

Multi-Modal Large Language Model-Guided Discovery of Gut Microbiota-Targeting Plant Polysaccharides for Precision Intervention in Metabolic Dysfunction-Associated Steatotic Liver Disease

Andries Heffman

Department of Computer Science, University of New Hampshire, Durham, NH, USA.
hoffman1996@unh.edu

Aarav Mathur

Department of Computer Science, Binghamton University, Binghamton, NY, USA.
aarav.mathur@binghamton.edu

Shandan He

Department of Computer Science, University of Central Florida, Orlando, FL, USA.
shandanwork@ucf.edu

Abstract

The rising global prevalence of metabolic dysfunction-associated steatotic liver disease (MASLD) demands innovative interventions that target its gut–liver pathophysiological axis. Plant polysaccharides have emerged as potent modulators of the gut microbiota, yet the vast chemical space and the complex multi-omics interactions render traditional discovery pipelines inefficient. This paper presents a system-level vision for multi-modal large language model (LLM)-guided discovery of gut microbiota-targeting plant polysaccharides as precision therapies for MASLD. We articulate a comprehensive architecture that integrates heterogeneous data streams—genomic, metabolomic, glycomic, clinical, and imaging data—into a unified multi-modal reasoning framework anchored by a generalist LLM. The system leverages cross-modal representation learning, knowledge graphs linking botanical sources to host–microbe metabolic networks, and in silico screening cascades to prioritize candidate polysaccharides. Beyond the computational engine, we examine the structural trade-offs between centralized and federated data governance models, the biases embedded in multi-modal training corpora that risk perpetuating inequitable interventions, and the sustainability costs of large-scale LLM training and inference. We discuss deployment pathways that span from cloud-based high-performance computing to edge inference within clinical and research settings, and we evaluate the robustness of such systems under data distribution shifts and adversarial perturbations. Policy dimensions are given sustained attention, particularly with respect to intellectual property of natural product sequences, regulatory approval of AI-discovered compounds, and international frameworks for equitable benefit sharing. The paper contends that responsible, system-scale design of multi-modal LLM architectures can not only accelerate the discovery of microbiota-directed polysaccharides but also recalibrate how AI is embedded in translational metabolic medicine, foregrounding fairness, environmental accountability, and resilient infrastructure. By grounding these arguments in concrete

biomedical challenges, we offer a blueprint for the next generation of AI-augmented natural product discovery systems.

Keywords

multi-modal large language models, gut microbiota, plant polysaccharides, metabolic dysfunction-associated steatotic liver disease, precision intervention, AI governance, sustainable infrastructure.

1. Introduction

Metabolic dysfunction-associated steatotic liver disease has rapidly become one of the most pressing metabolic epidemics of the twenty-first century, affecting approximately a quarter of the global adult population and driving an increasing burden of cirrhosis, hepatocellular carcinoma, and cardiovascular morbidity [1]. The liver's intimate bidirectional communication with the gut microbiota, often termed the gut–liver axis, is now recognized as a central pathogenic mechanism, wherein microbial dysbiosis alters bile acid profiles, short-chain fatty acid production, intestinal permeability, and systemic inflammatory cascades. Targeting the gut microbiota through dietary interventions and bioactive compounds represents a compelling strategy for MASLD management, yet the sheer dimensionality and inter-individual variability of host–microbiome interactions have outpaced the discovery capabilities of conventional hypothesis-driven research. Plant polysaccharides, a structurally diverse class of biopolymers, exhibit remarkable prebiotic and microbiota-reshaping properties, and recent studies have illuminated their potential to improve glycolipid metabolism through microbiome-mediated mechanisms [3]. However, navigating the vast combinatorial space of monosaccharide composition, glycosidic linkage patterns, molecular weight, and botanical origins to identify the most efficacious and safe candidates remains an open grand challenge.

Parallel to these biomedical roadblocks, the field of artificial intelligence has witnessed a profound transformation through large language models and their multi-modal extensions. Foundation models trained on internet-scale text, subsequently aligned with visual, chemical, and biological modalities, have demonstrated emergent capabilities in cross-domain reasoning, zero-shot inference, and the generation of scientifically plausible hypotheses [4,5,6]. In biomedicine, multi-modal architectures can now integrate pathology images, genomic sequences, protein structures, electronic health records, and molecular graphs within a shared representational space, opening the door to holistic system-level reasoning rather than siloed single-task prediction [7]. The application of such capabilities to natural product discovery, however, is still in its infancy, and systematizing their use for gut microbiota-targeting polysaccharides demands careful architectural design, robust data infrastructures, and anticipatory governance frameworks. This paper does not present empirical drug discovery results; instead, it develops an in-depth, system-level perspective on the design, deployment, and societal embedding of multi-modal LLM-guided platforms for this precision intervention domain. We analyze the structural trade-offs inherent in constructing multi-modal data fabrics, the epistemological risks of black-box recommendation pipelines, and the policy scaffolds required to ensure equitable, sustainable, and scientifically rigorous translation from computational prediction to clinical reality.

