

ARTIFICIAL INTELLIGENCE IN BIOMEDICAL RESEARCH: FROM DATA INTEGRATION TO PRECISION MEDICINE

Anonymous authors
Paper under double-blind review

ABSTRACT

This comprehensive review examines the transformative role of artificial intelligence in biomedical research, from foundational data integration to clinical applications. The paper explores how AI techniques facilitate multimodal data fusion across diverse biological data types, employing both traditional statistical methods and advanced deep learning architectures including variational autoencoders, graph neural networks, and transformer models. It evaluates AI applications in medical imaging, where convolutional neural networks have achieved remarkable diagnostic accuracy (up to 94% in COVID-19 detection) while enhancing segmentation and classification tasks across multiple imaging modalities. The review further investigates generative AI's impact on molecular design and drug discovery, highlighting transformer-based architectures like TransAntivirus that navigate vast chemical spaces to optimize therapeutic candidates. Finally, it examines AI-enabled precision medicine applications, including Clinical Decision Support Systems and federated learning approaches that balance analytical power with privacy preservation. Despite significant progress, implementation challenges persist, including data heterogeneity, model explainability, and ethical concerns regarding bias and privacy. The paper underscores the importance of developing interpretable AI systems that integrate seamlessly into clinical workflows while addressing regulatory, ethical, and economic considerations to realize the full potential of AI in advancing biomedical research and healthcare delivery.

1 AI-DRIVEN MULTIMODAL DATA INTEGRATION IN BIOMEDICAL RESEARCH

1.1 TECHNIQUES AND APPLICATIONS OF MULTIMODAL DATA FUSION

The convergence of multiple high-dimensional biological data types has created unprecedented opportunities to comprehensively characterize complex biological systems, transcending the limitations of single-modality approaches. Biomedical data fusion strategies span a spectrum from classical statistical methods to advanced deep learning architectures, each offering distinct advantages for integrating heterogeneous data types. Traditional integration approaches include joint non-negative matrix factorization (jNMF), partial least square (PLS), canonical correlation analysis (CCA), and multiple kernel learning (MKL). Among these, sparse canonical correlation analysis (CCA) has demonstrated particular utility in identifying relationships between different data modalities while maintaining interpretability. For instance, sparse CCA analysis of laboratory results and radiomics features in COVID-19 patients revealed significant correlations ($\text{cor}(X_{u1}, Z_{v1}) = 0.596$), linking elevated lactate dehydrogenase and acute phase reactants with specific radiomic signatures that indicated increased entropy and heterogeneity in lung imaging features. This integration revealed distinct clinical phenotypes from multimodal data, underscoring the value of statistical approaches in uncovering biologically meaningful patterns across data types. Deep learning has emerged as a transformative paradigm for multimodal biomedical data integration, offering superior capacity to model complex nonlinear relationships both within and across modalities. Variational autoencoders (VAEs) stand at the forefront of generative modeling approaches, providing flexible designs that balance dimensionality reduction

with generative capabilities for applications such as data imputation, denoising, and joint embedding creation. Recent methodological innovations have enhanced VAE performance through various regularization strategies, including adversarial training, cycle-consistency, contrastive learning, and disentangled representation learning. Intermediate fusion strategies with joint representation learning have proven particularly effective for capturing the intricate regulatory dynamics of biological systems, as they enable gradual fusion of modalities at different depths within the architecture, more closely reflecting true biological relationships. These approaches excel at bridging the heterogeneity gap between diverse data types—such as imaging, genomics, and clinical records—by applying modality-specific network architectures before integration, thereby mimicking the multilevel reasoning process used in clinical diagnosis and prognosis. Graph neural networks (GNNs) have further expanded the integration toolkit, proving especially valuable for analyzing complex relationships in graph-structured biomedical data, while transformer architectures with attention mechanisms have revolutionized cross-modality interaction learning.

