCAST Conference [58].
- Using knowledge graphs with an attention mechanism to explain how highlighted phrases and anticipated codes are connected.
- Simplifying knowledge graphs for visualization of attention mechanism results by providing word-to-word and word-to-code level connections.

1.1.3 Limitations

- The accuracy of the fine-tuned model with a larger problem domain is lower than that of the pretrained model. This was entirely predictable given that I increased the label count from 50 to 550 medical codes.
- The knowledge graphs that were created were limited to particular types of nodes and relationships. The semantic information supplied by ontologies was used for this. The solution could be the integration of several ontologies.
- The model predictions completely determine the results of attention-based explainability. The model's comparable low accuracy will undoubtedly influence the word-to-word and word-to-code level explanations.

1.1.4 Thesis Layout

Seven chapters make up the thesis. Chapter 2 discusses the background and introduces the tools and techniques used in our suggested techniques right after the introduction. The pertinent literature, the gaps in the prior study, and the solutions are all covered in Chapter 3. In order to help the reader comprehend our own suggested and developed ways, Chapter 4 described them in detail using graphics. The outcomes of our experiments are shown in Chapter 5, and the various charts demonstrate how effective our methods were. The topic of discussion in Chapter 6 is how novelties contribute to the advancement of various fields. The report is concluded in Chapter 7 with a few ideas and directions for the future.

Chapter 2

Domain Concepts

The thesis introduces the readers to a variety of new concepts and makes use of a wide range of tools to further research. The terms and tools of computer science are introduced in this chapter along with some background information. The descriptions of many technologies and relevant subjects are provided here.

2.1 Electronic Health Records (EHR)

A digitally stored systematized collection of patient and population health information is known as an electronic health record (EHR). The sharing of these records is possible between various healthcare facilities. EHRs may contain a variety of data, including billing information, demographics, medical history, prescription and allergy information, immunization records, laboratory test results, radiological pictures, and vital signs. Figure 2.1 [36] shows a glimpse of EHR.

Electronic health records (EHRs) have been hailed as essential to raising the standard of care for many years. Additional to patient charting, electronic health records are utilized for other purposes. Today, physicians leverage data from patient records to enhance quality outcomes through their care management programmes. EHR compiles all patient demographic data into one big pool and uses it to help develop "novel treatments or innovation in healthcare delivery," which ultimately advances healthcare aims. Clinicians have been able to recognise and categorise individuals with chronic illnesses by combining several forms of clinical data from the system's health records. By using the data and analytics to stop hospitalizations among high-risk patients, EHR can enhance the quality

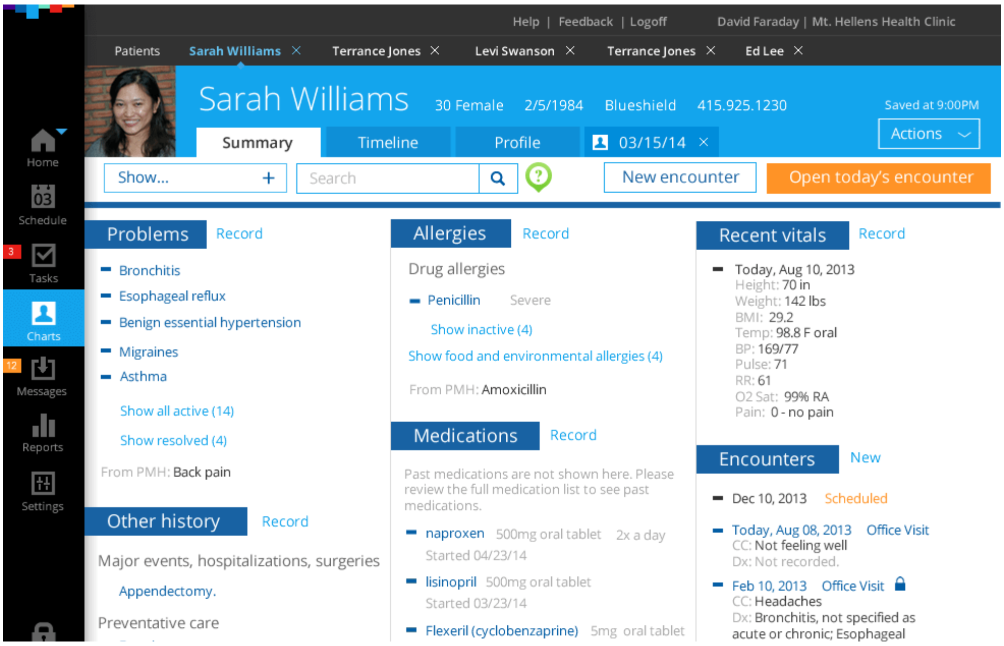

Figure 2.1: Electronic Health Record Example

of care.

EHR systems are made to accurately store data and record a patient's state over time. It makes it unnecessary to find a patient's earlier paper medical records, and it helps to guarantee that the data is current, accurate, and readable. Along with "privacy and security," it also enables open dialogue between the patient and the physician. As there is only one changeable file, there is less chance of data replication, which lowers the possibility of lost documents, increases cost-effectiveness, and increases the likelihood that the file is current [69].

2.1.1 Clinical Notes

EHRs include both structured and unstructured data. Clinical notes are a common term used to describe the written portion of medical reports. The doctors and the rest of the healthcare team can use them to determine the patient's condition and treatment options, making them the most crucial bits of information for medical professionals [31]. They serve as legal documentation for insurance and billing as well as a communication tool for medical professionals.

2.2 Medical Coding

The term "Medical Coding" was first used in England in the 17th century and has gained enormous popularity ever since. As the information on illnesses, disorders, and afflictions was gathered and arranged into numerical codes to determine the causes and reasons for deaths, it was known as the "London Bills of Mortality." Dr William Farr suggested a unified classification system in the late 1830s because bills and their terminology were inconsistent. By the 1930s, this approach had developed into the International List of Causes of Death. The World Health Organization (WHO) used this list as the medical profession grew to measure mortality rates and for other reasons. The International Classification of Diseases expanded this same list (ICD). In 1977, the ICD codes underwent further development to include not only causes of death but also illnesses and injuries [11].

Diagnoses, medical practices, services, and equipment data are transformed into alphanumeric codes by medical coding [8]. Medical professionals wrote the clinical notes that include these treatments, services, and diagnoses. The appropriate businesses and organisations use these codes to generate and produce claims and bills. Insurance firms and governmental organisations are spared from having to read lengthy discharge reports. Additionally, the patient's private information is kept on file. Disease outbreaks, research, mortality forecasting, ICU capacity forecasting, etc. all make use of medical coding. A medical coder, clinical coder, diagnostic coder, or nosologist is the individual who converts these clinical notes into codes [61].

Diagnostic and procedural coding systems are the two categories into which medical coding is divided. The International Classification of Diseases (ICD), Current Procedural Therapy (CPT), and Healthcare Standard Procedure Coding System are the medical codes used for billing purposes. Every year, these billing codes are changed and updated to reflect the circumstances. These are described in further detail below:

2.2.1 Diagnostic Codes

Diagnostic codes are used in medical classification to translate patient ailments, symptoms, illnesses, diseases, drug effects, etc. into alphanumeric codes. ICD is the most popular diagnostic coding system, and these codes are special types. These patients encounter coding systems concentrate on particular patient encounters such as accidents, emergencies, outpatient, inpatient, surgical care, mental health, etc. It is currently in its 12th version after being updated annually.

For specific diagnoses, the ICD coding system has chapters and blocks. Each character plays a significant role in the creation of these codes, which follow specific criteria. For instance, the first character of the 10th edition of ICD codes is always a number, there must be a dot after the first three characters, and no code may be less than three characters. The chapter and block details for the tenth version of the ICD are displayed in Figure 2.2 [71]. Compared to procedural codes, the number of diagnostic codes is enormous.

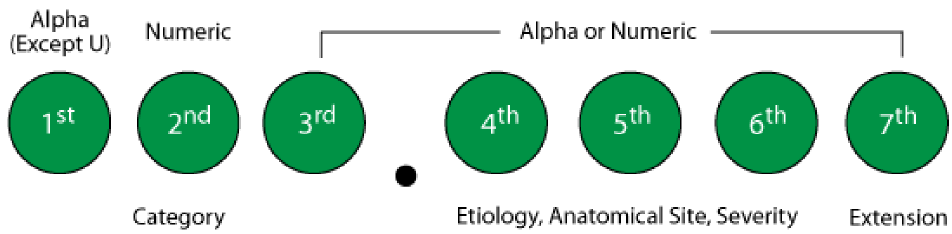

Figure 2.2: ICD-10 CM structure

2.2.2 Procedural Codes

Unlike diagnostic codes, prescription codes include details about the medical care given to specific patients. The procedure could involve a surgical procedure, medicine, lab tests, etc. CPT, ICD-10-PCS, and HCPCS are the coding systems used for procedure coding. They are seven-character alpha-numeric codes with rules, just like diagnostic codes. Procedure codes contain detailed information about the procedure being performed in each character.

The first character in Figure 2.3 [22], which depicts the ICD-10-PCS structure, stands in for a section, the second for a bodily system, the third for a root action, and so on. There are three groups of CPT codes. The most often used codes are found in category 1, performance tracking and management codes are found in category 2, and experimental or emergency procedures are found in category 3. There are around 10,000 CPT codes.

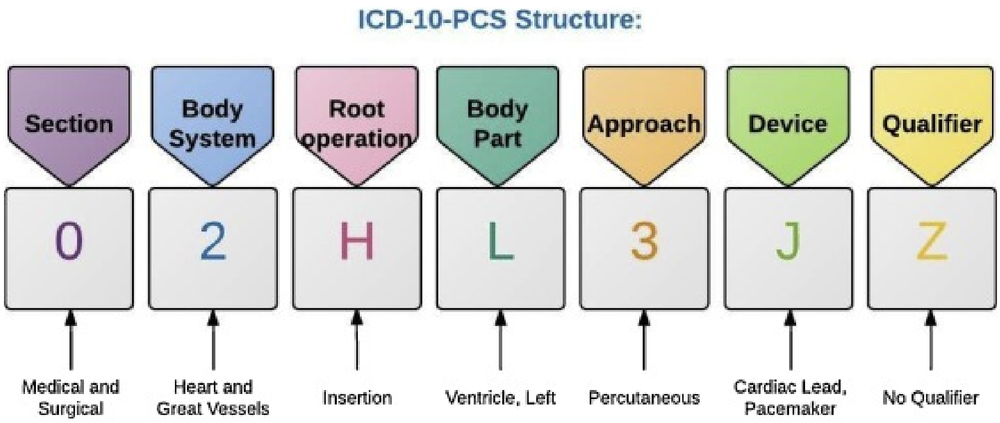

Figure 2.3: ICD-10 PCS structure

2.3 Computer-Assisted-Coding/ Automated Medical Coding

Medical coders with degrees and expertise in this field perform medical coding. Machine learning and artificial intelligence as a whole have attracted a tonne of interest lately. The process of medical coding known as Computer-Assisted-Coding (CAC) or Automated Medical Coding (AMC) has been automated thanks to the usage of AI models [54]. These systems have a user interface that takes input and transforms these summaries into codes.

Hierarchical Attention Networks (HAN) [14], Convolutional Neural Networks (CNN), Recurrent Neural Networks (RNN) [3], Gated Recurrent Units (GRU), Long Short-Term Memory (LSTM) [7], and others are some of the popular models for CAC. These models are used for multi-label classification, which involves assigning various labels or medical codes to a single clinical note. These systems are very helpful for medical coders because they can accept or reject model-predicted recommendations.

2.4 Hugging Face

An American business company called Hugging Face [12] develops tools for programmers who want to use machine learning in their applications. It is renowned for the transformer library it provides for applications involving Natural Language Processing (NLP). Hugging Face Hub is a platform that allows users to share datasets, pretrained models, demos of developed applications, etc. It is used to share code and resources for research and development and was inspired by Github. Hugging Face currently offers thousands of freely downloadable datasets and models.

2.5 BioPortal

Bioportal [5] is one of the largest repositories of freely available ontologies. The BioPortal ontologies include a variety of features, including suggestions based on dataset descriptions, word searching across various ontologies, annotating medical concepts with ontology terms, and more. The hub of more than a thousand ontologies, BioPortal, has 79,636,946 mappings and 14,427,459 classes. The ontologies are available in a variety of formats, including Comma Separated Files (CSV), Web Ontology Language (OWL), Resource Description Framework (RDF), and Extensible Markup Language (XML).

2.6 Artificial Intelligence

Since its invention in the 1950s, artificial intelligence has been one of the most intriguing and widely used technologies [52]. John McCarthy first used the term "artificial intelligence" at the first AI conference in 1956. Since then, there have been many improvements. After learning from the data, the AI model attempts to imitate human intelligence and predict outcomes. Different kinds of AI have emerged over time [1]. Below are some details.

2.6.1 Symbolic Artificial Intelligence

From the mid-1950s to the mid-1990s, symbolic AI was a popular field [70]. It is also referred to as the first wave of AI. The main principle of symbolic AI was that humans begin to learn symbols in early childhood and eventually begin to use them in decision-making. Why can't a machine be taught to think if humans think in symbols and machines operate in symbols? Symbols can be anything that humans in this world can understand. It might represent an operator in mathematics, a sign, a colour, an emotion, etc. Symbolic AI builds logic and facilitates some decision-making by using symbols and their relationships with other symbols. Triplets or facts are other names for the group of symbols and relationships. The "brain" of symbolic AI is also referred to as knowledge graphs or knowledge bases which are a hub of hundreds to thousands of triplets.

Knowledge Graphs and Ontologies

A unique kind of graphs known as knowledge graphs are directed, labelled, and heterogeneous multigraphs [49]. In knowledge graphs, a triplet is a grouping of two nodes and one relation. The triplets' hub contains related information about a certain field. A few examples of the domains are healthcare, finance, law, and education. Formally, they are employed to depict the semantics of a few domain ideas. The semantics of COVID-19 are displayed in the knowledge graph as shown in Figure 2.4. It displays the signs, variations, and traits of COVID. What are the symptoms of COVID-19, for example? can be answered using these knowledge graphs or knowledge bases [28].

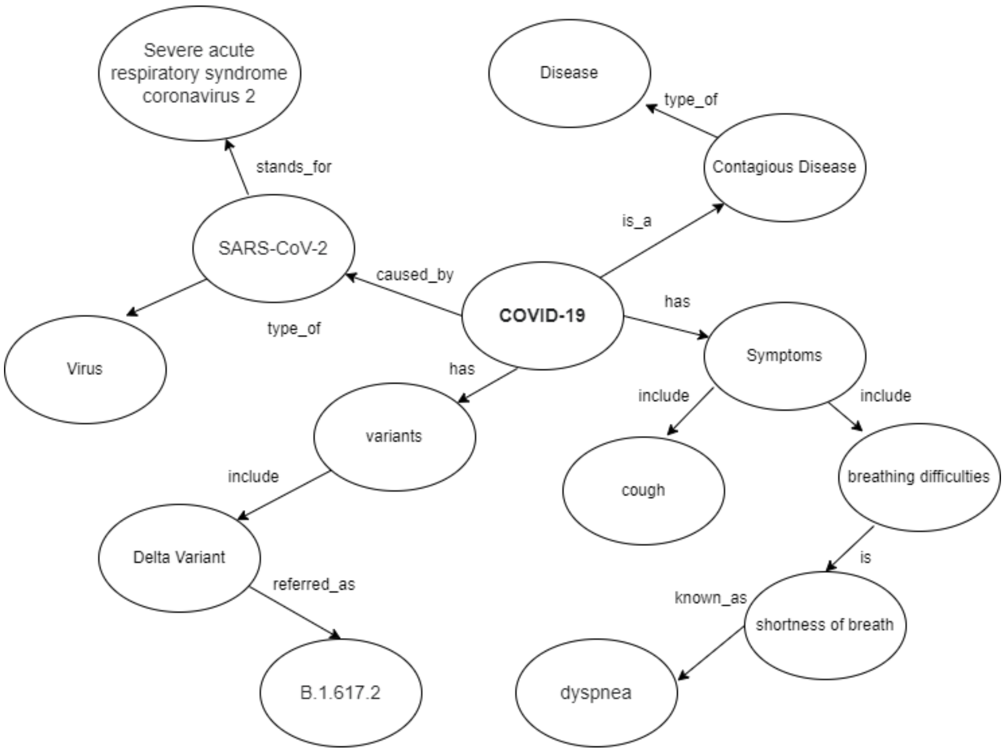

Figure 2.4: COVID-19 Knowledge Graph. This Knowledge base can be used to answer certain questions such as: What are the symptoms of COVID-19?

