

REVIEW

Open Access



Open challenges and opportunities in federated foundation models towards biomedical healthcare

Xingyu Li¹, Lu Peng^{1*}, Yu-Ping Wang² and Weihua Zhang³

*Correspondence:
lpeng3@tulane.edu

¹ Department of Computer Science, Tulane University, New Orleans, LA, USA

² Department of Biomedical Engineering, Tulane University, New Orleans, LA, USA

³ School of Computer Science, Fudan University, Shanghai, China

Abstract

This survey explores the transformative impact of foundation models (FMs) in artificial intelligence, focusing on their integration with federated learning (FL) in biomedical research. Foundation models such as ChatGPT, LLaMa, and CLIP, which are trained on vast datasets through methods including unsupervised pretraining, self-supervised learning, instructed fine-tuning, and reinforcement learning from human feedback, represent significant advancements in machine learning. These models, with their ability to generate coherent text and realistic images, are crucial for biomedical applications that require processing diverse data forms such as clinical reports, diagnostic images, and multimodal patient interactions. The incorporation of FL with these sophisticated models presents a promising strategy to harness their analytical power while safeguarding the privacy of sensitive medical data. This approach not only enhances the capabilities of FMs in medical diagnostics and personalized treatment but also addresses critical concerns about data privacy and security in healthcare. This survey reviews the current applications of FMs in federated settings, underscores the challenges, and identifies future research directions including scaling FMs, managing data diversity, and enhancing communication efficiency within FL frameworks. The objective is to encourage further research into the combined potential of FMs and FL, laying the groundwork for healthcare innovations.

Keywords: Foundation model, Federated learning, Healthcare, Biomedical, Large language model, Vision language model, Privacy, Multimodal

Introduction

Foundation models (FMs) [1, 2] have risen to prominence as pivotal elements in the field of artificial intelligence [3]. These models are distinguished by their deep learning architectures and a vast number of parameters, allowing them to excel in tasks ranging from text generation to video analysis-capabilities that surpass those of previous AI systems. FMs are developed using advanced training techniques, including unsupervised pretraining [4, 5], self-supervised training [6, 7], instructed fine-tuning [8], and reinforcement human preference feedback [9]. These methodologies equip them to generate



©The Author(s) 2025. **Open Access** This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

coherent text and realistic images with unprecedented accuracy, showcasing their transformative potential across various domains.

The potential of foundation models extends far beyond mere technical capabilities. These models mark a significant paradigm shift from how we utilize artificial intelligence for cutting-edge scientific problem-solving. As versatile tools, they can be rapidly adapted and fine-tuned for specific tasks, eliminating the need to develop new models from the ground up. This adaptability is crucial in fields where processing limited datasets to extract meaningful insights is essential. It is particularly transformative in biomedical healthcare, where the efficacy of AI must be balanced with stringent data privacy considerations [10–12]. In this domain, foundation models not only enhance our analytical capabilities but also ensure that sensitive health information is handled with the utmost integrity, thereby aligning technological advancement with ethical standards.

Federated Learning (FL) [13, 14], a method for training machine learning models across multiple decentralized devices or servers without exchanging local data samples, aligns well with the capabilities of foundation models in the biomedical healthcare sector. In this context, where data privacy and collaborative efforts are essential, FL enables the utilization of vast and varied datasets characteristic of the medical field while protecting sensitive patient information. By applying FL, foundation models can access and analyze extensive medical data without breaching privacy [15, 16], thus overcoming major obstacles in deploying AI technologies where data confidentiality is crucial. Existing applications of FL in conjunction with FMs typically involve training strategies that range from starting from scratch to prompt fine-tuning. FL enhances the application of FMs across both large language models and vision-language models, allowing for comprehensive and privacy-conscious analyses.

Integrating the privacy-preserving and decentralization features of FL with the robust, generalizable capabilities of FMs enables researchers to perform in-depth analyses using insights pooled from local datasets. This approach not only broadens the scope and accuracy of medical research but also complies with stringent data protection laws such as the General Data Protection Regulation (GDPR) [17] in Europe and the Health Insurance Portability and Accountability Act (HIPAA) [18] in the United States. The potential for healthcare is profound, facilitating more personalized medicine where treatment plans are precisely tailored to individual genetic profiles, lifestyles, and medical histories. Additionally, FMs that are pre-trained or fine-tuned via federated learning on diverse datasets can reveal new biomarkers and therapeutic targets, thereby significantly pushing the boundaries of medical research and improving patient care. The synergy between federated learning and foundation models heralds a significant leap forward in the use of medical data, driving innovation in medical technologies while rigorously protecting patient privacy.

This paper presents a comprehensive survey of the latest advancements in foundation models and federated learning within the biomedical and healthcare sectors, highlighting their implementations and addressing the persistent challenges encountered in these fields. A notable application of these technologies involves the use of federated foundation models to train pre-trained vision-language models, such as FedClip [19], which enhance both generalization and personalization in image classification tasks. Additionally, MedCLIP [20] employs vision-text contrastive learning

with 20K medical datasets to surpass current benchmarks in medical diagnostics. FedMed [21], a tailored federated learning framework, effectively counters performance degradation in federated settings, facilitating high-quality collaborative training. Another groundbreaking model, MedGPT [22], based on the GPT architecture, utilizes electronic health records to predict future medical events, offering the potential to detect early signs of critical illnesses, such as cancer or cardiovascular diseases [23], before they are typically diagnosable through conventional methods. Importantly, the utilization of federated learning ensures that sensitive patient data is processed on-site, never leaving the institution's local environment, thus significantly enhancing data security and maintaining strict patient confidentiality.

The integration of federated learning (FL) with foundation models (FMs) offers unprecedented potential to transform medical diagnostics and personalize treatments, greatly enhancing the capabilities of healthcare systems to deliver exceptional care while adhering to rigorous standards of data privacy and security. This technological advancement not only improves patient outcomes but also strengthens trust in the use of AI within critical sectors such as healthcare. However, deploying federated FMs in the biomedical domain comes with significant challenges, including ensuring data privacy and security, achieving model generalization across diverse datasets, and maintaining bias and fairness. Addressing these issues is essential for harnessing the full capabilities of FL FMs in healthcare and biomedical research.

Furthermore, this paper explores future directions and ongoing challenges in the field, emphasizing the importance of real-time learning and adaptation, fostering collaborative innovation, and the generation of synthetic data for both academic and industrial applications within FL frameworks. By overcoming these challenges, researchers and practitioners can fully realize the potential of federated foundation models, leading to revolutionary advancements in healthcare. These efforts will not only contribute to scientific progress but also to the practical, ethical, and efficient implementation of AI technologies in sensitive environments, ultimately benefiting global health outcomes.

- We provide a comprehensive review of existing literature on Federated Learning (FL) and Foundation Models (FM) within the biomedical and healthcare domains. This review meticulously categorizes and discusses various aspects such as biomedical and healthcare data sources, foundation models, federated privacy, and downstream tasks, offering a thorough synthesis of current knowledge and methodologies.
- We introduce a taxonomy of biomedical healthcare foundation models, classifying the existing representative FMs from diverse perspectives including model architecture, training strategy, and intended application purposes. This taxonomy aids in the systematic understanding and comparison of different models.
- We explore the open challenges and outline future research directions for the integration of FL with FMs in the biomedical and healthcare sectors, providing insights into unresolved issues and potential advancements.
- To the best of our knowledge, this is the first survey paper to extensively cover foundation models in federated learning specifically tailored for biomedical and

healthcare applications. Our survey uniquely addresses both large-language and vision-language models, highlighting their relevance and transformative potential in this context.

How do we collect papers? In this survey, we collect over two hundred related papers in the field of Federated Learning, Foundation Model, and Biomedical healthcare. We consider Google Scholar as our main literature search engine, where the MedPub, Web of Science, and IEEE Xplore are also used as essential tools. Moreover, we check most of the related top-tier conferences, such as NeurIPS, ICML, ICLR, CVPR, and ECCV, and Bioinformatics. The major keywords we use are “Biomedical Federated Learning, Medical Pretrained Foundation Model, Healthcare Federated Pretrain Training, etc”. The most representative papers like Med-BERT [24], FedClip [19], and MedClip [20] are regarded as seed papers for reference check.

Organization The rest of this survey is organized as follows. [Background](#) section describes the FM and FL literature relevant to our work. [Federated learning and foundation models](#) section details how to apply FMs with FL. The applications of FM on biomedical and healthcare is summarized in [Foundation models in biomedical healthcare](#) section. The challenges and future directions of Federated FMs in the biomedical and healthcare sectors are discussed in [Open challenges and opportunities in federated foundation biomedical research](#) section. Finally, we conclude our survey in [Conclusions](#) section.

Background

Background on foundation models

The latest wave of AI innovation sees the evolution of a new class of AI models often referred as foundation models (FMs) - a term popularized by the Stanford Institute for Human-Centered AI [25] which can be categorized into two model types: Large-language Model (LLM) and Vision-language Model (VLM). For example, LLMs including ChatGpt and Gpt-4 [26] from OpenAI demonstrate impressive capabilities to generate coherent text. VLMs such as DALL·E 2 [27] shows the ability to create realistic images and art from a text description. These models are trained with pretraining, self-supervised training, and reinforcement-instructed fine-tuning with broad data at immense scale and high resource costs, resulting in models with billions of parameters [25]. In this section, we will introduce the backbone of Foundation Models in [Backbone networks in foundation models](#) section, where the pre-trained large-language models and vision-language models are discussed in [Foundation on text: large language models](#) and [Foundation beyond text: vision language models](#) sections, respectively.

Backbone networks in foundation models

The significant advancements in foundation models are largely due to the evolution of their underlying architectures, transitioning from Long Short-Term Memory networks (LSTM) [28] to Transformers [29]. Initially, LSTMs served as the basic architecture for early pre-trained models, where the recurrent structure is computationally intensive

when scaled to deeper layers. In response to these limitations, the Transformer architecture was developed and quickly established itself as the standard for modern natural language processing (NLP) [30]. The superiority of Transformers over LSTMs can be attributed to two key factors: (1) Efficiency: Transformers eliminate recurrence, enabling parallel computation of tokens. (2) Effectiveness: The attention mechanism facilitates dynamic spatial interactions between tokens, contingent on the input itself. This section provides a brief overview of the evolution of backbone networks in foundation models, highlighting the transition from LSTMs to Transformers, followed by vision language model backbones from Convolutional neural networks (CNNs) [31] to Vision Transformers (ViTs) [32].

Backbone Networks in Texts. Transformer has become the backbone of most pre-trained language models, such as BERT [33], GPT[1], and T5[34], building upon self-attention module and feed-forward networks (FFNs). The self-attention module facilitates token interaction, while FFN refines token representations using non-linear transformations. The Transformer architecture is designed to process tokens efficiently in parallel, thanks to the elimination of recurrent units and the use of position embeddings. Additionally, the architecture includes residual connections, layer normalization, and other features that prevent saturation issues and enhance expressive power with large-scale data and deep layers. The input is linearly transformed into query, key, value (Q, K, V), and output spaces in the self-attention module: the attention scores between the query and key is computed, which are then used to weight the values. The FFN module processes the weighted values to generate the output. The Transformer architecture has proven to be superior in terms of capacity and scalability, enabling the development of increasingly sophisticated language models. Considering an input X the linear transformation of X into Q, K, V is computed as follows:

$$Q = XW^Q, K = XW^K, V = XW^V, \quad (1)$$

where the self-attention module is calculated with a softmax function as follows:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK}{\sqrt{d_k}}\right)V. \quad (2)$$

To this end, the FFN provides the non-linear features for the transformer architecture. Besides the self-attention and FFN modules, the transformer architecture also includes residual connections [35], layer normalization [36], and positional encoding [37] to enhance the model's performance. The transformer architecture has been widely adopted in various pre-trained language models, such as BERT, GPT, and T5, and has been instrumental in advancing the field of natural language processing (NLP).

Backbone Networks in Images. Convolutional Neural Networks (CNNs) [31] have long been the foundation for many vision-related tasks, characterized by their distinctive architecture comprising convolutional, pooling, activation, and fully connected layers. These layers work in unison where convolutional layers act as trainable filters identifying image patterns like edges and textures, pooling layers reduce data dimensionality, activation layers introduce non-linearity, and fully connected layers synthesize these features

into predictions. This architecture has been not only pivotal in vision applications but also adapted for language understanding tasks.

As the field evolves, there has been a notable shift towards incorporating Transformer architectures into vision tasks. This integration is exemplified by the development of Vision Transformers (ViT) [32], which apply the Transformer's self-attention mechanisms to image patches for feature extraction, representing a significant evolution from traditional CNN approaches. This concept has similarly impacted computational biology, as seen in models like AlphaFold2 [38], which leverages Transformer technology for protein structure prediction. These adaptations underscore the versatility and robustness of Transformer models across different scientific domains.

Foundation on text: large language models

In the field of Natural Language Processing (NLP), the evolution of methods to build token representations has been marked by significant advancements. Initially, typical approaches such as those proposed by [39, 40] focused on creating 'static word embeddings,' where a one-to-one mapping between words and their vector representations is established. These embeddings are termed 'static' because they do not account for the context in which a word is used, thus limiting their ability to reflect the diverse meanings words can have in different settings.

Recognizing the limitations of static embeddings, there has been a shift towards developing 'contextualized word embeddings.' These representations are dynamic, with the vector for a word varying according to its contextual usage. For instance, the word 'bank' would have different embeddings in 'river bank' compared to 'money bank.' This approach, exemplified by models like ELMo [41], GPT [42], and BERT [33], significantly enhances the quality of word representations by modeling bi-directional contexts, thereby improving performance across various NLP tasks.

Historically, neural language models [43, 44] served as foundational frameworks in NLP, utilizing relatively shallow neural architectures for efficient training. These models were primarily pre-trained on tasks like unidirectional language modeling, which involves predicting the next word based on previous words. However, subsequent innovations such as Skip-Gram [39] aimed to enrich word embeddings by predicting surrounding words or using bidirectional context, respectively. GloVe [40] extended this by focusing on word co-occurrence probabilities.

The advent of deep learning brought about more sophisticated approaches for learning word representations. ELMo [41] introduced a bidirectional language modeling task, utilizing both forward and backward context in its pre-training. GPT [42] continued with unidirectional modeling, while BERT [33] innovated with the Masked Language Model. This method involves masking words in a sentence and predicting them based on the remaining unmasked context, allowing for deeper bidirectional context modeling.

Further developments like the T5 model [34] introduced an encoder-decoder framework for generating text outputs, proving particularly effective in text generation tasks such as summarization and question-answering. These advancements have been integral to the development of versatile language models like OpenAI's GPT-3,

InstructGPT, Codex, and ChatGPT, which not only generate text but also engage in conversational exchanges, admit errors, and handle complex user interactions.

Representative Large Language Models Large Language Models (LLMs) have become pivotal in the evolution of natural language understanding and generation. The progression of the GPT series, from GPT-1 [42] to GPT-3 [45], and the subsequent release of GPT-4 [26], illustrates a remarkable expansion in model size and versatility. These models have profoundly impacted AI research and applications, heralding a new era of computational linguistics. Concurrently, BERT [33] revolutionized pre-training approaches by emphasizing bidirectional training, which significantly enhances language understanding capabilities. PaLM [46], another notable advancement, has achieved state-of-the-art results across diverse language tasks, highlighting the potential for scalability in LLMs. Recent innovations also include Bard [47], which integrates extensive world knowledge into a context-aware framework, and LLaMa [2], which prioritizes efficiency and practical applicability in language model design. Collectively, these models mark crucial developments in the field, each contributing distinctively to the enrichment and complexity of machine learning techniques that underpin contemporary AI systems.

Foundation beyond text: vision language models

Deep neural networks have exhibited remarkable success across a variety of vision tasks, such as image classification, object detection, and instance segmentation, largely attributable to the effectiveness of pre-training. Initially, pre-training in the vision domain involved training models on extensive annotated image datasets like ImageNet [48]. However, to address issues such as generalization errors and spurious correlations inherent in supervised learning, various self-supervised learning methods have been developed.

A significant area of advancement in AI research is the integration of vision and language models, which aims to develop systems capable of understanding and generating content that spans visual and textual modalities. The introduction of the Vision Transformer (ViT) [49] marked a pivotal shift by applying the transformer architecture—originally designed for natural language processing—directly to sequences of image patches. This approach fundamentally changed the paradigm of how models process visual information. Building on this, CLIP (Contrastive Language-Image Pre-training) [50] advanced the field by learning visual concepts through natural language supervision, enabling the model to adeptly handle various vision tasks with minimal task-specific training. Further extending these innovations, Stable Diffusion [51] ventured into generative art, providing tools to create intricate images from textual descriptions. The most recent breakthrough, Segment Anything [52], tackles image segmentation using deep learning to precisely identify and delineate multiple objects within images in ways that are contextually relevant. Collectively, these developments not only bridge the gap between visual data and language processing but also set the stage for more intuitive and interactive AI systems.

Challenges of foundation models

The paradigm of foundation models (FMs) represents a significant shift from traditional task-specific models that have long dominated the AI landscape. These pre-trained models are designed for adaptation to a variety of tasks they were not originally trained for [53]. Adaptation techniques include user or engineer prompts, continual learning, and fine-tuning-methods that expand their application to fields where data scarcity impedes the development of specialized algorithms. This flexibility introduces exciting possibilities for scalable, reusable AI models across diverse domains, including transformative potential in healthcare [54]. However, this shift also presents unique challenges, including the risk of over-generalization, the difficulty in fine-tuning for highly specialized tasks, and the ethical implications of deploying such versatile technologies in sensitive areas. These challenges necessitate rigorous validation, careful implementation, and ongoing monitoring to ensure that the deployment of foundation models aligns with ethical standards and practical requirements.

Over-trusting High Performance & Output Coherence: Ensuring Safe & Reliable Use Despite the high accuracy and broad capabilities of larger models, it's critical to address ethical and legal standards to ensure their use remains safe, fair, and privacy-conscious [53]. In healthcare, the necessity for accurate and reliable data for clinical decision-making cannot be overstated. However, verifying the correctness of outputs from FMs poses a challenge, as demonstrated by systems like ChatGPT, whose outputs can mimic human-like text, potentially leading to automation bias and misuse [55]. The complexity of these models often precludes a full understanding of their mechanisms, necessitating cautious deployment decisions, especially in sensitive fields like healthcare. This includes designing interfaces that clearly articulate the limitations and probabilistic nature of AI outputs and developing robust validation processes to ensure safety and fairness.

Building AI in a Vacuum: Decontextualized & Centralized AI development frequently takes place in isolation, focused on technological accuracy before considering real-world user needs [56]. This 'development in a vacuum' has drawn increasing scrutiny for failing to address the actual conditions and requirements of end-users [57]. Foundation models, in particular, suffer from this issue as they require significant adaptation to be truly effective outside of initial testing environments. A greater emphasis on ethnographic studies could provide deeper insights into the practical applications and challenges of AI within operational settings. Moreover, integrating AI technology into everyday use demands an understanding of specific user contexts, necessitating strategies for risk mitigation and a move towards more user-centered research directions. Validating the utility of AI in real-world settings and ensuring their integration into clinical practice remain a formidable challenge, but one that the human-computer interaction (HCI) community is well-equipped to tackle by bridging the 'last mile' of AI in healthcare [58].

Background of federated learning in foundation models

Background of conventional federated learning and frameworks

Federated Learning (FL) is a machine learning paradigm where multiple clients, such as mobile devices or entire organizations, collaboratively train a model under the orchestration of a central server, such as a service provider, while keeping the training data decentralized. This method not only adheres to the principles of focused collection and data minimization but also addresses many systemic privacy risks and costs associated with traditional centralized machine learning approaches. The concept of FL, first introduced by McMahan et al. in 2016 [13], has grown significantly in interest from both theoretical and practical perspectives. This approach is defined by challenges including unbalanced and non-IID data across numerous unreliable devices, limited communication bandwidth, and the complexities of model training and implementation across diverse and distributed environments.

Since its inception, the focus of federated learning has expanded beyond mobile and edge devices to include applications involving a small number of more reliable entities, such as multiple organizations collaborating to train a model. This has led to distinguishing between “cross-device” and “cross-silo” federated learning, each with unique challenges and requirements. In this survey, we delve into the specifics of cross-device federated learning, highlighting its practical aspects, challenges, and its potential to train and implement foundation models (FMs) in a distributed fashion.

Groundbreaking works of FL including McMahan et al. [13] laid the foundational framework for FL systems. Research in FL has since advanced, focusing on enhancing data privacy in applications such as medical image segmentation [59] and addressing ongoing challenges related to communication efficiency, scalability, and model robustness [60, 61]. Notable developments in FL include SCAFFOLD [62] and FedProx [61], which tackle issues such as client update variance and client drift in non-IID data environments. Further contributions from FedGSam [63], FedLGA [14], and LoMar [64] have advanced FL by developing generalized strategies and adaptive algorithms that enhance learning processes in federated settings.