The practical applications of multimodal data integration span the entire spectrum of biomedical research and clinical practice. In oncology, the integration of radiological images with genomic profiles has enhanced prognostic prediction and patient stratification, while the fusion of pathology whole-slide images with genomic features using attention-based frameworks like MCAT and PORPOISE has improved biomarker discovery. For COVID-19 analysis, cooperative learning combining clinical, laboratory, radiomics, and viral genome sequencing data achieved superior prediction performance (AUC = 0.87) compared to unimodal approaches. Beyond molecular omics data, modern integration frameworks increasingly incorporate diverse modalities including histopathology slides, MRI and PET images, electronic health records, and biosignals from wearable devices. This broader integration enables more holistic views of biological processes and diseases, particularly in fields such as oncology and neurodegenerative disorders, where linking molecular mechanisms to clinical manifestations supports precise disease subtyping and targeted therapy development. Despite these advances, significant challenges remain, including data heterogeneity, missing modalities, limited cohort sizes, privacy concerns, and model interpretation[1]. Future directions point toward foundation models inspired by large language models, which aim to unify multimodal inputs, incorporate medical knowledge, and support reasoning across diverse biomedical data types.

2 ADVANCED AI APPLICATIONS IN BIOMEDICAL ANALYSIS

2.1 DEEP LEARNING IN MEDICAL IMAGING AND DIAGNOSTICS

Deep learning has emerged as a transformative approach to medical image analysis, significantly enhancing diagnostic capabilities across multiple healthcare domains. Convolutional Neural Networks (CNNs) have become the dominant architecture for medical image processing due to their exceptional ability to extract meaningful features from complex visual data. These networks excel particularly in 2D data analysis, demonstrating rapid learning capabilities and strong performance when provided with sufficient labeled data. The application of CNNs and specialized frameworks such as AlexNet, VGG, Inception, and ResNet has substantially improved the efficacy of human clinicians in medical image interpretation, enabling more accurate and efficient diagnoses.

Medical image classification has achieved remarkable accuracy through deep learning implementations. For instance, on the Image CLEF benchmark containing a 31-class image dataset, evolved techniques demonstrated an average classification accuracy of 88%, representing an improvement of up to 11% compared to previous state-of-the-art methods using identical datasets. In the context of COVID-19 diagnosis, the DeTraC (Decompose, Transfer, Compose) deep convolutional neural network achieved an impressive 94% accuracy with a true positive rate of 100% in identifying positive COVID-19 cases from chest X-ray images. Beyond classification, deep learning has revolutionized object localization and segmentation in medical imaging. For anatomical structure localization, ConvNets equipped with spatial pyramid pooling can analyze sagittal, axial, and coronal slices from three-dimensional images. These capabilities are essential for numerous computer-aided diagnostic applications spanning microscopy,

ultrasound, computed tomography, dermoscopy, magnetic resonance imaging, positron emission tomography, and X-ray modalities .

108 Deep learning has demonstrated remarkable versatility across diverse medical domains, particularly
109 in oncology. In gastric cancer research, models like convolutional neural networks (CNNs) and arti-
110 ficial neural networks (ANNs) have achieved significant praise in their application, revolutionizing
111 diagnostic approaches . For cancer detection and characterization, deep neural networks can
112 extract patterns from histopathology images to infer genomic features without re- quiring tumor
113 sequencing. Notable examples include Image2TMB, which predicts tumor mutation burden in
114 lung cancer with accuracy comparable to large panel sequencing but with significantly less
115 variance, and HE2RNA, which infers gene expression from histopathology images to deter-
116 mine microsatellite instability status in colorectal cancer . These approaches enable genomic
117 inference from routine histopathological slides, potentially eliminating the need for expensive
molecular testing in some clinical scenarios .

118 The integration of multimodal data through deep learning frameworks has opened new frontiers in
119 medical image analysis. Advanced models can combine clinical data (diagnostic test results, pathol-
120 ogy reports), medical images (histopathology, computed tomography), and various omics data (ge-
121 nomic, transcriptomic, and proteomic profiles) to generate comprehensive insights . Autoencoder
122 architectures, comprising encoders that create low-dimensional representation vectors and decoders
123 that reconstruct the original input, have proven particularly effective for mul- timodal learning.

