

# Revisiting the MIMIC-IV Benchmark: Experiments Using Language Models for Electronic Health Records

Jesus Lovon-Melgarejo, Thouria Ben-Haddi, Jules Di Scala,  
Jose G. Moreno and Lynda Tamine

University of Toulouse, IRIT, 31000 Toulouse, France  
{firstname.lastname}@irit.fr

## Abstract

The lack of standardized evaluation benchmarks in the medical domain for text inputs can be a barrier to widely adopting and leveraging the potential of natural language models for health-related downstream tasks. This paper revisited an openly available MIMIC-IV benchmark for electronic health records (EHRs) to address this issue. First, we integrate the MIMIC-IV data within the Hugging Face *datasets* library to allow an easy share and use of this collection. Second, we investigate the application of templates to convert EHR tabular data to text. Experiments using fine-tuned and zero-shot LLMs on the mortality of patients task show that fine-tuned text-based models are competitive against robust tabular classifiers. In contrast, zero-shot LLMs struggle to leverage EHR representations. This study underlines the potential of text-based approaches in the medical field and highlights areas for further improvement.

**Keywords:** Large language models, MIMIC-IV benchmark, Text-based mortality classification

## 1. Introduction

Recent advancements in natural language processing (NLP) and information retrieval (IR) tasks have been significantly driven by Transformers-based models, such as BERT (Devlin et al., 2018) and RoBERTa (Liu et al., 2019). These models have been trained on raw linguistic information with minimal supervision. Furthermore, the emergence of large language models (LLMs), such as ChatGPT (Achiam et al., 2023) and Llama 2 (Touvron et al., 2023), has extended these capabilities by scaling in parameters size and training data. In the medical domain, applying LLMs has emerged as a novel tool for patients and healthcare practitioners (Meskó and Topol, 2023). For example, electronic health records (EHR), composed of non-linguistic information such as laboratory measurements, procedures, and medication codes, are translated into linguistic reports using these models (Van Veen et al., 2023). However, it is still unclear how useful EHR model representations are in non-linguistic tasks. Beyond privacy concerns, the critical issue preventing the broad adoption of LLMs in for this problem is effectively transforming patient structured information from the raw EHR format to a linguistic unstructured format that can leverage the potential of LLMs' text-based representations. Existing Transformer-based models for patient data, such as TransformEHR (Yang et al., 2023) and BEHRT (Li et al., 2020), have adapted their architecture to consider tabular input data. However, this process requires a costly pre-training step that does not take advantage of the advancements in improved LLMs and free EHR benchmarks such

as MIMIC IV (Johnson et al., 2023). The latter provides large-scale intensive care unit (ICU) patient data in a tabular form related to established cohorts used in different downstream tasks (e.g., mortality patient classification). Consequently, we argue that improving the accessibility of these resources to meet the models' evolution is crucial for the field.

In this paper, we propose a simple but effective methodology to standardize the MIMIC-IV benchmark towards using state of the art (SOTA) Transformer-based architectures (BERT, DistilBERT (Sanh et al., 2019) and RoBERTa), and LLMs (Llama 2, Meditron (Chen et al., 2023)) for health-related predictive tasks. For this purpose, we identify six main groups of features on the ICU data and propose a template-based data-to-text transformation. Thus, we are able to provide a text document input that summarizes the patient's ICU entry. Additionally, and for the sake of reproducibility, we provide a Hugging Face *datasets* object<sup>1</sup> that automatically produces a clinical cohort in the desired textual format<sup>2</sup>. Our main contributions are as follows: 1) A standard MIMIC-IV benchmark, integrated into the Hugging Face *datasets* library, allowing flexible use of the EHRs representation in health-related downstream tasks; 2) A comprehensive set of experiments using eight different models for evaluating the effectiveness of our revisited MIMIC-IV benchmark on the mortality classification task.

<sup>1</sup><https://huggingface.co/docs/datasets/index>

<sup>2</sup>Publicly available at <https://huggingface.co/datasets/thbndi/Mimic4Dataset>

## 2. Background and Related Work

### 2.1. MIMIC Collections and Benchmarks

The Medical Information Mart for Intensive Care (MIMIC) collection (Johnson et al., 2023; Johnson et al.) is one of the largest and most recent EHR datasets. It includes more than 250,000 patients admitted to intensive care at Beth Israel Deaconess at Boston’s Beth Israel Deaconess Medical Center. For each patient, details of the full journey in the ICU are available in a deidentified form for privacy concerns<sup>3</sup>. The current version is the MIMIC-IV collection (Gupta et al., 2022) which collect patient data between 2008-2019 and uses ICD-9 and ICD-10 versions of the International Classification of Diseases (ICD)<sup>4</sup> to list diagnoses and to link medical procedures to diagnoses.