Ontologies are generalised data models as opposed to knowledge graphs. It implies that they are aware of classes, how they relate to one another and their characteristics. They don't include any precise information about words, ideas, or objects. For instance, the information "Casper is Roger's pet dog. He barks whenever he encounters an unfamiliar person. He enjoys eating meat". This information can be shown in a knowledge graph, however, for ontology, it will be "Individual has a pet dog. Dog will bark when it sees an unfamiliar person. Dog enjoys eating meat". These ontologies are useful resources for building a particular knowledge graph.

2.6.2 Deep Learning

Lapa and Alexey Ivakhnenko began their study towards creating the learning algorithm in 1967. A functional supervised, deep, multilayer preceptrons was published by them. Kunihiko Fukushima introduced other comparable designs for computer vision in 1980. Rina Dechter initially introduced the phrase "Deep Learning" to the machine learning field in 1986. In 1989, Yann LeCun applied a standard backpropagation technique that required three training days. Similar to this, Wei Zhang independently used backpropagation in 1988 on a convolutional neural network for alphabets recognition. In the early 1990s, similar studies on breast cancer diagnosis and medical image segmentation were conducted [72].

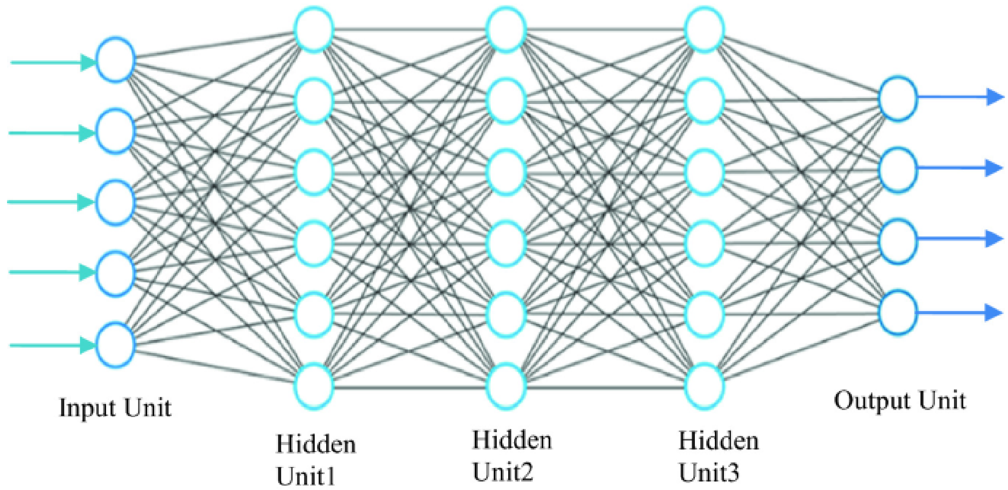

Figure 2.5: Deep Neural Network

Today's most popular technology is deep learning. Compared to the conventional symbolic AI technique, it has a number of benefits. Instead of developing rules, it uncovers patterns from data pixel by pixel and bit by bit. It gives the machine the intelligence to forecast outcomes for unobserved data, something symbolic AI cannot do. It has a variety of layers, including input, hidden, and output layers. Numerous neurons are found in each of these levels. A deep neural network, is made up of several hidden layers. The deep neural network's structure is depicted in Figure 2.5 [44], where the circles represent neurons and the lines represent the connections between them. Each layer has a unique activation function that is utilised to carry out calculations and make subsequent judgements. The sigmoid, RELU, Tanh, and other activation functions could be used.

Attention Mechanism

One of the most ground-breaking ideas in recent years is the Attention Mechanism. It functions according to how people pay attention to different stimuli and objects. For instance, we observe dark hues and large text in advertising, as well as the title on books or newspaper covers, etc. The attention Mechanism was proposed as an improvement to the traditional encoder decoder of machine translation in neural networks. Although it was designed for NLP tasks, its variant also functions for speech processing and computer vision applications.

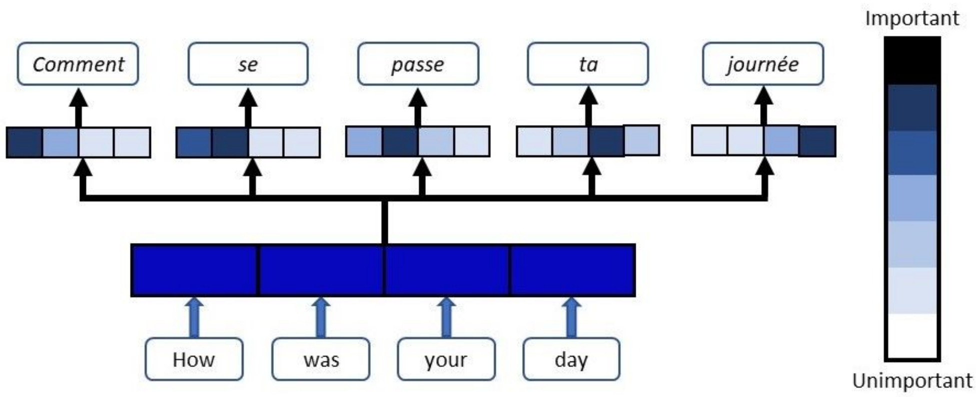

Figure 2.6: Attention Mechanism. The machine translation highlights some words depicting their importance for certain predictions.

The encoder processes the input text and encodes it into a context vector in a conventional encoder-decoder manner. Encoder sent the desired accurate summary to the decoder, which transforms the input into output. Therein lies the problem; if the encoder's summarization is subpar, the outcomes will suffer. When the input sentences are lengthy, this issue gets much worse. Additionally, recurrent neural networks are unable to retain so lengthy information. Accordingly, performance declines with lengthier inputs [10]. Although an LSTM is intended to better capture the long-range dependency than an RNN, it can occasionally overlook certain details. Another issue is that, when translating the text, there is no method to assign some input words greater weight than others.

Bahdanau [9] proposed a straightforward yet elegant idea, saying that each input word should be given a certain amount of relative value in addition to being taken into consideration by the context vector. Therefore, the suggested model looks for a set of points in the encoder's hidden states where the most pertinent information is present every time it generates a phrase. This concept is known as "Attention". Figure 2.6 shows the abstract working of the Attention Mechanism.

Hierarchical Attention Network

The Hierarchical Attention Network (HAN) functions in a manner comparable to how a human understands information, particularly documents. Even if we don't recognize many of the words in a sentence, we begin to read the words and form sentences. The same applies to paragraphs before moving on to documents.

As a result, the HAN operates under the principle that "Words make sentences, and sentences make documents." Its structure includes an attention mechanism and a hierarchy. The premise is that not all words in a sentence are equally essential, so attention is paid to both the word and sentence levels.

A hierarchical attention network is made up of bidirectional RNNs that take into account the context of each sentence or document before applying the appropriate attention, which highlights specific words and sentences to show how they relate to the prediction. The concept of classifying documents or texts was introduced by Zichao Yang in a study titled "Hierarchical Attention Networks for Document Classification" [14] published in 2016. Given that explainability was included, this concept represented a significant advancement in the classification fields. Figure 2.7 [21] depicts how HAN operates internally.

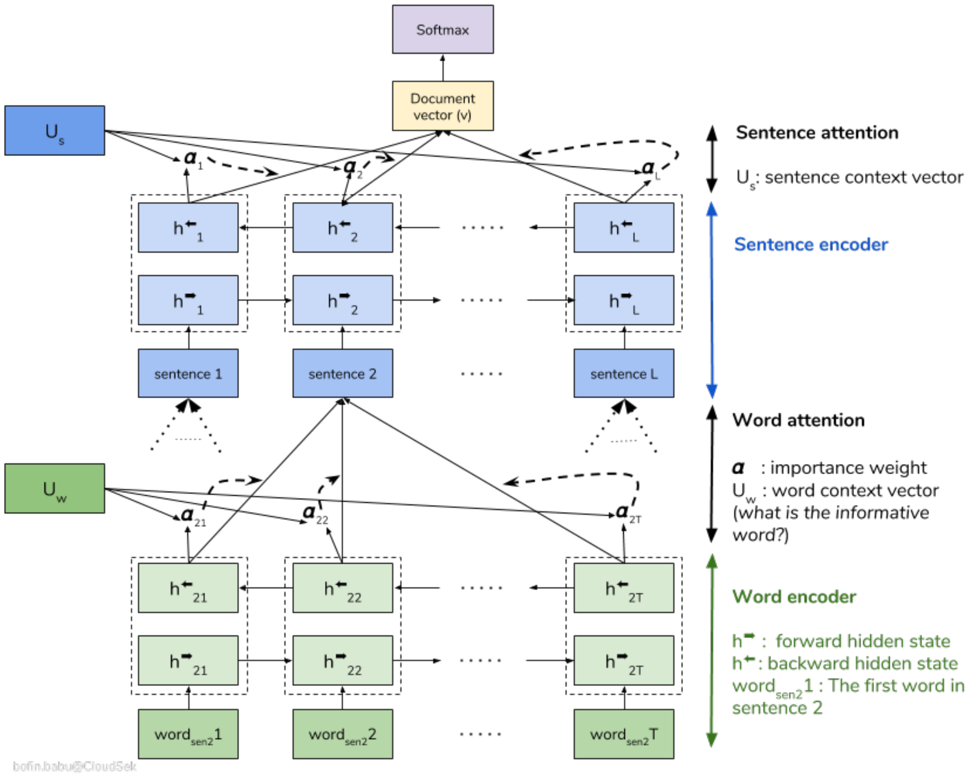

Figure 2.7: Hierarchical Attention Networks

Bidirectional Encoder Representations from Transformers (BERT)

An innovative approach for Natural Language Processing (NLP) called BERT (Bidirectional Encoder Representations from Transformers) was developed by Google AI Language researchers and publicly released at the end of 2018 [18]. Because of its outstanding performance, BERT was an instant hit in the Deep Learning community.

Prior to BERT, a language model would either look at the text sequence from left to right during training or from a combination of left to right and right to left. For creating sentences, this one-directional method works well. We may anticipate the subsequent word, add it to the sequence, and then predict the subsequent word until we have a complete sentence. BERT, a bi-directionally trained language model, now joins the conversation (this is also its key technical innovation). As opposed to the single-direction language models, we can now perceive the context and flow of language more deeply.

BERT employs a cutting-edge method called Masked LM (MLM), which randomly masks words in the phrase before attempting to predict them. This method is used in place of traditional word prediction. When a word is "masked," the model uses both left and right surrounds, as well as both directions of the sentence, to anticipate the hidden word. It considers the previous and following tokens simultaneously, in contrast to earlier language models. This "same-time portion" was absent from the existing combined left-to-right and right-to-left LSTM-based models. (BERT may be more accurately described as non-directional.)

BERT is dependent on a Transformer (the attention mechanism that learns contextual relationships between words in a text). An encoder to read the text input and a decoder to create a prediction for the task make up a basic Transformer. Since the objective of BERT is to produce a language representation model, just the encoder portion is required. A series of tokens that are first transformed into vectors and then processed by the neural network makes up the input to the BERT encoder. However, BERT requires the input to be modified and embellished with additional metadata before processing can begin. **Token embeddings:** At the start of the first sentence, a [CLS] token is added to the input word tokens, and at the conclusion of each sentence, a [SEP] token is added. **embeddings of segments:** Each token receives a marking designating Sentence A or Sentence B. Because of this, the encoder can tell which sentences are which. **Positional embeddings:** Each token is given a positional embedding to show where it belongs in the sentence. Figure 2.8 [32] shows the working. The next word in the sentence is not attempted by BERT. Two tactics are employed during training.

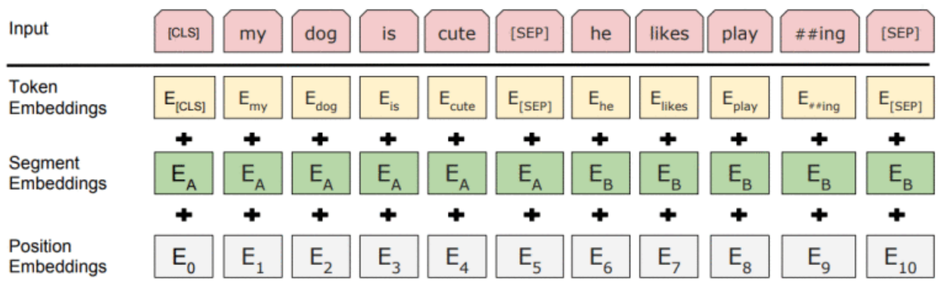

Figure 2.8: Bidirectional Encoder Representations from Transformers (BERT) Working

Masked LM (MLM): BERT runs the full sequence through the attention-based encoder before predicting only the masked words based on the context provided by the other non-masked words in the sequence. BERT then randomly masks off 15% of the words in the input, substituting them with a [MASK] token. This naive masking strategy has a drawback in that the model only attempts to predict the correct tokens when the [MASK] token is present in the input when what we really want is for the model to attempt to predict the proper tokens regardless of which token is present in the input. To address this problem, the following 15% of the tokens were chosen for masking. The token [MASK] actually replaces 80% of the other tokens. 10% of the time, a random token is used to replace a token. Tokens are unaltered 10% of the time. The BERT loss function ignores the prediction of the non-masked tokens during training and only takes into account the prediction of the masked tokens. As a result, the model converges significantly more gradually than right-to-left or left-to-right models. **Next Sentence Prediction (NSP):** The BERT training procedure additionally makes use of the next sentence prediction to comprehend the relationship between two sentences. When doing activities like answering questions, a pre-trained model with this kind of expertise is useful. The model learns to predict whether the second sentence is the next one in the original text when given pairs of sentences as input during training.

Chapter 3

Literature Review

Knowledge graphs, a rich semantic resource have several benefits. Yet, its creation is extremely difficult and laborious. There have been some advancements in its creation but total automation or automatic creation in few and easy steps are not done before. Knowledge graphs with a bundle of information can be procured for overcoming several limitations but their use have been remain restricted. Similarly, knowledge graphs in nature are explainable but their essence has not been used for the sake of results interpretation. I have proposed three different methodologies to overcome these limitations. The related work below shows the limitations of the research done before.

3.1 Knowledge Graph Automation

Knowledge Graphs have become a popular research area due to their several advantages over standard relational databases and deep learning models. They are useful in a variety of industries, including business, law, education, finance, and healthcare. They have been a fascinating area of study order domain as well [68]. The hub of triplets can be utilised to provide answers to a variety of queries and aid in the efficient administration of various data sources [28].

Despite their impactful significance, their creation is extremely difficult and limited as they need the collaboration of multidisciplinary teams. Besides, manual construction is extremely laborious and time-consuming as computer experts have to create hundreds to thousands of triplets. A manual effort was done for KG creation using the Influenza use case [42]. Data was collected from different sources such as Immport, ImmuneSpace, Reactome, and the Center for disease control and prevention (CDC).

Table 3.1: Knowledge Graph Creation Research with Limitations.