2. Metabolic Dysfunction-Associated Steatotic Liver Disease and the Gut Microbiota Axis

The pathophysiology of MASLD extends far beyond hepatic lipid accumulation, encompassing a multisystem disorder driven by insulin resistance, adipose tissue dysfunction, and a maladaptive gut environment. Seminal work established that the gut microbiota directly influences host energy harvest, intestinal permeability, and the production of endotoxins that fuel hepatic inflammation [2]. Subsequent multi-omic profiling has revealed that specific microbial consortia correlate with histological severity, with signatures involving reduced abundances of short-chain fatty acid producers such as *Faecalibacterium prausnitzii* and expansions of ethanol-producing *Enterobacteriaceae* contributing to steatohepatitis progression. This microbial ecosystem operates as a distributed metabolic organ, processing indigestible dietary fibers into bioactive metabolites, modulating bile acid receptor signaling, and educating mucosal immunity. Therapeutically intercepting this axis thus requires interventions that can selectively sculpt the microbiome's metabolic output without inducing broad ecological collapse, an objective that plant polysaccharides are uniquely positioned to achieve given their structural tunability.

Nevertheless, the path from microbial ecology to clinical outcome is nonlinear and encumbered by high-dimensional confounders. Host genetics, baseline diet, medication exposures, and circadian rhythms all modulate individual responses to a given polysaccharide. Conventional in vitro fermentation assays and animal models, while indispensable, offer limited scalability, and their findings often fail to replicate across human populations with heterogeneous genetic backgrounds and lifestyles. This creates an urgent need for computational systems that can simultaneously model the polysaccharide structure–microbiota interaction space, predict downstream host metabolic impacts, and stratify patient populations for precision application. Multi-modal AI platforms that ingest data across molecular, microbial, and clinical scales offer a route to deconvolve this complexity, provided they are designed with sufficient causal grounding and are embedded in ethically governed translational workflows.

3. Plant Polysaccharides as Gut Microbiota Modulators: A Multi-Dimensional Data Challenge

Plant polysaccharides are among the most structurally complex and information-dense molecules in nature. Their biological activity is determined by monosaccharide composition, anomeric configuration, branching patterns, degree of polymerization, charge density, and three-dimensional conformation, all of which can vary across plant species, tissue sources, and extraction protocols. Robust evidence demonstrates that specific polysaccharide fractions can preferentially promote the growth of beneficial *Bifidobacterium* and *Lactobacillus* species while simultaneously inhibiting opportunistic pathogens through competitive exclusion and immunomodulation [3]. A notable recent investigation into a polysaccharide extracted from *Phyllostachys nigra* demonstrated significant improvements in glycolipid metabolism and pronounced reshaping of the gut microbiome in a mouse model, underscoring the translational promise hidden within underexplored botanical taxa [12]. However, mapping the vast polysaccharide structural space onto the dynamic functional landscape of the gut microbiome requires a departure from reductionist structure–activity relationship frameworks toward integrative, data-driven system models.

The multi-dimensional data challenge encompasses several heterogeneous layers. Glycomic data, including liquid chromatography–mass spectrometry profiles and nuclear magnetic resonance spectra, encode structural fingerprints that are inherently high-dimensional and sparsely annotated. Microbiome data from metagenomic sequencing produce millions of

reads per sample, whose functional interpretation depends on reference databases that remain unevenly populated for non-Western populations. Host metabolomic readouts in plasma, liver, and feces add further dimensions capturing the physiological consequences of microbial remodeling. Integrating these layers demands computational architectures that can align disparate data modalities at biological resolution, learn joint latent representations that preserve causal structure, and reason over incomplete and noisy real-world evidence. Multi-modal LLMs, with their capacity for cross-modal alignment and in-context reasoning, present a promising but underexamined vehicle to tackle this integration challenge, provided that the inductive biases of training procedures are transparently evaluated.