In recent works, multiple benchmarks were proposed for the medical domain (Harutyunyan et al., 2019; Gupta et al., 2022; Wang et al., 2020) using MIMIC collections (Johnson et al., 2023, 2016). They appear as a mainstream mean of model comparability and reproducibility. The MIMIC-IV data pipeline (Gupta et al., 2022) is proposed to preprocess data for downstream tasks. This pipeline is able to transform raw data into a ready-to-use tabular representation of the patient’s data. Additionally, it provides the mapping to ICD as well as standard techniques for dimensionality reduction. Although a first step is the proposal of the benchmarks, we aim to go for two steps forward in this work by proposing the integration of the MIMIC IV benchmark into *datasets* of Hugging Face<sup>5</sup>, one of the largest hub ready-to-use datasets, as well as the possibility of using Transformer-based models (including LLMs) for predictive tasks on EHRs.

### 2.2. Transformers for EHRs

Transformers-based models of the general domain, such as BERT, have been adapted to the clinical domain using medical-related linguistic collections such as PubMed (BioBERT (Lee et al., 2020) and ClinicalBERT (Alsentzer et al., 2019)). Recently, efforts to encode non-linguistic information of EHRs to model patient data have emerged with models such as BEHRT (Li et al., 2020), MedBERT (Rasmy et al., 2021), and TransformEHR (Yang et al., 2023). These models encode different health modalities in flexible architectures. However, they require pre-training on large-scale datasets and do not benefit from the significant progress of

<sup>3</sup>With regard to the Safe Harbor provision of the HIPAA.

<sup>4</sup><https://www.who.int/standards/classifications/classification-of-diseases>

<sup>5</sup>Our implementation respects MIMIC’s access policies by asking the user to provide the original data.

Transformer-based models in the NLP literature. Furthermore, LLMs such as ChatGPT (Achiam et al., 2023), Llama 2 (Touvron et al., 2023), and its medical variant Meditron (Chen et al., 2023) have shown outstanding performance in different clinical tasks related to adapting non-linguistic health data, such as images and EHR diagnostics, into text (Meskó and Topol, 2023; Yeo et al., 2023). However, the exploration of this linguistic EHR representation for non-linguistic tasks, referred to as EHR downstream tasks, is limited. In order to bridge this gap, we present experiments on EHR data to explore their potential.

## 3. MIMIC-IV Benchmark Revisited

Here, we detail the pipeline and EHR data used, then we describe the templates proposed for transforming tabular EHR data into textual inputs.

### 3.1. The Pipeline

We rely on the MIMIC-IV benchmark to produce the standard evaluation framework for text. Thus, first, we integrated the recommended pre-processing guidelines in the *datasets* library and implemented all the features of the MIMIC-IV-Data-Pipeline<sup>6</sup> provided in a tabular form, as shown in the left side of Figure 1. After the preprocessing steps, we obtained a tabular representation that includes the demographic, current diagnosis features and time-series features related to labs, medications, procedures, and vitals, as show in Table 1.

Note that features like CHAR/LAB are given in time intervals, thus a reduction/expansion strategy must be applied to normalize the representation size. Data imputation is commonly applied by sampling data from a fixed number of time windows or even averaging values across a sequence of time windows. As shown in Section 4.2, we did not find large differences between sampling or averaging these features.

### 3.2. Proposed Templates

Finally, feature EHR data is transformed to text using a template-based strategy as shown in right side of Figure 1 and described below:

“The patient {ethnicity} {gender}, {age} years old, covered by {insurance} was diagnosed with {cond\_text}.” With {cond\_text} corresponding to the textual description from ICD10 of the diagnoses.

“The chart events measured were: {chart\_text}.” With chart\_text the list of biological measurements of the form: {mean\_val} for {feat\_label}, mean\_val

<sup>6</sup><https://github.com/healthylaife/MIMIC-IV-Data-Pipeline>

Figure 1: Dataset generation pipeline for the tabular format and the text format.