124 This architecture forces the model to encapsulate meaningful features from the input,
125 demonstrating good generalizability and the unique ability to integrate different data modal-
126 ities into a single “end-to-end optimized” model . Recent innovations like the Segment Anything
127 Model (SAM) have further advanced medical image segmentation capabil- ities, demonstrating
128 superior performance compared to other interactive methods in non-iterative prompting settings,
129 with notably better overall performance even when not used in its optimal box- prompting mode .
130 Despite these promising advances, significant challenges remain in implementing deep learning for
131 medical imaging diagnostics. Data scarcity presents a fundamental limitation, as medical imag-
132 ing datasets are typically much smaller than those used for general computer vision problems .
133 The “black box” nature of many deep learning models raises legal and ethical concerns, as
134 healthcare professionals may be reluctant to rely on systems whose decision-making processes
135 cannot be fully explained . Additionally, the computational re- sources required to train complex
136 deep learning models can be prohibitively expensive . Addressing these challenges will be crucial
137 for realizing the full potential of deep learning in medical imaging diagnostics and facilitating its
138 integration into clinical workflows .

139 2.2 GENERATIVE AI FOR MOLECULAR DESIGN AND DRUG DISCOVERY

141 While deep learning has revolutionized medical imaging, similar transformative advances are oc-
142 curring in molecular design and drug discovery, where generative AI models are redefining tradi-
143 tional approaches. The integration of artificial intelligence methodologies into the drug development
144 pipeline has yielded meaningful enhancements in both efficiency and effectiveness, particularly with
145 the emergence of large language models and generative AI technologies [10]. These approaches are
146 particularly valuable for navigating the vast chemical space—estimated to contain 10⁶⁰ drug-like
147 molecules—where traditional virtual screening techniques cannot feasibly explore the entirety of
148 potential therapeutic candidates .

149 Generative models employing transformer-based architectures have demonstrated remarkable capa-
150 bilities in molecular design tasks. The TransAntivirus framework represents a novel data-driven
151 self-supervised pretraining generative model capable of performing select-and-replace edits to or-
152 ganic molecules, optimizing them for desired properties in antiviral drug candidate development .
153 This approach leverages the International Union of Pure and Applied Chemistry (IUPAC)
154 nomenclature rather than Simplified Molecular Input Line Entry System (SMILES), pro- viding
155 human-readable and easily editable molecular representations that more closely align with
156 chemists’ knowledge-based design practices . Such representation choice is partic- ularly
157 advantageous for analogue-based drug design, where modifications typically involve altering

functional groups rather than individual atoms, enabling more intuitive manipulation of molecular structures. Evaluations of TransAntivirus have demonstrated superior performance compared to control models across multiple metrics including novelty, validity, uniqueness, and diversity of generated compounds.

Recent advances in structure-based generative design have further enhanced the field by incorporating protein structural information into de novo molecule optimization. These approaches can be categorized based on whether they employ distribution learning or goal-directed optimization, and whether they explicitly or implicitly incorporate protein structure information into the generative model. Structure-based approaches aim to maximize predicted on-target binding affinity of generated molecules, thereby increasing the likelihood of identifying viable drug candidates with desired therapeutic properties. This integration of structural information represents a significant advancement over earlier generative models that relied solely on small-molecule information for training and conditioning de novo molecule generators.

The future of molecular design and drug discovery increasingly points toward AI agent systems capable of skeptical learning and reasoning. These “AI scientists” combine human creativity and expertise with artificial intelligence’s capacity to analyze large datasets, navigate hypothesis spaces, and execute repetitive tasks. Rather than replacing human researchers, these biomedical AI agents function collaboratively, integrating various AI models and biomedical tools with experimental platforms. Such systems employ large language models and generative models with structured memory capabilities for continual learning, while incorporating scientific knowledge, biological principles, and theories through specialized machine learning tools. The potential applications span from virtual cell simulation and programmable control of phenotypes to the design of cellular circuits and development of novel therapeutics. As these AI-driven approaches continue to mature, they promise to accelerate the traditionally lengthy and resource-intensive drug discovery process, potentially addressing urgent health challenges such as emerging viral threats through rapid identification and optimization of candidate molecules.

3 FROM DISCOVERY TO APPLICATION

3.1 AI-ENABLED PRECISION MEDICINE AND PERSONALIZED HEALTHCARE

The evolution from broad-spectrum therapeutic approaches to precision medicine represents one of the most significant paradigm shifts in modern healthcare, with artificial intelligence serving as a critical enabler of this transition. The integration of AI methodologies with multimodal clinical data has accelerated the development of tailored treatment protocols that account for individual patient variability in genes, environment, and lifestyle. The novel Drug Intelligence Science (DIS®) platform exemplifies this advancement by combining single-cell technology with AI and machine learning to gain high-resolution insights into cell biology, thereby facilitating the discovery of disease-relevant targets, high-quality drug candidates, and predictive biomarkers. This innovative approach provides unprecedented mechanistic understanding of human diseases and enables in-depth pharmacological profiling of drug candidates, significantly increasing the probability of success in drug development and therapeutic interventions.

