

REVIEW

# Clinical insights: A comprehensive review of language models in medicine

Nikita Neveditsin<sup>1</sup>, Pawan Lingras<sup>1</sup>, Vijay Mago<sup>2\*</sup>

**1** Department of Mathematics and Computing Science, Saint Mary's University, Halifax, Nova Scotia, Canada, **2** School of Health Policy and Management, York University, Toronto, Ontario, Canada

\* [vmago@yorku.ca](mailto:vmago@yorku.ca)



## Abstract

This paper explores the advancements and applications of language models in health-care, focusing on their clinical use cases. It examines the evolution from early encoder-based systems requiring extensive fine-tuning to state-of-the-art large language and multimodal models capable of integrating text and visual data through in-context learning. The analysis emphasizes locally deployable models, which enhance data privacy and operational autonomy, and their applications in tasks such as text generation, classification, information extraction, and conversational systems. The paper also highlights a structured organization of tasks and a tiered ethical approach, providing a valuable resource for researchers and practitioners, while discussing key challenges related to ethics, evaluation, and implementation.

## OPEN ACCESS

**Citation:** Neveditsin N, Lingras P, Mago V (2025) Clinical insights: A comprehensive review of language models in medicine. PLOS Digit Health 4(5): e0000800. <https://doi.org/10.1371/journal.pdig.0000800>

**Editor:** Jie Xu, University of Florida, UNITED STATES OF AMERICA

**Received:** August 01, 2024

**Accepted:** February 25, 2025

**Published:** May 8, 2025

**Copyright:** © 2025 Neveditsin et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

**Funding:** The authors received no specific funding for this work.

**Competing interests:** The authors have declared that no competing interests exist.

## 1. Introduction

The advances in the area of artificial intelligence (AI) in recent years has opened countless opportunities for various sectors, including healthcare. The potential influence of AI is a subject of debate concerning its impact on humanity. Leading AI experts have called for caution, evidenced by an open letter urging a pause in the expansion of advanced AI models, which reflects growing concerns among policymakers and the public about the ethical, social, and economic ramifications of AI. While some argue that AI can bring substantial advances in efficiency and effectiveness across many sectors, others fear it could exacerbate inequalities, displace jobs, and challenge societal norms [1]. While AI in healthcare has an extensive history of research [2], the emergence of advanced foundational models [3] such as GPT family [4], Gemini [5], and a series of open models like the Llama [6], offers unprecedented perspectives for the transformation of the healthcare sector.

This review paper synthesizes and critically examines the landscape of language models in the medical domain, with a particular emphasis on their clinical applications. Compared to other reviews in this area, it highlights the following aspects:

1. A dedicated focus on clinical applications.
2. An emphasis on locally deployable models.
3. A structured organization based on NLP task categories, with appropriate justification for this framework.

4. A structured ethical assessment of reviewed tasks.
5. A detailed discussion of evaluation challenges.

Table 1 outlines how these distinguishing aspects compare to the focus areas of other reviews in the field.

This review navigates the advancements from early encoder-based systems to state-of-the-art large language and multimodal models. It places a particular emphasis on locally deployable models, which enhance data privacy and adaptability within healthcare settings. Additionally, the structured exploration of ethical considerations and evaluation challenges provides a critical perspective on the current state of these models. By offering both a clear roadmap and valuable insights, this review serves as a comprehensive resource for newcomers and experienced interdisciplinary researchers, fostering innovation and promoting responsible implementation in this vital domain.

## 1.1. Definitions

**Artificial Intelligence (AI):** A multidisciplinary field aimed at creating systems capable of tasks that typically require human intelligence, such as problem-solving, learning, reasoning, and perception [12].

**Machine Learning (ML):** A branch of AI focused on developing algorithms and statistical models that enable computer systems to improve performance on specific tasks by learning from data and identifying patterns, rather than relying on explicit programming [12,13].

**Natural Language Processing (NLP):** A branch of AI that deals with the interaction between computers and human language, allowing machines to understand, interpret, and generate human languages [12,14].

**Language Models:** Probabilistic models that learn the statistical structure of word sequences, enabling the prediction or generation of text based on patterns observed in previously seen data [14].

**Large Language Models (LLMs):** Language models trained on extensive text corpora using deep learning architectures, enabling them to acquire broad linguistic knowledge and perform diverse natural language tasks [14,15].

**Generative Large Language Models (Generative LLMs):** A subclass of LLMs that generate coherent and contextually relevant text by learning statistical patterns from large-scale language datasets, enabling tasks such as text generation, summarization, and dialogue creation.

**Open LLMs:** LLMs developed and distributed with open access to their architectures and weights, facilitating transparency, reproducibility, and community-driven improvements [16].

Table 1. Unique aspects of this review compared to others in the field.

| Criterion  | Luo et al. (2024) [7] | Meng et al. (2024) [8] | Wang et al. (2023) [9] | Thirunavukaras et al. (2023) [10] | Yang et al. (2023) [11] |
|--|-----------------------|------------------------|------------------------|-----------------------------------|-------------------------|
| Focus on clinical applications                       | Partially             | Partially              | No                     | Yes                               | Yes                     |
| Emphasis on locally deployable models                | Partially             | No                     | Yes                    | No                                | Partially               |
| Structured by NLP task categories with justification | Yes                   | Partially              | Yes                    | No                                | Partially               |
| Structured ethical assessment of reviewed tasks      | No                    | Partially              | No                     | Partially                         | Partially               |
| Detailed discussion of evaluation challenges         | Partially             | Partially              | Partially              | Partially                         | Partially               |

<https://doi.org/10.1371/journal.pdig.0000800.t001>

**Foundation Models:** Large-scale, general-purpose AI models, such as LLMs, trained on massive and diverse datasets. They serve as a base for a wide range of downstream tasks, often requiring minimal additional training or fine-tuning [3].

**Embeddings:** Dense vector representations of words, phrases, or other linguistic units that capture semantic relationships and similarity based on contextual usage in large corpora [14].

**API (Application Programming Interface):** In the context of LLMs, an API provides a set of tools and protocols that enable users to interact with a language model remotely, without requiring local deployment. Typically, API access abstracts away the model's internals, offering only input-output functionality.

**GPU (Graphics Processing Unit):** A specialized processor originally designed for rendering graphics, now widely used to accelerate parallelizable computations, particularly in AI and ML tasks. Locally deployed models often require GPUs to run efficiently due to their ability to handle large-scale matrix operations.

## 2. Background

### 2.1. Language models

A critical milestone in the field of Natural Language Processing (NLP) was the introduction of the attention mechanism in neural machine translation [17]. This mechanism, linking the encoder and decoder in sequence-to-sequence models, paved the way for subsequent advancements by enabling the model to focus on different parts of the input sequence for each step of the output, substantially improving the handling of longer input sequences and complex dependencies. Notably, it led to the creation of the Transformer model [18], which exclusively relies on attention mechanisms. This innovation revolutionized not only the field of NLP but also the broader realms of AI and machine learning (ML).

The emergence of Transformer-based models has spurred the development of a wide array of *language models*, both in commercial and open formats. Among these, the GPT series [19] has advanced text generation capabilities while establishing the foundation for interactive applications. A notable milestone was the introduction of ChatGPT [20], a conversational agent that revolutionized chat-based language models by facilitating intuitive, human-like dialogue generation. The success of ChatGPT and similar systems [5,21] underscored the adaptability of Transformer architectures and catalyzed the development of language models optimized for interactive and specialized applications, including those in the medical domain.

The overall task of language modeling can be expressed as estimating the joint probability of a sequence of words  $w_1, w_2, \dots, w_n$  in a sentence, drawn from large text corpora:

$$P(w_1, w_2, \dots, w_n) = \prod_{i=1}^n P(w_i \mid w_1, w_2, \dots, w_{i-1})$$

Despite the simplicity of statistical approaches to language modeling and their initial lack of attention to the underlying rules of language, as critiqued by Chomsky [22], contemporary large language models (LLMs) like GPT-4 have demonstrated remarkable proficiency in a wide range of language understanding tasks [4], exhibiting emergent abilities and an advanced capacity to learn natural language patterns.

Modern language models have undergone a notable paradigm shift, moving from a pre-training and fine-tuning approach to embracing in-context learning (ICL) [23] and zero-shot learning [24]. Traditionally, language model development has followed a two-step process: pre-training and fine-tuning. During pre-training, a model learns general language representations from large text corpora in an unsupervised manner. For example, BERT [25] predicts

masked tokens, while GPT-2 [19] predicts the next token in a sequence. The fine-tuning phase then adapts these pre-trained models to specific tasks (e.g., classification, summarization) using labeled data. Although effective, this approach often demands substantial task-specific datasets, which can limit its scalability for diverse downstream applications.

In contrast, zero-shot learning enables models to generalize to unseen tasks *without requiring task-specific fine-tuning*. Instead, the model leverages its pre-trained knowledge and interprets carefully designed prompts to perform tasks directly. For example, a model in a zero-shot setting can respond to a prompt such as “Summarize the key findings of this medical report” without being explicitly trained on clinical summarization datasets.

In-context learning enhances the capabilities of these models by enabling them to solve tasks using information provided within the input prompt. Unlike fine-tuning, ICL does not require updating the model’s weights. Instead, the model temporarily “learns” from examples in the prompt. For instance, the model can be provided with a series of medical cases that include patient symptoms, relevant background knowledge, and corresponding diagnoses. When presented with a new case, such as a patient with specific symptoms (e.g., coughing up phlegm and blood), the model can infer the most likely diagnosis based on the examples and background knowledge, without explicit task-specific fine-tuning [26].

Nevertheless, challenges remain, as highlighted by Mahowald et al. [27], who emphasize the gap between formal and functional linguistic competencies in LLMs. Formal linguistic competencies refer to a model’s ability to understand and generate syntactically and semantically correct language, enabling tasks like sentence completion, grammar correction, or summarization. In contrast, functional competencies involve applying language in practical, goal-oriented contexts, such as interpreting patient symptoms to suggest a diagnosis or deriving actionable steps from clinical guidelines. Achieving functional performance requires not only linguistic fluency but also *reasoning* and task-specific understanding, integrating broader contextual and intentional knowledge.

This gap can be analogized to human neuroscience, where distinct neural mechanisms underpin linguistic processing and reasoning [27]. Similarly, LLMs, trained predominantly on textual data, excel at linguistic mimicry but often lack the modularity and cognitive integration necessary for complex functional reasoning. To bridge this dissociation, Mahowald et al. propose integrating modular architectures or revising training processes to better align with these competencies.

For medical professionals, this gap underscores the limitations of LLMs in healthcare. While LLMs may summarize medical literature with linguistic precision, they may struggle to apply this information effectively in clinical scenarios, such as tailoring treatment plans based on patient-specific data. Recognizing these limitations is critical to ensuring LLMs are deployed as tools that complement, rather than replace, human expertise.

Notably, current research is shifting towards multimodal models that integrate multiple *modalities*, such as visual and textual data, into a single model [4,5,28], paving the way for more comprehensive AI solutions. While this review focuses on text-only language models, it briefly examines the ecosystem of multimodal models to highlight their emerging relevance. This exploration provides readers with starting points for further investigation into how these models expand the scope of AI applications, particularly in domains like medical imaging, where integrating textual and visual data is essential.

## 2.2. Locally-deployable models

Models can be categorized based on their data management and deployment strategies: locally-deployable models and API-based models. API-based models, such as GPT-4, require

data transfer to third-party servers via web interfaces or APIs. In contrast, locally-deployable models run on an organization's hardware, offering full data control and independence from external vendors.

In the medical domain, locally-deployable models provide critical advantages by ensuring sensitive medical data remains internal, complying with stringent data protection regulations, and enhancing operational autonomy by eliminating reliance on third-party providers.

Many locally-deployable models feature open weights and permissive usage terms [29]. This flexibility allows them to be customized for medical domain, optimized for specific types of inquiries, and seamlessly integrated with confidential datasets, while maintaining strict patient confidentiality.

The primary limitations of locally-deployable models are their substantial hardware requirements and the technical expertise needed for deployment. Advanced models often comprise a significant number of parameters, typically represented as floating-point values. While a higher parameter count generally improves reasoning, pattern recognition, and linguistic capabilities [61], it also increases storage demands and necessitates more powerful GPUs for both inference and fine-tuning.

Recent advancements have addressed the computational challenges of large models through parameter-efficient fine-tuning and memory optimization. Techniques like Low-Rank Adaptation (LoRA) [62] update only low-rank parameters, reducing the need to adjust all model weights. Quantization [63] further minimizes memory usage and accelerates inference, enabling broader hardware compatibility. Combining these, Quantized Low-Rank Adaptation (QLoRA) [64] efficiently fine-tunes large models while significantly lowering memory and computational requirements. For example, a 1-billion-parameter model using 16-bit precision requires around 2GB of GPU memory, whereas quantization can halve this demand. Additional innovations, such as Virtual Memory Stitching (VMS) [65], optimize GPU memory allocation, reducing fragmentation and usage. Flash Attention [66] further enhances hardware efficiency, expanding the feasibility of deploying larger models on on-premise systems.

Table 2 summarizes widely used locally-deployable textual and multimodal models (predominantly English-language) categorized by the number of parameters and their architectures. The table distinguishes between three primary architectures: encoder-only, decoder-only, and encoder-decoder, each of which brings its own strengths and use cases.

Encoder-only models (e.g., BERT [25], DeBERTa [32]) excel in understanding and classification tasks, consistently achieving strong performance on comprehension-oriented benchmarks [67–69]. Decoder-only models (e.g., GPT-Neo [30], Falcon [46]) are better suited for generative tasks, such as text completion and narrative generation. Encoder-decoder models (e.g., BART [35], Flan-T5 [37]) strike a balance between comprehension and generation. While multimodal models can extend the capabilities of any of these architectures by incorporating multiple modalities (e.g., text and vision), we deliberately omit these details to maintain the review's focus on text-based systems.

**General performance guidelines.** General evaluation of encoder-only language models typically focuses on tasks that assess deep language understanding without requiring generative capabilities. Benchmarks such as GLUE [70] and SuperGLUE [69], encompassing tasks like natural language inference, sentiment analysis, paraphrase detection, and coreference resolution, are widely adopted. Many of these benchmarks maintain leaderboards to track state-of-the-art performance [71]. Since encoder-only models often need fine-tuning for specific downstream tasks, it is advisable to prioritize evaluation on benchmarks that reflect the intended application.

**Table 2. Distribution of widely used pre-trained open-source general-purpose models by number of parameters and architecture.**

| Size  | Decoder-based   | Encoder-based   | Encoder-Decoder   | Multimodal  |
|-------|---|---|---|---|
| ≤ 1B  | GPTNeo [30] (125–350M)  | BERT [25] (110–340M), ALBERT [31] (12–235M), DeBERTa [32] (134M), ELECTRA [33] (14–335M), RoBERTa [34] (125–355M) | BART [35], BertGeneration [36] (140–400M), Flan-T5 [37] (77–783M), Pegasus [38] (568M), T5 [39] (60–770M) | BLIP-2 [40] (188M), CLIP [41] (428M), deplot [42] (300M), Donut [43] (200M), LayoutLMv3 [44] (133–368M)         |
| ≤ 10B | CTRL [45] (1.63B), Falcon [46] (7B) GPT-J [47] (6B), Gemma [48] (2–7B), Llama [6,16] (7B, 8B), Mistral [49] (7B), Phi-3 [50] (3.8B) | DeBERTa (1.5B)  | Flan-T5 (3B), LongT5 [51] (3B), T5 (3B)   | Fuyu [52] (8B), BLIP-2-Opt (3.8B), Llava [53] (7B), PaliGemma [54] (3B), Chameleon [28] (7B), LLaVa [53] (7–8B) |
| ≤ 20B | GPT-NeoX [55] (20B), Llama (13B)  | N/A   | Flan-T5 (11B), T5 (11B), UL2 [56] (20B)   |   |
| ≤ 80B | Cohere [57] (35B), Falcon (40B), Llama (65B, 70B)   | N/A   | N/A   | Chameleon (34B), LLaVa (34B)  |
| ≤ 80B | DBRX [58] (132B), Falcon (180B), OPT [59] (175B)  | N/A   | N/A   | BLOOM [60] (176B)   |

Note: B is billions, M is millions.

<https://doi.org/10.1371/journal.pdig.0000800.t002>

The performance of decoder-only models is typically measured by their ability to generate coherent text, comprehend and reason about given inputs, and follow instructions. Prominent benchmarks include MMLU [72] and MMLU-Pro [73], while the technical reports of models like GPT-3 [24] provide a comprehensive list of metrics tailored for generative, decoder-only models.

Encoder-decoder models, designed for sequence-to-sequence tasks, are usually evaluated on benchmarks that measure their capacity to transform input sequences into output sequences. Examples include SQuAD [67] for question answering and XSum [74] for summarization. Readers interested in a deeper understanding of encoder-decoder architectures and their performance considerations are encouraged to review the T5 paper [39].