Authors	**Year**	**Dataset**	**Limitations**
Hofmann, Alexandra and Perchani [15]	2017	Multiple small knowledge graphs	Manual Creation of Knowledge Graph using structured data
Wu, Tianxing and Wang, Haofen and Li [39]	2020	Multiple Encylopedias	Manual Creation of Knowledge Graph using structured data
Bukhari, Syed Ahmad Chan and Pawar [42]	2021	Reactome, (CDC), Immport, ImmuneSpace	Manual Creation of Knowledge Graph using structured data
Jozashoori, Samaneh and Vidal[23]	2019	Online data sources	Semi-automated Creation of Knowledge Graph
Wang, Qingyun and Li, Manling and Wang [38]	2020	Biomedical Knowledge in Literature	Semi-automated Creation of Knowledge Graph
Cheng, Binjie and Zhang, Jin and Liu [43]	2021	Crowd-sourcing website data	Semi-automated Creation of Knowledge Graph
Malik, Khalid Mahmood and Krishnamurthy [33]	2020	Inhouse Built ontology (ICO), Stroke-related Literature, BioPortal Ontologies	Multiple Complex steps

Knowledge Graphs can also be created by combining various data sources, allowing for the early stages of data cleansing to be skipped for the creation of a comprehensive KG. In order to locate and fetch intricate and unknown biological relations because humans are unable to understand that much knowledge at once, data from several sources was combined for KG development. Small knowledge graphs were combined manually to produce a vast and rich knowledge base [15]. Similar to this, authors [23] developed a rule-based mapping framework called "MapSDI" to combine various data sources and semantically enhance the data.

To facilitate the development of Knowledge Graphs and to utilise them for various purposes, some semi-automated methods have been suggested. Knowledge extraction and knowledge linking are the two main elements of the technique presented by Wu, Tianxing and Wang [39]. Two core modules, Regular Extraction and Live Extraction were proposed. Live extraction only detects fresh and updated articles, whereas the regular extraction sub-module seeks and downloads articles over several stages/time periods. By employing entity matching, the knowledge-linking module locates duplicate entities. Several encyclopedias were combined to create Knowledge Graph.

Similar to this, a thorough knowledge discovery framework known as COVID-KG was created to extract chemical structures and correlations between concepts utilising pertinent biomedical knowledge from scientific literature. Wang, Qingyun and Li [38] used the drug repurposing use case to extract chemical structures and relationships between concepts. She also acquired biomedical information from the literature. For the aim of answering questions, this framework was utilized.

A semi-automated methodology was used to create the stroke-related Knowledge Graph [43]. KG was developed by constructing a dictionary and ontology of stroke disease using crowd-sourcing website data and standard medical term sets. The nodes were next connected to external data by identifying similarities with entities.

The proposed approaches are useful but still total automation has been an unaccompanied task. Malik, Khalid Mahmood and Krishnamurthy [33] tried to automate the creation process with multiple steps and layers such as Semantic Knowledge Layer, statistical Knowledge layer, knowledge factory layer, and knowledge factory layer. Relevant ontologies from BioPortal and self created Intracranial aneurysm ontology (ICO) was used for Graph creation.

Above mentioned approaches are useful but still, a lot of them have put effort into manual creation and a lot of time is consumed. Semi-automated approaches also require human intervention and testing. Even for the case of total automation, the proposed approach is difficult as it is comprised of multiple layers. Most of these approaches relied on simple or limited data, which is not the true nature of Knowledge Graphs. Our approach of ‘ Automated Knowledge Graph creation’ is easy and simple to implement as it contains four simple steps of Natural Language Processing, Ontology mapping, Concept Enrichment, and Neo4j Knowledge Graph Creation. It also contains enriched information from multiple expert-created ontologies. It overcomes the above-mentioned limitations.

3.2 NeuroSymbolic AI Approaches

Two alternative NeuroSymbolic AI approaches were what we had suggested. Both Domain-Specific Knowledge Infusion and Knowledge Graph-based Explainable Predictions fall under this category. The pertinent literature is listed here, along with any downsides.

3.2.1 Domain-specific Knowledge Infusion

Computer Assisted Coding (CAC) is used to automate the medical billing process and save medical coders time and effort. Computer Assisted Coding (CAC) or Automated Medical Coding (AMC) fall into the category of Multi-label classification. For one clinical note, there are various codes or labels available. The Multi-label classification models are trained on a more condensed sample, such as the top-20 or top-50 codes. As a result, they are limited in scope and cannot be generalized. The accuracy of these classification models has also drawn the attention of researchers.

Convolutional Neural Network (CNN) was employed with an attention mechanism for the creation of an explainable clinical decision support system (DSS) [67]. Spanish corpus was used as a training dataset. The confidence and F-1 score were calculated on the top 50 or top 100 codes as the wider problem domain performed poorly.

The hierarchical label-wise attention Transformer model (HiLAT) [59] was presented for the explainable prediction of medical codes. To construct a code-specific clinical note/document representation that was sent to the feed-forward neural network for medical code prediction, a pretrained transformer model was fine-tuned and a bi-level hierarchical label-wise attention mechanism was used. ClinicalplusXLNet was created to perform continuous pretraining from XLNet-Base using Mimic-III data. On the top 50 codes, an F-1 score of 73.5% was obtained.

To increase the multi-label classification's accuracy in medical a Shallow and Wide

Attention Convolutional Mechanism (SWAM) was proposed to improve the ability of CNN-based models to learn from local and low-level medical code features [50]. MIMIC-III was used to put the method to the test. The results revealed that precision was enhanced from 0% to 53% on average for the 10% lowest-performing labels. Furthermore, the method outperformed previous results on the top 50 medical code predictions. For explainability, the attention mechanism was introduced.

Another comparable approach was presented, in which a Hierarchical Attention Network (HAN) was used to predict medical codes [45]. The MIMIC-III database was used to train the model on the top 20 and top 50 ICD codes. The technique was broken down into three layers: the embedding layer, the hidden layers, and the prediction layer. Along with Word2vec embedding, label embedding was included, giving the system the name Hierarchical Label-wise Attention Network (HLAN). As the name implies, the attention mechanism was induced at both the phrase and word levels. With the addition of label embedding, overall accuracy improved by nearly 1-2%, as measured by the F-1 score.

A similar approach of a transformer-based architecture called "TransICD" [41] was proposed to discover label interdependence, and a code-wise attention mechanism was used to learn the code-specific representation of the clinical note or document. The authors used the label distribution aware margin (LDAM) loss function to deal with the code frequency imbalance in the dataset. Mimic-III was procured, and an F-1 score of 64.4% on the top 50 labels was obtained.

For knowledge infusion, an approach known as "G-coder" was put forth, in which MIMIC-III data was used to train a multi-layer CNN for the prediction of medical codes [37]. In addition to this, adversarial learning and knowledge infusion were carried out. On the top 50 codes, the model had an accuracy of 69.2%. By mapping to freebase, a tiny knowledge graph was induced. The resulting graph has 1560 nodes and over 20,000 relationships.

All of the ideas given are useful, but there are some limits. All of these models have been trained for a specific issue area and cannot be generalised. Furthermore, the models are dependent on training data and cannot perform on their own. They fail because of a lack of domain knowledge. Some recent studies attempted to infuse information, but the knowledge injection was limited to a small portion of knowledge. By combining domain-specific knowledge graphs with input data, our knowledge consolidation/infusion

Table 3.2: Computer-Assisted-Coding (CAC) Research with Limitations.

Authors	**Year**	**Dataset**	**Limitations**
Teng, Fei and Yang, Wei and Chen [37]	2020	MIMIC-III	Trained for top-50 codes and dependence on training data
Dong, Hang and Suarez-Paniagua [45]	2021	MIMIC-III	Trained for top-50, top-20 codes and dependence on training data
Biswas, Biplob and Pham [41]	2021	MIMIC-III	Trained for top-50 codes and dependence on training data
Hu, Shuyuan and Teng [50]	2021	MIMIC-III	Trained for top-50 codes and dependence on training data
Trigueros, Owen and Blanco[67]	2022	Spanish corpus	Trained for top-50, top-100 codes and dependence on training data
Liu, Leibo and Perez-Concha [59]	2022	MIMIC-III	Trained for top-50 codes and dependence on training data

technique overcomes these limitations.

3.2.2 Knowledge Graph-based Explainable Predictions

Explainability is a contemporary criterion that may soon become a legal requirement in fields where human lives are at stake, such as education, healthcare, law, finance, and so on. Transparency inclusion in black-box deep learning models necessitates mathematical model manipulation, which necessitates professional intervention [30].

Neuro-Symbolic AI is an intriguing and rapidly emerging field of study that has attracted a great deal of attention. It utilises two distinct AI strategies to achieve accuracy and explainability [51]. Although the terms explainability and interpretability are frequently used interchangeably, they have separate meanings. The ability of an AI model to be explicit or transparent about its inner workings is referred to as interpretability, whereas the ability of an AI model to explain its results or predictions is referred to as explainability [47]. In this section, we will examine the many approaches developed on the subject of Neuro-Symbolic AI, as well as the amount of explainability they provide, as well as their strengths and weaknesses.

Because of their nature, knowledge graphs are considered a clean data source. The triplets' hub contains the subject, object, and predicate. When these graphs are combined with deep learning models, the results can be anticipated more precisely. Authors [27] coupled KGs with a long-short-term memory (LSTM) model [2] to identify thyroid illness. On mental health data, authors [56] performed a superficial infusion of knowledge graphs with neural networks. An automated knowledge graph-generating approach [33] was used to create a graph based on the use case of subarachnoid haemorrhage. For graph generation, a dataset of 1000 summaries was obtained, along with ensemble learning to add rupture probability as nodes.

Deep learning models are very reliant on training data. Due to a lack of domain-specific knowledge, either no forecasts or misleading positive findings are produced. Knowledge graphs are enriched data sources that can help with problem-solving. Knowledge was injected both superficially and thoroughly [47] in a variety of use cases. On the one hand, a self-supervised BERT model was given a brief infusion of domain information from a drug-abuse ontology [18], mostly to assist the model in comprehending the context. Shallow and deep knowledge graph infusions, on the other hand, were performed in educational contexts in an effort to comprehend a student's performance and identify his or her inadequate domain knowledge regions. Certain low-level visuals were generated for clarity in this case. Another strategy similar to process-knowledge infusion was developed to improve the accuracy of classifiers. It made use of psychometric surveys (PHQ-9) as well as process knowledge [64].

Some recent study on knowledge consolidation with inputs has been conducted. The goal is to demonstrate richer knowledge through attention mechanisms in order to increase explainability. Authors [37] presented a method called "G-coder" that combined a multi-layer CNN with an attention mechanism. The end result was a knowledge network with 1560 nodes and over 20,000 relations that mapped the ICD-9 description to Freebase ontology data. To make the terminology and coding results interpretable, the enhanced knowledge network was integrated with the attention mechanism. The model predicted the top 50 codes effectively, but the explainability remained limited to the attention mechanism. In the study of identifying depression symptoms, a graph attention embedding method was used. A psychometric questionnaire (PHQ-9) and patient-written text were used to generate a hypergraph, which then allowed embeddings to be created. An Internet-Psychological Treatment (IDPT) was created by combining a bidirectional LSTM [4] with an attention mechanism to assist people in dealing with depression while utilizing fewer resources [53]. Attention was paid to both the node and the edge levels.

A novel concept known as "graph embedding" generates vector representations of graph facts in a manner similar to word embedding. These embeddings may aid in the accuracy of model outputs. Graphs and neural networks (GNN) were combined for forecasting the risk of mental disease [60]. For the advantage of computer specialists, node embedding and visualization were used to see the model in operation before applying it to the prediction layer. Link prediction algorithms were utilised in a similar way to build a drug repurposing strategy. The study established the relationship(s) between a chemical and a specific target while preserving interpretability and transparency [55].

Knowledge graphs were employed in ways to promote explainability on rare occasions. Researchers [24] developed a "Knowledge-aware path recurrent network (KPRN)" that generated recommendations based on knowledge graphs. The graphs' networks and connections between diverse items can be utilized to understand not only user preferences but also the semantics of entities and relationships. It also produced explainable predictions.

Finally, the papers reviewed in this section coupled knowledge graphs with either machine learning models or deep learning techniques. They have a few things in common. In the bulk of these experiments, KGs were utilized prior to model training with the goal of either adding domain information to improve model performance or showing the enriched knowledge as an output in the attention layer. The restricted explainability of knowledge graphs is not their fundamental nature or goal.

Some studies [24, 40, 65], using "name entity recognition" (NER), datasets, and trained corpora, utilised graphs to explain the results, but they were not in the medical sectors, such as movies and music. Instead, they were in other fields. None of the aforementioned research made an attempt to provide predictability-level explainability. None of them used neuro-symbolic AI to explain the model's predictions or the attention mechanisms. Table 3.3 presents a summary of our studied approaches.

Table 3.3: Neuro-Symbolic AI approaches with Level of explainability.

Authors	**Neuro-Symbolic Approach**	**Deep Learning Model**	**Graph-Type**	**Level of Explainability**
Chai, Xuqing [27]	Knowledge Graph embeddings were employed as training data	LSTM	Knowledge Graph	None
Gaur, Manas and Gunaratna [56]	Shallow Infusion of Knowledge Graphs	Neural Network	Knowledge Graph	None
Malik, Khalid Mahmood and Krishnamurthy [33]	Emplyed Ensemble Learning and used its predications as nodes of graphs	Ensemble Learning	Knowledge Graph	Low
Drance, Martin [55]	Knowledge Graph Embeddings were used	GNN	Knowledge Graph	Low
Sheth, Amit and Gaur [64]	Infusion of Knowledge	Neural Network	Knowledge Graph	Low
Lu, Haohui and Uddin [60]	Graphs employed as input to Neural Network	GNN	Bipartite	Low
Gaur, Manas and Faldu [47]	Knowledge Graph were used as shallow and deep infusion	BERT	Knowledge Graph	Moderate
Wang, Xiang and Wang [24]	Knowledge Graph embeddings used	LSTM	Knowledge Graph	High
Teng, Fei and Yang [37]	Knowledge Graph + Data Infusion	Multi-Layer CNN	Knowledge Graph	Attention Mechanism
Ahmed, Usman and Lin [53]	Knowledge Graph embedding with Bidirectional LSTM	LSTM	HyperGraph	Attention Mechanism

Chapter 4

Design and Methodology

Chapter 3's literature highlights pertinent findings and highlights the limitations in the areas of symbolic AI and deep learning. Three strategies were put forth. The Automated Creation of a Knowledge Graph and two Neuro-Symbolic methods are used to get over Deep Learning's shortcomings. I developed my own dataset and trained the model that was applied on the approaches I suggest. Additionally, a web application is developed for end-user comprehension.

The Chapter is set up as follows: The model's training and dataset preparation processes are described in Section 4.1. The Automated Knowledge Graph Creation Approach is illustrated in Section 4.2. The proposed Neuro-Symbolic Approaches are defined in Sections 4.3 and 4.4, and our developed web application is displayed in Section 4.5.

4.1 Dataset preparation and Model Training

4.1.1 Dataset

For the purpose of model training and technique validation, I procured different datasets. The datasets were I2B2 (Informatics for Integrating Biology the Bedside) Challenge 2010 [6], MIMIC (Medical Information Mart for Intensive Care) [13], and Codisep (Clinical Case Coding in Spanish Shared Task (eHealth CLEF 2020) corpus) [34]. While MIMIC was utilized for model training and testing, Codisep and I2B2 datasets were employed for testing and validation.

More than 40k health records of patients admitted to the intensive care units of "Beth

Israel Deaconess Medical Center" between 2001 and 2012 are available in MIMIC, one of the largest openly accessible databases. The 26 tables that make up MIMIC provide data on patients' lab results, diagnoses, demographics, and vital signs. Patients, Procedures, Services, Lab events, ICU stays, medications, transfers, etc. are all included in the tables. Each of these tables includes various data, including patient admittance, ICU stay, disease diagnosis, prescribed medications, and much more.

The clinical free text note is located in the table Noteevents, which is the most crucial one for us. These contain things like radiological reports, ECGs, and doctor's prescriptions. This table served as the primary training tool for our model. There is a CSV (comma-separated values) version of each MIMIC table. The MIMIC-III database's overview is shown in Figure 4.1. MIMIC is originally annotated with ICD-9 Medical codes. ICD-9, the International Classification of Diseases' ninth revision, is no longer in use. I created the dataset for the ICD-10 edition, which has 11 times as many codes. The step of preparing the dataset is described in the following section.

