

Methodological Review

Deep representation learning of patient data from Electronic Health Records (EHR): A systematic review



Yuqi Si ^a, Jingcheng Du ^a, Zhao Li ^a, Xiaoqian Jiang ^a, Timothy Miller ^b, Fei Wang ^c, W. Jim Zheng ^a, Kirk Roberts ^{a,*}

^a School of Biomedical Informatics, The University of Texas Health Science Center at Houston, TX, USA

^b Computational Health Informatics Program (CHIP), Boston Children's Hospital and Harvard Medical School, MA, USA

^c Department of Population Health Sciences. Weill Cornell Medicine, Cornell University, NY, USA

ARTICLE INFO

Keywords:

Systematic review
Electronic health records
Patient representation
Deep learning

ABSTRACT

Objectives: Patient representation learning refers to learning a dense mathematical representation of a patient that encodes meaningful information from Electronic Health Records (EHRs). This is generally performed using advanced deep learning methods. This study presents a systematic review of this field and provides both qualitative and quantitative analyses from a methodological perspective.

Methods: We identified studies developing patient representations from EHRs with deep learning methods from MEDLINE, EMBASE, Scopus, the Association for Computing Machinery (ACM) Digital Library, and the Institute of Electrical and Electronics Engineers (IEEE) Xplore Digital Library. After screening 363 articles, 49 papers were included for a comprehensive data collection.

Results: Publications developing patient representations almost doubled each year from 2015 until 2019. We noticed a typical workflow starting with feeding raw data, applying deep learning models, and ending with clinical outcome predictions as evaluations of the learned representations. Specifically, learning representations from structured EHR data was dominant (37 out of 49 studies). Recurrent Neural Networks were widely applied as the deep learning architecture (Long short-term memory: 13 studies, Gated recurrent unit: 11 studies). Learning was mainly performed in a supervised manner (30 studies) optimized with cross-entropy loss. Disease prediction was the most common application and evaluation (31 studies). Benchmark datasets were mostly unavailable (28 studies) due to privacy concerns of EHR data, and code availability was assured in 20 studies.

Discussion & Conclusion: The existing predictive models mainly focus on the prediction of single diseases, rather than considering the complex mechanisms of patients from a holistic review. We show the importance and feasibility of learning comprehensive representations of patient EHR data through a systematic review. Advances in patient representation learning techniques will be essential for powering patient-level EHR analyses. Future work will still be devoted to leveraging the richness and potential of available EHR data. Reproducibility and transparency of reported results will hopefully improve. Knowledge distillation and advanced learning techniques will be exploited to assist the capability of learning patient representation further.

1. Introduction

In Electronic Health Records (EHRs), information regarding patient status is extensively documented. Therefore, EHR data provides a feasible mechanism to track patient health information and to make better decisions based on data-driven technologies. Unlike data in clinical trials or other biomedical studies, secondary data extracted from EHRs are not designed to answer a specific hypothesis. Instead, their

primary goal is to monitor a patient. This results in the issue that EHR data have many challenging characteristics such as *uncurated* (data are not carefully chosen and thoughtfully organized or presented), *poor-quality* (data are rarely subject to data quality audits), *high-dimensional* (thousands of distinct medical events), *sparse* (lots of zero values), *heterogeneous* (drawn from different resources), *temporal* (data are collected over time), *incomplete* (missing values), *large-scale* (a large volume of data), and *multimodal* (multiple data modalities). A wide variety of

* Corresponding author at: School of Biomedical Informatics, University of Texas Health Science Center at Houston, 7000 Fannin St. #600, Houston TX 77030.
E-mail address: kirk.roberts@uth.tmc.edu (K. Roberts).

studies have conducted predictive modeling of EHR data, which is a machine learning task that applies EHR data to construct a statistical model for the purpose of predicting a given clinical outcome of interest [1]. However, the complexity of EHR data discussed as above makes it difficult to directly use EHR raw data in machine learning models to achieve predictive modeling.

A critical element in predictive modeling of EHR data is to effectively convert patient data from the raw EHR format to a machine learning representation—in other words, to transform patient data to meaningful information that can be further understood algorithmically. The effectiveness of predictive models for improving disease diagnosis, phenotyping, and prognosis heavily depends on the quality of this feature representation. In machine learning, the task of representation learning is to learn and extract good feature representations from raw data automatically [2]. Patient representation learning is one particularly promising direction of combining representation learning and large EHR datasets. It refers to learning a mathematical description of patient data to find an appropriate way of transforming raw data into meaningful features. The patient representations built from many EHR data modalities (including clinical narratives, lab tests, treatments, etc.) should be organized in a form that enables machine learning to learn effective prediction models for many tasks.

This work investigates the methods of representation learning and the field of patient representation learning from EHRs through a methodological review of the literature. We collected data on 28 patient representation variables from 49 papers published in a diverse variety of venues, published up to December 2019.

We seek to understand the following research questions:

1. **Resources:** What are the resources available for learning patient representations? How is patient data transformed from raw input to important features?
2. **Methods:** What representation learning methods are being contributed? What kind of models and algorithms are used to develop patient representations?
3. **Applications:** What types of clinical problems and outcomes are addressed?
4. **Potential:** How could these methods potentially contribute to diverse research communities?

2. Background

2.1. Patient learning data pipeline

First, we briefly summarize the methods and resources that are commonly used in related studies (Fig. 1). The learning starts with raw patient data from either structured codes such as ICD-9, or unstructured data like clinical notes. After some initial embedding techniques to transform patient data into input features, deep learning models are

utilized to create the patient representation. The learning model can be unsupervised, supervised, or self-supervised. The final step is to evaluate the learned representations by applying it to some clinical task(s), such as mortality, readmission, or a specific disease prediction. There are also some ways to visualize patient representations to provide some intuitive understanding or interpretation of the representation. The majority of studies in this field follow this pipeline to learn patient representations from EHR data and further evaluate on downstream tasks.

2.2. Patient representation methods

According to Bengio et al. [2], deep learning models are representation learning methods with multiple layers, each with their own representations. The model starts with raw input, constructs non-linear modules that transform the representation at each layer into high-level, abstract representations. As already discussed, EHR data is high-dimensional and sparse. Thus, deep learning models are particularly well-suited to encode EHR data to learn patient representations. In this section, we will introduce the background of representation learning methods and how they are applied to learn patient representations. We describe five methods of representing a patient: vector-based, temporal matrix-based, graph-based, sequence-based, and tensor-based representations. For more detailed and technical interpretations of deep learning models, we recommend the review by LeCun et al. [3].

2.2.1. Vector-based patient representation

In vector-based patient representation, every patient is represented by a mathematical vector. Models that attempt to construct vector-based patient representations are as follows:

Fully connected deep neural network (fully connected DNN), originally inspired by biological neural networks, contain one or more hidden layers connected in a feed-forward manner [4]. Fully connected DNN are also known as multi-layer perceptrons (MLPs) or dense neural networks. Like the other neural network methods described below, they use backpropagation for supervised learning with non-linear activations. Some early attempts for learning patient representation with neural networks have applied fully connected DNN approaches [5]. However, the majority of recent studies have considered this architecture as a baseline method [6].

Autoencoders are an unsupervised deep learning model that learns abstract representations from high-dimensional data, which is also a way of performing dimensionality reduction. An autoencoder learns to predict its input, but must compress the input signal through progressively smaller intermediate layers. The inner layer can be used as a low-dimensional representation. The variants of autoencoders include denoising [7], stacked denoising [8], variational [9], and contractive autoencoder [10]. Given the prior success of autoencoders for representation learning, their application to patient representation learning was first introduced by Miotto et al. [11], where a dense and general

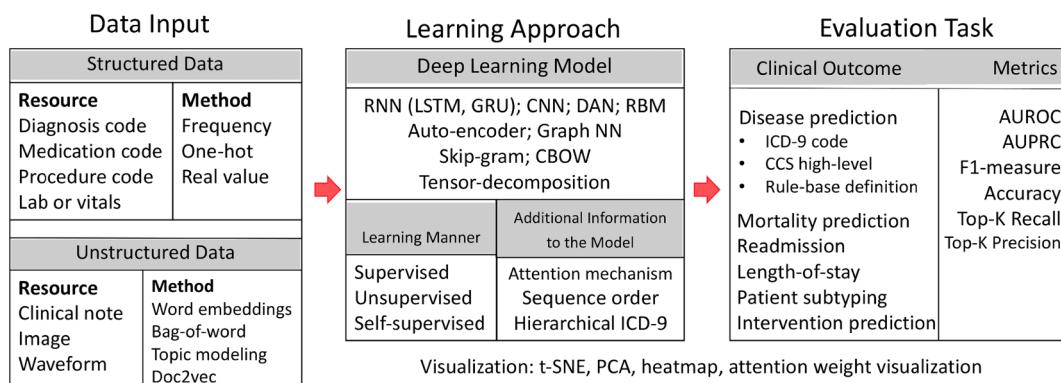


Fig. 1. Workflow of patient representation learning.

patient representation was derived with unsupervised three-layer stacked denoising autoencoders.

Convolutional Neural Networks (CNNs) were specifically developed for image processing [12]. Each module of the CNN contains a convolutional layer and a pooling layer. Those modules are stacked to construct a deep CNN. The convolutional layer uses a small neural network that slides across the image, allowing for the capture of local image properties irrespective of their location in the full image. However, CNNs have been applied to more data types than just images, including text [13] and waveforms [14]. More advanced variants of CNN also attempt to incorporate temporal information in addition to the convolutional layer to model longitudinal patient data [15]. This is different from the variants of CNN in image processing fields such as VGGNet [16] and ResNet [17], which add more complex layers in the architecture.

Word2vec originated from natural language processing to learn word embeddings from large-scale text resources [18]. There are two algorithms in word2vec: continuous bag of words (CBOW) predicts a target word given the surrounding context, and skip-gram predicts the surrounding context given a target word [19]. Both algorithms have just one hidden layer; thus, word2vec techniques are shallow networks compared with other deep learning methods. The variants of word2vec have been applied to learn representations from clinical codes [20], characterized by different assumptions and aimed to capture relationships between code sequences.

2.2.2. Temporal matrix-based patient representation

Temporal matrix-based patient representation constructs a two-dimensional matrix with one dimension related to time and the other dimension related to clinical events from the EHR. Nonnegative matrix factorization (NMF) is an algorithm for decomposing high-dimensional data from a set of nonnegative elements. NMF has been widely used in bioinformatics for clustering sources of variations [21,22]. There were also early attempts of applying NMF or its variants into EHR patient data representation [23–25,15], which are some of the earliest examples of constructing a mathematical representation for each patient. They also showed the challenges of using EHR data and demonstrated the feasibility of encoding the latent factors of temporal patient data by providing a one-to-one identifiable mapping between the patient matrix and the target label.

2.2.3. Graph-based patient representation

Graph-based patient representation constructs a compact graph for each patient where the nodes in the graph encode clinical events and the edges between the nodes encode relationships among the clinical events. One of the early works developing graph-based EHR representation was proposed by Liu et al. [26], where they designed a novel graph representation algorithm to learn distinct clinical events from EHR data that included temporal relationships among the events. Although their work was not developed based on deep learning, one such emerging application of deep learning into graph representation is the Graph Neural Network.