Evaluation of multimodal models often involves visual question answering (VQA) benchmarks. For instance, TextVQA [75] focuses on text reading in natural images, DocVQA [76] targets document understanding, and ChartQA [77] assesses chart comprehension. The Gemini technical report [5] offers a comprehensive list of metrics for vision-and-language multimodal models.

### 2.3. Domain-specific language models in healthcare

Domain-specific language models are models tailored for a narrow field or topic, enhancing their ability to produce more accurate and relevant responses within that specific area. Generally, there are two common methods to create a domain-specific model: one approach is to pre-train a model using a set of domain-specific documents, such as medical papers; another is to take a generically trained model and fine-tune or adapt it to the target domain.

The medical field, rich in unstructured textual data from sources such as Electronic Health Records (EHR) and other medical documentation, presents an ideal research domain for language models to address a wide array of challenges. These applications range from tasks previously tackled by NLP techniques, such as clinical acronym disambiguation, to those unattainable a decade ago. For example, medical conversational agents that can assist both patients and healthcare professionals are now emerging. In response to these advancements,



the research community has been actively developing language models specifically designed for medical applications. The evolution of these models, from pre-training and fine-tuning strategies to creating open models capable of zero-shot learning, highlights the growing synergy between artificial intelligence and healthcare [78–81].

Early medical language models were predominantly based on BERT and followed the common trend of the pre-training and fine-tuning paradigm. ClinicalBERT [82], pre-trained on the MIMIC-III [83] dataset and fine-tuned to predict hospital readmissions, was one of the pioneering models in this domain. Another notable model, BioBERT, was pre-trained on the PubMed library and fine-tuned for tasks such as named entity recognition (NER), relation extraction (RE), and medical question answering (QA) [84]. The proven efficacy of pre-training language models on biomedical corpora and fine-tuning them for specific downstream tasks, like clinical concept extraction or measuring semantic text similarity (STS), led to the development of many other BERT-based models, including BiomedBERT [85], PubMedBERT [86], BEHRT [87], and GatorTron [88].

The debut of BioGPT has marked another milestone in biomedical NLP. This decoder-based model, based on GPT-2 and trained on millions of PubMed abstracts, has demonstrated proficiency in specialized tasks, including text classification, and biomedical text generation [89].

AlpaCare, a medical instruction-tuned LLM trained on more than fifty thousand instruction-response pairs generated artificially using GPT-4, represents a notable example from the current generation of medical LLMs that can be utilized in the medical domain without the need for fine-tuning on downstream tasks [78]. Clinical Camel [81], another medical LLM fine-tuned from the Llama-2 model, excels in various medical tasks, ranging from clinical note creation to medical triaging. Despite its limitations, such as the potential for generating misleading content and the need for continuous updates, it marked a remarkable advancement in medical LLMs. The issue of needing continuous updates was addressed by the ChatDoctor [79] model. Based on the Llama family and fine-tuned with real-world patient-doctor conversations, this model, as part of an integrated framework, can retrieve information from external sources, a crucial capability in the medical field, particularly for addressing emerging diseases.

Recent advancements in the field, particularly the development of multimodal models that work with both text and images, have substantially impacted the medical domain as well. The LLaVA-Med model [90] can answer open-ended questions about biomedical images, while the CheXagent model [91], designed specifically for chest X-ray interpretation, exemplifies the successful transition from general medical models to specialist models capable of tackling problems in narrow medical fields.

Table 3 provides a comprehensive, though not exhaustive, summary of existing medical language models. It highlights current trends in the development of medical models and the evolution of training and experimental datasets utilized. The experimental datasets listed in the table offer valuable references for readers interested in evaluating performance on specific tasks within the medical domain.

### 3. Medical applications of language models

Medical applications of language models can be defined as the intersection of two sets: tasks that language models can accomplish and potential healthcare needs where predominantly textual language models can add value. Although both sets are finite, the cardinality of the first set is arguably smaller given the vast scope of the medical domain. Thus, our approach

**Table 3. Overview of prevalent medical language models.**

| Name                | Year (Appx.) | Arch.              | Training Data   | Experimental Datasets  |
|---------------------|--------------|--------------------|---|--|
| ClinicalBERT [82]   | 2019         | BERT               | MIMIC-III [83]  | MIMIC-III  |
| BioBERT [84]        | 2019         | BERT               | PubMed abstracts and PMC full-text articles   | NER: NCBI [92], i2b2/VA [93], BC5 [94], BC4CHEMD [95], BC2GM [96], JNLPBA [97], LINNAEUS [98], Species-800 [99]; RE: GAD [100], EU-ADR [101], CHEMPROT [102]; QA: BioASQ [103] 4b, 5b, 6b, and 7b. |
| BiomedBERT [85]     | 2020         | BERT               | Pre-trained on the BREATHE dataset [85]   | NER: NCBI, BC5CDR, BC4CHEMD, BC2GM, JNLPBA; RE: GAD and EU-ADR; QA: SQuAD [67] v1.1 and v2.0, BioASQ 4b, 5b, 6b, and 7b.   |
| PubMedBERT [86]     | 2020         | BERT               | PubMed only   | BC5, NCBI, BC2GM, JNLPBA, EBM PICO [104], CHEMPROT, DDI [105], GAD, BIOSSES [106], HoC [107], PubMedQA [108], BioASQ   |
| BEHRT [87]          | 2020         | BERT               | Clinical Practice Research Datalink (CPRD) [109]  | Predict diseases in future patients' visits with CPRD  |
| GatorTron [88]      | 2022         | BERT-style         | UF Health [110], PubMed articles, Wikipedia   | Clinical Concept Extraction: i2b2, n2c2 [111]; RE: n2c2; STS: n2c2/OHNP Clinical STS; NLI: MedNLI [112]; QA: emrQA [113].  |
| BioGPT [89]         | 2022         | GPT-2 XL           | 15 million PubMed items   | RE: BC5CDR, KD-DTI [114], DDI [105]. MQA: PubMedQA. Document classification: HoC. Text generation: custom dataset [89].  |
| ClinicalT5 [115]    | 2022         | T5                 | MIMIC-III (textual notes)   | Document classification: HoC; NER: NCBI [92], BC5CDR-disease; NLI: MedNLI; Real-world evaluation based on MIMIC-III to predict ICU readmission and mortality risks                                 |
| AlpaCare [78]       | 2023         | LLaMA with IFT     | MedInstruct-52 (introduced in the paper): 52000 instruction-response pairs generated using GPT-4 prompted with a clinician-crafted seed set   | iCliniq2 [78], MedInstruct-test  |
| BioInstruct [116]   | 2023         | LLaMA with IFT     | BioInstruct dataset (introduced in the paper): 25000 instructions in natural language collected with GPT4   | QA: MedQA-USMLE [117], MedMCQA [118], PubMedQA, BioASQ MCQA; NLI: MedNLI. Text Generation: Conv2note [119], iCliniq [120].   |
| ChatDoctor [79]     | 2023         | LLaMA              | Conversations from HealthCareMagic-100k [79]  | HealthCareMagic100k, iCliniq   |
| Clinical Camel [81] | 2023         | LLaMA              | Clinical articles converted into synthetic dialogues, data from ShareGPT [81], MedQA [117]  | MMLU [121], MedMCQA, MedQA, PubMedQA, USMLE Sample Exam [122]  |
| MedAlpaca [123]     | 2023         | LLaMA              | Medical Meadow dataset (introduced in paper).   | USMLE assessment   |
| PMC-LLaMA [124]     | 2023         | LLaMA with IFT     | MedC-K (introduced in paper): based on S2ORC [125] with emphasis on biomedical papers, 30000 textbooks. MedC-I (introduced, for IFT): based on MedAlpaca and ChatDoctor datasets, USMLE, PubMedQA, MedMCQA, UMLS medical knowledge graphs | MCQA: PubMedQA, MedMCQA, USMLE   |
| LLaVA-Med [90]      | 2023         | Fine-tuned LLaVA   | Based on PMC-15M [126]  | Visual QA: VQA-RAD, SLAKE [127], PathVQA   |
| BioMistral [128]    | 2024         | Mistral            | PubMed Central Open Access  | MMLU, MedQA, MedMCQA, PubMedQA   |
| CheXagent [91]      | 2024         | custom, multimodal | CheXinstruct (introduced in paper), MIMIC-CXR [129], PadChest [130], BICMV-COVID-19 [131]   | MIMIC-CXR, CheXpert [132], SIIM [133], RSNA [134], OpenI [135], SLAKE  |

<https://doi.org/10.1371/journal.pdig.0000800.t003>

to categorizing medical applications is based on language models' tasks of various granularity, which, in turn, have a large area of intersection with NLP tasks in general. We focus on



practically significant tasks from an application perspective, as well as notable macro applications that encompass multiple NLP tasks, deliberately excluding specific fine-grained NLP tasks such as coreference resolution and dependency parsing.

Table 4 provides an overview of selected applications of language models in the medical domain, focusing on clinical practice and related research. Subsequent subsections delve into each task, highlighting key applications, notable datasets, and exemplary implementations. Within the scope of this study, we intentionally exclude applications in medical education and highly specialized areas of biomedical research. Readers interested in a wider range of applications of language models in medicine may refer to the paper by Thirunavukarasu et al. [10].

### 3.1. Text generation

The task of text generation in the medical domain involves creating contextually accurate and relevant medical texts based on a sequence of prior tokens and specific context. This task may include generating clinical notes, patient reports, or treatment plans. The primary challenge is ensuring that the generated text is precise, medically accurate, and adheres to relevant privacy and ethical standards.

**Table 4. Summary of major applications of language models in medical domain.**

| LLM Application                               | Examples of Medical Applications  | Notable Datasets   | Notable Solutions   |
|---|---|--|---|
| Text Generation                               | Medical Report Generation [136,137], Clinical Note Generation [138], Generating Summaries For Laypersons [139], Generating Summaries for Patient-Provider Dialogues [140], Generating Textual Descriptions From Graph Models [141]                                  | CTRG-Chest-548K, CTRG-Brain-263K [142], IU-Xray [143], MIMIC-CXR [144], CheXpert [145]   | Dia-LLaMA [137], Talk2Care [146], MEDSUM-ENT [140]                          |
| Token Classification                          | Clinical Acronym Disambiguation [147,148], Eponyms Disambiguation [149]   | CASI [150], NLM-WSD [151]  | SciBERT, BioBERT, ClinicalBERT [152]  |
| Sequence Classification                       | Phenotyping [153], Medical Coding [154], Modeling Patient Timeline [155–157], Social Media Monitoring [158]   | Suicide Watch [159], CSSRS [160], MIMIC-III [83], MIMIC-IV [161], eICU-CRD [162], HPO-GS [163], BIOC-GS [164], CAMS [165], Wellness-Reddit [166], Mental Disturbance [167] | Foresight [155], Bio_ClinicalBERT [168], BioMedLM [156]                     |
| Question Answering and Information Extraction | Querying Data from Electronic Health Records [169], Extracting Information from Clinical Narrative Reports [170], Extracting Information From Medical Articles [171]  | CASI [150], n2c2 [111], i2b2 [93], PubMedQA [108], MedMCQA [118], emrQA [113], BIOASQ [103]  | quEHRy [172], BiomedBERT [85], PubMedBERT [86], BioGPT [89], Llava-med [90] |
| Summarization and Paraphrasing                | Summarizing Clinical Study Reports [173], Summarizing Patient-Provider Dialogues [174], Simplification of Medical Texts [175,176], Simplification of Radiology Reports [177], Improving Biomedical Text Readability [178], Anonymization of Medical Documents [179] | PLS-Cochrane Reviews [175], CELLS [180], Pfizer Clinical Trial Data [173], MultiCochrane [181]   | fine-tuned BART [175], RALL [180], fine-tuned Llama [173,179]               |
| Conversation                                  | Mental Health Bots [182–184], Medical Chatbots and Health Assistants [80,185,186], Triaging [187], Differential Diagnosis [188]   | MotiVate [189], Depression_Reddit [190], CLPsych [191], Dreddit [192], Clinical Vignettes [187]  | MedPaLM, DRG-LLaMA [193], openCHA [80]                                      |

<https://doi.org/10.1371/journal.pdig.0000800.t004>

To achieve this, text generation models rely on a probabilistic framework. The probability of generating the next token  $x_t$ , given a sequence of previous tokens  $x_1, x_2, \dots, x_{t-1}$  and additional context  $c$ , can be defined as:

$$P(x_t | x_1, x_2, \dots, x_{t-1}; c)$$

In this formulation, context  $c$  represents various factors depending on the nature of the task. For instance, when generating patient reports, the context would include the available information about the patient, such as vital statistics, previous diagnoses and treatments.

Decoder-based language models inherently generate text by predicting the probability of the next token based on preceding tokens and contextual information. The methods for providing context vary. The simplest approach is using prompts given directly to the language model. However, because the knowledge within language models is static, integrating external knowledge sources can be advantageous. This integration is often accomplished using variations of Retrieval-Augmented Generation (RAG) techniques [194,195].

Advancements in transformer-based models have led to innovative approaches in generating clinical documentation from patient-provider interactions. A study by Brake and Schaaf [196] compares two model designs for generating clinical notes from doctor-patient conversations using the encoder-decoder-based PEGASUS-X model [197]. The first design, GENMOD, generates the entire note in one step, while the second design, SPECMOD, generates each section independently. The study aims to evaluate the consistency of the generated notes in terms of age, gender, body part, and coherence. Evaluations were performed using ROUGE [198] and Factuality [199] metrics, human reviewers, and the Llama2 LLM. Results indicate that GENMOD improves consistency in age, gender, and body part references, while SPECMOD may have advantages in coherence depending on the interpretation. The study uses a proprietary dataset with 10,859 doctor-patient conversations for training and testing [196]. Nair et al. present MEDSUM-ENT [140], a multi-stage approach to generating medically accurate summaries from patient-provider dialogues. Using GPT-3 as the backbone, the approach first extracts medical entities and their affirmations from conversations and then constructs summaries based on these extractions through prompt chaining. The model leverages ICL and dynamic example selection to improve entity extraction and summarization. The dataset used for evaluation consists of 100 de-identified clinical encounters from a telehealth platform. MEDSUM-ENT demonstrates improved clinical accuracy and coherence in summaries compared to a zero-shot, single-prompt baseline, as evidenced by both qualitative physician assessments and quantitative metrics designed to capture medical correctness [140].

The integration of LLMs with visual models facilitates the automatic generation of medical imaging reports by aligning visual features from models like Swin Transformer [200] with the LLM's word embedding space, enabling effective multi-modal understanding. Chen et al. developed Dia-LLaMA [137], a framework that utilizes the LLaMA2-7B model combined with a pre-trained ViT3D [201] to manage high-dimensional CT data. It features a disease prototype memory bank and a disease-aware attention module to counteract the imbalance in disease occurrence. The framework, tested on the CTRG-Chest-548K dataset [142], outperformed other methods in various natural language generation metrics presented in the study. Another approach, R2GenGPT [136], enables radiology report generation by utilizing a visual alignment module that aligns visual features from chest X-ray images with the word embedding space of LLMs, thereby enhancing the capability of static LLMs to process visual data. It explores three alignment strategies: shallow, deep, and delta, each varying in trainable parameters. Evaluated on the IU-Xray and MIMIC-CXR datasets, R2GenGPT achieved impressive

results in model efficiency and clinical metrics, leveraging the Swin Transformer [200] and Llama2-7B model for enhanced integration.

Evaluating commercial models for clinical tasks inevitably draws the interest of the scientific community. Ali et al. evaluated the use of ChatGPT for generating patient clinic letters, focusing on its readability, factual correctness, and human-like quality by testing the model with shorthand instructions simulating clinical input for creating letters addressing skin cancer scenarios. The study involved 38 hypothetical clinical scenarios, including basal cell carcinoma, squamous cell carcinoma, and malignant melanoma. Readability was assessed using the online tool Readable [202], targeting a sixth-grade reading level. Two independent clinicians evaluated the letters' correctness and human-like quality using a Likert scale. The study found that ChatGPT-generated letters scored highly in correctness and human-like quality [203].

Overall, while we observe the significant potential of text-only LLMs in text generation tasks, the emergence of multimodal models will inevitably bear fruit in this class of tasks in the medical domain. For example, fully multimodal models will be able to accurately generate clinical documentation based on patient-provider verbal dialogues and summarize them based on the end user's needs, while generating medical imaging reports will become less complicated and more accurate when done by a single multimodal model.

### 3.2. Token classification

Token classification tasks in the medical domain involve labeling individual words or phrases within a text with specific medical annotations, such as identifying and disambiguating medical conditions, medications, dosages, and symptoms from clinical text.