Moreover, the development of open-source FL frameworks such as TsingTao [65], Flower [66], FedML [67], FATE [68], and FederatedScope [69] has significantly advanced the accessibility and standardization of FL practices. Designing specialized FL systems and benchmarks is imperative to meet the unique needs and challenges of foundation models (FM). Although current FL frameworks have made significant strides in both academic and industrial settings [66–71], they may not fully satisfy the specific requirements for optimizing memory, communication, and computational demands associated with FMs. Platforms like FedML [67] and FATE [68] are adapted to better support FMs, but extensive research is still needed to thoroughly explore system requirements and integration strategies for these models.

Motivations of federated learning for foundation models

Scarcity of Compliant Large-Scale Data The shortage of large-scale, high-quality, legally compliant data has become a critical driver for the adoption of federated learning

in the context of foundation models. This scarcity is particularly acute in sectors such as technology and social media, where data compliance and privacy issues are increasingly foregrounded [72–74].

High Computational Resource Demand Training large-scale foundation models demands significant computational resources. For instance, training LLaMa with 65 billion parameters required 2048 NVIDIA A100 GPUs over 21 days [2], while the smaller 1.3 billion parameter GPT-3 model needed 64 Tesla V100 GPUs for a week [45]. The development of GPT-4 also highlighted these intensive demands, utilizing substantial resources over several months at considerable financial costs [26]. Federated learning can help alleviate these demands by distributing computational tasks across multiple devices, thereby optimizing resource utilization.

Continuous Model Updating Challenges As data continually evolves, particularly from sources like IoT sensors and edge devices, keeping foundation models updated becomes a significant challenge [75, 76]. Federated learning offers a dynamic solution by enabling ongoing, incremental updates to FMs with new data, which allows these models to adapt to emerging data landscapes without the need to reinitiate training processes. This approach not only enhances the models' accuracy and relevance but also ensures their adaptability to real-world changes [77].

Reducing Response Delays and Enhancing FM Services One of the foremost benefits of applying federated learning to foundation models is the potential to deliver nearly instant responses, thus significantly improving user experience. Traditional central server deployments often face latency and privacy issues due to the required network communications between users and servers [78]. Federated learning addresses these concerns by enabling models to operate directly on local devices, minimizing network dependencies, reducing latency, and improving privacy protections. This approach not only enhances response times but also ensures a seamless, privacy-conscious interaction, maintaining user trust and satisfaction in the services provided by foundation models.

Motivations of foundation models for federated learning

Foundation Models can significantly contribute to enhancing the efficacy of Federated Learning. This section explores the motivations behind leveraging FM within FL, examines the challenges posed by this integration, and discusses the potential opportunities it offers to the field.

Data Privacy and Shortage Dilemma in FL In federated settings, clients often grapple with limited or imbalanced datasets, especially in federated few-shot learning contexts [79]. Such data scarcity can result in suboptimal model performance, as it may not fully capture the diversity of the data distribution [80]. Moreover, privacy concerns are intensified due to the potential for sensitive information recovery from model updates in FL [81, 82]. These issues are particularly acute in sectors like healthcare or finance, where data privacy regulations or the inherent sensitivity of the data restrict availability, thus

complicating the training process and limiting FL's effectiveness in these crucial areas. One promising solution is the use of synthetic data generated by FMs. Being extensively pre-trained on vast datasets and further refined through techniques such as fine-tuning and prompt engineering, FMs possess a deep understanding of complex data distributions, enabling them to produce synthetic data that closely mirrors real-world diversity.

Performance Dilemma in FL FL can mitigate issues related to non-IID and biased data by leveraging the advanced capabilities of FMs, thus enhancing performance across various tasks and domains [83]. FMs can improve FL's efficiency in several ways. (1). Starting Point Advantage: FMs provide a robust starting point for FL. Clients can begin fine-tuning directly on their local data instead of starting from scratch, leading to faster convergence and enhanced performance while reducing the need for extensive communication rounds [84, 85]. (2). Data Diversity Enhancement: FMs act as powerful generators that can synthesize diverse data, enriching the training dataset in FL. An example is GPT-FL [86], which utilizes generative models to produce synthetic data that improves downstream model training on servers. This approach not only boosts test accuracy but also enhances communication and client sampling efficiency. (3). Knowledge Distillation: FMs can address performance issues in FL by acting as knowledgeable teachers through techniques like knowledge distillation [87].

New Sharing Paradigm Empowered by FM Unlike traditional FL, which involves sharing high-dimensional model parameters, FMs use a new paradigm through prompt tuning. PROMPTFL [88] showcases how FM capabilities can be leveraged to efficiently combine global aggregation with local training on sparse data. This approach focuses on training prompts rather than the entire model, thereby optimizing resource use and enhancing performance. Building on this concept, FedPrompt [89] introduces an innovative prompt tuning method specifically designed for FL, while a recent study FedTPG [90] explores a scalable prompt generation network that learns across multiple clients, aiming to generalize to unseen classes effectively.

Machine learning in biomedical and health care

Biomedical ML: data fusion

Data is the cornerstone of sense-making in artificial intelligence (AI), playing a crucial role in various sectors, including healthcare, where they are from diverse sources like care providers, insurers, and academic publications [53, 91]. They vary in form (e.g., clinical notes, medical images), scale (e.g., patient versus population level), and style (professional versus lay language), posing both opportunities and challenges for the application and training of AI models. Despite the proficiency of machine learning methods in managing and extracting insights from vast, multi-dimensional data [92], it is vital to address how societal biases and inequalities are embedded in the data. Disparities can manifest in various aspects of healthcare, such as the prioritization of certain medical issues and the exclusion or misrepresentation of specific population groups. These issues often stem from barriers like limited healthcare access, restrictive criteria for clinical trial participation, or the risk of inaccurate data due to documentation errors and systemic discrimination [93]. For example, in California, the

mandate to verify citizenship at hospitals has reduced autism diagnosis rates among Hispanic children in the context of stringent federal immigration policies.

Healthcare and biomedicine are major sectors within the U.S. economy, accounting for about 17% of the Gross Domestic Product (GDP) [94–97]. These fields require substantial financial investments and extensive medical knowledge, encompassing everything from patient care to scientific exploration of diseases and the development of new therapies [54, 98]. We envision machine learning models as central repositories of medical knowledge, trained on a diverse array of data sources and modalities within medicine [99, 100]. These models could serve as dynamic platforms that medical professionals and researchers use to access and contribute to the latest findings, enhancing their ability to make informed decisions [101].

Biomedical Data Fusion In the field of biomedical research, a significant challenge lies in deciphering the complex interactions within and between the cellular and organismal levels, characterized by diverse components that exhibit emergent behaviors [102]. The data collected through various sensors, while rich, often provide limited insights when examined in isolation due to the specificity of each measurement modality [103]. Data fusion, the process of integrating data from multiple views, aims to provide a holistic view of biological phenomena by combining disparate data sources that offer unique perspectives on the same subject [104]. This approach is generally advantageous in several ways, categorized into complementary, redundant, and cooperative features of the data [105]. These features are not mutually exclusive but interact synergistically, enhancing the robustness and accuracy of the insights gained.

Data fusion requires the use of sophisticated machine learning (ML) methods capable of integrating both structured and unstructured data while accommodating their varied statistical properties, sources of non-biological variation, high dimensionality, and distinct patterns of missing values [103, 106]. A comprehensive examination of these strategies is presented in a review of multimodal deep learning approaches with potential advancements and methodologies in the medical field [107].

Categories Summary of Data Fusion The categories of data fusion techniques can be broadly summarized into three main approaches: easy fusion, intermediate fusion, and late fusion. Easy fusion typically involves direct modeling techniques where different types of neural networks are used to process the input data. This includes fully connected networks for a straightforward integration of features across modalities [108], convolutional networks that are effective in handling spatial data [109], and recurrent networks suited for sequential data integration [110]. Autoencoders also play a significant role in easy fusion, with variations such as regular [111], denoising [112], stacked [113], and variational autoencoders [114] being employed to refine the fusion process.

Intermediate fusion, on the other hand, involves branching strategies that can be homogeneous, focusing on either marginal [115] or joint representations [116], or heterogeneous, which also targets both marginal [117] and joint data representations [118]. These strategies optimize the integration by selectively focusing on how data from the same or different modalities are fused.

Late fusion utilizes aggregation methods to combine features at a higher level, often after initial independent processing. Techniques in this category include simple averaging [119] and weighted averaging [120], where weights might be assigned based on the reliability or importance of each modality. Furthermore, meta-learning approaches are utilized to dynamically adjust these weights for optimal performance [121], thus enhancing the fusion's effectiveness by incorporating learning-based adjustments. These methods ensure that the final model output maximally benefits from the diverse characteristics of all data modalities involved.

FM in biomedical healthcare

Motivations Foundation models hold transformative potential for biomedical research, particularly in the realms of drug discovery and disease understanding, thereby enhancing healthcare solutions [122]. Biomedical discovery processes are currently characterized by intensive demands on human resources, lengthy experimental timelines, and substantial financial outlays. For example, the drug development journey includes stages from basic research, such as protein target identification and potent molecule discovery, through clinical development involving clinical trials, to the final drug approval stage. This extensive process typically spans more than a decade and incurs costs often exceeding one billion dollars [123]. Thus, the ability to expedite biomedical discovery by harnessing existing data and published findings becomes crucial, especially during critical times like the COVID-19 outbreak, which resulted in significant loss of life and economic damage [124].

Foundation models contribute to biomedical advancements in two primary ways. Firstly, these models exhibit strong generative capabilities, as seen with coherent text generation in models such as GPT-3. These capabilities can be utilized for generating experimental protocols in clinical trials and in designing novel molecules for drug discovery [125, 126]. Secondly, foundation models excel at integrating diverse data modalities in medicine, facilitating the exploration of biomedical concepts across various scales—from molecular to patient and population levels—and integrating multiple knowledge sources, including imaging, textual, and chemical data [127–132]. This integrated approach enables discoveries that might be challenging with single-modality data alone.

Additionally, foundation models are adept at transferring knowledge across different data modalities. For instance, research by Lu et al. [133] demonstrated how a transformer model, initially trained on natural language, a data-rich modality, could be adapted for other sequence-based tasks, such as protein folding predictions, a long-standing challenge in biomedicine. These capabilities highlight the potential applications of foundation models in addressing complex biomedical tasks.

Applications Foundation models (FMs) hold significant potential for revolutionizing healthcare applications through their adaptability and efficiency in performing specific healthcare and biomedical tasks. They have been proposed for use in a variety of areas including disease prediction [24], triage or discharge recommendations [98], and health administration tasks such as clinical notes summarization [134] and medical text

simplification [1]. These applications leverage the unique capabilities of FMs, such as fine-tuning and prompting [45], to tailor solutions to specific needs, enhancing both the accuracy and efficiency of medical services.

FMs are particularly effective in patient-facing roles, such as question-answering systems and clinical trial matching applications, benefiting both researchers and patients by simplifying access to information and streamlining patient recruitment processes [126, 135–137]. As central interfaces, FMs facilitate interactions among data, tasks, and individuals, improving the operational efficiency of healthcare services. This is further explored in subsequent sections focusing on specific healthcare and biomedical tasks.

Additionally, FMs serve as repositories of extensive medical knowledge, accessible by healthcare professionals and the public for purposes like medical question-answering and interactive chatbot applications. Innovations such as ChatGPT [86] and Bard [138] provide conversational user interfaces that assist users in navigating complex health information and obtaining relevant health advice.

The implementation of FMs also promises to accelerate healthcare application development and research. These models can automate processes such as structured dataset generation, data labeling, and synthetic data creation [139]. Looking forward, there is considerable scope for developing new FM-enabled capabilities, particularly through the use of multimodal data, a feature characteristic of the healthcare domain. Beyond natural language processing, breakthroughs are already evident in areas like biomedical research, where tools like AlphaFold [140] have made significant advances in predicting human protein structures to aid drug development. Similarly, innovations in genome sequencing are hastening the detection of disease-causing genetic variants [141], and new methods are being developed for optimized clinical trial design [142]. This multidisciplinary integration highlights the transformative potential of FMs in enhancing and expanding the capabilities of the healthcare sector.

Federated learning and foundation models

Federated learning (FL) and foundation models represent two cutting-edge approaches in the field of machine learning. Federated learning offers a decentralized approach to training models across multiple nodes or devices, ensuring privacy and maintaining data locality. In contrast, foundation models, due to their vast size and generalized pre-training, provide significant adaptability and scalability for a variety of tasks. Integrating these technologies poses unique challenges but also opens up exciting opportunities for innovation in AI training and application. This section explores the relationship between federated learning and foundation models, highlighting key research directions and recent advancements in this domain. Depending on the training paradigm, foundation models can either be trained from scratch or fine-tuned on top of pre-trained models within a federated learning framework. Additionally, the application of foundation models in federated learning can extend to large language models or vision language models.

Related surveys further enrich our understanding of this integration. A recent survey by Yu et al. [83] discusses the intersection of foundation models with federated learning, exploring the motivations behind their integration, the challenges faced, and future directions for research in this domain. This survey serves as a crucial resource for

comprehending the current landscape and the potential of combining federated learning strategies with the robust capabilities of foundation models. Another pivotal work by Zhuang et al. [143] provides an in-depth analysis of how foundational models can be effectively adapted and optimized within a federated learning framework, discussing both the technical hurdles and the potential breakthroughs. Additionally, Kairouz et al. [60] offers a comprehensive overview of the advancements and persistent challenges in federated learning, highlighting issues such as algorithm efficiency, data heterogeneity, and security concerns. These surveys collectively offer a rich tapestry of insights into the evolving field of federated learning and foundation models, emphasizing their complexities and transformative potential.

Federated learning and foundation models

The integration of pre-training techniques within federated learning (FL) setups, especially for large-scale models, is increasingly viewed as essential for boosting model performance and broadening their applicability. Extensive research underscores the importance of pre-training in preparing large models to face the unique challenges presented by the decentralized nature of federated datasets. Chen et al. highlight the vital role of pre-training in readying large models for these challenges, emphasizing its necessity for effective performance within federated learning frameworks [84]. In a similar vein, Nguyen et al. investigate how the initial conditions of model training, such as the starting points of pre-training and model initialization, critically influence the effectiveness and convergence of FL models [144]. These studies stress the importance of meticulous pre-training phases to ensure that large models are fully equipped to navigate the complexities of federated learning, thereby maximizing their performance and utility in diverse applications. The application of federated learning to foundation models not only addresses these preparatory needs but also leverages the inherent strengths of both paradigms to offer several key advantages: (1). **Efficient Distributed Learning:** Federated learning enables models to learn from data distributed across multiple devices or servers without needing to centralize the data, thus preserving privacy and reducing data movement costs. (2). **Parameter-efficient Training:** By utilizing techniques such as model compression and prompt tuning within a federated framework, the training process becomes more parameter-efficient. This is particularly beneficial in environments where computational resources are limited. (3). **Prompt Tuning:** This method involves fine-tuning a model on a specific task by adjusting a small set of parameters, and when combined with federated learning, it allows for personalized model tuning on decentralized data. (4). **Model Compression:** Techniques like quantization and pruning that reduce the model size can be effectively applied in federated settings, enhancing the feasibility of deploying large models on edge devices with limited storage and processing capabilities.

Efficient Distributed Learning Algorithms Efficient distributed learning algorithms are critical for optimizing foundation models within the constraints of limited resources [89, 145]. These algorithms are specifically engineered to address the twin challenges of enhancing communication and computation efficiency during the training and deployment of large FMs across a network of devices, which may vary in capabilities and

network conditions. Two pivotal techniques in this regard are model parallelism and pipeline parallelism.

Model parallelism [146] involves dividing the model into different segments and distributing these segments across multiple devices. This allows for simultaneous processing and can significantly expedite the computation process by leveraging the combined power of multiple devices. On the other hand, pipeline parallelism [147] focuses on enhancing the overall system’s efficiency and scalability by organizing the computation process in stages. Each stage can be processed on different devices in a pipeline manner, thus optimizing the workflow and reducing idle times.

An illustrative example of these parallelism strategies in federated learning (FL) for FMs is demonstrated in Fig. 1, where participants train distinct layers of a model using their own private, local data. Note that Fig. 1a is the illustration of model parallelism and Fig. 1b demonstrates the pipeline parallelism. This approach not only maintains the privacy of the data but also contributes to the efficiency of the learning process. Recent studies, such as the research conducted by Yuan et al., validate the practicality of utilizing pipeline parallelism for decentralized FM training across heterogeneous devices [148].

Parameter-efficient Training Methods Parameter-efficient training methods are increasingly critical in optimizing foundation models for specific domains or tasks. These methods typically involve integrating adapters—a technique where the core parameters of the FM are frozen, and only a small, task-specific section of the model is fine-tuned. This approach is illustrated in Fig. 1c, which shows how adapters can be effectively incorporated into the federated learning framework for FMs. Recent implementations such as FedCLIP [50] and FFM [83] utilize this method to fine-tune FMs, achieving substantial performance improvements.

By focusing adjustments on small adapters rather than the entire model, these training methods greatly reduce the computational and communication demands typically required [149, 150]. This is particularly beneficial in FL environments where conserving bandwidth and processing power is crucial due to the distributed nature of the data and

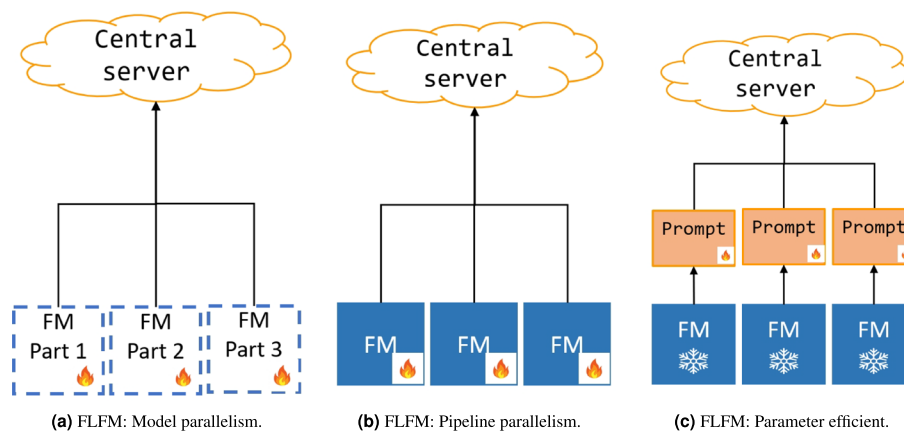


Fig. 1 Efficient distributed learning and parameter efficient strategies for foundation models in federated learning

the varying capacities of participating devices. However, despite these efficiencies, the underlying requirement for substantial computational resources to manage the FM and execute the fine-tuning process remains significant.

Prompt Tuning Prompt tuning has rapidly gained traction as a communication-efficient alternative to full model tuning, demonstrating effectiveness comparable to more resource-intensive methods [151]. This technique involves fine-tuning lightweight, additional tokens while keeping the foundational model's main parameters frozen, which avoids the necessity of sharing large model parameters across the network. In federated learning scenarios, this approach enables leveraging the collective knowledge from multiple participants to refine the prompts used in FM training, potentially enhancing the performance of the FM.

The integration of prompt tuning in FL, similar to the parameter efficient approach depicted in Fig. 1c, has been explored in recent research. Studies such as FedPrompt [152] and PROMPTFL [88] have shown promising results by improving the quality and effectiveness of prompt-based training methods through FL frameworks. These methods enable efficient and targeted tuning of model behaviors without requiring extensive data transfer or the deployment of large-scale models on each participant's device, thereby conserving bandwidth and computational resources.

Moreover, a recent study, FedTPG [90], investigates a scalable prompt generation network that learns across multiple clients, aiming to generalize effectively to unseen classes. This approach demonstrates the potential of FL to enhance the sophistication of prompt tuning methodologies by distributing the learning process across a wide array of devices and data sources.

However, the implementation of prompt tuning in FL is not without challenges. Concerns include the assumptions that large FMs are readily available on user devices, which may not always be feasible in resource-constrained environments. Additionally, there are potential privacy risks associated with utilizing cloud-based FM APIs, which could compromise the security of sensitive data.

Model Compression Model compression has emerged as a vital strategy to mitigate the substantial memory, communication, and computational demands of large foundation models. By minimizing the size of these models, model compression enables more practical deployments within federated learning frameworks without significantly compromising performance. Prominent compression techniques include knowledge distillation, where a smaller model is trained to emulate the performance of a larger one [153], and quantization, which reduces the numerical precision of model parameters to decrease both size and computational complexity [154]. Additionally, pruning eliminates superfluous or redundant model parameters, significantly lowering the resource requirements of the model [155].