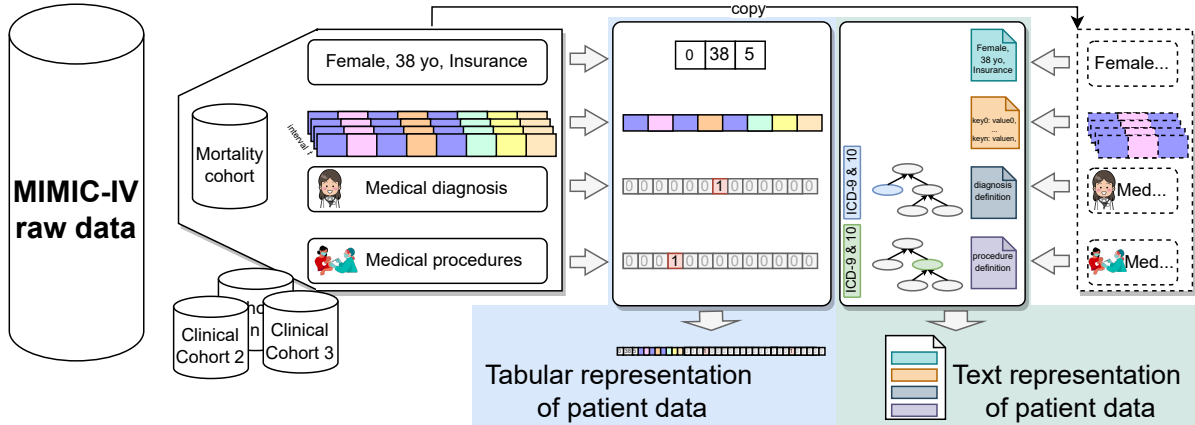


Table 1: Features list for the MIMIC-IV benchmark.

Name	Description
<i>Demo</i> <i>graphics</i> ( <i>DEMO</i> ):	The list of demographic data is a tiny vector corresponding to the patient’s gender, ethnicity, medical insurance, and age category. This data is encoded to obtain a numerical vector.
<i>Diagnosis</i> ( <i>COND</i> ):	The list of diagnoses established on a patient’s admission is encoded using a one-hot vector of all ICD codes including the patient’s identified diseases. Note that this vector could be large w.r.t. other features.
<i>Chart</i> <i>Events/Lab</i> ( <i>CHART/</i> <i>LAB</i> ):	Gives the value of the biological <i>item_id</i> performed in time interval <i>t</i> .
<i>Medications</i> ( <i>MEDS</i> ):	For each <i>item_id</i> corresponding to a medication the quantity administered in time interval <i>t</i> or zero if not administered.
<i>Procedures</i> ( <i>PROC</i> ):	The list of medical procedures performed is given as a form of a one-hot vector setting to 1 the <i>item_id</i> of procedures performed in time interval <i>t</i> .
<i>Output</i> <i>Events</i> ( <i>OUTE</i> ):	The list of biological samples taken is encoded using a one-hot vector of each <i>item_id</i> of the samples performed in time interval <i>t</i> .

is the mean value of the {feat\_label} measurement over the episode.

“The mean amounts of medications administered during the episode were: {meds\_text}.” With {meds\_text} the list of quantities of drugs administered of the form: {mean\_val} for {feat\_label}, {mean\_val} the average value over the episode of the quantity of drug {feat\_label}.

“The procedures performed were: {proc\_text}.” With {proc\_text} the list of medical procedures performed during the episode.

“The outputs collected were: {out\_text}.” With {out\_text} the list of biological prebiological samples taken during the episode.

Table 2 shows a sample of the textual input.

## 4. Experiments

### 4.1. Experimental Setup

To ensure a fair reproducibility of our experiments, we develop a *datasets* object that is able to produce tabular information as well as template-based textual data.

For **tabular data**, we create *Representation 1*, which follows the default configuration used in (Gupta et al., 2022), but other configurations are available in our implementation. Similarly, *Representation 2* is an aggregated representation of the same data. The main difference is the number of final features as the former uses 2766 features (as result of the concatenation of each window representation) and the latter 1110 features (as the values of all windows are averaged). We evaluated our revised MIMIC-IV benchmark on patient mortality classification as a pilot downstream task as provided in (Gupta et al., 2022). Evaluation focuses on both benchmark reproducibility (Cf. Section 4.2) and both feasibility and effectiveness using representative models (Cf. Section 4.3). Model parameters were selected using a 5-fold cross validation for classical machine learning algorithms available on Scikitlearn library<sup>7</sup>. We used algorithms for tabular data, such as Gradient Boosting (default parameters), XGBoost (objective=“binary:logistic”), Random Forest (n\_estimators=300, criterion=“gini”), and Logistic regression (default parameters).