Given a sequence of tokens  $X = (x_1, x_2, \dots, x_n)$  and context  $c$ , the task is to assign a label  $y_i \in \mathcal{Y}$  to each token  $x_i$ , where  $\mathcal{Y}$  is the set of possible categories. This can be expressed within a probabilistic framework as:

$$y_i = \arg \max_{k \in \mathcal{Y}} P(k | x_i, c),$$

where  $P(k | x_i, c)$  is the probability of token  $x_i$  belonging to class  $k$ . In this formulation, context  $c$  may refer to the surrounding tokens or encompass information outside of the sequence  $X$ , such as external vocabulary.

Encoder-based models, such as BERT, typically excel in this category of tasks. However, they require fine-tuning on a labeled dataset to adapt to specific tasks, and fine-tuning may need to be reiterated if new information or data becomes available. In contrast, models with a decoder component, can use prompts to generate text annotated with labels, thereby directly producing labeled sequences without additional task-specific fine-tuning.

The widespread use of medical abbreviations and acronyms often leads to misunderstandings, necessitating accurate disambiguation of these terms to safeguard against misinterpretations that could jeopardize patient care [204]. The process of mapping the short forms of medical terms to their full expressions is referred to as clinical acronym disambiguation. Wang and Khanna [152] evaluated the performance of various clinical BERT-based language models on the Clinical Acronym Sense Inventory (CASI) dataset and found that ClinicalBert addresses the task effectively, achieving an F1 score of 0.915. In contrast, Sivarajkumar et al. [205] assessed the capabilities of generative LLMs, including GPT3.5, Bard (Gemini), and Llama2, in acronym disambiguation using the same dataset. Their study revealed that these models perform well in acronym disambiguation without fine-tuning, with GPT3.5 achieving the highest accuracy of 0.96.

While the problem of clinical acronym disambiguation might seem largely addressed, several challenges persist. For instance, Kugic et al. [147] conducted a study on clinical acronym disambiguation in German using ChatGPT and Bing Chat (Copilot), achieving an F1 score of 0.679, which highlights the need for improvement. Another concern involves the datasets used in experiments. There is a possibility that LLMs may memorize specific terms, which could misrepresent their true disambiguation capabilities [206,207]. This issue necessitates further investigation to ensure the reliability of LLMs in clinical abbreviation disambiguation tasks with realistic datasets.

### 3.3. Sequence classification

Sequence classification tasks in the medical domain involve assigning a label or multiple labels to an *entire sequence* of text, rather than to individual tokens. These tasks may include classifying clinical documents or patient notes into categories such as diagnosis, treatment recommendation, or urgency level.

Most sequence classification tasks fall into the category of single-label classification, which can be formulated as follows: given a sequence of tokens  $X = (x_1, x_2, \dots, x_n)$ , the task is to assign a label  $y \in \mathcal{Y}$ , where  $\mathcal{Y}$  is the set of possible categories:

$$y = \arg \max_{k \in \mathcal{Y}} P(k | X),$$

where  $P(k | X)$  represents the probability of the sequence  $X$  belonging to class  $k$ .

Encoder-based models are well-suited for sequence classification tasks. For example, BERT can be fine-tuned for classification tasks by adding a fully connected layer on top of the embedding of the [CLS] token, which serves as a contextual representation of the entire sequence. Alternatively, pooling functions such as average pooling or max pooling can be applied to the token embeddings of the sequence to aggregate information across the entire input. Similar to token classification tasks, sequence classification often requires re-fine-tuning when new data or updated label distributions become available.

Generative language models, which incorporate a decoder component, can be either simply prompted or fine-tuned to predict class labels as part of the sequence. Fine-tuning approach typically involves training the model on inputs where a special token or delimiter demarcates the end of the input sequence and the beginning of the output label. In some cases, integrating a linear layer after the final token output in generative models can refine the logits corresponding to class predictions, enhancing the classification boundaries produced by the model's generative framework. However, this approach introduces additional complexity and is less commonly used.

The array of tasks that fall under sequence classification is vast. The following subsections detail a few of the most prominent applications.

**3.3.1. Suicidal behavior prediction.** Suicidal behavior prediction tasks predominantly focus on analyzing individuals' social media activities. Dus and Nefedov [157] proposed an automated tool for identifying potential self-harm indications in social media posts, framing the task as a binary classification problem. The set of possible categories  $\mathcal{Y}$  includes two elements: 1, indicating the presence of suicidal behavior, and 0, indicating its absence. The input, a sequence of tokens  $X$ , is derived from the text of social media posts. The model estimates  $P(y = 1 | X)$ , the probability that the input suggests suicidal behavior, by leveraging a fine-tuned ELECTRA model. Training data included samples from Kaggle's "Suicide Watch" dataset [159], supplemented with additional social media sources. The proposed method achieved an accuracy of 0.93 and an F1 score of 0.93 [157].

Beyond mere social media post analysis, Levkovich et al. [208] assessed ChatGPT-3.5 and ChatGPT-4's ability to evaluate suicide risk based on perceived burdensomeness and thwarted belongingness. By comparing ChatGPT's assessments to those made by mental health professionals using vignettes, they discovered that ChatGPT-4's evaluations were in close alignment with professional judgments. In contrast, ChatGPT-3.5 tended to underestimate suicide risk, underscoring the limitations of these models in this specific area.

In summary, while treating suicidal behavior identification as a straightforward classification task on social media posts can lead to impressive scores using standard classification metrics, the practical and ethical implications of such approaches, including potential breaches of autonomy and principles of non-maleficence, are debatable [209]. Well-structured vignette studies on the effectiveness of LLMs and other models can further advance research in this area. Additionally, exploring the potential of human-AI collaboration represents another promising research direction in this field.

**3.3.2. Modeling patient timeline.** The task of modeling patient timelines is multifaceted, involving forecasting future medical events, understanding patient trajectories, and predicting medical outcomes. This endeavor employs deep learning, transformers, and generative models to analyze data from various medical records, both structured and unstructured.

Kraljevic et al. introduced Foresight [155], a GPT-2-based pipeline developed for modeling biomedical concepts extracted from clinical narratives. This pipeline employs NER and linking tools to transform unstructured text into structured, coded concepts. Utilizing datasets from three hospitals, covering over 800,000 patients, Foresight showed promise in forecasting future medical events. Its effectiveness was manually validated by clinicians on synthetic patient timelines, highlighting its potential in real-world risk forecasting and clinical research.

Among different types of models, generative adversarial networks (GANs) have gained popularity, extending their applications beyond the initial domain of image generation. Shankar et al. proposed Clinical-GAN [161], which merges Transformer and GAN methodologies to model patient timelines, focusing on predicting future medical events based on past diagnosis, procedure, and medication codes. Tested on the MIMIC-IV dataset, Clinical-GAN outperformed baseline methods in trajectory forecasting and sequential disease prediction [210]. Another study [211] employed GAN for predicting the length of stay in emergency departments. The learning process was done in multiple stages. Initially, an unsupervised training phase used a generator and discriminator to approximate the probability distribution and perform feature discovery and reconstruction. Discriminator was then fine-tuned to optimize its parameters for global optimum. A predictor layer, initially randomly initialized, was added and optimized during fine-tuning, enabling the model to map observations to their lengths of stay. The model was trained on data from the Pediatric Emergency Department in CHRU-Lille and proved the potential of GANs in this field.

*Medical outcome prediction* can be seen as a subtask of modeling patient timeline and is often scoped to either predict mortality, outcomes of a specific disease, or risk of progression from one disease to another. A recent study by Shoham and Rappoport [156] examined data related to chronic kidney disease, acute and unspecified renal failure, and adult respiratory failure from the MIMIC-IV and eICU-CRD datasets. Using this data, the team generated labeled datasets for disease diagnosis prediction based on patient histories. They introduced a method named Clinical Prediction with Large Language Models (CPLLM) by fine-tuning LLMs (Llama2 and BioMedLM) using medical-specific prompts to help the models understand complex medical concept relationships. Xie et al. [168] used EHR analysis to predict epilepsy seizures, leveraging Bio\_ClinicalBERT, RoBERTa, and T5, achieving an F1 score of 0.88 in outcome classification.

A notable approach to predict outcomes of COVID-19 patients was proposed by Henriksson et al. [212]. The authors created a model that combines structured data and unstructured clinical notes in a multimodal fashion, leveraging a clinical KB-BERT model for multimodal fine-tuning. Trained on data from six hospitals in Stockholm, Sweden, their model effectively predicted 30-day mortality, safe discharge, and readmission of COVID-19 patients in the emergency department, as measured by AUC.

**3.3.3. Phenotyping and medical coding.** The phenotyping task primarily involves identifying phenotypic abnormalities from a patient's various medical records, which aids in the identification of rare diseases. There exists the Human Phenotype Ontology (HPO) project that systematically categorizes human phenotypes with detailed annotations [213]. The phenotyping task can be framed as a **multi-label classification task**, which extends the single-label classification task as follows. Given a sequence of tokens  $X = (x_1, x_2, \dots, x_n)$ , the task is to assign a subset of labels  $Y \subseteq \mathcal{Y}$ , where  $\mathcal{Y}$  is the set of all possible labels.

For each label  $k \in \mathcal{Y}$ , the model predicts a probability  $P(k | X)$  representing the likelihood of the sequence  $X$  being associated with label  $k$ . The set of labels  $Y$  is typically determined by applying a predefined threshold  $\theta$  to these probabilities:

$$Y = \{k \in \mathcal{Y} \mid P(k | X) \geq \theta\}.$$

Thus, in the phenotyping task, the sequence  $X$  represents a medical record of a patient, and  $\mathcal{Y}$  denotes the set of all possible phenotype labels from the HPO. The objective is to identify a subset  $Y \subseteq \mathcal{Y}$  of HPO labels associated with the medical record.

Traditionally, the task of phenotyping has relied on named entity recognition, where models similar to BERT have demonstrated proficiency. However, recent studies have started exploring in-context learning and zero-shot learning with contemporary LLMs, yielding promising results [153].

Medical coding is another multi-label classification task that involves identifying a set of International Classification of Diseases (ICD) codes associated with a medical record. This task can be formulated similarly to phenotyping, with the key difference being that  $\mathcal{Y}$  represents a set of ICD codes rather than HPO labels. Besides observing trends similar to those in the phenotyping subdomain, it is also noteworthy that there is a shift towards explainable medical coding, as highlighted in [154].

### 3.4. Question answering and information extraction

Question Answering (QA) task can be formulated as finding the answer  $A$  from a possible set of answers  $\mathcal{A}$ , given a question  $q$  and a context  $c$  (often a document or set of documents containing information relevant to the question). This can be expressed as:

$$A = \operatorname{argmax}_{a \in \mathcal{A}} P(a | q, c)$$

where  $P(a | q, c)$  is the probability of  $a$  being the correct answer given the question  $q$  and the context  $c$ .

Information Extraction (IE) task involves identifying specific pieces of information (entities, relationships, events) within documents. This can be described as a function  $f$  that maps a set of documents  $D$  to a set of structured attributes  $S$ , which includes entities  $E$ , relationships  $R$ , and other attributes of interest:

$$S = f(D) = \{E, R, \dots\}$$



Here,  $D$  is the input document, and  $S$  represents the structured output containing extracted elements.

Encoder-based models are well-suited for question answering tasks. They can be fine-tuned on specific QA datasets, where the input is a concatenation of the question and context (a paragraph or document containing the answer). The model is then trained to identify the span of text that answers the question, typically by adding a start and end token classifier to the output embeddings of the model. These classifiers predict the beginning and end positions of the answer in the text. On the other hand, generative models with a decoder component leverage their extensive pre-training on diverse data. By inputting a question (along with the document of interest when necessary) and following it with a prompt that encourages the model to generate an answer, these models can produce responses without needing explicit pointers to answer spans.

In the medical domain, IE and QA systems are instrumental for extracting data from electronic health records, such as medication lists and diagnostic details, essential for patient management and treatment planning. A notable example is quEHRY, a QA system designed to query EHRs using natural language. The primary goal of quEHRY is to provide precise and interpretable answers to clinicians' questions from structured EHR data [172]. Beyond the successful applications of BERT-based models like BioBERT, BiomedBERT, and PubMedBERT for QA and IE, generative models also show proficiency. Agrawal et al. [214] demonstrated the effectiveness of generative LLMs such as InstructGPT and GPT-3 in zero-shot and few-shot information extraction from clinical texts, measured through accuracy and F1 score. When tested on the re-annotated CASI dataset, these models showed considerable potential in tasks requiring structured outputs. Furthermore, Ge et al. [215] compared the effectiveness of LLMs versus manual chart reviews for extracting data elements from EHRs, focusing specifically on hepatocellular carcinoma imaging reports. Using the GPT-3.5-turbo model, implemented as "Versa Chat" within a secure UCSF environment to protect patient health information, the study analyzed 182 CT or MRI abdominal imaging reports from the Functional Assessment in Liver Transplantation study. It extracted six distinct data elements, including the maximum LI-RADS score [216], number of hepatocellular carcinoma lesions, and presence of macrovascular invasion. The performance was evaluated by calculating accuracy, precision, recall, and F1 scores, showing high overall accuracy (0.889) with variations depending on the complexity of the data elements.

### 3.5. Summarization and paraphrasing

Paraphrasing involves rewriting a text  $T$  into a new form  $P$ , ensuring that  $P$  maintains the same meaning as  $T$  but utilizes different vocabulary and potentially altered sentence structures. Summarization, on the other hand, entails generating a brief version of a text  $T$  that preserves its core information. Abstractive summarization can be considered a specific case of paraphrasing.

Encoder-based models are proficient at extractive summarization. They evaluate sentences within a text to determine their relevance and informativeness. By scoring each sentence, these models identify and concatenate the most important sentences to form a coherent summary. In contrast, abstractive summarization and paraphrasing typically employ decoder-based or sequence-to-sequence (encoder-decoder) models. These models are trained to understand the entire narrative or document and then recreate its essence in a different form.

Summarization and paraphrasing tools are used in the medical domain for managing extensive documentation and enhancing communication. Summarization helps healthcare

professionals quickly grasp essential details from lengthy clinical notes, generate concise abstracts of medical research papers, and craft clear patient discharge summaries, thereby improving patient comprehension and adherence to medical advice. Paraphrasing makes complex information more accessible by translating medical jargon into simpler language for patient education. It also enhances the clarity and consistency of electronic health records, aiding healthcare providers in better understanding and utilizing the data effectively.

Summarization and paraphrasing in the medical domain are largely driven by advancements in these tasks in general. Devaraj et al. [175] introduce a new dataset derived from the Cochrane Database of Systematic Reviews [217], featuring pairs of technical abstracts and plain language summaries. They propose a novel metric based on encoder-based language models to better distinguish between technical and simplified texts. The study utilizes baseline encoder-decoder Transformer models for text simplification and introduces an innovative approach to penalize the generation of jargon terms. The code and data are publicly available for further research.

The paper titled “Biomedical Text Readability After Hypernym Substitution with Fine-Tuned Large Language Models” investigates simplifying biomedical text using LLMs to enhance patient understanding. The authors fine-tuned three LLM variants to replace complex biomedical terms with their hypernyms. The models used include GPT-J-6b, SciFive T5 [218], and an approach combining sequence-to-sequence and sciBERT [219] models. The study processed 1,000 biomedical definitions from the Unified Medical Language System (UMLS) and evaluated readability improvements using metrics such as the Flesch-Kincaid Reading Ease and Grade Level, Automated Readability Index, and Gunning Fog Index. Results showed substantial readability improvements, with the GPT-J-6b model performing best in reducing sentence complexity [178].

Another interesting application of paraphrasing is the anonymization of medical documents, which is crucial for balancing ethical principles and research needs. Wiest et al. [179] present an approach to de-identify medical free text using LLMs. The authors benchmarked eight locally deployable LLMs, including Llama-3 8B, Llama-3 70B, Llama-2 7B, Llama-2 70B, and Mistral 7B, on a dataset of 100 clinical letters from a German hospital. They developed the LLM-Anonymizer pipeline, which achieved a success rate of 98.05% in removing personal identifying information using Llama-3 70B. The tool is open-source, operates on local hardware, and does not require programming skills, making it accessible and practical for use in medical institutions. The study demonstrates the potential of LLMs to effectively de-identify medical texts, outperforming traditional NLP methods and providing a robust solution for privacy-preserving data sharing in healthcare.

