# Papers to Cover


1. Wasif Khan, Seowung Leem, Kyle B. See, Joshua K. Wong, Shaoting Zhang, Ruogu Fang: “A Comprehensive Survey of Foundation Models in Medicine”, 2024; [http://arxiv.org/abs/2406.10729 arXiv:2406.10729].
2. He Y, Huang F, Jiang X, Nie Y, Wang M, Wang J, Chen H. Foundation Model for Advancing Healthcare: Challenges, Opportunities and Future Directions. IEEE Rev Biomed Eng. 2024 Nov 12;PP. doi: 10.1109/RBME.2024.3496744. Epub ahead of print. PMID: 39531565.
3. Guo, L.L., Fries, J., Steinberg, E. et al. A multi-center study on the adaptability of a shared foundation model for electronic health records. npj Digit. Med. 7, 171 (2024). https://doi.org/10.1038/s41746-024-01166-w

# Foundation Models Driving Healthcare Advancements

Joseph Moretto (jam5cq), Jing-Ning Su (fzf9mg), Jingyi Cui (cau8rc), Sheharyar Khalid (fsr5wf)

GenAI-overview

January 24th, 2025

# What is a foundation model?

- ➢ Very large deep learning models
  - Massive broad datasets
- ➢ What is the purpose?
  - Acts a foundation to build upon
  - Can quickly be adapted for a task
  - Cost effective

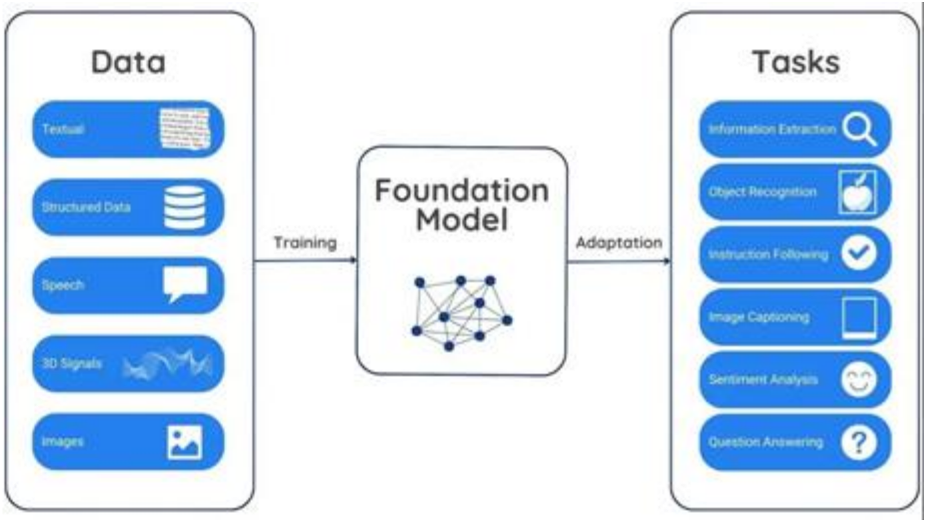


https://www.google.com/url?sa=i&url=https%3A%2F%2Fviso.ai%2Fdeep-learning%2Ffoundation-models%2F&psig=AOvVaw0P4gSt45xU0C3hXZwQQbYd&ust=1737938482424000&source=images&cd=vfe&opi=89978449&ved=0CBQQjRxqFwoTCKjmvsCTkosDFQAAAAAdAAAAABBF

# AI in Healthcare

- AI in healthcare can be used for several purposes
  - Diagnosis
  - Image analysis
  - Predictive
  - Personalization
- How do foundation models relate to AI in healthcare?

https://www.google.com/url?sa=i&url=https%3A%2F%2Fwww.news-medical.net%2Fnews%2F20231219%2FAI-in-healthcare-A-double-edged-sword-Study-reveals-impact-on-diagnostic-accuracy.aspx&psig=AOvVaw0WASwxstHW7nfa7QFLwWt3&ust=1737938699915000&source=images

# Foundation Models (FM)

- ➢ FMs are large-scale, pre-trained models fine-tuned for various downstream tasks, leveraging extensive training datasets.
- ➢ FMs utilize self-supervised learning to autonomously generate pre-training tasks from unlabeled data.
- ➢ FMs are versatile and can be applied to various fields including text, video, speech, and tabular data.