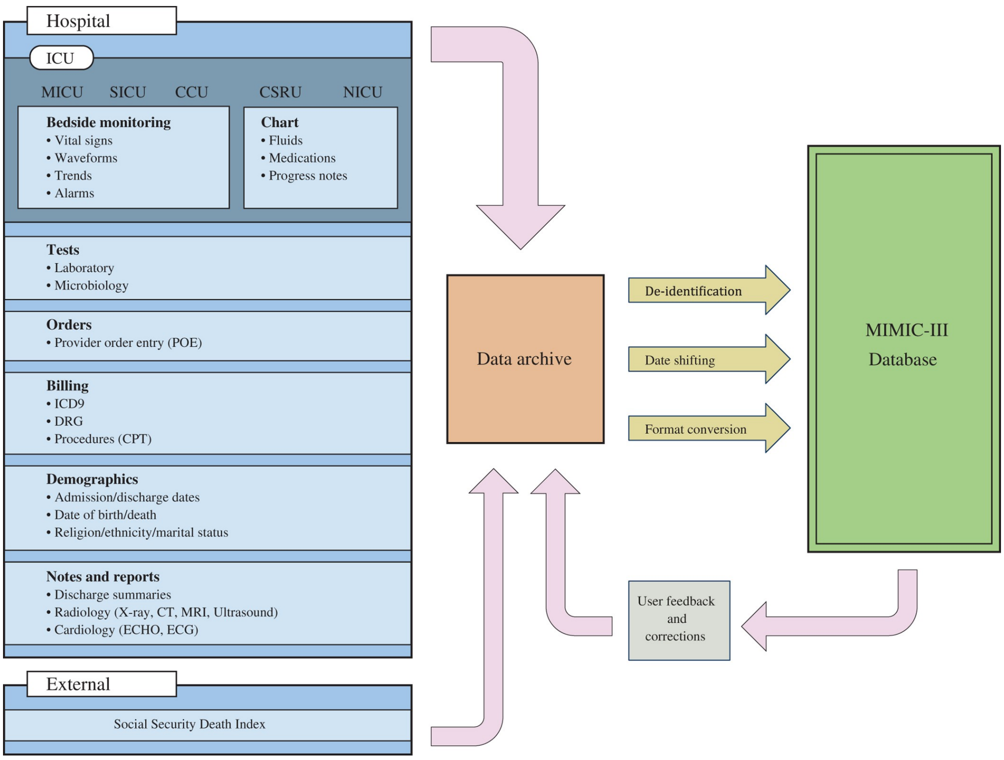

Figure 4.1: Overview of MIMIC-III Database

4.1.2 Dataset preparation

I searched for publicly available datasets for ICD-10 but there is no source for them. The codisep dataset has ICD-10 annotation but there were few discharge summaries, and also I analyzed wrong annotations too. By wrong annotations, I mean that the annotated code was not present in the ICD-10-cm library. Hugging Face is an AI community that contains hundreds to thousands of datasets and models [12]. I procured an already-trained clinical BERT [66] for annotation of the MIMIC-III dataset with ICD-10 codes.

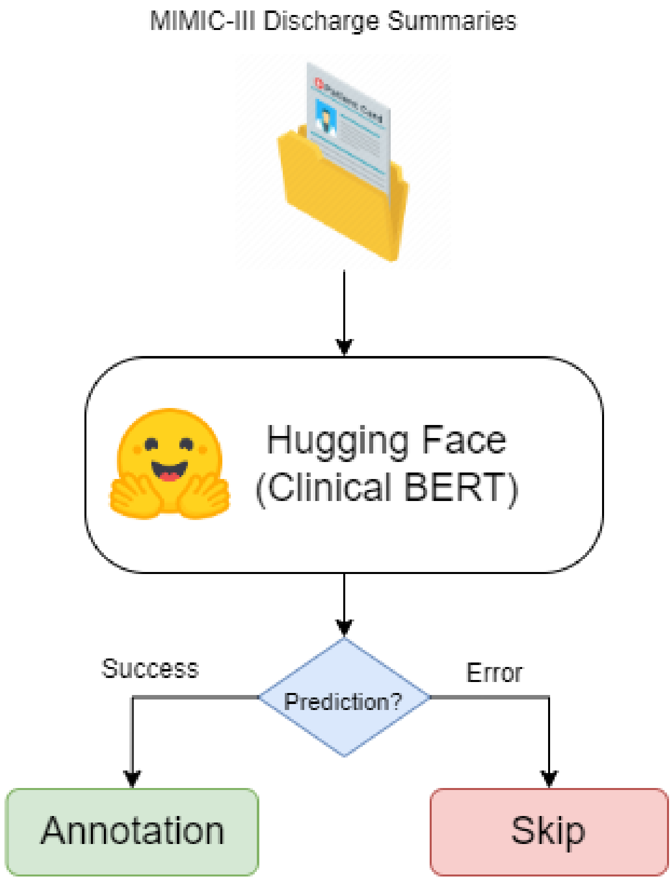

Figure 4.2: Dataset Preparation

With excellent outcomes, the transformer model was trained on a sizable corpus. Utilizing a real medical coder, I tested the outcomes of this model. The medical coder gave these outcomes a really good rating. The problem of scalability was the motivation behind not using this model for predictions. For large summaries, the BERT model was not scalable in terms of output. This problem served as a barrier and constraint in the development of an explainable medical code prediction application.

I used a pre-trained BERT and an algorithm to automatically annotate the MIMIC-III summaries. Our code was programmed to skip the lengthy summaries and annotate the ones it could. The MIMIC noteevents table was primarily utilized. The algorithm received the 40k+ notes for annotation. The dataset's creation and annotation took more than an hour. The dataset was created in the format of "summary___label___code1 code2 code3".

The trained model might yield some inaccurate results (codes which are not present in ICD-10 classification). Perhaps the incorrect manual annotation is to blame. To remove the erroneous codes, I further filtered the results using the ICD-10-cm library. After annotation, I received nearly 5,000 summaries, with each summary receiving an average of four codes. Nearly 22k labels were assigned to summaries, compared to 550 unique labels. The codes weren't specific to any one block or area of the coding hierarchy; rather, they were generic.

4.1.3 Model Training

The majority of deep learning models were trained using the out-of-date ICD 9th version. Since the 10th version is the one that is currently in use, I must change the predictions to reflect this. I examined the performance of various models [20, 29, 41] and replicated the findings for each, before selecting a baseline model called the Hierarchical Label-wise Attention Network (HLAN) [45]. HLAN was chosen because of its predictions and results that could be explained using the Attention mechanism. The top 20 ICD-10 and top 50 ICD-9 codes were used to train the model. On both a word- and sentence-level, the model was delivering explainable results using Attention Mechanism in addition to predictions.

Instead of using ICD-9 predictions, I adjusted the model for ICD-10 code prediction. The model was limited to making predictions for COVID-19-related notes because it was only trained on a limited set of categories. I used 11x–27x more codes to train the model, which can now predict almost all summaries. For much better outcomes, I used the hierarchical structure of the medical coding system. For training purposes, about 82k ICD diagnostic codes were obtained. I was able to successfully train the model and switch it to the ICD 10th version. 3266 summaries were used for training, 800 for testing, and 800 for validation. Table 5.1 shows a comparison between pretrained and fine-tuned models.

4.2 Automated Neo4j Knowledge Graph Creation Framework

It takes a lot of effort and time to create a knowledge graph [63], which calls for assistance from humans as well as specialists in both the medical and computer science fields. There are hundreds to thousands of triplets in a knowledge graph, and each triplet consists of two nodes and one relation. Such a large knowledge base must be created carefully, and manual creation is likely to contain errors. Any semantic language, including RDF, OWL, XML, etc., could be used to create a knowledge graph. With the help of a few simple steps, including data preprocessing, ontology-based information retrieval, concept enrichment, and Neo4j Knowledge Graph Creation, I was able to automate the difficult and time-consuming task of creating graphs. For the development of KG, I used the expert-created Bioportal Ontologies. Below is a detailed definition of each step.

4.2.1 Natural Language Processing

Data from multiple medical concepts, such as diagnosis, prognosis, and prescriptions, are included in discharge summaries. We must process the data and eliminate words that are unnecessary for machine processing or KG creation but make sense to humans. To perform processing, I procured a Natural Language toolkit. I used N-grams, lemmatization, stemming, stop-word removal, sentence detection, word/token detection, and other natural language processing techniques. The steps in relation to medical summaries are described in more detail below.

Sentence and Word Detection

The sentences in medical notes and discharge summaries typically span multiple lines. In order to move on, we must find sentences. It is necessary for the following data preprocessing steps. Sentence completion with characters like ".," "?," "!" and other punctuation marks was used for sentence detection. Figure 4.3 demonstrates how sentences are created from a paragraph. Words or tokens were detected after sentences were found. This is accomplished by using white space or other special characters, such as commas and brackets. The output includes some information that is necessary and some that is not.

Stop-word Removal

The stop-word removal module eliminated extraneous and unnecessary words after the detected words were run through it. The creation of a graph was useless with these words. The words that were deleted included "am," "is," "a," "to," "through," I "his," "with," "some,", "was", etc. The filtered medical words are visible in Figure 4.3 following stop-word removal.

Lemmatization and Stemming

Natural language processing methods like lemmatization and stemming are used to determine a word's root. For instance, "run" serves as the base for "ran". This improves the results and makes them simpler to comprehend. Figure 4.3 demonstrates how the words were reduced to their original root form. Words like "wheezing" became "wheeze," "rales" became "rale," and vice versa.

N-Grams

All feasible word combinations were obtained using n-grams. We must map these words to ontologies developed by experts, therefore this step is essential. A four-node N-gram was used. For instance, there are four words "severe," "acute," "respiratory," and "syndrome". Together, these words make meaning; apart, they will point in opposite ways. These words after word detection were in their n=1 form in their original sense. When n exceeded 2, the form changed to "severe acute," "acute respiratory," "respiratory syndrome," etc. Acute respiratory syndrome, severe acute respiratory, etc. are the forms when n=3. Following a scenario analysis, I used N=4 because it is appropriate.

4.2.2 Ontology-based Information Retrieval

The Bioportal Ontologies were then used to map these terms. The technique was left open-ended. It indicates that I did not limit our approach to mapping and extracting data from a small number of ontologies. There are benefits and drawbacks to this general strategy. The benefit is that we have a lot of data, which makes it easier to create a comprehensive graph. And the drawback is that the reader may become confused as a result of the time and complexity of this vast amount of information. I used the bioportal's 1000+ ontologies for information retrieval.

The data that results from mapping was voluminous and includes data from various sources and in numerous languages. These factors influence the final outcomes in some way, but the amount of information encourages thorough knowledge graph building and greater comprehension. Figure 4.3 displays the ontologies to which I mapped the n-grams. Information retrieval and ontology mapping were performed using the Bioportal REST API.

4.2.3 Concept Enrichment

Medical ontologies and enriched medical concepts can be found in Bioportal. The ontologies provide information about medical concepts such as definitions, synonyms, hierarchies, perLabel, etc. For instance, Figure 4.3 displays richer data around the concept of "wheeze." The terms used to describe the symptoms included "asthmatic breath noises," "bronchospasm," "rhonchus," "asthmatic breathing," "stridor," "bronchospasm finding," "asthmatoid," "increasing exercise," "nocturnal cough," "asthmatoid wheeze," "jackson's sign," and "rhonchi," among others. Definitions included "a disorder characterised by a high-pitched, whistling sound during breathing," "it results from the narrowing or obstruction of the respiratory airways," "it is a symptom and a finding during physical examination," etc., and hierarchy included ['wheezing'. 'is a', 'adverse effect'], ['wheezing'. 'is a', 'Respiratory, thoracic and mediastinal disorders], ['Respiratory, thoracic and mediastinal disorders', 'is a', 'Physical health-related concept']. In order to grasp medical terminology, hierarchy is crucial. For instance, COVID is a virus, and a virus is a type of microorganism. This results in the concept's general meaning.

4.2.4 Neo4j Knowledge Graph Creation

For the development and presentation of our graphs, I employed Neo4j, a scalable and potent visualization tool. Neo4j offers various visualization formats for viewing nodes and relations. Additionally, it has tools and procedures for handling and processing graphs. Neo4j Graph query creation was automated. Neo4j creates graphs using the

Cypher language. The cypher language has its own rules and methods for constructing nodes and the connections between them [19]. The use of the application of ID to nodes is likewise subject to various limitations. These constraints were implemented, the query building process was automated, and the process of creating the Graph was also automated. The diagram displays a small Neo4j graph built for the "wheeze" concept. The actual procedure of creating a query is very challenging and time-consuming. I developed our algorithm to create triplets while adhering to the query structure.

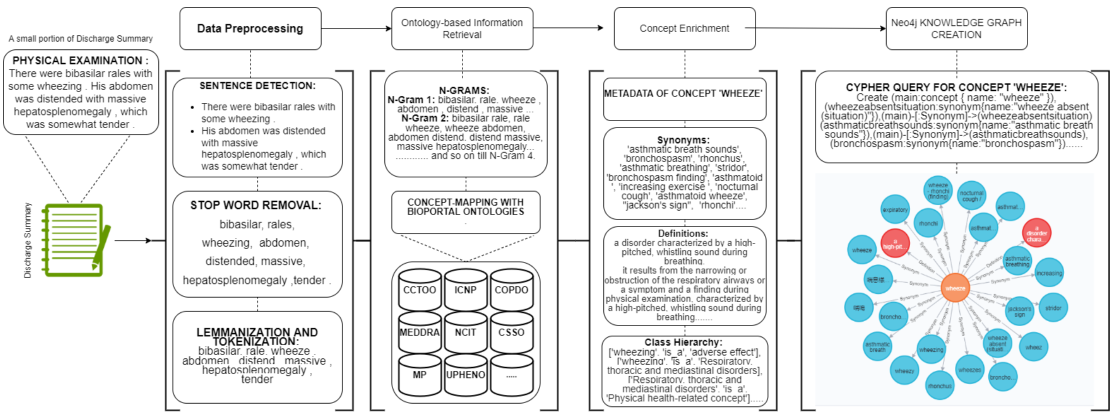

Figure 4.3: Automated Neo4j Knowledge Graph Creation Framework

4.3 Neuro-Symbolic Approach-I

4.3.1 Problem Formulation

Deep learning models forecast with great accuracy, but they have several limits. The models are accurate, but they rely on training data. They use training data to understand patterns but are unable to predict anything out of the box. As a result, either a false positive or no findings are obtained. This issue can be overcome by combining domain-specific knowledge Graphs with discharge summaries, allowing the model to better analyze the data and make more accurate predictions.

4.3.2 Domain-Specific Knowledge Infusion

We understand things in everyday life based on our prior knowledge. We begin learning at birth and create a foundation for understanding complex architectures. This indicates that our prior knowledge aids us in understanding advanced issues. For example, a man went to a village and was taking some medicine. Someone sitting nearby inquired as to why you were taking this medication. I have hypoglycemia, he said. What, the person exclaimed. It is Diabetes Mellitus, he retorted. The question, "What is this?" was repeated. I have sugar, is the reply. The first individual then comprehends what this means. This suggests that in order to help someone understand the context, we must first understand them or their degree of knowledge.