Despite advancements in summarization and paraphrasing, some challenges persist, particularly in preserving factual accuracy and precision. Jeblick et al. [177] explored the effectiveness of using ChatGPT (version December 15th, 2022) to simplify radiology reports into language understandable by non-experts. A radiologist created three hypothetical radiology reports, which were then simplified by prompting ChatGPT. Fifteen radiologists evaluated the quality of these simplified reports based on criteria such as factual correctness, completeness, and potential harm to patients. The study used Likert scale analysis and inductive free-text categorization to assess the simplified reports. Overall, the radiologists found the simplified reports to be factually correct and complete, with minimal potential for harm. However, some issues were noted, including incorrect information, omissions of relevant medical data, and occasionally misleading or vague statements. These issues highlight the need for careful supervision by medical professionals when using language models to simplify complex medical texts. A recent study by Landman et al. [173] discusses a challenge organized by Pfizer to explore the use of LLMs for automating the summarization of safety tables in clinical study

reports. Various teams employed GPT models with prompt engineering techniques to generate summary texts. The datasets included safety outputs from 72 reports from recent clinical studies, split into 70% for training and 30% for testing. The study concluded that while LLMs show promise in automating the summarization of clinical study report tables, human involvement and further research are necessary to optimize their application.

### 3.6. Conversation

The task of conversation, or dialogue generation, can be formulated as follows. Given a dialogue history  $H = (h_1, h_2, \dots, h_n)$ , where each  $h_i$  represents an utterance in the conversation, the objective is to generate an appropriate response  $R$ . This can be expressed as:

$$R = \operatorname{argmax}_r P(r | H),$$

Where  $P(r | H)$  represents the conditional probability of generating the response  $r$  given the dialogue history  $H$ . Typically, pre-trained decoder-based large language models are fine-tuned using specialized datasets to develop their conversational capabilities. In the medical domain, conversational applications facilitate interactive communication with patients. For example, conversational AI can be deployed as virtual health assistants that provide initial consultations based on symptoms described by patients. These systems can ask relevant follow-up questions, assess symptoms, and offer preliminary advice or direct patients to seek professional care when necessary. Additionally, these conversational tools can be utilized for patient education, explaining complex medical conditions and treatments in simple language to enhance understanding and compliance. Another notable application is in mental health support, where conversational AI can offer coping strategies and basic support, thereby augmenting traditional therapy sessions.

**Chatbots and health assistants.** The proficiency of LLMs in generating coherent text and finding patterns in natural language makes them excellent candidates for Conversational Health Agents (CHAs) or chatbots. The impressive capabilities of systems like ChatGPT have sparked researchers' interest in evaluating them as out-of-the-box medical chatbots. These chatbots are capable of holding conversations on medical topics and providing valid, science-based responses, akin to human doctors.

Cung et al. [185] assessed the performance of three commercial systems - ChatGPT, Bing (Copilot), and Bard (Gemini) in the context of skeletal biology and disorders. The study involved posing 30 questions across three categories, with the responses graded for accuracy by four reviewers. While ChatGPT 4.0 had the highest overall median score, the study revealed that the quality and relevance of responses from all three chatbots varied widely, presenting issues such as inconsistency and failure to account for patient demographics. Another study explored using ChatGPT for patient-provider communication. A survey of 430 participants found that ChatGPT responses were often indistinguishable from those of healthcare providers, indicating a level of trust in chatbots for answering lower-risk health questions [186].

Despite the success of chatbots in general and low-risk medical interactions, studies suggest that chatbots are not yet suitable for high-risk subdomains. For instance, a study focusing on resuscitation advice provided by Bing (Copilot) and Bard (Gemini) chatbots revealed that the responses frequently lacked guideline-consistent instructions and occasionally contained potentially harmful advice. Only a small fraction of responses from Bing (9.5%) and Bard

(11.4%) completely met the checklist criteria ( $P > .05$ ), underscoring the current limitations of LLM-based chatbots in critical healthcare scenarios [220].

Another research direction in the field of medical chatbots is the implementation of conversational agents specifically designed for the medical domain. Abbasian et al. [80] proposed a complex LLM-based multimodal framework for CHAs, concentrating on critical thinking, knowledge acquisition, and multi-step problem-solving. This framework aims to enable CHAs to provide personalized healthcare responses and handle intricate tasks such as stress level estimation.

Domain-specific LLMs, such as the ChatDoctor model [79], can also serve as chatbots. This model integrates a self-directed information retrieval mechanism, allowing it to access up-to-date information from online and curated offline medical databases. Evaluated using BERTScore [221], ChatDoctor exhibited a higher F1 score compared to ChatGPT-3.5, demonstrating the effectiveness of smaller domain-specific models as alternatives to large commercial solutions.

Overall, chatbots show promise, particularly in low-risk consultation areas. However, concerns such as potential confabulations, lack of explainability, and biases highlight challenges in their application in real-case scenarios [222]. Additionally, the absence of a robust, comprehensive, and universally accepted evaluation metric for chatbots is notable. Human evaluation lacks scalability, and similarity metrics like BERTScore may overlook critical factual inaccuracies.

**Mental health bots.** The idea of using a machine as a personal psychologist dates back to at least the 1960s when Weizenbaum proposed a simple rule-based system called ELIZA [223]. Contemporary advancements in mental health chatbots are largely driven by LLMs. Yang et al. [183] investigated the capabilities of current LLMs in automated mental health analysis. Their study involved evaluating LLMs across diverse datasets for tasks such as emotional reasoning and detecting mental health conditions, employing various similarity metrics including BLEU, ROUGE family, BERTScore derivatives, BART-score [224], and human assessments. They discovered that while ChatGPT displays robust in-context learning abilities, it still encounters challenges in emotion-related tasks and requires careful prompt engineering to improve performance to enhance its performance.

Saha et al. [182] introduced a Virtual Assistant for supporting individuals with Major Depressive Disorder, using a dataset called MotiVate. Their system, based on modified GPT-2 model and reinforced learning, shows promising results in generating empathetic and motivational responses, as evidenced by both automated evaluations based on text similarity and human evaluations based on fluency, adaptability, and degree of motivation. Sharma et al. [225] introduced a dataset for training a GPT-3-based model for generating reframes with controlled linguistic attributes. Deployed on the Mental Health America website [226], this allowed for a randomized field study to gather findings on human preferences. Another team explored the fine-tuning of open-source LLMs on psychotherapy assistant instructions, using a dataset from Alexander Street Press [227] therapy and counseling sessions. Their results indicated that LLMs fine-tuned on domain-specific instructions surpassed their non-fine-tuned counterparts in psychotherapy tasks, underscoring the significance of professional and context-specific training for these models [184].

Promising outcomes have been observed through collaborations between humans and AI. A recent study [228] conducted a randomized controlled trial involving human peer supporters, demonstrating that an AI-in-the-loop agent led to a 19.60% increase in conversational empathy in interactions between individuals seeking mental health support and support specialists. This was achieved by providing suggestions for response improvements to peer

supporters. This research reveals that human-AI collaboration is a crucial area for potential exploration, particularly in the medical domain.

**Other applications.** The evolving reasoning capabilities of LLMs have sparked interest in their use for disease diagnostics. Levine et al. [187] conducted experiments with the GPT-3 model to assess its diagnostic and triage accuracy. Their results indicate that GPT-3's diagnostic accuracy is comparable to that of physicians but lags in triage accuracy. GPT-3 correctly identified the diagnosis in its top three choices for 88% of the cases, surpassing non-experts (54%) but slightly underperforming compared to professional physicians (96%). In triage performance, GPT-3 achieved an accuracy of 70%, on par with non-experts (74%) but significantly lower than physicians (91%). Despite GPT-3's notable performance, the study raises ethical concerns, particularly regarding the model's potential to perpetuate existing data biases, exhibiting racial and gender biases and occasionally producing misleading or incorrect information.

A recent study by Liu et al. [229] introduced a framework named PharmacyGPT, which leverages the current GPT family models to emulate the role of clinical pharmacists. This research utilized real data from the ICU at the University of North Carolina Chapel Hill (UNC) Hospital. PharmacyGPT was applied to tackle various challenges in the realm of pharmacy, encompassing patient outcome studies, AI-based medication prescription generation, and interpretable patient clustering analysis. The study revealed that the GPT-4 model, when provided with dynamic context and similar samples, attained the highest accuracy among all models tested. However, the precision and recall scores were not notably high across the approaches. This outcome may be caused by the binary nature of mortality prediction, a significant imbalance in the dataset, and the complex, individualized nature of ICU pharmacy regimens. The research highlights the need for custom evaluation metrics to assess the performance of AI-generated medication plans, enhancing understanding of the models' strengths and limitations.

## 4. Discussion

This section explores the challenges and opportunities arising from the integration of large language models in healthcare.

### 4.1. Evaluation challenges

The evaluation of language models in the medical domain is a multifaceted challenge. One major issue arises from the technical complexity of assessing model performance on tasks with minimal human supervision. For instance, while classification tasks benefit from well-established metrics such as accuracy, precision, recall, and F-measure, evaluating models on more complex tasks, such as medical conversations, remains technically challenging. These tasks often require human assessment or intricate evaluation frameworks [185,186,230,231].

Another challenge arises from unique domain-specific issues, where standard model evaluation from a purely technical perspective is often insufficient. Abbasian et al. [232] categorize evaluation metrics into *intrinsic* metrics, which measure internal language proficiency and coherence, and *extrinsic* metrics, which assess real-world impact and the model's ability to meet human-centric expectations. Within extrinsic metrics, the authors identify general metrics applicable to all tasks, such as trustfulness, bias, and toxicity, as well as domain-specific metrics, such as up-to-dateness and empathy. While the intrinsic/extrinsic categorization provides a useful theoretical framework for understanding evaluation metrics, in practice, metrics are often organized based on their functional role in model evaluation. Table 5 presents a task-oriented categorization, where metrics are grouped by their evaluation purpose.

**Table 5. Evaluation metrics for major tasks in medical domain using language models.**

| Tasks  | Task-Specific Metrics  | Extrinsic Generic Metrics and Benchmarks   | Extrinsic Domain-Specific Metrics and Benchmarks    |
|--|--|--|---|
| <b>Classification and Information Extraction:</b> Clinical Acronym Disambiguation, Eponyms Disambiguation, Phenotyping, Medical Coding, Modeling Patient Timeline, Social Media Monitoring, Querying Data from Electronic Health Records, Extracting Information from Clinical Narrative Reports, Extracting Information From Medical Articles   | Accuracy, Precision, Recall, F-Score, AUC  | Necessary only when using generative models: Metrics for evaluating LLM performance in general (e.g., Academic Benchmarks, Factuality, Complex Reasoning) [5,233]. | Privacy [234] (when patient data is involved)       |
| <b>Generation and Summarization:</b> Medical Report Generation, Clinical Note Generation, Generating Summaries For Laypersons, Generating Summaries for Patient-Provider Dialogues, Generating Textual Descriptions From Graph Models, Summarizing Clinical Study Reports, Summarizing Patient-Provider Dialogues, Simplification of Medical Texts, Simplification of Radiology Reports, Improving Biomedical Text Readability | BLEU, ROUGE, METEOR, Perplexity [14], BERTScore  | Metrics for evaluating LLM performance in general (e.g., Academic Benchmarks, Factuality, Complex Reasoning).  | Reliability [232], Up-to-dateness [235], Privacy    |
| <b>System-Patient Conversation:</b> Mental Health Bots, Medical Chatbots and Health Assistants   | Match Rate, Dialogue Accuracy, Average Request Turn, Complex Chatbot-Specific Metrics (e.g., DEAM [231], DynaEval [230]) | Metrics for evaluating LLM performance in general (e.g., Academic Benchmarks, Factuality, Complex Reasoning), Safety and Bias [232,236–239].                       | Reliability, Up-to-dateness, Privacy, Empathy [232] |

<https://doi.org/10.1371/journal.pdig.0000800.t005>

Unlike task-specific metrics, which are widely adopted and often have clear mathematical formulations, extrinsic metrics involving evaluation of LLMs are significantly more technically complex due to their reliance on specialized frameworks, human alignment, or external models. For example, DynaEval [230] is a framework-based metric that employs a graph convolutional network to evaluate model performance. Benchmarks designed for specific goals, such as TruthfulQA [239], RealToxicityPrompts [238], and BBQ [237], may fail to generalize well to multilingual tasks or domain-specific applications due to their reliance on carefully crafted datasets tailored to specific scenarios. Furthermore, some extrinsic metrics rely on external models for evaluation [231], introducing stochasticity into the process.

Moreover, the probabilistic nature of LLMs makes it challenging to evaluate their performance with respect to domain-specific metrics, especially for unseen input combinations or highly specialized, out-of-distribution data.

Notably, the demand for domain-specific extrinsic metrics grows as models are opened to a broader audience performing more complex tasks. In the healthcare sector, domain-specific evaluation is particularly critical due to its unique ethical and regulatory requirements. Hond et al. [240] highlight the importance of complementing general and domain-specific evaluation with *clinical impact validation*, a process that assesses outcomes such as improved health results, higher patient satisfaction, or reduced administrative burdens. However, to the best of our knowledge, no existing frameworks can automatically perform clinical impact validation with reasonable accuracy.

## 4.2. Ethical issues

The application of language models in clinical settings raises substantial ethical concerns. To address these, foundational principles such as respect of persons, beneficence, and justice, as outlined in the Belmont Report [241], provide a guiding framework. Solomonides et al. [242] further expand this by emphasizing technical principles like fairness, interpretability, and



explainability, alongside organizational principles such as transparency, accountability, and benevolence. While these principles offer a comprehensive ethical blueprint, implementing them effectively in real-world systems remains a considerable challenge. To navigate these challenges, we propose a *tiered approach* that establishes progressively stricter *levels* of compliance with ethical principles. Each level represents a set of actionable system properties derived from ethical principles, ensuring they are translated into concrete requirements. Systems must satisfy the foundational properties of lower levels before progressing to higher ones. This structure enables a practical approach to balancing ethical integrity with technological feasibility while accommodating gradual improvements in system capabilities.

**Level 1** establishes a baseline set of system properties required for safe, minimal-risk usage. Given the inherent risk of generating misleading or factually incorrect content [177, 187,220], the first Level 1 requirement focuses on the principle of *nonmaleficence*. Additionally, ensuring the confidentiality and security of patient data is critical, as emphasized by Ong et al. [243], to uphold patient autonomy and system dependability, both central to Level 1 requirements. These safeguards can be achieved through strict data management protocols and local deployments that minimize data leakage. Robust anonymization techniques [179] play a key role in securing patient privacy when using data externally.

Systems that comply with this level may be acceptable in lower-risk contexts, where errors can be quickly identified and corrected by qualified professionals. Adhering to Level 1 properties enables the ethical use of language models for a variety of tasks, including most classification and information extraction tasks, as well as many summarization and paraphrasing tasks.

**Level 2** builds upon Level 1 compliance and introduces additional principles: fairness, interpretability, auditability, and knowledge management. These enhanced requirements enable systems to perform more sensitive tasks, such as decision support, while avoiding direct system-patient interactions. However, modern LLMs often lack interpretability and remain susceptible to bias [244–246], a challenge further exacerbated by the absence of well-established, widely applicable metrics to measure these properties. This makes the integration of Level 2 principles an open research area. As a result, the deployment of LLMs in high-stakes applications, such as differential diagnosis, triaging, or modeling patient timelines, should generally involve human-in-the-loop supervision to mitigate risks.

**Level 3** encompasses the principles of beneficence, justice, explainability, and benevolence, which are critical for systems capable of interacting with patients without supervision from healthcare professionals. A system can achieve Level 3 only after fully satisfying the requirements of Levels 1 and 2. Attaining this tier would enable widespread, ethically responsible deployment across all reviewed medical tasks. However, to the best of our knowledge, no current models or systems meet this standard. Table 6 summarizes the proposed hierarchical levels, outlines scenarios of use for models and systems that comply with the corresponding ethical principles adopted from Solomonides et al. [242], and provides representative tasks for each level.

From an ethical standpoint, current language model systems are far from fully ready for deployment in high-stakes clinical settings. While many systems can meet the baseline requirements of Level 1 for low-risk tasks under professional supervision, substantial challenges remain in achieving compliance with higher levels. The lack of widely adopted interpretability, fairness, and accountability metrics hampers progress toward Level 2. To our best knowledge, no existing systems currently meet the stringent requirements of Level 3, which are essential for patient-facing applications.