For **textual data**, we fine-tuned six different Transformer-based models. We used optimal hyperparameters including learning rate of  $5e - 5$ ,

<sup>7</sup><https://scikit-learn.org/>

Table 2: Example of a text-based representation of a patient from the MIMIC-IV benchmark dataset. Values were changed to avoid leaking the example.

Feature	Example text
DEMO	The patient white male, 55 years old, covered by Other
COND	was diagnosed with Streptococcal sepsis; Acute pancreatitis; resistance to anti-microbial drugs.
CHAR/LAB	The chart events measured were: 73.655 for Heart Rate; 116.859 for Heart rate Alarm - High; ...
MEDS	The mean amounts of medications administered during the episode were: 44.778 of Albumin 5%; ...
PROC	The procedures performed were: Dialysis Catheter; 18 Gauge; EKG; ...
OUTE	The outputs collected were: OR EBL; OR Urine; Pre-Admission; ...

Table 3: Comparative evaluation of our standardized MIMIC-IV vs. original benchmark (Gupta et al., 2022) on the patient mortality classification task.

Algorithm	Representation 1		Representation 2		(Gupta et al., 2022)	
	AU-ROC	AU-PRC	AU-ROC	AU-PRC	AU-ROC	AU-PRC
<b>Gradient Boosting</b>	0.86	0.53	0.86	0.53	0.85	0.48
<b>XGBoost</b>	0.86	0.51	0.85	0.51	0.84	0.47
<b>Random Forest</b>	0.82	0.49	0.84	0.50	0.79	0.39
<b>Logistic Regression</b>	0.77	0.36	0.77	0.37	0.67	0.24

AdamW optimization and 3 epochs. For our zero-shot setup with LLMs, we explored multiple prompts. In the following, we report two of these prompts, which provide the highest number of valid responses for the task. We limited the output generation to 2 tokens.

We refer as *P1* for the prompt:

Prompt P1: "You are an extremely helpful health-care assistant. You answer the question using only 'yes' or 'no' and considering a patient hospital profile: '[textual EHR]'.  
Question: Is the patient dead?.  
Answer (only yes or no): "

Similarly, we refer as *P2* for the prompt:

Prompt P2: "Analyze the provided ICU data for a patient. The data covers the first 48 hours of the ICU stay, including vital statistics, lab test results, and treatments administered. Answer only Yes for a prediction of survival or No for a prediction of mortality. The patient ICU data is: '[textual EHR]'. Based on this data, answer.  
Question: Will the patient survive in the next 24 hours?.  
Answer (use only yes or no): "

We set a limit of 512 tokens for input length for

fine-tuned models and 1024 tokens for zero-shot models. It should be noted that this truncation only affected the fine-tuned models, and at times, it removed relevant information related to MEDS, PROC, and OUTE features. In Section 4.3, we discuss an ablation study that looks into the impact of these features.

## 4.2. Evaluation with Tabular EHR Data

Our results on tabular data and the reference values from the original benchmark (Gupta et al., 2022) are presented in Table 3. Note that our results are presented for two different aggregation strategies, *Representation 1* and *Representation 2*. In both cases, our results are slightly higher than those of the approach proposed in (Gupta et al., 2022) and used as a starting point. This is mainly due to our careful pre-processing of the data. As an important result, note that the *Representation 2* column performs similarly to *Representation 1* but uses significantly fewer features. Additionally, 1,034 values among 1,110 from the vector representation are sparse as they are dedicated to the diagnosis representation. These results lead us to pursue the text-based representation as only 66 values from biological signals combined with textual data (diagnosis) are enough to achieve state-of-the-art results on tabular data.