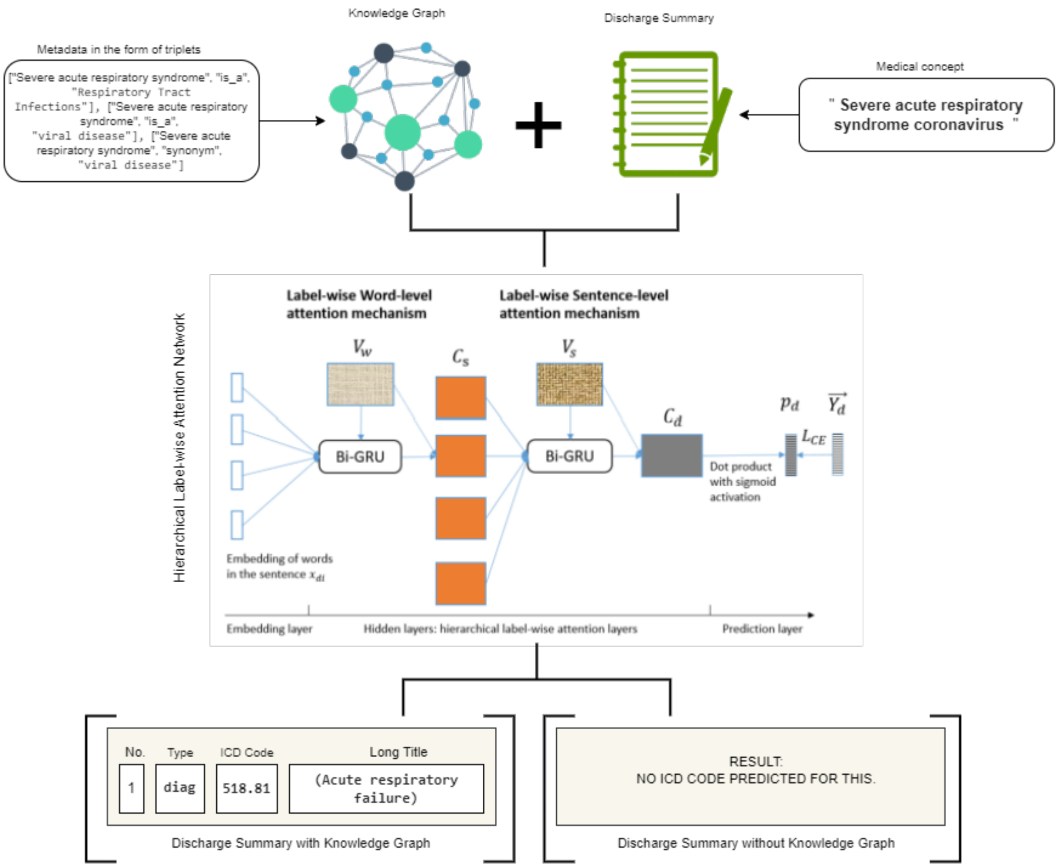

Figure 4.4: Knowledge Consolidation with Deep Learning

It's a universal occurrence that everyone observes. I made use of this fundamental aspect of human nature to improve deep learning models' accuracy and get around their drawbacks. The Knowledge Graphs are an augmented source of knowledge since they have a wealth of connections between the information via rules or triplets. These information-dense graphs can aid the model in comprehending the semantics of input discharge summaries. Our system has semantically enhanced each medical term in the discharge summary.

The HLAN model [45] was trained on the top 50 ICD-9 codes. I examined the model and discovered that it was incapable of producing any results for numerous discharge summaries and, in particular, single medical concepts. It is possible that this is because the training data was missing this information and the model is unable to grasp the input. I consolidated the graph to assist the model in comprehending the results. For example, the synonyms and definitions of a medical concept derived from KG can assist the model in determining a relevant direction.

Figure 4.4 depicts the effect of KG consolidation with the deep learning model HLAN. Prior to the incorporation of KG, the model was unable to predict any results. The semantically enriched KG aided the model in understanding the input and producing outputs. The illustration depicts the outcomes of the medical concept of "Severe Acute Respiratory Syndrome" with and without knowledge graph consolidation.

4.4 Neuro-Symbolic Approach-II

4.4.1 Problem Formulation

The introduction of the Attention mechanism is one of the most recent advances in explainability. In machine translation, the encoding and decoding system receives some data and highlights some attention words. The question is why those terms were highlighted by the model. One possible reason is that they lead to the discovery/prediction of some output. But why those specific words, and why together? Similarly, what is the relationship between these words and their output? Unfortunately, the answer is nothing. Deep learning's black-box nature prevents it from answering such concerns. That is, the explainable system is not itself explainable. By combining Graphs with a Deep Learning model and attention mechanism output, I presented another neuro-symbolic technique to bridge this explainability gap. I suggested a method to use a knowledge graph to show word-to-word and word-to-code level connections.

4.4.2 Explainable Deep Learning Results

I used fine-tuned HLAN model for ICD-10 forecasts. The discharge summaries served as the model's inputs, and the projected labels or medical codes with attention results served as the model's output. These terms are highlighted by the model, indicating that they all have an impact on this prediction. Yet how? Similar to how they will have some links if they are highlighted together. Such questions cannot be answered by deep learning due to their complexity and black-box nature, but they can be by Symbolic AI! I combined the HLAN model's output data and used the automatic graph construction method to do so. This time, the entire summary is represented via a graph rather than just a single medical concept. Figure 4.5 shows the inputs and ontologies used in Graph creation.

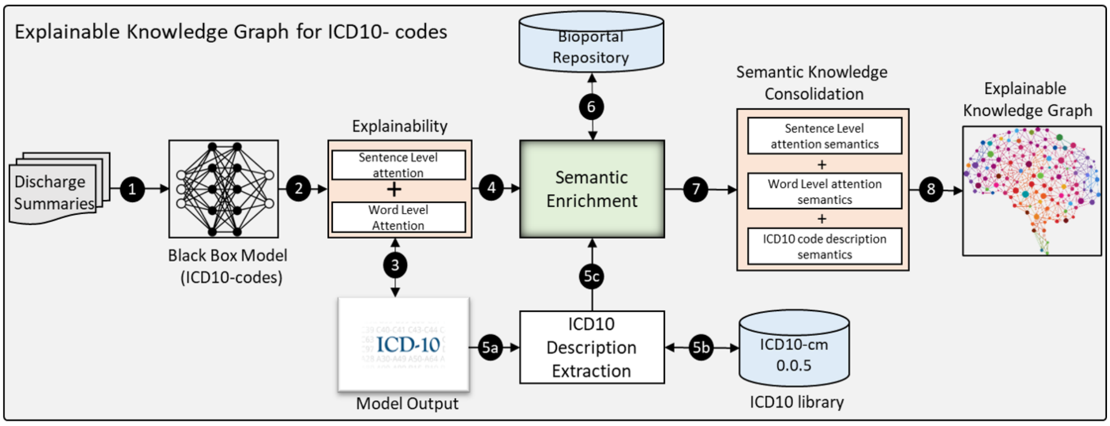

Figure 4.5: Explainable Knowledge Graph Creation Approach

Semantic Enrichment

Understanding the context and boundaries of certain words is made easier by semantics. For our second technique, we must comprehend how the highlighted terms relate to one another and to the predicted medical codes. The inputs for semantic enrichment were model predict labels and words that were highlighted by an attention method. To extract the predicted code descriptions, I made use of the ICD-10-cm package [62]. The word and phrase detection process was skipped. Synonyms, definitions, parent / hierarchy-level-1, parent / hierarchy-level-2, till parent / hierarchy-level-5 were added to words to enhance their meaning. For the end user, such rich information can offer a deeper level of explainability.

Semantic Knowledge Consolidation

The redundant synonyms, definitions, and hierarchical descriptions were removed as different sorts of knowledge pertinent to attention weights and description words were combined. Additionally, I divided the nodes into categories based on their kinds, such as patient_summary, medical_code, medical_code_desp_words, model_attention_words, synonyms, definitions, parent_1, parent_2, parent_3, parent_4, and parent_5. Similar to this, duplicate relationships and relationship kinds were eliminated. Synonyms, definitions, highlighted_words, description_words, parent_level_1, parent_level_2, parent_level_3, parent_level_4, parent_level_5, and connected were the several kinds of relationships.

Explainable Knowledge Graph Creation

I applied our method for automatically creating knowledge graphs. Due to the diversity and complexity of the data, knowledge graph generation is a computationally intensive task, with run-time creation being even more challenging. As each summary is input for ICD-10 code recommendations in the deep learning model, a specific graph is generated for that summary. The relationships between the graph nodes aid in our understanding of context and semantics in general, but the 'connected relationship' focuses specifically on how the attention mechanism and model prediction work. In the medical billing codes prediction, words are given attention weights based on how similar they are. By offering two different types of explainability—word-to-word connections and word-to-code connections—I have made it easier for users to understand through graph visualization.

Word-Word Level Connection

Due to the attention mechanism, I received some highlighted words in our situation. Why does the model highlight them together? Are these words related to one another? What link exists between these two words? Is it appropriate for the model to highlight them both together? because they all ultimately result in the predicted label. If they are predicting one label, then we deduce that they must also be connected to one another.

One patient visited a doctor, for instance, after suffering a knee injury. The doctor bandaged him up and gave him some painkillers to take. The patient then remarked, "I also need some hair loss medicine, maybe some vitamins." That was also noted by the physician on the same prescription. Only words with similar meanings should be highlighted by the model. The word "hair loss" shouldn't be highlighted in the knee injury diagnostic code. It's incorrect if the model behaves that way. How then should we evaluate that? possibly manually. Considering that deep learning cannot explain why such words are highlighted? And why do they work together?

Symbolic AI enters the picture at this point. To unlock this nature's mystery, I employed symbolic AI. I mapped these terms to expert-made ontologies and retrieved the links. It begins to make sense why the model behaved in this way, even though the linkages may not be direct. Our method allows end-users to interpret the results by visualizing links between them, which is the essence of symbolic AI. The strength of such connections can be utilized to evaluate the performance of the model.

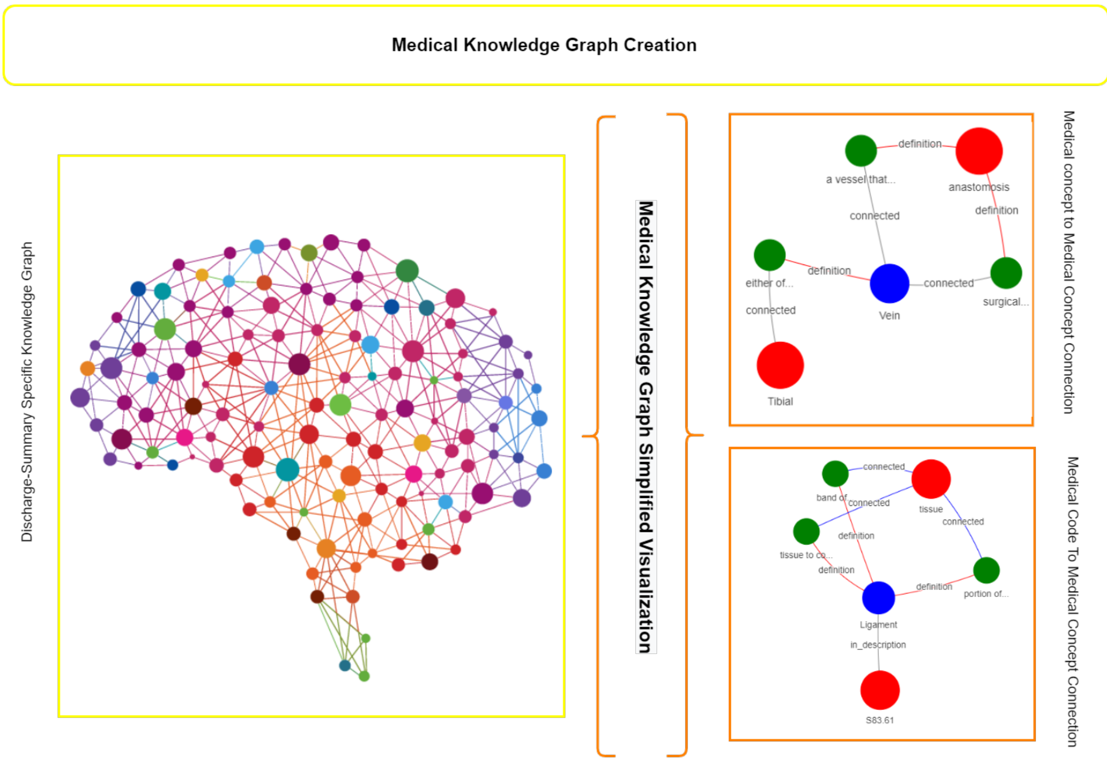

Figure 4.6: Knowledge Graph Simplified Visualization

Word-Code Level Connection

Using attention weights, the deep learning model draws attention to the words and illustrates their significance. The model aims to convey to the user that these phrases are interconnected and contribute to the prediction of the medical code in some way. The questions are, is the model working correctly? If so, how does this prediction result from these words? What association exists between the label and the words? Such inquiries have no solutions. Symbolic AI can provide answers to these inquiries. A medical code's description draws a line around its specific domain. I preprocessed the data, applied ontologies to enrich it, and fetched relationships using words from medical code descriptions.

4.5 Web Application

We are unable to offer our techniques and methodologies for the benefit of end users. Not that outside of this field, but computer scientists, can understand these methods. A web-based tool for code prediction and attention visualisation was made for the benefit of medical coders and other relevant people. A discharge summary is entered on the application's first page, trained model checkpoints are called, and results are displayed on the second page. The first page of the web application is shown in Figure 4.7, and the second page is shown in Figure 4.8. The flask-based web application is made in visual studio.

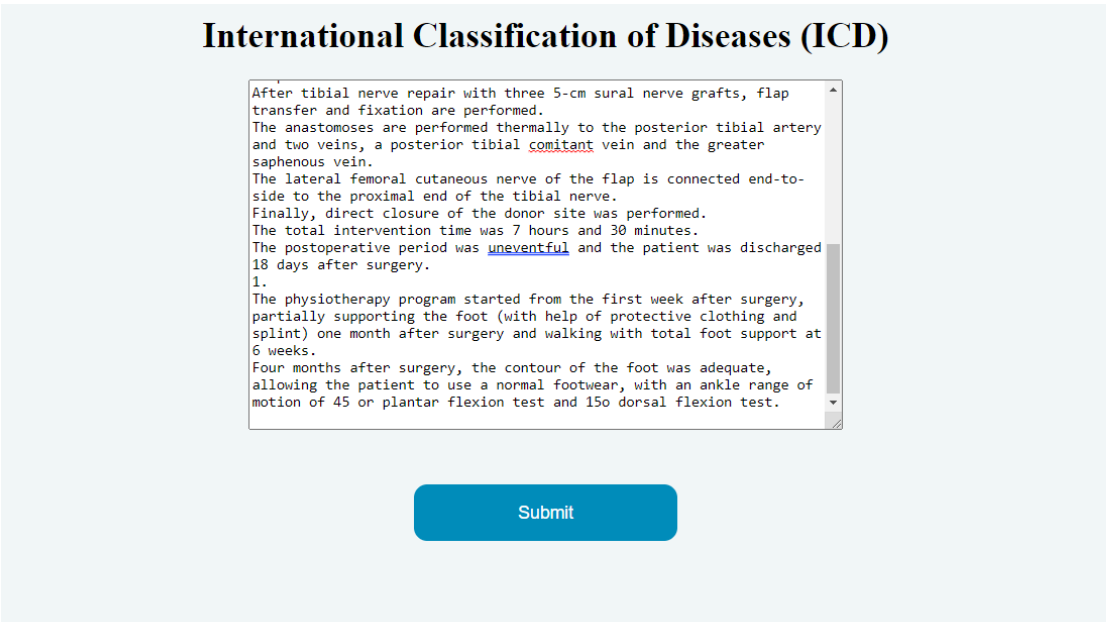

Figure 4.7: First/Input Page of Web Application

For the explainable visualization of attention results and medical code predictions, the Knowledge Graphs are consolidated and other pages were added. The process of creating a graph is quite helpful for comprehending how these deep-learning models forecast.

Although they can aid in several diagnoses, projections, and related duties, these are particularly difficult to understand. The knowledge graphs that were actually built for the various summaries had, on average, 2000 nodes and 3000 relationships. They are extremely complex and challenging for the end user to view and analyze. They will make life extremely difficult for the medical coder rather than making the model prediction simpler.

I employed visualization strategies to leverage Knowledge Graphs to describe the outcomes. Certain human behaviours were considered when creating the depiction. Humans tend to notice larger objects first, and colours can also be used to symbolize different concepts. For instance, red denotes danger, white denotes calm, and bold results indicate significance. These were utilized to produce the visualization.