In the clinical context, AI-powered Clinical Decision Support Systems (CDSSs) have demonstrated substantial value in personalizing patient care. These systems incorporate features such as risk level estimation, diagnosis recommendations, and tailored treatment suggestions, collectively contributing to more effective healthcare delivery. The integration of AI into clinical workflows has shown positive outcomes across various implementation models, with both electronic medical record-integrated and stand-alone CDSSs demonstrating benefits to healthcare providers. For instance, MilleDSS, an Italian CDSS, illustrates practical implementation with its four domains of general practitioner-software interaction covering clinical management, prescribing appropriateness, prevention strategies, and medical computerized stewardship. Despite these advances, the economic valuation of personalized medicine approaches remains complex. Studies suggest that while many personalized medicine tests are relatively cost-effective, fewer have been found to be cost-saving, and many available or emerging tests still require economic evaluation. This economic dimension underscores the

need for more evidence to inform decision-making and assessment of genomic priorities in health-care resource allocation .

The success of AI in precision medicine ultimately depends on clinician trust and adoption. Research confirms that both accuracy and understandability are crucial for fostering clinician trust in predictive CDSSs, with the degree of reliance on these systems within clinical workflows potentially influencing trust requirements . Addressing this challenge, recent advances in interpretable machine learning have enabled the development of models that not only provide predictions but also identify features that drive these predictions, as demonstrated in studies predicting physiological and perceived stress in pregnant women . Furthermore, AI models have proven valuable in optimizing clinical trials, patient selection, appropriate dosing regimens, and biomarker identification—all critical components of the personalized medicine ecosystem. These applications hold promise for streamlining clinical investigations and improving patient outcomes through more targeted therapeutic approaches.

The implementation of AI-enabled precision medicine also presents significant data privacy and ethical considerations. Federated learning has emerged as a promising solution to data-sharing challenges, allowing models to be trained across multiple institutions without centralizing sensitive patient data. This approach facilitates cross-institutional collaboration while preserving privacy, an essential consideration given that the development of AI algorithms typically requires extensive processing of big data in biobanks. Legal frameworks such as the EU’s General Data Protection Regulation (GDPR) provide guidance on ensuring compliance with data protection requirements when handling various categories of health data. Additionally, the Medical Device Regulation (EU 2017/745) stipulates that clinical evidence must be provided for any software intended for medical purposes, necessitating rigorous epidemiological studies to validate the effectiveness of AI systems in clinical practice.

3.2 EXPLAINABILITY, TRUST, AND IMPLEMENTATION CHALLENGES

Despite the transformative potential of AI in precision medicine, significant challenges persist in its widespread implementation. The deployment of these technologies necessitates careful attention to a complex development process that balances automation with human expertise. Deep learning models, while powerful, are not general-purpose AI systems but specialized tools that extract patterns from inputs and compute probabilities of class labels, requiring both representative training data and an understanding of their inherent limitations.. This technical reality underscores the importance of fostering intuitive understanding of these models among domain experts in areas such as health, education, and agriculture to facilitate effective translation to practice.. The potential for bias represents a particularly pressing concern in biomedical AI applications, with multiple sources of bias including insufficient data, sampling bias, and the use of health-irrelevant features or race-adjusted algorithms. Addressing these challenges requires sophisticated debiasing approaches that can be broadly categorized as distributional (data augmentation, perturbation, reweighting, and federated learning) or algorithmic (unsupervised representation learning, adversarial learning, disentangled representation learning, and causality-based methods). The rapidly growing scale and variety of biomedical data repositories further complicate implementation by raising important privacy concerns that conventional data-sharing frameworks inadequately address. Privacy-enhancing technologies (PETs) have emerged as promising solutions that safeguard sensitive data while enabling broader usage and analysis. These technologies facilitate data sharing across institutional boundaries through mechanisms that provide formal privacy guarantees, with statistical disclosure control (SDC) and differential privacy (DP) representing two dominant frameworks that address the fundamental statistical problem of balancing disclosure risk against data utility . Despite their different formulations, both approaches share core statistical challenges in designing optimal release mechanisms that satisfy bounds on disclosure risk while maximizing analytical utility . Beyond technical considerations, successful AI integration requires holistic attention to ethical and regulatory requirements. A comprehensive perspective on applications, opportunities, and challenges from a programmatic viewpoint is essential for ethical and sustainable implementation of AI solutions in medical contexts . This multifaceted approach ensures that

AI-based algorithms enhance outcomes, quality, and efficiency while respecting the complex social, ethical, and regulatory landscape in which they operate .