Graph Neural Networks consist of a finite number of nodes and edges to connect data. Nodes contain information about entities, and edges contain relations between entities. Prominent methods of GNN include Directed Acyclic Graph [27], Graph Convolutional Network [28], Graph Attention Network [29], and node2vec embeddings [30]. GNNs attempt to learn graphical structures of EHR data that can infer the missing information through other representation mechanisms, therefore resulting in a more explainable representation. GNNs are particularly useful to introduce domain knowledge into the architecture. For instance, a few studies have employed hierarchical ICD-9 knowledge graphs as the graph model to enhance the interpretability and performance of the proposed method [27,31]. Other knowledge-based resources include adverse drug-drug interactions [32], and comorbidities [33].

2.2.4. Sequence-based patient representation

Sequence-based patient representation produces a timestamped event sequential features for each patient.

Recurrent Neural Networks (RNNs) were developed for processing sequential inputs such as language [34,35]. RNNs deal with a sequence of inputs one item at a time and transfer the hidden state of each input unit to the next input unit, so the current state implicitly contains information about the entire sequence history. The common variants of RNN include Gated Recurrent Unit (GRU) and Long Short-Term Memory (LSTM), both designed to minimize the vanishing gradient problem. GRU adds a gating mechanism into the RNN [36]. LSTM is particularly useful for dealing with long-term dependencies [37]. Thanks to the advances in RNN architecture, many studies have applied RNNs to develop patient representations where they have focused on representing patients using combinations or sequences of clinical codes. One such RNN-inspired approach is Doctor AI [38], where distributed vector representations of clinical codes are derived to represent patient trajectories.

Although RNNs enable to sequential modeling, especially some variants of RNNs like LSTMs are effective in dealing with long dependencies, limitations of RNNs still persist. One notable limitation is that RNNs cannot be trained in parallel (which will increase training time). Also, RNNs only process information from one direction, and even bi-directional RNNs which encode from two directions are a simple concatenation of two directions. Therefore, the **Transformer** architecture equipped with self-attention and positional embeddings was introduced to achieve true bi-directional representation [39].

Recently, with the wide adoption of Transformer-based language models in NLP (e.g., GPT [40], BERT [41]), Transformers continue to be applied for patient representation learning and for clinical sequence modeling [42–45]. Similar to the use of RNNs to enable sequence modeling of EHR data, Transformers also encode each clinical event at every time stamp as a unit and an entire patient trajectory as a whole sequence. Unlike a RNN using recurrence to predict the next unit, the Transformer architecture considers the sequence of units as a whole and employs self-attention to learn the essential information from the entire sequence.

2.2.5. Tensor-based patient representation

Tensor-based patient representation is a method that represents each patient with a three-dimensional (or more) tensor consisting of events such as diagnoses and treatments.

Tensor Decomposition is the high-order extension of matrix decomposition that seeks to decompose high-dimensional tensors into products of low-dimensional factors [46]. The CANDECOMP/PARAFAC alternating Poisson regression (CP-APR) algorithm is one of the representative methods that are well-studied for tensor decomposition, trying to express a tensor as the sum of a finite number of rank-one tensors [47]. Compared with traditional dimensionality reduction methods, tensor decomposition is unique in having the capability of managing multi-aspect features in multiple dimensions and is versatile enough to incorporate domain knowledge into the operation [48]. These advantages have made tensor decomposition a promising modeling approach to learn abstract patient representations from EHR data and provide good interpretability and scalability [49]. Each tensor is constructed of patient-level data, with three different factors including a patient factor, and two other factors related to clinical events (i.e., diagnosis with medication, or diagnosis with procedure). Each tensor constitutes a phenotype with a weighted sum of rank-one tensor from the outer product of three factor vectors. A tensor-based patient representation allows for the capture of complex interactions and relationships between clinical events (especially phenotypes, comorbidities, and medications) that are not evident in flattened EHR data [50].

3. Materials and methods

Papers that are eligible for inclusion in this review are characterized by (a) focus on patient representation learning, (b) use of patient data from longitudinal EHRs, and (c) utilization of deep learning or neural network models. Notably, we consider studies using deep learning as one criterion because the unique characteristics of EHR data (as discussed in Introduction) make deep learning a feasible data-driven solution to generate robust representation simultaneously from diverse resources compared with other techniques. It is likely that, at least in the next few years, most advances in patient representation will utilize deep learning approaches.

These criteria were applied to generate keyword searching queries for literature database search. Five databases consisting of PubMed, EMBASE, Scopus, the Association for Computing Machinery (ACM) Digital Library, and the Institute of Electrical and Electronics Engineers (IEEE) Xplore Digital Library were queried on December 31, 2019. The search consists of the following combination of keywords:

(“deep learning” OR “neural network”) AND patient AND (vector OR representation OR embedding OR vec) AND (“electronic health records” OR ehr OR “electronic medical records” OR emr).

As for publication characteristics, we limit our search to English-language research articles with the publication time range from 2000 to 2019 and only focus on peer-reviewed journals and conference proceedings (i.e., posters and preprints are not included). For study characteristics, only study designs related to developing a deep learning- or neural network-based patient representation from electronic health records are considered. The study outcome should also apply the learned representation to downstream evaluations of clinical predictions. In terms of exclusion criteria, studies such as review papers and proceeding summaries are excluded. Studies that use patient data other than EHRs are also excluded.

Duplicates were removed and additional records were added using a snowballing strategy. The additional records were subsequently traced from references of the included papers and personal readings. Next, the articles were uploaded to Rayyan [51] and screened for title and abstract. Two authors (YS, JD) screened the articles under a blind format. At the full-text screening step performed on Zotero [52], each study was read in full by one author (YS) to validate the final relevancy.

As a result, 363 papers were queried at the beginning, including 349 articles retrieved from the five databases, and 14 additional records identified through snowballing. At deduplication, 54 duplicate records were removed and 309 records remained. At the step of title and abstract screening, 199 records were excluded because the topic is not relevant ($n = 114$), they are traditional predictive modeling of the single task (28), they are not peer-reviewed or original research papers (26 studies are not research, 5 preprints), no deep learning is used (12), no evaluation of the patient representations (7), and there is no EHR data (7).

After these exclusions, 110 papers went through to the next screening step (full text screening) to determine the relevancy of each study. At this step, 61 papers were excluded due to no patient representation being developed ($n = 37$), no deep learning method being used (8), no EHR data being used (7), no evaluation (6), and only containing abstracts (3). Finally, this left 49 articles to be included in this work. A detailed literature selection procedure (PRISMA diagram) [53] is shown in Fig. 2. Based on the variables extracted from 49 publications, some interesting findings are worthy of being highlighted and discussed in the following sections.

We performed data extraction from the final included papers considering the following variables:

- Publication date: the date on which this paper was published by the conference or journal.

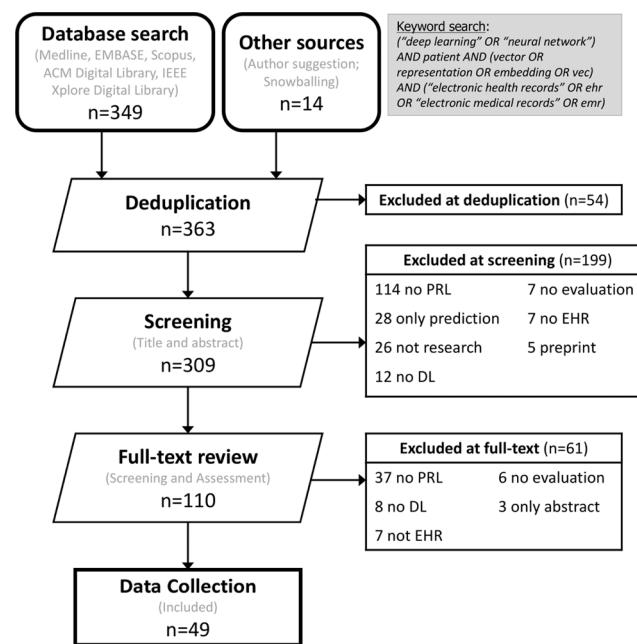


Fig. 2. PRISMA flowchart for including articles in this review. (PRL: Patient Representation Learning. DL: Deep Learning. EHR: Electronic Health Records.)

- Publication venue: the conference or journal where this paper was published.
- Representation learning approach: this contains two items including the proposed deep learning model(s) and additional attributes to enhance the model.
- Patient data type: what modality of EHR data was used? For instance, structured code from diagnoses/procedures/medications tables, numerical values from lab measurements, or unstructured data from clinical notes/images/waveforms.
- EHR resources: where does the patient EHR data come from? For example, publicly available datasets or a private local hospital data warehouse? What is the size/duration?
- Preprocessing methods: what is the format of the input data that was fed into the model architecture? In other words, how the EHR raw data was preprocessed into numerical data?
- Clinical outcome tasks: main clinical tasks for evaluating learned patient representation.
- Evaluation metrics: quantitative measure of the performance of prediction/clustering/regression on clinical outcome tasks.
- Interpretability: the degree to which a human can understand the model’s result or learned patient representations.
- Objective functions: the loss function when optimizing the models.
- Computational resources: for instance, deep learning frameworks and computational platforms (GPU).
- Code reproducibility: reproducibility of the results/instructions on open-source platforms.

4. Results

4.1. Publication characteristics

First, we look at the publication numbers over time in Fig. 3. The volume of studies is currently growing rapidly each year (almost 3 times in 2017 over 2016). Although there is a slight decrease from 2018 to 2019, we think this is because we exclude preprints at the first step of screening. Apart from the publication date, we also outline the major research communities that have contributed to the topic. Table 1 lists the top conferences or journals that have published this topic in three main communities. The conferences or journals include, but are not

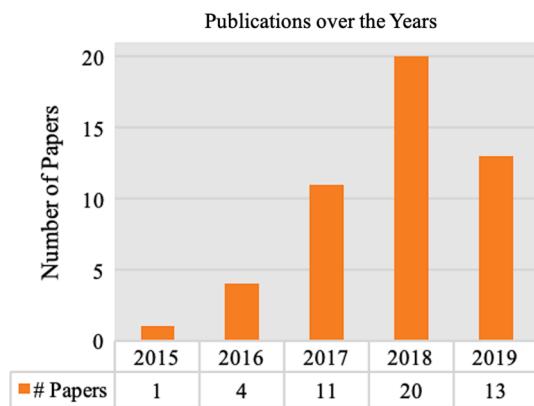


Fig. 3. Publication counts over time.

Table 1
Selection of conferences/journals publishing the included studies.

Community	Conference/ Journal	Name (# Papers)
Informatics	Conference	AMIA Summit (1), AMIA Symposium (1), others (4)
	Journal	JBI (3), JAMIA (1)
Medicine	Journal	NPJ Digital Medicine (1), Scientific Reports (2), PLoS ONE (1)
	Conference	KDD (8), AAAI (4), CIKM (3), NeurIPS_ML4H (4), IEEE-related (7), others (7)
Computer Science	Journal	JMLR (2)

limited to, AMIA Symposium, AMIA Joint Summit, Conference on Information and Knowledge Management (CIKM), Journal of the American Medical Informatics Association (JAMIA), Journal of Biomedical Informatics (JBI), BMC Medicine, Scientific Data, Scientific Reports, NPJ Digital Medicine, Neural Information Processing Systems (NeurIPS), Knowledge Discovery and Data Mining (KDD), Machine Learning for Health Care (NeurIPS_ML4H), and Association for the Advancement of Artificial Intelligence (AAAI).

4.2. Study characteristics

Pertinent study characteristics of the 49 reviewed publications are reported in the following sub-sections, and also summarized in [Supplementary Table 1](#). These study characteristics enable to qualify the scientific validity and applicability of each study, which reveals the inadequacies and gaps of current research fields, and would be potentially useful to inform the need for new study designs for future research in patient representation learning.