# Introduction and Outline

Joseph Moretto (jam5cq)

# Outline

- Introduction (Joseph)
  - Importance of AI in healthcare
  - Brief overview of foundational models and what is being presented
- Overview of foundation models (Mati)
  - What is a foundation model
  - Type of foundation models
- Applications in healthcare (Jing-Ning)
  - Examples of how they are used in healthcare
  - Advancements, include specific examples/ real world uses
- AI challenges in healthcare, include specific examples/ real world cases (Jingyi)
  - Data (getting good data, privacy)
  - Algorithmic
  - Resources cost/infrastructure
- Future work/directions (Sheharyar)
- Conclusion (Sheharyar)
- Discussion(Joseph)

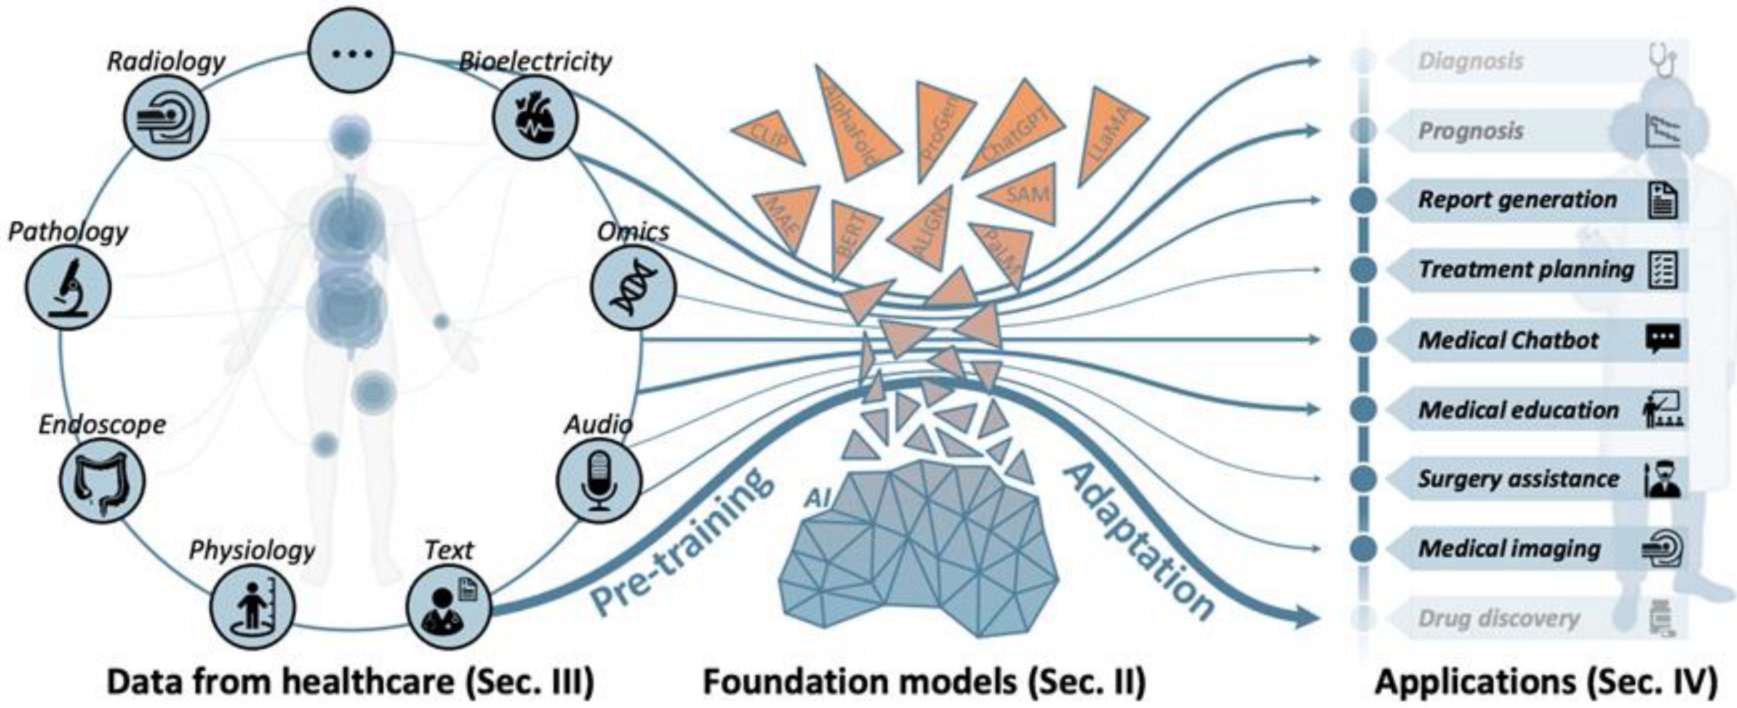


Fig. 1. The pipeline of the healthcare foundation models (HFMs) including the methods (Sec.II), datasets (Sec.III), and applications (Sec.IV).

# Foundation models driving AI in healthcare

- AI has achieved specialist level performance
- Many diverse scenarios and requirements
- What do foundation models add?
  - Cost effective and fast
- Efficient at creating models for a variety of healthcare tasks
- AI in healthcare is moving away from specific into general purpose

# FMs in Medical Domain

- **General-Purpose FMs:** Challenging to adopt in healthcare due to specialized needs.
- **Text-Based Models:** Word2Vec, ELMo, and BERT perform poorly on biomedical texts due to differing word distributions.
- **Vision-Text Pre-Training:** Models like CLIP struggle with nuanced differences in medical versus general images.
- **Segment Anything Model (SAM):** Ineffective for 3D medical images with its 2D design.
- **Task-Agnostic Model Development:** New models inspired by BERT and GPT are tailored for medical data.

# Clinical Large Language Models (CLLM)

- **BioBERT:** Pioneering biomedical NLP model, excelling in NER, RE, and QA tasks across 15 datasets.
- **BioMegatron:** Enhanced with up to 1.2 billion parameters; superior performance in biomedical benchmarks, trained on extensive PubMed and PMC data.
