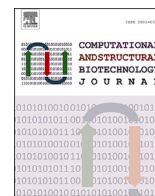




Contents lists available at ScienceDirect

Computational and Structural Biotechnology Journal

journal homepage: www.elsevier.com/locate/csbj

Review Article

From explainable to interpretable deep learning for natural language processing in healthcare: How far from reality?

Guangming Huang^a, Yingya Li^b, Shoaib Jameel^c, Yunfei Long^{a,*}, Giorgos Papanastasiou^{d,*}^a School of Computer Science and Electronic Engineering, University of Essex, Colchester, CO4 3SQ, United Kingdom^b Harvard Medical School and Boston Children's Hospital, Boston, 02115, United States^c Electronics and Computer Science, University of Southampton, Southampton, SO17 1BJ, United Kingdom^d Archimedes Unit, Athena Research Centre, Athens, 15125, Greece

ARTICLE INFO

Keywords:

Explainable
Interpretable
Deep learning
NLP
Healthcare

ABSTRACT

Deep learning (DL) has substantially enhanced natural language processing (NLP) in healthcare research. However, the increasing complexity of DL-based NLP necessitates transparent model interpretability, or at least explainability, for reliable decision-making. This work presents a thorough scoping review of explainable and interpretable DL in healthcare NLP. The term “eXplainable and Interpretable Artificial Intelligence” (XIAI) is introduced to distinguish XAI from IAI. Different models are further categorized based on their functionality (model-, input-, output-based) and scope (local, global). Our analysis shows that attention mechanisms are the most prevalent emerging IAI technique. The use of IAI is growing, distinguishing it from XAI. The major challenges identified are that most XIAI does not explore “global” modelling processes, the lack of best practices, and the lack of systematic evaluation and benchmarks. One important opportunity is to use attention mechanisms to enhance multi-modal XIAI for personalized medicine. Additionally, combining DL with causal logic holds promise. Our discussion encourages the integration of XIAI in Large Language Models (LLMs) and domain-specific smaller models. In conclusion, XIAI adoption in healthcare requires dedicated in-house expertise. Collaboration with domain experts, end-users, and policymakers can lead to ready-to-use XIAI methods across NLP and medical tasks. While challenges exist, XIAI techniques offer a valuable foundation for interpretable NLP algorithms in healthcare.

1. Introduction

Recently, deep learning (DL) has been instrumental in enhancing the performance of natural language processing (NLP) in a wide range of healthcare settings [1–10]. DL has notably advanced the analysis of NLP in areas like medical diagnosis [11,12,7], patient monitoring [13] and drug discovery [14,15].

The evolution of DL models in NLP has led to substantial performance improvements over traditional machine learning methods, with advancements ranging from recurrent neural networks (RNNs) [16] and long short-term memory (LSTM) networks [17] to convolutional neural networks (CNNs) [18] and the latest self-attention mechanisms and transformers [19]. However, these gains in performance come at the cost of increasingly complex model architectures. In healthcare settings, where trustworthy and transparent decision-making is paramount, there is a growing need for interpretable or explainable models [20]. Both

explainability and interpretability pose significant challenges, as understanding how NLP embeddings translate into DL decision-making mechanisms remains complex [21,22]. Fortunately, recent research in explainable [23,24] and interpretable [25,26] DL for NLP in healthcare shows promise. Furthermore, the rise of large language models (LLMs) highlights the growing importance of evaluating which explainable and interpretable methods are most beneficial for healthcare, especially as data and model complexity increase over time.

In our work, the evolution of eXplainable and Interpretable Artificial Intelligence (denoted with the new term “XIAI”) is reviewed [27–30]. Since there is no systematic definition of IAI and XAI [27,28], there is an ongoing inconsistency regarding how these terms are used in DL [29]. To define XIAI in this review, the statistical definition by Rudin [30] is followed: IAI focuses on designing inherently interpretable models; XAI aims to provide post hoc model explanations. Further, XIAI method designs were evaluated, by organizing them into model-based,

* Corresponding author.

E-mail addresses: gh22231@essex.ac.uk (G. Huang), yingya.li@childrens.harvard.edu (Y. Li), m.s.jameel@southampton.ac.uk (S. Jameel), y120051@essex.ac.uk (Y. Long), g.papanastasiou@athenarc.gr (G. Papanastasiou).

<https://doi.org/10.1016/j.csbj.2024.05.004>

Received 10 November 2023; Received in revised form 3 May 2024; Accepted 3 May 2024

Available online 9 May 2024

2001-0370/Crown Copyright © 2024 Published by Elsevier B.V. on behalf of Research Network of Computational and Structural Biotechnology. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

input-based, and output-based approaches. To improve clarity and understanding, XIAI is further grouped based on their scope (local, global): local XIAI yields insights derived from particular inputs, whereas global XIAI grants a wider understanding based on the entire predictive mechanism of the model [20]. Our analysis seeks to enhance clarity and insight into the most effective ways to combine DL and XIAI methods in healthcare NLP by examining and revealing essential architectural designs. Through a quantitative assessment of IAI and XAI in both NLP and medical tasks, we pinpoint current and future opportunities, challenges, advantages, and limitations. In line with the most recent LLM developments, our review evaluates whether attention mechanisms can reinforce the implementation of IAI against XAI. By carefully investigating XIAI DL in healthcare NLP, our review aims to provide insights into the scientific and clinical impact that can potentially be important for XIAI democratization in healthcare NLP.

This review is organised as follows. Initially, the key related work in the domain is reviewed by analysing the important XIAI research trends, methodology designs, paradigms and scope. Subsequently, our review analyses challenges and opportunities for translating XIAI into clinical practice. Finally, we assess the gap between current XIAI capabilities and real-world healthcare implementation, highlighting areas for future research.

2. Methods

2.1. Literature review strategy

A comprehensive scoping review of explainable and interpretable deep learning (DL)-based natural language processing (NLP) techniques applied to healthcare was conducted, considering publications between January 1, 2018, and December 31, 2022. We searched the following databases: Scopus, Web of Science, PubMed, and the Association for Computing Machinery (ACM), following the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) guidelines [31]. To address real-world data needs and assess the potential for the democratization of XIAI methods, our review focused on electronic health record (EHR) and electronic medical record (EMR) NLP data. The emergence of Transformers as an effective approach to learning long-range interactions via attention modules in 2017 [19] marked the beginning of a new era in NLP, with Transformer variants becoming the most prevalent and promising techniques [32]. To focus our review on these significant developments, our observation window was selected to start in 2018. Fig. 1(a) illustrates the PRISMA flow.

Initial Filtering: To broaden the search, publications were initially retrieved using the following keywords in the title, abstract, and manuscript keywords: (“*natural language processing*” OR “*NLP*”) AND (“*healthcare*” OR “*medic*” OR “*clinic*” OR “*EHR*” OR “*EMR*”). This led to a preliminary list of 4,147 papers. Subsequently, to refine the focus on XIAI methods, irrelevant papers were removed, resulting in 234 papers. In the subsequent step, removing duplicate records across databases resulted in 146 unique papers selected for title and abstract screening.

Title and Abstract Screening: All authors independently reviewed the titles and abstracts of all 146 papers. Articles demonstrably irrelevant to the field of study were excluded, resulting in 110 papers eligible for full-text review.

Full-Text Review: During the full-text review, 68 papers were excluded for the following reasons:

- 25 papers did not utilize DL techniques.
- 5 papers were not relevant to XIAI.
- 4 papers lacked a medical focus.
- 21 papers were review articles.
- 13 papers fell outside the scope of DL, NLP, or healthcare.

This resulted in a final selection of 42 journal and conference papers for our review analysis.

2.2. Review aspects

Full-Text Evaluation: During the full-text review, we focused on evaluating the following aspects of each paper:

1. **Publication Date:** This provided context for the development of the methods.
2. **DL Model:** The specific DL model employed was identified (see Results for DL model categorization).
3. **IAI vs. XAI:** The paper’s focus on interpretability or explainability was distinguished.
4. **Scope of XIAI Method:** We categorized the XIAI method as local or global based on its ability to inform about the model predictions regarding a specific input, or the entire modelling process (end-to-end), respectively [20].
5. **XIAI Method Type:** The type of XIAI method used (model-based, input-based, or output-based) was identified.
6. **Medical Task:** The specific medical task addressed by the XIAI approach was categorized.
7. **NLP Task:** The type of NLP task involved (e.g., classification, information extraction) was identified.
8. **Data Source:** We categorized the data source used (public or private data).
9. **Source Code Availability:** The availability of the source code for the method was assessed.
10. **XIAI Evaluation:** The presence or absence of an evaluation for the XAI method was noted.

Furthermore, we analyzed the XIAI methods based on three key aspects to identify challenges, opportunities, advantages, and drawbacks:

XAI and IAI methods with DL Models, NLP Tasks, and Medical Tasks: This analysis explored how XAI and IAI methods were combined with a) the specific DL model used, b) the type of NLP task analyzed, and c) the medical task addressed.