**Table 6. Levels of ethical compliance and corresponding scenarios in system or model applications.**

| Level          | Ethical Principles Compliance  | Scenarios of Use  | Sample Tasks  |
|----------------|--|---|---|
| <b>Level 1</b> | Nonmaleficence, Autonomy, Dependability  | <ul style="list-style-type: none"> <li>- Non-critical tasks with low risk to patients</li> <li>- Tasks under supervision of professionals</li> <li>- Tasks that exclude direct interaction of patients with a system</li> </ul> | Phenotyping, Medical Coding, Eponyms Disambiguation, Clinical Acronym Disambiguation, Generating Textual Descriptions from Graph Models, Generating Summaries for Patient-Provider Dialogues, Clinical Note Generation, Medical Report Generation, Querying Data from Electronic Health Records, Extracting Information from Clinical Narrative Reports, Extracting Information from Medical Articles, Summarizing Clinical Study Reports, Summarizing Patient-Provider Dialogues, Anonymization of Medical Documents |
| <b>Level 2</b> | Fairness, Interpretability, Auditability, Knowledge Management, Accountability | <ul style="list-style-type: none"> <li>- Tasks that involve decision-making without direct interaction of patients with a system</li> </ul>   | Triaging, Differential Diagnosis, Social Media Monitoring, Modeling Patient Timeline  |
| <b>Level 3</b> | Beneficence, Justice, Explainability, Benevolence                              | <ul style="list-style-type: none"> <li>- Patient-facing applications</li> </ul>   | Medical Chatbots and Health Assistants  |

<https://doi.org/10.1371/journal.pdig.0000800.t006>

### 4.3. Datasets

With new applications of textual AI emerging in areas like medication plan generation, triaging, extracting structured data from medical records, and providing medical consultations, the development of novel, open, and de-identified datasets becomes increasingly necessary. Many existing datasets were created before the advent of LLMs, which may inflate study results and lead to an overestimation of the current models' efficacy. Moreover, access to many existing datasets requires special approvals, which hinders widespread research in this area. Future efforts should focus on creating and utilizing open datasets specifically designed to evaluate LLMs in the medical domain to more accurately reflect their true capabilities.

### 4.4. Human-AI collaboration

Further research is required to enhance our understanding and optimization of human-AI collaboration in healthcare. This includes exploring how medical professionals can best interact with and leverage AI tools for improved decision-making and patient care, as well as reducing routine work to help prevent burnout. An example of this could be the further exploration of AI-in-the-loop agents, similar to those described in [228].

### 4.5. Necessity for empirical studies

Empirical research on real-world use-cases of AI in healthcare is essential. Theoretical studies have broadened our understanding, but practical challenges in real healthcare environments, such as hospitals and clinics, are less understood. Research should focus on how AI applications integrate with healthcare systems, their impact on workflows and healthcare professionals, and the long-term effects on patient outcomes, staff efficiency, and costs. Additionally, addressing AI implementation challenges, including data privacy, ethical concerns, and the need for ongoing system training and updates, is vital. This will guide best practices for AI integration, reduce risks, and ensure these technologies effectively enhance patient care and healthcare delivery.

## 5. Conclusion

This study provides an in-depth examination of recent advancements in language models within the medical domain, with a particular emphasis on clinical applications and locally deployable solutions. It traces the development of language models, exploring both general-purpose and domain-specific architectures, and evaluates their role in medical contexts. The study highlights key tasks performed by these models, such as text generation, token classification, and question answering, demonstrating their practical utility through real-world healthcare scenarios.

Recent advancements in the field have introduced more comprehensive approaches, particularly through multimodal models that seamlessly integrate visual and textual data. These innovations enable holistic AI solutions and are bolstered by techniques such as parameter-efficient fine-tuning and flash attention, which significantly reduce computational requirements. The rise of generative LLMs with in-context learning capabilities marks a pivotal evolution, unlocking new possibilities in specialized medical domains like radiology report generation and medical chatbots - tasks that were considered unattainable just a decade ago.

However, deploying language models in healthcare presents several challenges. The first challenge lies in evaluating generative models, especially for complex tasks like medical conversations or summarizations, where task-specific metrics are insufficient, and extrinsic evaluations require intricate frameworks or human alignment. Additionally, critical domain-specific requirements, such as privacy, up-to-dateness, and empathy, are difficult to quantify or standardize in healthcare applications. Lastly, the lack of automated frameworks for clinical impact validation complicates the evaluation process further, hindering the ability to assess real-world outcomes, such as improved patient care or administrative efficiency.

The second challenge revolves around the ethical deployment of language models in clinical settings. While many systems meet baseline compliance for low-risk tasks under professional supervision, advancing to higher levels of ethical compliance remains demanding. Requirements such as fairness and interpretability are hindered by the absence of widely adopted metrics and the persistent biases in modern LLMs. Fully satisfying the comprehensive set of ethical standards essential for patient-facing applications is particularly daunting, and, to the best of our knowledge, no current systems meet these stringent requirements.

Mitigating these challenges requires both technological advancements and focused research. Data privacy concerns can be addressed through locally deployed models or robust anonymization techniques, while the up-to-dateness requirement can be managed using retrieval-augmented generation techniques for generative models or continuous adaptation for encoder-based models. These advancements make a wide range of tasks, those that exclude direct patient interaction and unsupervised decision-making, technically viable for broader adoption within medical organizations. However, even these tasks may require additional regulatory approvals and evaluation in real-world clinical settings before widespread implementation [247].

For tasks involving unsupervised decision-making, addressing privacy and model up-to-dateness alone may not suffice. This is especially true for generative models producing unstructured text, as there is a lack of widely adopted automated metrics for assessing both ethical compliance and performance. A practical approach for handling such tasks is to reformulate them into well-defined tasks, such as classification or information extraction, along with incorporating the latest developments in reasoning and interpretability [248–250]. In these contexts, established task-specific metrics are sufficient to evaluate performance, and mathematical formulations for assessing ethical components like fairness [209] are readily available. If the approach with reformulating the tasks is not feasible, research exploring the

integration of LLMs with ontologies, graph attention networks, and other more deterministic and interpretable models represents another promising direction in moving the adoption of models feasible for unsupervised decision-making.

Finally, we find no existing solutions that are ethically and technically prepared for tasks involving direct interaction with patients beyond purely administrative functions.

In general, the specific demands of the medical field, where precision is paramount due to the high cost of errors and ethical compliance must address a wide range of principles such as interpretability, fairness, and accountability, result in an unavoidable *temporal gap* between technological advancements and their adoption. To address this gap, future research, besides purely technological advancements, should prioritize empirical studies in real-world settings to explore how AI can be seamlessly integrated into healthcare workflows without increasing the burden on healthcare professionals. Additionally, it should assess the impact of AI on patient care and its long-term implications for outcomes and costs. We also advocate for the development of **Medical Model Cards**, inspired by the generic Model Cards proposed by Mitchell et al. [251]. These Medical Model Cards should accompany all models intended for use in clinical settings, providing detailed information about compliance with ethical principles, validated through appropriate metrics or benchmarks. They should also include the intended tasks and corresponding performance benchmarks. This framework will facilitate quicker and more informed model selection for adoption in the healthcare domain.

## Author contributions

**Conceptualization:** Nikita Neveditsin.

**Investigation:** Nikita Neveditsin.

**Methodology:** Nikita Neveditsin.

**Supervision:** Pawan Lingras, Vijay Mago.

**Writing – original draft:** Nikita Neveditsin.

**Writing – review & editing:** Nikita Neveditsin.

## References

1. Goldfarb A. Pause artificial intelligence research? Understanding AI policy challenges. *Canadian J Econ/Revue canadienne d'économie*. 2024;57(2):363–77. <https://doi.org/10.1111/caje.12705>
2. Secinaro S, Calandra D, Secinaro A, Muthurangu V, Biancone P. The role of artificial intelligence in healthcare: a structured literature review. *BMC Med Inform Decis Mak*. 2021;21(1):125. <https://doi.org/10.1186/s12911-021-01488-9> PMID: 33836752
3. Bommasani R, Hudson D, Adeli E, Altman R, Arora S, von Arx S. On the opportunities and risks of foundation models. *arXiv preprint*. 2021. <https://arxiv.org/abs/2108.07258>
4. OpenAI. GPT-4 technical report. 2023.
5. Team G, Anil R, Borgeaud S, Alayrac JB, Yu J, Soricut R, et al. Gemini: a family of highly capable multimodal models; 2024. Available from: <https://arxiv.org/abs/2312.11805>
6. Touvron H, Martin L, Stone K, Albert P, Almahairi A, Babaei Y, et al. Llama 2: open foundation and fine-tuned chat models. 2023.
7. Luo X, Deng Z, Yang B, Luo MY. Pre-trained language models in medicine: a survey. *Artif Intell Med*. 2024;154102904. <https://doi.org/10.1016/j.artmed.2024.102904> PMID: 38917600
8. Meng X, Yan X, Zhang K, Liu D, Cui X, Yang Y, et al. The application of large language models in medicine: a scoping review. *iScience*. 2024;27(5):109713. <https://doi.org/10.1016/j.isci.2024.109713> PMID: 38746668
9. Wang B, Xie Q, Pei J, Chen Z, Tiwari P, Li Z. Pre-trained language models in biomedical domain: a systematic survey. *ACM Comput Surv*. 2023;56(3):1–52.

10. Thirunavukarasu AJ, Ting DSJ, Elangovan K, Gutierrez L, Tan TF, Ting DSW. Large language models in medicine. *Nat Med*. 2023;29(8):1930–40. <https://doi.org/10.1038/s41591-023-02448-8> PMID: 37460753
11. Yang R, Tan TF, Lu W, Thirunavukarasu AJ, Ting DSW, Liu N. Large language models in health care: development, applications, and challenges. *Health Care Sci*. 2023;2(4):255–63. <https://doi.org/10.1002/hcs2.61> PMID: 38939520
12. Russell SJ, Norvig P. *Artificial intelligence: a modern approach*. Pearson; 2016.
13. IBM. AI vs. Machine Learning vs. Deep Learning vs. Neural Networks; n.d. Available from: <https://www.ibm.com/think/topics/ai-vs-machine-learning-vs-deep-learning-vs-neural-network>
14. Jurafsky D, Martin JH. *Speech and language processing: an introduction to natural language processing, computational linguistics, and speech recognition with language models*; 2024. Available from: <https://web.stanford.edu/~jurafsky/slp3/>
15. Roberts A, Raffel C, Shazeer N. How much knowledge can you pack into the parameters of a language model? In: Webber B, Cohn T, He Y, Liu Y, editors. *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Online: Association for Computational Linguistics; 2020. p. 5418–26. Available from: <https://aclanthology.org/2020.emnlp-main.437>
16. Touvron H, Lavril T, Izacard G, Martinet X, Lachaux M, Lacroix T, et al. LLaMA: open and efficient foundation language models. 2023.
17. Bahdanau D, Cho K, Bengio Y. Neural machine translation by jointly learning to align and translate. In: Bengio Y, LeCun Y, editors. *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, 2015 May 7–9, Conference Track Proceedings*; 2015. Available from: <http://arxiv.org/abs/1409.0473>
18. Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez A, et al. Attention is all you need. *Adv Neural Inf Process Syst*. 2017;30:5998–6008. <https://doi.org/10.5555/3295222.3295349>
19. Radford A, Wu J, Child R, Luan D, Amodei D, Sutskever I. Language models are unsupervised multitask learners; 2019. Available from: <https://api.semanticscholar.org/CorpusID:160025533>
20. OpenAI. ChatGPT; 2023. <https://openai.com/chatgpt>
21. Anthropic. Model Card for Claude 3; 2024. Available from: [https://www-cdn.anthropic.com/de8ba9b01c9ab7cbabf5c33b80b7bbc618857627/Model\\_Card\\_Claude\\_3.pdf](https://www-cdn.anthropic.com/de8ba9b01c9ab7cbabf5c33b80b7bbc618857627/Model_Card_Claude_3.pdf)
22. Norvig P. On Chomsky and the two cultures of statistical learning. *Berechenbarkeit der Welt? Philosophie und Wissenschaft im Zeitalter von Big Data*. 2017; p. 61–83.
23. Liu Z, Huang Y, Yu X, Zhang L, Wu Z, Cao C, et al. DeID-GPT: Zero-shot Medical Text De-Identification by GPT-4; 2023. Available from: <http://arxiv.org/abs/2303.11032>
24. Brown T, Mann B, Ryder N, Subbiah M, Kaplan J, Dhariwal P. Language models are few-shot learners. *CoRR*. 2020. <https://arxiv.org/abs/2005.14165>
25. Kenton JDMWC, Toutanova LK. Bert: pre-training of deep bidirectional transformers for language understanding. In: *Proceedings of naacL-HLT*. vol. 1; 2019. p. 2.
26. Wu J, Wu X, Qiu Z, Li M, Lin S, Zhang Y, et al. Large language models leverage external knowledge to extend clinical insight beyond language boundaries. *J Am Med Inform Assoc: JAMIA*. 2023;ocae079.
27. Mahowald K, Ivanova A, Blank I, Kanwisher N, Tenenbaum J, Fedorenko E. Dissociating language and thought in large language models. *Trends Cognit Sci*. 2023.
28. Team C. Chameleon: mixed-modal early-fusion foundation models; 2024. Available from: <https://arxiv.org/abs/2405.09818>
29. Grattafiori A, Dubey A, Jauhri A, Pandey A, Kadian A, Al-Dahle A, et al. The Llama 3 Herd of Models; 2024. Available from: <https://arxiv.org/abs/2407.21783>
30. Black S, Gao L, Wang P, Leahy C, Biderman S. GPT-Neo: large scale autoregressive language modeling with mesh-tensorflow; 2021. Available from: <https://doi.org/10.5281/zenodo.5297715>
31. Lan Z, Chen M, Goodman S, Gimpel K, Sharma P, Soricut R. ALBERT: a lite BERT for self-supervised learning of language representations; 2020. Available from: <https://arxiv.org/abs/1909.11942>
32. He P, Liu X, Gao J, Chen W. DeBERTa: decoding-enhanced BERT with disentangled attention; 2021. Available from: <https://arxiv.org/abs/2006.03654>
33. Clark K, Luong MT, Le QV, Manning CD. ELECTRA: pre-training text encoders as discriminators rather than generators. In: *ICLR*; 2020. Available from: <https://openreview.net/pdf?id=r1xMH1BtvB>
34. Liu Y, Ott M, Goyal N, Du J, Joshi M, Chen D, et al. RoBERTa: a robustly optimized BERT pretraining approach; 2019. Available from: <https://arxiv.org/abs/1907.11692>

35. Lewis M, Liu Y, Goyal N, Ghazvininejad M, Mohamed A, Levy O. BART: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. CoRR. 2019. <https://arxiv.org/abs/1910.13461>
36. Rothe S, Narayan S, Severyn A. Leveraging pre-trained checkpoints for sequence generation tasks. Trans Assoc Comput Linguist. 2020;8:264–80. [https://doi.org/10.1162/tacl\\_a\\_00313](https://doi.org/10.1162/tacl_a_00313)
37. Chung HW, Hou L, Longpre S, Zoph B, Tay Y, Fedus W, et al.. Scaling instruction-finetuned language models; 2022. Available from: <https://arxiv.org/abs/2210.11416>
38. Zhang J, Zhao Y, Saleh M, Liu P. Pegasus: pre-training with extracted gap-sentences for abstractive summarization. 2019.
39. Raffel C, Shazeer N, Roberts A, Lee K, Narang S, Matena M. Exploring the limits of transfer learning with a unified text-to-text transformer. CoRR. 2019. <https://arxiv.org/abs/1910.10683>
40. Li D, Li J, Le H, Wang G, Savarese S, Hoi SCH. LAVIS: a library for language-vision intelligence. 2022.
41. Radford A, Kim JW, Hallacy C, Ramesh A, Goh G, Agarwal S, et al. Learning transferable visual models from natural language supervision; 2021. Available from: <https://arxiv.org/abs/2103.00020>
42. Liu F, Eisenschlos JM, Piccinno F, Krichene S, Pang C, Lee K, et al.. DePlot: one-shot visual language reasoning by plot-to-table translation; 2023. Available from: <https://arxiv.org/abs/2212.10505>
43. Kim G, Hong T, Yim M, Nam J, Park J, Yim J, et al.. OCR-free document understanding transformer; 2022. Available from: <https://arxiv.org/abs/2111.15664>
44. Huang Y, Lv T, Cui L, Lu Y, Wei F. LayoutLMv3: pre-training for document AI with unified text and image masking; 2022. Available from: <https://arxiv.org/abs/2204.08387>
45. Keskar NS, McCann B, Varshney LR, Xiong C, Socher R. CTRL: A Conditional Transformer Language Model for Controllable Generation. CoRR. 2019; <https://arxiv.org/abs/1909.05858>
46. Almazrouei E, Alobeidli H, Alshamsi A, Cappelli A, Cojocaru R, Debbah M, et al.. The falcon series of open language models; 2023. Available from: <https://arxiv.org/abs/2311.16867>
47. Wang B, Komatsuzaki A. GPT-J-6B: a 6 billion parameter autoregressive language model. 2021. <https://github.com/kingoflolz/mesh-transformer-jax>
48. Team G. Gemma: open models based on gemini research and technology; 2024. Available from: <https://arxiv.org/abs/2403.08295>
49. AI M. Announcing Mistral 7B; <https://mistral.ai/news/announcing-mistral-7b/>
50. Abidin M, Jacobs SA, Awan AA, Aneja J, Awadallah A, Awadalla H, et al. Phi-3 technical report: a highly capable language model locally on your phone; 2024. Available from: <https://arxiv.org/abs/2404.14219>
51. Guo M, Ainslie J, Uthus D, Ontanon S, Ni J, Sung YH, et al. LongT5: efficient text-to-text transformer for long sequences; 2022. Available from: <https://arxiv.org/abs/2112.07916>
52. Adept. Fuyu-8B;. <https://www.adapt.ai/blog/fuyu-8b>
53. Liu H, Li C, Wu Q, Lee YJ. Visual instruction tuning; 2023. Available from: <https://arxiv.org/abs/2304.08485>
54. Beyer L, Steiner A, Pinto AS, Kolesnikov A, Wang X, Salz D, et al. PaliGemma: a versatile 3B VLM for transfer; 2024. Available from: <https://arxiv.org/abs/2407.07726>
55. Black S, Biderman S, Hallahan E, Anthony Q, Gao L, Golding L, et al. GPT-NeoX-20B: An open-source autoregressive language model. In: Proceedings of the ACL Workshop on Challenges & Perspectives in Creating Large Language Models; 2022. Available from: <https://arxiv.org/abs/2204.06745>
56. Tay Y, Dehghani M, Tran VQ, Garcia X, Wei J, Wang X, et al. UL2: unifying language learning paradigms; 2023. Available from: <https://arxiv.org/abs/2205.05131>
57. Cohere. Cohere Command R. <https://cohere.com/command>
58. Databricks. Introducing DbrX: a new state-of-the-art open LLM;. <https://www.databricks.com/blog/introducing-dbrx-new-state-art-open-llm>
59. Zhang S, Roller S, Goyal N, Artetxe M, Chen M, Chen S, et al. OPT: open pre-trained transformer language models; 2022. Available from: <https://arxiv.org/abs/2205.01068>
60. Workshop B, Team B. BLOOM: a 176B-parameter open-access multilingual language model; 2023. Available from: <https://arxiv.org/abs/2211.05100>
61. Chung H, Hou L, Longpre S, Zoph B, Tay Y, Fedus W. Scaling instruction-finetuned language models. J Mach Learn Res. 2024;25(70):1–53.
62. Hu EJ, Shen Y, Wallis P, Allen-Zhu Z, Li Y, Wang S, et al.. LoRA: Low-Rank Adaptation of Large Language Models; 2021