Implementing these compression techniques effectively requires striking a balance between reducing model size and preserving essential capabilities. This balance ensures that the compressed model performs robustly in real-world applications, maintaining the functionality of the foundation model while reducing operational demands. Therefore, research and development in model compression focus not only on shrinking

model dimensions but also on enhancing efficiency and intelligence, tailored for specific deployment scenarios. [156] introduces ResFed, a framework that leverages model compression in federated learning to significantly cut down on bandwidth and storage needs while maintaining high model accuracy. [153] presents the concept of knowledge distillation, which allows a compact “student” model to learn effectively from a larger “teacher” model, thus enabling the student to achieve similar performance with much lower computational costs. [154] explores quantization techniques for training neural networks that perform inference using only integer arithmetic, substantially lightening model load without sacrificing accuracy. [155] provides a thorough review of neural network pruning techniques, showing their potential to significantly reduce model size while maintaining or improving performance. [61] discusses the integration of model compression into federated learning, tackling challenges related to efficiency and scalability in privacy-preserving, decentralized machine learning.

FL on large language models and vision language models

In this section, we delve into the integration of federated learning with foundation models on the two main model applications: large language models and vision language models, exploring the unique challenges and opportunities presented by these advanced AI systems.

FL on large language models

Federated learning applied to large language models (LLMs) represents a transformative approach to harnessing decentralized datasets for model training, while prioritizing data privacy and security [21, 157]. This method is especially crucial for LLMs because of their inherent requirement for vast and varied data inputs to accurately capture and interpret the complexities of human language.

This nature of FL effectively addresses privacy concerns by ensuring that sensitive or proprietary data does not leave its original location, thereby reducing the risk of data breaches. Additionally, this decentralized approach allows LLMs to learn from a wider array of linguistic inputs, reflecting regional dialects, colloquialisms, and cultural nuances that might not be present in a centralized dataset [158].

Moreover, the application of FL to LLMs facilitates the development of models that are not only linguistically comprehensive but also more personalized and responsive to local contexts. By training on diverse datasets that are geographically dispersed, LLMs can develop a deeper understanding of language variations and user-specific preferences, leading to improved performance in tasks such as language translation, sentiment analysis [159], and contextual understanding.

This method also helps in mitigating biases that are often present in centralized training datasets. Since FL involves multiple datasets that are not centrally collected, the resulting model is trained on a broader spectrum of data sources, which can contribute to more balanced and fair outputs. Thus, federated learning not only enhances the privacy and security of data used in training LLMs but also boosts the models’ ability to decipher and utilize the full richness of human language, making them more accurate and effective in real-world applications.

Practical Applications of FL on LLMs The integration of federated learning with large language models is yielding groundbreaking frameworks and methodologies that significantly enhance language model training while adhering to data privacy and security protocols. In this survey, we highlight several notable applications and advancements in this domain, including:

- *Privacy-preserving Federated Learning and its application to natural language processing*: [157] explores privacy-preserving techniques in federated learning for training large language models. It particularly focuses on models such as BERT and GPT-3, providing insights into how federated learning can be leveraged to maintain privacy without sacrificing the performance of language models in NLP applications.
- *FedMed: A federated learning framework for language modeling*: [21] introduces “FedMed”, a novel federated learning framework designed specifically for enhancing language modeling. The framework addresses the challenge of performance degradation commonly encountered in federated settings and showcases effective strategies for collaborative training without compromising on model quality.
- *Efficient Federated Learning with Pre-Trained Large Language Model Using Several Adapter Mechanisms*: [160] highlights a method to enhance federated learning efficiency by integrating adapter mechanisms into pre-trained large language models. The study emphasizes the benefits of using smaller transformer-based models to alleviate the extensive computational demands typically associated with training large models in a federated setting. The approach not only preserves data privacy but also improves learning efficiency and adaptation to new tasks.
- *OpenFedLLM*: This contribution is a seminal effort in federated learning specifically designed for large language models. The “OpenFedLLM” framework facilitates the federated training of language models across diverse and geographically distributed datasets. A standout feature of this framework is its capability to ensure data privacy during collaborative model training. It also incorporates federated value alignment, a novel approach that promotes the alignment of model outputs with human ethical standards, ensuring that the trained models adhere to desirable ethical behaviors [161]. Moreover, OpenFedLLM is open-source¹, making it accessible to the broader research community and fostering collaboration in the development of federated language models.
- *Pretrained Models for Multilingual Federated Learning*: This study addresses the complex challenges of utilizing pretrained language models within a federated learning context across multiple languages. Weller et al.’s work is crucial for understanding how multilingualism impacts federated learning algorithms, particularly exploring the effects of non-IID (independently and identically distributed) data inherent in natural language processing tasks across different languages. The research explores three main tasks: language modeling, machine translation, and text classification, providing valuable insights into the adaptability of federated learning to diverse linguistic datasets [162].

¹ <https://github.com/rui-ye/OpenFedLLM>

- *GPT-fl*: This innovative approach integrates federated learning with prompt-based techniques to train large language models. “GPT-fl” employs prompt learning within a federated framework, which allows for efficient learning from decentralized data sources while maintaining data privacy. This method enhances model adaptability and performance across various linguistic tasks, making it a promising solution for applications requiring high levels of customization and responsiveness to user-specific contexts [86].

In summary, LLMs in FL focus on balancing privacy preservation with maintaining high performance in NLP applications. Models like Privacy-preserving Federated Learning, e.g., FedMed, explore strategies to mitigate performance degradation in federated settings and enhance training efficiency using techniques such as adapter mechanisms. These approaches are particularly adept at managing the significant computational overhead associated with LLMs while ensuring that sensitive data remains secure within its local environment in FL. OpenFedLLM introduces an open-source framework that emphasizes ethical alignment in model outputs, advocating for responsible AI practices that adhere to human ethical standards, which can be crucial as it addresses the growing concern over AI alignment with societal values. Meanwhile, research on Pretrained Models for Multilingual Federated Learning tackles the challenges of multilingualism and non-IID data in federated learning, offering insights into effectively managing diverse linguistic data and enhancing the robustness of language models across different languages. GPT-fl combines prompt-based techniques within a federated framework, improving model adaptability and customization across linguistic tasks, which allows for personalized and contextually relevant responses that are essential in dynamic real-world applications.

FL on vision language models

The integration of federated learning with vision language models (VLMs) marks a significant advancement in multimodal learning where both visual and textual data are processed in a privacy-preserving, distributed learning environment. These models are crucial for tasks that necessitate a deep understanding and generation of information from visual cues and textual descriptions. Federated learning enhances the capability of VLMs by enabling them to learn from a diverse set of decentralized data sources, including images and associated annotations from various geographic and demographic distributions without the need to centralize sensitive data.

VLMs integrated with FL are particularly beneficial in scenarios where data privacy is paramount, such as in healthcare for patient image data or in surveillance where personal data protection is critical. By processing data locally and only sharing model updates, FL preserves the privacy and security of the underlying data, while still benefiting from the diverse data attributes necessary for robust model training.

This approach also allows for the training of more personalized and region-specific models, capturing a wide array of cultural and contextual nuances in visual-textual datasets. For example, a VLM trained via federated learning can better understand and generate language descriptions for regional landmarks or culturally specific events, enhancing its applicability across different global contexts.

Moreover, the decentralized nature of FL helps in mitigating dataset bias, a common issue in centralized training datasets. Since the training data in FL comes from a wide range of sources, the models are less likely to overfit to the biases present in a single dataset, leading to more generalizable and fair VLMs.

This section underscores the crucial role of prompt learning in expanding the capabilities of both language and vision models trained in federated environments. By facilitating efficient task adaptation and maintaining data privacy, prompt learning represents a significant step forward in the development of AI systems that can operate across diverse and distributed data landscapes.

Practical Applications of FL on VLMs

- *FedCLIP*: Pioneering the field of federated vision-language models, FedCLIP [19] adapts the powerful CLIP (Contrastive Language-Image Pre-training) architecture [50] to operate in a federated setting. Unlike traditional learning models that centralize data, FedCLIP enables collaborative learning across decentralized image datasets with accompanying text descriptions. Crucially, this approach safeguards data privacy by eliminating the need for sensitive user data to leave local devices.
- *PromptFL*: [88] demonstrates the power of combining federated learning with prompt learning techniques for training models on distributed visual and textual data. Prompt learning injects flexibility into model training. In PromptFL, federated learning preserves privacy while prompt learning improves training effectiveness and efficiency across diverse datasets.
- *FedPrompt*: Communication-Efficient and Privacy-Preserving Prompt Tuning in Federated Learning [152] addresses two critical aspects of federated learning for vision-language models: efficiency and privacy. Prompt tuning offers adaptability but can be communication-intensive. FedPrompt explores methods to reduce communication overhead while still reaping the benefits of prompt tuning, all while ensuring that sensitive data remains protected.
- *pFedPrompt*: [163] addresses personalization challenges in federated vision-language models. “Personalized Prompt for Vision-Language Models in Federated Learning” investigates how to learn personalized prompts. These prompts are tailored to individual clients or datasets within the federated system. The aim is to unlock performance gains by having the model adapt its behavior for specific local data distributions.
- *FedTPG*: (Text-driven Prompt Generation for Vision-Language Models in Federated Learning) [90] aims to enhance prompt generation techniques in a federated context. It introduces the idea of learning a prompt generator network which can produce context-aware prompts that guide vision-language models to tackle a variety of tasks. This has potential benefits for scenarios where a model must adapt to new classes or data it hasn't encountered previously, aligning well with the distributed nature of federated learning.
- *FedMM*: In computational pathology, fusing information from multiple modalities can significantly improve diagnostic accuracy. However, centralized training approaches raise privacy concerns due to the sensitive nature of medical images.

FedMM introduces a federated framework designed specifically to handle multi-modal data in this context. The key idea of [164] is to train individual feature extractors for each modality in a federated manner. Because only these learned feature extractors are shared, raw image data remains protected within each institution. FedMM can accommodate the situation where different institutions or hospitals may have different sets of available modalities. It enables collaborative learning even with this data heterogeneity. Subsequent tasks like classification can be performed locally using the features extracted by the federated models.

- *FedDAT*: Foundation models offer impressive performance across many tasks but often require substantial amounts of data for finetuning. FedDAT [165] addresses the challenge of finetuning these models in a federated context where the goal is to protect data privacy. To handle heterogeneity, FedDAT leverages a Dual-Adapter Teacher technique to regularize how model updates are made on each client. Furthermore, it employs Mutual Knowledge Distillation to facilitate efficient knowledge transfer across clients in the federated system.
- *CLIP2FL*: Real-world data is often messy, and client devices in a federated system might have data with different characteristics or class imbalances. CLIP2FL [166] tackles this by using a pre-trained CLIP model as guidance. On the client-side, CLIP is used for knowledge distillation to improve the local feature representations. On the server-side, CLIP is employed to generate features which help retrain the server's classifier, mitigating the negative impact of the long-tailed data problem.
- *FedAPT*: [167] introduces FedAPT, a novel method for collaborative learning in federated settings where data resides on multiple clients with varying domains (e.g., different image styles or categories). FedAPT aims to improve model generalization across these domains while maintaining data privacy. The key innovation lies in adaptive prompt tuning within the federated learning framework. Instead of directly sharing raw data, FedAPT trains a meta-prompt and adaptive network to personalize text prompts for each specific test sample. This allows the model to better adjust to domain-specific characteristics.
- *General Commerce Intelligence*: [168] discusses the development of a novel NLP-based engine designed for commerce applications. This engine leverages federated learning to provide personalized services while ensuring privacy preservation across multiple merchants. The authors focus on creating a “glocally” (globally and locally) optimized system that balances global optimization needs with local data privacy requirements.

In summary, federated learning models for vision-language tasks have demonstrated significant innovation in enhancing privacy and efficiency while leveraging the synergy between textual and visual data. Models like FedCLIP, PromptFL, and FedPrompt focus on privacy and decentralization, with trade-offs in model accuracy and communication overhead. Personalization and adaptability are key in pFedPrompt and FedAPT, which aim to tailor learning to local datasets but introduce complexity in managing local and global optimization. FedTPG and CLIP2FL enhance adaptability to new tasks and data variability, although at the cost of increased computational demands. FedMM and FedDAT tackle challenges in multi-modal and heterogeneous data integration, crucial for

applications like medical diagnostics. Lastly, General Commerce Intelligence optimizes federated learning for commercial applications, balancing local privacy with global optimization needs.

Framework for federated foundation models in biomedical

Conceptual framework

We simulate a hierarchical multi-tier architecture for the integration of FL and FMs to handle biomedical challenges:

- *Data Centers*: Each compute node in FL hosts its private biomedical datasets, which is stored locally and only communicated with the server to ensure privacy.
- *Model Host*: Foundation Models, pretrained on large-scale, public datasets, serve as the backbone on the server, which can be fine-tuned with the distributed data for the targeted biomedical challenges. Note that the large-scale model can be trained or transferred with model distillation or finetuning methods like Parameter-Efficient Fine-Tuning [169].
- *Aggregation*: FL algorithms like FedAvg [13], FedProx [61] are performed to aggregate updates from the nodes while addressing data heterogeneity and fairness concerns.
- *Feedback*: Explainable metrics and evaluation pipelines will be introduced post-aggregation to make the system robust and trustworthy.

Algorithm

In the following algorithm, we introduce the algorithm of Federated Foundation Model in biomedical:

Algorithm 1 Federated Foundation Model in Biomedical

```

1: Input: Pretrained Foundation Model parameters  $\theta_0$ , local datasets  $\{D_c\}$  for each client  $c$ , number of training rounds  $T$ , client selection strategy.
2: Output: Fine-tuned global model  $\theta_T$ .
3: Initialize global model parameters  $\theta_0$ .
4: for each round  $t = 1, 2, \dots, T$  do
5:   Select a subset of clients  $C_t$  based on client selection strategy.
6:   for each client  $c \in C_t$  in parallel do
7:     Client  $c$  downloads distilled global model parameters  $\hat{\theta}_t$  from the server.
8:     Fine-tune  $\hat{\theta}_t$  on local dataset  $D_c$  using task-specific loss  $\mathcal{L}_{\text{task}}$ .
9:     Perform fairness-aware updates and local explainability checks.
10:    Client  $c$  sends updated parameters  $\hat{\theta}_t^c$  to the server.
11:   end for
12:   Aggregate updated parameters on the server:
        $\theta_{t+1} = \text{Aggregate}(\{\hat{\theta}_t^c \mid c \in C_t\})$ .
13:   Evaluate the global model on the developed explainable metrics.
14: end for
15: return Final global model parameters  $\theta_T$ .

```

Practical applications

Training foundation models within a federated learning framework presents distinct challenges, particularly due to the disparate nature of data sources and the varied

computational resources across participating devices. The overarching goal is to cultivate effective and inclusive training strategies that can efficiently manage device heterogeneity and ensure data privacy, all while maintaining high model performance.

Training foundation models from scratch within a federated learning context is an ambitious endeavor that involves complex coordination and robust algorithmic strategies. Unlike traditional centralized training environments, federated learning necessitates handling data that remains on local devices, preventing the direct sharing of raw data. This scenario demands sophisticated techniques to efficiently aggregate learning from disparate data sources, which are often uneven in size and diversity. The primary challenge lies in ensuring that the model learns effectively from each node without requiring extensive computational resources or compromising the integrity and privacy of the data. To overcome these hurdles, training strategies must be carefully designed to optimize the learning process across the network, allowing for both model convergence and performance retention. Such strategies often involve advanced algorithms for secure multi-party computation, differential privacy, or decentralized optimization methods. By training foundation models from scratch in this way, the federated approach not only safeguards data privacy but also harnesses the unique insights embedded in local data distributions, leading to more robust and generalizable models.

Prompt learning is emerging as a pivotal approach in both natural language processing and computer vision fields, enabling models to adapt to new tasks with minimal changes to their architecture or weights. This section explores the integration of prompt learning with federated learning (FL) across different domains, highlighting recent advancements and unique applications.

Furthermore, beyond merely fostering participation, it is crucial to consider how profits and costs associated with deploying FMs via APIs are distributed. Ensuring a fair allocation of rewards and benefits is imperative to maintain trust and promote sustained cooperation among stakeholders. Mechanisms need to be established to define the distribution of profits derived from the use of FMs, guaranteeing a fair share of economic benefits. This equitable distribution is essential not only for fostering a sense of fairness but also for encouraging continued participation and investment in the FL ecosystem for FMs.

The concept of Federated Foundation Models is at the forefront of federated learning, enabling the training of large-scale models across distributed networks. This methodology is particularly effective in dealing with the challenges related to synchronizing and updating model parameters in environments where data quality and quantity are inconsistent across nodes. It ensures that learning is continuous and effective, even when network conditions and data availability vary significantly [83].

Additionally, the work titled “Heterogeneous Ensemble Knowledge Transfer for Training Large Models in Federated Learning” by Cho et al. [170] explores innovative techniques for transferring knowledge in federated settings. This study is crucial for the development of robust models capable of performing well across diverse network conditions. By facilitating knowledge transfer, this approach allows for the aggregation of insights from different data distributions and device capabilities, which is essential for building comprehensive and resilient models.

Furthermore, “No One Left Behind: Inclusive Federated Learning over Heterogeneous Devices” by Liu et al. [171] focuses on creating federated learning algorithms that integrate every participating device, regardless of its computational capabilities or the quality of the data it holds. This inclusivity ensures that every device contributes to and benefits from the collaborative learning process, thus maximizing the utilization of available data and enhancing the overall performance of the model. This approach is fundamental to achieving equity in model training and ensuring that the advantages of sophisticated model learning are universally accessible.

These studies provide a foundation for further research into strategies that enhance the combination of foundation models in federated learning frameworks, prioritizing inclusivity and efficiency.

Real-world applications of federated fm in healthcare

Recently, the integration of Federated Learning with Foundation Models has begun to demonstrate transformative potential in real-world healthcare applications. In this part, we highlight specific case studies and practical examples where these technologies have been successfully deployed.

- *Predicting Parkinson’s Disease Progression*: [172] applies FL to train explainable AI (XAI) [173] models for predicting the progression of Parkinson’s disease. By collaborating across multiple medical centers without sharing raw patient data, they developed models that maintained patient privacy while achieving high predictive accuracy.
- *Mammography Analysis*: [174] focuses on using FL for mammography analysis, enabling different healthcare providers to collaboratively train deep learning models without centralizing sensitive patient data.
- *Intensive Care Unit (ICU) Mortality Prediction*: FLICU [175] framework utilizes FL to predict ICU mortality rates. By training models on decentralized data from multiple ICUs, the study demonstrated that FL could achieve performance comparable to centralized models.

Foundation models in biomedical healthcare

This section is dedicated to a comprehensive exploration of the application of foundation models within the biomedical healthcare domain, focusing on both language and vision-language models. It will delve into the benchmarks and setups employed to evaluate these models, highlighting the specialized frameworks and metrics used to assess their performance in healthcare-specific downstream tasks in foundation models.

Biomedical foundation models

The application of foundation models has revolutionized numerous fields, particularly in natural language processing (NLP) and vision-language multimodal tasks. This transformation is largely attributed to several pivotal factors. Firstly, extensive pre-training on large text corpora allows these models to develop comprehensive universal language representations, which significantly enhance performance on various downstream tasks [176]. Secondly, such pre-training provides an improved initialization for models, which

not only boosts generalization capabilities but also speeds up convergence on specific target tasks. Thirdly, this method acts as a powerful form of regularization, crucial for preventing overfitting, particularly when training data is scarce.

Meanwhile, Vision-language multimodal models [50, 177, 178] are emerging as a powerful subset of foundation models, particularly in the field of image classification tasks in biomedical healthcare. These models synergistically combine the capabilities of image processing and language understanding to tackle complex tasks that require the integration of visual and textual data. In the healthcare sector, this ability is invaluable, as it enables the models to interpret medical imagery, such as scans and X-rays, alongside associated clinical notes or diagnostic information [179]. For example, a vision-language model might analyze an MRI scan while simultaneously considering a patient's written medical history to provide a more accurate diagnosis. This dual capability enhances the model's precision in identifying disease markers, understanding patient symptoms, and suggesting appropriate medical interventions. The integration of these two modalities in a single model not only streamlines the diagnostic process but also improves the accuracy of treatment recommendations, paving the way for more personalized and effective healthcare solutions. By leveraging these advanced models, medical professionals can gain deeper insights into patient conditions, leading to better outcomes and more efficient management of healthcare resources.

The training of foundation models (FMs) in the biomedical domain involves several crucial phases that enhance their applicability and effectiveness. Initially, unsupervised pretraining [5] plays a pivotal role, where models learn from large corpora without labeled data. This phase emphasizes the discovery of inherent structures and abstract relationships within the data, without the need for specific predictive tasks, making it invaluable for identifying complex patterns. Subsequently, self-supervised learning forms the backbone of foundation models, traditionally utilizing unstructured text from general-domain sources such as Wikipedia or web-crawled pages. Recent advancements, however, have steered the customization of pre-trained FMs towards specific fields to better meet domain-specific requirements. For instance, CodeBERT [180] is meticulously trained on programming languages to proficiently comprehend and generate code, whereas SciBERT [58] is tailored for parsing scientific publications and biological sequences, addressing the unique challenges of academic and medical research. Following this, reinforcement learning from human feedback (RLHF) [9] introduces a novel fine-tuning approach where models are adjusted based on rewards derived from human feedback rather than traditional labels. This method significantly aligns model outputs with human values and preferences, essential for applications demanding high engagement and accuracy in user interactions. Lastly, in-context learning [181], especially effective in models like GPT, leverages the model's capacity to generalize from a few examples. By presenting models with specific examples of the desired task at inference time, they dynamically adapt their responses to the context, enhancing their flexibility and utility without the need for additional training. This sequence of training methods collectively enhances the adaptability and performance of FMs, making them highly suitable for sophisticated tasks in the biomedical domain.