3. Results

The results section is organised as follows: first the research trends are detailed, then the XAI against IAI methods, the scope of XIAI methods, the three XIAI method paradigms (i.e., model-, input- and output-based methods), the DL models explained/interpreted and the NLP and medical tasks addressed. Further, the evaluation metrics of XIAI methods and open code are described. Finally, the dataset sources (public and/or private) are assessed.

3.1. Research trends

Published work on XAI and IAI NLP in healthcare was reviewed. A rapid increase in the number of studies was observed from 2021 to 2022, compared to the first years of our observation window (Fig. 1(b)). This trend coincides with the dominance of Transformer DL models (Fig. 1(b)). In addition, this trend reflects the rapid rise of attention mechanisms, which have rapidly overtaken other XIAI methods during this time (Fig. 1(d–e)).

3.2. IAI against XAI

All methods across papers were examined to assess whether they focused on IAI or XAI. More studies involved IAI methods ($n = 29$) than XAI methods ($n = 15$), with this trend being evident in the period 2020–2022. Fig. 1(d) illustrates the publication record of XAI methods used per year.

Among XAI methods, t-SNE ($n = 4$) was the most widely used, followed by LIME ($n = 3$), SHAP ($n = 3$) and other methods ($n = 5$) such as maximum aggregator (MaXi) [33], SKET eXplained (SKET X) [34], syntax tree-guided semantic explanation (STEP) [35], evidence-based [36] and sentiment intensity score [37].

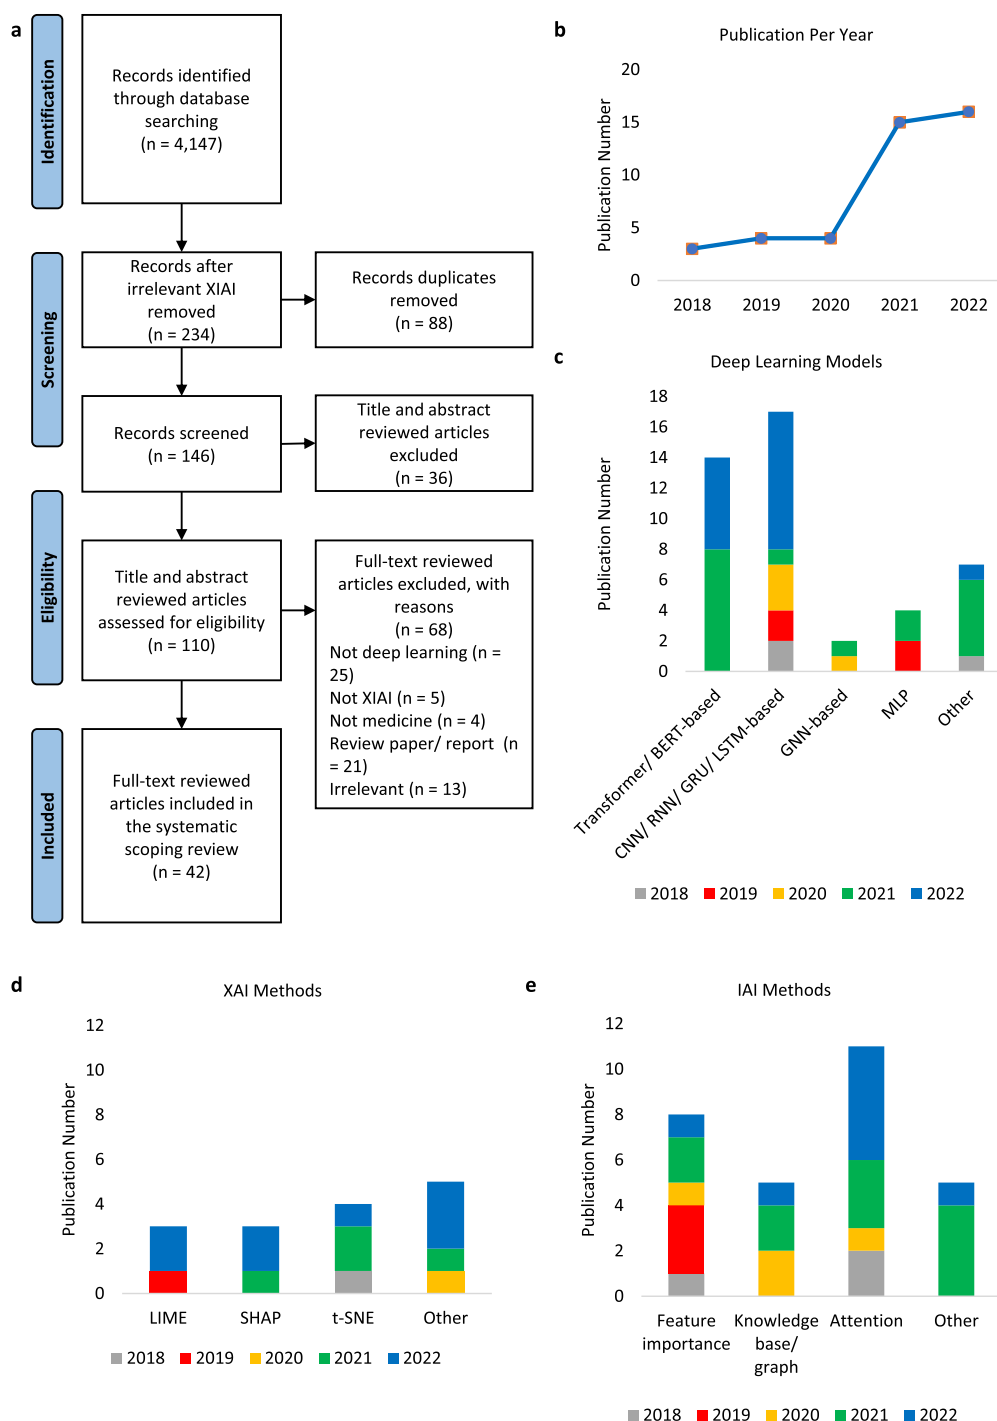


Fig. 1. (a) PRISMA flowchart. The flowchart presents the inclusion and exclusion of papers at each review stage. (b) Publication records of XIAI, (c) main DL model category, (d) main XAI and (e) main IAI methods, per year.

As illustrated in Fig. 1(e), most IAI methods involved attention mechanisms ($n = 11$), followed by feature importance ($n = 8$), knowledge base/graph ($n = 5$) and other methods ($n = 5$), including causal graphs [38], logistic regression-based parametric predictions [39], opinion aggregator [40], case-based reasoning [41] and interactive classification [42].

3.3. Scope of XIAI methods

XIAI method scopes were further grouped into local and global [20]. The majority of papers employed local methods ($n = 37$), with only $n = 5$ studies involving global XIAI (Table 1).

3.4. Three XIAI method paradigms

Three distinct XIAI categories were introduced: model-, input- and output-based methods. These definitions were conceptualized based on whether an XIAI method relies on internal/external modules (model) to perform XIAI, measures how the input features affect model decisions or explains/interprets model behaviour through analyzing prediction outcomes, respectively. Table 1 presents a comprehensive overview of sixteen different XIAI approaches identified in our review and their distinct categories (model-, input-, and output-based methods). There were eighteen, thirteen and thirteen model-, input- and output-based methods identified, respectively.

Table 1
Explainable and interpretable artificial intelligence (XIAI) methods, their category, and scope.

XIAI category	XIAI methods	Explainable/Interpretable	Local/Global	Articles
Model-based	Causal graph	Interpretable	Global	[38]
	Logistic regression-based parametric predictor	Interpretable	Global	[39]
	Opinion aggregator	Interpretable	Global	[40]
	Case-based reasoning	Interpretable	Global	[41]
	Interactive classification	Interpretable	Global	[42]
	LIME	Explainable	Local	[43,23,44]
	MAXi	Explainable	Local	[33]
	SHAP	Explainable	Local	[24,45,46]
	SKET X	Explainable	Local	[34]
	STEP	Explainable	Local	[35]
t-SNE	Explainable	Local	[47–50]	
Input-based	Feature importance	Interpretable	Local	[51–55,48,56,57]
	Knowledge base/graph	Interpretable	Local	[58–61,25]
Output-based	Attention	Interpretable	Local	[26,25,62–70]
	Evidence-based	Explainable	Local	[36]
	Sentiment intensity score	Explainable	Local	[37]

According to Table 1, the most popular model-based XIAI was t-SNE ($n = 4$), followed by LIME ($n = 3$), SHAP ($n = 3$) and other methods that have been sparsely used. Among input-based XIAI, studies either used feature importance-based ($n = 8$) or knowledge graph-based ($n = 5$) methods. Regarding output-based methods, the attention mechanism was the most frequently used approach (Table 1). While benefiting from their performance advancements, attention mechanisms may therefore be an important IAI approach.

3.4.1. Model-based XIAI

Model-based XIAI focuses on describing how DL functions through the use of internal or external modules. The following section presents the most promising XAI and IAI methods that we have identified.

LIME. Ribeiro et al. [71] introduced the local interpretable model-agnostic explanation (LIME) method. To produce explanations, LIME perturbs an instance, generates neighbouring data, and learns linear models within that neighbourhood. For instance Zhang et al. [43] analyzed physician ordering for magnetic resonance imaging (MRI) scans and corresponding radiology reports (outcomes), by developing an RNN model integrated with a LIME module. The authors showed that it is possible to justify whether a patient should undergo MRI or not and if any brain abnormalities are likely to be detected, based on certain keywords from the physician order text entries.