63. Jacob B, Kligys S, Chen B, Zhu M, Tang M, Howard A, et al. Quantization and training of neural networks for efficient integer-arithmetic-only inference; 2017. Available from: <https://arxiv.org/abs/1712.05877>
64. Dettmers T, Pagnoni A, Holtzman A, Zettlemoyer L. QLoRA: efficient finetuning of quantized LLMs. In: Oh A, Naumann T, Globerson A, Saenko K, Hardt M, Levine S, editors. Advances in neural information processing systems. vol. 36. Curran Associates, Inc.; 2023. p. 10088–115. Available from: [https://proceedings.neurips.cc/paper\\_files/paper/2023/file/1feb87871436031bdc0f2beaa62a049b-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/1feb87871436031bdc0f2beaa62a049b-Paper-Conference.pdf)
65. Guo C, Zhang R, Xu J, Leng J, Liu Z, Huang Z, et al. GMLake: efficient and transparent GPU memory defragmentation for large-scale DNN training with virtual memory stitching. In: Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2. ASPLOS '24. New York, NY, USA: Association for Computing Machinery; 2024. p. 450–66. Available from: <https://doi.org/10.1145/3620665.3640423>
66. Dao T, Fu DY, Ermon S, Rudra A, Re C. FlashAttention: fast and memory-efficient exact attention with IO-awareness; 2022. Available from: <https://arxiv.org/abs/2205.14135>
67. Rajpurkar P, Zhang J, Lopyrev K, Liang P. Squad: 100,000+ questions for machine comprehension of text. arXiv preprint 2016. <https://arxiv.org/abs/1606.05250>
68. Lai G, Xie Q, Liu H, Yang Y, Hovy E. RACE: Large-scale ReAding comprehension dataset from examinations. In: Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing; 2017. p. 785–94.
69. Wang A, Pruksachatkun Y, Nangia N, Singh A, Michael J, Hill F, et al. SuperGlue: a stickier benchmark for general-purpose language understanding systems. Adv Neural Inf Process Syst. 2019;32.
70. Wang A. Glue: a multi-task benchmark and analysis platform for natural language understanding. arXiv preprint. 2018. <https://arxiv.org/abs/1804.07461>
71. Benchmark G. GLUE benchmark leaderboard; 2024. <https://gluebenchmark.com/leaderboard/>
72. Hendrycks D, Burns C, Basart S, Zou A, Mazeika M, Song D, et al. Measuring massive multitask language understanding. arXiv preprint 2020. <https://arxiv.org/abs/2009.03300>
73. Wang Y, Ma X, Zhang G, Ni Y, Chandra A, Guo S, et al. Mmlu-pro: A more robust and challenging multi-task language understanding benchmark. arXiv preprint 2024. <https://arxiv.org/abs/2406.01574>
74. Narayan S, Cohen SB, Lapata M. Don't give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization; 2018. Available from: <https://arxiv.org/abs/1808.08745>
75. Singh A, Natarajan V, Shah M, Jiang Y, Chen X, Batra D, et al. Towards VQA models that can read; 2019. Available from: <https://arxiv.org/abs/1904.08920>
76. Mathew M, Karatzas D, Jawahar CV. DocVQA: a dataset for VQA on document images; 2021. Available from: <https://arxiv.org/abs/2007.00398>
77. Masry A, Long DX, Tan JQ, Joty S, Hoque E. ChartQA: a benchmark for question answering about charts with visual and logical reasoning; 2022. Available from: <https://arxiv.org/abs/2203.10244>
78. Zhang X, Tian C, Yang X, Chen L, Li Z, Petzold L. AlphaCare: instruction-tuned large language models for medical application. 2023. <https://doi.org/10.48550/ARXIV.2310.14558>
79. Li Y, Li Z, Zhang K, Dan R, Jiang S, Zhang Y. ChatDoctor: a medical chat model fine-tuned on a large language model Meta-AI (LLaMA) using medical domain knowledge. 2023.
80. Abbasian M, Azimi I, Rahmani AM, Jain R. Conversational health agents: a personalized LLM-powered agent framework; 2023. Available from: <http://arxiv.org/abs/2310.02374>
81. Toma A, Lawler PR, Ba J, Krishnan RG, Rubin BB, Wang B. Clinical camel: an open expert-level medical language model with dialogue-based knowledge encoding. 2023. <https://doi.org/10.48550/ARXIV.2305.12031>
82. Huang K, Altosaar J, Ranganath R. ClinicalBERT: modeling clinical notes and predicting hospital readmission; 2020.
83. Johnson A, Pollard T, Shen L, Lehman LW, Feng M, Ghassemi M, et al. MIMIC-III, a freely accessible critical care database; 2016. Available from: <https://paperswithcode.com/dataset/mimic-iii>
84. Lee J, Yoon W, Kim S, Kim D, Kim S, So CH, et al. BioBERT: a pre-trained biomedical language representation model for biomedical text mining. Bioinformatics. 2020;36(4):1234–40. <https://doi.org/10.1093/bioinformatics/btz682> PMID: 31501885
85. Chakraborty S, Bisong E, Bhatt S, Wagner T, Elliott R, Mosconi F. BioMedBERT: a pre-trained biomedical language model for QA and IR. In: Scott D, Bel N, Zong C, editors. Proceedings of the 28th International Conference on Computational Linguistics. Barcelona, Spain (Online):

- International Committee on Computational Linguistics; 2020. p. 669–79. Available from: <https://aclanthology.org/2020.coling-main.59>
86. Gu Y, Tinn R, Cheng H, Lucas M, Usuyama N, Liu X. Domain-specific language model pretraining for biomedical natural language processing. *ACM Trans Comput Healthc*. 2021;3(1):1–23. <https://doi.org/10.1145/3458754>
  87. Li Y, Rao S, Solares JRA, Hassaine A, Ramakrishnan R, Canoy D, et al. BEHRT: transformer for electronic health records. *Sci Rep*. 2020;10(1):7155. <https://doi.org/10.1038/s41598-020-62922-y> PMID: 32346050
  88. Yang X, Chen A, PourNejatian N, Shin HC, Smith KE, Parisien C, et al. GatorTron: a large clinical language model to unlock patient information from unstructured electronic health records; 2022. Available from: <http://arxiv.org/abs/2203.03540>
  89. Luo R, Sun L, Xia Y, Qin T, Zhang S, Poon H, et al. BioGPT: generative pre-trained transformer for biomedical text generation and mining. *Brief Bioinform*. 2022;23(6):bbac409. <https://doi.org/10.1093/bib/bbac409> PMID: 36156661
  90. Li C, Wong C, Zhang S, Usuyama N, Liu H, Yang J. Llava-med: Training a large language-and-vision assistant for biomedicine in one day. *Adv Neural Inf Process Syst*. 2024;36.
  91. Chen Z, Varma M, Delbrouck JB, Paschali M, Blankemeier L, Veen DV, et al. CheXagent: towards a foundation model for chest X-Ray interpretation; 2024. Available from: <https://arxiv.org/abs/2401.12208>
  92. Doğan RI, Leaman R, Lu Z. NCBI disease corpus: a resource for disease name recognition and concept normalization. *J Biomed Inform*. 2014;47:1–10. <https://doi.org/10.1016/j.jbi.2013.12.006> PMID: 24393765
  93. Uzuner Ö, South BR, Shen S, DuVall SL. 2010 i2b2/VA challenge on concepts, assertions, and relations in clinical text. *J Am Med Inform Assoc*. 2011;18(5):552–6. <https://doi.org/10.1136/amiajnl-2011-000203> PMID: 21685143
  94. Li J, Sun Y, Johnson RJ, Sciaky D, Wei CH, Leaman R. BioCreative V CDR task corpus: a resource for chemical disease relation extraction. *Database*. 2016;2016.
  95. Krallinger M, Leitner F, Rabal O, Vazquez M, Oyarzabal J, Valencia A. CHEMDNER: the drugs and chemical names extraction challenge. *J Cheminform*. 2015;7(Suppl 1 Text mining for chemistry and the CHEMDNER track):S1. <https://doi.org/10.1186/1758-2946-7-S1-S1> PMID: 25810766
  96. Smith L, Tanabe L, Ando Rj, Kuo C, Chung I, Hsu C, et al. Overview of BioCreative II gene mention recognition. *Genome Biology*. 2008;9(1):1–19.
  97. Huang M, Lai P, Tsai R, Hsu W. Revised JNLPBA corpus: a revised version of biomedical NER corpus for relation extraction task. *arXiv preprint*. 2019. <https://arxiv.org/abs/1901.10219>
  98. Gerner M, Nenadic G, Bergman CM. LINNAEUS: a species name identification system for biomedical literature. *BMC Bioinformatics*. 2010;11:85. <https://doi.org/10.1186/1471-2105-11-85> PMID: 20149233
  99. Pafilis E, Frankild SP, Fanini L, Faulwetter S, Pavloudi C, Vasileiadou A, et al. The SPECIES and ORGANISMS resources for fast and accurate identification of taxonomic names in text. *PLoS One*. 2013;8(6):e65390. <https://doi.org/10.1371/journal.pone.0065390> PMID: 23823062
  100. Bravo Á, Piñero J, Queralt-Rosinach N, Rautschka M, Furlong LI. Extraction of relations between genes and diseases from text and large-scale data analysis: implications for translational research. *BMC Bioinformatics*. 2015;16:55. <https://doi.org/10.1186/s12859-015-0472-9> PMID: 25886734
  101. Coloma PM, Schuemie MJ, Trifirò G, Gini R, Herings R, Hippisley-Cox J, et al. Combining electronic healthcare databases in Europe to allow for large-scale drug safety monitoring: the EU-ADR Project. *Pharmacoepidemiol Drug Saf*. 2011;20(1):1–11. <https://doi.org/10.1002/pds.2053> PMID: 21182150
  102. Kringelum J, Kjaerulff SK, Brunak S, Lund O, Oprea TI, Taboureaux O. ChemProt-3.0: a global chemical biology diseases mapping. *Database (Oxford)*. 2016;2016:bav123. <https://doi.org/10.1093/database/bav123> PMID: 26876982
  103. Tsatsaronis G, Balikas G, Malakasiotis P, Partalas I, Zschunke M, Alvers MR, et al. An overview of the BIOASQ large-scale biomedical semantic indexing and question answering competition. *BMC Bioinformatics*. 2015;16:138. <https://doi.org/10.1186/s12859-015-0564-6> PMID: 25925131
  104. Nye B, Jessy Li J, Patel R, Yang Y, Marshall IJ, Nenkova A, et al. A corpus with multi-level annotations of patients, interventions and outcomes to support language processing for medical literature. *Proc Conf Assoc Comput Linguist Meet*. 2018;2018:197–207. PMID: 30305770
  105. Herrero-Zazo M, Segura-Bedmar I, Martínez P, Declerck T. The DDI corpus: an annotated corpus with pharmacological substances and drug-drug interactions. *J Biomed Inform*. 2013;46(5):914–20. <https://doi.org/10.1016/j.jbi.2013.07.011> PMID: 23906817
  106. Sogancioglu G, Öztürk H, Özgür A. BIOSSES: a semantic sentence similarity estimation system for the biomedical domain. *Bioinformatics*. 2017;33(14):i49–58. <https://doi.org/10.1093/bioinformatics/btx238> PMID: 28881973

107. Hanahan D, Weinberg RA. The hallmarks of cancer. *Cell*. 2000;100(1):57–70. [https://doi.org/10.1016/S0092-8674\(00\)81683-9](https://doi.org/10.1016/S0092-8674(00)81683-9) PMID: 10647931
108. Jin Q, Dhingra B, Liu Z, Cohen W, Lu X. Pubmedqa: a dataset for biomedical research question answering. *arXiv preprint*. 2019. <https://arxiv.org/abs/1909.06146>
109. Herrett E, Gallagher AM, Bhaskaran K, Forbes H, Mathur R, van Staa T, et al. Data resource profile: clinical practice research datalink (CPRD). *Int J Epidemiol*. 2015;44(3):827–36. <https://doi.org/10.1093/ije/dyv098> PMID: 26050254
110. University of Florida Health. Integrated Data Repository (IDR). Available from: <https://idr.uflhealth.org/>
111. Henry S, Buchan K, Filannino M, Stubbs A, Uzuner O. 2018 n2c2 shared task on adverse drug events and medication extraction in electronic health records. *J Am Med Inform Assoc*. 2020;27(1):3–12. <https://doi.org/10.1093/jamia/ocz166> PMID: 31584655
112. Romanov A, Shivade C. Lessons from natural language inference in the clinical domain. *arXiv preprint*. 2018. <https://arxiv.org/abs/1808.06752>
113. Pampari A, Raghavan P, Liang J, Peng J. emrqa: a large corpus for question answering on electronic medical records. *arXiv preprint* 2018. <https://arxiv.org/abs/1809.00732>
114. Hou Y, Xia Y, Wu L, Xie S, Fan Y, Zhu J, et al. Discovering drug-target interaction knowledge from biomedical literature. *Bioinformatics*. 2022;38(22):5100–7. <https://doi.org/10.1093/bioinformatics/btac648> PMID: 36205562
115. Lu Q, Dou D, Nguyen T. ClinicalT5: A Generative Language Model for Clinical Text. In: Findings of the Association for Computational Linguistics: EMNLP 2022. Abu Dhabi, United Arab Emirates: Association for Computational Linguistics; 2022. p. 5436–5443. Available from: <https://aclanthology.org/2022.findings-emnlp.398>
116. Tran H, Yang Z, Yao Z, Yu H. BioInstruct: instruction tuning of large language models for biomedical natural language processing. *J Am Med Inform Assoc*. 2024;31(9):1821–32. <https://doi.org/10.1093/jamia/ocae122> PMID: 38833265
117. Jin D, Pan E, Oufattole N, Weng W, Fang H, Szolovits P. What disease does this patient have? a large-scale open domain question answering dataset from medical exams. *Appl Sci*. 2021;11(14):6421.
118. Pal A, Umapathi LK, Sankarasubbu M. Medmcqa: a large-scale multi-subject multi-choice dataset for medical domain question answering. *Conference on Health, Inference, and Learning*. PMLR; 2022. p. 248–60.
119. Ben Abacha A, Yim Ww, Adams G, Snider N, Yetisgen M. Overview of the MEDIQA-Chat 2023 shared tasks on the summarization & generation of doctor-patient conversations. In: Naumann T, Ben Abacha A, Bethard S, Roberts K, Rumshisky A, editors. *Proceedings of the 5th Clinical Natural Language Processing Workshop*. Toronto, Canada: Association for Computational Linguistics; 2023. p. 503–513. Available from: <https://aclanthology.org/2023.clinicalnlp-1.52>
120. iCliniq. iCliniq: Online Doctor Consultation;. Available from: <https://www.icliniq.com/>
121. Hendrycks D, Burns C, Basart S, Zou A, Mazeika M, Song D, et al.. Measuring massive multitask language understanding; 2021. Available from: <https://arxiv.org/abs/2009.03300>
122. Examination USML. United States Medical Licensing Examination (USMLE). Available from: <https://www.usmle.org/>
123. Han T, Adams LC, Papaioannou JM, Grundmann P, Oberhauser T, Loser A, et al. MedAlpaca – an open-source collection of medical conversational AI models and training data. 2023. <https://arxiv.org/abs/2304.08247>
124. Wu C, Lin W, Zhang X, Zhang Y, Wang Y, Xie W. PMC-LLaMA: towards building open-source language models for medicine. 2023. <https://arxiv.org/abs/2304.14454>
125. Lo K, Wang LL, Neumann M, Kinney RM, Weld DS. S2ORC: the semantic scholar open research corpus. In: *Annual Meeting of the Association for Computational Linguistics*; 2020. Available from: <https://api.semanticscholar.org/CorpusID:215416146>
126. Zhang S, Xu Y, Usuyama N, Xu H, Bagga J, Tinn R, et al.. BiomedCLIP: a multimodal biomedical foundation model pretrained from fifteen million scientific image-text pairs; 2024. Available from: <https://arxiv.org/abs/2303.00915>
127. Liu B, Zhan LM, Xu L, Wu XM. SLAKE: a semantically-labeled knowledge-enhanced dataset for medical visual question answering; 2021. Available from: <https://www.med-vqa.com/slake/>
128. Labrak Y, Bazoge A, Morin E, Gourraud PA, Rouvier M, Dufour R. BioMistral: A Collection of Open-Source Pretrained Large Language Models for Medical Domains; 2024. Available from: <https://arxiv.org/abs/2402.10373>