- **GatorTron:** A clinical NLP powerhouse with 8.9 billion parameters, trained on a vast corpus including UF Health clinical texts, achieving top results in five clinical NLP tasks.
- **GatorTronGPT:** GPT-3 architecture adaptation with 20 billion parameters, generating medically aligned synthetic data and performing comparably to human clinicians in NER, QA, RE, NLI, and semantic tasks.

# Medical Image Analysis

- **Distinct Characteristics:** Medical images feature unique patterns and characteristics compared to natural images, necessitating specialized models.
- **Segment Anything Model (SAM):** Originally designed for general-purpose segmentation, SAM was later adapted for medical imaging tasks.
- **MedSAM Architecture:** Utilizes a vision transformer-based (ViT) image encoder and a mask decoder, specifically tailored for medical image segmentation.
- **Training and Performance:** MedSAM was trained on 1,570,263 image-text pairs from online medical datasets, demonstrating superior performance compared to models like SAM and U-Net.

# CLIP-based FMs

- **CLIP Architecture:** CLIP is a neural network that classifies images using natural language, trained on a dataset of 400 million image-text pairs.
- **Medical Limitations:** In medical applications, CLIP struggles with smaller datasets, often misclassifying similar images.
- **MedCLIP Adaptation:** MedCLIP modifies CLIP for medical use by separating image-text inputs and adding a medical-specific semantic loss, improving data efficiency.

# Text-to-Image

- ➢ Text-to-image models are applied to generate text-conditional MRI scans, demonstrating their utility in medical imaging.
- ➢ MedXChat, a new text-to-image model, excels in synthesizing accurate x-ray images and medical reports, surpassing existing models in adaptability and precision.
- ➢ These technologies improve health literacy and comprehension of medical texts, enabling more effective diagnosis, reducing the need for repeated scans, and minimizing radiation exposure for patients.