4.2.1. Learning scenarios

We find that, among the 49 included studies, the majority ($n = 30$) applied supervised approaches to learning the representation. Here one would train and optimize the representation with some target or objective, usually focusing on phenotyping or other clinical outcomes. Unsupervised learning was applied in 11 studies. Unlike supervised learning, unsupervised learning does not require labels. This would include the use of autoencoders, which function by reconstructing the data from its raw input. In all forms of unsupervised learning, patient representations are learned separately from the prediction task. Since no labels are required, unsupervised learning is more likely to benefit from large EHR datasets.

Self-supervised methods were used by 8 studies. Unlike supervised learning that relies on labels for a target task, or unsupervised learning that has no label-driven objectives, self-supervised learning frames a learning scenario as predicting a subset of information using the

remaining data. Essentially this enables the use of models that are commonly used for supervised tasks (including many deep learning models) in a setting that does not have the labels of interest for the target task. With self-supervised learning, the patient representation is obtained similar to word embeddings in NLP studies, where a target word is predicted with its contextual information. Since patient representations are themselves largely intended as intermediate steps to a final prediction task (usually done via supervised learning), learning the representations in a self-supervised manner provides many of the benefits of unsupervised learning (no manual curation of labels) while still allowing for the use of powerful deep learning methods (e.g., transformers) that can otherwise be considered supervised. The self-supervised tasks included in this review are mainly used for clinical code representation learning, where clinical codes are ordered sequentially, and the codes in between are masked and predicted using the context of clinical codes nearby, similar to the algorithms of skip-gram and CBOW used to learn word embeddings [\[19\]](#).

4.2.2. Patient data types

In terms of patient data types in EHRs, among the 49 studies, 37 papers used structured codes (i.e., diagnosis codes, procedure codes, medication codes) to build the representation, while 6 papers used unstructured notes. The remaining 6 papers used both structured codes and unstructured notes jointly. Though combining heterogeneous resources is promising, we noticed these studies were not really combining the raw data fully from the two modalities into the same model. Instead, they only used a subset of data from one modality [\[6,11,54–57\]](#). One such method is to simplify clinical notes by using topic models to represent unstructured resources [\[6,11,55\]](#).

4.2.3. Preprocessing methods

Although not many studies emphasized the preprocessing steps used to transfer raw input data into the model, we assume this is because deep learning does not necessarily need heavy manual work for preprocessing. However, we still extract the preprocessing approach and workflow of each study because we consider this an essential step towards representing a patient. Among 37 studies learning from structured codes, the majority ($n = 27$) generate one-hot vectors for the input data, and the frequency of a given code is used in 5 studies [\[11,48,58–60\]](#). For inputs from multiple resources, input vectors are concatenated before feeding them into the architecture [\[14\]](#). Normalized formats of patient data are also important to learn an effective representation. For instance, z-score is applied to convert data from continuous variables to discrete one-hot vectors [\[55,57,61\]](#). For both structured codes and unstructured notes, standardized vocabularies including FHIR [\[6\]](#) and UMLS CUIs [\[62,63\]](#) are introduced to transfer raw tokens into unique phrases to reduce the sparsity of the input data.

4.2.4. Learning models and architectures

A primary interest of this study is the types of deep learning models that the proposed methods used as the foundation to learn patient representations. We plot the number of papers that have applied each deep learning model as the foundation architecture, and also differentiate studies with different data input types, shown in [Fig. 4](#) and [Table 2](#). In total, there are 13 studies (structured: 8; unstructured: 1; combined: 4) that have applied an LSTM model to develop patient representations, which is the most common architecture. CNNs are the next with 11 studies (structured: 7; unstructured: 2; combined: 2). GRU is also common with 11 studies (all structured EHR data). This result is reasonable since both CNN and RNN architectures are utilized to learn local or sequential information from large-scale data and to find the nuance of variance within the data.

Although the methods of learning patient representations are all built upon deep learning models, many unique architectures were introduced, and additional learning modules have been applied to capture the characteristics of EHR data. To improve performance and

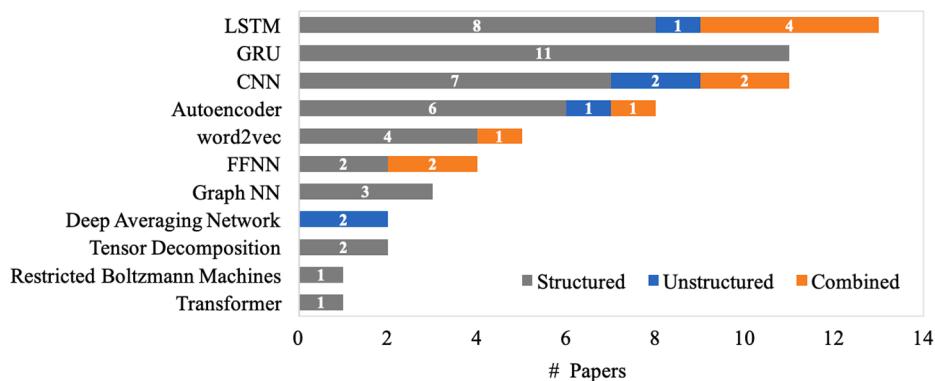


Fig. 4. Deep learning architectures used in the included studies.

Table 2
Correspondence between patient data types and deep learning models.

DL Models	Patient Input Data Type		
	Structured	Unstructured	Combined
LSTM	[14,58,61,64–68]	[69]	[6,54,55,57]
CNN	[14,33,66,68,70–72]	[63,73]	[54,55]
GRU	[31,38,72,74–81]	0	0
Autoencoder	[61,67,74,75,82,83]	[84]	[11]
word2vec	[59,85–87]	0	[56]
FFNN	[88,89]	0	[6,54]
Tensor Decomposition	[48,58]	0	0
Deep Averaging Network	0	[62,90]	0
Graph NN	[27,32,91]	0	0
Restricted Boltzmann Machines	[60]	0	0
Transformer	[43]	0	0

provide a degree of interpretability to the model, an attention mechanism is adopted in 13 studies [6,14,27,31,43,57,64,72,76–78,80,88]. A hierarchical knowledge graph has been incorporated into learning representations in 6 studies [27,31–33,81,91].

Sequential order between clinical events has been considered in the architecture in 27 studies. The majority of them have applied RNN (i.e., LSTM, GRU) to model sequential order with time-series patient data from diagnoses or measurements. Specifically, RNN adopts hidden states to carry information from the previous step to the next, which attempts to transmit sequential order between clinical events. Clinical events at each timestamp (i.e., diagnoses at each visit, or lab measures at every hour) are fed into every unit of the RNN, and the hidden state at this time point is updated with the hidden states from the previous timestamp and additional input from the current timestamp. Along these lines, sequential information from discrete event time-series can be transmitted to the final prediction target. The remaining 22 studies have not considered sequential information, based on varied ways of pre-processing patient data. For instance, deep learning models that fail to process sequential orders such as CNNs or autoencoders are applied in the architecture [73,98], the patient EHR data are only derived from clinical notes [62,90], and the patient EHR data are represented with one-hot encoding or frequency counts [11].

An interesting observation when speculating about what the “ideal” architecture for learning patient representations would be is that deep learning models are not mutually exclusive from each other. The best aspects of each model can be combined into a joint deep learning architecture. For instance, RNNs can assist with autoencoders to learn unsupervised representations while improving their ability to capture sequential clinical events [61,67,74,75]. CNNs can also be combined with RNNs to learn multifaceted patient representations [14,55,72].

4.2.5. Patient data resource

We have observed the recent success in the image processing field built upon the curation of a very large image dataset, ImageNet [92]. In this review, we investigated the resources of available EHR data that are large-scale and publicly available to advance research. Although it may be hard to access EHR data due to privacy concerns, a large publicly available clinical dataset can push representation performance to a scale similar to that of image processing by giving researchers a common benchmark to build on prior work. We observed that the prominent dataset in 49 included studies is the Medical Information Mart for Intensive Care (MIMIC-III) [93], where 22 out of 49 studies have used this dataset. Because MIMIC-III has both structured and unstructured data, clinical outcomes studied using MIMIC-III are widely varied, including mortality prediction, disease diagnoses, phenotype predictions, readmission, and length-of-stay forecasting. Other publicly available datasets that have been used to develop patient representations are: the Parkinson Progression Marker Initiative (PPMI) [67,74,94], the Alzheimer’s disease neuroimaging initiative (ADNI) [74,95], the i2b2 obesity challenge [62,63,90,96], and eICU Collaborative Database [42,97]. We summarize the public datasets and their corresponding clinical applications in Table 3. Considering the public datasets for each clinical task, the most popular is MIMIC-III for disease predictions ($n = 10$). Private datasets from local clinical data warehouses were used in 27 of the 49 studies. The detailed description of local hospital datasets can be found in Supplementary Table 2. A total of 9 studies have applied their methods on two or more datasets (i.e., private or public).

4.2.6. Clinical outcomes and applications

Studies involving patient representation learning evaluate their representations based on risk predictions. The general assumption is that a more advanced and robust patient representation would improve the performance of predictive models. Table 4 illustrates the clinical outcomes addressed among the 49 papers. We categorize main clinical tasks for each paper into seven types: disease prediction ($n = 31$), mortality prediction (12), length-of-stay forecasting (8), admission prediction (7), patient subtyping (4), intervention prediction (5), and medical cost forecasting (3). To show the generalizability of the study, the majority of works have applied two or more clinical outcomes to evaluate the representations, resulting in a number of studies that have sizable overlap in terms of outcomes studied. The most common clinical tasks (31 studies) are disease diagnosis predictions, which is to predict whether the patient will develop a given condition. Mortality prediction is the second most common task with 12 studies.

Rather than focusing on predictive tasks, patient representations can be potentially useful for clustering patients based on similarity, which can be used to identify patient subtypes, among other things. Patient subtyping is helpful for studying certain types of diseases, including Parkinson’s Disease. However, very few studies have focused on patient

Table 3

Benchmark dataset and clinical outcome applications.

Public Resource	Description	Patient Data Type	Patient Count	Application Tasks	Paper
MIMIC-III	A critical care database, including de-identified patient data with ~ 60,000 Intensive Care Unit (ICU) admissions [93]	Structured; Unstructured	~46,000	Mortality Prediction	[14,43,57,58,61,65,73,77]
				Disease Prediction	[27,31,43,54,56,57,59,91,98,99]
				Admission Prediction	[77,88]
				Length-of-stay Prediction	[14,43,73]
				Patient Similarity	[58]
				24-hour decompensation	[14,43,61]
				Intervention Prediction (ventilation, prescription, lab test order, etc.)	[32,54,55,65,98]
PPMI	A longitudinal patient dataset comprising clinical and behavioral variables, imaging, and specimen data of Parkinson's disease patients. [94]	Structured	~1000	Patient subtyping	[67,74]
ADNI	A longitudinal patient dataset consisting of assessments collected from selected patients in varied stages of Alzheimer's Disease. [95]	Structured	~600	Patient subtyping	[74]
i2b2 obesity challenge	1237 discharge summaries from the Partners HealthCare Research Patient Data Repository. Each discharge summary was annotated with patient disease status corresponding to obesity and fifteen comorbidities of obesity. [96]	Unstructured	~1000	Phenotype Prediction	[62,63,90]
eICU Collaborative Database	A combination of multiple critical care units across the United States. The data covers patients who were admitted to critical care units in 2014 and 2015. [97]	Structured	~139,000	Mortality medication, and diagnosis predictions	[42]

MIMIC-III: <https://mimic.physionet.org/>.PPMI: <http://www.ppmi-info.org/>.ADNI: <http://ida.loni.usc.edu/>.eICU: <https://eicu-crd.mit.edu/about/eicu/>.**Table 4**

Clinical outcome of interests with evaluation metric(s) and prediction period level.