Foundation models can greatly benefit from training on expanded, domain-specific corpora [182]. For achieving peak performance in specialized downstream tasks, it is

increasingly recognized that integrating in-domain data during the training phase is imperative. This targeted approach not only refines the model's understanding of complex biomedical terminologies but also significantly enhances its practical applications in healthcare. By tailoring the training process to incorporate specific biomedical vocabulary and contextual nuances, FMs can be transformed into more effective tools, offering substantial improvements in processing and understanding medical texts, which is vital for advancing innovations and solutions within the healthcare industry.

Biomedical healthcare ML applications and benchmarks

Applications The application of FMs in the biomedical domain is propelled by a range of compelling reasons, each underscoring the unique challenges and opportunities this field presents.

- *Complexity of Sequential Biomedical Data:* Biomedical information, including electronic health records and biomedical texts, often comes in the form of sequential tokens lacking annotations. Historically, this complexity posed significant hurdles for effective data modeling. However, advancements in FMs have enabled effective training on such data in a self-supervised manner, significantly expanding the possibilities for processing and understanding biomedical information using these sophisticated models.
- *Scarcity of Annotated Data:* In the biomedical field, annotated data is typically scarce and expensive to produce, often leading to “zero-shot” or “few-shot” learning scenarios. Recent developments in language models, notably GPT-3 [45], have showcased remarkable capabilities in few-shot and even zero-shot learning. This evolution means that a well-trained FM can act as a powerful feature extractor in the biomedical domain, reducing the dependency on large volumes of annotated data and easing the barriers to entry for complex biomedical analysis.
- *Knowledge Intensity:* The biomedical sector is densely packed with specialized knowledge, much more so than general domains, often necessitating expert-level understanding. FMs serve as an accessible, soft knowledge base [183], which can assimilate and replicate expert knowledge from vast biomedical texts without direct human annotation. For example, GPT-3 has shown an impressive ability to recall and apply extensive, intricate common knowledge in practical applications [45], demonstrating its utility as a tool for knowledge dissemination and decision support in healthcare.
- *Diversity of Biological Data:* The scope of biomedical data extends beyond textual information to include diverse biological sequences, such as proteins and DNA. The application of FMs to these types of data has been notably successful, particularly in tasks like protein structure prediction. This success underlines the potential of FM to tackle a broader array of biological challenges, suggesting a promising future where FMs contribute substantially to critical tasks in genomics, proteomics, and other areas of biological research.
- *Speed of Knowledge Synthesis* [184]: The rapid pace at which biomedical knowledge evolves makes it challenging to keep up with the latest research and clinical practices manually. FMs, trained on the latest corpus of literature and clinical

guidelines, can quickly synthesize new information, making them invaluable tools for healthcare professionals who need to stay informed about the latest developments in real-time.

- *Enhanced Predictive Analytics* [185]: FMs have the potential to revolutionize predictive analytics in healthcare by integrating diverse data types—from patient records to research articles—to predict disease outbreaks, patient outcomes, and treatment efficacy. This capability can lead to more personalized medicine, where treatments are tailored to individual patients based on predictions made by these models.
- *Automated Reasoning and Decision Support*: FMs can be employed to automate reasoning processes and support decision-making in clinical environments. By processing and analyzing large volumes of medical data, these models can suggest diagnostic options, propose treatment plans, and even predict possible complications, thereby assisting medical professionals in making better-informed decisions.
- *Reduction in Diagnostic Errors* [186]: By providing comprehensive, data-driven insights, FMs can help reduce diagnostic errors, one of the significant challenges in healthcare. Their ability to learn from vast datasets and identify patterns that may be overlooked by human experts can contribute to more accurate diagnoses and, consequently, more effective treatments.

These factors collectively motivate the integration of foundation model into biomedical research and healthcare operations, indicating a robust pathway for leveraging AI to manage and utilize complex biomedical data more effectively.

Benchmarks The application of pre-trained FMs in the biomedical field exploits a diverse array of unstructured data sources, including electronic health records, scientific publications, social media texts, biomedical image-text pairs, and various biological sequences such as proteins. For a comprehensive review of mining electronic health records (EHR), please refer to the previous survey [187]. Discussions on the integration of health records and social media texts are explored in [188], while a systematic overview of biomedical textual corpora is presented in [189].

Key Benchmarks in Biomedical Research:

- *Electronic Health Records (EHR)*: EHRs encapsulate a comprehensive digital record of patient information, including demographics, medical history, medications, laboratory test results, and billing details. They are pivotal for longitudinal studies, allowing researchers to track patient outcomes over time and identify patterns and predictors of diseases. The vast amount of data contained within EHRs makes them invaluable for training FMs to recognize and predict medical conditions accurately, although access is tightly regulated to protect patient privacy [190, 191].
- *MIMIC-III (Medical Information Mart for Intensive Care III)*: This critical care database contains detailed information from over 58,976 ICU admissions, including 2,083,180 vital signs, medications, laboratory measurements, observations, and notes. This richness makes MIMIC-III ideal for developing models that predict patient outcomes, tailor treatments, and conduct epidemiological studies in critical care settings [192].

- *CPRD (Clinical Practice Research Datalink)*: A comprehensive dataset that provides a complete medical record from GP practices in the UK. It includes diagnoses, prescriptions, and clinical events, making it highly suitable for observational studies and clinical trials. The linkage to secondary care data enhances its utility in comprehensive healthcare research [193].
- *Reddit and Tweets*: These datasets are increasingly used for public health monitoring and sentiment analysis. Reddit's COMETA corpus and Twitter's COVID-tweet-BERT provide real-time data on public health trends, misinformation patterns, and community response to health crises, which are crucial for understanding public health behavior and improving communication strategies [194, 195].
- *MIMIC-CXR*: This dataset of chest x-rays and accompanying radiological reports is crucial for developing automated diagnostic tools that assist radiologists in detecting and diagnosing pathologies from imaging studies. The textual descriptions help train models to correlate visual signs with diagnostic language [196].
- *DNA Dataset*: This genomic dataset facilitates the training of models on genetic sequences to predict gene functions, understand genetic variations, and assist in personalized medicine strategies. It is essential for advancing genomics research and integrating genetic information with clinical data [197].
- *fMRI datasets*: These datasets comprise data from functional magnetic resonance imaging (fMRI) studies, which are invaluable in providing detailed insights into brain activity. Utilized extensively in neuroscience, fMRI data helps researchers understand brain functions, diagnose neurological disorders, and predict outcomes of therapeutic interventions. Notable datasets like the Philadelphia Neurodevelopmental Cohort (PNC) [198], Autism Brain Imaging Data Exchange (ABIDE) [199], and UK Biobank [200] include both functional and structural brain imaging data. These resources are critical for advancing our understanding of the brain, enhancing the accuracy of neurological diagnoses, and improving the efficacy of treatments by enabling a deeper analysis of the brain's response to various stimuli and conditions.
- *The Human Protein Atlas*: Contains high-resolution images detailing the spatial distribution of proteins in human tissues and cells [201]. This atlas is used for bioinformatics studies that integrate protein expression with gene expression data to elucidate cellular functions and disease mechanisms.
- *GEUVADIS RNA sequencing dataset* [202]: Provides RNA sequencing data from multiple populations, which is crucial for understanding how genetic variation affects gene expression across different human populations. This dataset is instrumental in studying population genetics, evolutionary biology, and disease susceptibility.
- *ImageCLEFmed* [203]: A benchmark dataset for multimodal biomedical information retrieval that includes medical images, captions, and text descriptions. It supports tasks such as medical image classification, annotation, and retrieval, which are crucial for medical informatics applications.

These datasets exemplify the diverse types of biomedical data available for research, each offering unique insights and challenges that can be leveraged to train more

effective and nuanced FMs for varied applications in healthcare and medical research. Note that the numerical details of the datasets are demonstrated in Table 1.

Biomedical healthcare on large language models

As introduced in [Backbone networks in foundation models](#) section, the backbone of most pre-trained foundation models, including prominent ones like BERT, GPT, T5, and their variants, is founded on the Transformer architecture, which framework is characterized by its reliance on self-attention networks and feed-forward networks (FFNs). The benign enables dynamic interactions between tokens, enhancing the model's ability to handle complex input relationships, while FFNs perform non-linear transformations to deepen token representations, bolstering feature extraction capabilities.

In parallel, the evolution of text representation through FMs has significantly advanced from initial static word embedding methods to sophisticated models capable of understanding contextual nuances [204, 205]. Historical neural language models laid the groundwork by predicting word contexts in a unidirectional manner, but modern approaches like ELMO [41], GPT [206], and BERT [207] have transformed the landscape with bi-directional and context-aware strategies. These models, through methodologies such as bidirectional language modeling and masked language model tasks, offer dynamic, context-sensitive word representations that vastly enhance performance across diverse NLP applications, making them fundamental to contemporary language processing tasks.

How to tailor LLMs to the biomedical domain

The adaptation of large language models to the biomedical domain involves specialized methodologies tailored to enhance their functionality for this sector's unique tasks. Initially crafted for general natural language processing (NLP) tasks, models like BERT [207] typically undergo a two-stage training process: initial training through a self-supervised meta-task (such as a masked language model or causal language model) on a broad, task-agnostic corpus, followed by fine-tuning on more

Table 1 Overview of key biomedical healthcare benchmarks

Dataset	Size	Types
MIMIC-III	58,976 admissions	Text, Numeric, Categorical
CPRD	11.3 million patients	Text, Numeric, Categorical
Reddit and Tweets	800K reddit posts and up-to-date tweets	Text
MIMIC-CXR	77,110 images	Images, Text
DNA Dataset	106 DNA sequences	Genetic Sequences
PNC	9,500 participants	Imaging (Functional, Structural)
ABIDE	Over 1,100 individuals	Imaging (Functional, Structural)
UK Biobank	Over 500,000 participants	Imaging (Functional, Structural), Genetic, Text
Human Protein Atlas	12,003 proteins	Images, Text
GEUVADIS RNA sequencing	462 individuals	Genetic Sequences

specialized, often smaller-scale, downstream tasks relevant to specific fields. The two strategies have been developed to better integrate LLMs into the biomedical field are as follows:

- *Continual Pre-training*: This method involves taking general LLMs such as BERT, initially pre-trained on extensive general corpora like Wikipedia or BookCorpus, and continuing their training on domain-specific corpora, such as PubMed texts and MIMIC-III data. For instance, BioBERT [208] extends BERT's training to include PubMed abstracts and articles, while BlueBERT [209] is further trained on both PubMed and MIMIC-III texts. These adaptations often retain the original model's vocabulary, which may not fully capture the specialized terminology of biomedical texts [182].
- *Pre-training from Scratch*: Some research advocates starting anew with domain-specific corpora to tailor PLMs more closely to biomedical needs [182, 210]. SciBERT [210] is an example of this approach, where a novel vocabulary of 30,000 terms specific to the domain was developed, and the model was trained on a corpus comprising both computer science (18%) and biomedical (82%) texts. However, recent findings suggest that mixed-domain pre-training might not be optimal for applications requiring high domain specificity. Instead, exclusive pre-training on biomedical corpora is recommended to ensure maximum relevance and efficacy.

Related Literatures Before exploring the applications of Large Language Models in the biomedical healthcare domain, it is essential to recognize several representative surveys and peer-reviewed publications that have thoroughly reviewed the landscape of biomedical language models. These resources provide invaluable insights into the development, applications, and future prospects of LLMs within this specialized field, laying a foundational understanding for ongoing and future research. Several key surveys and publications have extensively discussed the current state and potential advancements of transformer-based biomedical models and general prompting methods in natural language processing:

- *AMMU: A Survey of Transformer-Based Biomedical Pretrained Language Models* [211]: This comprehensive survey examines the evolution and impact of transformer-based models that have been specifically developed for the biomedical field. The survey details various approaches to adapting general language models to address the unique challenges posed by biomedical texts, such as the high specificity of vocabulary and the critical nature of the accuracy needed in medical contexts. The authors discuss multiple models that have been successfully implemented, highlighting their methodologies, the datasets they were trained on, and their performance on different biomedical NLP tasks. It provides a critical analysis of the strengths and limitations of these models, offering insights into how the field might evolve and suggesting directions for future research to enhance model accuracy and applicability.
- *Pre-train, Prompt, and Predict: A Systematic Survey of Prompting Methods in Natural Language Processing* [212]: by Liu et al. [212]: This survey explores the

relatively new technique of prompting, which adapts pre-trained models to specific tasks using minimal task-specific data. Prompting involves modifying the input to pre-trained models in such a way that the task is reformulated to leverage the model's existing knowledge. The survey systematically categorizes different types of prompts, discusses their applications in various NLP tasks, and evaluates their effectiveness across several benchmarks. It provides a detailed look at how prompting can reduce the need for large annotated datasets, which is particularly beneficial in domains like biomedicine where acquiring such data can be costly and time-consuming. The paper also considers the future of prompting in NLP, suggesting that further refinement of prompting strategies could lead to more generalizable and efficient NLP systems.

- *Foundation Models in Healthcare: Opportunities, Risks & Strategies Forward* [213]: This survey delves into the dual-edged nature of applying foundation models within the healthcare sector. It discusses the substantial opportunities these models present, such as enhancing diagnostic accuracy, predicting patient outcomes, and personalizing treatment plans. However, it also addresses the significant risks involved, particularly concerning data privacy, model bias, and the ethical implications of automated decision-making in healthcare. The authors propose a framework of strategies to mitigate these risks while capitalizing on the potential benefits. These strategies include developing robust governance frameworks, ensuring transparency in model workings, and engaging with a broad range of stakeholders to ensure that the deployment of these models in healthcare settings is both ethical and effective.
- *On the Opportunities and Risks of Foundation Models* [25]: This broad survey provides an extensive overview of the application of foundation models across various domains, with a particular focus on their transformative potential and the risks they pose. In the context of healthcare, the survey highlights how these models can revolutionize medical research and practice by providing new insights into disease patterns and patient care strategies. However, it also raises critical concerns about the reliability, fairness, and transparency of these models, especially given their potential to impact patient outcomes directly. The paper calls for a balanced approach to harnessing the power of foundation models, advocating for rigorous testing, ethical considerations, and regulatory oversight to ensure they benefit society as a whole.

Practical LLMs in biomedical healthcare

Since the introduction of BERT, a variety of biomedical pre-trained language models have been developed, enhancing the capabilities of NLP applications within the biomedical field. These models have been adapted either by further training on specialized in-domain corpora or by being built from scratch to cater specifically to the needs of medical and scientific communication. Below is a detailed summary of several existing pre-trained language models, where the specialized corpora, LLM backbone and released date are highlighted in Table 2.

- *BioBert* [208]: A pioneering work represents a significant advance in the application of language models to the biomedical domain. By adapting the BERT architecture,

Table 2 Overview of pre-trained language models in biomedicine with release dates

Model Name	Corpora	LLM Backbone	Release Date
BioBert	PubMed abstracts, PMC articles	BERT	2020
MedBert	Medical texts, EHRs	BERT	2021
ClinicalBERT	MIMIC-III clinical notes	BERT	2019
SciBERT	Scientific papers (82% biomedical)	BERT	2019
COVID-twitter-BERT	Tweets about COVID-19	BERT	2023
MedGPT	Electronic health records (EHRs)	GPT	2021
SCIFIVE	Biomedical corpora	T5	2021
LLMBiomedicine	Biomedical texts (NER [214] tasks)	GPT-4	2024
ClinicalGPT	Diverse medical data	GPT	2023
MultiMedQA	Medical QA datasets	PaLM [46]	2023
Chatdoctor	Patient-physician conversations	LLaMa	2023
Taiyi	Biomedical texts, multilingual	Qwen [215]	2024

originally designed for general language understanding, BioBert is fine-tuned with biomedical texts sourced from extensive databases such as PubMed abstracts and PMC full-text articles. This adaptation is not merely a continuation of training but a targeted effort to align the model's learning with the intricacies and terminologies unique to biomedical literature. As a result, BioBert excels in several biomedical text mining tasks including named entity recognition, relation extraction, and question answering over biomedical knowledge bases. The strength of BioBert lies in its ability to capture deep semantic connections between biomedical concepts, significantly improving the model's utility for researchers and healthcare professionals who rely on swift and accurate interpretations of medical texts.

- *MedBert* [24]: MedBert is an innovative approach to creating a language model that is steeped from the outset in the medical context. Unlike models that are adapted from general-purpose architectures, MedBert is pre-trained from scratch on a large and diverse corpus of medical texts, including electronic health records and other clinical documents. This ground-up approach allows MedBert to develop a nuanced understanding of medical language, including jargon, abbreviations, and the complex relationships between medical concepts. The model has shown significant improvements in tasks such as patient phenotyping and diagnostic prediction, making it a vital tool for healthcare analytics. MedBert's design addresses the challenges of applying general language models to medical data, ensuring that the nuances and critical details of medical communication are not lost in translation.
- *ClinicalBERT* [216]: Tailored for understanding and processing clinical notes, ClinicalBERT was trained exclusively on data from the MIMIC-III database [192], which includes around 2 million clinical notes. This specialized training prepares ClinicalBERT to handle a variety of clinical documentation styles and medical shorthand, making it an invaluable tool for applications like patient outcome prediction and automated documentation review, which require a deep understanding of clinical narratives.
- *SciBERT* [210]: Developed from scratch, SciBERT focuses on scientific text, primarily from the biomedical field, leveraging a corpus of papers available through the

Semantic Scholar database. With 82% of its training corpus composed of biomedical research articles, SciBERT is adept at deciphering complex scientific terminology and extracting relevant information from scholarly articles, thereby facilitating advanced text mining and information retrieval tasks in scientific research.

- *COVID-twitter-BERT* [195]: This model was specifically developed to analyze and understand discourse about COVID-19 on Twitter. It was trained during the initial stages of the pandemic on a dataset comprising approximately 160 million tweets related to the virus. The model is designed to capture the nuances of public sentiment, misinformation, and evolving topics related to COVID-19, providing valuable insights for public health officials and researchers studying communication patterns during health crises.
- *MedGPT* [22]: Inspired by the GPT architecture, MedGPT was trained on electronic health records (EHRs) and is designed to predict future medical events based on patients' medical histories. Its training allows it to model and predict various outcomes, such as diagnoses and complications, making it a potential tool for prognostic assessments in clinical settings.
- *SCIFIVE* [217]: This model is a domain-specific adaptation of the T5 model, trained under the Seq2seq framework on extensive biomedical corpora. SCIFIVE is engineered to transform complex biomedical queries into concise answers, facilitating tasks such as summarizing scientific texts and generating explanatory notes from dense medical data.
- *LLM-Biomedicine* [218]: This research highlights the effectiveness of meticulously designed prompts and the strategic selection of in-context examples to enhance the performance of LLMs on biomedical NER tasks. By adjusting prompts and examples to better fit the context of biomedical data, the study demonstrates significant improvements in model performance, making LLMs more adept at identifying and classifying medical entities in text.
- *ClinicalGPT* [219]: ClinicalGPT, a model that has been fine-tuned with a diverse set of medical data to enhance its performance and reliability in clinical scenarios. The model undergoes rigorous evaluations to ensure it meets the high standards required for medical applications, focusing particularly on its ability to maintain factual accuracy and provide contextually appropriate responses in simulated clinical interactions. ClinicalGPT represents a significant advancement in the use of LLMs in medicine, offering potential improvements in automated patient interaction, diagnostic support, and personalized treatment planning. By leveraging a vast corpus of medical texts for fine-tuning, the model is better equipped to handle the nuanced and highly specialized language found in clinical notes and patient interactions.
- *MultiMedQA* [220]: MultiMedQA is a comprehensive benchmark combining six existing medical question answering datasets, which span a variety of contexts from professional medicine to consumer health inquiries. The benchmark is enhanced by a newly developed dataset, HealthSearchQA, which consists of medical questions frequently searched online. This diverse collection of datasets is utilized to test the LLMs' ability to understand and process complex medical information across different facets of healthcare and patient inquiries. The authors discuss the significant challenge of assessing LLMs in clinical settings, where the accuracy of information

and the models' understanding of nuanced medical language are crucial. By employing MultiMedQA, the authors aim to provide a more nuanced and thorough evaluation of LLMs than previous benchmarks allowed.