4 CONCLUSION

Artificial intelligence has fundamentally transformed biomedical research by enabling unprecedented integration and analysis of heterogeneous data modalities. This review demonstrates AI’s significant contributions across the biomedical spectrum—from multimodal data fusion using variational autoencoders and graph neural networks to enhanced diagnostic capabilities in medical imaging and accelerated therapeutic discovery through generative models. The practical implementations in precision medicine highlight AI’s potential to personalize healthcare delivery while improving clinical decision-making.

However, critical limitations persist that require focused attention. The ‘black box’ nature of many deep learning models undermines clinician trust and regulatory compliance, while data scarcity and quality issues compromise model generalizability. Privacy concerns and potential algorithmic biases further complicate clinical implementation. Economic evaluations of AI-enabled precision medicine approaches remain inadequate, complicating healthcare resource allocation decisions.

Future research should prioritize developing inherently explainable AI architectures that maintain high performance while providing interpretable insights. Federated learning and privacy-enhancing technologies deserve further investigation to enable collaborative model training without compromising data security. Additionally, standardized frameworks for evaluating AI systems’ economic impact and clinical utility are essential. Most importantly, the field must evolve toward collaborative human-AI systems that augment rather than replace clinical expertise, ensuring that technological advancement serves the fundamental goal of improving patient outcomes through more targeted, effective, and accessible healthcare interventions.

REFERENCES

- Ana R Bai~ao, Zhaoxiang Cai, Rebecca C Poulos, Phillip J Robinson, Roger R Reddel, Qing Zhong, Susana Vinga, and Emanuel Gonçalves. A technical review of multi-omics data integration methods: from classical statistical to deep generative approaches. *arXiv preprint arXiv:2501.17729*, 2025a.
- Ana R Bai~ao, Zhaoxiang Cai, Rebecca C Poulos, Phillip J Robinson, Roger R Reddel, Qing Zhong, Susana Vinga, and Emanuel Gonçalves. A technical review of multi-omics data integration methods: from classical statistical to deep generative approaches. *arXiv preprint arXiv:2501.17729*, 2025b.
- Hyunghoon Cho, David Froelicher, Natnatee Dokmai, Anupama Nandi, Shuvom Sadhuka, Matthew M Hong, and Bonnie Berger. Privacy-enhancing technologies in biomedical data science. *Annual review of biomedical data science*, 7(1):317–343, 2024.
- Iacopo Cricelli, Ettore Marconi, and Francesco Lapi. Clinical decision support system (cdss) in primary care: from pragmatic use to the best approach to assess their benefit/risk profile in clinical practice. *Current Medical Research and Opinion*, 38(5):827–829, 2022.
- Geoff Currie, K Elizabeth Hawk, Eric Rohren, Alanna Vial, and Ran Klein. Machine learning and deep learning in medical imaging: intelligent imaging. *Journal of medical imaging and radiation sciences*, 50(4):477–487, 2019.
- Ahmet Gorkem Er, Daisy Yi Ding, Berrin Er, Mertcan Uzun, Mehmet Cakmak, Christoph Sadee, Gamze Durhan, Mustafa Nasuh Ozmen, Mine Durusu Tanriover, Arzu Topeli, et al. Multimodal data fusion using sparse canonical correlation analysis and cooperative learning: a covid-19 cohort study. *NPJ digital medicine*, 7(1):117, 2024.
- Shanghai Gao, Ada Fang, Yepeng Huang, Valentina Giunchiglia, Ayush Noori, Jonathan Richard Schwarz, Yasha Ektefaie, Jovana Kondic, and Marinka Zitnik. Empowering biomedical discovery with ai agents. *Cell*, 187(22):6125–6151, 2024.