Evaluation Tasks (# paper)	Evaluation Metrics	Time periods	Papers
Disease Prediction (31)	AUROC, AUPRC, Accuracy, Top K Recall/Precision	Year-level, visit-level	[6,11,27,31,33,38,43,54,56,57,59,60,62–64,66,67,69–72,75,76,80,81,83,86,89–91,98]
Mortality Prediction (12)	AUROC, AUPRC, F1, Accuracy	24-hour, month-level	[6,14,43,57,58,61,65,73,75,77,85,98]
Length-of-stay Prediction (8)	AUROC, AUPRC, F1, Accuracy, MSE, MAPE	12-hour, 24-hour	[6,14,43,73,82,83,85,87]
Admission Prediction (7)	AUROC, AUPRC, Accuracy, Top K Recall/Precision	Month-level, year-level	[6,38,48,77–79,88]
Patient Subtyping (4)	Significant difference among different groups; NMI, MSE, Sparsity and Similarity	Visit-level	[67,70,74,79]
Intervention Prediction (5)	Jaccard, Recall, Precision, F1	24-hour	[32,54,55,65,98]
Medical Cost (3)	R-squared, RMSE	Year-level	[48,85,87]

subtyping ($n = 4$). We assume this is because patient subtyping is not like other prediction tasks such as mortality and disease predictions that have straightforward label information. Typically, a data-driven technique is required to differentiate between different subtypes. For instance, the technique of measuring the distance of two patient representation vectors is to calculate the similarity; thus, to decide whether the two patients belong to the same subtype. Zhang et al. [100] proposed a data-driven subtyping method with LSTM and Dynamic Time Warping to first transform patients into temporal representations and then calculate the similarity of every patient pair.

Evaluation metrics are mainly dependent on the prediction task. We summarize the metrics for the above clinical outcomes in Table 4. Typically, for prediction tasks, AUROC, AUPRC, Precision, Recall, F1-measure, and Accuracy are the most used. Metrics including Top_K_Recall or Top_K_Precision (ranking metrics from information retrieval), are adopted particularly for predicting full sets of diagnoses

where K is the number of diseases. In addition, for regression tasks that involve predicting a real-valued quantity (i.e., medical cost, readmission forecasting), Root Mean Squared Error (RMSE), Mean Absolute Percentage Error (MAPE) and Goodness-of-Fit R-squared are utilized.

Another evident observation stands out when we look into the time periods for the evaluation, that is how to select appropriate observation and prediction periods. Normally, clinical events in the observation period are accumulated to construct the samples for training, while clinical events in the prediction period are considered as the gold standard for evaluation. We discover an interesting connection that the design of observation and prediction periods are highly correlated with the clinical outcome tasks that are evaluated. Table 4 shows the time period for prediction with corresponding clinical tasks. Prediction tasks, such as disease, readmission, patient subtyping, and medical cost forecasting, generally extend observation windows for long time periods. Month-level [69,88], year-level [11,48,66,78,80], and visit-level

predictions [31,38,56,59,64,70,71,76] are prevalent for those tasks. Take disease prediction as an example, a number of studies extracted diagnoses from the first visit and then predicted diagnoses in the next visit (i.e., visit-level observation) [31,38,56,59,64,70,71,76]. On the contrary, for those outcomes that happen in hospital, such as inpatient mortality [14,57,61,65], length-of-stay [14,83], and intervention prediction [32,55,65], observation windows at hour level are developed. Specifically, 24 hours are common as these tasks are immediate and physicians make decisions instantly. Rajkomar et al., [6] assessed two periods for mortality prediction (i.e., 12-hour and 24-hour) and showed that, if the deep learning model and other parameters remained the same, the prediction made 24 hours after admission attained a slightly better performance than that of 12 hours. This outcome is reasonable because more training features were accumulated with a longer time of observation. However, we assume this is a trade-off between time and accuracy, because the decision for this type of demanding task should be up-to-date as well.

4.2.7. Objective function

Only a few papers have reported the loss function applied in the study. We find that cross-entropy loss (or logarithmic loss) was the most common objective function for optimization. Aside from traditional cross-entropy loss, a weighted loss function was applied to account for an imbalanced-class dataset [55]. For regression tasks, mean squared error (MSE) loss and mean absolute error (MAE) loss were mainly performed [43]. What's more, studies developing patient representations with several learning objectives conduct a unified loss function with different targets being optimized jointly. For instance, a masked binary cross-entropy loss was employed for a multi-task architecture [43,69]. Doctor AI [38] has a joint loss function containing the cross-entropy loss from diagnosis predictions and the squared loss for the prediction of time duration for a single patient.

4.2.8. Interpretability

While patient representations may be relatively low-dimensional (often less than 500 dimensions), these collections of continuous numbers can be particularly abstruse and difficult for humans to interpret [101].

Visualization techniques are often performed to construct qualitative clusters of patient cohorts to provide an intuitive interpretation of the learned representations [102]. Such techniques for patient representations include t-Distributed Stochastic Neighbor Embedding (t-SNE) [103], Multidimensional Scaling (MDS) [104], Principal Component Analysis (PCA) [105], and Uniform Manifold Approximation and Projection (UMAP) [106]. Among these four methods, t-SNE has developed into the standard tool in this field due to its capability of revealing clusters in data. Ideally, patients typically from test sets associated with different labels, are projected and clustered into distinct groups. We observed 8 papers visualized with t-SNE to interpret the representations [14,27,31,60,66,73,75,79], and one study visualized with MDS [70]. Apart from projecting representations, interpretability can be provided by visualization methods including a heatmap of features [33,67,78], visualizing attention weights [57,72,76,80,88], and case study (explanation by examples) [32].

4.2.9. Computational resources and reproducibility

With the rapid improvements in computational resources, GPU-accelerated techniques have been exploited to train deep learning models faster and more efficiently. Because the literature in this study is from recent years (2015–2019), we observe the majority (29 out of 49) have reported the computational resources along with the deep learning frameworks. The reported deep learning frameworks include Theano (n = 12), TensorFlow (7), PyTorch (6), and Keras (4).

The need for reproducibility of scientific results has been growing in awareness in biomedicine in recent years. This is also true in this field of research because we want to ensure that the patient representations are

generalizable and reproducible across multiple institutions. Unfortunately, reproducibility in patient representation learning is difficult due to the variety of data resources, learning methods and architectures, and data preprocessing steps. Another major difficulty in data sharing is related to privacy concerns as inappropriate access to sensitive patient data might lead to patients' identity disclosure. While there are machine learning techniques that preserve data privacy [107–111], it is often impossible to share patient data due to privacy issues. Thus, providing workflow and computer code associated with publications is becoming increasingly common to enable the reproducibility of the method, if not the full work. One such open-source platform, GitHub, manages a robust infrastructure that is appropriate for sharing the full suite of intermediate experiments such as programming code, statistical results, etc. Therefore, we consider papers that provide code as one measurement of reproducibility. Among the 49 included papers, 20 of them have provided code links. Notably, the source code of RETAIN [80], based on Theano, has been starred on Github more than 100 times and forked over 50 times, and it has been redeveloped in other deep learning frameworks (TensorFlow, Keras). RETAIN was also commonly used as baselines in the subsequent series of advanced models (i.e., GRAM [27], Dipole [76], Health-ATM [72]). Thus, we anticipate that sharing open-source resources and code would motivate researchers to further explore innovations in a more scientific manner.

5. Discussion

5.1. Growing importance

In this review, we provide an overview of the current research into EHR patient representation learning. A total of 363 articles were assessed, with 49 studies from 2015 to 2019 meeting the full criteria for review. We observe a growing trend of developing deep learning-based patient representations from EHRs. The earliest attempt of learning patient representations with deep learning used Restricted Boltzmann Machines, proposed in 2015 [60]. An increasing number of works continued to contribute to improving the representation power of deep learning-derived representations.

Notable observations are identified throughout the review and can be addressed to the four questions we proposed at the end of Introduction.

1. Resources: We investigated that the resources of available EHR data are quite limited due to privacy concerns. The most prominent dataset is the Medical Information Mart for Intensive Care (MIMIC-III), where 22 out of 49 studies used this dataset.
2. Methods: Learning was mainly performed in a supervised manner optimized with cross-entropy loss. Recurrent Neural Networks were the most common deep learning model.
3. Applications: Disease prediction was the most common application and evaluation. We also identified a wide variety of clinical outcomes being addressed, including mortality prediction, admission prediction, length-of-stay forecasting, patient subtyping, intervention prediction, and medical cost forecasting.
4. Potential: Enhancements in patient representation learning techniques will be continuously growing for powering patient-level EHR analyses. We anticipate that researchers from diverse communities will leverage the richness and potential of EHR data, and will assist the capability of learning patient representation further.

Furthermore, analyses and evidence from data extractions of the 49 included papers suggest the importance and feasibility of learning comprehensive representations of patient EHR data. Studies that forego any attempt to utilize patient representations as an intermediate structure lose out on several advantages. Most notably, many representation learning models utilize—through unsupervised or self-supervised learning—very large amounts of EHR data. In contrast, traditional

predictive models focus on just the labeled information needed for the single task they are trained on. This is similar in essence to the difference between discriminative and generative machine learning models, where the latter attempt to model the full distribution of the data. Since learned representations attempt to gain a cohesive picture of a patient's data, these models may very well be more robust to small changes in the data distribution (e.g., when porting a model to a new institution without re-training). The extent to which this may be true requires further investigation, but it is worth noting that the field of natural language processing has become dominated recently by transformer-based models [41] that essentially follow this approach. Contextual word representations are learned on large corpora in a self-supervised manner ("pre-training") and then those representations form the basis on which prediction models are built ("fine-tuning").

The capability to develop effective patient representations from EHRs derives primarily from how to model the pertinent characteristics of EHR data. For instance, recent clinical events are more likely to contribute to the final label compared to earlier events [80]. Another challenge is that clinical events may be correlated with each other (i.e., diagnoses lead to treatments, disease co-occurrence) [32,81]. Leveraging this inherent structure is a key factor to improving the representation. This has resulted in a wide range of methods that can complement deep learning models to enhance representation power. These methods include developing additional components to models for capturing a hierarchical representation from event-level, to visit-level, to patient-level [81,112]; for encoding the longitudinal factors between clinical events [65,77]; for incorporating domain knowledge into the representation [27,31–33,91]; and for adopting multiple modalities of EHR data [11,14,57].

5.2. Future directions

Building upon current promising trends in learning meaningful patient representations, we believe the growth will continue, and the challenges of EHR data can be addressed in the following directions. We highlight a limited number of works published after the end of the review period (i.e., since January 2020) to exemplify these promising directions.