- *Chatdoctor* [221]: Chatdoctor improves the performance and relevance of responses in medical conversational systems. To achieve this, the model was fine-tuned using a substantial dataset of 100,000 real-world patient-physician conversations sourced from online medical consultations. This approach ensures that the model not only understands medical terminology and procedures but also grasps the nuances of patient interactions and inquiries.
- *Taiyi* [222]: Taiyi highlights the limitations of existing fine-tuned biomedical LLMs, which are predominantly monolingual and focused on question answering and conversation tasks within the biomedical field. Taiyi, by contrast, is designed to enhance performance across a broader spectrum of NLP applications, including entity extraction, relation extraction, and information retrieval, catering to both English and non-English texts. The development and evaluation of Taiyi involve rigorous fine-tuning processes that adjust the model to grasp the nuances and specific terminology used in various biomedical contexts, significantly improving its utility and applicability in a global healthcare context. This model represents a substantial advancement in the field of biomedical NLP by supporting multilingual capabilities and addressing the critical need for diverse language processing in medical research and healthcare delivery.

Biomedical LLMs like BioBert, MedBert, and ClinicalBERT have been developed to enhance tasks such as named entity recognition, relation extraction, and patient outcome prediction by training on specialized datasets like PubMed and MIMIC-III. While these models excel in capturing deep semantic relationships within specific domains, they often struggle with generalizability outside their specialized fields and require rigorous ongoing updates. Models like COVID-twitter-BERT, which focus on specific events or datasets, face challenges in maintaining relevance over time due to the dynamic nature of their data sources. Innovations like MedGPT and ClinicalGPT show promise in clinical settings, yet they must navigate significant challenges related to data privacy and the need for extensive, diverse training data to ensure accuracy and utility across varying medical scenarios. Furthermore, approaches like MultiMedQA and Taiyi aim to broaden the applicability of biomedical LLMs across different languages and medical contexts, yet they must balance the breadth of language coverage with the depth of medical understanding to be truly effective in global healthcare applications.

Biomedical healthcare on vision language models

This section elaborates on the training of vision language models for biomedical imaging, and their practical applications in the biomedical healthcare sector.

How to train vision language models for biomedical imaging

Deep neural networks demonstrate outstanding performance in various vision tasks, including image classification, object detection, and instance segmentation. A key to

this success in the foundation model era is the concept of pre-training, which, unlike in NLP where it usually involves language models, traditionally meant training on extensive labeled image datasets like ImageNet [48]. More recently, diverse learning methods have been introduced to overcome limitations of conventional supervised learning, such as generalization errors and spurious correlations. We examine several methodologies suitable for imaging applications as follows:

- *Unsupervised Pre-training*: Unsupervised pre-training leverages large volumes of unlabeled image data to learn rich feature representations without the guidance of explicit annotations. Techniques such as autoencoders [223] and generative adversarial networks (GANs) [224] train models to generate or reconstruct images, enabling them to capture the underlying data distributions and learn complex patterns within the visual inputs. This approach is particularly useful in domains where labeled data is scarce or expensive to obtain.
- *Contrastive Self-supervised Learning*: Contrastive self-supervised learning techniques [225–227] train models to differentiate between various modifications of a given input image, such as determining whether two images are rotated versions of each other or entirely distinct. This method enables the model to develop features applicable to diverse vision tasks, including object detection and semantic segmentation.
- *Masked Self-supervised Learning*: Drawing inspiration from BERT's approach in NLP, masked self-supervised learning [228–230] is gaining popularity in computer vision. This generative pre-training method trains models to reconstruct images from partially obscured inputs, aiding in understanding the underlying structure of visual data.
- *Contrastive Language-image Pre-training*: An innovative method, contrastive language-image pre-training [50] (CLIP), involves training a vision model using diverse image-text datasets. The model learns to match images with corresponding texts within a mini-batch through contrastive learning. CLIP shows impressive zero-shot capabilities, performing on par with traditional models like ResNet [35] on ImageNet without task-specific training. Text descriptions enhance understanding of the visual content, facilitating the model's comprehension of visual elements and their interrelations, which is essential for effective learning.
- *Instructed fine-tuning*: Instructed fine-tuning involves explicitly guiding the model during the fine-tuning process with task-specific instructions. This method builds upon the foundation established during pre-training by aligning the model's learning objectives closely with the nuances of the target task [50]. For example, in biomedical imaging, models can be instructed to identify specific medical conditions from images using detailed descriptions of symptoms or expected imaging features. This approach helps the model to focus on relevant aspects of the data, enhancing its performance on specialized tasks such as diagnosing diseases from medical scans.

Note that a significant challenge in harnessing FMs for vision-language tasks lies in overcoming the “task gap” and the “domain gap”. The task gap refers to the differences between the generic meta-tasks used in FMs, such as masked language modeling in

BERT or causal language modeling in GPT, and the specialized requirements of downstream vision-language tasks, such as medical image annotation or diagnostic interpretation. The domain gap further highlights the disparity between the general training corpora used, and the highly specialized datasets needed for tasks in specific fields like biomedicine. To effectively deploy a pre-trained language model in vision-language applications within a specific domain, it is crucial to undertake both domain and task adaptations [182, 231–233]. Domain adaptation involves additional training of a model—originally pre-trained on broad, general datasets within a targeted domain, such as biomedicine. This step ensures that the model becomes attuned to the specific terminologies and data types characteristic of the domain.

Practical VLMs in biomedical healthcare

Biomedical vision-and-language models have largely been shaped by influential self-supervised pre-training techniques, such as SimCLR [225] in computer vision and BERT [207] in natural language processing. These foundational approaches have paved the way for the adoption of advanced text-to-image diffusion models [27, 51, 234] in the medical field [235, 236], enhancing tasks ranging from diagnostic imaging to patient interaction. This subsection provides a detailed overview of the existing vision-and-language models (VLMs) within the biomedical sector and elucidates their functionalities. In this survey, VLMs in the biomedical healthcare sector are categorized into three primary types: dual-encoder, fusion encoder, and hierarchical structures. Each model type offers distinct advantages and limitations, tailored to specific application needs within the healthcare context.

Dual-Encoder Models process visual and textual inputs independently through separate encoders before merging the resulting vectors for final task execution. This architecture is particularly effective for tasks that require robust single-modal or crossmodal representation, such as image classification, image captioning, and cross-modal retrieval. However, the dual-encoder approach may fall short in fully capturing the intricate interplays between visual and linguistic elements, which can limit its effectiveness in more complex multimodal tasks. Fusion-encoder models integrate visual and linguistic data early in the processing pipeline, utilizing a single encoder to manage both modalities. This method facilitates the capture of complex interactions between text and image, proving advantageous for tasks that demand a deep multimodal understanding, such as visual question answering and complex diagnostic reasoning. While fusion encoders excel in multimodal integration, they may encounter challenges in scenarios where a clear distinction between modalities is necessary.

Besides the dual-encoder and fusion-encoder models, the field also explores innovative biomedical FMs that combine vision and language, such as hierarchical encoder alignment [237, 238] and medical text-to-image diffusion models [234, 239]. Hierarchical alignment constructs input pyramids on both visual and linguistic sides, enhancing the model's ability to match features across modalities at multiple abstraction levels, which not only improves feature correspondence and model generalization but also optimizes the learning process, making it more efficient and adaptable to complex tasks like medical diagnosis from imaging and textual data. Such structured FMs offers significant advantages in terms of computational efficiency, robustness, and scalability,

demonstrating potential for broad applications, especially in the biomedical field. Diffusion models are generative models inspired by non-equilibrium thermodynamics. They operate by defining a Markov chain of diffusion steps that gradually add random noise to data. The model then learns to reverse this diffusion process, reconstructing the desired data samples from the noise. This approach is particularly powerful in medical applications where generating high-fidelity images from textual descriptions can assist in diagnostic visualizations and treatment planning.

We summarize recent significant developments in adapting vision-language models for biomedical healthcare as follows, where the model type, encoder details, training corpora, and released dates are detailed at Table 3.

- *ConVIRT* [253]: ConVIRT utilizes contrastive learning techniques to simultaneously train ResNet and BERT encoders on paired image and text data. This approach significantly improves the model's performance in image classification tasks by effectively reducing the dependency on large volumes of labeled data. By optimizing feature extraction and enhancing the semantic alignment between images and their textual descriptions, ConVIRT enables more accurate and efficient classification, making it particularly useful in scenarios where annotated datasets are limited.
- *GLoRIA* [254]: GLoRIA advances the field by employing both global and local contrastive learning strategies to finely align words in radiology reports with corresponding sub-regions within images. This method enhances local representation learning, allowing for more precise identification and classification of localized features in medical images. Such detailed alignment improves diagnostic accuracy and aids in the development of more sophisticated automated radiology analysis tools.
- *MedCLIP* [20]: MedCLIP leverages the innovative architecture of the CLIP model, specifically tailored for the medical domain. By utilizing pre-computed matching

Table 3 Overview of vision-language models in biomedical healthcare

Model Name	Type	Image Encoder	Text Encoder	Training Corpora	Release Date
ConVIRT	Dual	ResNet	ClinicalBERT	MIMIC-CXR	2022
GLoRIA	Dual	ResNet	BioClinicalBERT	Chexpert [240]	2021
MedCLIP	Dual	ResNet/ViT	BioClinicalBERT	Chexpert, MIMIC-CXR	2022
CheXZero	Dual	CLIP-Image	CLIP-Text	Chest X-rays	2022
LoVT	Dual	ResNet	ClinicalBERT	MIMIC-CXR	2022
Adapted VLMs	Hierarchical	Diffusion, VAE [241]	Bert, CLIP	Chexpert, MIMIC-CXR	2022
VisualBERT	Fusion	Varies	BERT	MIMIC-CXR	2020
MedViLL	Fusion	ResNet	BERT	MIMIC-CXR	2022
ARL	Fusion	CLIP-Image	RoBERTa [242]	MediCaT [243], MIMIC-CXR, ROCO[244]	2022
LViT	Fusion	ViT	BERT	QaTa-COV19 [245], MoNuSeg [246]	2023
RoentGen	Hierarchical	Diffusion	CLIP-Text	MIMIC-CXR	2022
CLIPSyntel	Dual	CLIP	GPT-3.5	MMQS [247]	2024
Med-unic	Dual	ResNet/ViT	CXR-BERT [248]	MIMIC-CXR, PadChest [249]	2024
EchoCLIP	Dual	ConvNeXt [250]	CLIP-Text	Echocardiogram videos	2024
Llava-med	Fusion	Llava [8]	Llava	PubMed [251], PMC-15M [252]	2024

scores, MedCLIP enhances the alignment between medical images and their corresponding textual descriptions. This capability facilitates effective zero-shot learning, allowing MedCLIP to accurately classify medical conditions without the need for extensive fine-tuning on large annotated datasets. The model's ability to directly apply learned representations from diverse medical contexts makes it a valuable tool for rapid and efficient disease diagnosis, particularly in environments where labeled medical data is scarce.

- *CheXZero* [255]: CheXZero is another adaptation of the CLIP model, focused on zero-shot learning for medical imaging, specifically in chest radiography. Unlike traditional models that require detailed annotations for each new disease classification task, CheXZero applies the powerful zero-shot capabilities of the CLIP model to accurately identify pathologies in chest X-rays without additional model training. This approach is particularly beneficial for rapidly evolving medical scenarios, such as new disease outbreaks or rare conditions, where the availability of comprehensive labeled datasets might be limited. CheXZero's innovative use of CLIP for direct application in medical diagnosis demonstrates its potential to significantly streamline diagnostic processes in healthcare settings.
- *LoVT* [256]: LoVT specifically targets localized medical imaging tasks by implementing a local contrastive loss that aligns representations of sentences or specific image regions. This alignment is crucial for tasks that require detailed understanding of small, localized anatomical structures or pathological features, enhancing the model's accuracy in specialized medical imaging applications.
- *Adapting Pretrained Vision-Language Foundational Models to Medical Imaging Domains* [236]: This study explores the effectiveness of adapting general vision-language models to the medical imaging domain. It demonstrates how foundational models, originally designed for broad applications, can be fine-tuned to meet the specific needs of medical diagnostics and research, thus broadening their applicability and improving performance in specialized tasks.
- *VisualBERT* [257]: This study focuses on adapting general-domain vision-language models, such as LXMERT and VisualBERT, for the integration of medical images and texts. The effectiveness of these adapted models in disease classification showcases their potential in clinical settings, where they can support diagnostic processes and enhance the accuracy of medical assessments.
- *MedViLL* [258]: MedViLL enhances the multimodal interaction between medical images and associated textual data through an innovative vision-language model framework by incorporating extensive medical knowledge and utilizing tailored masking schemes, MedViLL is specifically designed to improve understanding and generation tasks within the medical field, which excels in synthesizing comprehensive medical reports and generating detailed medical annotations, crucial for assisting healthcare professionals in making informed decisions. The approach to integrating complex medical datasets ensures a deeper contextual understanding and a more nuanced interpretation of both visual and textual medical data.
- *ARL* [259]: ARL (Align and Reasoning Language model) introduces a unique alignment strategy that specifically targets the challenges of medical imaging and text analysis. By aligning sentence or image region representations through a localized

contrastive loss, ARL effectively bridges the gap between visual features and their corresponding textual annotations. This model is particularly adept at tasks that require precise localization of medical findings within images, supporting detailed diagnostic processes. ARL's focus on enhancing the correlation between detailed image regions and descriptive text makes it an invaluable tool for advanced medical imaging applications where accuracy and detail are paramount.

- *LViT* [260]: LViT leverages medical text annotations to significantly improve segmentation results, particularly in semi-supervised settings where labeled data may be scarce. By integrating rich textual information, LViT enhances its understanding of medical imagery, leading to more accurate segmentation and analysis of medical scans.
- *RoentGen* [235]: RoentGen introduces a pioneering approach by applying text-to-image diffusion models to medical imaging. This novel methodology holds promise for generating detailed and accurate medical images from textual descriptions, potentially revolutionizing the way medical imagery is produced and understood.
- *CLIPSyntel* [247]: CLIPSyntel represents a synergistic application of CLIP and large language models to address the challenge of multimodal question summarization in healthcare. This model harnesses the strengths of both visual and textual data processing to provide concise and relevant summaries of complex medical inquiries, aiding healthcare professionals. CLIP and LLM Synergy for Multimodal Question Summarization in Healthcare [261].
- *Med-unic* [262]: Med-unic presents an approach to enhance the performance of medical vision-language pre-training models across different languages. The authors focus on reducing bias in these models, which often perform better in languages with abundant training data (like English) compared to languages with less data. They introduce a unified framework that integrates multilingual textual features and visual content effectively. Their method involves using a debiasing technique that ensures more equitable learning from visual and textual data across various languages. This is achieved by carefully balancing the dataset and incorporating cross-lingual adaptation [263] techniques to improve model performance uniformly across different linguistic contexts.
- *EchoCLIP* [264]: EchoCLIP is a specialized vision-language foundation model designed to improve echocardiography interpretation. EchoCLIP leverages the relationship between cardiac ultrasound images and expert cardiologist interpretations across diverse patient groups and diagnostic scenarios. The development of this model addresses the critical challenge of limited availability of annotated clinical data in cardiac imaging. By training on over one million cardiac ultrasound images, EchoCLIP aims to enhance the accuracy and efficiency of echocardiogram, offering a robust tool for cardiac diagnostics that can adapt to various clinical conditions and imaging indications.
- *Llava-med* [265]: Llava-med presents a novel approach to training a vision-language conversational assistant tailored for the biomedical field. The study introduces a cost-efficient method for rapidly developing a multimodal conversational AI that can understand and discuss biomedical images alongside textual data. Unlike previous models that rely extensively on large-scale image-text pairs from general domains,

Llava-med is trained specifically with biomedical data to better address the unique needs of the medical community.

VLMs in the biomedical domain offer diverse applications but also face certain limitations. Models like ConVIRT and MedCLIP leverage contrastive and zero-shot learning to improve semantic alignment between medical images and texts, reducing reliance on extensive labeled datasets and enhancing diagnostic accuracy. However, these models may struggle with generalization outside their training specifics and have limited ability for continual learning. GLoRIA and ARL focus on fine-grained alignment of medical data, improving localized feature identification but potentially lacking in broader application flexibility. Models such as CheXZero and RoentGen introduce innovative approaches to medical imaging by applying zero-shot learning and text-to-image diffusion models, streamlining diagnostic processes and even generating medical images from textual descriptions. Yet, these models do not fully address the varying complexities of medical conditions across diverse datasets. EchoCLIP and Llava-med exemplify specialized applications, targeting cardiac imaging and conversational biomedical AI, respectively, but must overcome challenges like limited annotated data and the need for specialized training to ensure widespread applicability.

Open challenges and opportunities in federated foundation biomedical research

The integration of AI technologies, particularly large pre-trained foundation models in the biomedical field, presents a range of future challenges and opportunities that must be tackled to unlock their full potential. This section delves into the key issues and potential avenues for progress concerning the application of federated foundation models in biomedical research. It underscores the need for robust solutions that ensure privacy, enhance model generalizability, improve computational efficiency, and address regulatory and ethical considerations. As we explore these challenges, we also highlight promising strategies that may pave the way for more effective and equitable AI-driven healthcare solutions.

Challenges of foundation models in biomedical healthcare

Research on foundation models in the biomedical healthcare domain presents several challenges and directions for future exploration:

- *Data Privacy and Security:* The primary challenge in foundation model, especially in healthcare, revolves around maintaining patient confidentiality and adhering to stringent data protection regulations like HIPAA [266] in the U.S and GDPR [17] in Europe. Future research needs to focus on developing robust encryption methods and privacy-preserving algorithms that allow for the secure sharing of insights without exposing sensitive patient data [267, 268].
- *Model Generalization across Diverse Datasets:* FMs involve training models on highly heterogeneous data sources, often leading to challenges in model generalization. Research should explore techniques to enhance the generalizability of

foundation models across diverse healthcare systems and varied patient demographics without compromising performance.

- *Scalability and Computational Efficiency*: The computational demand for training large-scale LLMs and VLMs is significant. Optimizing resource allocation, reducing communication overhead, and proposing efficient model updating mechanisms are crucial areas for future development to ensure scalability and practicality in real-world healthcare settings.
- *Bias and Fairness*: Ensuring that foundation models do not perpetuate or amplify biases present in downstream tasks is critical [26, 269, 270], especially under the biomedical domain, where the targeted problem for each patient can be narrow. Future research should include developing methodologies for bias detection and mitigation in model training and deployment phases. This also involves designing fair algorithms that provide equitable healthcare outcomes across different populations [271, 272].
- *Interoperability and Standardization*: There is a need for standardized protocols to ensure interoperability among different healthcare systems participating in foundation model learning.
- *Personalization*: Medical treatments and diagnostics often require high degrees of personalization. AI models must be capable of adapting to individual patient needs and conditions, which pose challenges in model design and data utilization without compromising generalizability.
- *Scaling*: Deploying AI solutions on a large scale, particularly in diverse healthcare settings, presents logistical and computational challenges. Scalability involves not only the expansion of AI systems to handle larger datasets but also ensuring these systems are accessible across different regions and healthcare infrastructures.
- *Biomedical Requirements of Accuracy*: Biomedical applications demand extremely high levels of accuracy and reliability. AI models used in diagnostics or treatment recommendation must meet rigorous standards to prevent errors that could adversely affect patient health.
- *Robustness to Adversarial Attacks*: As the applications of a foundation model in biomedical scenarios can be distributed, they are susceptible to various types of adversarial attacks that can compromise model integrity. Enhancing the robustness of foundation models against such attacks, and ensuring secure and reliable model performance, are a significant direction for ongoing research.
- *Regulatory and Ethical Considerations*: As foundation models evolve, there will be increased scrutiny from regulatory bodies concerning their use in clinical settings. Research must address these regulatory challenges by developing models that are not only effective but also transparent and explainable to satisfy regulatory requirements and maintain public trust.
- *Longitudinal Studies and Continuous Learning*: Implementing models that can adapt over time to new data and evolving biomedical conditions is crucial. Research into continuous learning mechanisms that allow FMs to update without forgetting [273] previously learned knowledge while integrating new insights is essential for maintaining the relevance and accuracy of biomedical models.