4. Multi-Modal Large Language Models: Architectural Foundations for Biological Discovery

Large language models have evolved from unimodal text generators into multi-modal reasoning engines that can process and generate across images, chemical structures, genomic sequences, and structured knowledge bases. Architecturally, such systems typically comprise modality-specific encoders that project raw data into a shared latent space, a transformer-based core that performs cross-attention and sequence modeling, and decoder heads that output modality-aligned predictions or generate free-text explanations [6,7]. In the biomedical domain, models like BioGPT have demonstrated state-of-the-art performance in biomedical text mining and evidence synthesis, while generalist medical AI models like Med-PaLM exhibit the capacity to interpret clinical images and laboratory data within a unified reasoning process [5,7]. Extending this paradigm to polysaccharide discovery involves constructing modality-specific encoders for glycan representations, microbial taxonomic and functional profiles, host metabolic signatures, and botanical taxonomic metadata, and aligning them through contrastive learning objectives that pull functionally related samples together in the embedding space.

A critical design choice concerns whether to train a single monolithic multi-modal model or to orchestrate an ensemble of specialized sub-models coordinated by a reasoning LLM through tool-use and chain-of-thought protocols. The monolithic approach offers tighter integration and potentially deeper cross-modal emergence, but it demands enormous computational resources and is prone to catastrophic forgetting when new data modalities are added. The orchestrated approach, modeled after retrieval-augmented generation and tool-augmented LLMs, allows modular updating, easier regulatory scrutiny of individual components, and a more transparent reasoning chain, though it introduces interface brittleness and error propagation risks. Both paradigms must contend with the limited availability of high-quality paired multi-omics data for gut microbiota interventions, which forces reliance on semi-supervised domain adaptation and synthetic data augmentation strategies that carry their own validity risks. The trade-off between integration depth and system modularity thus becomes a central architectural theme whose resolution must be guided by the specific demands of biomedical translation, including reproducibility, interpretability, and safety.

5. System Architecture for LLM-Guided Polysaccharide Discovery

We envision a multi-layered system architecture organized into four interconnected tiers: data ingestion and harmonization, cross-modal knowledge representation, hypothesis generation and prioritization, and validation and feedback integration. The data ingestion tier federates access to heterogeneous sources, including public glycomic databases, curated gut microbiome resources such as MGnify and the Human Microbiome Project, botanical chemotaxonomic repositories, and clinical registries with de-identified MASLD patient

records. Ingested data are harmonized through large-scale entity resolution pipelines that align chemical identifiers, microbial taxonomy, and phenotypic ontologies, leveraging the LLM itself to resolve semantic ambiguities and to standardize unstructured text descriptions. The use of knowledge graphs plays a pivotal role in this tier; a graph that connects plant polysaccharide structures to specific microbial enzymatic degradation pathways and host signaling receptors enables structured reasoning that complements the LLM's statistical pattern recognition, referencing prior work on systematic biomedical knowledge integration [9].

The cross-modal knowledge representation tier trains modality-specific encoders and employs contrastive objectives to map different data views into a shared embedding space that supports semantic search, similarity-based retrieval, and the construction of polysaccharide-centric multi-omic fingerprints. A key innovation is the inclusion of a dynamic memory module that continuously ingests newly published experimental evidence, automatically extracts structured findings using the LLM, and updates the graph and embedding models, thus enabling the system to remain current with the rapidly expanding literature on gut microbiota and natural products. The hypothesis generation tier combines the shared representations with an LLM-based reasoning core that proposes candidate polysaccharides for specific MASLD sub-phenotypes. Reasoning unfolds through a multi-step process in which the model retrieves analogous cases from the embedding space, consults the knowledge graph for mechanistic plausibility, and generates a ranked list of candidates accompanied by natural language justifications that cite supporting evidence. Specialized LLM implementations for traditional Chinese medicine formula recommendations [14] have already demonstrated the feasibility of in-context reasoning over complex herbal compositions, and these approaches can be adapted to the structurally defined domain of purified polysaccharides. The top-ranked candidates then proceed through an *in silico* screening cascade that incorporates molecular docking simulations with gut microbial enzymes, predicted fermentation products with host metabolic models, and population stratification models that estimate efficacy variability across demographic groups, directly addressing fairness concerns early in the discovery pipeline.