5.2.1. Methodology

With the pace of deep learning models diffusing into patient representation learning, algorithm and model development continues to enhance in order to adequately consider and model EHR data [113]. Some pilot works have applied advanced methods to model contextualized information into the learning architecture, such as the transformer model [39]. For instance, Choi et al. [42] proposed the Graph Convolutional Transformer (GCT) to learn hidden patterns of EHR data. This work is an advancement over works like MiME [81] in that GCT can capture hidden logical relationships between structured EHR data (e.g., which diagnoses lead to which treatment) that are completely missing in the raw data. There is also a recent NLP study attempting to show the expressive power of transformers in that they can be universal approximators of sequence-to-sequence aggregators with compact support [114], though this study was not in the clinical domain. Another groundbreaking technique fueled by language models highlights the potential to learn generic representations from raw data. Ethan et al. [115] proposed language model-based representations to learn structured data from EHRs. Si et al. [116] applied pre-training and fine-tuning techniques to learn general and transferable clinical language representations. Li et al. [44], inspired from the architecture of BERT, proposed BHERT, a transformer model for structured data in electronic health records. BHERT is scalable to perform well across a wide range of downstream predictions.

Advanced learning techniques including transfer learning, multi-task learning, and meta-learning can combat data scarcity and label insufficiency to some extent. One of the earliest attempts to learn patient

representations, Doctor AI [38], also demonstrated the impact of transfer learning from large-scale clinical data (i.e., 263,706 patients in Sutter Health) to tasks with relatively small amounts of data (i.e., 2,695 patients in MIMIC-III). Furthermore, with the success of multi-task learning in the open domain, such as in image processing and NLP, a growing number of studies also highlighted the potential and feasibility of properly integrating multiple related tasks to learn meaningful patient representations [73,89]. The concept of meta-learning is similar to transfer learning in that it takes advantage of information from other related high-resource domains. We have included a recent novel study in this review work, known as MetaPred [66]. MetaPred applied an optimization-based meta learning built on Model Agnostic Meta Learning (MAML) [117] in clinical scenarios. Because insufficient data is an obstacle for machine learning to solve many clinical problems, such as diagnosing rare conditions, we hope to leverage as many related datasets and knowledge resources as possible. Progress of advanced techniques for learning patient representations are responsible to take knowledge base into account to ensure the clinical relevance of problems and solutions [118].

5.2.2. Data sources

Despite the richness and potential of EHR data, few studies take full advantage of the heterogeneity of EHR data. Instead, many only use a subset of the data. For instance, many studies use diagnosis codes to represent patients, and ignore the real-valued measurements associated with the diagnoses, such as lab test values. Unstructured data, additionally, are often overlooked when developing patient representations (only 12 out of 49 studies used unstructured data). Future work will still be devoted to leveraging rich, yet varied, information in EHRs. One technique known as deep learning-based multimodal representation learning can be exploited to narrow the heterogeneity of EHR data among different modalities. Recently, a joint representation learning strategy named HORDE was proposed by Lee et al. [119], which applies several graphical modules to embed different types of data sources dynamically from EHR data.

The cause of many diseases is very complicated. Many factors, including inherited genetics, living environments, daily diets, and habits, have an impact on disease development. Just relying on the phenotype and medical history information in the EHR would generally fail to effectively represent the patient's full health status. Therefore, integrating multi-source information, e.g. genomics and clinical imaging, would be critical for accurate patient representation learning as well as the associated downstream clinical tasks [120]. Although most current work focuses on a single type of data, there are also some pilot works that are trying to harmonize multiple types of data. Cheerla et al. [121] developed a multimodal deep learning method to predict the survival of pan-cancer prognosis using clinical, multi-omics, and histopathology data together. Some researchers are also investigating the linkage of imaging to genetic data [122] and EHR to genetic data [123] for clinical outcome prediction and disease heritability estimation. These works show the enormous potential to integrate multi-source information to effectively model patient health status [22]. All these data sources can expand the set of features that potentially contribute to fine-grained patient representations.

5.2.3. Multi-Institutional data

Ethical and legal issues with regards to the use of EHR patient data for research, including AI, are well known. The problems lie in the trade-off between the privacy of patient data and data access for research. Yet learning representations from multiple institutions is critical to achieving both scalability and generalizability. Privacy-preserving algorithms and constrained domain adaptation settings are developed to compensate for this tradeoff to some degree [124]. Distributed learning and federated learning are such paradigms that enable the safe use of data from multiple institutions by only learning mathematical parameters [125] and avoiding any instance that might trace back to a specific

patient. Distributed learning is an algorithm that, in a parallelized fashion, iteratively learns a single model on separate datasets and obtains the shared model as if data were centralized [126]. Federated learning is a machine learning technique to train different models on multiple decentralized datasets, in contrast to traditional machine learning in which all samples are combined [127,128]. They both attempt to collaboratively train while keeping all the training data at their original institutions. The difference between the two learning paradigms is that distributed learning trains a separate model for each split of the data, and each separate model is transferred to a centralized model, while federated learning trains individual models on heterogeneous datasets. Many existing studies have exploited distributed learning and federated learning across multiple institutions in fields such as medical image processing [129,130], clinical decision support [131], and clinical oncology [132], to facilitate privacy-preserving multi-centric rapid learning of health care. We also observe some pilot studies applying federated learning and distributed learning to learn patient representations in this review. Li et al. [59] proposed a method that can learn clinical code representations from different EHR databases by adding distributed noise contrastive estimation. Liu et al. [90] proposed a federated learning framework to pre-train on MIMIC-III notes and predict clinical outcomes related to obesity. The adoption of distributed and federated learning shows that there are emerging technologies to manage privacy concerns and facilitate scientific research without directly sharing data.

5.3. Limitations

A potential limitation of this review is incomplete retrieval of relevant studies that meet the inclusion criteria. As a wide variety of methods to represent patients are included, it is challenging to conduct an inclusive search strategy with an automatic query by keyword searching. There is, unfortunately, no MeSH term that covers patient representation. We mitigated this issue by working cooperatively with experienced librarians and applying snowballing search strategies from the included publications. It is also challenging to differentiate our included studies with numerous works of predicting one clinical outcome end-to-end with deep learning. To achieve this, we ensured at the full-text review step that all the papers that are included specifically mentioned patient representation or embedding, and that their goal was learning representations at a patient level, not merely a model that maximized performance on a specific disease. In addition, we do not include manuscripts [133–136] that can only be found as preprints on arXiv or bioRxiv. All these limitations would pose a potential threat to selective bias in publication trends.

Another limitation is that our review fails to answer all technical questions. Due to the lack of transparency and reproducibility of reported results, many studies in our review claim state-of-the-art results, but few can be verified by external parties. This might be a barrier for future model development and would slow the pace of improvement. Even though MIMIC-III is widely used as a benchmark dataset, not many generic pipelines for benchmarking machine learning studies are available. Prior works have proposed different components in addition to the deep learning models to improve representation ability and predictive performance. However, it is still unclear which approach works best for representing EHR data. Therefore, one future direction to mitigate this limitation is to more comprehensively investigate different methods and conduct comparative studies using a common set of shared clinical benchmark datasets. Some recent comparative studies such as Ayala Solares et al. [137], Sadati et al. [138], and Min et al. [139] provided some surprising observations in terms of how to choose the best representation learning methods. For instance, Min et al. [139] conducted a case study of applying different machine learning methods to represent Chronic Obstructive Pulmonary Disease (COPD) patient claim data to predict readmission. They have shown some contradictory observations that medical problems are unlike problems in NLP and

image processing, merely applying complex deep learning without incorporating medical knowledge does not necessarily result in better performance. Therefore, in the future, we would encourage such studies to be conducted across a variety of patient representation use cases.

6. Conclusion

Deep representation learning has led to a wide variety of innovations in the process of modeling EHR patient data. As deep learning models benefit largely from the model capabilities to address the challenges of EHR data, deep patient representation learning is a promising direction to acquire powerful, robust, and precise representations. By adopting advanced learning techniques in addition to the model architecture, patient representation learning attempts to further address issues related to patient data and promote scientific research. We conducted a systematic review of this work and discussed the current research scenarios pertinent to patient representation learning. We believe a growing number of advanced methods will be continuously developed to learn meaningful patient representations, and that these representations will play a greater and greater role in clinical prediction tasks.

Author contributions statement

YS and KR conceived the review. YS, JD, ZL, and KR completed the initial screening. YS and ZL performed data extraction. YS conducted the data analysis and drafted the initial manuscript. XJ, TM, FW, and WJZ contributed to the design of the study and editing of the manuscript. KR provided overall leadership for the study and revised the manuscript. All authors approved the final version of the article.

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.

Acknowledgments

This work was supported by the U.S. National Library of Medicine, National Institutes of Health (NIH) under award R00LM012104, R01LM012973, UL1TR00316701, U01TR002062, R01GM118609, and R01AG06674901, the National Science Foundation under award NSF IIS-1750326, the Cancer Prevention Research Institute of Texas (CPRIT) under awards RP170668 and RR180012, as well as the Patient-Centered Outcomes Research Institute (PCORI) under award ME-2018C1-10963.

Appendix A. Supplementary material

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.jbi.2020.103671>.

References

- [1] J. Wu, J. Roy, W.F. Stewart, Prediction modeling using EHR data: Challenges, strategies, and a comparison of machine learning approaches, *Med. Care* 48 (2010) S106–S113, <https://doi.org/10.1097/MLR.0b013e3181de9e17>.
- [2] Y. Bengio, A. Courville, P. Vincent, Representation learning: A review and new perspectives, *IEEE Trans. Pattern Anal. Mach. Intell.* 35 (2013) 1798–1828, <https://doi.org/10.1109/TPAMI.2013.50>.
- [3] Y. LeCun, Y. Bengio, G. Hinton, Deep learning, *Nature* 521 (2015) 436–444, <https://doi.org/10.1038/nature14539>.
- [4] D. Svozil, V. Kvasnicka, J. Pospichal, Introduction to multi-layer feed-forward neural networks, *Chemomet. Intell. Lab. Syst.* 39 (1997) 43–62, [https://doi.org/10.1016/S0169-7439\(97\)00061-0](https://doi.org/10.1016/S0169-7439(97)00061-0).
- [5] Z. Che, D. Kale, W. Li, M.T. Bahadori, Y. Liu, Deep computational phenotyping, in: Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, 2015, pp. 507–516.
- [6] A. Rajkomar, E. Oren, K. Chen, A.M. Dai, N. Hajaj, M. Hardt, P.J. Liu, X. Liu, J. Marcus, M. Sun, P. Sundberg, H. Yee, K. Zhang, Y. Zhang, G. Flores, G. E. Duggan, J. Irvine, Q. Le, K. Litsch, A. Mossin, J. Tansuwan, D. Wang, J. Wexler,

J. Wilson, D. Ludwig, S.L. Volchenboum, K. Chou, M. Pearson, S. Madabushi, N. H. Shah, A.J. Butte, M.D. Howell, C. Cui, G.S. Corrado, J. Dean, Scalable and accurate deep learning with electronic health records, *Npj Digit. Med.* 1 (2018), <https://doi.org/10.1038/s41746-018-0029-1>.

[7] P. Vincent, H. Larochelle, Y. Bengio, P.-A. Manzagol, Extracting and composing robust features with denoising autoencoders, in: *Proceedings of the 25th International Conference on Machine Learning - ICML'08*, ACM Press, Helsinki, Finland, 2008, pp. 1096–1103. <https://doi.org/10.1145/1390156.1390294>.

[8] P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, P.-A. Manzagol, Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion, *J. Mach. Learn. Res.* 11 (2010) 3371–3408. <http://jmlr.org/papers/v11/vincent10a.html>.

[9] C. Doersch, Tutorial on variational autoencoders, *ArXiv Preprint ArXiv:* 1606.05908 (2016).

[10] S. Rifai, P. Vincent, X. Muller, X. Glorot, Y. Bengio, Contractive auto-encoders: explicit invariance during feature extraction, in: *Proceedings of the 28th International Conference on International Conference on Machine Learning*, 2011, pp. 833–840.