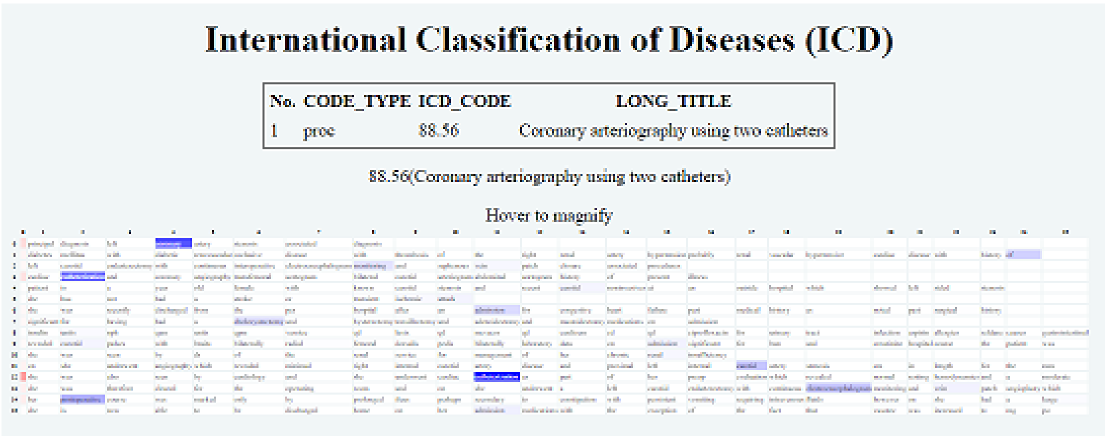

Figure 4.8: Second/Output Page of Web Application

I used user-entered terms and codes to query the graph and retrieve results. The links between them were all shown in the results, and users may limit their view to only direct or strong connections. You can see from the approach diagram that the trinity library was used to query the actual graph using Cypher. Using the kglab library, a basic visualization was created. The shortest path was determined using a variety of path detection techniques, including Neo4j's depth-first search and breath-first search algorithms. The diversity of generated graphs caused these algorithms to perform poorly. In order to obtain the findings, I used basic Cypher functions. This application will soon be finished and the results will be reported in a journal paper. The entire application flow and results are ready. Figure 4.6 provides an overview of this Knowledge Graph simplification for word-to-word and word-to-code level connections.

Chapter 5

Implementation and Results

This chapter shows the outcomes of our three innovative approaches: Automated Knowledge Graph Creation, Domain-specific Knowledge Infusion, and Explainable Deep Learning Results.

5.1 Model Training

The dataset was prepared in a specific format. The codes were appended to the last of a discharge summary with the term "__label__" in between. Figure 5.1 shows a glimpse of the prepared dataset format.

```
An 8-year-old girl who has a left ankle sprain. CRPS of the ankle and
left foot is diagnosed and remits completely within a couple of weeks
with medical treatment. Two months later, without any previous trauma,
the patient developed ankle and right foot pain, which was also
diagnosed with CRPS. He was referred to the Pain Unit by the Pediatric
Traumatology Service two months after diagnosis. Pregabalin plus
tramadol was started, plus capsaicin 8 % patch application due to
severe allodynia, without adequate pain relief or other accompanying
symptoms. Therefore, it was decided to remove (after paterno consent)
the implantation, in the operating room and under general anesthesia,
of a lumbar epidural catheter and its connection to an external pump
for the administration of continuous infusion of bupivacaine for two
weeks. Complete remission of symptoms that continues __label__M89.072
M25.571 R52 T14.90 S93.402
```

Figure 5.1: Dataset Format

I trained the baseline HLAN model on ICD-10 codes and obtained an F-1 score of 67.2%

on the top five codes, with 800 summaries used for testing and 800 for validation. Table 5.1 displays the comparing results. The trained model was slightly less accurate than the pretrained models, but this was completely expected given that I expanded the problem domain and acquired 11-27x more medical codes.

Table 5.1: Comparison between pretrained and fine-tuned HLAN model.

Dataset	**Train**	**Validation**	**Test**	**Label Count**	**Score calculation**	**F-1 Score**
MIMIC-III (pre-trained)	8066	1573	1729	50	Top-50	64.1%
MIMIC-III (pre-trained)	4574	153	322	20	Top-20	74.6%
MIMIC-III (fine-tuned)	3266	800	800	550	Top-5	67.2%

5.2 Automated Neo4j Knowledge Graph Creation Framework

Our Automated Knowledge Graph creation approach was tested on 80 different medical concepts. Neo4j was used for graph creation because of its scalability and visualization. The Bioportal REST API dealt with the standard format of ontologies e.g XML, OWL, RDF, etc and fetched the enriched data.

There were an average of 10 unique definitions, 15 unique classes, and 20 unique synonyms fetched from BioPortal Ontologies. A knowledge graph of this size can definitely help us to increase the overall performance of a deep learning model.

The "Neo4j Knowledge Graph Creation" step of our automated approach created the query for the respective medical concept graph. Our approach has automated the Cypher query creation process and thus automated the Knowledge graph process. The query creation is an extremely laborious and burdensome task. A small chunk of the query (3,749 characters) created for the term "Insomnia" is shown in Figure 5.3. The original query was made up of nearly 19,881 characters.

```
Create (main:concept { name: "insomnia" }),(disordersinsomniasleep:synonym{name:"disorders
insomniasleep"}),(main)[:Synonym]>(disordersinsomniasleep),(delayedsleepphasesyndromesusceptibilit
yto:synonym{name:"delayed sleep phase syndrome, susceptibility to"}),(main)-[:Synonym]-
>(delayedsleepphasesyndromesusceptibilityto),(agrypnia:synonym{name:"agrypnia"}),(main)-
[:Synonym]->(agrypnia),(sleepdifficult:synonym{name:"sleep difficult"}),(main)-[:Synonym]-
>(sleepdifficult),(sleepdisturbed:synonym{name:"sleep disturbed"}),(main)-[:Synonym]-
>(sleepdisturbed),(sleeprestless:synonym{name:"sleep restless"}),(main)-[:Synonym]-
>(sleeprestless),(fragmentedsleep:synonym{name:"fragmented sleep"}),(main)-[:Synonym]-
>(fragmentedsleep),(insomniadisorder:synonym{name:"insomnia (disorder)"}),(main)-[:Synonym]-
>(insomniadisorder),(insomniaadverseevent:synonym{name:"insomnia adverse event"}),(main)-
[:Synonym]->(insomniaadverseevent),(omnia:synonym{name:"asomnia"}),(main)-[:Synonym]-
>(omnia),(sleeplessness:synonym{name:"sleeplessness"}),(main)-[:Synonym]-
>(sleeplessness),(insomniasleepdisorder:synonym{name:"insomnia sleep disorder"}),(main)-[:Synonym]-
>(insomniasleepdisorder),(sleepdecreased:synonym{name:"sleep decreased"}),(main)-[:Synonym]-
>(sleepdecreased),(mysleepqualitywas:synonym{name:"my sleep quality was"}),(main)-[:Synonym]-
>(mysleepqualitywas),(不眠症:synonym{name:"不眠症"}),(main)-[:Synonym]->(不眠症
),(vigil:synonym{name:"vigil"}),(main)-[:Synonym]->(vigil),(insomniactcae50:synonym{name:"insomnia,
ctcae 5.0"}),(main)-[:Synonym]->(insomniactcae50),(insomniadisease:synonym{name:"insomnia
(disease)"}),(main)-[:Synonym]->(insomniadisease),(insomnias:synonym{name:"insomnias"}),(main)-
[:Synonym]->(insomnias),(dspd:synonym{name:"dspd"}),(main)-[:Synonym]-
>(dspd),(insomnia:synonym{name:"insomnia"}),(main)-[:Synonym]->(insomnia),(不眠:synonym{name:"
不眠"}),(main)-[:Synonym]->(不眠),(dsps:synonym{name:"dsps"}),(main)-[:Synonym]-
>(dsps),(delayedsleepphasedisordersusceptibilityto:synonym{name:"delayed sleep phase disorder,
susceptibility to"}),(main)-[:Synonym]
>(delayedsleepphasedisordersusceptibilityto),(difficultystayingorfallingasleep:synonym{name:"difficulty
staying or falling asleep"}),(main)-[:Synonym]-
>(difficultystayingorfallingasleep),(insomniadistress:synonym{name:"insomnia distress"}),(main)-
[:Synonym]->(insomniadistress),(sleepquality:synonym{name:"sleep quality"}),(main)-[:Synonym]-
>(sleepquality),(wakefulness:synonym{name:"wakefulness"}),(main)-[:Synonym]-
>(wakefulness),(hyposomia:synonym{name:"hyposomia"}),(main)-[:Synonym]-
>(hyposomia),(awakeningearly:synonym{name:"awakening early"}),(main)-[:Synonym]-
>(awakeningearly),(insomniasymptom:synonym{name:"insomnia symptom"}),(main)-[:Synonym]-
>(insomniasymptom),(insomniasymptoms:synonym{name:"insomnia symptoms"}),(main)-[:Synonym]-
>(insomniasymptoms),(ad040:definition{name:"ad040"}),(main)-[:Definition]-
```

Figure 5.2: Cypher Query for Concept "Insomnia"

Knowledge Graphs are useful in many aspects and one of their major advantages is their explainable nature through their simple visualisation of nodes and the connections among them. These graphs apart from other uses can be a source of domain understanding.

Figure 5.2 shows a generated knowledge graph for the concept of insomnia which is a sleep disorder. In the figure, the synonyms are presented as blue nodes, classes (till depth one) as orange, definitions as red, and ontologies as brown nodes. These all nodes contain enriched information about some medical term. These nodes and relations can be used to understand this concept in detail and vice versa.

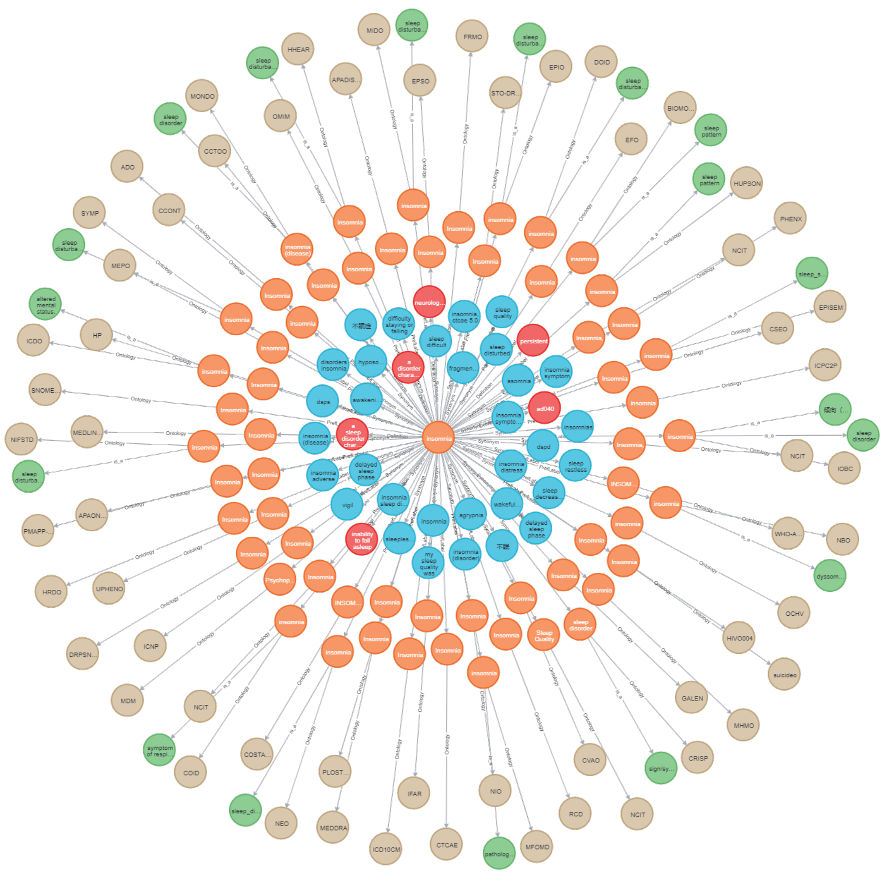

Figure 5.3: Knowledge Graph for the concept "Insomnia (sleep disorder)"

5.3 Domain-Specific Knowledge Infusion

I assessed HLAN's performance by combining the enriched medical ideas with the discharge summary. For example, I input the term "Severe Acute Respiratory Syndrome Coronavirus" and received no results, claiming that "No ICD code predicted," but this changed when the concept was enriched. The HLAN model predicted "diag code 518.81 (Acute respiratory failure)" after adding the Knowledge Graph. Similarly, HLAN could not predict any results for the notion of "Polyuria," but with the addition of the knowledge graph, results were "diag code 599.0. (Urinary tract infection, site not specified)".

For them, visualisation also demonstrated the significance of Graph technology. I put HLAN to the test using over 100 discharge summaries from MIMIC-III (Medical Infor-

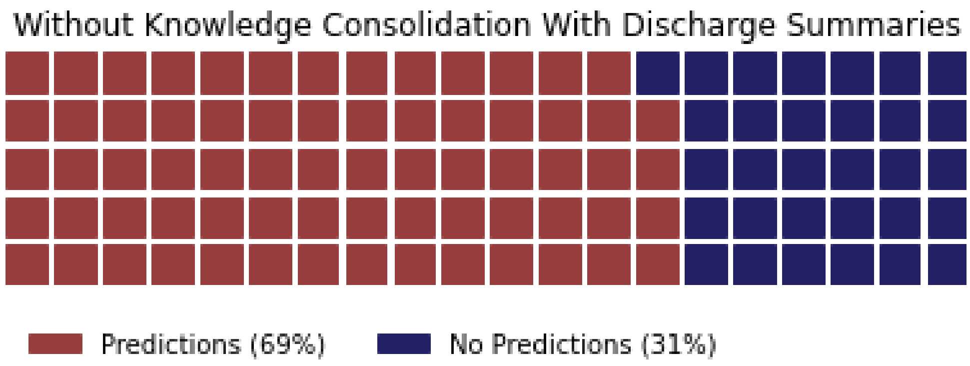

Figure 5.4: Without Knowledge Infusion

mation Mart for Intensive Care) [13] and I2B2 (Informatics for Integrating Biology and the Bedside) Challenge 2010 [6]. Without the Knowledge Graph Fusion, HLAN gave findings for 69 summaries. The Knowledge Graph Fusion assisted HLAN in predicting relevant codes for 82 discharge reports. I used the same amount of synonyms, meanings, and hierarchies for all ideas to ensure that HLAN remains unbiased and does not favour certain terms.

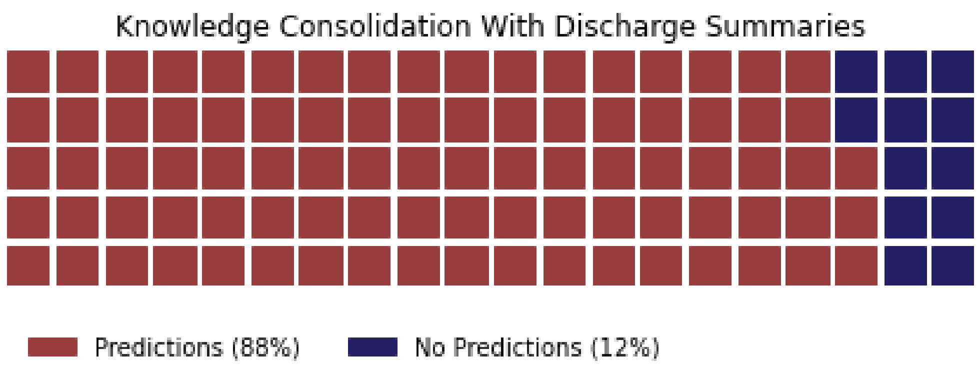

Figure 5.5: With Knowledge Infusion

HLAN is unable to predict any ICD codes for specific medical phrases like "hyperglycemia," "sclerosis," "hypothyroidism," and "auditory hallucinations," among others. After using a knowledge graph to enrich these concepts, I looked up 80 distinct medical terminologies. HLAN was successful in predicting outcomes for 66 of these 80 medical words. Figure 5.4 and 5.5 shows the results of the model with and without knowledge infusion with medical terms. To verify the veracity of the findings, I utilised this this link. Our accuracy was close to 81%.