SHAP. Shapley Additive exPlanations (SHAP) is an algorithm used to obtain local, post-hoc explanations [72]. The SHAP method is a model agnostic approach, that allows to elucidate how individual input features contribute to model decisions. The method is based on deriving Shapley values from cooperative game theory [73], which aims to quantify the contribution of each input, both individually as well as collectively in combination with all other inputs. Naseem et al. [46] introduced a Transformer-based encoder with a SHAP module for suicide risk identification. The SHAP module collected the contribution of important words that were most relevant to mental health and suicide.

t-SNE. t-SNE can effectively capture local interactions within data, although it is less effective in identifying long-range interactions [74]. Using t-SNE, it is possible to visualize local clusters consisting of specific word embeddings, towards enhancing model explainability [47–50].

Although less explored, model-based techniques may be promising for IAI (Table 1). Lu et al. [40] employed an opinion aggregator in combination with BERT, which involved k-means clustering to disentangle responses for COVID-19 questions. Another study developed a Bayesian network-derived causal graph to identify causal relationships between NLP variables extracted from mammography reports [38]. The authors fed the causal graph into a TabNet model and showed that they could improve model interpretability. Causal modelling can be a promising technique to perform global IAI.

3.4.2. Input-based XIAI

Input-based XIAI methods focus on understanding how specific inputs influence network decisions. All input-based methods identified were either under the feature importance or knowledge base/graphs categories.

Feature Importance. Lindsay et al. [51] extracted semantic, syntactic and paralinguistic features using NLP from spontaneous picture descriptions to identify cognitive impairments for Alzheimer’s disease (AD) detection. The authors quantified statistically significant (“generalizable”) NLP features, across two different languages. Subsequently, they performed binary (AD vs healthy) patient classification using DL and showed more accurate results when the generalizable features were used, against when all features were considered. Holderness et al. [54] integrated clinically relevant multiword expressions during preprocessing which improved the accuracy and interpretability of their DL models (multilayer perceptrons and radial basis function DL) in identifying psychosis patients at risk of hospital readmission. Mellado et al. [55] engineered clinically interpretable NLP-based features from topic extraction and clinical sentiment analysis to predict early readmission risk in psychiatry patients. In a different setting, Xue and Chuah [56] designed a knowledge extraction framework from heterogeneous medical data sources, capable of generating an aggregated dataset to characterize diseases. Based on this dataset, they introduced an end-to-end DL-based medical multi-disease diagnosis system.

Knowledge graphs. Knowledge graphs (KGs) are employed to structure knowledge representations from either internal or external sources [75–77]. Teng et al. [25] introduced an end-to-end KG-based framework for diagnosis code predictions. Gu et al. [61] proposed a novel method to solve complex question-answering from a medical dataset by combining KG information. The authors used Wikipedia as an external text source to extract documents related to questions and then mined triples (subject, predicate, object) from the documents to construct a KG. They subsequently extracted evidence graphs from the KG and embedded them into an Attention-based Graph Neural Network to perform complex question answering.

3.4.3. Output-based XIAI

Output-based XIAI methods focus on explaining/interpreting DL models by unravelling how internal computations within DL converge to output decisions.

Attention Mechanisms. The attention mechanism was initially designed to learn long-range interactions in NLP and improve machine translation [78]. The attention mechanism allows to search for a set of positions in a source sentence and encourages the model to predict a target word based on the context vectors associated with these source positions and all the previously generated target words [78].

Following attention, the self-attention mechanism was designed to further enhance the modelling of long-range interactions [79]. Within a self-attention layer, each position (known as query, key, and value) can attend to all positions in the output of the previous layer. Attention and self-attention mechanisms are the building blocks of Transformers. Transformer models have been recently applied in biomedical NLP [65,66] and image analysis [80]. Next to performance gains, attention mechanisms and Transformers have demonstrated their applicability in enhancing interpretability [81]. Attention mechanisms were the most widely used XIAI techniques in our review.

Dong et al. [26] designed a Hierarchical Label-wise Attention Network to perform medical coding. The proposed model was able to provide interpretations in the form of attention weights quantified from both words and sentences. Duarte et al. [62] engineered attention mechanisms inside a neural network to derive death certifications through the analysis of diagnosis codes. Their method produced accurate and interpretable results and can potentially be transferrable to clinical settings. Zhang et al. [63] incorporated a series of self-attention modules into an RNN architecture. They showed accurate and interpretable predictions of patients at risk for all-cause hospitalizations through the analysis of EHR data. Chen et al. [65] trained BERT-based models via federated learning for ICD-10 coding. To interpret the model outputs, the authors added a label attention mechanism to the BERT model, which was able to display federated learning outcomes.

3.5. Deep learning

Table 2 presents all DL models identified across papers. DL models were classified as 1) CNN/RNN/GRU/LSTM-based models; 2) Transformer/BERT-based models, featuring pre-trained language models; 3) Multilayer Perceptrons (MLPs); 4) Graph neural networks (GNNs) and; 5) other which includes models used once like unsupervised learning [39] and stacked denoising autoencoder [48].

The publication trends across DL models were examined (Fig. 1(c)). Figs. 2(a) and (b) illustrate DL model combinations with XAI and IAI, respectively. CNN/RNN/GRU/LSTM-based models were the most frequently used ($n = 17$), followed by Transformer/BERT-based ($n = 14$), MLP ($n = 4$), GNN-based ($n = 2$) and other models ($n = 7$). Considering the period 2021–2022, Transformer/BERT-based models were the most dominant backbone architecture used. Although Transformer-based techniques were introduced in 2017 [19], our analysis shows that their adoption in the healthcare domain was not immediately adopted.

3.6. XIAI and NLP

As shown in Table 2, seven NLP tasks were identified: 1) text classification (which includes sentiment analysis [55]), 2) information extraction, including knowledge extraction [34,56] and relation extraction [60], 3) named entity recognition, 4) question answering, 5) document clustering, 6) information retrieval and 7) topic extraction. Text classification was the main NLP task ($n = 35$), followed by information extraction ($n = 5$), named entity recognition ($n = 4$) and other 4 tasks ($n = 5$).

Fig. 2(c–d) presents NLP tasks against XAI and IAI, respectively. Of note, IAI methods have been employed across all NLP tasks defined. In contrast, XAI methods have been exclusively used for text classification and information extraction.

3.7. XIAI and medical tasks

Medical tasks across all studies were extracted and categorized into sixteen categories (see Table 2). Fig. 3 illustrates the primary medical tasks, including medical diagnosis ($n = 10$), medical coding ($n = 6$), medical decision making ($n = 6$), mental health diagnosis ($n = 5$), disease classification ($n = 3$) and other ($n = 12$).

Fig. 3(a–b) illustrates all medical tasks addressed using XAI and IAI methods, respectively. IAI methods were broadly applied to a more diversified set of medical tasks, against XAI methods which were mainly utilized in mental health diagnosis, medical diagnosis and decision making.

3.8. Evaluation metrics and open code

In our review, only 3 out of all 42 articles involved dedicated evaluation processes and metrics of XIAI methods. These studies either involved qualitative evaluations by developers [26] and end-users (i.e., clinicians) [25], or developed a quantitative explainability index [35].

Limited code availability (only 5 out of 42 studies provided code access) was identified across publications, which restricts the ability to assess if these XIAI methods can be replicated on public or future benchmark datasets [39,26,52,58,50].

3.9. Datasets

Finally, methods were categorized based on whether they were developed using *public datasets* (open source and accessible), or *private datasets* (otherwise). According to Table 2, the majority of papers ($n = 32$ out of 42 papers) used private data and only 10 papers analyzed public datasets. EHR from hospitals was the most prevalent type of private data analyzed (Table 2). The most widely used public dataset was MIMIC ($n = 5$), which is a large clinical database consisting of multi-modal (text and imaging) data (Table 2).

4. Discussion

4.1. Challenges and opportunities for clinical translation

In this section we discuss the key opportunities and challenges surrounding the clinical translation of XIAI. XIAI presents significant opportunities such as leveraging attention mechanisms to combine and interpret multi-modal information (text, images, genetic data, clinical history) for personalized medicine and combining DL with causal modelling towards further enhancing inherent interpretability. However, our findings also suggest that all these methods require strong in-house technical expertise to infer XIAI. Other key challenges include the lack of established best practices for XIAI selection based on data and problem type, as well as the unmet need for systematic evaluation methods and high-quality benchmarks. In the following paragraphs, we expand on these challenges and opportunities, identified from the perspective of clinical translation.

Based on our literature review, there is no previous survey focusing on these topics. Considering their early development phase, more extensive studies are needed in the future to develop XIAI methods that will allow “global” (end-to-end) interpretations that go beyond visualization maps, devise more transparent XIAI that will require less technical oversight and design XIAI evaluation metrics that are critical to establish best practices for XIAI selection. The following points can therefore guide future work and systematic reviews towards strengthening and democratizing XIAI in healthcare NLP.