[11] R. Miotto, L. Li, B.A. Kidd, J.T. Dudley, Deep patient: An unsupervised representation to predict the future of patients from the electronic health records, *Sci. Rep.* 6 (2016) 26094, <https://doi.org/10.1038/srep26094>.

[12] Y. LeCun, B.E. Boser, J.S. Denker, D. Henderson, R.E. Howard, W.E. Hubbard, L. D. Jackel, Handwritten digit recognition with a back-propagation network, *Adv. Neural Inf. Process. Syst.* (1990) 396–404.

[13] Y. Kim, Convolutional Neural Networks for Sentence Classification, in: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2014, pp. 1746–1751.

[14] Y. Xu, S. Biswal, S.R. Deshpande, K.O. Maher, J. Sun, RAIM: Recurrent attentive and intensive model of multimodal patient monitoring data, in: *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, ACM, New York, NY, USA, 2018, pp. 2565–2573. <https://doi.org/10.1145/3219819.3220051>.

[15] Y. Cheng, F. Wang, P. Zhang, J. Hu, Risk prediction with electronic health records: A deep learning approach, in: *Proceedings of the 2016 SIAM International Conference on Data Mining*, SIAM, 2016, pp. 432–440.

[16] K. Simonyan, A. Zisserman, Very deep convolutional networks for large-scale image recognition, in: *International Conference on Learning Representations*, 2015.

[17] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.

[18] T. Mikolov, I. Sutskever, K. Chen, G.S. Corrado, J. Dean, Distributed representations of words and phrases and their compositionality, *Adv. Neural Inf. Process. Syst.* (2013) 3111–3119.

[19] T. Mikolov, K. Chen, G. Corrado, J. Dean, Efficient estimation of word representations in vector space, *ArXiv Preprint ArXiv:1301.3781* (2013).

[20] Y. Choi, C.Y.-I. Chiu, D. Sontag, Learning low-dimensional representations of medical concepts, in: *AMIA Summits on Translational Science Proceedings*, 2016, p. 41.

[21] L. Taslaman, B. Nilsson, A framework for regularized non-negative matrix factorization, with application to the analysis of gene expression data, *PLoS ONE* 7 (2012) e46331, <https://doi.org/10.1371/journal.pone.0046331>.

[22] G.L. Stein-O'Brien, R. Arora, A.C. Culhane, A.V. Favorov, L.X. Garmire, C. S. Greene, L.A. Goff, Y. Li, A. Ngom, M.F. Ochs, Y. Xu, E.J. Fertig, Enter the Matrix: Factorization Uncovers Knowledge from Omics, *Trends Genet.* 34 (2018) 790–805, <https://doi.org/10.1016/j.tig.2018.07.003>.

[23] F. Wang, N. Lee, J. Hu, J. Sun, S. Ebadollahi, A.F. Laine, A framework for mining signatures from event sequences and its applications in healthcare data, *IEEE Trans. Pattern Anal. Mach. Intell.* 35 (2012) 272–285.

[24] F. Wang, N. Lee, J. Hu, J. Sun, S. Ebadollahi, Towards heterogeneous temporal clinical event pattern discovery: a convolutional approach, in: *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2012, pp. 453–461.

[25] J. Zhou, F. Wang, J. Hu, J. Ye, From micro to macro: data driven phenotyping by densification of longitudinal electronic medical records, in: *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2014, pp. 135–144.

[26] C. Liu, F. Wang, J. Hu, H. Xiong, Temporal phenotyping from longitudinal electronic health records: A graph based framework, in: *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ACM, 2015, pp. 705–714.

[27] E. Choi, M.T. Bahadori, L. Song, W.F. Stewart, J. Sun, GRAM: Graph-based Attention Model for Healthcare Representation Learning, in: *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '17*, ACM Press, Halifax, NS, Canada, 2017, pp. 787–795. <https://doi.org/10.1145/3097983.3098126>.

[28] M. Niepert, M. Ahmed, K. Kutzkov, Learning convolutional neural networks for graphs, in: *International Conference on Machine Learning*, 2016, pp. 2014–2023.

[29] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Lio, Y. Bengio, Graph attention networks, *ArXiv Preprint ArXiv:1710.10903* (2017).

[30] A. Grover, J. Leskovec, node2vec: Scalable feature learning for networks, in: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016, pp. 855–864.

[31] F. Ma, Q. You, H. Xiao, R. Chitta, J. Zhou, J. Gao, KAME: Knowledge-based Attention Model for Diagnosis Prediction in Healthcare, in: *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, ACM, New York, NY, USA, 2018, pp. 743–752. <https://doi.org/10.1145/326920.3271701>.

[32] S. Wang, P. Ren, Z. Chen, Z. Ren, J. Ma, M. de Rijke, Order-free Medicine Combination Prediction with Graph Convolutional Reinforcement Learning, in: *Proceedings of the 28th ACM International Conference on Information and Knowledge Management - CIKM '19*, ACM Press, Beijing, China, 2019, pp. 1623–1632. <https://doi.org/10.1145/3357384.3357965>.

[33] J. Zhang, J. Gong, L. Barnes, HCNN: Heterogeneous Convolutional Neural Networks for Comorbid Risk Prediction with Electronic Health Records, in: *Proceedings of the Second IEEE/ACM International Conference on Connected Health: Applications, Systems and Engineering Technologies*, IEEE Press, Piscataway, NJ, USA, 2017, pp. 214–221. <https://doi.org/10.1109/CHASE.2017.80>.

[34] D. Bahdanau, K. Cho, Y. Bengio, Neural machine translation by jointly learning to align and translate, *ArXiv Preprint ArXiv:1409.0473* (2014).

[35] I. Sutskever, O. Vinyals, Q.V. Le, Sequence to sequence learning with neural networks, *Adv. Neural Inf. Process. Syst.* (2014) 3104–3112.

[36] K. Cho, B. Van Merriënboer, D. Bahdanau, Y. Bengio, On the properties of neural machine translation: Encoder-decoder approaches, *ArXiv Preprint ArXiv:1409.1259* (2014).

[37] S. Hochreiter, J. Schmidhuber, Long short-term memory, *Neural Comput.* 9 (1997) 1735–1780.

[38] E. Choi, M.T. Bahadori, A. Schuetz, W.F. Stewart, J. Sun, Doctor AI: Predicting Clinical Events via Recurrent Neural Networks, *ArXiv:1511.05942* [Cs]. (2015). <http://arxiv.org/abs/1511.05942> (accessed April 10, 2019).

[39] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, L. Kaiser, I. Polosukhin, Attention is all you need, *Adv. Neural Inf. Process. Syst.* (2017) 598–6008.

[40] A. Radford, K. Narasimhan, T. Salimans, I. Sutskever, Improving language understanding by generative pre-training, *Technical Report*, OpenAI, 2018.

[41] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding, in: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 2019, pp. 4171–4186.

[42] E. Choi, Z. Xu, Y. Li, M.W. Dusenberry, G. Flores, E. Xue, A.M. Dai, Learning the Graphical Structure of Electronic Health Records with Graph Convolutional Transformer, *Proceedings of the AAAI Conference on Artificial Intelligence*. (2020).

[43] H. Song, D. Rajan, J.J. Thiagarajan, A. Spanias, Attend and diagnose: Clinical time series analysis using attention models, in: *32nd AAAI Conference on Artificial Intelligence, AAAI 2018*, 2018, pp. 4091–4098.

[44] Y. Li, S. Rao, J.R.A. Solares, A. Hassaine, R. Ramakrishnan, D. Canoy, Y. Zhu, K. Rahimi, G. Salimi-Khorshidi, BEHRT: Transformer for electronic health records, *Sci. Rep.* 10 (2020) 7155, <https://doi.org/10.1038/s41598-020-62922-y>.

[45] L. Rasmy, Y. Xiang, Z. Xie, C. Tao, D. Zhi, Med-BERT: pre-trained contextualized embeddings on large-scale structured electronic health records for disease prediction, *ArXiv Preprint ArXiv:2005.12833* (2020).

[46] T.G. Kolda, B.W. Bader, Tensor decompositions and applications, *SIAM Rev.* 51 (2009) 455–500, <https://doi.org/10.1137/07070111X>.

[47] E.C. Chi, T.G. Kolda, On tensors, sparsity, and nonnegative factorizations, *SIAM J. Matrix Anal. Appl.* 33 (2012) 1272–1299, <https://doi.org/10.1137/110859063>.

[48] K. Yang, X. Li, H. Liu, J. Mei, G. Xie, J. Zhao, B. Xie, F. Wang, TaGiTeD: Predictive task guided tensor decomposition for representation learning from electronic health records, in: *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, AAAI Press, 2017, pp. 2824–2830.

[49] J.C. Ho, J. Ghosh, J. Sun, Marble: High-throughput phenotyping from electronic health records via sparse nonnegative tensor factorization, in: *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '14*, ACM Press, New York, New York, USA, 2014, pp. 115–124. <https://doi.org/10.1145/2623330.2623658>.

[50] H. He, J. Henderson, J.C. Ho, Distributed Tensor Decomposition for Large Scale Health Analytics, in: *The World Wide Web Conference on - WWW '19*, ACM Press, San Francisco, CA, USA, 2019, pp. 659–669. <https://doi.org/10.1145/330855.83313548>.

[51] M. Ouzzani, H. Hammady, Z. Fedorowicz, A. Elmagarmid, Rayyan—a web and mobile app for systematic reviews, *Syst. Rev.* 5 (2016) 210, <https://doi.org/10.1186/s13643-016-0384-4>.

[52] K.M. Ahmed, B. Al Dhubaib, Zoter: A bibliographic assistant to researcher, *J. Pharmacol. Pharmacotherap.* 2 (2011) 303.

[53] D. Moher, A. Liberati, J. Tetzlaff, D.G. Altman, Preferred reporting items for systematic reviews and meta-analyses: The PRISMA statement, *Int. J. Surg.* 8 (2010) 336–341, <https://doi.org/10.1016/j.ijssu.2010.02.007>.

[54] E. Zhang, R. Robinson, B. Pfahringer, Deep Holistic Representation Learning from EHR, in: 2018. <https://doi.org/10.1109/ISMICIT.2018.8573698>.

[55] H. Suresh, N. Hunt, A. Johnson, L.A. Celi, P. Szolovits, M. Ghassemi, Clinical intervention prediction and understanding with deep neural networks, *Mach. Learn. Healthc. Conf.* (2017) 322–337.

[56] T. Bai, A.K. Ch, B.L. Egleson, S. Vucetic, EHR phenotyping via jointly embedding medical concepts and words into a unified vector space, *BMC Med. Inf. Decis. Mak.* 18 (2018) 123. <http://www.embase.com/search/results?subaction=viewrecord&from=export&id=L625525781>.

[57] J. Kemp, A. Rajkomar, A.M. Dai, Improved Hierarchical Patient Classification with Language Model Pretraining over Clinical Notes, *ArXiv:1909.03039* [Cs, Stat]. (2019). <http://arxiv.org/abs/1909.03039> (accessed November 22, 2019).

[58] K. Yin, D. Qian, W.K. Cheung, B.C.M. Fung, J. Poon, Learning phenotypes and dynamic patient representations via RNN regularized collective non-negative tensor factorization, *AAAI* 33 (2019) 1246–1253, <https://doi.org/10.1609/aaai.v33i01.33011246>.