# Omics

- ➢ Traditional FMs fails to generalize across genomics, proteomics, metabolomics, and other omics domains.
- ➢ Specialized BERT models like scBERT were developed for gene-level analysis, RNABERT for RNA sequence alignment, and DNABERT for decoding non-coding DNA, enhancing scalability and accuracy.
- ➢ Ongoing issues include the necessity for curated datasets, high computational costs, and model limitations in capturing comprehensive genomic information, pointing to areas for future enhancement.

# Applications in Healthcare

Jing-Ning Su (fzf9mg)

https://doi.org/10.1038/s41746-024-01166-w

## A multi-center study on the adaptability of a shared foundation model for electronic health records

Check for updates

Lin Lawrence Guo[1,7], Jason Fries[2,7], Ethan Steinberg[2], Scott Lanyon Fleming[2], Keith Morse[3], Catherine Aftandilian[4], Jose Posada[5], Nigam Shah[2,8] & Lillian Sung[1,6,8]

# Applications in Healthcare

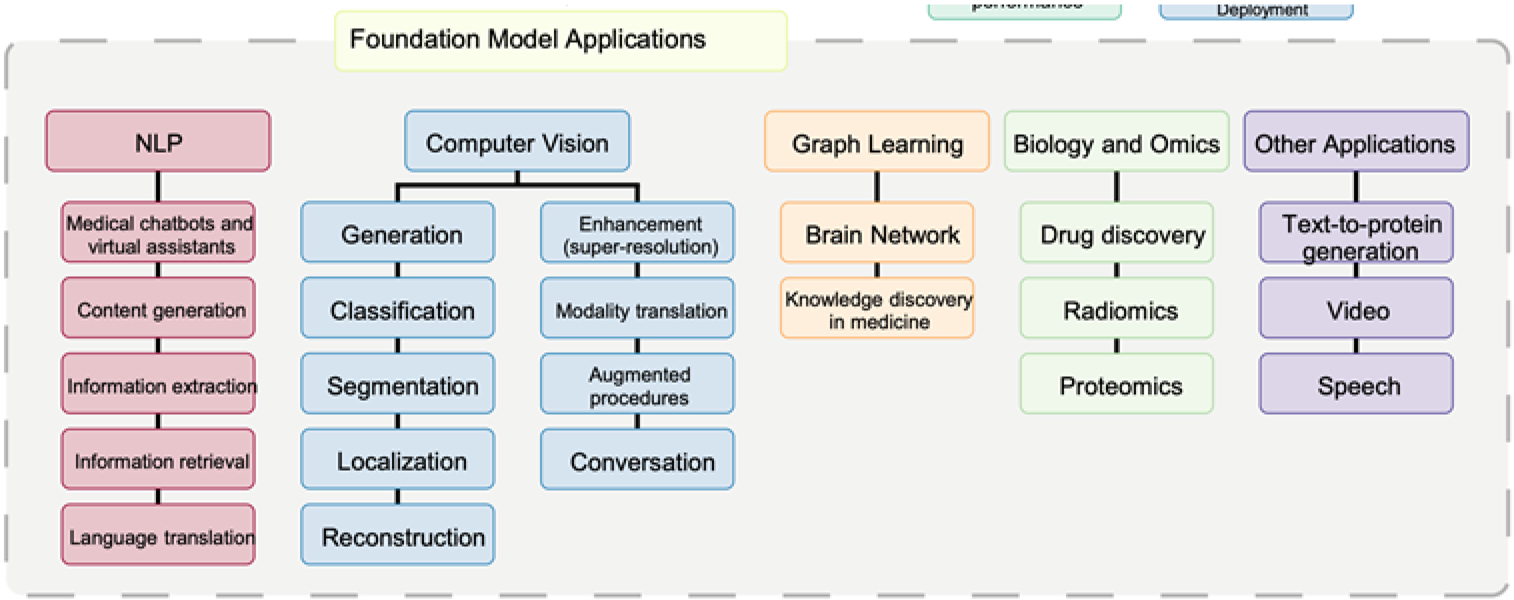


References:
[1] A Comprehensive Survey of Foundation Models in Medicine

# Applications in Healthcare

- **Clinical Natural Language Processing (NLP)**
  - Virtual Assistants: Answer patient queries, provide health advice
  - Report Generation: Automate medical report creation
  - Information Extraction: Extract key insights from EHRs
- **Medical Imaging**
  - Disease Detection: Identify tumors from medical images
  - Image Generation: Create images from text descriptions
  - Educational Tools: Visual aids for patient education