Challenges: The majority of publications focused on local XIAI, with only 5 papers involving a global XIAI approach [38–42]. Moving beyond local XIAI (which provides relatively limited insights based on specific inputs or features) towards “globally” enhancing the transparency of the entire modelling process might be necessary to reduce the requirement of in-house technical expertise and to develop readily translational XIAI methods for end-users.

In our review, the evaluation of XIAI methods was substantially limited (only 3 papers involved XIAI evaluation metrics). Developing robust evaluation metrics for XIAI is challenging given the limited availability of ground truths in healthcare data [80] and the complexity of

Table 2

All the articles grouped based on the XIAI method along with the associated medical task, NLP task and the dataset used for.

DL Model	Medical task category	Medical task	NLP task	Dataset	Article
Transformer/BERT-based	Medical entity disambiguation	Biomedical named entity disambiguation	NER	ClinWikiNED	[52]
	Medical decision making	Chinese medicine instruction parsing	NER;	Private dataset	[60]
	Medical bias mitigation	Reducing genderization	Information extraction	MIMIC-III	[49]
	Medical decision making	Drug indication classification	Text classification	curated drug indication corpus	[50]
	Medical research	Accelerate COVID-19 research	Question Answering	Private dataset	[40]
	Management of health services	Automatically scoring request for proposals (RFP)	Information extraction;	Private dataset	[45]
	Medical note generation	Radiology report generation assistance	Text classification		
	Medical coding	Medical coding	NER	MIMIC-CXR	[41]
	Medical diagnosis	Long-term COVID effects	Text classification	Private dataset	[65]
	Medical fake news detection	Fake news detector to covid-19	Text classification	Private dataset	[66]
	Mental health diagnosis	Suicide risk identification	Text classification	Private dataset	[36]
	Medical diagnosis	Chronic cough	Text classification	Private dataset	[46]
	Medical diagnosis	Alzheimer's disease	Text classification	Private dataset	[69]
	Mental health diagnosis	Depression symptoms detection (DSD)	Text classification	Private dataset	[70]
CNN/RNN/GRU/LSTM-based	Medical coding	Medical coding	Text classification	MIMIC-III	[25]
	Medical diagnosis	Ischemic stroke	Text classification	Private dataset	[53]
	Disease classification	Predict adherence to ACR guidelines	Text classification	Private dataset	[43]
	Medical coding	Causes of death	Text classification	Private dataset	[62]
	Medical diagnosis	Diagnosis	Text classification	Private dataset	[63]
	Medical decision making	Entrustment assessment	Text classification	Private dataset	[57]
	Medical diagnosis	Cancer pathology	Text classification	Private dataset	[33]
	Survival detection	Discrete-time survival estimates	Text classification	Private dataset	[24]
	Medical coding	ICD multi-label classification	Text classification	MIMIC-III;	[64]
				Spanish Osa Corpus (Private dataset)	
	Management of health services	Automatically scoring request for proposals (RFP)	Information extraction;	Private dataset	[45]
	Medical diagnosis	Disease diagnosis	Text classification	Private dataset	[56]
	Monitor public sentiments during pandemics	COVID-19 pandemic	Information extraction;	Private dataset	[37]
	Telehealth services	Telehealth services	Text classification	Private dataset	[23]
Mental health diagnosis	Internet-delivered psychological treatment (IDPT)	Text classification	Private dataset	[67]	
Mental health diagnosis	Depression	Text classification	Amazon Mechanical Turk	[68]	
Mental health diagnosis	Depression	Text classification	Private dataset	[44]	
Disease classification	Diabetes-related biomedical literature exploration	Document clustering;	Private dataset	[42]	
		Text classification			
GNN-based	Medical diagnosis	Breast cancer	Text classification	Private dataset	[38]
	Medical question-answering	Medical QA	Question Answering	Head-QA	[61]
MLP	Medical diagnosis	Language impairment in Alzheimer's disease	Text classification	Private dataset	[51]
	Readmission prediction	Psychiatric readmission prediction	Text classification	Private dataset	[54]
	Readmission prediction	Psychiatric readmission risk prediction	Text classification;	Private dataset	[55]
		Topic extraction			
Medical diagnosis	Alzheimer's disease	Text classification	Private dataset	[70]	
Other	Medical coding	Patient medical conditions	NER;	Private dataset	[39]
	Medical decision making	Clustering of patients' visits	Text classification	Private dataset	[47]
	Medical coding	Medical coding	Text classification	MIMIC-III-50;	[26]
				MIMIC-III-shielding;	
	Medical decision making	Esophageal achalasia	Information retrieval	MIMIC-III	[58]
	Disease classification	Cardiovascular diseases	Text classification	OHSUMED	[59]
Medical decision making	Patient mortality prediction; Primary diagnostic category prediction; Primary procedural category prediction; Gender prediction	Text classification	Private dataset	[48]	
Medical diagnosis	Pathology	Information extraction	Private dataset	[34]	

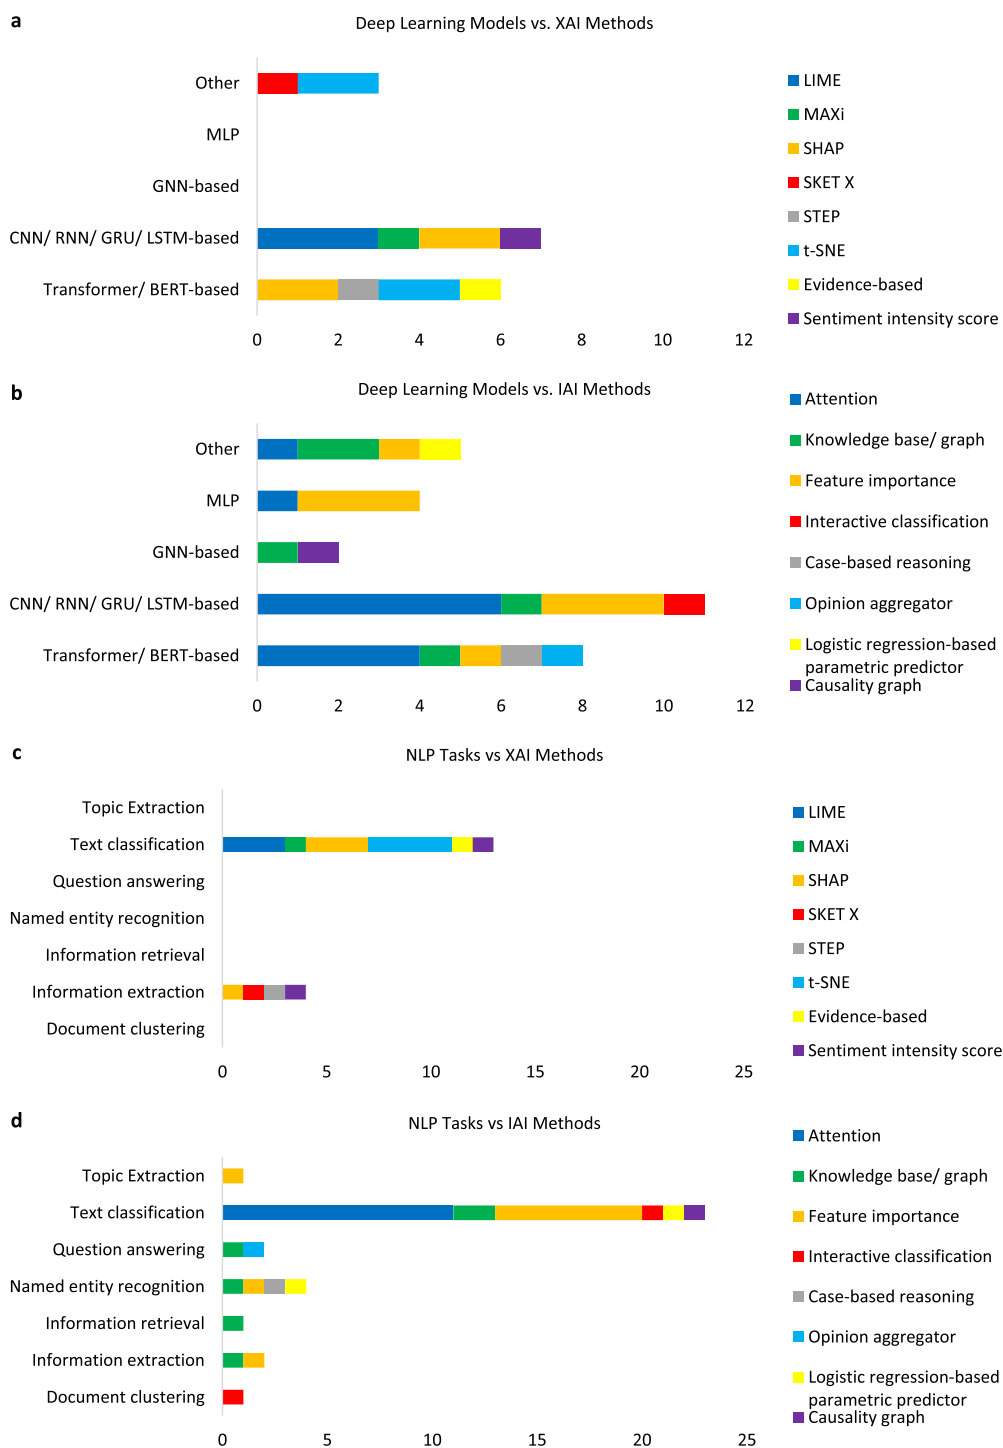


Fig. 2. Publication record showing combinations between (a) deep learning model and XAI method, (b) deep learning model and IAI method, (c) NLP task and XAI method and (d) NLP task and IAI method.

rapidly evolving models (like LLMs) [82]. The majority of papers analyzed private data (32 papers) and did not provide source code (37 papers). Legal implications around private data pose another critical challenge in terms of democratizing data and assessing reproducibility for these XIAI techniques. Hence, it is particularly challenging to assess the reproducibility of these XIAI NLP methods. Significant variability in the literature was observed across model-, input- and output-based methodologies used (Table 1) and NLP and medical tasks addressed (Figs. 2 and 3). Each of these method types presents different advantages and drawbacks (see subsection 4.2). It is therefore substantially

challenging to develop generalization frameworks regarding which XIAI method is most optimal across different NLP and medical applications. In that context, there is an unmet yet pivotal need to develop benchmark data/metrics to extensively evaluate XIAI methods across different NLP and medical tasks, with associated open-access codes.