The validation and feedback tier closes the loop by interfacing with experimental platforms—from *in vitro* continuous fermentation systems to gnotobiotic animal models—and absorbing their results to fine-tune predictions. Critically, the architecture supports an uncertainty calibration layer that quantifies epistemic and aleatoric uncertainties, enabling risk-informed go/no-go decisions for costly wet-lab validation. Such a feedback loop transforms the system from a static predictive engine into a continuously learning scientific partner, but it also introduces new governance challenges that we examine next. The entire pipeline is conceived as a socio-technical infrastructure in which human domain experts retain agency through interactive interfaces that allow them to interrogate, critique, and override model suggestions, preserving the necessary balance between algorithmic acceleration and scientific judgment.

6. Governance, Bias, and Fairness Considerations in AI-Driven Natural Product Research

Embedding a multi-modal LLM system in the discovery of medical interventions raises acute governance concerns that go well beyond generic AI ethics. First, the training corpora for biomedical LLMs are overwhelmingly derived from English-language sources, Western research institutions, and commercially funded studies, which can systematically underrepresent polysaccharides from traditional medical systems in Africa, South America,

and Asia, as well as microbiome reference data from non-European populations. Such representational skew risks channeling discovery efforts toward compounds and microbial targets already favored by dominant scientific infrastructures, thereby reproducing and amplifying global health inequities under the guise of algorithmic objectivity. Addressing this requires deliberate data curation policies that mandate proportional inclusion of globally diverse genomic, botanical, and clinical datasets, alongside fairness metrics that evaluate candidate recommendations across ancestry-specific virtual clinical cohorts [10].

Second, the LLM-guided discovery process complicates existing intellectual property frameworks. When a model proposes a polysaccharide that was previously described in Indigenous pharmacopeias but never formally characterized in the patent literature, questions of prior art, public domain status, and benefit sharing become legally and ethically fraught. System designers must therefore incorporate provenance tracing modules that track the knowledge sources contributing to each recommendation and that can flag culturally sensitive leads for community consultation before advancing them through proprietary development pipelines. Such modules can be aided by the same LLM’s ability to mine historical and ethnobotanical literature, transforming the AI from a pure discovery engine into a machine that also surfaces ethical meta-information. Regulatory agencies such as the US Food and Drug Administration and the European Medicines Agency are beginning to articulate frameworks for AI-generated drug candidates, yet they have not yet developed specific guidance for natural products whose molecular complexity and batch variability challenge conventional chemical entity definitions. Close collaboration between AI developers, natural product chemists, and international regulatory bodies is essential to craft approval pathways that are neither excessively permissive nor fossilize outdated standards that stymie innovation [11].

Finally, the opacity of multi-modal LLM reasoning poses unique risks in a biomedical context where clinical decisions can have life-altering consequences. Explainable AI techniques, including attention visualization, concept-based explanations, and counterfactual generation, must be woven into the system’s output layer so that domain experts can trace which structural motifs, microbial taxa, or metabolic pathways drove a particular recommendation [17]. Robustness to adversarial data manipulations, such as intentionally distorted polysaccharide annotations in public databases, must also be systematically evaluated, and ongoing monitoring for input distribution shifts—spatial, temporal, or demographic—must be operationally mandated. Governance frameworks should thereby impose continuous post-market surveillance and algorithmic auditing protocols analogous to pharmacovigilance, ensuring that the AI system remains safe and equitable as both the underlying data and the disease landscape evolve.

7. Infrastructure, Deployment, and Environmental Sustainability

The computational infrastructure required to train, fine-tune, and serve multi-modal LLMs at biomedical scale is anything but negligible. Training a single large transformer model can consume hundreds of megawatt-hours of electricity, with associated carbon emissions that, if not accounted for, turn a project aimed at improving metabolic health into an aggravating factor for environmental determinants of disease. Recent work has drawn attention to the energy and policy considerations of deep learning in natural language processing, and the same calculus applies with even greater force to multi-modal architectures that ingest high-resolution imaging and sequencing data [15]. Sustainability must therefore be designed into the system from the outset through strategies such as model distillation, sparse architectures,

mixture-of-experts deployments that activate only relevant parameters per modality, and geographically distributed training that exploits low-carbon energy grids. Federated learning paradigms can further reduce the need for centralized data lakes, instead allowing model updates to be computed locally at partner clinical sites and aggregated through privacy-preserving protocols that also lower data transfer overhead [16]. Such an approach not only cuts energy consumption but also aligns with data sovereignty requirements in multinational collaborations.