129. Johnson AEW, Pollard TJ, Greenbaum NR, Lungren MP, ying Deng C, Peng Y, et al.. MIMIC-CXR-JPG, a large publicly available database of labeled chest radiographs; 2019. Available from: <https://arxiv.org/abs/1901.07042>
130. Bustos A, Pertusa A, Salinas J-M, de la Iglesia-Vayá M. PadChest: A large chest x-ray image dataset with multi-label annotated reports. *Med Image Anal.* 2020;66:101797. <https://doi.org/10.1016/j.media.2020.101797> PMID: 32877839
131. Vayá M, Saborit J, Montell J, Pertusa A, Bustos A, Cazorla M. BIMCV COVID-19: a large annotated dataset of RX and CT images from COVID-19 patients. *arXiv preprint.* 2020. <https://arxiv.org/abs/2006.01174>
132. Irvin J, Rajpurkar P, Ko M, Yu Y, Ciurea-Illcus S, Chute C. Chexpert: a large chest radiograph dataset with uncertainty labels and expert comparison. *Proc AAAI Conf Artif Intell.* 2019;33:590–7.
133. SIIM-ACR. SIIM-ACR pneumothorax segmentation challenge. Available from: <https://www.kaggle.com/c/siim-acr-pneumothorax-segmentation>
134. Saharia C, Chan W, Saxena S, Li L, Whang J, Denton E, et al.. Photorealistic text-to-image diffusion models with deep language understanding; 2022. Available from: <https://arxiv.org/abs/2205.11487>
135. Demner-Fushman D, Antani S, Simpson M, Thoma G. Design and development of a multimodal biomedical information retrieval system. *J Comput Sci Eng.* 2012;6(2):168–77.
136. Wang Z, Liu L, Wang L, Zhou L. R2GenGPT: Radiology report generation with frozen LLMs. *Meta-Radiology.* 2023;1(3):100033. <https://doi.org/10.1016/j.metrad.2023.100033>
137. Chen Z, Luo L, Bie Y, Chen H. Dia-LLaMA: towards large language model-driven CT report generation; 2024
138. Brake N, Schaaf T. Comparing two model designs for clinical note generation; Is an LLM a useful evaluator of consistency? 2024.
139. Eppler MB, Ganjavi C, Knudsen JE, Davis RJ, Ayo-Ajibola O, Desai A, et al. Bridging the gap between urological research and patient understanding: the role of large language models in automated generation of layperson's summaries. *Urol Pract.* 2023;10(5):436–43. <https://doi.org/10.1097/UPJ.0000000000000428> PMID: 37410015
140. Nair V, Schumacher E, Kannan A. Generating medically-accurate summaries of patient-provider dialogue: a multi-stage approach using large language models. In: *Proceedings of the 5th Clinical Natural Language Processing Workshop.* 2023. p. 200–17.
141. Phatak A, Mago VK, Agrawal A, Inbasekaran A, Giabbanelli PJ. Narrating causal graphs with large language models; 2024. Available from: <https://arxiv.org/abs/2403.07118>
142. Tang Y, Yang H, Zhang L, Yuan Y. Work like a doctor: Unifying scan localizer and dynamic generator for automated computed tomography report generation. *Exp Syst Appl.* 2024;237:121442. <https://doi.org/https://doi.org/10.1016/j.eswa.2023.121442>
143. Demner-Fushman D, Kohli MD, Rosenman MB, Shooshan SE, Rodriguez L, Antani S, et al. Preparing a collection of radiology examinations for distribution and retrieval. *J Am Med Inform Assoc.* 2016;23(2):304–10. <https://doi.org/10.1093/jamia/ocv080> PMID: 26133894
144. Johnson A, Pollard T, Mark R, Berkowitz S, Horng S. MIMIC-CXR Database (version 2.0.0); 2019. Available from: <https://physionet.org/content/mimic-cxr/2.0.0/>
145. Irvin J, Rajpurkar P, Ko M, Yu Y, Ciurea-Illcus S, Chute C, et al.. CheXpert: a large chest radiograph dataset with uncertainty labels and expert comparison; 2019. Available from: <https://arxiv.org/abs/1901.07031>
146. Yang Z, Xu X, Yao B, Rogers E, Zhang S, Intille S, et al. Talk2Care: an LLM-based voice assistant for communication between healthcare providers and older adults. *Proc ACM Interact Mob Wearable Ubiquit Technol.* 2024;8(2). <https://doi.org/10.1145/3659625>
147. Kugic A, Kreuzthaler M, Schulz S. Clinical acronym disambiguation via ChatGPT and BING. In: Giacomini M, Stoicu-Tivadar L, Balestra G, Benis A, Bonacina S, Bottrighi A, et al., editors. *Studies in health technology and informatics.* IOS Press; 2023. Available from: <https://ebooks.iospress.nl/doi/10.3233/SHTI230743>
148. Liu Y, Melton GB, Zhang R. Exploring large language models for acronym, symbol sense disambiguation, and semantic similarity and relatedness assessment. *AMIA Jt Summits Transl Sci Proc.* 2024;2024:324–33. PMID: 38827102
149. Toddenroth D. Classifiers of medical eponymy in scientific texts. In: Hagglund M, Blusi M, Bonacina S, Nilsson L, Cort Madsen I, Pelayo S, et al., editors. *Studies in health technology and informatics.* IOS Press; 2023. Available from: <https://ebooks.iospress.nl/doi/10.3233/SHTI230271>
150. Moon S, Pakhomov S, Liu N, Ryan JO, Melton GB. A sense inventory for clinical abbreviations and acronyms created using clinical notes and medical dictionary resources. *J Am Med Inform Assoc.* 2014;21(2):299–307. <https://doi.org/10.1136/amiainl-2012-001506> PMID: 23813539

151. Weeber M, Mork JG, Aronson AR. Developing a test collection for biomedical word sense disambiguation. In: Proceedings of the AMIA Symposium. American Medical Informatics Association; 2001. p. 746
152. Wagh A, Khanna M. Clinical abbreviation disambiguation using clinical variants of BERT. In: Morusupalli R, Dandibhotla TS, Atluri VV, Windridge D, Lingras P, Komati VR, editors. Multi-disciplinary trends in artificial intelligence. Cham: Springer; 2023. p. 214–24
153. Groza T, Caufield H, Gration D, Baynam G, Haendel MA, Robinson PN, et al. An evaluation of GPT models for phenotype concept recognition. arXiv preprint 2023. <https://arxiv.org/abs/2309.17169>
154. López-García G, Jerez JM, Ribelles N, Alba E, Veredas FJ. Explainable clinical coding with in-domain adapted transformers. *J Biomed Inform.* 2023;139:104323. <https://doi.org/10.1016/j.jbi.2023.104323> PMID: 36813154
155. Kraljevic Z, Bean D, Shek A, Bendayan R, Hemingway H, Yeung JA, et al. Foresight-a generative pretrained transformer for modelling of patient timelines using electronic health records: a retrospective modelling study. *Lancet Digit Health.* 2024;6(4):e281–90. [https://doi.org/10.1016/S2589-7500\(24\)00025-6](https://doi.org/10.1016/S2589-7500(24)00025-6) PMID: 38519155
156. Ben Shoham O, Rappoport N. CPLLM: Clinical prediction with large language models. *PLOS Digit Health.* 2024;3(12):e0000680. <https://doi.org/10.1371/journal.pdig.0000680> PMID: 39642102
157. Dus Y, Nefedov G. An automated tool to detect suicidal susceptibility from social media posts; 2023. Available from: <http://arxiv.org/abs/2310.06056>
158. Fisher A, Young MM, Payer D, Pacheco K, Dubeau C, Mago V. Automating detection of drug-related harms on social media: machine learning framework. *J Med Internet Res.* 2023;25:e43630. <https://doi.org/10.2196/43630> PMID: 37725410
159. Komati N. Suicide Watch Dataset; 2019. Available from: <https://www.kaggle.com/datasets/nikhileswarkomati/suicide-watch>
160. Gaur M, Alambo A, Sain JP, Kursuncu U, Thirunarayan K, Kavuluru R, et al. Knowledge-aware assessment of severity of suicide risk for early intervention. In: The World Wide Web Conference. WWW '19. New York, NY, USA: Association for Computing Machinery; 2019. p. 514–25. Available from: <https://doi.org/10.1145/3308558.3313698>
161. Johnson A, Bulgarelli L, Pollard T, Horng S, Celi LA, Mark R. MIMIC-IV (version 2.2); 2023. Available from: <https://physionet.org/content/mimiciv/2.2/>
162. Pollard TJ, Johnson AEW, Raffa JD, Celi LA, Mark RG, Badawi O. The eICU collaborative research database, a freely available multi-center database for critical care research. *Sci Data.* 2018;5:180178. <https://doi.org/10.1038/sdata.2018.178> PMID: 30204154
163. Groza T, Köhler S, Doelken S, Collier N, Oellrich A, Smedley D, et al. Automatic concept recognition using the human phenotype ontology reference and test suite corpora. *Database (Oxford).* 2015;2015:bav005. <https://doi.org/10.1093/database/bav005> PMID: 25725061
164. Weissenbacher D, Rawal S, Zhao X, Priestley JRC, Szigety KM, Schmidt SF, et al. PhenOLD, a language model normalizer of physical examinations from genetics clinical notes. *medRxiv.* 2024:2023.10.16.23296894. <https://doi.org/10.1101/2023.10.16.23296894> PMID: 37904943
165. Garg M, Saxena C, Saha S, Krishnan V, Joshi R, Mago V. CAMS: an annotated corpus for causal analysis of mental health issues in social media posts. In: Proceedings of the Thirteenth Language Resources and Evaluation Conference. Marseille, France: European Language Resources Association; 2022. p. 6387–96. Available from: <https://aclanthology.org/2022.lrec-1.686>
166. Liyanage C, Garg M, Mago V, Sohn S. Augmenting reddit posts to determine wellness dimensions impacting mental health. In: Demner-fushman D, Ananiadou S, Cohen K, editors. The 22nd Workshop on Biomedical Natural Language Processing and BioNLP Shared Tasks. Toronto, Canada: Association for Computational Linguistics; 2023. p. 306–12. Available from: <https://aclanthology.org/2023.bionlp-1.27>
167. Garg M, Shahbandegan A, Chadha A, Mago V. An annotated dataset for explainable interpersonal risk factors of mental disturbance in social media posts. In: Rogers A, Boyd-Graber J, Okazaki N, editors. Findings of the Association for Computational Linguistics: ACL 2023. Toronto, Canada: Association for Computational Linguistics; 2023. p. 11960–9. Available from: <https://aclanthology.org/2023.findings-acl.757>
168. Xie K, Gallagher RS, Shinohara RT, Xie SX, Hill CE, Conrad EC, et al. Long-term epilepsy outcome dynamics revealed by natural language processing of clinic notes. *Epilepsia.* 2023;64(7):1900–9. <https://doi.org/10.1111/epi.17633> PMID: 37114472
169. Goel A, Gueta A, Gilon O, Liu C, Erell S, Nguyen LH, et al. LLMs accelerate annotation for medical information extraction; 2023
170. Bhagat N, Mackey O, Wilcox A. Large language models for efficient medical information extraction. *AMIA Jt Summits Transl Sci Proc.* 2024;2024:509–14. PMID: 38827084



171. Cao M, Wang H, Liu X, Wu J, Zhao M. LLM collaboration PLM improves critical information extraction tasks in medical articles. In: China Health Information Processing Conference. Springer; 2023. p. 178–85.
172. Soni S, Datta S, Roberts K. quEHRy: a question answering system to query electronic health records. *J Am Med Inform Assoc*. 2023;30(6):1091–102. <https://doi.org/10.1093/jamia/ocad050> PMID: 37087111
173. Landman R, Healey SP, Loprinzo V, Kochendoerfer U, Winnier AR, Henstock PV, et al. Using large language models for safety-related table summarization in clinical study reports. *JAMIA Open*. 2024;7(2):ooae043. <https://doi.org/10.1093/jamiaopen/ooae043> PMID: 38818116
174. Mishra N, Sahu G, Calixto I, Abu-Hanna A, Laradji I. LLM aided semi-supervision for efficient extractive dialog summarization. In: Bouamor H, Pino J, Bali K, editors. Findings of the Association for Computational Linguistics: EMNLP 2023. Singapore: Association for Computational Linguistics; 2023. p. 10002–9. Available from: <https://aclanthology.org/2023.findings-emnlp.670>
175. Devaraj A, Marshall I, Wallace B, Li JJ. Paragraph-level simplification of medical texts. In: Toutanova K, Rumshisky A, Zettlemoyer L, Hakkani-Tur D, Beltagy I, Bethard S, et al., editors. Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. Online: Association for Computational Linguistics; 2021. p. 4972–84. Available from: <https://aclanthology.org/2021.naacl-main.395>
176. Artsi Y, Sorin V, Konen E, Glicksberg BS, Nadkarni G, Klang E. Large language models in simplifying radiological reports: systematic review. *medRxiv*. 2024. <https://doi.org/10.1101/2024.01.05.24300884>
177. Jeblick K, Schachtner B, Dexl J, Mittermeier A, Stüber AT, Topalis J, et al. ChatGPT makes medicine easy to swallow: an exploratory case study on simplified radiology reports. *Eur Radiol*. 2024;34(5):2817–25. <https://doi.org/10.1007/s00330-023-10213-1> PMID: 37794249
178. Swanson K, He S, Calvano J, Chen D, Televizian T, Jiang L, et al. Biomedical text readability after hypernym substitution with fine-tuned large language models. *PLOS Digit Health*. 2024;3(4):e0000489. <https://doi.org/10.1371/journal.pdig.0000489> PMID: 38625843
179. Wiest I, Leßmann M, Wolf F, Ferber D, Van Treeck M, Zhu J. Anonymizing medical documents with local, privacy preserving large language models: The LLM-Anonymizer. *medRxiv*. 2024. <https://doi.org/10.1101/2024.06.11.24308355>
180. Guo Y, Qiu W, Leroy G, Wang S, Cohen T. Retrieval augmentation of large language models for lay language generation. *J Biomed Inform*. 2024;149:104580. <https://doi.org/10.1016/j.jbi.2023.104580> PMID: 38163514
181. Joseph S, Kazanas K, Reina K, Ramanathan VJ, Xu W, Wallace BC, et al.. Multilingual simplification of medical texts; 2023. Available from: <https://arxiv.org/abs/2305.12532>
182. Saha T, Gakhreja V, Das AS, Chakraborty S, Saha S. Towards motivational and empathetic response generation in online mental health support. In: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval. Madrid Spain: ACM; 2022. p. 2650–6. Available from: <https://dl.acm.org/doi/10.1145/3477495.3531912>
183. Yang K, Ji S, Zhang T, Xie Q, Kuang Z, Ananiadou S. Towards interpretable mental health analysis with large language models. In: Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing. Singapore: ACL; 2023. p. 6056–77. Available from: <https://aclanthology.org/2023.emnlp-main.370>
184. Kang C, Cheng Y, Urbanovad K, Hu L, Zhang Y, Hu Y, et al. Domain-Specific Assistant-Instruction on Psychotherapy Chatbot. *SSRN*; 2023. Available from: <https://www.ssrn.com/abstract=4616282>
185. Cung M, Sosa B, Yang HS, McDonald MM, Matthews BG, Vluga AG, et al. The performance of artificial intelligence chatbot large language models to address skeletal biology and bone health queries. *J Bone Miner Res*. 2024;39(2):106–15. <https://doi.org/10.1093/jbmr/zjad007> PMID: 38477743
186. Nov O, Singh N, Mann D. Putting ChatGPT's medical advice to the (turing) test: survey study. *JMIR Med Educ*. 2023;9:e46939. <https://doi.org/10.2196/46939> PMID: 37428540
187. Levine DM, Tuwani R, Kompa B, Varma A, Finlayson SG, Mehrotra A, et al. The diagnostic and triage accuracy of the GPT-3 Artificial Intelligence Model. *medRxiv*. 2023;2023.01.30.23285067. <https://doi.org/10.1101/2023.01.30.23285067> PMID: 36778449
188. Kim S, Schramm S, Berberich C, Rosenkranz E, Schmitzer L, Serguen K. Human-AI collaboration in large language model-assisted brain MRI differential diagnosis: a usability study. *medRxiv*. 2024;2024–02.
189. Saha T, Chopra S, Saha S, Bhattacharyya P, Kumar P. A large-scale dataset for motivational dialogue system: an application of natural language generation to mental health. 2021 International Joint Conference on Neural Networks (IJCNN). 2021. p. 1–8.