While attention-based IAI holds promise, it primarily relies on visualizing important attention weights as heat maps [78,79]. Heat maps may reveal patterns in the way the model prioritizes specific types of words or grammatical structures, providing some insight into which parts of the processed text are most influential. However, heat maps

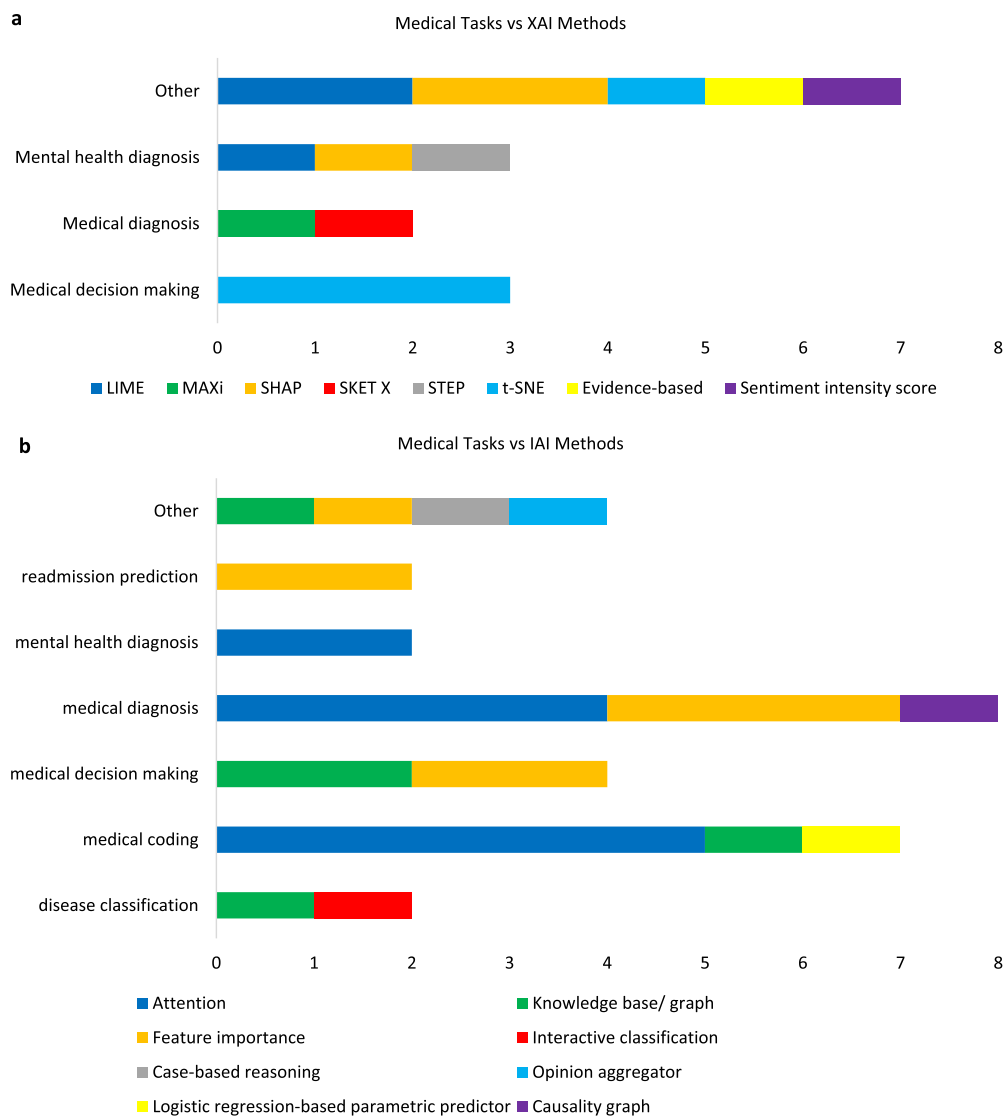


Fig. 3. Publication record showing combinations between medical task and (a) XAI and (b) IAI method.

highlight what the model “looks at,” but not necessarily how it interprets that information. Moreover, attention weights do not linearly correlate with model outcomes [21,70]. These two main limitations currently compromise attention-based IAI. The adoption of current XIAI in healthcare and industry systems is therefore challenging, given the limited access to global XIAI techniques and the absence of robust XIAI metrics, ground truths and benchmark data/studies, which form an important barrier to their systematization.

XIAI systematization and deployment necessitate further thorough work. A possible practical mitigation to accelerate the systematization of robust and transparent XIAI developments for NLP in healthcare is to bring “humans into the DL loop” [83,84]: domain experts, end-users, policymakers and patients will be able to contribute to the XIAI method design, development and evaluation. This collective approach can potentially lead to the emergence of robust and ready-to-use XIAI methods across different NLP and medical tasks.

Opportunities: Attention mechanisms were the most diversified XIAI techniques in terms of DL used (Fig. 2 a, b) as well as NLP (Fig. 2 c, d) and medical tasks (Fig. 3) addressed. It is noteworthy that although Transformers were the dominant DL backbones in the period 2021–2022, individual attention mechanisms were the most frequently used IAI technique. Attention and self-attention mechanisms have also recently been used in combination with CNN models for medical image

analysis, as lighter alternatives of Transformers [80]. This demonstrates their diverse applicability across data domains and tasks. A major opportunity identified is the versatility of attention mechanisms to be combined with multiple DL structures (e.g., CNN [64], RNN [63], BERT [65] and full Transformers) [70]. Another translational opportunity identified via the use of attention mechanisms is to simultaneously model and interpret information from variable multi-modal data (e.g., text, images, genetics, clinical history) [80]. Based on the previous, developing multi-modal XIAI can potentially support personalized medicine which benefits from combining patient-level information from multiple sources [85]. Despite the challenges associated with attention-based heat maps described above, combining information from multiple sources e.g., images and text, can potentially enhance interpretability. A characteristic example are DL models developed for MIMIC data analysis, in which features extracted from X-ray images and radiology reports are combined to create a unified representation that fuses information from both modalities [25,26,41,47,49]. Attention heat maps can highlight which parts of the image and text were most influential in the model decision. Combining image and text data can offer insights into a model’s decision-making that go beyond the limitations of single-modal (text-only) attention heat maps.

A clear trend towards IAI against XAI was observed across studies, particularly in the period 2021–2022. While our review highlights

the benefits of using IAI against XAI, their combination can be useful as an auxiliary assessment to cross-evaluate IAI and XAI outcomes. For instance, the SHAP method that can provide both global and local explanations [72], has been effectively combined with a Transformer encoder for suicide risk prediction [46]. Although this work focused on the SHAP method to perform XAI, future work could aim to also co-extract IAI information from specific attention mechanisms within Transformers.

Causality is an emerging topic in DL which aims to improve model interpretability, fairness and generalization [86,87]. The fundamental aim of causal DL is to unravel causal relationships between variables which determine the model's decision-making process [86]. Understanding how the data are causally related can help us design better DL models. This leads to more reliable predictions and a deeper understanding of how our models work [88]. Next to identifying causal paths between data, central to causality is also to understand the relationships between cause and effect [87]. Understanding cause and effect is crucial for many important decisions. In clinical trials for example, doctors need to know if a new drug actually improves patient outcomes (not just whether there is a statistical correlation). In our review, a Bayesian Network-derived causal graph has been fed into a TabNet, showing promising IAI results [38]. Further work in this field can significantly enhance IAI in NLP for healthcare. We endeavour to inspire and guide relevant benchmark studies to thoroughly examine XIAI in terms of strengthening NLP applications in healthcare. For example, leveraging causal mechanisms to better understand how foundation models such as Transformers interact with data or prompts, may be a critical path forward.