Opportunities in federated foundation models

Federated learning offers a unique framework for addressing several challenges associated with foundation models, particularly in the sensitive and data-intensive field of biomedical healthcare. In this survey, we would like to highlight how federated learning can help overcome the challenges of FMs and what opportunities it presents in both academic research and industrial applications:

- *Data Privacy and Security:* Federated learning enables the collaborative training of predictive models by sharing model updates rather than raw data. Each participating institution retains its data locally, significantly minimizing the risk of data breaches and unauthorized access. This method is especially beneficial in healthcare, where patient data is highly sensitive and subject to strict privacy regulations. Note that upholding data privacy is crucial not only for complying with laws like GDPR and HIPAA but also for maintaining patient trust. Federated learning's ability to train models without compromising data privacy helps healthcare organizations implement AI solutions without risking patient confidentiality or facing legal penalties.
Increased Model Robustness and Trustworthy: Trust in medical AI systems is essential for their acceptance by both medical professionals and patients. Systems known for their reliability and backed by a transparent, accountable training process are more likely to be trusted and thus more widely adopted, which makes it crucial to further investigate the trustworthy federated foundation models.
- *Bias and Fairness:* Addressing bias and ensuring fairness is critical for Federated Foundation Models, particularly in the healthcare domain where imbalanced data or algorithmic biases can lead to unequal patient outcomes. FL frameworks must consider demographic diversity and institutional disparities in data availability to create equitable models. Techniques such as fairness-aware aggregation methods, adversarial training for debiasing, and reweighting of underrepresented data during local training can mitigate biases. Additionally, the use of synthetic data generation to balance class distributions can further improve trust and fairness in healthcare applications.
- *Real-time Learning and Adaptation:* In federated learning frameworks, models can be updated continually as new data becomes available across the network. This dynamic learning process allows the models to adapt to emerging health trends or new strains of diseases. The ability to update and adapt foundation models in real-time is vital for keeping pace with the fast-evolving nature of diseases and treatments, ensuring that healthcare providers have the most current tools at their disposal.
- *Collaborative Innovation:* By aggregating insights from diverse healthcare environments and patient demographics, federated learning facilitates the development of models that perform well across different settings. This heterogeneous data input helps the model learn more comprehensive patterns and reduces the risk of bias towards any particular group or condition. Specifically, how to use federated learning to efficiently establish a cooperative ecosystem where different healthcare entities can contribute to and benefit from shared AI advancements without compromising their data sovereignty is worth for real-world application, which can lead to more rapid development and refinement of AI technologies.

- *Multimodality*: Medical data are often multimodal, encompassing an extensive array of data types-including text, images, videos, databases, and molecular structures-across various scales from molecules to populations [129, 274], and presented in both professional and/or lay language [275, 276]. While current self-supervised models excel within individual modalities, such as text [208], images [277], genes [197], and proteins [278], they typically lack the capability to integrate and learn from these diverse sources simultaneously. To truly leverage the rich information available across different modalities, there is an urgent need to develop models that can perform both feature-level and semantic-level fusion. Successfully integrating these varied data types could revolutionize how biomedical knowledge is unified and significantly accelerate discovery processes in biomedicine. Federated learning frameworks can capture a richer and more nuanced understanding of patient conditions, which is crucial for multimodal tasks like diagnosing complex diseases that may require correlating symptoms, radiology images, and genetic information.
- *Synthetic Data Generation for Further Training*: Federated learning can facilitate the generation of synthetic training data, which helps address the scarcity of annotated datasets, particularly in specialized medical fields. By learning from diverse sources, federated models can generate new, synthetic examples that preserve the statistical properties of real data without revealing any individual patient's information. This synthetic data can then be used to further train and refine FMs across the network. The generation of synthetic data is a critical solution for overcoming data limitations in biomedical research, where privacy concerns and the rarity of certain conditions can significantly constrain the availability of training data.

Conclusions

This survey has delved into the transformative potential of foundation models and federated learning within the biomedical healthcare domains. Foundation models represent a significant advancement in artificial intelligence, offering robust, adaptable tools that can be fine-tuned for specific applications without constructing new models from scratch. These models, trained on expansive datasets, are capable of performing a wide array of tasks-from text generation to video analysis-that were previously beyond the reach of earlier AI systems. In the biomedical and healthcare sectors, where the efficacy of AI and the integrity of data privacy are crucial, foundation models play a pivotal role by enabling the extraction of valuable insights from constrained datasets. This survey has highlighted the current applications of FMs in these sectors, particularly focusing on large language models and vision-language models.

Meanwhile, Federated learning, characterized by its privacy-preserving and decentralized approach, complements the capabilities of foundation models perfectly. By combining the robust, generalizable nature of FMs with the privacy-centric, decentralized attributes of federated learning, researchers can perform deep analyses using globally pooled insights from locally held datasets. This synergy holds immense potential to meet the specific needs of biomedical AI applications, offering scalable solutions that accommodate the continuous updating of foundation models with new, relevant data.

Additionally, this survey has outlined various challenges and opportunities that arise with the adoption of federated foundation models in healthcare. Federated learning

addresses critical issues such as data privacy, model generalization, scalability, and inherent biases within AI models. By allowing multiple institutions to collaboratively train models while keeping their data localized, federated learning not only complies with strict data privacy laws but also enhances the diversity and efficacy of medical AI applications. Key areas where federated foundation models could notably impact biomedical research and practice include enhancing model robustness and fairness, enabling real-time model updates and adaptations, and facilitating cross-institutional and international collaborations without compromising data security.

Acknowledgements

Not applicable.

Authors' contributions

Xingyu Li wrote the main manuscript, Lu Peng and YuPing Wang revised the main manuscript, prepared figures and tables. All authors reviewed the manuscript.

Funding

There was no external source of funding.

Data availability

No datasets were generated or analysed during the current study.

Declarations**Ethics approval and consent to participate**

Not applicable.

Consent for publication

Not applicable.

Competing interests

The authors declare no competing interests.

Received: 19 May 2024 Accepted: 9 December 2024

Published online: 04 January 2025

References

1. 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.* 2023;34:1–9.
2. Touvron H, Lavril T, Izacard G, Martinet X, Lachaux MA, Lacroix T, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971.* 2023.
3. LeCun Y, Bengio Y, Hinton G. Deep learning. *New York: IEEE Corporate Headquarters.* 2015;521(7553):436–44.
4. Erhan D, Courville A, Bengio Y, Vincent P. Why does unsupervised pre-training help deep learning? In: *Proceedings of the thirteenth international conference on artificial intelligence and statistics. JMLR Workshop and Conference Proceedings. Italy: Chia Laguna Resort; 2010.* p. 201–8.
5. Caron M, Bojanowski P, Mairal J, Joulin A. Unsupervised pre-training of image features on non-curved data. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision.* 2019. pp. 2959–68.
6. Chen T, Frankle J, Chang S, Liu S, Zhang Y, Carbin M, et al. The lottery tickets hypothesis for supervised and self-supervised pre-training in computer vision models. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition.* New York: IEEE Corporate Headquarters. 2021. p. 16306–16.
7. Wang X, Zhang R, Shen C, Kong T, Li L. Dense contrastive learning for self-supervised visual pre-training. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition.* New York: IEEE Corporate Headquarters. 2021. p. 3024–33.
8. Liu H, Li C, Wu Q, Lee YJ. Visual instruction tuning. *Adv Neural Inf Process Syst.* 2024;36.
9. Christiano PF, Leike J, Brown T, Martic M, Legg S, Amodei D. Deep reinforcement learning from human preferences. *Adv Neural Inf Process Syst.* 2017;30.
10. Park C, Took CC, Seong JK. Machine learning in biomedical engineering. *Biomed Eng Lett.* 2018;8:1–3.
11. Habehh H, Gohel S. Machine learning in healthcare. *Curr Genomics.* 2021;22(4):291.
12. Kaur D, Uslu S, Rittichier KJ, Durrresi A. Trustworthy artificial intelligence: a review. *ACM Comput Surv (CSUR).* 2022;55(2):1–38.
13. McMahan B, Moore E, Ramage D, Hampson S, Aguera y Arcas B. Communication-efficient learning of deep networks from decentralized data. In: *Artificial intelligence and statistics. PMLR; 2017.* pp. 1273–82.
14. Li X, Qu Z, Tang B, Lu Z. FedLGA: Toward System-Heterogeneity of Federated Learning via Local Gradient Approximation. *IEEE Trans Cybern.* 2023.

15. Li Q, Wen Z, Wu Z, Hu S, Wang N, Li Y, et al. A survey on federated learning systems: Vision, hype and reality for data privacy and protection. *IEEE Trans Knowl Data Eng.* 2021;35(4):3347–66.
16. Wei K, Li J, Ding M, Ma C, Yang HH, Farokhi F, et al. Federated learning with differential privacy: Algorithms and performance analysis. *IEEE Trans Inf Forensic Secur.* 2020;15:3454–69.
17. Li H, Yu L, He W. The impact of GDPR on global technology development. Hershey: Taylor & Francis; 2019.
18. of Medicine I. Beyond the HIPAA Privacy Rule: Enhancing Privacy, Improving Health Through Research. Nass SJ, Levit LA, Gostin LO, editors. Washington, DC: The National Academies Press; 2009. <https://doi.org/10.17226/12458>.
19. Lu W, Xixu H, Wang J, Xie X. Fedclip: fast generalization and personalization for clip in federated learning. In: *ICLR 2023 Workshop on Trustworthy and Reliable Large-Scale Machine Learning Models*. Appleton: International Conference on Learning Representations. 2023.
20. Wang Z, Wu Z, Agarwal D, Sun J. MedCLIP: Contrastive Learning from Unpaired Medical Images and Text. In: Goldberg Y, Kozareva Z, Zhang Y, editors. *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. Abu Dhabi: Association for Computational Linguistics; 2022. pp. 3876–87. <https://doi.org/10.18653/v1/2022.emnlp-main.256>.
21. Wu X, Liang Z, Wang J. Fedmed: A federated learning framework for language modeling. *Sensors.* 2020;20(14):4048.
22. Kraljevic Z, Shek A, Bean D, Bendayan R, Teo J, Dobson R. MedGPT: Medical concept prediction from clinical narratives. *arXiv preprint arXiv:2107.03134*. 2021.
23. Subramanian AAV, Venugopal JP. A deep ensemble network model for classifying and predicting breast cancer. *Comput Intell.* 2023;39(2):258–82.
24. Rasmy L, Xiang Y, Xie Z, Tao C, Zhi D. Med-BERT: pretrained contextualized embeddings on large-scale structured electronic health records for disease prediction. *NPJ Digit Med.* 2021;4(1):86.
25. Bommasani R, Hudson DA, Adeli E, Altman R, Arora S, von Arx S, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*. 2021.
26. Achiam J, Adler S, Agarwal S, Ahmad L, Akkaya I, Aleman FL, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*. 2023.
27. Ramesh A, Dhariwal P, Nichol A, Chu C, Chen M. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*. 2022;1(2):3.
28. Hochreiter S, Schmidhuber J. Long short-term memory. *Neural Comput.* 1997;9(8):1735–80.
29. Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, et al. Attention is all you need. *Adv Neural Inf Process Syst.* 2017;30.
30. Chowdhary K, Chowdhary K. Natural language processing. *Fundam Artif Intell.* 2020;603–49. <https://dl.acm.org/doi/10.5555/1074100.1074630>.
31. LeCun Y, Bottou L, Bengio Y, Haffner P. Gradient-based learning applied to document recognition. *Proc IEEE.* 1998;86(11):2278–324.
32. 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. In: *International Conference on Learning Representations*. Appleton: International Conference on Learning Representations; 2020.
33. Kenton JMW, Toutanova LK. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In: *Proceedings of NAAACL-HLT*. Stroudsburg: Association for Computational Linguistics (ACL). 2019. p. 4171–86.
34. Raffel C, Shazeer N, Roberts A, Lee K, Narang S, Matena M, et al. Exploring the limits of transfer learning with a unified text-to-text transformer. *J Mach Learn Res.* 2020;21(140):1–67.
35. He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. New York: IEEE Corporate Headquarters. 2016. p. 770–8.
36. Ba JL, Kiros JR, Hinton GE. Layer Normalization. *Stat.* 2016;1050:21.
37. Shrivastava A, Proukakis C. Novel positional encodings to enable tree-based transformers. *Adv Neural Inf Process Syst.* 2019;32.
38. Chen Y. Convolutional neural network for sentence classification. University of Waterloo; 2015.
39. Mikolov T. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*. 2013;3781.
40. Pennington J, Socher R, Manning CD. Glove: Global vectors for word representation. In: *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*. Stroudsburg: Association for Computational Linguistics (ACL). 2014. p. 1532–43.
41. Peters ME, Neumann M, Iyyer M, Gardner M, Clark C, Lee K, et al. Deep Contextualized Word Representations. In: Walker M, Ji H, Stent A, editors. *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. vol. 1 (Long Papers). New Orleans: Association for Computational Linguistics; 2018. pp. 2227–37. <https://doi.org/10.18653/v1/N18-1202>.
42. Radford A, Wu J, Child R, Luan D, Amodei D, Sutskever I, et al. Language models are unsupervised multitask learners. *OpenAI Blog.* 2019;1(8):9.
43. Bengio Y, Ducharme R, Vincent P. A neural probabilistic language model. *Adv Neural Inf Process Syst.* 2000;13.
44. Mikolov T, Sutskever I, Chen K, Corrado GS, Dean J. Distributed representations of words and phrases and their compositionality. *Adv Neural Inf Process Syst.* 2013;26.
45. Brown T, Mann B, Ryder N, Subbiah M, Kaplan JD, Dhariwal P, et al. Language models are few-shot learners. *Adv Neural Inf Process Syst.* 2020;33:1877–901.
46. Chowdhery A, Narang S, Devlin J, Bosma M, Mishra G, Roberts A, et al. Palm: Scaling language modeling with pathways. *J Mach Learn Res.* 2023;24(240):1–113.
47. Manyika J, Hsiao S. An overview of Bard: an early experiment with generative AI. *AI Google Static Doc.* 2023;2.
48. Deng J, Dong W, Socher R, Li LJ, Li K, Fei-Fei L. Imagenet: A large-scale hierarchical image database. In: *2009 IEEE conference on computer vision and pattern recognition*. Washington, DC: IEEE Computer Society. 2009. pp. 248–55.

49. 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. In: International Conference on Learning Representations. 2021. <https://openreview.net/forum?id=YicbFdNTTy>. Accessed 12 Jan 2021.
50. Radford A, Kim JW, Hallacy C, Ramesh A, Goh G, Agarwal S, et al. Learning transferable visual models from natural language supervision. In: International conference on machine learning. Somerville: Microtome Publishing. 2021. p. 8748–63.
51. Rombach R, Blattmann A, Lorenz D, Esser P, Ommer B. High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. New York: IEEE Corporate Headquarters. 2022. p. 10684–95.
52. Kirillov A, Mintun E, Ravi N, Mao H, Rolland C, Gustafson L, et al. Segment anything. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. 2023. p. 4015–26.
53. Wójcik MA. Foundation Models in Healthcare: Opportunities, Biases and Regulatory Prospects in Europe. In: International Conference on Electronic Government and the Information Systems Perspective. Heidelberg: Springer Nature. 2022. pp. 32–46.
54. Yu KH, Beam AL, Kohane IS. Artificial intelligence in healthcare. *Nat Biomed Eng*. 2018;2(10):719–31.
55. Bender EM, Gebru T, McMillan-Major A, Shmitchell S. On the dangers of stochastic parrots: Can language models be too big? In: Proceedings of the 2021 ACM conference on fairness, accountability, and transparency. New York: Association for Computing Machinery (ACM). 2021. p. 610–23.
56. Osman Andersen T, Nunes F, Wilcox L, Kazianus E, Matthiesen S, Magrabi F. Realizing AI in healthcare: challenges appearing in the wild. In: Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems. New York: Association for Computing Machinery (ACM). 2021. p. 1–5.
57. Liao QV, Zhang Y, Luss R, Doshi-Velez F, Dhurandhar A. Connecting algorithmic research and usage contexts: a perspective of contextualized evaluation for explainable AI. In: Proceedings of the AAAI Conference on Human Computation and Crowdsourcing. Association for the Advancement of Artificial Intelligence (AAAI), vol. 10. Washington, DC: Association for the Advancement of Artificial Intelligence (AAAI). 2022. p. 147–59.
58. Zając HD, Li D, Dai X, Carlsen JF, Kensing F, Andersen TO. Clinician-facing AI in the Wild: Taking Stock of the Socio-technical Challenges and Opportunities for HCI. *ACM Trans Comput Hum Interact*. 2023;30(2):1–39.
59. Li W, Milletari F, Xu D, Rieke N, Hancox J, Zhu W, et al. Privacy-preserving federated brain tumour segmentation. In: Machine Learning in Medical Imaging: 10th International Workshop, MLMI 2019, Held in Conjunction with MICCAI 2019, Shenzhen, China, October 13, 2019, Proceedings 10. Heidelberg: Springer Nature. 2019. pp. 133–41.
60. Kairouz P, McMahan HB, Avent B, Bellet A, Bennis M, Nitin Bhagoji A, et al. Advances and Open Problems in Federated Learning. *Found Trends Mach Learn*. 2021;14(1–2):1–210. <https://doi.org/10.1561/22000000083>.
61. Li T, Sahu AK, Talwalkar A, Smith V. Federated learning: Challenges, methods, and future directions. *IEEE Signal Process Mag*. 2020;37(3):50–60.
62. Karimireddy SP, Kale S, Mohri M, Reddi S, Stich S, Suresh AT. Scaffold: Stochastic controlled averaging for federated learning. In: International conference on machine learning. Somerville: Microtome Publishing. 2020. p. 5132–43.
63. Qu Z, Li X, Duan R, Liu Y, Tang B, Lu Z. Generalized federated learning via sharpness aware minimization. In: International Conference on Machine Learning. PMLR; 2022. pp. 18250–80.
64. Li X, Qu Z, Zhao S, Tang B, Lu Z, Liu Y. Lomar: A local defense against poisoning attack on federated learning. *IEEE Trans Dependable Secure Comput*. 2021;20(1):437–50.
65. Zhang J, Hua Y, Wang H, Song T, Xue Z, Ma R, et al. FedCP: Separating Feature Information for Personalized Federated Learning via Conditional Policy. In: Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. New York: Association for Computing Machinery (ACM); 2023.
66. Beutel DJ, Topal T, Mathur A, Qiu X, Fernandez-Marques J, Gao Y, et al. Flower: A Friendly Federated Learning Research Framework. arXiv preprint arXiv:2007.14390. 2020.
67. He C, Li S, So J, Zeng X, Zhang M, Wang H, et al. Fedml: A research library and benchmark for federated machine learning. arXiv preprint arXiv:2007.13518. 2020.
68. Liu Y, Fan T, Chen T, Xu Q, Yang Q. Fate: An industrial grade platform for collaborative learning with data protection. *J Mach Learn Res*. 2021;22(1):10320–5.
69. Xie Y, Wang Z, Gao D, Chen D, Yao L, Kuang W, et al. FederatedScope: A Flexible Federated Learning Platform for Heterogeneity. *Proc VLDB Endowment*. 2023;16(5):1059–72.
70. Bonawitz K, Eichner H, Grieskamp W, Huba D, Ingerman A, Ivanov V, et al. Towards Federated Learning at Scale: System Design. In: Talwalkar A, Smith V, Zaharia M, editors. Proceedings of Machine Learning and Systems. vol. 1. 2019. pp. 374–88.
71. Huba D, Nguyen J, Malik K, Zhu R, Rabbat M, Yousefpour A, et al. Papaya: Practical, private, and scalable federated learning. *Proc Mach Learn Syst*. 2022;4:814–32.
72. Hays K. Elon Musk's plan to charge for Twitter API access is unraveling — businessinsider.com. <https://www.businessinsider.com/elon-musk-plan-to-charge-for-twitter-api-access-unraveling-2023-5>. Accessed 23 Feb 2024.
73. Fung B. Reddit sparks outrage after a popular app developer said it wants him to pay \$20 million a year for data access | CNN Business — cnn.com. <https://www.cnn.com/2023/06/01/tech/reddit-outrage-data-access-charge/index.html>. Accessed 23 Feb 2024.
74. Nast C. Stack Overflow Will Charge AI Giants for Training Data — wired.com. <https://www.wired.com/story/stack-overflow-will-charge-ai-giants-for-training-data/>. Accessed 23 Feb 2024.
75. Hadsell R, Rao D, Rusu AA, Pascanu R. Embracing change: Continual learning in deep neural networks. *Trends Cogn Sci*. 2020;24(12):1028–40.
76. Li X, Tang B, Li H. AdaER: An adaptive experience replay approach for continual lifelong learning. *Neurocomputing*. 2024;572:127204.
77. Yoon J, Jeong W, Lee G, Yang E, Hwang SJ. Federated continual learning with weighted inter-client transfer. In: International Conference on Machine Learning. Somerville: Microtome Publishing. 2021. pp. 12073–86.
78. Li Y, Wang H, Jin Q, Hu J, Chemerys P, Fu Y, et al. Snapfusion: Text-to-image diffusion model on mobile devices within two seconds. *Adv Neural Inf Process Syst*. 2024;36.