### 4.3. Evaluation of Using Template-based Text Inputs

Our main results on using text-based models for patient mortality classification tasks are presented in Table 4. For the fine-tuned models, we used the three general purpose trained models, namely DistilBERT (distilbert-base-uncased (Sanh et al., 2019)), BERT (bert-base-uncased (Devlin et al., 2018)), and RoBERTa (roberta-base (Liu et al., 2019)) (top three), and three others from the medical domain, namely BioClinicalBERT (Bio\_ClinicalBERT (Alsentzer et al., 2019)), BioBERT (dmis-lab/biobert-v1.1 (Lee et al., 2020)), and BiomedNLP (microsoft/BiomedNLP (Gu et al., 2021)) (bottom three). We reported only results with oversampling<sup>8</sup>. Our results show that the general purpose and domain-specific models behave similarly regarding AU-ROC, with all models getting close values (between 0.87 and 0.88). However, AU-PRC values differ as models from the medical domain outperform the general-purpose ones. Although a slight improvement was observed for general-purpose models in terms of AU-PRC, this is not enough to achieve the performance of the domain-specific models. Unsurprisingly, there is a clear interest in fine-tuning medical texts. However, general-purpose models, such as RoBERTa, closely follow top performances.

Furthermore, we explored using two LLMs, Llama2 (13b) (meta-llama/Llama-2-7b-hf (Touvron et al., 2023)) and its medical variant Meditron (7b) (epfl-llm/meditron-7b (Chen et al., 2023)) in a zero-shot setup considering two different prompts named *P1* and *P2*. We generally observed a lower performance from the Zero-shot section (as shown in Table 4) compared to Fine-tuned models. After analyzing the Zero-shot section, we found that prompt *P1* received better scores than *P2*. These results indicate that models are sensitive to the query format for this task. In addition, we noticed that domain-specific models, such as Meditron, performed better than general ones like Llama 2, using both prompts, similar to the fine-tuned setup. These findings suggest that SOTA LLMs struggle to encode and transfer EHR representations to downstream tasks within the explored prompts. A possible development towards using LLMs with tabular data is to define better translation methods to integrate this structured knowledge into the language models. Also, these findings motivate further research and experimentation by applying alternative techniques such as in-context learning (Dong et al., 2022) or prompt-tuning (Lester et al., 2021).

Moreover, in this setup, in addition to right or wrong answers, we also consider unanswered questions. Such questions occur when the LLM

Table 4: Results of the general purpose and medical domain models on the patient mortality task using text representations of patient data.

Models	AU-ROC	AU-PRC
<b>Fine-tuned</b>		
DistilBERT	0.87	0.42
BERT	0.87	0.43
RoBERTa	0.88	0.47
BioClinicalBERT	0.87	0.43
BioBERT	0.88	0.45
BiomedNLP	0.88	0.46
<b>Zero-shot with prompt P1</b>		
Llama 2 (13b)	0.50	0.38
Meditron (7b)	0.61	0.39
<b>Zero-shot with prompt P2</b>		
Llama 2 (13b)	0.50	0.13
Meditron (7b)	0.51	0.23

Table 5: Number of answered and unanswered samples by the LLMs for the zero-shot setup.

Model	# answered	# unanswered
<b>With prompt P1</b>		
Llama 2 (13b)	5952 (96.70%)	203 (3.30%)
Meditron (7b)	6152 (99.96%)	3 (0.04%)
<b>With prompt P2</b>		
Llama 2 (13b)	1885 (30.63%)	4270 (69.37%)
Meditron (7b)	6155 (100.0%)	0 (0.00%)

fails to provide an output from the expected tokens, which are “Yes” or “No” in our case. For our experiments, we consider “No” the default answer for the results reported in Table 4. To provide further details, we display the number of answered and unanswered questions per model in Table 5. Upon analysis, we found that the Llama 2 model left 3.30% of the dataset unanswered, while the Meditron model left only 0.04% unanswered using prompt *P1*. In contrast, the prompt *P2* obtained 69.37% of unanswered questions with Llama 2 and no unanswered questions by Meditron.

By comparing the different prompts used to describe the task, we can observe that Llama 2 (general domain model) struggles to understand the task while making some modifications. In contrast, Meditron (domain-specific models) is more stable when using different reformulations of the task.

We further our analysis by performing an ablation study with two representative models, BERT and BiomedNLP, to study the accumulative effect of the features. Results are presented in Table 6. As a main feature, we can easily identify COND as a clear buster in performance. This feature alone

<sup>8</sup>We found similar results without oversampling.