4.2. The three XIAI paradigms

Model-based XIAI methods, such as SHAP [72], LIME [71], and t-SNE [74], offer important advantages by providing both global and local DL interpretability options. These methods are designed to be transparent and accessible as ready-to-use entities thus, potentially being able to democratize XIAI tools for a wide audience [24]. Intuitive visualizations can to some extent elucidate complex model predictions, enhancing trust and facilitating the communication of findings. However, they also have drawbacks, such as the instability of methods like LIME, which is prone to minor data variations that can affect XAI reliability [30,70]. Moreover, relying on separate modules for XIAI may be subject to misinterpretations, in case these modules have inherent biases or limitations. Hence, despite their ease of use, they require technical expertise to evaluate any biases or implementation incompatibilities.

Input-based XIAI methods allow for DL model interpretation by leveraging specific input feature importance [51–53] and medical KG [58–60], to discern how different inputs influence model decisions. These methods are intuitive and accessible to computer scientists and NLP researchers, due to their relative ease of being combined with DL architectures. In that context, the incorporation of “important features” or medical KGs into AI can enrich models with domain-specific insights, which in turn allows them to interpret even complex medical concepts that are common in clinical practice [58].

Input methods based on feature importance requires the incorporation of techniques that are not readily transparent to non-technical medical professionals (end-users). Hence, it may not always be straightforward for end-users to understand “how an important feature was derived and/or evaluated” [51,54,55]. The effectiveness of KG methods relies on medical professionals with expertise and capacity to develop comprehensive medical KG, posing a challenge if such expertise is scarce. Integrating medical KG into DL presents challenges in terms of ontology construction as well as knowledge extraction [25], which are labour-intensive techniques.

Output-based XIAI methods, are important to interpret DL outputs and can offer computational insights through mechanisms like

attention [81]. Output-based methods focus on explaining/interpreting DL models by uncovering how internal computations within DL converge to output decisions, which can be mainly useful to computer scientists/modellers. Nevertheless, attention mechanism-based IAI faces debates [89–93], as high attention weights don't necessarily linearly correlate with model predictions [21]. This can lead to ambiguity while emphasizes the need for further research on IAI methods and their evaluation.

4.3. Inspired by the future

Go Large: Recently, LLMs have attracted significant attention in AI [94–96]. The intersection of LLMs and healthcare creates unique opportunities towards designing future studies, from drug discovery to personalized diagnosis and treatment [97–99]. Healthcare data analysis is one of the high gain-high risk domains for LLMs [82]. One of the limitations of LLMs is that they sometimes tend to “hallucinate” results [94–96]. This can add considerable barriers to their utilization in healthcare settings [97–99]. As discussed in subsection 4, causal inference can be an effective solution towards enhancing IAI [86]. Although designing causal logic inside LLMs is challenging due to their architectural complexity and the fact that models are already “trained” on a causal-agnostic mode, there have been recent attempts which aim to develop causal reasoning between prompts and responses [100,101].

Go Small: An ongoing discussion in the community is whether LLMs or domain-specific smaller models can be more robust solutions for healthcare data [102]. Recently, smaller-parameter domain-specific LMs have outperformed larger LMs when examined on clinical notes from large public databases (MIMIC) [103]. This approach has several possible benefits: a) small-parameter LMs can be trained using in-house computational capabilities, which b) minimizes the risks associated with transmitting sensitive patient data to cloud-based or external servers (thus, adhering to privacy regulations). Moreover, c) causal DL techniques can be more optimally designed and validated since they can be trained on domain-specific data from scratch.

5. Conclusions

There is a growing trend of using IAI techniques alongside XAI, with attention mechanisms emerging as the most dominant approach for interpretation. Our analysis reveals a major opportunity: leveraging attention mechanisms to combine and interpret multi-modal data (text, images, genetics, etc.), potentially leading to advances in XIAI approaches for personalized medicine solutions [80]. Another promising XIAI opportunity lies in combining DL with causality, which can further enhance inherent interpretability. While the benefits of model-, input-, and output-based XIAI have been demonstrated, our findings suggest that the drawbacks of all these methods is that they require in-house technical expertise to infer XIAI. The key challenges include the lack of established best practices for XIAI selection based on data and problem type, as well as the unmet need for systematic evaluation metrics and high-quality benchmarks. Moreover, to effectively adopt DL in healthcare NLP, moving from “local” to “global” XIAI techniques that can inform the entire modelling process is crucial. Finally, our discussion encourages the integration of XIAI into the ongoing development of LLMs and domain-specific smaller language models.

Our results indicate that XIAI adoption in healthcare is not possible without dedicated in-house technical expertise. Bringing “humans into the DL loop” (domain experts, end-users, policymakers and patients) may be an effective solution to collectively design, develop and evaluate ready-to-use XIAI methods across different NLP and medical tasks. Despite the challenges identified, the XIAI techniques detailed in this review offer a valuable foundation for further research and benchmark studies. This will ultimately lead to enhanced inherent interpretability and facilitate the use of complex NLP algorithms in healthcare settings.

Funding

This work was supported by the MIS under Grant 5154714 of the National Recovery and Resilience Plan Greece 2.0 funded by EU under the NextGenerationEU Program.

CRediT authorship contribution statement

Guangming Huang: Visualization, Investigation, Formal analysis, Data curation, Software, Writing – original draft. **Yingya Li:** Writing – review & editing, Resources, Methodology, Investigation, Data curation. **Shoaib Jameel:** Writing – review & editing. **Yunfei Long:** Writing – review & editing, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition, Data curation, Conceptualization. **Giorgos Papanastasiou:** Writing – review & editing, Writing – original draft, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

YL and GH would like to acknowledge financial support from the School of Computer science and Electrical Engineering, University of Essex, UK. This work was supported by the MIS under Grant 5154714 of the National Recovery and Resilience Plan Greece 2.0 funded by EU under the NextGenerationEU Program.