Deployment choices span from cloud-hosted high-performance computing clusters to edge inference on specialized accelerators within hospital data centers or field laboratories in low-resource settings. Cloud-native architectures offer elastic scaling and ease of maintenance but can introduce latency that hinders interactive hypothesis exploration and may become cost-prohibitive for large academic consortia. Edge deployment puts computational capability closer to the point of care, enabling real-time decision support during microbiome sampling and polysaccharide formulation development, yet it demands compact model formats and energy-efficient hardware, which in turn requires aggressive optimization that must not compromise accuracy in understudied populations. Hybrid architectures that centralize computationally intensive training while distributing lightweight inference models across a network of trusted nodes represent a pragmatic middle ground. The resilience of such infrastructure against hardware failures, connectivity disruptions, and cyberattacks must be a first-class design requirement, because interruptions in a discovery pipeline that is integrated into clinical trial logistics could lead to costly delays or compromised data integrity. Architects must therefore incorporate redundant data storage, model versioning, and disaster recovery protocols as integral system services rather than afterthoughts.

8. Policy Implications and Socio-Technical Integration

Translating an AI-discovered polysaccharide from *in silico* recommendation to regulatory approval and public health deployment necessitates a policy ecosystem that is currently fragmented. The discovery system must operate within a framework that defines the evidentiary standards for AI-generated leads, clarifies liability when an AI-recommended compound is later associated with adverse events, and establishes mechanisms for transparent reporting of development processes. The World Health Organization's ethics and governance guidelines for AI in health offer foundational principles, but they require domain-specific elaboration for natural product discovery where biological materials cross jurisdictional boundaries [19]. International agreements such as the Nagoya Protocol on access and benefit-sharing directly intersect with the operations of the system, because the digital sequence information of plant polysaccharides and the associated microbial genomes may fall under evolving legal definitions that have not yet been consistently interpreted in the context of machine learning. System designers should engage proactively with policy experts to embed compliance by design, for instance through metadata tags that record geographical provenance and consent status of all training data.

The role of AI in revolutions in healthcare and drug delivery [13] is increasingly recognized, but the focus has largely been on synthetic small molecules and biologics rather than structurally heterogeneous natural products. This gap creates both an opportunity and a responsibility. Policy incentives should encourage the creation of open-access, curated data commons that lower the barrier for multi-modal AI research while protecting the rights of data donors and communities, avoiding the replication of extractive relationships that have historically marred natural product bioprospecting. Additionally, workforce development

must keep pace so that clinical teams, regulatory reviewers, and research ethics committees possess sufficient literacy to critically evaluate LLM-generated evidence, a capacity that is not yet standard in medical and scientific education. The long-term vision is not an autonomous discovery pipeline but a collaborative human–AI intelligence infrastructure that is subject to democratic oversight, continuous value alignment, and public deliberation about the desirable directions of metabolic medicine.

9. Conclusion

Multi-modal large language model-guided discovery of gut microbiota-targeting plant polysaccharides offers a transformative paradigm for addressing the complex, multifactorial nature of metabolic dysfunction-associated steatotic liver disease. This paper has articulated a system-level architecture that harmonizes glycomic, microbiomic, and clinical data within a coherent reasoning framework, emphasizing the structural trade-offs between monolithic integration and modular orchestration, and between centralized computational power and federated sustainable infrastructure. We have foregrounded the governance, bias, and fairness challenges that arise when AI is superimposed on domains marked by deep global inequities and contested intellectual property landscapes, insisting that these must be treated as core system requirements rather than after-the-fact mitigations. Robustness, transparency, and environmental accountability were examined as inseparable pillars of responsible system design. Looking forward, the field must move from speculative frameworks to concrete implementations that undergo rigorous prospective evaluation, and the broader academic and policy communities must collaboratively construct the regulatory, ethical, and educational scaffolding that such implementations demand. When designed with deliberate care, multi-modal LLM systems have the potential to unlock the vast therapeutic potential of plant polysaccharides in a manner that is not only scientifically powerful but also socially just and ecologically sustainable.