Table 6: Ablation study using different text features. ‘✓’ indicates that the feature was used in the patient representation. CH/LA stands for CHART/LAB. Results in bold indicate best performance.

Features	COND	✓	✓	✓	✓	✓	✓	
	DEMO			✓	✓	✓	✓	✓
CH/LA				✓	✓	✓	✓	✓
MEDS					✓	✓	✓	
PROC						✓	✓	
OUTE							✓	

	AU-ROC							
BERT	0.87	0.88	0.86	<b>0.89</b>	<b>0.89</b>	0.88	0.75	
BiomedNLP	0.87	0.88	<b>0.88</b>	0.87	0.86	0.88	0.63	

	AU-PRC							
BERT	0.41	0.44	0.40	<b>0.46</b>	<b>0.46</b>	<b>0.46</b>	0.24	
BiomedNLP	0.45	0.46	<b>0.50</b>	0.42	0.39	0.47	0.13	

achieves a close value to top performance, indicating that it is a clear signal of the patient profile representation. However, other no-expert-based features, such as *CHAR/LAB*, are also reliable. Note that this is an encouraging result as the features are given in an aggregated form. In fact, compared to the best model, the model can perform correctly in terms of AU-ROC. Also, note that both models achieve the best performance before integrating all the features. In particular, *MEDS*, *PROC*, and *OUTE* (only for BiomedNLP) do not improve previous combinations. This indicates that more elaborated templates are worth investigating to integrate these features.

## 5. Conclusion

In this paper, we presented a publicly available Hugging Face *datasets* object that allows a reproducible way to use the MIMIC-IV benchmark for EHR representation and use in health-related tasks based on text-based model. Using MIMIC IV in the proposed object, we aim to facilitate the experimentation with a comprehensive public EHR dataset in its original tabular form and text format through text-based templates. Our experiments showed that fine-tuned text-based models perform similarly to the strongest tabular-based alternatives regarding AU-ROC. On the contrary, LLMs in a zero-shot setup suggested limitations when encoding EHR information. This evaluation provides a starting point for a new family of large language models toward improving the current SOTA on health-related predictive tasks.

## 6. Bibliographical References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altmenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Emily Alsentzer, John Murphy, William Boag, Weihung Weng, Di Jindi, Tristan Naumann, and Matthew McDermott. 2019. Publicly available clinical bert embeddings. In *Proceedings of the 2nd Clinical Natural Language Processing Workshop*, pages 72–78.
- Zeming Chen, Alejandro Hernández-Cano, Angelika Romanou, Antoine Bonnet, Kyle Matoba, Francesco Salvi, Matteo Pagliardini, Simin Fan, Andreas Köpf, Amirkeivan Mohtashami, Alexandre Sallinen, Alireza Sakhaeirad, Vinitra Swamy, Igor Krawczuk, Deniz Bayazit, Axel Marmet, Syrielle Montariol, Mary-Anne Hartley, Martin Jaggi, and Antoine Bosselut. 2023. *Meditron-70b: Scaling medical pretraining for large language models*. *arXiv preprint arXiv:2311.16079*.
- E. Choi, Zhen Xu, Yujia Li, Michael W. Dusenberry, Gerardo Flores, Yuan Xue, and Andrew M. Dai. 2019. *Learning the graphical structure of electronic health records with graph convolutional transformer*. In *AAAI Conference on Artificial Intelligence*.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Sidhartha Brahma, et al. 2022. Scaling instruction-