References

- [1] Nguyen T, Yates A, Zirikly A, Desmet B, Cohan A. Improving the generalizability of depression detection by leveraging clinical questionnaires. In: Proceedings of the 60th annual meeting of the association for computational linguistics (volume 1: long papers); 2022. p. 8446–59.
- [2] Koleck TA, Dreisbach C, Bourne PE, Bakken S. Natural language processing of symptoms documented in free-text narratives of electronic health records: a systematic review. *J Am Med Inform Assoc* 2019;26:364–79.
- [3] Kim B-H, Deng Z, Philip SY, Ganapathi V. Can current explainability help provide references in clinical notes to support humans annotate medical codes? In: Proceedings of the 13th international workshop on health text mining and information analysis; 2022. p. 26–34.
- [4] Zhang F, Laish I, Benjamini A, Feder A. Section classification in clinical notes with multi-task transformers. In: Proceedings of the 13th international workshop on health text mining and information analysis; 2022. p. 54–9.
- [5] Cahyawijaya S, Wilie B, Lovenia H, Zhong H, Zhong M, Ip YYN, et al. How long is enough? Exploring the optimal intervals of long-range clinical note language modeling. In: Proceedings of the 13th international workshop on health text mining and information analysis; 2022. p. 160–72.
- [6] Michalopoulos G, Williams K, Singh G, Lin T. Medicalsum: a guided clinical abstractive summarization model for generating medical reports from patient-doctor conversations. In: Findings of the association for computational linguistics: EMNLP 2022; 2022. p. 4741–9.
- [7] Wu X, Yang S, Qiu Z, Ge S, Yan Y, Wu X, et al. Deltanet: conditional medical report generation for covid-19 diagnosis. In: Proceedings of the 29th international conference on computational linguistics; 2022. p. 2952–61.
- [8] You J, Li D, Okumura M, Suzuki K. Jpg-jointly learn to align: automated disease prediction and radiology report generation. In: Proceedings of the 29th international conference on computational linguistics; 2022. p. 5989–6001.
- [9] Moro G, Ragazzi L, Valgimigli L, Freddi D. Discriminative marginalized probabilistic neural method for multi-document summarization of medical literature. In: Proceedings of the 60th annual meeting of the association for computational linguistics (volume 1: long papers); 2022. p. 180–9.
- [10] Otmakhova J, Verspoor K, Baldwin T, Lau JH. The patient is more dead than alive: exploring the current state of the multi-document summarisation of the biomedical literature. In: Proceedings of the 60th annual meeting of the association for computational linguistics (volume 1: long papers); 2022. p. 5098–111.
- [11] Grundmann P, Oberhauser T, Gers F, Löser A. Attention networks for augmenting clinical text with support sets for diagnosis prediction. In: Proceedings of the 29th international conference on computational linguistics; 2022. p. 4765–75.
- [12] Yan B, Pei M. Clinical-bert: vision-language pre-training for radiograph diagnosis and reports generation. *Proc AAAI Conf Artif Intell* 2022;36:2982–90.
- [13] Liu Z, Lim H, Suhaimi NFA, Tong SC, Ong S, Ng A, et al. Fast prototyping a dialogue comprehension system for nurse-patient conversations on symptom monitoring. In: Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 2 (industry papers); 2019. p. 24–31.
- [14] Yang J, Ding Y, Long S, Poon J, Han SC. Ddi-mug: multi-aspect graphs for drug-drug interaction extraction. In: Proceedings of the 13th international workshop on health text mining and information analysis; 2022. p. 127–37.
- [15] Iinuma N, Miwa M, Sasaki Y. Improving supervised drug-protein relation extraction with distantly supervised models. In: Proceedings of the 21st workshop on biomedical language processing; 2022. p. 161–70.
- [16] Rumelhart DE, Hinton GE, Williams RJ. Learning representations by back-propagating errors. *Nature* 1986;323:533–6.
- [17] Hochreiter S, Schmidhuber J. Long short-term memory. *Neural Comput* 1997;9:1735–80.
- [18] Kim Y. Convolutional neural networks for sentence classification. In: Proceedings of the 2014 conference on empirical methods in natural language processing; 2014. p. 1746–51.
- [19] Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, et al. Attention is all you need. *Adv Neural Inf Process Syst* 2017;30.
- [20] Danilevsky M, Qian K, Aharonov R, Katsis Y, Kawas B, Sen P. A survey of the state of explainable ai for natural language processing. In: Proceedings of the 1st conference of the Asia-Pacific chapter of the association for computational linguistics and the 10th international joint conference on natural language processing; 2020. p. 447–59.
- [21] Sun X, Yang D, Li X, Zhang T, Meng Y, Qiu H, et al. Interpreting deep learning models in natural language processing: a review. *arXiv preprint. arXiv:2110.10470*, 2021.
- [22] Payrovnaziri SN, Chen Z, Rengifo-Moreno P, Miller T, Bian J, Chen JH, et al. Explainable artificial intelligence models using real-world electronic health record data: a systematic scoping review. *J Am Med Inform Assoc* 2020;27:1173–85.
- [23] Ozyegen O, Kabe D, Cevik M. Word-level text highlighting of medical texts for telehealth services. *Artif Intell Med* 2022;127:102284.
- [24] Thorsen-Meyer H-C, Placido D, Kaas-Hansen BS, Nielsen AP, Lange T, Nielsen AB, et al. Discrete-time survival analysis in the critically ill: a deep learning approach using heterogeneous data. *npj Digit Med* 2022;5:142.
- [25] Teng F, Yang W, Chen L, Huang L, Xu Q. Explainable prediction of medical codes with knowledge graphs. *Front Bioeng Biotechnol* 2020;8:867.
- [26] Dong H, Suárez-Paniagua V, Whiteley W, Wu H. Explainable automated coding of clinical notes using hierarchical label-wise attention networks and label embedding initialisation. *J Biomed Inform* 2021;116:103728.
- [27] Lipton ZC. The mythos of model interpretability: in machine learning, the concept of interpretability is both important and slippery. *Queue* 2018;16:31–57.
- [28] Doshi-Velez F, Kim B. Towards a rigorous science of interpretable machine learning. *arXiv preprint. arXiv:1702.08608*, 2017.
- [29] Carvalho DV, Pereira EM, Cardoso JS. Machine learning interpretability: a survey on methods and metrics. *Electronics* 2019;8:832.
- [30] Rudin C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nat Mach Intell* 2019;1:206–15.
- [31] Moher D, Liberati A, Tetzlaff J, Altman DG, PRISMA Group. Preferred reporting items for systematic reviews and meta-analyses: the prisma statement. *Ann Intern Med* 2009;151:264–9.
- [32] Lin T, Wang Y, Liu X, Qiu X. A survey of transformers. *AI Open* 2022;3:111–32. <https://doi.org/10.1016/j.aiopen.2022.10.001>.
- [33] Martina S, Ventura L, Frasconi P. Classification of cancer pathology reports: a large-scale comparative study. *IEEE J Biomed Health Inform* 2020;24:3085–94.
- [34] Marchesin S, Giachelle F, Marini N, Atzori M, Boytcheva S, Buttafuoco G, et al. Empowering digital pathology applications through explainable knowledge extraction tools. *J Pathol Inform* 2022;13:100139.
- [35] Farruque N, Goebel R, Zaiane OR, Sivapalan S. Explainable zero-shot modelling of clinical depression symptoms from text. In: 2021 20th IEEE international conference on machine learning and applications. IEEE; 2021. p. 1472–7.
- [36] De Magistris G, Russo S, Roma P, Starczewski JT, Napoli C. An explainable fake news detector based on named entity recognition and stance classification applied to covid-19. *Information* 2022;13:137.
- [37] Boukobza A, Burgun A, Roudier B, Tsopra R. Deep neural networks for simultaneously capturing public topics and sentiments during a pandemic: application on a covid-19 tweet data set. *JMIR Med Inform* 2022;10:e34306.
- [38] Chen D, Zhao H, He J, Pan Q, Zhao W. An causal xai diagnostic model for breast cancer based on mammography reports. In: 2021 IEEE international conference on bioinformatics and biomedicine. IEEE; 2021. p. 3341–9.
- [39] Féré T, Cossin S, Schaefferbeke T, Barnette T, Jouhet V, Hejblum BP. Automatic phenotyping of electronic health record: Phevis algorithm. *J Biomed Inform* 2021;117:103746.
- [40] Lu Z-H, Wang JX, Li X. Revealing opinions for covid-19 questions using a context retriever, opinion aggregator, and question-answering model: model development study. *J Med Internet Res* 2021;23:e22860.
- [41] Amador-Domínguez E, Serrano E, Manrique D, Bajo J. A case-based reasoning model powered by deep learning for radiology. report recommendation. 2021.