References:
[1] A Comprehensive Survey of Foundation Models in Medicine
[2] Foundation Model for Advancing Healthcare
[3] A Multi-center Study on the Adaptability of a Shared Foundation Model for Electronic Health Records

# Applications in Healthcare

- **Electronic Health Records (EHR)**
  - Risk Prediction: Predict patient risks like readmission
  - Summarization: Summarize patient histories
  - Decision Support: Aid in treatment decisions
- **Bioinformatics and Genomics**
  - Genetic Analysis: Identify disease markers in genetic sequences
  - Drug Discovery: Design new drugs, predict protein structures
  - Personalized Medicine: Tailor treatments based on genetic profiles

References:
[1] A Comprehensive Survey of Foundation Models in Medicine
[2] Foundation Model for Advancing Healthcare
[3] A Multi-center Study on the Adaptability of a Shared Foundation Model for Electronic Health Records

# Applications in Healthcare

- **Multimodal Integration**
  - Data Fusion: Combine text, images, and genetic data
  - Enhanced Diagnostics: Provide a comprehensive understanding of conditions
  - Personalized Treatment: Develop more effective treatment plans
- **Healthcare Service Optimization**
  - Patient Prioritization: Triage patients based on urgency
  - Automated Reporting: Streamline healthcare processes
  - Efficiency Improvement: Reduce administrative burden

References:
[1] A Comprehensive Survey of Foundation Models in Medicine
[2] Foundation Model for Advancing Healthcare
[3] A Multi-center Study on the Adaptability of a Shared Foundation Model for Electronic Health Records