190. Pirina I, Coltekin C. Identifying depression on Reddit: the effect of training data. *Proceedings of the 2018 EMNLP Workshop SMM4H: The 3rd Social Media Mining for Health Applications Workshop & Shared Task*. 2018. p. 9–12.
191. Zirikly A, Resnik P, Uzuner O, Hollingshead K. CLPsych 2019 Shared task: predicting the degree of suicide risk in reddit posts. In: Niederhoffer K, Hollingshead K, Resnik P, Resnik R, Loveys K, editors. *Proceedings of the Sixth Workshop on Computational Linguistics and Clinical Psychology*. Minneapolis, Minnesota: Association for Computational Linguistics; 2019. p. 24–33. Available from: <https://aclanthology.org/W19-3003>
192. Turcan E, McKeown K. Dreddit: A Reddit dataset for stress analysis in social media. In: Holderness E, Jimeno Yepes A, Lavelli A, Minard AL, Pustejovsky J, Rinaldi F, editors. *Proceedings of the Tenth International Workshop on Health Text Mining and Information Analysis (LOUHI 2019)*. Hong Kong: Association for Computational Linguistics; 2019. p. 97–107. Available from: <https://aclanthology.org/D19-6213>
193. Zhou S, Ding S, Wang J, Lin M, Melton G, Zhang R. Interpretable differential diagnosis with dual-inference large language models. *arXiv preprint*. 2024. <https://arxiv.org/abs/2407.07330>
194. Singh A, RG E, et al. RAG-based medical assistant especially for infectious diseases. In: 2024 International Conference on Inventive Computation Technologies (ICICT). 2024. p. 1128–33.
195. Jeong M, Sohn J, Sung M, Kang J. Improving medical reasoning through retrieval and self-reflection with retrieval-augmented large language models. *Bioinformatics*. 2024;40(Suppl 1):i119–29. <https://doi.org/10.1093/bioinformatics/btae238> PMID: 38940167
196. Brake N, Schaaf T. Comparing two model designs for clinical note generation; is an LLM a useful evaluator of consistency? 2024. Available from: <https://arxiv.org/abs/2404.06503>
197. Phang J, Zhao Y, Liu PJ. Investigating efficiently extending transformers for long input summarization. In: *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*; 2023. p. 3946–61
198. Lin C. Rouge: A package for automatic evaluation of summaries. *Text summarization branches out*. 2004. p. 74–81.
199. Glover J, Fancellu F, Jagannathan V, Gormley MR, Schaaf T. Revisiting text decomposition methods for NLI-based factuality scoring of summaries. In: Bosselut A, Chandu K, Dhole K, Gangal V, Gehrmann S, Jernite Y, et al., editors. *Proceedings of the 2nd Workshop on Natural Language Generation, Evaluation, and Metrics (GEM)*. Abu Dhabi, United Arab Emirates (Hybrid): Association for Computational Linguistics; 2022. p. 97–105. Available from: <https://aclanthology.org/2022.gem-1.7>
200. Liu Z, Lin Y, Cao Y, Hu H, Wei Y, Zhang Z, et al.. Swin transformer: hierarchical vision transformer using shifted windows; 2021. Available from: <https://arxiv.org/abs/2103.14030>
201. Dosovitskiy A, Beyer L, Kolesnikov A, Weissenborn D, Zhai X, Unterthiner T, et al.. An image is worth 16x16 words: transformers for image recognition at scale; 2021. Available from: <https://arxiv.org/abs/2010.11929>
202. Readable. Readable – content analysis and readability tools; 2024. <https://readable.com/>
203. Ali SR, Dobbs TD, Hutchings HA, Whitaker IS. Using ChatGPT to write patient clinic letters. *Lancet Digit Health*. 2023;5(4):e179–81. [https://doi.org/10.1016/S2589-7500\(23\)00048-1](https://doi.org/10.1016/S2589-7500(23)00048-1) PMID: 36894409
204. Parry D, Odedra A, Fagbohun M, Oeppen RS, Davidson M, Brennan PA. Abbreviation use decreases effective clinical communication and can compromise patient safety. *Br J Oral Maxillofac Surg*. 2023;61(8):509–13. <https://doi.org/10.1016/j.bjoms.2023.07.004> PMID: 37563053
205. Sivarajkumar S, Kelley M, Samolyk-Mazzanti A, Visweswaran S, Wang Y. An empirical evaluation of prompting strategies for large language models in zero-shot clinical natural language processing. *arXiv preprint* 2023. <https://arxiv.org/abs/2309.08008>
206. Bordt S, Nori H, Caruana R. Elephants never forget: testing language models for memorization of tabular data; 2024.
207. Ranaldi L, Ruzzetti ES, Zanzotto FM. PreCog: exploring the relation between memorization and performance in pre-trained language models; 2023.
208. Levkovich I, Elyoseph Z. Suicide risk assessments through the eyes of ChatGPT-3.5 versus ChatGPT-4: vignette study. *JMIR Ment Health*. 2023;10:e51232. <https://doi.org/10.2196/51232> PMID: 37728984
209. Singhal A, Neveditsin N, Tanveer H, Mago V. Toward fairness, accountability, transparency, and ethics in AI for social media and health care: scoping review. *JMIR Med Inform*. 2024;12:e50048. <https://doi.org/10.2196/50048> PMID: 38568737
210. Shankar V, Yousefi E, Manashty A, Blair D, Teegapuram D. Clinical-GAN: trajectory forecasting of clinical events using transformer and generative adversarial networks. *Artif Intell Med*. 2023;138:102507. <https://doi.org/10.1016/j.artmed.2023.102507> PMID: 36990584

211. Kadri F, Dairi A, Harrou F, Sun Y. Towards accurate prediction of patient length of stay at emergency department: a GAN-driven deep learning framework. *J Ambient Intell Humaniz Comput.* 2022; 14(9):11481–11495. <https://doi.org/10.1007/s12652-022-03717-z> PMID: 35132336
212. Henriksson A, Pawar Y, Hedberg P, Nauc  r P. Multimodal fine-tuning of clinical language models for predicting COVID-19 outcomes. *Artif Intell Med.* 2023;146:102695. <https://doi.org/10.1016/j.artmed.2023.102695> PMID: 38042595
213. Consortium THPO. The human phenotype ontology; 2024. Available from: <https://hpo.jax.org/>
214. Agrawal M, Hegselmann S, Lang H, Kim Y, Sontag D. Large language models are few-shot clinical information extractors. In: *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. Abu Dhabi, United Arab Emirates: Association for Computational Linguistics; 2022. p. 1998–2022. Available from: <https://aclanthology.org/2022.emnlp-main.130>
215. Ge J, Li M, Delk MB, Lai JC. A comparison of a large language model vs manual chart review for the extraction of data elements from the electronic health record. *Gastroenterology.* 2024;166(4):707–709.e3. <https://doi.org/10.1053/j.gastro.2023.12.019> PMID: 38151192
216. American College of Radiology. Liver Imaging Reporting and Data System (LI-RADS); 2024. Available from: <https://www.acr.org/Clinical-Resources/Clinical-Tools-and-Reference/Reporting-and-Data-Systems/LI-RADS>
217. Collaboration C. Cochrane database of systematic reviews; 2024. Available from: <https://www.cochranelibrary.com/cdsr/about-cdsr>
218. Phan LN, Anibal JT, Tran H, Chanana S, Bahadroglu E, Peltekian A, et al.. SciFive: a text-to-text transformer model for biomedical literature; 2021. Available from: <https://arxiv.org/abs/2106.03598>
219. Beltagy I, Lo K, Cohan A. SciBERT: a pretrained language model for scientific text; 2019. Available from: <https://arxiv.org/abs/1903.10676>
220. Birkun AA, Gautam A. Large Language Model (LLM)-powered chatbots fail to generate guideline-consistent content on resuscitation and may provide potentially harmful advice. *Prehosp Disaster Med.* 2023;38(6):757–63. <https://doi.org/10.1017/S1049023X23006568> PMID: 37927093
221. Zhang T, Kishore V, Wu F, Weinberger KQ, Artzi Y. BERTScore: evaluating text generation with BERT; 2020. Available from: <https://arxiv.org/abs/1904.09675>
222. Schwartz IS, Link KE, Daneshjou R, Cort  s-Penfield N. Black box warning: large language models and the future of infectious diseases consultation. *Clin Infect Dis.* 2024;78(4):860–6. <https://doi.org/10.1093/cid/ciad633> PMID: 37971399
223. Weizenbaum J. ELIZA—a computer program for the study of natural language communication between man and machine. *Commun ACM.* 1966;9(1):36–45. <https://doi.org/10.1145/365153.365168>
224. Yuan W, Neubig G, Liu P. BARTScore: evaluating generated text as text generation; 2021. Available from: <https://arxiv.org/abs/2106.11520>
225. Sharma A, Rushton K, Lin I, Wadden D, Lucas K, Miner A, et al. Cognitive reframing of negative thoughts through human-language model interaction. In: *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Toronto, Canada: Association for Computational Linguistics; 2023. p. 9977–10000. Available from: <https://aclanthology.org/2023.acl-long.555>
226. America MH. Mental Health America; 2025. Available from: <https://mhanational.org/>
227. Street A. Publisher of streaming video, audio, and text library databases; 2025. Available from: <https://alexanderstreet.com/>
228. Sharma A, Lin IW, Miner AS, Atkins DC, Althoff T. Human-AI collaboration enables more empathic conversations in text-based peer-to-peer mental health support; 2022.
229. Liu Z, Wu Z, Hu M, Zhao B, Zhao L, Zhang T, et al.. PharmacyGPT: the AI pharmacist; 2023.
230. Zhang C, Chen Y, D'Haro LF, Zhang Y, Friedrichs T, Lee G, et al. DynaEval: unifying turn and dialogue level evaluation. In: *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Online: Association for Computational Linguistics; 2021. p. 5676–89. Available from: <https://aclanthology.org/2021.acl-long.441>
231. Ghazarian S, Wen N, Galstyan A, Peng N. DEAM: dialogue coherence evaluation using AMR-based semantic manipulations. In: *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*; 2022. p. 771–85.
232. Abbasian M, Khatibi E, Azimi I, Oniani D, Shakeri Hossein Abad Z, Thieme A, et al. Foundation metrics for evaluating effectiveness of healthcare conversations powered by generative AI. *NPJ Digit Med.* 2024;7(1):82. <https://doi.org/10.1038/s41746-024-01074-z> PMID: 38553625

233. Devaraj A, Sheffield W, Wallace BC, Li JJ. Evaluating factuality in text simplification. *Proc Conf Assoc Comput Linguist Meet.* 2022;2022:7331–45. <https://doi.org/10.18653/v1/2022.acl-long.506> PMID: 36404800
234. Lukas N, Salem A, Sim R, Tople S, Wutschitz L, Zanella-Béguelin S. Analyzing leakage of personally identifiable information in language models. In: 2023 IEEE Symposium on Security and Privacy (SP). IEEE; 2023. p. 346–63.
235. Petroni F, Rocktaschel T, Riedel S, Lewis P, Bakhtin A, Wu Y, et al. Language models as knowledge bases? In: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*; 2019. p. 2463–73.
236. Dhamala J, Sun T, Kumar V, Krishna S, Pruksachatkun Y, Chang K. Dataset and metrics for measuring biases in open-ended language generation. *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency.* 2021. p. 862–72.
237. Parrish A, Chen A, Nangia N, Padmakumar V, Phang J, Thompson J, et al. BBQ: a hand-built bias benchmark for question answering. In: *Findings of the Association for Computational Linguistics: ACL 2022*; 2022. p. 2086–105.
238. Gehman S, Gururangan S, Sap M, Choi Y, Smith NA. RealToxicityPrompts: evaluating neural toxic degeneration in language models. In: *Findings of the Association for Computational Linguistics: EMNLP 2020*; 2020. p. 3356–69.
239. Lin S, Hilton J, Evans O. Truthful QA: measuring how models mimic human falsehoods. In: *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*; 2022. p. 3214–52.
240. de Hond A, Leeuwenberg T, Bartels R, van Buchem M, Kant I, Moons KG, et al. From text to treatment: the crucial role of validation for generative large language models in health care. *Lancet Digit Health.* 2024;6(7):e441–3. [https://doi.org/10.1016/S2589-7500\(24\)00111-0](https://doi.org/10.1016/S2589-7500(24)00111-0) PMID: 38906607
241. Protection of Human Subjects of Biomedical TNC, Research B. The Belmont Report: Ethical Principles and Guidelines for the Protection of Human Subjects of Research; 1979. Available from: <https://www.hhs.gov/ohrp/regulations-and-policy/belmont-report/index.html>
242. Solomonides AE, Koski E, Atabaki SM, Weinberg S, McGreevey JD, Kannry JL, et al. Defining AMIA's artificial intelligence principles. *J Am Med Inform Assoc.* 2022;29(4):585–91. <https://doi.org/10.1093/jamia/ocac006> PMID: 35190824
243. Ong JCL, Chang SY-H, William W, Butte AJ, Shah NH, Chew LST, et al. Ethical and regulatory challenges of large language models in medicine. *Lancet Digit Health.* 2024;6(6):e428–32. [https://doi.org/10.1016/S2589-7500\(24\)00061-X](https://doi.org/10.1016/S2589-7500(24)00061-X) PMID: 38658283
244. Zhao H, Chen H, Yang F, Liu N, Deng H, Cai H. Explainability for large language models: a survey. *ACM Trans Intell Syst Technol.* 2024;15(2). <https://doi.org/10.1145/3639372>
245. Gallegos IO, Rossi RA, Barrow J, Tanjim MM, Kim S, Derroncourt F, et al. Bias and fairness in large language models: a survey. *Computat Linguist.* 2024; p. 1–79.
246. Chu Z, Wang Z, Zhang W. Fairness in large language models: a taxonomic survey. *SIGKDD Explorations Newsletter.* 2024;26(1):34–48. <https://doi.org/10.1145/3682112.3682117>
247. Sblendorio E, Dentamaro V, Lo Cascio A, Germini F, Piredda M, Cicolini G. Integrating human expertise & automated methods for a dynamic and multi-parametric evaluation of large language models' feasibility in clinical decision-making. *Int J Med Inform.* 2024;188:105501. <https://doi.org/10.1016/j.ijmedinf.2024.105501> PMID: 38810498
248. Jie YW, Satapathy R, Goh R, Cambria E. How interpretable are reasoning explanations from prompting large language models? In: *Findings of the Association for Computational Linguistics: NAACL 2024*; 2024. p. 2148–64.
249. Liu L, Zhang D, Li S, Zhou G, Cambria E. Two heads are better than one: zero-shot cognitive reasoning via multi-LLM knowledge fusion. In: *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management.* 2024. p. 1462–72.
250. Ong K, Mao R, Satapathy R, Shiota Filho R, Cambria E, Sulaeman J. Explainable natural language processing for corporate sustainability analysis. *Information Fusion.* 2025;115:102726.
251. Mitchell M, Wu S, Zaldivar A, Barnes P, Vasserman L, Hutchinson B, et al. Model cards for model reporting. In: *Proceedings of the Conference on Fairness, Accountability, and Transparency.* 2019. p. 220–9.