- 324 Klaudia Grechuta, Pedram Shokouh, Ahmad Alhusein, Dirk Müller-Wieland, Juliane Meyerhoff,
325 Jeremy Gilbert, Sneha Purushotham, Catherine Rolland, et al. Benefits of clinical decision support
326 systems for the management of noncommunicable chronic diseases: targeted literature review.
327 *Interactive Journal of Medical Research*, 13(1):e58036, 2024.
- 328 Hao Guan, Pew-Thian Yap, Andrea Bozoki, and Mingxia Liu. Federated learning for medical image
329 analysis: A survey. *Pattern recognition*, 151:110424, 2024.
- 330 Arnulf Jentzen, Benno Kuckuck, and Philippe von Wurtemberg. Mathematical introduction to
331 deep learning: Methods, implementations, and theory. *arXiv preprint arXiv:2310.20360*, 2023.
- 332 Junwei Liu, Xiaoping Cen, Chenxin Yi, Feng-ao Wang, Junxiang Ding, Jinyu Cheng, Qinhua Wu,
333 Baowen Gai, Yiwen Zhou, Ruikun He, et al. Challenges in ai-driven biomedical multimodal data
334 fusion and analysis. *Genomics, Proteomics & Bioinformatics*, 23(1):qzaf011, 2025.
- 335 Jiashun Mao, Jianmin Wang, Amir Zeb, Kwang-Hwi Cho, Haiyan Jin, Jongwan Kim, Onju Lee,
336 Yunyun Wang, and Kyoung Tai No. Transformer-based molecular generative model for antiviral
337 drug design. *Journal of chemical information and modeling*, 64(7):2733–2745, 2023.
- 338 Ada Ng, Boyang Wei, Jayalakshmi Jain, Erin A Ward, S Darius Tandon, Judith T Moskowitz, Sheila
339 Krogh-Jespersen, Lauren S Wakschlag, and Nabil Alshurafa. Predicting the next-day perceived
340 and physiological stress of pregnant women by using machine learning and explainability: algo-
341 rithm development and validation. *JMIR mHealth and uHealth*, 10(8):e33850, 2022.
- 342 Ajita Paliwal, Smita Jain, Sachin Kumar, Pranay Wal, Madhusmriti Khandai, Prasanna Shama
343 Khandige, Vandana Sadananda, Md Khalid Anwer, Monica Gulati, Tapan Behl, et al. Predictive
344 modelling in pharmacokinetics: from in-silico simulations to personalized medicine. *Expert
345 Opinion on Drug Metabolism & Toxicology*, 20(4):181–195, 2024.
- 346 Malhar Patel, Ittai Dayan, Elliot K Fishman, Mona Flores, Fiona J Gilbert, Michal Guindy, Eu-
347 gene J Koay, Michael Rosenthal, Holger R Roth, and Marius G Linguraru. Accelerating artificial
348 intelligence: How federated learning can protect privacy, facilitate collaboration, and improve
349 outcomes. *Health informatics journal*, 29(4):14604582231207744, 2023.
- 350 Kathryn A Phillips, Julie Ann Sakowski, Julia Trosman, Michael P Douglas, Su-Ying Liang, and
351 Peter Neumann. The economic value of personalized medicine tests: what we know and what we
352 need to know. *Genetics in Medicine*, 16(3):251–257, 2014.
- 353 Simon JD Prince. *Understanding deep learning*. MIT press, 2023.
- 354 Jessica M Schwartz, Maureen George, Sarah Collins Rossetti, Patricia C Dykes, Simon R Minshall,
355 Eugene Lucas, and Kenrick D Cato. Factors influencing clinician trust in predictive clinical
356 decision support systems for in-hospital deterioration: qualitative descriptive study. *JMIR Human
357 Factors*, 9(2):e33960, 2022.
- 358 Liang Schweizer. Drug intelligence science (dis®): Pioneering a high-resolution translational plat-
359 form to enhance the probability of success for drug discovery and development. *Drug Discovery
360 Today*, 28(11):103795, 2023.
- 361 Aleksandra Slavkovi'c and Jeremy Seeman. Statistical data privacy: A song of privacy and utility.
362 *Annual Review of Statistics and Its Application*, 10(1):189–218, 2023.
- 363 Sören Richard Stahlschmidt, Benjamin Ulfenborg, and Jane Synnergren. Multimodal deep learning
364 for biomedical data fusion: a review. *Briefings in bioinformatics*, 23(2):bbab569, 2022.
- 365 S Suganyadevi, V Seethalakshmi, and Krishnasamy Balasamy. A review on deep learning in medical
366 image analysis. *International Journal of Multimedia Information Retrieval*, 11(1):19–38, 2022.
- 367 Morgan Thomas, Andreas Bender, and Chris de Graaf. Integrating structure-based approaches in
368 generative molecular design. *Current Opinion in Structural Biology*, 79:102559, 2023.
- 369 Takis Vidalis. Artificial intelligence in biomedicine: a legal insight. *BioTech*, 10(3):15, 2021.
- 370
371
372
373
374
375
376
377