# Overview of Model Training and Evaluation

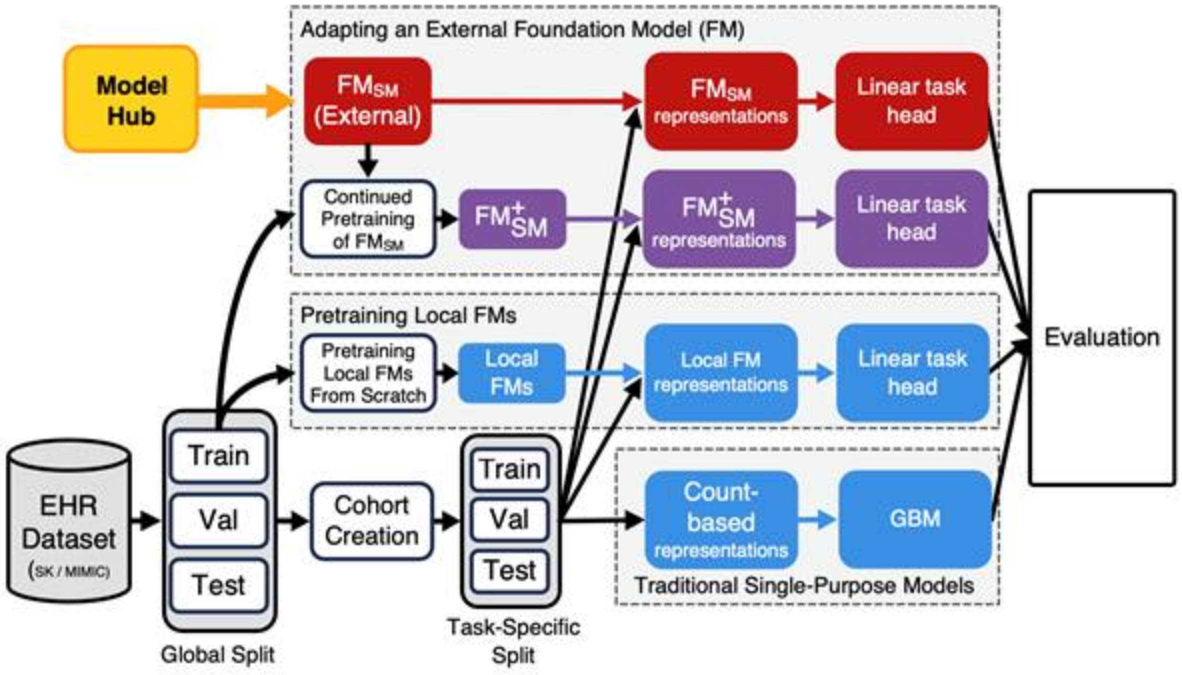


References:
[1] A Multi-center Study on the Adaptability of a Shared Foundation Model for Electronic Health Records

# Examples of Foundation Models in Healthcare

- **Clinical Prediction Tasks**: In-hospital mortality, long length of stay, 30-day readmission, abnormal lab results.
- **Structured EHR Data**: Summarize patient medical history for diagnostics.
- **Model Architecture**: FMSM, 141M parameter Transformer model, pretrained on 2.57M patient records.
- **Data Mapping**: OMOP CDM for compatibility.

**EHR**: Electronic Health Records, digital versions of patients' paper charts containing medical history, diagnoses, medications, treatment plans, and more.

**FMSM**: Foundation Model Stanford Medicine, A model trained on patient records from Stanford Medicine.

**OMOP CDM**: Observational Medical Outcomes Partnership Common Data Model, standardizes healthcare data for analysis.

References:
[1] A Multi-center Study on the Adaptability of a Shared Foundation Model for Electronic Health Records

# In-Depth Study of Clinical Prediction Tasks

- **Methodology:**
  - Data Sources: SickKids, MIMIC-IV.
  - Evaluation Metrics: AUROC, ECE.
- **Results**:
  - In-hospital mortality: AUROC 0.957.
  - Long length of stay: AUROC 0.839.
  - Hypoglycemia, anemia: Improved prediction accuracy.
- **Technical Insights:**
  - Feature Representations: Dense vectors for patient timelines.
  - Continued Pretraining: +3% performance improvement.

**SickKids**: The Hospital for Sick Children

**MIMIC-IV**: Medical Information Mart for Intensive Care

**AUROC**: Area Under the Receiver Operating Characteristic, measures model's ability to distinguish between classes.

**ECE**: Expected Calibration Error, assesses model's prediction accuracy and confidence alignment.

References:
[1] A Multi-center Study on the Adaptability of a Shared Foundation Model for Electronic Health Records

# Impact on Healthcare Processes

- **Efficiency**:
  - Few-shot learning: 128 examples match GBM performance.
  - Reduced label acquisition costs.
- **Robustness**:
  - Cross-site adaptability: SickKids and MIMIC-IV.
  - Continued pretraining: Enhanced local adaptation.
- **Technical Insights**:
  - Hyperparameter tuning.
  - Hierarchical bootstrapping for evaluation.

**GBM**: Gradient Boosting Machines, ensemble method using weak models for predictions.

References:
[1] A Multi-center Study on the Adaptability of a Shared Foundation Model for Electronic Health Records

# Advancements

- **Adaptability**:
  - Strong performance across datasets.
  - Continued pretraining: +3% performance.
- **Efficiency**:
  - 60-90% fewer training examples needed.
  - Faster deployment.
- **Accuracy**:
  - High AUROC scores.
  - Improved calibration (lower ECE).
- **Technical Insights**:
  - Decoder-only Transformer architecture.
  - Next-code prediction task for pretraining.

References:
[1] A Multi-center Study on the Adaptability of a Shared Foundation Model for Electronic Health Records

# Multi-Center Applicability and Scalability

- **Real-World Use Cases**:
  - SickKids: Pediatric healthcare.
  - MIMIC-IV: Adult ICU settings.
- **Technical Insights**:
  - OMOP CDM for data consistency.
  - Hierarchical bootstrapping for robust evaluation.
- **Cohort Characteristics**:
  - Detailed patient demographics.
  - Few-shot learning experiments.