- finetuned language models. *arXiv preprint arXiv:2210.11416*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and Zhifang Sui. 2022. A survey on in-context learning. *arXiv preprint arXiv:2301.00234*.
- Yu Gu, Robert Tinn, Hao Cheng, Michael Lucas, Naoto Usuyama, Xiaodong Liu, Tristan Naumann, Jianfeng Gao, and Hoifung Poon. 2021. Domain-specific language model pretraining for biomedical natural language processing. *ACM Transactions on Computing for Healthcare (HEALTH)*, 3(1):1–23.
- Mehak Gupta, Brennan Gallamoza, Nicolas Cutrona, Pranjal Dhakal, Raphael Poulain, and Rahmatollah Beheshti. 2022. [An Extensive Data Processing Pipeline for MIMIC-IV](#). In *Proceedings of the 2nd Machine Learning for Health symposium*, volume 193 of *Proceedings of Machine Learning Research*, pages 311–325. PMLR.
- Hrayr Harutyunyan, Hrant Khachatryan, David C. Kale, Greg Ver Steeg, and Aram Galstyan. 2019. [Multitask learning and benchmarking with clinical time series data](#). *Scientific Data*, 6(1):96.
- Alistair EW Johnson, Lucas Bulgarelli, Tom J Pollard, Steven Horng, L A Celi, and R Mark. [MIMIC-IV version 2.2](#).
- Alistair EW Johnson, Lucas Bulgarelli, Lu Shen, Alvin Gayles, Ayad Shammout, Steven Horng, Tom J Pollard, Sicheng Hao, Benjamin Moody, Brian Gow, et al. 2023. MIMIC-IV, a freely accessible electronic health record dataset. *Scientific data*, 10(1):1.
- Alistair EW Johnson, Tom J Pollard, Lu Shen, Li-wei H Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. 2016. MIMIC-III, a freely accessible critical care database. *Scientific data*, 3(1):1–9.
- Mihir Kale and Abhinav Rastogi. 2020. Text-to-text pre-training for data-to-text tasks. *arXiv preprint arXiv:2005.10433*.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2020. Biobert: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4):1234–1240.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*.
- Yikuan Li, Shishir Rao, José Roberto Ayala Solares, Abdelaali Hassaine, Rema Ramakrishnan, Dexter Canoy, Yajie Zhu, Kazem Rahimi, and Ghulamreza Salimi-Khorshidi. 2020. Behrt: transformer for electronic health records. *Scientific reports*, 10(1):7155.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Sean MacAvaney, Andrew Yates, Sergey Feldman, Doug Downey, Arman Cohan, and Nazli Goharian. 2021. Simplified data wrangling with `ir_datasets`. In *SIGIR*.
- Bertalan Meskó and Eric J Topol. 2023. The imperative for regulatory oversight of large language models (or generative ai) in healthcare. *NPJ digital medicine*, 6(1):120.
- Laila Rasmy, Yang Xiang, Ziqian Xie, Cui Tao, and Degui Zhi. 2021. Med-bert: pretrained contextualized embeddings on large-scale structured electronic health records for disease prediction. *NPJ digital medicine*, 4(1):86.
- Alexey Romanov and Chaitanya Shivade. 2018. Lessons from natural language inference in the clinical domain. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1586–1596.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *ArXiv*, abs/1910.01108.
- Irene Solaiman, Miles Brundage, Jack Clark, Amanda Askill, Ariel Herbert-Voss, Jeff Wu, Alec Radford, Gretchen Krueger, Jong Wook Kim, Sarah Kreps, et al. 2019. Release strategies and the social impacts of language models. *arXiv preprint arXiv:1908.09203*.
- Huan Song, Deepta Rajan, Jayaraman J. Thiagarajan, and Andreas Spanias. 2017. Attend and diagnose: Clinical time series analysis using attention models. In *Proceedings of the 2018 AAAI Association for the Advancement of Artificial Intelligence*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava,

Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.

Dave Van Veen, Cara Van Uden, Louis Blanke-meier, Jean-Benoit Delbrouck, Asad Aali, Christian Bluethgen, Anuj Pareek, Malgorzata Polacin, Eduardo Pontes Reis, Anna Seehofnerova, et al. 2023. Clinical text summarization: Adapting large language models can outperform human experts. *Research Square*.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.

Shirly Wang, Matthew BA McDermott, Geeticka Chauhan, Marzyeh Ghassemi, Michael C Hughes, and Tristan Naumann. 2020. Mimic-extract: A data extraction, preprocessing, and representation pipeline for mimic-iii. In *Proceedings of the ACM conference on health, inference, and learning*, pages 222–235.

Zhichao Yang, Avijit Mitra, Weisong Liu, Dan Berlowitz, and Hong Yu. 2023. Transformehr: transformer-based encoder-decoder generative model to enhance prediction of disease outcomes using electronic health records. *Nature Communications*, 14(1):7857.

Yee Hui Yeo, Jamil S Samaan, Wee Han Ng, Peng-Sheng Ting, Hirsh Trivedi, Aarshi Vipani, Walid Ayoub, Ju Dong Yang, Omer Liran, Brennan Spiegel, et al. 2023. Assessing the performance of chatgpt in answering questions regarding cirrhosis and hepatocellular carcinoma. *Clinical and molecular hepatology*, 29(3):721.