Yifan Yang, Mingquan Lin, Han Zhao, Yifan Peng, Furong Huang, and Zhiyong Lu. A survey of recent methods for addressing ai fairness and bias in biomedicine. *Journal of Biomedical Informatics*, 154:104646, 2024.

378 Chaoran Yu and Ernest Johann Helwig. Artificial intelligence in gastric cancer: A translational
379 narrative review. *Annals of translational medicine*, 9(3):269, 2021.

380 Emerging Research Directions in AI-Empowered Biomedical Research

381
382 Based on the provided paper content, I have identified three promising research directions that rep-
383 resent significant opportunities for advancing AI applications in biomedical research. Each direction
384 addresses critical gaps in current approaches while leveraging emerging technological capabilities.

385 1. Multimodal Foundation Models for Biomedical Research

386 What is the direction?

387
388 This research direction focuses on developing large-scale foundation models specifically designed
389 for biomedical applications that can seamlessly integrate and reason across heterogeneous biomed-
390 ical data types. Similar to how large language models revolutionized natural language processing,
391 biomedical foundation models would serve as versatile platforms capable of processing and gener-
392 ating insights from genomics, proteomics, imaging, clinical records, and other modalities simulta-
393 neously.

394 What are the innovations and challenges?

395
396 Innovations: - Unified representation learning frameworks that bridge the heterogeneity gap be-
397 tween diverse biomedical data types - Pre-training strategies that incorporate domain-specific med-
398 ical knowledge and relationships - Transfer learning capabilities that allow adaptation to multiple
399 downstream tasks with minimal fine-tuning - Integration of attention mechanisms to identify cross-
400 modality interactions and contextual relationships

401 Challenges: - Extreme data heterogeneity across modalities (e.g., structured EHR data vs. un-
402 structured imaging data) - Computational requirements for training models on massive multimodal
403 datasets - Missing modalities and incomplete data in real-world biomedical datasets - Difficulty in
404 establishing ground truth for complex multimodal relationships - Ensuring interpretability of model
405 predictions for clinical acceptance

406 Significance to existing research Current approaches to multimodal data integration in biomedicine
407 typically employ specialized models for specific modality combinations or tasks. Foundation models
408 would transform this landscape by offering a unified framework that can serve as a base for numerous
409 downstream applications. This would accelerate research across multiple fields simultaneously, from
410 disease subtyping to drug discovery, by enabling knowledge transfer across domains and reducing
411 the need for modality-specific model development.

412 Suggested research steps

413 1. Develop novel architecture designs that efficiently handle the unique characteristics of different
414 biomedical data types

415 2. Create specialized pre-training objectives that capture biological relationships across modalities

416 3. Establish benchmark datasets that span multiple biomedical modalities with well-defined ground
417 truth

418 4. Design modular components that allow for missing modality handling through data imputation
419 techniques

420 5. Implement interpretability mechanisms that provide biological insights alongside predictions

421 6. Validate model performance on specific downstream tasks like disease prediction and biomarker
422 discovery

423 7. Build model distillation techniques to create lightweight versions for clinical deployment

424 2. AI Agents for Molecular Design and Drug Discovery

425 What is the direction?

This direction aims to develop autonomous AI agent systems that combine multiple AI capabilities (generative models, reinforcement learning, reasoning systems) to actively participate in the drug discovery process. These systems would move beyond passive prediction to actively propose, test, and refine hypotheses about molecular designs with minimal human intervention, functioning as “AI scientists” that collaborate with human researchers.