References:
[1] A Multi-center Study on the Adaptability of a Shared Foundation Model for Electronic Health Records

# AI challenges in healthcare

Jingyi Cui (cau8rc)

# AI challenges in healthcare

- **Healthcare Data Is Unique:**
  - Highly sensitive and personalized (e.g., patient medical records).
  - Diverse formats: EHRs, medical imaging, genomic data, etc.

- **Challenges:**
  - Data quality issues.
  - Privacy and security concerns.
  - Limited access to diverse datasets.

[1] Foundation Models in Medicine: Limited public datasets make generalization challenging.
[2] Sharing constraints: Privacy laws like GDPR, HIPAA hinder global data sharing

# Data Quality Issues

- **Inconsistencies in Data**
  - Hospitals adopt different EHR standards, leading to compatibility issues.
- **Example:**
  - SickKids (Canada): Pediatric-focused data.
  - MIMIC-IV (USA): Primarily adult ICU data.
  - The lack of a unified standard complicates training and testing across institutions.

[1] A Multi-Center Study on the Adaptability of a Shared Foundation Model for Electronic Health Records
[2] Foundation Models for Advancing Healthcare: Challenges, Opportunities, and Future Directions

# Data Quality Issues

- **Incomplete Data**
  - Missing values in EHR datasets (e.g., unrecorded lab results, treatment histories).
  - Fragmented patient data across departments limits its utility for AI models.
- **Bias in Data**
  - Underrepresentation of certain populations in training datasets.
  - SickKids data is leaned towards children, reducing generalizability to adult healthcare systems.

[1] A Multi-Center Study on the Adaptability of a Shared Foundation Model for Electronic Health Records
[2] Foundation Models for Advancing Healthcare: Challenges, Opportunities, and Future Directions

# Privacy and Security

- **Strict Privacy Laws:**
  - GDPR (Europe): Restricts data sharing across borders.
  - HIPAA (USA): Ensures protection of patient health information (PHI).
  - Institutions face legal and financial risks if they fail to comply.
- **Data Anonymization Issues**
  - De-identification techniques remove personal identifiers but can lead to: Loss of contextual richness in medical text. And reduced performance for AI models in tasks like prediction and diagnosis.
  - GatorTronGPT: Used de-identified text for training but experienced limitations in tasks requiring nuanced clinical context

[1] A Multi-Center Study on the Adaptability of a Shared Foundation Model for Electronic Health Records
[2] A Comprehensive Survey of Foundation Models in Medicine

# Privacy and Security

- **Challenges of Cross-Institutional Collaboration**
  - Hospitals and institutions lack standardized frameworks for secure data sharing.
  - SickKids and MIMIC-IV data were used separately due to privacy concerns, despite their complementary nature
- **Emerging Security Risks**
  - AI models are vulnerable to: Data breaches and cyber-attacks.
  - Unauthorized use of AI models trained on sensitive medical datasets could lead to privacy violations and ethical dilemmas

[1] A Multi-Center Study on the Adaptability of a Shared Foundation Model for Electronic Health Records
[2] Foundation Models for Advancing Healthcare: Challenges, Opportunities, and Future Directions

# Algorithmic Challenges in Healthcare AI

- **Why Algorithms Matter in Healthcare AI**
  - Algorithms drive decision-making: From diagnosis to treatment recommendations.

- **Challenges arise from model complexity and healthcare's dynamic nature.**
  - Explainability and interpretability.
  - Handling bias and fairness.
  - Reliability in clinical settings.

[1] A Comprehensive Survey of Foundation Models in Medicine
[2] Foundation Models for Advancing Healthcare: Challenges, Opportunities, and Future Directions

# Responsibility in Algorithmic Decision-Making

- **Explainability Challenges:**
    - Neural networks often function as "black boxes." GPT-4 in medical contexts provides no clear rationale for its conclusions.