79. Shome D, Kar T. FedAffect: Few-shot federated learning for facial expression recognition. In: Proceedings of the IEEE/CVF international conference on computer vision. New York: IEEE Corporate Headquarters. 2021. p. 4168–75.
80. Gao D, Yao X, Yang Q. A survey on heterogeneous federated learning. arXiv preprint arXiv:2210.04505. 2022.
81. Lyu L, Yu J, Nandakumar K, Li Y, Ma X, Jin J, et al. Towards fair and privacy-preserving federated deep models. *IEEE Trans Parallel Distrib Syst*. 2020;31(11):2524–41.
82. Chen C, Lyu L, Yu H, Chen G. Practical attribute reconstruction attack against federated learning. *IEEE Trans Big Data*. 2022;10(6):851–63.
83. Yu S, Muñoz JP, Jannesari A. Federated Foundation Models: Privacy-Preserving and Collaborative Learning for Large Models. arXiv preprint arXiv:2305.11414. 2023.
84. Chen HY, Tu CH, Li Z, Shen HW, Chao WL. On the importance and applicability of pre-training for federated learning. In: The Eleventh International Conference on Learning Representations. Appleton: International Conference on Learning Representations. 2022.
85. Tan Y, Long G, Ma J, Liu L, Zhou T, Jiang J. Federated learning from pre-trained models: A contrastive learning approach. *Adv Neural Inf Process Syst*. 2022;35:19332–44.
86. Zhang T, Feng T, Alam S, Dimitriadis D, Lee S, Zhang M, et al. Gpt-fl: Generative pre-trained model-assisted federated learning. arXiv preprint arXiv:2306.02210. 2023.
87. Li D, Wang J. Fedmd: Heterogeneous federated learning via model distillation. arXiv preprint arXiv:1910.03581. 2019.
88. Guo T, Guo S, Wang J, Tang X, Xu W. Promptfl: Let federated participants cooperatively learn prompts instead of models-federated learning in age of foundation model. *IEEE Trans Mob Comput*. 2023;23(5):5179–94.
89. Zhao H, Du W, Li F, Li P, Liu G. Reduce communication costs and preserve privacy: Prompt tuning method in federated learning. arXiv preprint arXiv:2208.12268. 2022;1(2).
90. Qiu C, Li X, Mummadi CK, Ganesh MR, Li Z, Peng L, et al. Federated Text-driven Prompt Generation for Vision-Language Models. In: The Twelfth International Conference on Learning Representations. 2024. <https://openreview.net/forum?id=NW31gAylm>. Accessed 16 Jan 2024.
91. Crawford K. The atlas of AI: Power, politics, and the planetary costs of artificial intelligence. New Haven: Yale University Press; 2021.
92. Thieme A, Hanratty M, Lyons M, Palacios J, Marques RF, Morrison C, et al. Designing human-centered AI for mental health: Developing clinically relevant applications for online CBT treatment. *ACM Trans Comput Hum Interact*. 2023;30(2):1–50.
93. Chen IY, Pierson E, Rose S, Joshi S, Ferryman K, Ghassemi M. Ethical machine learning in healthcare. *Ann Rev Biomed Data Sci*. 2021;4:123–44.
94. Swensen SJ, Kaplan GS, Meyer GS, Nelson EC, Hunt GC, Pryor DB, et al. Controlling healthcare costs by removing waste: what American doctors can do now. *BMJ Qual Saf*. 2011;20(6):534–7.
95. Van Hartskamp M, Consoli S, Verhaegh W, Petkovic M, Van de Stolpe A, et al. Artificial intelligence in clinical health care applications. *Interact J Med Res*. 2019;8(2):e12100.
96. Keehan SP, Cuckler GA, Poisal JA, Sisko AM, Smith SD, Madison AJ, et al. National Health Expenditure Projections, 2019–28: Expected Rebound In Prices Drives Rising Spending Growth: National health expenditure projections for the period 2019–2028. *Health Affairs*. 2020;39(4):704–14.
97. Easwaran S, Venugopal JP, Subramanian AAV, Sundaram G, Naseeba B. A comprehensive learning based swarm optimization approach for feature selection in gene expression data. *Heliyon*. 2024;10(17). <https://www.sciencedirect.com/science/article/pii/S2405844024131969>.
98. Korngiebel DM, Mooney SD. Considering the possibilities and pitfalls of Generative Pre-trained Transformer 3 (GPT-3) in healthcare delivery. *NPJ Digit Med*. 2021;4(1):93.
99. Krumholz HM, Terry SF, Waldstreicher J. Data acquisition, curation, and use for a continuously learning health system. *Jama*. 2016;316(16):1669–70.
100. Suresh A, Udendhran R, Vimal S. Deep neural networks for multimodal imaging and biomedical applications. Hershey: IGI Global; 2020.
101. Ionescu D. Deep learning algorithms and big health care data in clinical natural language processing. *Linguist Philos Investig*. 2020;19:86–92.
102. Ma'ayan A. Complex systems biology. *J R Soc Interface*. 2017;14(134):20170391.
103. Ramachandram D, Taylor GW. Deep multimodal learning: A survey on recent advances and trends. *IEEE Signal Process Mag*. 2017;34(6):96–108.
104. Hall DL, Llinas J. An introduction to multisensor data fusion. *Proc IEEE*. 1997;85(1):6–23.
105. Castanedo F, et al. A review of data fusion techniques. *Sci World J*. 2013;2013.
106. Li Y, Wu FX, Ngom A. A review on machine learning principles for multi-view biological data integration. *Brief Bioinforma*. 2018;19(2):325–40.
107. Stahlschmidt SR, Ulfenborg B, Synnergren J. Multimodal deep learning for biomedical data fusion: a review. *Brief Bioinforma*. 2022;23(2):bbab569.
108. Park C, Ha J, Park S. Prediction of Alzheimer's disease based on deep neural network by integrating gene expression and DNA methylation dataset. *Expert Syst Appl*. 2020;140:112873.
109. Peng C, Zheng Y, Huang DS. Capsule network based modeling of multi-omics data for discovery of breast cancer-related genes. *IEEE/ACM Trans Comput Biol Bioinforma*. 2019;17(5):1605–12.
110. Bichindaritz I, Liu G, Bartlett C. Integrative survival analysis of breast cancer with gene expression and DNA methylation data. *Bioinformatics*. 2021;37(17):2601–8.
111. Chaudhary K, Poirion OB, Lu L, Garmire LX. Deep learning-based multi-omics integration robustly predicts survival in liver cancer. *Clin Cancer Res*. 2018;24(6):1248–59.
112. Franco EF, Rana P, Cruz A, Calderon VV, Azevedo V, Ramos RT, et al. Performance comparison of deep learning autoencoders for cancer subtype detection using multi-omics data. *Cancers*. 2021;13(9):2013.
113. Islam MM, Huang S, Ajwad R, Chi C, Wang Y, Hu P. An integrative deep learning framework for classifying molecular subtypes of breast cancer. *Comput Struct Biotechnol J*. 2020;18:2185–99.

114. Albaradei S, Napolitano F, Thafar MA, Gojobori T, Essack M, Gao X. MetaCancer: A deep learning-based pan-cancer metastasis prediction model developed using multi-omics data. *Comput Struct Biotechnol J*. 2021;54(1):401–14.
115. Lee G, Nho K, Kang B, Sohn KA, Kim D. Predicting Alzheimer's disease progression using multi-modal deep learning approach. *Sci Rep*. 2019;9(1):1952.
116. Suk HI, Lee SW, Shen D, Initiative ADN, et al. Hierarchical feature representation and multimodal fusion with deep learning for AD/MCI diagnosis. *NeuroImage*. 2014;101:569–82.
117. Xu M, Ouyang L, Han L, Sun K, Yu T, Li Q, et al. Accurately differentiating between patients with COVID-19, patients with other viral infections, and healthy individuals: multimodal late fusion learning approach. *J Med Internet Res*. 2021;23(1):e25535.
118. Wang X, Liu M, Zhang Y, He S, Qin C, Li Y, et al. Deep fusion learning facilitates anatomical therapeutic chemical recognition in drug repurposing and discovery. *Brief Bioinforma*. 2021;22(6):bbab289.
119. Soto JT, Weston Hughes J, Sanchez PA, Perez M, Ouyang D, Ashley EA. Multimodal deep learning enhances diagnostic precision in left ventricular hypertrophy. *Eur Heart J Digit Health*. 2022;3(3):380–9.
120. Sun D, Wang M, Li A. A multimodal deep neural network for human breast cancer prognosis prediction by integrating multi-dimensional data. *IEEE/ACM Trans Comput Biol Bioinforma*. 2018;16(3):841–50.
121. Huang SC, Pareek A, Zamanian R, Banerjee I, Lungren MP. Multimodal fusion with deep neural networks for leveraging CT imaging and electronic health record: a case-study in pulmonary embolism detection. *Sci Rep*. 2020;10(1):22147.
122. Hanney SR, Castle-Clarke S, Grant J, Guthrie S, Henshall C, Mestre-Ferrandiz J, et al. How long does biomedical research take? Studying the time taken between biomedical and health research and its translation into products, policy, and practice. *Health Res Policy Syst*. 2015;13(1):1–18.
123. Wouters OJ, McKee M, Luyten J. Estimated research and development investment needed to bring a new medicine to market, 2009–2018. *Jama*. 2020;323(9):844–53.
124. Lalmuanawma S, Hussain J, Chhakchhuak L. Applications of machine learning and artificial intelligence for Covid-19 (SARS-CoV-2) pandemic: A review. *Chaos Solitons Fractals*. 2020;139:110059.
125. Kadurin A, Nikolenko S, Khrabrov K, Aliper A, Zhavoronkov A. druGAN: an advanced generative adversarial autoencoder model for de novo generation of new molecules with desired molecular properties in silico. *Mol Pharm*. 2017;14(9):3098–104.
126. Harrer S, Shah P, Antony B, Hu J. Artificial intelligence for clinical trial design. *Trends Pharmacol Sci*. 2019;40(8):577–91.
127. Lanckriet GR, De Bie T, Cristianini N, Jordan MI, Noble WS. A statistical framework for genomic data fusion. *Bioinformatics*. 2004;20(16):2626–35.
128. Aerts S, Lambrechts D, Maity S, Van Loo P, Coessens B, De Smet F, et al. Gene prioritization through genomic data fusion. *Nat Biotechnol*. 2006;24(5):537–44.
129. Kong J, Cooper LA, Wang F, Gutman DA, Gao J, Chisolm C, et al. Integrative, multimodal analysis of glioblastoma using TCGA molecular data, pathology images, and clinical outcomes. *IEEE Trans Biomed Eng*. 2011;58(12):3469–74.
130. Ribeiro RT, Marinho RT, Sanches JM. Classification and staging of chronic liver disease from multimodal data. *IEEE Trans Biomed Eng*. 2012;60(5):1336–44.
131. Wu KE, Yost KE, Chang HY, Zou J. BABEL enables cross-modality translation between multiomic profiles at single-cell resolution. *Proc Natl Acad Sci*. 2021;118(15):e2023070118.
132. Karthikeyan N, et al. A novel attention-based cross-modal transfer learning framework for predicting cardiovascular disease. *Comput Biol Med*. 2024;170:107977.
133. Lu K, Grover A, Abbeel P, Mordatch I. Frozen Pretrained Transformers as Universal Computation Engines. *Proc AAAI Conf Artif Intell*. 2022;36(7):7628–36. <https://doi.org/10.1609/aaai.v36i7.20729>.
134. Krishna K, Khosla S, Bigham JP, Lipton ZC. Generating SOAP Notes from Doctor-Patient Conversations Using Modular Summarization Techniques. In: *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*. vol. 1 (Long Papers). Stroudsburg: Association for Computational Linguistics (ACL). 2021. p. 4958–72.
135. Klasnja P, Pratt W. Healthcare in the pocket: mapping the space of mobile-phone health interventions. *J Biomed Inform*. 2012;45(1):184–98.
136. Zhu M, Ahuja A, Wei W, Reddy CK. A hierarchical attention retrieval model for healthcare question answering. In: *The World Wide Web Conference*. New York: Association for Computing Machinery (ACM). 2019. p. 2472–82.
137. Beck JT, Rammage M, Jackson GP, Preininger AM, Dankwa-Mullan I, Roebuck MC, et al. Artificial intelligence tool for optimizing eligibility screening for clinical trials in a large community cancer center. *JCO Clin Cancer Inform*. 2020;4:50–9.
138. Pichai S. An important next step on our AI journey — [blog.google](https://blog.google/intl/en-africa/products/explore-get-answers/an-important-next-step-on-our-ai-journey/). 2023. <https://blog.google/intl/en-africa/products/explore-get-answers/an-important-next-step-on-our-ai-journey/>. Accessed 06 Feb 2023.
139. Chen RJ, Lu MY, Chen TY, Williamson DF, Mahmood F. Synthetic data in machine learning for medicine and healthcare. *Nat Biomed Eng*. 2021;5(6):493–7.
140. Tunyasuvunakool K, Adler J, Wu Z, Green T, Zielinski M, Židek A, et al. Highly accurate protein structure prediction for the human proteome. *Nature*. 2021;596(7873):590–6.
141. Goenka SD, Gorzynski JE, Shafin K, Fisk DG, Pesout T, Jensen TD, et al. Accelerated identification of disease-causing variants with ultra-rapid nanopore genome sequencing. *Nat Biotechnol*. 2022;40(7):1035–41.
142. Chien I, Deliu N, Turner R, Weller A, Villar S, Kilbertus N. Multi-disciplinary fairness considerations in machine learning for clinical trials. In: *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. Heidelberg: Springer Nature. 2022. p. 906–24.
143. Zhuang W, Chen C, Lyu L. When foundation model meets federated learning: Motivations, challenges, and future directions. *arXiv preprint arXiv:2306.15546*. 2023.

144. Nguyen J, Wang J, Malik K, Sanjabi M, Rabbat M. Where to begin? on the impact of pre-training and initialization in federated learning. In: Workshop on Federated Learning: Recent Advances and New Challenges (in Conjunction with NeurIPS). 2022. <https://openreview.net/forum?id=zE9ctlWm5lx>.
145. Wu C, Wu F, Lyu L, Huang Y, Xie X. Communication-efficient federated learning via knowledge distillation. *Nat Commun*. 2022;13(1):2032.
146. Liao Y, Xu Y, Xu H, Yao Z, Wang L, Qiao C. Accelerating federated learning with data and model parallelism in edge computing. *IEEE/ACM Trans Network*. 2023;32(6):904–18.
147. Wang Z, Che B, Guo L, Du Y, Chen Y, Zhao J, et al. PipeFL: Hardware/software co-design of an FPGA accelerator for federated learning. *IEEE Access*. 2022;10:98649–61.
148. Yuan B, He Y, Davis J, Zhang T, Dao T, Chen B, et al. Decentralized training of foundation models in heterogeneous environments. *Adv Neural Inf Process Syst*. 2022;35:25464–77.
149. Houlsby N, Giurghi A, Jastrzebski S, Morrone B, De Laroussilhe Q, Gesmundo A, et al. Parameter-efficient transfer learning for NLP. In: International Conference on Machine Learning. Somerville: Microtome Publishing. 2019. pp. 2790–9.
150. Hu EJ, yelong shen, Wallis P, Allen-Zhu Z, Li Y, Wang S, et al. LoRA: Low-Rank Adaptation of Large Language Models. In: International Conference on Learning Representations. 2022. <https://openreview.net/forum?id=nZvKeeFYf9>. Accessed 28 Jan 2022.
151. Lester B, Al-Rfou R, Constant N. The Power of Scale for Parameter-Efficient Prompt Tuning. In: Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing. Stroudsburg: Association for Computational Linguistics (ACL). 2021. p. 3045–59.
152. Zhao H, Du W, Li F, Li P, Liu G. FedPrompt: Communication-Efficient and Privacy-Preserving Prompt Tuning in Federated Learning. In: ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). New York: IEEE Corporate Headquarters. 2023. pp. 1–5.
153. Hinton G. Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531. 2015.
154. Yang J, Shen X, Xing J, Tian X, Li H, Deng B, et al. Quantization networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2019. pp. 7308–16.
155. Blalock D, Gonzalez Ortiz JJ, Frankle J, Gutttag J. What is the state of neural network pruning? *Proc Mach Learn Syst*. 2020;2:129–46.
156. Song R, Zhou L, Lyu L, Festag A, Knoll A. Resfed: Communication efficient federated learning with deep compressed residuals. *IEEE Internet Things J*. 2023;11(6):9458–72.
157. Nagy B, Hegedűs I, Sándor N, Egedi B. Privacy-preserving Federated Learning and its application to natural language processing. *Knowl Based Syst*. 2023;264:109693.
158. Liu M, Ho S, Wang M, Gao L, Jin Y, Zhang H. Federated learning meets natural language processing: A survey. arXiv preprint arXiv:2107.12603. 2021.
159. Jothi Prakash V, Arul Antran Vijay S. A multi-aspect framework for explainable sentiment analysis. *Pattern Recogn Lett*. 2024;178:122–9.
160. Kim G, Yoo J, Kang S. Efficient Federated Learning with Pre-Trained Large Language Model Using Several Adapter Mechanisms. *MDPI*. 2024;11(21):4479.
161. Ye R, Wang W, Chai J, Li D, Li Z, Xu Y, et al. OpenFedLLM: Training Large Language Models on Decentralized Private Data via Federated Learning. In: Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining; 2024. p. 6137–47.
162. Weller O, Marone M, Braverman V, Lawrie D, Van Durme B. Pretrained Models for Multilingual Federated Learning. In: Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. Stroudsburg: Association for Computational Linguistics (ACL). 2022. p. 1413–21.
163. Guo T, Guo S, Wang J. Pfedprompt: Learning personalized prompt for vision-language models in federated learning. In: Proceedings of the ACM Web Conference 2023. New York: Association for Computing Machinery (ACM). 2023. p. 1364–74.
164. Peng Y, Bian J, Xu J. FedMM: Federated Multi-Modal Learning with Modality Heterogeneity in Computational Pathology. In: ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE. 2024. p. 1696-700.
165. Chen H, Zhang Y, Krompass D, Gu J, Tresp V. Feddat: An approach for foundation model finetuning in multi-modal heterogeneous federated learning. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 38. Washington, DC: Association for the Advancement of Artificial Intelligence (AAAI); 2024. p. 11285–93.
166. Shi J, Zheng S, Yin X, Lu Y, Xie Y, Qu Y. CLIP-Guided Federated Learning on Heterogeneity and Long-Tailed Data. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 38. Washington, DC: Association for the Advancement of Artificial Intelligence (AAAI); 2024. p. 14955–63.
167. Su S, Yang M, Li B, Xue X. Federated Adaptive Prompt Tuning for Multi-Domain Collaborative Learning. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 38. 2024. pp. 15117–25.
168. Lee KJ, Jeong B, Kim S, Kim D, Park D. General Commerce Intelligence: Glocally Federated NLP-Based Engine for Privacy-Preserving and Sustainable Personalized Services of Multi-Merchants. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 38. Washington, DC: Association for the Advancement of Artificial Intelligence (AAAI); 2024. p. 22752–60.
169. Ding N, Qin Y, Yang G, Wei F, Yang Z, Su Y, et al. Parameter-efficient fine-tuning of large-scale pre-trained language models. *Nat Mach Intell*. 2023;5(3):220–35.
170. Cho YJ, Manoel A, Joshi G, Sim R, Dimitriadis D. Heterogeneous Ensemble Knowledge Transfer for Training Large Models in Federated Learning. In: Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence (IJCAI) Main Track. Vienna: Center for Computer Science; 2022.
171. Liu R, Wu F, Wu C, Wang Y, Lyu L, Chen H, et al. No one left behind: Inclusive federated learning over heterogeneous devices. In: Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. New York: Association for Computing Machinery (ACM); 2022. p. 3398–406.