What are the innovations and challenges?

Innovations: - Integration of skeptical learning and reasoning capabilities that allow agents to question and validate their own hypotheses - Closed-loop systems connecting in silico predictions with automated experimental platforms - Structured memory architectures enabling continual learning from successes and failures - Incorporation of scientific knowledge, biological principles, and theoretical constraints in agent reasoning processes - Multi-objective optimization capabilities that balance efficacy, toxicity, synthesizability, and other drug properties

Challenges: - Creating reliable simulation environments that accurately predict real-world molecular behavior - Developing frameworks for effective human-AI collaboration in hypothesis generation - Ensuring reproducibility and reliability of agent-designed experiments - Bridging the gap between computational predictions and wet-lab validation - Managing the vast hypothesis space efficiently while avoiding common chemical pitfalls

Significance to existing research Current generative models for drug discovery typically focus on passive molecule generation or optimization for specific properties in isolation. AI agent systems represent a paradigm shift by actively driving the discovery process through hypothesis generation, testing, and refinement. This approach could dramatically accelerate the traditionally lengthy drug discovery pipeline by navigating the vast chemical space more efficiently than either humans or traditional computational methods alone.

Suggested research steps

1. Develop integrated frameworks that combine generative chemistry models with reasoning capabilities
2. Create simulation environments that accurately reflect pharmacological principles and constraints
3. Design protocols for agent-driven hypothesis generation with built-in validation mechanisms
4. Implement feedback loops connecting computational predictions with experimental validation
5. Establish metrics to evaluate the novelty, diversity, and biological relevance of agent-proposed molecules
6. Build specialized knowledge bases that encode domain expertise in chemistry and biology
7. Validate the system by targeting specific therapeutic areas with unmet medical needs

3. Privacy-Preserving Collaborative AI for Precision Medicine

What is the direction?

This research direction focuses on developing AI frameworks that enable cross-institutional collaboration for precision medicine while maintaining strict privacy guarantees. These frameworks would leverage privacy-enhancing technologies such as federated learning, differential privacy, and secure multi-party computation to allow model training across distributed datasets without centralizing sensitive patient information.

What are the innovations and challenges?

Innovations: - Federated learning architectures optimized for heterogeneous biomedical data types - Differential privacy mechanisms that provide formal privacy guarantees while preserving data utility - Secure multi-party computation protocols for collaborative model training across institutions - Privacy-preserving techniques for multimodal data integration in clinical settings - Distributed validation frameworks for model performance assessment

Challenges: - Performance degradation in privacy-preserving settings compared to centralized approaches - Statistical challenges in balancing disclosure risk against analytical utility - Regulatory compliance across different jurisdictions and healthcare systems - Non-uniform data distributions

across institutions (data heterogeneity) - Computational overhead of secure protocols in resource-constrained clinical environments - Trust establishment among participating institutions

Significance to existing research While AI has demonstrated remarkable potential in precision medicine, its widespread adoption is limited by data siloing and privacy concerns. Privacy-preserving collaborative AI addresses this fundamental limitation by enabling model development on much larger and more diverse patient populations without compromising confidentiality. This approach could significantly accelerate the translation of AI advances into clinical practice by facilitating evidence generation at scale while respecting ethical and regulatory requirements.

Suggested research steps

1. Develop federated learning algorithms specifically optimized for heterogeneous clinical data
2. Design privacy budgeting frameworks that maximize utility while maintaining privacy guarantees
3. Create benchmark datasets and evaluation metrics for privacy-preserving biomedical AI
4. Implement efficient secure computation protocols suitable for resource-constrained environments
5. Establish governance frameworks for multi-institutional collaboration
6. Validate approaches on real-world use cases such as rare disease diagnosis or treatment response prediction
7. Develop tools to quantify and communicate privacy-utility tradeoffs to stakeholders

In summary, these three research directions represent complementary approaches to advancing AI in biomedical research. Multimodal foundation models offer the broadest potential impact by creating versatile platforms for diverse applications. AI agents for drug discovery could revolutionize therapeutic development but require more time to mature and validate. Privacy-preserving collaborative AI addresses immediate barriers to clinical implementation and could see faster adoption in practice. Together, these directions address the key challenges identified in the paper while leveraging emerging AI capabilities to transform biomedical research and healthcare delivery.

486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539