[59] Z. Li, K. Roberts, X. Jiang, Q. Long, Distributed learning from multiple EHR databases: Contextual embedding models for medical events, *J. Biomed. Inform.* 92 (2019) 103138, <https://doi.org/10.1016/j.jbi.2019.103138>.

[60] T. Tran, T.D. Nguyen, D. Phung, S. Venkatesh, Learning vector representation of medical objects via EMR-driven nonnegative restricted Boltzmann machines (eNRBM), *J. Biomed. Inform.* 54 (2015) 96–105, <https://doi.org/10.1016/j.jbi.2015.01.012>.

[61] H. Suresh, J.J. Gong, J. Guttag, Learning tasks for multitask learning: heterogeneous patient populations in the ICU, in: *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining - KDD '18*, (2018) 802–810, <https://doi.org/10.1145/3219819.3219930>.

[62] D. Dligach, T. Miller, Learning Patient Representations from Text, in: *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, 2018, pp. 119–123.

[63] D. Dligach, M. Afshar, T. Miller, Toward a clinical text encoder: pretraining for clinical natural language processing with applications to substance misuse, *J. Am. Med. Inform. Assoc.* (2019) ocz072, <https://doi.org/10.1093/jamia/ocz072>.

[64] T. Bai, B.L. Egleston, S. Zhang, S. Vučetić, Interpretable representation learning for healthcare via capturing disease progression through time, in: 2018, pp. 43–51, <https://doi.org/10.1145/3219819.3219904>.

[65] L. Liu, J. Shen, M. Zhang, Z. Wang, J. Tang, Learning the joint representation of heterogeneous temporal events for clinical endpoint prediction, in: 2018, pp. 109–116, <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85060476955&partnerID=40&md5=e3cf1382f1464164edc3f0dd4ab7baa>.

[66] X.S. Zhang, F. Tang, H.H. Dodge, J. Zhou, F. Wang, MetaPred: Meta-Learning for Clinical Risk Prediction with Limited Patient Electronic Health Records, in: *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, ACM, New York, NY, USA, 2019, pp. 2487–2495, <https://doi.org/10.1145/3292500.3330779>.

[67] I.M. Baytas, C. Xiao, X. Zhang, F. Wang, A.K. Jain, J. Zhou, Patient subtyping via time-aware LSTM networks, in: *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ACM, 2017, pp. 65–74.

[68] M. Rafiq, G. Keel, P. Mazzocato, J. Spaak, C. Savage, C. Guttmann, Deep learning architectures for vector representations of patients and exploring predictors of 30-day hospital readmissions in patients with multiple chronic conditions, 2019, https://doi.org/10.1007/978-3-030-12738-1_17.

[69] J. Liu, Z. Zhang, N. Razavian, Deep EHR: Chronic Disease Prediction Using Medical Notes, *ArXiv:1808.04928 [Cs, Stat.]* (2018). <http://arxiv.org/abs/1808.04928> (accessed April 9, 2019).

[70] Q. Suo, F. Ma, Y. Yuan, M. Huai, W. Zhong, J. Gao, A. Zhang, Deep patient similarity learning for personalized healthcare, *IEEE Trans. Nanobiosci.* 17 (2018) 219–227, <https://doi.org/10.1109/TNB.2018.2837622>.

[71] Z. Che, Y. Cheng, Z. Sun, Y. Liu, Exploiting Convolutional Neural Network for Risk Prediction with Medical Feature Embedding, *ArXiv:1701.07474 [Cs, Stat.]* (2017). <http://arxiv.org/abs/1701.07474> (accessed April 10, 2019).

[72] T. Ma, C. Xiao, F. Wang, Health-ATM: A deep architecture for multifaceted patient health record representation and risk prediction, in: 2018, pp. 261–269, <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85045143398&partnerID=40&md5=b63d906c942e23cb2d49ae4fd27fd>.

[73] Y. Si, K. Roberts, Deep patient representation of clinical notes via multi-task learning for mortality prediction, in: *AMIA Jt Summits Transl Sci Proc.*, 2019, 2019, pp. 779–788.

[74] Y. Zhang, H. Zhou, J. Li, W. Sun, Y. Chen, A Time-Sensitive Hybrid Learning Model for Patient Subgrouping, in: *2018 International Joint Conference on Neural Networks (IJCNN)*, IEEE, Rio de Janeiro, 2018, pp. 1–8, <https://doi.org/10.1109/IJCNN.2018.8488991>.

[75] L. Lei, Y. Zhou, J. Zhai, L. Zhang, Z. Fang, P. He, J. Gao, An Effective Patient Representation Learning for Time-series Prediction Tasks Based on EHRs, in: 2019, pp. 885–892, <https://doi.org/10.1109/BIBM.2018.8621542>.

[76] F. Ma, R. Chitta, J. Zhou, Q. You, T. Sun, J. Gao, Dipole: Diagnosis Prediction in Healthcare via Attention-based Bidirectional Recurrent Neural Networks, *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '17*. (2017) 1903–1911. <https://doi.org/10.1145/3097983.3098088>.

[77] L. Liu, H. Li, Z. Hu, H. Shi, Z. Wang, J. Tang, M. Zhang, Learning Hierarchical Representations of Electronic Health Records for Clinical Outcome Prediction, *ArXiv Preprint ArXiv:1903.08652*. (2019).

[78] J. Zhang, K. Kowsari, J.H. Harrison, J.M. Lobo, L.E. Barnes, Patient2Vec: A Personalized Interpretable Deep Representation of the Longitudinal Electronic Health Record, *IEEE Access* 6 (2018) 65333–65346, <https://doi.org/10.1109/ACCESS.2018.2875677>.

[79] C. Xiao, T. Ma, A.B. Dieng, D.M. Blei, F. Wang, Readmission prediction via deep contextual embedding of clinical concepts, *PLoS ONE* 13 (2018) e0195024, <https://doi.org/10.1371/journal.pone.0195024>.

[80] E. Choi, M.T. Bahadori, J. Sun, J. Kulas, A. Schuetz, W. Stewart, RETAIN: An interpretable predictive model for healthcare using reverse time attention mechanism, *Adv. Neural Inf. Process. Syst.* (2016) 3504–3512.

[81] E. Choi, C. Xiao, W. Stewart, J. Sun, Mime: Multilevel medical embedding of electronic health records for predictive healthcare, *Adv. Neural Inf. Process. Syst.* (2018) 4547–4557.

[82] C. Zhou, Y. Jia, M. Motani, J. Chew, Learning Deep Representations from Heterogeneous Patient Data for Predictive Diagnosis, in: *Proceedings of the 8th ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics - ACM-BCB '17*, ACM Press, Boston, Massachusetts, USA, 2017, pp. 115–123. <https://doi.org/10.1145/3107411.3107433>.

[83] C. Zhou, Y. Jia, M. Motani, Optimizing Autoencoders for Learning Deep Representations from Health Data, *IEEE J. Biomed. Health. Inf.* 23 (2019) 103–111, <https://doi.org/10.1109/JBHI.2018.2856820>.

[84] M. Sushil, S. Šuster, K. Luyckx, W. Daelemans, Patient representation learning and interpretable evaluation using clinical notes, *J. Biomed. Inform.* 84 (2018) 103–113, <https://doi.org/10.1016/j.jbi.2018.06.016>.

[85] J. Stojanovic, D. Gligorijevic, V. Radosavljevic, N. Djuric, M. Grbovic, Z. Obradovic, Modeling Healthcare Quality via Compact Representations of Electronic Health Records, in: *IEEE/ACM Trans. Comput. Biol. Bioinformatics*, 14, 2017, pp. 545–554, <https://doi.org/10.1109/TCBB.2016.2591523>.

[86] E. Choi, M.T. Bahadori, E. Seales, C. Coffey, M. Thompson, J. Bost, J. Tejedor-Sojó, J. Sun, Multi-layer Representation Learning for Medical Concepts, in: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '16*, ACM Press, San Francisco, California, USA, 2016, pp. 1495–1504. <https://doi.org/10.1145/2939672.2939823>.

[87] L. Cui, X. Xie, Z. Shen, Prediction task guided representation learning of medical codes in EHR, *J. Biomed. Inform.* 84 (2018) 1–10, <https://doi.org/10.1016/j.jbi.2018.06.013>.

[88] S. Barbieri, J. Kemp, O. Perez-Concha, S. Kotwal, M. Gallagher, A. Ritchie, L. Jorm, Benchmarking Deep Learning Architectures for Predicting Readmission to the ICU and Describing Patients-at-Risk, *Sci Rep.* 10 (2020) 1111, <https://doi.org/10.1038/s41598-020-58053-z>.

[89] D.Y. Ding, C. Simpson, S. Pfohl, D.C. Kale, K. Jung, N.H. Shah, The Effectiveness of Multitask Learning for Phenotyping with Electronic Health Records Data, in: *Par Symp Biocomput.*, 24, 2019, pp. 18–29.

[90] D. Liu, D. Dligach, T. Miller, Two-stage Federated Phenotyping and Patient Representation Learning, *ArXiv:1908.05596 [Cs]*. (2019). <http://arxiv.org/abs/1908.05596> (accessed September 20, 2019).

[91] A. Hosseini, T. Chen, W. Wu, Y. Sun, M. Sarrafzadeh, HeteroMed: Heterogeneous Information Network for Medical Diagnosis, in: *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, Association for Computing Machinery, New York, NY, USA, 2018, pp. 763–772, <https://doi.org/10.1145/3269206.3271805>.

[92] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, L. Fei-Fei, Imagenet: A large-scale hierarchical image database, in: *2009 IEEE Conference on Computer Vision and Pattern Recognition*, Ieee, 2009, pp. 248–255.

[93] A.E. Johnson, T.J. Pollard, L. Shen, H.L. Li-wei, M. Feng, M. Ghassemi, B. Moody, P. Szolovits, L.A. Celi, R.G. Mark, MIMIC-III, a freely accessible critical care database, *Sci. Data* 3 (2016) 160035.

[94] K. Marek, D. Jennings, S. Lasch, A. Siderowf, C. Tanner, T. Simuni, C. Coffey, K. Kieburtz, E. Flagg, S. Choudhury, The Parkinson progression marker initiative (PPMI), *Prog. Neurobiol.* 95 (2011) 629–635.

[95] S.G. Mueller, M.W. Weiner, L.J. Thal, R.C. Petersen, C. Jack, W. Jagust, J. Q. Trojanowski, A.W. Toga, L. Beckett, The Alzheimer's disease neuroimaging initiative, *Neuroimaging. Clin.* 15 (2005) 869–877.

[96] Ö. Uzuner, Recognizing obesity and comorbidities in sparse data, *J. Am. Med. Inform. Assoc.* 16 (2009) 561–570.

[97] T.J. Pollard, A.E.W. Johnson, J.D. Raffa, L.A. Celi, R.G. Mark, O. Badawi, The eICU Collaborative Research Database, a freely available multi-center database for critical care research, *Sci. Data* 5 (2018) 180178, <https://doi.org/10.1038/sdata.2018.178>.

[98] M. Sushil, S. Šuster, K. Luyckx, W. Daelemans, Unsupervised patient representations from clinical notes with interpretable classification decisions, *ArXiv:1711.05198 [Cs]*. (2017). <http://arxiv.org/abs/1711.05198> (accessed April 10, 2019).

[99] W. Wang, C. Guo, J. Xu, A. Liu, Bi-Dimensional Representation of Patients for Diagnosis Prediction, in: *2019 IEEE 43rd Annual Computer Software and Applications Conference (COMPSAC)*, IEEE, Milwaukee, WI, USA, 2019, pp. 374–379, <https://doi.org/10.1109/COMPSAC.2019.10235>.