References

1. Younossi, Z. M., Koenig, A. B., Abdelatif, D., Fazel, Y., Henry, L., & Wymer, M. (2019). Global epidemiology of nonalcoholic fatty liver disease—Meta-analytic assessment of prevalence, incidence, and outcomes. *Hepatology*, 69(6), 2672–2682. <https://doi.org/10.1002/hep.30251>
2. Cani, P. D., & Delzenne, N. M. (2009). The role of the gut microbiota in energy metabolism and metabolic disease. *Nature Reviews Endocrinology*, 5(6), 283–290. <https://doi.org/10.1038/nrendo.2009.57>
3. Chen, G., Chen, D., Zhou, W., Zhang, M., Huang, J., & Xie, M. (2022). Dietary polysaccharides as modulators of the gut microbiota: A review. *Critical Reviews in Food Science and Nutrition*, 62(23), 6422–6447. <https://doi.org/10.1080/10408398.2021.1908959>
4. Zhang, Y., Li, Y., Wu, J., Wang, Y., & Liu, Z. (2023). Large language models for drug discovery and development: A survey. *arXiv preprint arXiv:2306.13592*.
5. Luo, R., Sun, L., Xia, Y., Qin, T., Zhang, S., Poon, H., & Liu, T.-Y. (2022). BioGPT: Generative pre-trained transformer for biomedical text generation and mining. *Briefings in Bioinformatics*, 23(6), bbac409. <https://doi.org/10.1093/bib/bbac409>
6. Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., ... & Sutskever, I. (2021). Learning transferable visual models from natural language supervision.

Proceedings of the 38th International Conference on Machine Learning (pp. 8748–8763). PMLR.

7. Moor, M., Banerjee, O., Abad, Z. S. H., Krumholz, H. M., Leskovec, J., Topol, E. J., & Rajpurkar, P. (2023). Foundation models for generalist medical artificial intelligence. *Nature*, 616(7956), 259–265. <https://doi.org/10.1038/s41586-023-05881-4>
8. Ekins, S., Puhl, A. C., Zorn, K. M., Lane, T. R., Russo, D. P., Klein, J. J., ... & Clark, A. M. (2019). Exploiting machine learning for end-to-end drug discovery and development. *Nature Materials*, 18, 435–441. <https://doi.org/10.1038/s41563-019-0338-z>
9. Himmelstein, D. S., Lizee, A., Hessler, C., Brueggeman, L., Chen, S. L., Hadley, D., ... & Baranzini, S. E. (2017). Systematic integration of biomedical knowledge prioritizes drugs for repurposing. *eLife*, 6, e26726. <https://doi.org/10.7554/eLife.26726>
10. Chen, I. Y., Pierson, E., Rose, S., Joshi, S., Ferryman, K., & Ghassemi, M. (2021). Ethical machine learning in healthcare. *Annual Review of Biomedical Data Science*, 4, 123–144. <https://doi.org/10.1146/annurev-biodatasci-092820-114757>
11. Gerke, S., Minssen, T., & Cohen, G. (2020). Ethical and legal challenges of artificial intelligence-driven healthcare. In A. Bohr & K. Memarzadeh (Eds.), *Artificial Intelligence in Healthcare* (pp. 295–336). Academic Press.
12. Zhao K, et al. Phyllostachys nigra (Lodd ex Lindl) derived polysaccharide with enhanced glycolipid metabolism regulation and mice gut microbiome. *Int J Biol Macromol*. 2024;257:128588.
13. Liu, S. (2025). Revolutionizing healthcare: The role of artificial intelligence in drug discovery and delivery. *Journal of Angiotherapy*, 9(1), 1-16.
14. Li, X., He, H., Lu, G., Yue, P., Chen, J., Yang, Z., & Hon, C. (2025). TCM-DS: a large language model for intelligent traditional Chinese medicine edible herbal formulas recommendations. *Chinese Medicine*, 20(1), 191.
15. Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics* (pp. 3645–3650). ACL. <https://doi.org/10.18653/v1/P19-1355>
16. Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., ... & Cardoso, M. J. (2020). The future of digital health with federated learning. *NPJ Digital Medicine*, 3, 119. <https://doi.org/10.1038/s41746-020-00323-1>
17. Jiménez-Luna, J., Grisoni, F., & Schneider, G. (2021). Drug discovery with explainable artificial intelligence. *Nature Machine Intelligence*, 3, 573–584. <https://doi.org/10.1038/s42256-021-00336-1>
18. Sorbara, M. T., & Pamer, E. G. (2022). Microbiome-based therapeutics. *Nature Reviews Microbiology*, 20, 365–380. <https://doi.org/10.1038/s41579-021-00667-9>
19. World Health Organization. (2021). Ethics and governance of artificial intelligence for health. World Health Organization.
20. Topol, E. J. (2019). High-performance medicine: The convergence of human and artificial intelligence. *Nature Medicine*, 25, 44–56. <https://doi.org/10.1038/s41591-018-0300-7>