5.4 Explainable Deep Learning Results

The outputs of the fine-tuned model were then fed to the "Explainable Knowledge Graph Creation Approach." The outcomes include expected codes and words that were highlighted using the attention mechanism.

The BioPortal ontologies were used by the semantic enrichment module to extract the level 5 hierarchies, synonyms, and definitions for the biomedical concepts. After that, the semantic data were combined. I found 736 synonyms for each summary. According to figure 5.4, there were 159 definitions, 562 nodes for Hierarchy 1, 473 nodes for Hierarchy 2, 379 nodes for Hierarchy 3, 327 nodes for Hierarchy 4, and 284 nodes for Hierarchy 5.

On average, the explainable knowledge graph constructed for each summary contained 2900 nodes and 3340 relationships. A graph of this scale can be utilised effectively for reasoning and explanation. To that purpose, the model's performance was assessed using the strength of its connections. The proposed semantic enrichment technique (synonyms, definitions, hierarchy 1 to hierarchy 5) is critical in detecting word-to-word and word-to-code level links or relations. They were generated either by string matching or by considering the semantic relevance of medical ideas.

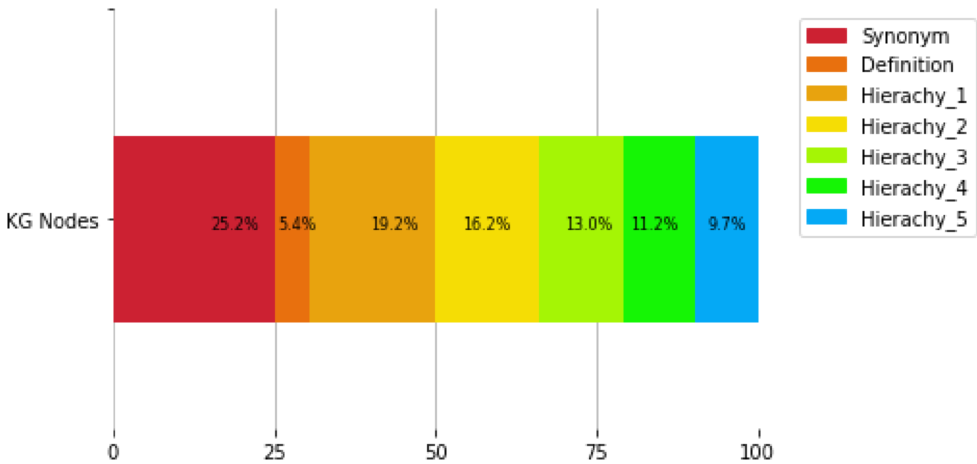

Figure 5.6: Semantic Enrichment Results. The different types of nodes in the Knowledge Graph are Synonyms, Definition, Hierarchy_1 to Hierarchy_5.

The model highlights the words that contributed to the prediction of ICD-10 codes when processing 100 summaries. When I examined these data for word-to-word linkages, I discovered an average of 176 connections based on synonyms and 75 connections based on the definition per summary, as shown in Figure 5.5. Similarly, I discovered that hierarchical levels of semantic information play an important role in word-to-word linkages. Hierarchy 1 generates an average of 39 links or relations, which is much larger than hierarchy 2 through hierarchy 5. Because of the nature of BioPortal Ontologies, which is not important at the word-to-word level, there are no or fewer linkages on some hierarchy levels.

I discovered an average of 52 links based on synonyms and 27 connections based on definitions while examining the word-to-code level interactions. Furthermore, I discovered that hierarchy levels 1 and 2 contribute more to the word-to-code level linkages than hierarchy levels 3, 4, and 5, as shown in Figure 5.5. As a result, the connections are significantly less than at the word-to-word level, although this is to be expected given that the average number of words per discharge summary is 21 but only 4 for code descriptions. Hierarchy 5 revealed 0 links for both levels, but they are included to examine the differences and decrease in the number of connections.

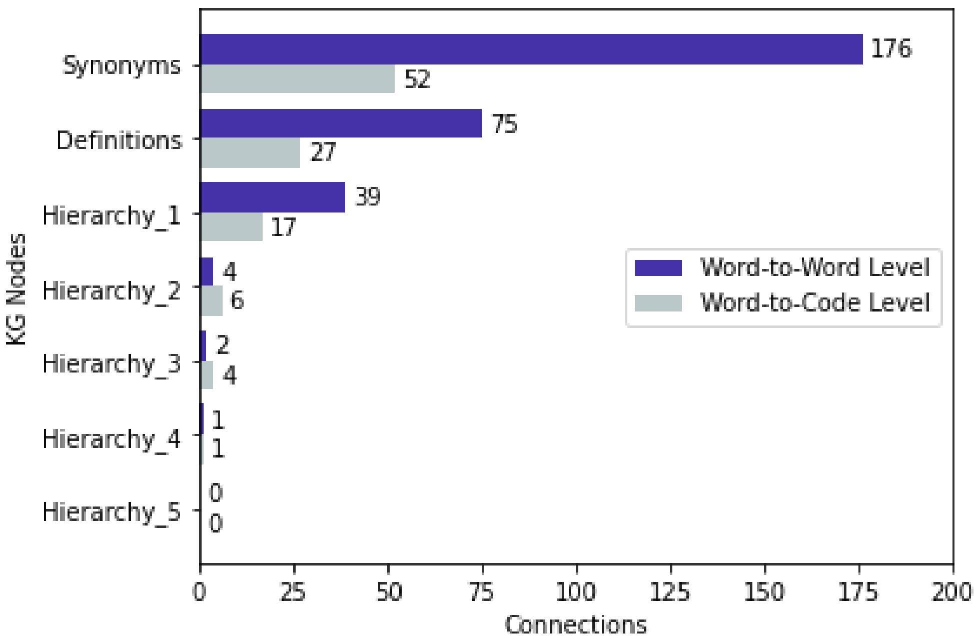

Figure 5.7: Word-to-Word and Word-to-Code Level Connections

Following that, I assessed the deep learning model's performance in terms of predictions and attention mechanism outcomes. The strength of the association between labels-with-words and words-with-words was evaluated. Experts determined whether the ties were weak or strong. More than 30 connections were deemed strong, whereas fewer than

(a) Word-to-Word level connection (b) Word-to-Code level connection

Figure 5.8: Connections Marked by Experts

10 were deemed weak.

The obtained 100 summaries include over 2500 medical ideas. Manual evaluation of their strong and weak connections is not practicable. To get findings, I used a tiny chunk of 40 randomly generated medical codes and roughly 100 words (Nearly 10 summaries). Figures 5.6 and 5.7 show the concepts marked by experts. Out of them, 64 of the 100 words had a correct strong, weak, or no connection with the corresponding words. 21 of the 40 labels included were also correct. Equation 5.4.1 has been used to calculate connection accuracy at various levels, where accuracy equals the number of correct connection instances divided by the total number of relationships. As a result, I assess accuracy in terms of strong and weak links.

$$Accuracy = Correct\_Relations/Total\_Relations \tag{5.4.1}$$

It has been discovered that words can have direct or indirect associations (or linkages). For example, the biomedical notion 'flap' has a direct association with the word 'graft' but an indirect tie with the biomedical concept 'anastomosis'.I divided the count in half for indirect relationships and kept it full for direct relationships. For example, if 'flap' and 'anastomosis' are linked by four nodes, we would consider it as two (in terms of equivalence to direct relationships). The outcomes of our method were entirely contingent on the model output. I expected low accuracy because I trained the model on generic ICD-10 codes rather than unique top 50 labels. In the accuracy calculation, the average number of words was roughly 200.

Chapter 6

Discussion

6.1 Dataset Preparation

One of the most important tools for all researchers, but particularly for computer scientists, is the dataset. Using the developed and suggested methodologies on the datasets, all of their efforts and creative ideas can be demonstrated. We require ICD-10 annotated data because they are currently used by healthcare professionals. Medical coders must invest their time and effort into manually annotating the data, which takes time and effort. We discovered a simpler, more practical solution.

We used a model that had already been trained, and we tested and confirmed its findings using actual medical coders. As a result, we were able to quickly and inexpensively create our own dataset without having to buy it. The incorrect annotations were also removed with the aid of medical code filtration. These things are helpful because they demonstrate that our dataset is accurate and free of errors, which will encourage others to trust it in the future.

We trained our own model on this dataset, and the outcomes were used to support the other Neuro-Symbolic AI methods we had previously suggested. This dataset could be made available for use by other researchers who are encountering similar difficulties due to the dataset's unavailability. This data can be used for many other advantages, as well as the training and validation of many other cutting-edge methods, in addition to our own benefits.

6.2 Model Training

We trained a baseline model for ICD-10 codes. Given that they are currently in use, this is helpful. In addition, we used a general problem domain to train the model. Other trained models have the drawback of having been trained on a particular issue. For instance, training on the top-20 and top-50 codes for specific diseases and code blocks like COVID-19, injuries, heart disease, influenza, etc. Because of this, the model is unable to predict the occurrence of other diseases. This criterion is good for accuracy calculation, but its application to end users is limited.

Since medical coders are the intended users of these applications, they can benefit from them in order to assign the proper codes. Instead of assisting, this specification may limit their ability to do their jobs. Additionally, these applications can be used for claim generation by government agencies, insurance companies, and hospital authorities. They deal with a variety of medical issues on a daily basis, and if these top-20 or top-50 codes are obtained, this will eventually result in the market's rejection of CAC applications.

Our trained model, which was trained on a larger problem domain, is presented here. It has 550 labels, which will aid all pertinent authorities in accepting this application. Our trained model's accuracy and F-1 score are somewhat lower than those of other trained models, but there is a tradeoff between these two factors.

6.3 Automated Knowledge Graph Creation

Knowledge Graphs are semantically enhanced data sources that have numerous benefits in all spheres of human endeavor. These triplet's hub can be used for a variety of things, including fraud detection, entity search, chatbots, knowledge management, drug repurposing, and recommendation systems. Additionally, because they are NOSQL forms of data, these have a number of advantages over conventional relational databases. Despite the fact that they are useful in every aspect of human life, creating them is very difficult. This calls for expertise outside of computer science and the pertinent field. Our method for creating them automatically, with only a few simple steps and no human involvement, is novel and exceptional and has automated this laborious task.

6.3.1 Data Preprocessing

A filtered source of data with links between various nodes is a knowledge graph. We need data preprocessing because we used the automatic graph construction method on discharge summaries, which are unstructured in nature. We used many Natural Language Processing (NLP) approaches, including sentence detection, word detection, stop-word removal, stemming, lemmatization, and N-grams, as shown in Figure 4.3. The outcomes of these were the necessary medical concepts, which are the Knowledge Graph nodes. This stage is advantageous since we are utilizing a machine, saving people the trouble of manually cleaning and structuring the data. Additionally, they provide the data needed for knowledge graphs.

6.3.2 Ontology and Semantic Enrichment

Knowledge graphs are typically mapped to ontologies for semantic augmentation since they are the data source that provides general information. Due to their professional creation, we used BioPortal Ontologies in this instance. This is helpful because we are using data from a reliable source, which will relieve users' worry.

Our method of conceptually mapping to all 1000+ ontologies has advantages and drawbacks of its own. The advantages are that we have a wealth of knowledge that will ultimately assist the users in many ways, but the disadvantage is that the ontologies supply that enriched information in a number of languages, as illustrated in Figure 5.2. This presents a challenge for the formulation of cypher queries as well as potential problems for readers who are unfamiliar with it. However, there are numerous methods to boost and improve different existing technologies using these data sources.

6.3.3 Cypher Query Creation

Knowledge graphs must be created using a certain language and format. We chose to utilise Cypher since Neo4j, a fantastic visualisation tool, uses it. In addition, Neo4j offers access to a number of machine learning and graph embedding algorithms. The cypher query has a certain format, and creating one is challenging. Each node has a unique id that cannot contain spaces, non-English characters, special characters, or numbers as the first character. Our strategy adheres to the format and automates its construction, automating this laborious process for computer professionals.

6.4 Domain-Specific Knowledge Infusion

CAC's applications lie in multi-label classification problems. In artificial intelligence, classification refers to grouping objects into clusters and dividing them with a line or hyperplane. They could be linear, non-linear, or one to many dimensions. As the number of classes or labels increases, the dimensions are increased, and vice versa. This leads to eventually lowering the accuracy of the classification model. That's the reason, that the CAC's models are trained on top-20, top-50, or top-100 codes. Also, the models trained on these restricted domains are unable to perform out of the box.

As mentioned in section 3.2.1, a few strategies have been put out in recent years to address this accuracy issue. Take the term "label interdependence" as an example. We have developed a novel method for semantically enhancing discharge summaries, which will eventually aid the model's comprehension of the context and improve performance. Our method may be used to enhance the classification accuracy of any deep learning or machine learning model and can be applied to any medical application. In order to benefit the relevant users, such as medical coders, healthcare professionals, hospital staff, insurance companies, and governmental entities, this strategy can be used to get around some of the constraints of these AI models.

6.5 Artificial Intelligence Explainable Predictions

The most fascinating and noteworthy inventions of the 20th century were artificial intelligence and its subfields, machine learning and deep learning. All age groups are currently using these AI models, which are applicable to practically all disciplines. Even those who are unaware of artificial intelligence use it in their regular activities. Among the many use cases are face lock, route detection, social media suggestions, number plate detection, etc. Despite their widespread use and acceptance across many industries, they continue to lack credibility in a number of sectors, including healthcare, law, security, human employment, etc. What then is the obstacle? The primary barrier is these AI systems' lack of reasoning abilities. When a person goes to the doctor and is given a diagnosis of a condition, the doctor can explain why; unfortunately, a machine cannot.

The biggest obstacle is Black-box deep learning models' inability to describe how they operate inside and how they produce outcomes.

Researchers have been particularly interested in explainability, and several novel ideas, such the attention mechanism, have been put forth. Although the results are very convincing, certain questions remain. These issues are brought up by the uncertain outcomes of attention mechanisms. The approach makes an attempt to highlight certain words to show how they relate to the output, but it is still unable to explain how the output is related to other highlighted keywords. This calls for the manipulation of deep learning mathematical models, yet symbolic AI may show to be a superior alternative strategy. To discover the relationship between these concepts and predictions, we can employ knowledge graphs that have been semantically enriched.

Our approach of explaining an already explainable system is novel and unique. To the best of our knowledge, no one has ever conducted research in this area. The use of knowledge graphs produced from expert-authored ontologies to describe a deep learning model will assist the relevant agencies to have confidence in these AI applications, which will result in their widespread acceptance. Any attention mechanism or input for medical applications can employ the resonance mechanism we proposed.

Chapter 7

Conclusion and Future Work

7.1 Conclusion

The brain behind Symbolic AI, Knowledge Graphs, has many practical applications and offers many benefits. However, creating them is a laborious and exceedingly difficult task. Although their automation has been worked on previously, their creation in a few simple stages has not been accomplished. We processed the data, used ontology mapping to enrich it, built graphs using query generation, and automated the entire process in a few straightforward steps.

The many drawbacks of deep learning can be overcome by knowledge graphs and vice versa. By using the Medical Coding use case, we offered two distinct NeuroSymbolic AI methodologies to demonstrate the efficacy of hybrid AI. We incorporated input from the CAC model to the semantically enhanced knowledge graphs and examined the outcomes. Using knowledge infusion, we were able to increase the model's accuracy and address a number of its shortcomings.

Similarly, we explained the model predictions using the graphing technique. The method was effectively applied to the results of attention mechanisms, and Symbolic AI was used to provide answers to numerous open-ended queries. The comprehensive graph on the model output for reasoning was produced using the explainable approach, and the graph connections were simplified for end-user understanding. These methods can be used to create graphs, boost accuracy, and produce results that are understandable in any medical application. Additionally, a modification in the relevant ontologies will make them applicable to all domains. To the best of our knowledge, no such research has been done before, and all of our suggested strategies are novel.