[100] X. Zhang, J. Chou, J. Liang, C. Xiao, Y. Zhao, H. Sarva, C. Henchcliffe, F. Wang, Data-Driven Subtyping of Parkinson's Disease Using Longitudinal Clinical Records: A Cohort Study, *Sci. Rep.* 9 (2019) 797, <https://doi.org/10.1038/s41598-018-37545-z>.

[101] G.E. Hinton, R.R. Salakhutdinov, Reducing the dimensionality of data with neural networks, *Science* 313 (2006) 504–507.

[102] F. Wang, R. Kaushal, D. Khullar, Should Health Care Demand Interpretable Artificial Intelligence or Accept "Black Box" Medicine? *Ann. Intern. Med.* 172 (2020) 59, <https://doi.org/10.7326/M19-2548>.

[103] L. van der Maaten, G. Hinton, Visualizing data using t-SNE, *J. Mach. Learn. Res.* 9 (2008) 2579–2605.

[104] M.A. Cox, T.F. Cox, Multidimensional scaling, in: *Handbook of Data Visualization*, Springer, 2008, pp. 315–347.

[105] M. Ringnér, What is principal component analysis? *Nat. Biotechnol.* 26 (2008) 303–304.

[106] L. McInnes, J. Healy, J. Melville, UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction, *ArXiv:1802.03426 [Cs, Stat.]*. (2018). <http://arxiv.org/abs/1802.03426> (accessed April 21, 2020).

[107] M.N. Sadat, M.M. Al Aziz, N. Mohammed, F. Chen, X. Jiang, S. Wang, SAFETY: secure gwAs in federated environment through a hYbrid solution, in: *IEEE/ACM Trans. Comput. Biol. Bioinf.*, 16, 2018, pp. 93–102.

[108] H. Yu, X. Jiang, J. Vaidya, Privacy-preserving SVM using nonlinear kernels on horizontally partitioned data, in: Proceedings of the 2006 ACM Symposium on Applied Computing, 2006, pp. 603–610.

[109] Y. Kim, J. Sun, H. Yu, X. Jiang, Federated tensor factorization for computational phenotyping, in: Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2017, pp. 887–895.

[110] W. Dai, S. Wang, H. Xiong, X. Jiang, Privacy preserving federated big data analysis, in: Guide to Big Data Applications, Springer, 2018, pp. 49–82.

[111] J. Lee, J. Sun, F. Wang, S. Wang, C.-H. Jun, X. Jiang, Privacy-preserving patient similarity learning in a federated environment: development and analysis, *JMIR Med. Inf.* 6 (2018) e20.

[112] Y. Si, K. Roberts, Patient Representation Transfer Learning from Clinical Notes based on Hierarchical Attention Network, in: AMIA Joint Summits on Translational Science Proceedings. AMIA Joint Summits on Translational Science, 2020.

[113] F. Wang, L.P. Casalino, D. Khullar, Deep Learning in Medicine—Promise, Progress, and Challenges, *JAMA, Intern Med.* 179 (2019) 293, <https://doi.org/10.1001/jamainternmed.2018.7117>.

[114] C. Yun, S. Bhojanapalli, A.S. Rawat, S.J. Reddi, S. Kumar, Are Transformers universal approximators of sequence-to-sequence functions?, *ArXiv:1912.10077 [Cs, Stat.]* (2020). <http://arxiv.org/abs/1912.10077> (accessed May 5, 2020).

[115] E. Steinberg, K. Jung, J.A. Fries, C.K. Corbin, S.R. Pfahl, N.H. Shah, Language models are an effective representation learning technique for electronic health record data, *J. Biomed. Inf.* 113 (2020) 103637.

[116] Y. Si, E.V. Bernstein, K. Roberts, Generalized and Transferable Patient Language Representation for Phenotyping with Limited Data, *arXiv* (2021).

[117] C. Finn, P. Abbeel, S. Levine, Model-agnostic Meta-learning for Fast Adaptation of Deep Networks, in: Proceedings of the 34th International Conference on Machine Learning - Volume 70, JMLR.org, 2017, pp. 1126–1135. <http://dl.acm.org/citation.cfm?id=3305381.3305498>.

[118] J. Wiens, S. Saria, M. Sendak, M. Ghassemi, V.X. Liu, F. Doshi-Velez, K. Jung, K. Heller, D. Kale, M. Saeed, P.N. Ossorio, S. Thadaney-Israni, A. Goldberg, Do no harm: a roadmap for responsible machine learning for health care, *Nat. Med.* 25 (2019) 1337–1340, <https://doi.org/10.1038/s41591-019-0548-6>.

[119] D. Lee, X. Jiang, H. Yu, Harmonized representation learning on dynamic EHR graphs, *J. Biomed. Inform.* 106 (2020) 103426, <https://doi.org/10.1016/j.jbi.2020.103426>.

[120] S. Shilo, H. Rossman, E. Segal, Axes of a revolution: challenges and promises of big data in healthcare, *Nat. Med.* 26 (2020) 29–38, <https://doi.org/10.1038/s41591-019-0727-5>.

[121] A. Cheerla, O. Gevaert, Deep learning with multimodal representation for pancancer prognosis prediction, *Bioinformatics* 35 (2019) 1446–1454.

[122] X. Zhu, J. Yao, G. Xiao, Y. Xie, J. Rodriguez-Canales, E.R. Parra, C. Behrens, I.I. Wistuba, J. Huang, Imaging-genetic data mapping for clinical outcome prediction via supervised conditional gaussian graphical model, in: 2016 IEEE International Conference on Bioinformatics and Biomedicine (BIBM), IEEE, 2016, pp. 455–459.

[123] G. Jia, Y. Li, H. Zhang, I. Chattopadhyay, A.B. Jensen, D.R. Blair, L. Davis, P. N. Robinson, T. Dahlén, S. Brunak, others, Estimating heritability and genetic correlations from large health datasets in the absence of genetic data, *Nat. Commun.* 10 (2019) 1–11.

[124] E. Laparra, S. Bethard, T.A. Miller, Rethinking domain adaptation for machine learning over clinical language, *JAMIA Open* (2020) ooaa010, <https://doi.org/10.1093/jamiaopen/ooaa010>.

[125] J. Konečný, B. McMahan, D. Ramage, Federated Optimization: Distributed Optimization Beyond the Datacenter, *ArXiv:1511.03575 [Cs, Math.]* (2015). <http://arxiv.org/abs/1511.03575> (accessed May 7, 2020).

[126] F. Zerka, S. Barakat, S. Walsh, M. Bogowicz, R.T.H. Leijenaar, A. Jochems, B. Miraglio, D. Townend, P. Lambin, Systematic Review of Privacy-Preserving Distributed Machine Learning From Federated Databases in Health Care, *JCO Clin. Cancer Inf.* (2020) 184–200, <https://doi.org/10.1200/JCO.19.00047>.

[127] J. Xu, F. Wang, Federated Learning for Healthcare Informatics, *ArXiv:1911.06270 [Cs.]* (2019). <http://arxiv.org/abs/1911.06270> (accessed May 7, 2020).

[128] N. Rieke, J. Hancock, W. Li, F. Milletari, H. Roth, S. Albarqouni, S. Bakas, M.N. Galtier, B. Landman, K. Maier-Hein, S. Ourselin, M. Sheller, R.M. Summers, A. Trask, D. Xu, M. Baust, M.J. Cardoso, The Future of Digital Health with Federated Learning, *ArXiv:2003.08119 [Cs.]* (2020). <http://arxiv.org/abs/2003.08119> (accessed May 7, 2020).

[129] P. McClure, C.Y. Zheng, J. Kaczmarzyk, J. Rogers-Lee, S. Ghosh, D. Nielson, P.A. Bandettini, F. Pereira, Distributed Weight Consolidation: A Brain Segmentation Case Study, in: S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, R. Garnett (Eds.), Advances in Neural Information Processing Systems 31, Curran Associates, Inc., 2018, pp. 4093–4103. <http://papers.nips.cc/paper/7664-distributed-weight-consolidation-a-brain-segmentation-case-study.pdf>.

[130] K. Chang, N. Balachandar, C. Lam, D. Yi, J. Brown, A. Beers, B. Rosen, D.L. Rubin, J. Kalpathy-Cramer, Distributed deep learning networks among institutions for medical imaging, *J. Am. Med. Inform. Assoc.* 25 (2018) 945–954, <https://doi.org/10.1093/jamia/ocy017>.

[131] T.M. Deist, A. Jochems, J. van Soest, G. Nalbantov, C. Oberije, S. Walsh, M. Eble, P. Bulens, P. Coucke, W. Dries, A. Dekker, P. Lambin, Infrastructure and distributed learning methodology for privacy-preserving multi-centric rapid learning health care: euroCAT, *Clin. Transl. Radiat. Oncol.* 4 (2017) 24–31, <https://doi.org/10.1016/j.ctro.2016.12.004>.

[132] G. Price, M. van Herk, C. Faivre-Finn, Data Mining in Oncology: The ukCAT Project and the Practicalities of Working with Routine Patient Data, *Clin. Oncol.* 29 (2017) 814–817, <https://doi.org/10.1016/j.clon.2017.07.011>.

[133] S. Darabi, M. Kachuee, S. Fazeli, M. Sarrafzadeh, TAPER: Time-Aware Patient EHR Representation, *ArXiv:1908.03971 [Cs, Stat.]* (2019). <http://arxiv.org/abs/1908.03971> (accessed September 20, 2019).

[134] B. Hettige, Y.-F. Li, W. Wang, S. Le, W. Buntine, MedGraph: Structural and Temporal Representation Learning of Electronic Medical Records, *ArXiv:1912.03703 [Cs, Stat.]* (2020). <http://arxiv.org/abs/1912.03703> (accessed May 7, 2020).

[135] S. Darabi, M. Kachuee, M. Sarrafzadeh, Unsupervised Representation for EHR Signals and Codes as Patient Status Vector, *ArXiv:1910.01803 [Cs, Stat.]* (2019). <http://arxiv.org/abs/1910.01803> (accessed May 7, 2020).

[136] S. Dubois, N. Romano, D.C. Kale, N. Shah, K. Jung, Effective Representations of Clinical Notes, *ArXiv:1705.07025 [Cs, Stat.]* (2017). <http://arxiv.org/abs/1705.07025> (accessed April 15, 2019).

[137] J.R. Ayala Solares, F.E. Dilettia Raimondi, Y. Zhu, F. Rahimian, D. Canoy, J. Tran, A.C. Pinho Gomes, A.H. Payberah, M. Zottoli, M. Nazarzadeh, N. Conrad, K. Rahimi, G. Salimi-Khorshidi, Deep learning for electronic health records: A comparative review of multiple deep neural architectures, *J. Biomed. Inf.* 101 (2020) 103337, <https://doi.org/10.1016/j.jbi.2019.103337>.

[138] N. Sadati, M.Z. Nezhad, R.B. Chinnam, D. Zhu, Representation Learning with Autoencoders for Electronic Health Records: A Comparative Study, *ArXiv:1801.02961 [Cs, Stat.]* (2018). <http://arxiv.org/abs/1801.02961> (accessed October 21, 2019).

[139] X. Min, B. Yu, F. Wang, Predictive Modeling of the Hospital Readmission Risk from Patients' Claims Data Using Machine Learning: A Case Study on COPD, *Sci. Rep.* 9 (2019) 2362, <https://doi.org/10.1038/s41598-019-39071-y>.