172. Ducange P, Marcelloni F, Renda A, Ruffini F. Federated Learning of XAI Models in Healthcare: A Case Study on Parkinson's Disease. *Cogn Comput*. 2024;16:1–26.
173. Ali S, Abuhmed T, El-Sappagh S, Muhammad K, Alonso-Moral JM, Confalonieri R, et al. Explainable Artificial Intelligence (XAI): What we know and what is left to attain Trustworthy Artificial Intelligence. *Inf Fusion*. 2023;99:101805.
174. Zielinski K, Kowalczyk N, Kocejko T, Mazur-Milecka M, Neumann T, Ruminski J. Federated Learning in Healthcare Industry: Mammography Case Study. In: 2023 IEEE International Conference on Industrial Technology (ICIT). New York: IEEE Corporate Headquarters. 2023. pp. 1–6.
175. Mondrejevski L, Miliou I, Montanino A, Pitts D, Hollmén J, Papapetrou P. FLICU: A Federated Learning Workflow for Intensive Care Unit Mortality Prediction. 2022. <https://arxiv.org/abs/2205.15104>. Accessed 313 Aug 2022.
176. Qiu X, Sun T, Xu Y, Shao Y, Dai N, Huang X. Pre-trained models for natural language processing: A survey. *Sci China Technol Sci*. 2020;63(10):1872–97.
177. Yuan L, Chen D, Chen YL, Codella N, Dai X, Gao J, et al. Florence: A new foundation model for computer vision. *arXiv preprint arXiv:2111.11432*. 2021.
178. Liu F, Zhu T, Wu X, Yang B, You C, Wang C, et al. A medical multimodal large language model for future pandemics. *NPJ Digit Med*. 2023;6(1):226.
179. Litjens G, Kooi T, Bejnordi BE, Setio AAA, Ciompi F, Ghafoorian M, et al. A survey on deep learning in medical image analysis. *Med Image Anal*. 2017;42:60–88.
180. Feng Z, Guo D, Tang D, Duan N, Feng X, Gong M, et al. Codebert: A pre-trained model for programming and natural languages. In: Findings of the Association for Computational Linguistics: EMNLP 2020; 2020. p. 1536–47.
181. Min S, Lyu X, Holtzman A, Artetxe M, Lewis M, Hajishirzi H, et al. Rethinking the role of demonstrations: What makes in-context learning work? *arXiv preprint arXiv:2202.12837*. 2022.
182. Gu Y, Tinn R, Cheng H, Lucas M, Usuyama N, Liu X, et al. Domain-specific language model pretraining for biomedical natural language processing. *ACM Trans Comput Healthc (HEALTH)*. 2021;3(1):1–23.
183. Petroni F, Rocktäschel 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). Stroudsburg: Association for Computational Linguistics (ACL); 2019. p. 2463–73.
184. Saxena S, Sangani R, Prasad S, Kumar S, Athale M, Awhad R, et al. Large-scale knowledge synthesis and complex information retrieval from biomedical documents. In: 2022 IEEE International Conference on Big Data (Big Data). New York: IEEE Corporate Headquarters. 2022. p. 2364–9.
185. Gupta V, Choudhary K, Tavazza F, Campbell C, Liao WK, Choudhary A, et al. Cross-property deep transfer learning framework for enhanced predictive analytics on small materials data. *Nat Commun*. 2021;12(1):6595.
186. Horiuchi D, Tatekawa H, Shimono T, Walston SL, Takita H, Matsushita S, et al. Accuracy of ChatGPT generated diagnosis from patient's medical history and imaging findings in neuroradiology cases. *Neuroradiology*. 2024;66(1):73–9.
187. Eichelberg M, Aden T, Riesmeier J, Dogac A, Laleci GB. A survey and analysis of electronic healthcare record standards. *Acm Comput Surv (CSUR)*. 2005;37(4):277–315.
188. Gonzalez-Hernandez G, Sarker A, O'Connor K, Savova G. Capturing the patient's perspective: a review of advances in natural language processing of health-related text. *Yearb Med Inform*. 2017;26(01):214–27.
189. Kalyan KS, Sangeetha S. SECNLP: A survey of embeddings in clinical natural language processing. *J Biomed Inform*. 2020;101:103323.
190. Solares JRA, Raimondi FED, Zhu Y, Rahimian F, Canoy D, Tran J, et al. Deep learning for electronic health records: A comparative review of multiple deep neural architectures. *J Biomed Inform*. 2020;101:103337.
191. Weng WH, Szolovits P. Representation learning for electronic health records. *arXiv preprint arXiv:1909.09248*. 2019.
192. Johnson AE, Pollard TJ, Shen L, Lehman LWH, Feng M, Ghassemi M, et al. MIMIC-III, a freely accessible critical care database. *Sci Data*. 2016;3(1):1–9.
193. 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.
194. Basaldella M, Liu F, Shareghi E, Collier N. COMETA: A corpus for medical entity linking in the social media. In: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP). 2020. p. 3122–37.
195. Müller M, Salathé M, Kummervold PE. Covid-twitter-bert: A natural language processing model to analyse covid-19 content on twitter. *Front Artif Intell*. 2023;6:1023281.
196. Johnson AE, Pollard TJ, Greenbaum NR, Lungren MP, Deng Cy, Peng Y, et al. MIMIC-CXR-JPG, a large publicly available database of labeled chest radiographs. *arXiv preprint arXiv:1901.07042*. 2019.
197. Ji Y, Zhou Z, Liu H, Davuluri RV. DNABERT: pre-trained Bidirectional Encoder Representations from Transformers model for DNA-language in genome. *Bioinformatics*. 2021;37(15):2112–20.
198. Satterthwaite TD, Connolly JJ, Ruparel K, Calkins ME, Jackson C, Elliott MA, et al. The Philadelphia Neurodevelopmental Cohort: A publicly available resource for the study of normal and abnormal brain development in youth. *Neuroimage*. 2016;124:1115–9.
199. Heinsfeld AS, Franco AR, Craddock RC, Buchweitz A, Meneguzzi F. Identification of autism spectrum disorder using deep learning and the ABIDE dataset. *NeuroImage Clin*. 2018;17:16–23.
200. Alfaro-Almagro F, Jenkinson M, Bangerter NK, Andersson JL, Griffanti L, Douaud G, et al. Image processing and Quality Control for the first 10,000 brain imaging datasets from UK Biobank. *Neuroimage*. 2018;166:400–24.
201. Digre A, Lindskog C. The human protein atlas-Integrated omics for single cell mapping of the human proteome. *Protein Sci*. 2023;32(2):e4562.
202. Frazee AC, Jaffe AE, Langmead B, Leek JT. Polyester: simulating RNA-seq datasets with differential transcript expression. *Bioinformatics*. 2015;31(17):2778–84.
203. Pelka O, Friedrich CM, García Seco de Herrera A, Müller H. Overview of the ImageCLEFmed 2020 concept prediction task: Medical image understanding. *CLEF2020 Working Notes*. 2020;2696.

204. Wang B, Shang L, Lioma C, Jiang X, Yang H, Liu Q, et al. On position embeddings in bert. In: International Conference on Learning Representations. Appleton: International Conference on Learning Representations; 2020.
205. Wang B, Zhao D, Lioma C, Li Q, Zhang P, Simonsen JG. Encoding word order in complex embeddings. In: International Conference on Learning Representations. Appleton: International Conference on Learning Representations; 2019.
206. Radford A, Narasimhan K. Improving Language Understanding by Generative Pre-Training. Appleton: International Conference on Learning Representations; 2018. <https://api.semanticscholar.org/CorpusID:49313245>.
207. Kenton JD, Chang MW, Toutanova LK. Bert: Pre-training of deep bidirectional transformers for language understanding. In: Proceedings of naacl-HLT. Minneapolis, Minnesota. 2018;1:2.
208. 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.
209. Peng Y, Yan S, Lu Z. Transfer learning in biomedical natural language processing: an evaluation of BERT and ELMO on ten benchmarking datasets. In: Proceedings of the 18th BioNLP Workshop and Shared Task. 2019. p. 58–65.
210. Beltagy I, Lo K, Cohan A. SciBERT: A Pretrained Language Model for Scientific Text. 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). Stroudsburg: Association for Computational Linguistics (ACL); 2019. p. 3615–20.
211. Kalyan KS, Rajasekharan A, Sangeetha S. AMMU: a survey of transformer-based biomedical pretrained language models. *J Biomed Inform*. 2022;126:103982.
212. Liu P, Yuan W, Fu J, Jiang Z, Hayashi H, Neubig G. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Comput Surv*. 2023;55(9):1–35.
213. Thieme A, Nori A, Ghassemi M, Bommasani R, Andersen TO, Luger E. Foundation Models in Healthcare: Opportunities, Risks & Strategies Forward. In: Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems. CHI EA '23. New York: Association for Computing Machinery; 2023. <https://doi.org/10.1145/3544549.3583177>.
214. Sang ETK, De Meulder F. Introduction to the CoNLL-2003 Shared Task: Language-Independent Named Entity Recognition. In: Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003. Stroudsburg: Association for Computational Linguistics (ACL). 2003. p. 142–7.
215. Bai J, Bai S, Chu Y, Cui Z, Dang K, Deng X, et al. Qwen technical report. arXiv preprint arXiv:2309.16609. 2023.
216. Huang K, Altaosaar J, Ranganath R. Clinicalbert: Modeling clinical notes and predicting hospital readmission. arXiv preprint arXiv:1904.05342. 2019.
217. Phan LN, Anibal JT, Tran H, Chanana S, Bahadroglu E, Peltekian A, et al. Scifive: a text-to-text transformer model for biomedical literature. arXiv preprint arXiv:2106.03598. 2021.
218. Monajatipoor M, Yang J, Stremmel J, Emami M, Mohaghegh F, Rouhsedaghat M, et al. LLMs in Biomedicine: A study on clinical Named Entity Recognition. arXiv preprint arXiv:2404.07376. 2024.
219. Wang G, Yang G, Du Z, Fan L, Li X. ClinicalGPT: large language models finetuned with diverse medical data and comprehensive evaluation. arXiv preprint arXiv:2306.09968. 2023.
220. Singhal K, Azizi S, Tu T, Mahdavi SS, Wei J, Chung HW, et al. Large language models encode clinical knowledge. *Nature*. Nature Publishing Group. 2023;620(7972):172–80.
221. Yunxiang L, Zihan L, Kai Z, Ruilong D, You Z. Chatdoctor: A medical chat model fine-tuned on llama model using medical domain knowledge. arXiv preprint arXiv:2303.14070. 2023;2(5):6.
222. Luo L, Ning J, Zhao Y, Wang Z, Ding Z, Chen P, et al. Taiyi: a bilingual fine-tuned large language model for diverse biomedical tasks. *J Am Med Inform Assoc*. 2024;ocae037. arxiv preprint paper
223. Vincent P, Larochelle H, Bengio Y, Manzagol PA. Extracting and composing robust features with denoising autoencoders. In: Proceedings of the 25th international conference on Machine learning. New York: Association for Computing Machinery (ACM); 2008. p. 1096–103.
224. Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, et al. Generative adversarial nets. *Adv Neural Inf Process Syst*. 2014;27.
225. Chen T, Kornblith S, Norouzi M, Hinton G. A simple framework for contrastive learning of visual representations. In: International conference on machine learning. Somerville: Microtome Publishing. 2020. p. 1597–607.
226. Grill JB, Strub F, Altché F, Tallec C, Richemond P, Buchatskaya E, et al. Bootstrap your own latent—a new approach to self-supervised learning. *Adv Neural Inf Process Syst*. 2020;33:21271–84.
227. He K, Fan H, Wu Y, Xie S, Girshick R. Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. New York: IEEE Corporate Headquarters; 2020. p. 9729–38.
228. Bao H, Dong L, Piao S, Wei F. BEiT: BERT Pre-Training of Image Transformers. In: International Conference on Learning Representations. Appleton: International Conference on Learning Representations. 2021.
229. He K, Chen X, Xie S, Li Y, Dollár P, Girshick R. Masked autoencoders are scalable vision learners. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. New York: IEEE Corporate Headquarters; 2022. pp. 16000–9.
230. Xie Z, Zhang Z, Cao Y, Lin Y, Bao J, Yao Z, et al. Simmim: A simple framework for masked image modeling. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. New York: IEEE Corporate Headquarters; 2022. p. 9653–63.
231. Gururangan S, Marasović A, Swayamdipta S, Lo K, Beltagy I, Downey D, et al. Don't Stop Pretraining: Adapt Language Models to Domains and Tasks. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics. Stroudsburg: Association for Computational Linguistics (ACL); 2020. p. 8342–60.
232. Rongali S, Jagannatha A, Rawat BPS, Yu H. Continual domain-tuning for pretrained language models. arXiv preprint arXiv:2004.02288. 2020.
233. Zhang R, Reddy RG, Sultan MA, Castelli V, Ferritto A, Florian R, et al. Multi-Stage Pre-training for Low-Resource Domain Adaptation. In: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP). Stroudsburg: Association for Computational Linguistics (ACL); 2020. p. 5461–8.

234. Saharia C, Chan W, Saxena S, Li L, Whang J, Denton EL, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Adv Neural Inf Process Syst.* 2022;35:36479–94.
235. Chambon P, Bluethgen C, Delbrouck JB, Van der Sluijs R, Polacin M, Chaves JMZ, et al. RoentGen: vision-language foundation model for chest x-ray generation. *arXiv preprint arXiv:2211.12737.* 2022.
236. Chambon PJM, Bluethgen C, Langlotz C, Chaudhari A. Adapting Pretrained Vision-Language Foundational Models to Medical Imaging Domains. In: *NeurIPS 2022 Foundation Models for Decision Making Workshop.* 2022. <https://openreview.net/forum?id=QtxbYdJVT8Q>. Accessed 04 Oct 2022.
237. Nguyen DK, Okatani T. Multi-task learning of hierarchical vision-language representation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition.* New York: IEEE Corporate Headquarters; 2019. p. 10492–501.
238. Gao Y, Liu J, Xu Z, Zhang J, Li K, Ji R, et al. Pyramidclip: Hierarchical feature alignment for vision-language model pretraining. *Adv Neural Inf Process Syst.* 2022;35:35959–70.
239. Ruiz N, Li Y, Jampani V, Pritch Y, Rubinstein M, Aberman K. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition.* New York: IEEE Corporate Headquarters; 2023. pp. 22500–10.
240. Irvin J, Rajpurkar P, Ko M, Yu Y, Ciurea-Ilicus S, Chute C, et al. Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In: *Proceedings of the AAAI conference on artificial intelligence*, vol. 33. Washington, DC: Association for the Advancement of Artificial Intelligence (AAAI); 2019. p. 590–7.
241. Doersch C. Tutorial on variational autoencoders. *arXiv preprint arXiv:1606.05908.* 2016;364.
242. Liu Y. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692.* 2019.
243. Subramanian S, Wang LL, Bogin B, Mehta S, van Zuylen M, Parasa S, et al. Medicat: A dataset of medical images, captions, and textual references. In: *Findings of the Association for Computational Linguistics: EMNLP 2020.* 2020. p. 2112–20.
244. Pelka O, Koitka S, Rückert J, Nensa F, Friedrich CM. Radiology objects in context (roco): a multimodal image dataset. In: *Intravascular Imaging and Computer Assisted Stenting and Large-Scale Annotation of Biomedical Data and Expert Label Synthesis: 7th Joint International Workshop, CVII-STENT 2018 and Third International Workshop, LABELS 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 16, 2018, Proceedings 3.* Heidelberg: Springer Nature. 2018. p. 180–9.
245. Degerli A, Ahishali M, Kiranyaz S, Chowdhury ME, Gabbouj M. Reliable covid-19 detection using chest x-ray images. In: *2021 IEEE International Conference on Image Processing (ICIP).* New York: IEEE Corporate Headquarters. 2021. p. 185–9.
246. Kumar N, Verma R, Sharma S, Bhargava S, Vahadane A, Sethi A. A dataset and a technique for generalized nuclear segmentation for computational pathology. *IEEE Trans Med Imaging.* 2017;36(7):1550–60.
247. Ghosh A, Acharya A, Jain R, Saha S, Chadha A, Sinha S. Clipsyntel: clip and llm synergy for multimodal question summarization in healthcare. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38. Washington, DC: Association for the Advancement of Artificial Intelligence (AAAI); 2024. p. 22031–9.
248. Boecking B, Usuyama N, Bannur S, Castro DC, Schwaighofer A, Hyland S, et al. Making the most of text semantics to improve biomedical vision–language processing. In: *European conference on computer vision.* Heidelberg: Springer Nature. 2022. p. 1–21.
249. Bustos A, Pertusa A, Salinas JM, De La Iglesia-Vaya M. Padchest: A large chest x-ray image dataset with multi-label annotated reports. *Med Image Anal.* 2020;66:101797.
250. Woo S, Debnath S, Hu R, Chen X, Liu Z, Kweon IS, et al. Convnext v2: Co-designing and scaling convnets with masked autoencoders. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition.* New York: IEEE Corporate Headquarters; 2023. p. 16133–42.
251. Zhang S, Xu Y, Usuyama N, Bagga J, Tinn R, Preston S, et al. Large-scale domain-specific pretraining for biomedical vision-language processing. *arXiv preprint arXiv:2303.00915.* 2023;2(3):6.
252. 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. *arXiv preprint arXiv:2303.00915.* 2023.
253. Zhang Y, Jiang H, Miura Y, Manning CD, Langlotz CP. Contrastive learning of medical visual representations from paired images and text. In: *Machine Learning for Healthcare Conference.* Somerville: Microtome Publishing. 2022. p. 2–25.
254. Huang SC, Shen L, Lungren MP, Yeung S. Gloria: A multimodal global-local representation learning framework for label-efficient medical image recognition. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision.* New York: IEEE Corporate Headquarters; 2021. p. 3942–51.
255. Tiu E, Talius E, Patel P, Langlotz CP, Ng AY, Rajpurkar P. Expert-level detection of pathologies from unannotated chest X-ray images via self-supervised learning. *Nat Biomed Eng.* 2022;6(12):1399–406.
256. Müller P, Kaissis G, Zou C, Rueckert D. Joint learning of localized representations from medical images and reports. In: *European Conference on Computer Vision.* Heidelberg: Springer Nature. 2022. p. 685–701.
257. Li Y, Wang H, Luo Y. A comparison of pre-trained vision-and-language models for multimodal representation learning across medical images and reports. In: *2020 IEEE international conference on bioinformatics and biomedicine (BIBM).* New York: IEEE Corporate Headquarters. 2020. p. 1999–2004.
258. Moon JH, Lee H, Shin W, Kim YH, Choi E. Multi-modal understanding and generation for medical images and text via vision-language pre-training. *IEEE J Biomed Health Inform.* 2022;26(12):6070–80.
259. Chen Z, Li G, Wan X. Align, reason and learn: Enhancing medical vision-and-language pre-training with knowledge. In: *Proceedings of the 30th ACM International Conference on Multimedia.* 2022. pp. 5152–61.
260. Li Z, Li Y, Li Q, Wang P, Guo D, Lu L, et al. Lvit: language meets vision transformer in medical image segmentation. *IEEE Trans Med Imaging.* 2023.
261. Paul A, Nayyar A, et al. A context-sensitive multi-tier deep learning framework for multimodal sentiment analysis. *Multimedia Tools Appl.* 2024;83(18):54249–78.
262. Wan Z, Liu C, Zhang M, Fu J, Wang B, Cheng S, et al. Med-unic: Unifying cross-lingual medical vision-language pre-training by diminishing bias. *Adv Neural Inf Process Syst.* 2024;36.

263. Prakash J, Vijay AAS. Cross-lingual Sentiment Analysis of Tamil Language Using a Multi-stage Deep Learning Architecture. *ACM Trans Asian Low Resour Lang Inf Process*. 2023;22(12):1–28.
264. Christensen M, Vukadinovic M, Yuan N, Ouyang D. Vision–language foundation model for echocardiogram interpretation. *Nat Med*. 2024;1–8. <https://www.nature.com/articles/s41591-024-02959-y#citeas>.
265. Li C, Wong C, Zhang S, Usuyama N, Liu H, Yang J, et al. Llava-med: Training a large language-and-vision assistant for biomedicine in one day. *Adv Neural Inf Process Syst*. 2024;36.
266. Act A. Health insurance portability and accountability act of 1996. *Public Law*. 1996;104:191.
267. Challen R, Denny J, Pitt M, Gompels L, Edwards T, Tsaneva-Atanasova K. Artificial intelligence, bias and clinical safety. *BMJ Qual Saf*. 2019;28(3):231–7.
268. Chamikara MAP, Bertok P, Khalil I, Liu D, Camtepe S. Privacy preserving distributed machine learning with federated learning. *Comput Commun*. 2021;171:112–25.
269. Wiens J, Saria S, Sendak M, Ghassemi M, Liu VX, Doshi-Velez F, et al. Do no harm: a roadmap for responsible machine learning for health care. *Nat Med*. 2019;25(9):1337–40.
270. Chen IY, Joshi S, Ghassemi M. Treating health disparities with artificial intelligence. *Nat Med*. 2020;26(1):16–7.
271. Kaushal A, Altman R, Langlotz C. Geographic distribution of US cohorts used to train deep learning algorithms. *Jama*. 2020;324(12):1212–3.
272. Zhao Q, Adeli E, Pohl KM. Training confounder-free deep learning models for medical applications. *Nat Commun*. 2020;11(1):6010.
273. Li Z, Hoiem D. Learning without forgetting. *IEEE Trans Pattern Anal Mach Intell*. 2017;40(12):2935–47.
274. Ruiz C, Zitnik M, Leskovec J. Identification of disease treatment mechanisms through the multiscale interactome. *Nat Commun*. 2021;12(1):1796.
275. Lavertu A, Altman RB. RedMed: Extending drug lexicons for social media applications. *J Biomed Inform*. 2019;99:103307.
276. Li I, Yasunaga M, Nuzumlalı MY, Caraballo C, Mahajan S, Krumholz H, et al. A neural topic-attention model for medical term abbreviation disambiguation. *arXiv preprint arXiv:1910.14076*. 2019.
277. Chaitanya K, Erdil E, Karani N, Konukoglu E. Contrastive learning of global and local features for medical image segmentation with limited annotations. In: Larochelle H, Ranzato M, Hadsell R, Balcan MF, Lin H, editors. *Advances in Neural Information Processing Systems*. vol. 33. Curran Associates, Inc.; 2020. pp. 12546–58. https://proceedings.neurips.cc/paper_files/paper/2020/file/9496886ecef4ee20a62d16b4a2d7ccca3-Paper.pdf. Accessed 16 Oct 2020.
278. Jumper J, Evans R, Pritzel A, Green T, Figurnov M, Tunyasuvunakool K, et al. High accuracy protein structure prediction using deep learning. *Fourteenth Crit Assessm Tech Protein Struct Prediction (Abstr Book)*. 2020;22(24):2.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.