7.2 Future Work

NeuroSymbolic AI also known as Explainable AI (XAI) or Hybrid AI is an emerging area of research. The extension could be done in various directions but there are a few things in which we are really interested. Knowledge Graphs in their nature are explainable but are shallow. They can deliver the knowledge they have, but besides that, they fall short. The results of DL models can be improved by their integration with deep learning, however, further research on the intelligence of knowledge graphs needs to be done. There are several ways to do this. Knowledge Graph embedding is one of them. The graphs using the embedding methods can be used to answer out of the box questions in chatbots, can help in intelligent recommendations, work in effective drug repurposing and vice versa.

The approaches we suggested were compatible with the deep learning model's input and output. It was beneficial in many ways. However, Deep Infusion is our goal. Deep infusion refers to a method of addressing the deep learning model's internal operations in order to get around some constraints, clarify internal operations, and reveal black-box models. The intricacy of these deep learning models may be handled in some way as a result of the deep infusion, which will produce extremely reliable applications.

User preferences could be taken into account in the development of programmes that are easy to understand. Our proposed methodologies will be integrated into an EHR application that we are currently creating, giving us a neutral perspective to improve them and the ability to add deep data infusions for improved outcomes..

Chapter 8

Glossary

Table 8.1: Terms used in Thesis.

Term	**Abbreviation**
AMC	Automated Medical Coding
CAC	Computer Assisted Coding
KG	Knowledge Graphs
HAN	Hierarchical Attention Network
HLAN	Hierarchical Label-wise Attention Network
LSTM	Long Short-term Memory
OWL	Web Ontology Language
RDF	Resource Description Framework
API	Application Programming Interface
ICD	International Classification of Diseases
CPT	Current Procedural Therapy
HCPCS	Healthcare Standard Procedure Coding System
NLP	Natural Language Processing
BERT	Bidirectional Encoder Representations from Transformers
EHR	Electronic Health Record

Bibliography

[1] Margaret A Boden. *Artificial intelligence.* Elsevier, 1996.

[2] Sepp Hochreiter and Jürgen Schmidhuber. "Long short-term memory". In: *Neural computation* 9.8 (1997), pp. 1735–1780.

[3] Larry R Medsker and LC Jain. "Recurrent neural networks". In: *Design and Applications* 5 (2001), pp. 64–67.

[4] Alex Graves, Santiago Fernández, and Jürgen Schmidhuber. "Bidirectional LSTM networks for improved phoneme classification and recognition". In: *International conference on artificial neural networks.* Springer. 2005, pp. 799–804.

[5] Natalya F Noy et al. "BioPortal: ontologies and integrated data resources at the click of a mouse". In: *Nucleic acids research* 37.suppl_2 (2009), W170–W173.

[6] Özlem Uzuner et al. "2010 i2b2/VA challenge on concepts, assertions, and relations in clinical text". In: *Journal of the American Medical Informatics Association* 18.5 (2011), pp. 552–556.

[7] Alex Graves. "Long short-term memory". In: *Supervised sequence labelling with recurrent neural networks* (2012), pp. 37–45.

[8] Patricia Aalseth. *Medical Coding: What it is and how it Works.* Jones & Bartlett Publishers, 2014.

[9] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. "Neural machine translation by jointly learning to align and translate". In: *arXiv preprint arXiv:1409.0473* (2014).

[10] Kyunghyun Cho et al. "Learning phrase representations using RNN encoder-decoder for statistical machine translation". In: *arXiv preprint arXiv:1406.1078* (2014).

[11] pdnseek. *Medical Coding History – The Past, Present and Future*. 2015. URL: https://www.pdnseek.com/medical-coding-history-the-past-present-and-future/#:~:text=The%5C%20official%5C%20coding%5C%20of%5C%20diseases,most%5C%20frequent%5C%20causes%5C%20of%5C%20death.. (accessed: 23.01.2023).

[12] Clem Delangue. *Hugging Face – The AI community building the future*. 2016. URL: https://huggingface.co/. (accessed: 01.06.2022).

[13] Alistair EW Johnson et al. “MIMIC-III, a freely accessible critical care database”. In: *Scientific data* 3.1 (2016), pp. 1–9.

[14] Zichao Yang et al. “Hierarchical attention networks for document classification”. In: *Proceedings of the 2016 conference of the North American chapter of the association for computational linguistics: human language technologies*. 2016, pp. 1480–1489.

[15] Alexandra Hofmann et al. “DBkWik: Towards Knowledge Graph Creation from Thousands of Wikis.” In: *ISWC (Posters, Demos & Industry Tracks)*. 2017.

[16] Nick Tate. *4 in 5 Medical Bills Contain Errors: Here’s What You Can Do*. 2017. URL: https://www.newsmax.com/Health/Headline/medical-bill-error-mistake/2017/08/04/id/805882/. (accessed: 01.07.2022).

[17] Finneas Catling, Georgios P Spithourakis, and Sebastian Riedel. “Towards automated clinical coding”. In: *International journal of medical informatics* 120 (2018), pp. 50–61.

[18] Jacob Devlin et al. “Bert: Pre-training of deep bidirectional transformers for language understanding”. In: *arXiv preprint arXiv:1810.04805* (2018).

[19] Nadime Francis et al. “Cypher: An evolving query language for property graphs”. In: *Proceedings of the 2018 international conference on management of data*. 2018, pp. 1433–1445.

[20] James Mullenbach et al. “Explainable prediction of medical codes from clinical text”. In: *arXiv preprint arXiv:1802.05695* (2018).

[21] Heet Sankesara. *Hierarchical Attention Networks*. 2018. URL: https://medium.com/analytics-vidhya/hierarchical-attention-networks-d220318cf87e. (accessed: 3.02.2023).

[22] Gema Hernandez-Ibarburu et al. "ICD-10-PCS extension with ICD-9 procedure codes to support integrated access to clinical legacy data". In: *International journal of medical informatics* 122 (2019), pp. 70–79.

[23] Samaneh Jozashoori and Maria-Esther Vidal. "MapSDI: A scaled-up semantic data integration framework for knowledge graph creation". In: *On the Move to Meaningful Internet Systems: OTM 2019 Conferences: Confederated International Conferences: CoopIS, ODBASE, C&TC 2019, Rhodes, Greece, October 21–25, 2019, Proceedings*. Springer. 2019, pp. 58–75.

[24] Xiang Wang et al. "Explainable reasoning over knowledge graphs for recommendation". In: *Proceedings of the AAAI conference on artificial intelligence*. Vol. 33. 01. 2019, pp. 5329–5336.

[25] Vera Alonso et al. "Problems and barriers during the process of clinical coding: a focus group study of coders' perceptions". In: *Journal of medical systems* 44.3 (2020), pp. 1–8.

[26] Sharon Campbell and Katrina Giadresco. "Computer-assisted clinical coding: A narrative review of the literature on its benefits, limitations, implementation and impact on clinical coding professionals". In: *Health Information Management Journal* 49.1 (2020), pp. 5–18.

[27] Xuqing Chai. "Diagnosis method of thyroid disease combining knowledge graph and deep learning". In: *IEEE Access* 8 (2020), pp. 149787–149795.

[28] Xiaojun Chen, Shengbin Jia, and Yang Xiang. "A review: Knowledge reasoning over knowledge graph". In: *Expert Systems with Applications* 141 (2020), p. 112948.

[29] Gaurav Desai. *gauravkdesai/MIDS-W210-Medical_Insurance_Payment_Assistant*. 2020. URL: https://github.com/gauravkdesai/MIDS-W210-Medical%5C_Insurance%5C_Payment%5C_Assistant. (accessed: 01.05.2022).

[30] Giuseppe Futia and Antonio Vetrò. "On the integration of knowledge graphs into deep learning models for a more comprehensible AI—Three challenges for future research". In: *Information* 11.2 (2020), p. 122.

[31] Kids Health. *Viewing Clinical Notes in Your Child's Electronic Medical Record (for Parents) - Nemours KidsHealth*. 2020. URL: https://kidshealth.org/en/parents/clinical-notes.html#:~:text=Clinical%5C%20notes%5C%20will%

5C%20vary%5C%20depending,symptoms%5C%20and%5C%20when%5C%20they%5C%20started.. (accessed: 01.02.2023).

[32] Samia Khalid. *BERT Explained: A Complete Guide with Theory and Tutorial.* 2020. URL: https://medium.com/@samia.khalid/bert-explained-a-complete-guide-with-theory-and-tutorial-3ac9ebc8fa7c. (accessed: 3.02.2023).

[33] Khalid Mahmood Malik et al. "Automated domain-specific healthcare knowledge graph curation framework: Subarachnoid hemorrhage as phenotype". In: *Expert Systems with Applications* 145 (2020), p. 113120.

[34] Antonio Miranda-Escalada et al. "Overview of Automatic Clinical Coding: Annotations, Guidelines, and Solutions for non-English Clinical Cases at CodiEsp Track of CLEF eHealth 2020." In: *CLEF (Working Notes)* 2020 (2020).

[35] Elias Moons et al. "A comparison of deep learning methods for ICD coding of clinical records". In: *Applied Sciences* 10.15 (2020), p. 5262.

[36] The App Solutions. *EHR vs. EMR vs. PHR Differences: how to provide digital patient records?* 2020. URL: https://theappsolutions.com/blog/development/ehr-vs-emr-vs-phr/. (accessed: 01.02.2023).

[37] Fei Teng et al. "Explainable prediction of medical codes with knowledge graphs". In: *Frontiers in Bioengineering and Biotechnology* 8 (2020), p. 867.

[38] Qingyun Wang et al. "COVID-19 literature knowledge graph construction and drug repurposing report generation". In: *arXiv preprint arXiv:2007.00576* (2020).

[39] Tianxing Wu et al. "Knowledge graph construction from multiple online encyclopedias". In: *World Wide Web* 23 (2020), pp. 2671–2698.

[40] Yikun Xian et al. "CAFE: Coarse-to-fine neural symbolic reasoning for explainable recommendation". In: *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. 2020, pp. 1645–1654.

[41] Biplob Biswas, Thai-Hoang Pham, and Ping Zhang. "Transicd: Transformer based code-wise attention model for explainable icd coding". In: *International Conference on Artificial Intelligence in Medicine*. Springer. 2021, pp. 469–478.

[42] Syed Ahmad Chan Bukhari et al. "LinkedImm: a linked data graph database for integrating immunological data". In: *BMC bioinformatics* 22 (2021), pp. 1–14.

[43] Binjie Cheng et al. "Research on medical knowledge graph for stroke". In: *Journal of Healthcare Engineering* 2021 (2021).

[44] P Deepan and LR Sudha. "Deep Learning Algorithm and Its Applications to IoT and Computer Vision". In: *Artificial Intelligence and IoT*. Springer, 2021, pp. 223–244.

[45] Hang Dong et al. "Explainable automated coding of clinical notes using hierarchical label-wise attention networks and label embedding initialisation". In: *Journal of biomedical informatics* 116 (2021), p. 103728.

[46] Warren J von Eschenbach. "Transparency and the black box problem: Why we do not trust AI". In: *Philosophy & Technology* 34.4 (2021), pp. 1607–1622.

[47] Manas Gaur, Keyur Faldu, and Amit Sheth. "Semantics of the black-box: Can knowledge graphs help make deep learning systems more interpretable and explainable?" In: *IEEE Internet Computing* 25.1 (2021), pp. 51–59.

[48] Ping Gu et al. "Disease Correlation Enhanced Attention Network for ICD Coding". In: *2021 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*. IEEE. 2021, pp. 1325–1330.

[49] Claudio Gutierrez and Juan F Sequeda. "Knowledge graphs". In: *Communications of the ACM* 64.3 (2021), pp. 96–104.

[50] Shuyuan Hu et al. "An explainable CNN approach for medical codes prediction from clinical text". In: *BMC Medical Informatics and Decision Making* 21 (2021), pp. 1–12.

[51] Md Kamruzzaman Sarker et al. "Neuro-symbolic artificial intelligence: Current trends". In: *arXiv preprint arXiv:2105.05330* (2021).

[52] Caiming Zhang and Yang Lu. "Study on artificial intelligence: The state of the art and future prospects". In: *Journal of Industrial Information Integration* 23 (2021), p. 100224.

[53] Usman Ahmed, Jerry Chun-Wei Lin, and Gautam Srivastava. "Hyper-graph-based attention curriculum learning using a lexical algorithm for mental health". In: *Pattern Recognition Letters* 157 (2022), pp. 135–143.

[54] Hang Dong et al. "Automated clinical coding: what, why, and where we are?" In: *NPJ digital medicine* 5.1 (2022), p. 159.

[55] Martin Drancé. "Neuro-Symbolic XAI: Application to Drug Repurposing for Rare Diseases". In: *International Conference on Database Systems for Advanced Applications*. Springer. 2022, pp. 539–543.

[56] Manas Gaur et al. "Knowledge-Infused Learning: A Sweet Spot in Neuro-Symbolic AI". In: *IEEE Internet Computing* 26.4 (2022), pp. 5–11.

[57] Pascal Hitzler et al. "Neuro-symbolic approaches in artificial intelligence". In: *National Science Review* 9.6 (2022), nwac035.

[58] Mutahira Khalid et al. "Explainable Prediction of Medical Codes through Automated Knowledge Graph Curation Framework". In: *2022 19th International Bhurban Conference on Applied Sciences and Technology (IBCAST)*. IEEE. 2022, pp. 331–336.

[59] Leibo Liu et al. "Hierarchical label-wise attention transformer model for explainable ICD coding". In: *Journal of biomedical informatics* 133 (2022), p. 104161.

[60] Haohui Lu et al. "Predictive risk modelling in mental health issues using machine learning on graphs". In: *Australasian Computer Science Week 2022*. 2022, pp. 168–175.

[61] pdnseek. *Clinical coder*. 2022. URL: https://en.wikipedia.org/wiki/Clinical_coder. (accessed: 23.01.2023).

[62] PYPI. *icd10-cm*. 2022. URL: https://pypi.org/project/icd10-cm/. (accessed: 01.05.2022).

[63] Vetle Ryen, Ahmet Soylu, and Dumitru Roman. "Building semantic knowledge graphs from (semi-) structured data: a review". In: *Future Internet* 14.5 (2022), p. 129.

[64] Amit Sheth et al. "Process Knowledge-Infused AI: Toward User-Level Explainability, Interpretability, and Safety". In: *IEEE Internet Computing* 26.5 (2022), pp. 76–84.

[65] Giuseppe Spillo et al. "Knowledge-aware Recommendations Based on Neuro-Symbolic Graph Embeddings and First-Order Logical Rules". In: *Proceedings of the 16th ACM Conference on Recommender Systems*. 2022, pp. 616–621.

[66] Akshat Surolia. *AkshatSurolia/ICD-10-Code-Prediction · Hugging Face*. 2022. URL: https://huggingface.co/AkshatSurolia/ICD-10-Code-Prediction. (accessed: 01.06.2022).

[67] Owen Trigueros et al. "Explainable ICD multi-label classification of EHRs in Spanish with convolutional attention". In: *International Journal of Medical Informatics* 157 (2022), p. 104615.

[68] Farkhanda Zafar et al. "Carpooling in connected and autonomous vehicles: current solutions and future directions". In: *ACM Computing Surveys (CSUR)* 54.10s (2022), pp. 1–36.

[69] Wikipedia. *Electronic health record*. 2023. URL: https://en.wikipedia.org/wiki/Electronic_health_record. (accessed: 01.02.2023).

[70] Wikipedia. *Symbolic artificial intelligence*. 2023. URL: https://en.wikipedia.org/wiki/Symbolic_artificial_intelligence#:~:text=Symbolic%5C%20AI%5C%20was%5C%20the%5C%20dominant,ultimate%5C%20goal%5C%20of%5C%20their%5C%20field.. (accessed: 01.02.2023).

[71] Unbound Medicine. *About ICD-10-CM Coding Guide: ICD-10-CM*. URL: https://www.unboundmedicine.com/icd/view/ICD-10-CM/860000/all/About_ICD_10_CM_Coding_Guide?q=2022. (accessed: 24.01.2023).

[72] Wikipedia. *Deep learning*. URL: https://en.wikipedia.org/wiki/Deep_learning#:~:text=The%5C%20term%5C%20Deep%5C%20Learning%5C%20was,context%5C%20of%5C%20Boolean%5C%20threshold%5C%20neurons.. (accessed: 24.01.2023).