- [42] Ahne A, Fagherazzi G, Tannier X, Czernichow T, Orchard F. Improving diabetes-related biomedical literature exploration in the clinical decision-making process via interactive classification and topic discovery: methodology development study. *J Med Internet Res* 2022;24:e27434.
- [43] Zhang AY, Lam SSW, Ong MEH, Tang PH, Chan LL. Explainable ai: classification of mri brain scans orders for quality improvement. In: Proceedings of the 6th IEEE/ACM international conference on big data computing, applications and technologies; 2019. p. 95–102.
- [44] Uddin MZ. Depression detection in text using long short-term memory-based neural structured learning. In: 2022 international conference on innovations in science, engineering and technology. IEEE; 2022. p. 408–14.
- [45] Maji S, Appe A, Bali R, Chowdhury AG, Raghavendra VC, Bhandaru VM. An interpretable deep learning system for automatically scoring request for proposals. In: 2021 IEEE 33rd international conference on tools with artificial intelligence. IEEE; 2021. p. 851–5.
- [46] Naseem U, Khushi M, Kim J, Dunn AG. Hybrid text representation for explainable suicide risk identification on social media. In: IEEE transactions on computational social systems; 2022.
- [47] Dobrakowski AG, Mykowiecka A, Marciniak M, Jaworski W, Biecek P. Interpretable segmentation of medical free-text records based on word embeddings. *J Intell Inf Syst* 2021;57:447–65.
- [48] Sushil M, Šuster S, Luyckx K, Daelemans W. Patient representation learning and interpretable evaluation using clinical notes. *J Biomed Inform* 2018;84:103–13.
- [49] Minot JR, Cheney N, Maier M, Elbers DC, Danforth CM, Dodds PS. Interpretable bias mitigation for textual data: reducing genderization in patient notes while maintaining classification performance. *ACM Trans Comput Healthc* 2022;3:1–41.
- [50] Bhatt A, Roberts R, Chen X, Li T, Connor S, Hatim Q, et al. Dice: a drug indication classification and encyclopedia for ai-based indication extraction. *Front Artif Intell* 2021;4:711467.
- [51] Lindsay H, Tröger J, König A. Language impairment in Alzheimer's disease—robust and explainable evidence for ad-related deterioration of spontaneous speech through multilingual machine learning. *Front Aging Neurosci* 2021:228.
- [52] Garcia-Olano D, Onoe Y, Baldini I, Ghosh J, Wallace BC, Varshney K. Biomedical interpretable entity representations. In: Findings of the association for computational linguistics: ACL-IJCNLP 2021; 2021. p. 3547–61.
- [53] Ong CJ, Orfanoudaki A, Zhang R, Caprasse FPM, Hutch M, Ma L, et al. Machine learning and natural language processing methods to identify ischemic stroke, acuity and location from radiology reports. *PLoS ONE* 2020;15:e0234908.
- [54] Holderness E, Miller N, Cawkwell P, Bolton K, Meteere M, Pustejovsky J, et al. Analysis of risk factor domains in psychosis patient health records. *J Biomed Semant* 2019;10:1–10.
- [55] Mellado EÁ, Holderness E, Miller N, Dhang F, Cawkwell P, Bolton K, et al. Assessing the efficacy of clinical sentiment analysis and topic extraction in psychiatric readmission risk prediction. In: Proceedings of the tenth international workshop on health text mining and information analysis; 2019. p. 81–6.
- [56] Xue Q, Chuah MC. Explainable deep learning based medical diagnostic system. *Smart Health* 2019;13:100068.
- [57] Gin BC, Ten Cate O, O'Sullivan PS, Hauer KE, Boscardin C. Exploring how feedback reflects entrustment decisions using artificial intelligence. *Med Educ* 2022;56:303–11.
- [58] Frisoni G, Moro G. Phenomena explanation from text: unsupervised learning of interpretable and statistically significant knowledge. In: Data management technologies and applications: 9th international conference, revised selected papers 9. Springer; 2021. p. 293–318.
- [59] Mandalios A, Chortaras A, Stamou G, Vazirgiannis M. Enriching graph representations of text: application to medical text classification. In: Complex networks & their applications IX: volume 2, proceedings of the ninth international conference on complex networks and their applications COMPLEX NETWORKS 2020. Springer; 2021. p. 92–103.
- [60] Zhang Y, Yang S, Wen H, Xie P, Wu Y. Unified framework for ner and re tasks with strong interpretability on Chinese medicine instruction parsing. In: 2022 international conference on computer engineering and artificial intelligence. IEEE; 2022. p. 460–6.
- [61] Gu G, Li B, Gao H, Wang M. Learning to answer complex questions with evidence graph. In: Web and big data: 4th international joint conference, proceedings, part I 4. Springer; 2020. p. 257–69.
- [62] Duarte F, Martins B, Pinto CS, Silva MJ. Deep neural models for icd-10 coding of death certificates and autopsy reports in free-text. *J Biomed Inform* 2018;80:64–77.
- [63] Zhang J, Kowsari K, Harrison JH, Lobo JM, Barnes LE. Patient2vec: a personalized interpretable deep representation of the longitudinal electronic health record. *IEEE Access* 2018;6:65333–46.
- [64] Trigueros O, Blanco A, Lebeña N, Casillas A, Pérez A. Explainable icd multi-label classification of ehrs in Spanish with convolutional attention. *Int J Med Inform* 2022;157:104615.
- [65] Chen P-F, He T-L, Lin S-C, Chu Y-C, Kuo C-T, Lai F, et al. Training a deep contextualized language model for international classification of diseases, 10th revision classification via federated learning: model development and validation study. *JMIR Med Inform* 2022;10:e41342.
- [66] Zhu Y, Mahale A, Peters K, Mathew L, Giuste F, Anderson B, et al. Using natural language processing on free-text clinical notes to identify patients with long-term covid effects. In: Proceedings of the 13th ACM international conference on bioinformatics, computational biology and health informatics; 2022. p. 1–9.
- [67] Ahmed U, Srivastava G, Yun U, Lin JC-W. Eandc: an explainable attention network based deep adaptive clustering model for mental health treatment. *Future Gener Comput Syst* 2022;130:106–13.
- [68] Ahmed U, Lin JC-W, Srivastava G. Multi-aspect deep active attention network for healthcare explainable adoption. *IEEE J Biomed Health Inform* 2022.
- [69] Luo X, Gandhi P, Zhang Z, Shao W, Han Z, Chandrasekaran V, et al. Applying interpretable deep learning models to identify chronic cough patients using ehr data. *Comput Methods Programs Biomed* 2021;210:106395.
- [70] Balagopalan A, Eyre B, Robin J, Rudzicz F, Novikova J. Comparing pre-trained and feature-based models for prediction of Alzheimer's disease based on speech. *Front Aging Neurosci* 2021;13:635945.
- [71] Ribeiro MT, Singh S, Guestrin C. Model-agnostic interpretability of machine learning. arXiv preprint. arXiv:1606.05386, 2016.
- [72] Lundberg SM, Lee S-I. A unified approach to interpreting model predictions. *Adv Neural Inf Process Syst* 2017;30.
- [73] Shapley LS. A value for n-person games. In: Classics in game theory; 1997. p. 69.
- [74] Van der Maaten L, Hinton G. Visualizing data using t-sne. *J Mach Learn Res* 2008;9.
- [75] Hamilton WL. Graph representation learning. Synthesis lectures on artificial intelligence and machine learning, vol. 14. 2020. p. 1–159.
- [76] Wang Q, Pan X, Huang L, Zhang B, Jiang Z, Ji H, et al. Describing a knowledge base. In: Proceedings of the 11th international conference on natural language generation; 2018. p. 10–21.
- [77] Fu B, Qiu Y, Tang C, Li Y, Yu H, Sun J. A survey on complex question answering over knowledge base: recent advances and challenges. arXiv preprint. arXiv:2007.13069, 2020.
- [78] Bahdanau D, Cho K, Bengio Y. Neural machine translation by jointly learning to align and translate. arXiv preprint. arXiv:1409.0473, 2014.
- [79] Cheng J, Dong L, Lapata M. Long short-term memory-networks for machine reading. In: Proceedings of the 2016 conference on empirical methods in natural language processing; 2016. p. 551–61.
- [80] Papanastasiou G, Dikaios N, Huang J, Wang C, Yang G. Is attention all you need in medical image analysis? A review. *IEEE J Biomed Health Inform* 2023:1–14. <https://doi.org/10.1109/JBHI.2023.3348436>.
- [81] Mascharka D, Tran P, Soklaski R, Majumdar A. Transparency by design: closing the gap between performance and interpretability in visual reasoning. In: Proceedings of the IEEE conference on computer vision and pattern recognition; 2018. p. 4942–50.
- [82] Ouyang L, Wu J, Jiang X, Almeida D, Wainwright C, Mishkin P, et al. Training language models to follow instructions with human feedback. *Adv Neural Inf Process Syst* 2022;35:27730–44.
- [83] Mosqueira-Rey E, Hernández-Pereira E, Alonso-Ríos D, Bobes-Bascarán J, Fernández-Leal Á. Human-in-the-loop machine learning: a state of the art. *Artif Intell Rev* 2023;56:3005–54.
- [84] Wu X, Xiao L, Sun Y, Zhang J, Ma T, He L. A survey of human-in-the-loop for machine learning. *Future Gener Comput Syst* 2022;135:364–81.
- [85] Salvi M, Loh HW, Seoni S, Barua PD, García S, Molinari F, et al. Multi-modality approaches for medical support systems: a systematic review of the last decade. *Inf Fusion* 2023:102134.
- [86] Schölkopf B, Locatello F, Bauer S, Ke NR, Kalchbrenner N, Goyal A, et al. Toward causal representation learning. *Proc IEEE* 2021;109:612–34.
- [87] Pearl J. An introduction to causal inference. *Int J Biostat* 2010;6. <https://doi.org/10.2202/1557-4679.1203>.
- [88] Feder A, Keith KA, Manzoor E, Pryzant R, Sridhar D, Wood-Doughty Z, et al. Causal inference in natural language processing: estimation, prediction, interpretation and beyond. *Trans Assoc Comput Linguist* 2022;10:1138–58.
- [89] Serrano S, Smith NA. Is attention interpretable? In: Proceedings of the 57th annual meeting of the association for computational linguistics. Association for Computational Linguistics; 2019.
- [90] Pruthi D, Gupta M, Dhingra B, Neubig G, Lipton ZC. Learning to deceive with attention-based explanations. In: Proceedings of the 58th annual meeting of the association for computational linguistics; 2020. p. 4782–93.
- [91] Wiegrefe S, Pinter Y. Attention is not not explanation. In: Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing; 2019. p. 11–20.
- [92] Vashishth S, Upadhyay S, Tomar GS, Faruqui M. Attention interpretability across nlp tasks. arXiv preprint. arXiv:1909.11218, 2019.
- [93] Brunner G, Liu Y, Pascual D, Richter O, Ciaramita M, Wattenhofer R. On identifiability in transformers. In: International conference on learning representations; 2019.
- [94] OpenAI. Chatgpt: optimizing language models for dialogue. <https://openai.com/blog/chatgpt/>, 2022.
- [95] Brown T, Mann B, Ryder N, Subbiah M, Kaplan JD, Dhariwal P, et al. Language models are few-shot learners. *Adv Neural Inf Process Syst* 2020;33:1877–901.
- [96] Achiam J, Adler S, Agarwal S, Ahmad I, Akkaya I, Aleman FL, et al. Gpt-4 technical report. arXiv preprint. arXiv:2303.08774, 2023.
- [97] Meskó B, Topol EJ. The imperative for regulatory oversight of large language models (or generative ai) in healthcare. *npj Digit Med* 2023;6:120.
- [98] Tu T, Azizi S, Driess D, Schaeckermann M, Amin M, Chang P-C, et al. Towards generalist biomedical ai. *NEJM AI* 2024;1. A10a2300138.

- [99] Singhal K, Azizi S, Tu T, Mahdavi SS, Wei J, Chung HW, et al. Publisher correction: large language models encode clinical knowledge. *Nature* 2023;620:E19.
- [100] Jin Z, Chen Y, Leeb F, Gresele L, Kamal O, Zhiheng L, et al. Cladder: assessing causal reasoning in language models. In: *Thirty-seventh conference on neural information processing systems*; 2023.
- [101] Hobbhahn M, Lieberum T, Seiler D. Investigating causal understanding in llms. In: *NeurIPS 2022 workshop on causality for real-world impact*; 2022.
- [102] Wang Y, Zhao Y, Petzold L. Are large language models ready for healthcare? A comparative study on clinical language understanding. In: *Machine learning for healthcare conference*. PMLR; 2023. p. 804–23.
- [103] Hernandez E, Mahajan D, Wulff J, Smith MJ, Ziegler Z, Nadler D, et al. Do we still need clinical language models? In: *Conference on health, inference, and learning*. PMLR; 2023. p. 578–97.