- **Fairness Issues:**
    - Models trained on biased datasets produce unfair outcomes.
    - Found inherent biases in GPT-based medical systems when diagnosing rare diseases

[1] A Comprehensive Survey of Foundation Models in Medicine
[2] Foundation Models for Advancing Healthcare: Challenges, Opportunities, and Future Directions

# Ensuring Reliability in AI Models

- **Model Hallucinations:**
  - Generating plausible but incorrect outputs.
  - A model suggested non-existent drug interactions during a clinical query.

- **Outdated Knowledge:**
  - Healthcare evolves rapidly; static models can't keep up.
  - FMs trained on outdated datasets missed new treatment guidelines.
  - EHR Systems and Reliability: CLMBR-T-base model showed limitations in handling rare conditions due to static pretraining .

[1] A Comprehensive Survey of Foundation Models in Medicine
[2] Foundation Models for Advancing Healthcare: Challenges, Opportunities, and Future Directions

# Resource and Infrastructure Challenges in Healthcare AI

- **AI models require significant resources:**
  - Training foundation models (FMs) demands high-end hardware.
  - Deployment in real-world healthcare settings adds further constraints.

- **Core Challenges:**
  - High computational and financial costs.
  - Limited scalability for resource-constrained environments.

[1] A Comprehensive Survey of Foundation Models in Medicine
[2] Foundation Models for Advancing Healthcare: Challenges, Opportunities, and Future Directions

# Computational Demands in AI Model Training

- **Training Foundation Models:**
  - Training LLaMA with 65B parameters took 21 days on 2048 A100 GPUs with 80GB RAM per GPU.
  - Estimated energy costs for similar models run into millions

- **Expensive Hardware Requirements:**
  - Models like MedSAM require 20 Nvidia A100 GPUs (1600 GB memory).
  - Limited availability of GPUs globally affects researchers and smaller hospitals

[1] Foundation Models for Advancing Healthcare: Challenges, Opportunities, and Future Directions

# Scalability in Real-World Settings

- **Adaptation for Resource-Limited Settings:**
  - Many models cannot run efficiently on smaller devices or low-resource environments (e.g., rural clinics).
  - Scaling MedSAM for smaller hospitals remains impractical without significant modifications.
- **Balancing Cost and Efficiency:**
  - Methods like model compression and pruning can reduce resource needs but may compromise accuracy.
  - We need some ongoing research into lightweight architectures.

[1] Foundation Models for Advancing Healthcare: Challenges, Opportunities, and Future Directions

# Future Directions

Sheharyar

# Future of FMs in Healthcare

- **Scaling FMs:**
  - Addressing specialized medical domains by leveraging task-specific fine-tuning and domain adaptation techniques.
- **Interoperability**:
  - Achieving seamless integration and collaboration across multi-institutional datasets through federated learning and standardized data-sharing protocols.
- **Personalized Medicine:**
  - Enabling real-time adaptive models capable of understanding patient-specific data, such as genomics, lifestyle, and historical medical records.
- **Rare Diseases:**
  - Fine-tune model trained on common data
- **Artificial General Intelligence (AGI)**

# Overcoming Current Challenges

- **Data Scarcity:**
  - Limited annotated datasets for model training can be addressed through synthetic data generation, augmentation strategies, and leveraging unstructured data via self-supervised learning.
- **Bias and Fairness:**
  - Disparities in performance across patient demographics necessitate the development of fairness-aware algorithms and bias evaluation metrics.
- **Scalability Issues:**
  - High computational costs for training and deployment require energy-efficient training methods, such as sparsity techniques and model compression.
- **Model Interpretability:**
  - Empowering clinicians with tools to intervene in the model's decision-making process can enable personalized and context-aware healthcare solutions.

# Overcoming Current Challenges

- **Model Size and Practical Use:**
  - Balance between performance and computational cost needs careful consideration.
- **Privacy Preserving:**
  - Ensuring data security and confidentiality while collaborative model training across multiple institutions.
- **Security:**
  - Incorporate robust FMs before deploying them into clinical settings.

# Multimodal Learning

- **Integration of Multimodal Data:**
  - Combining EHRs, medical images, genomic data, and wearable device signals to create comprehensive patient profiles.
- **Real-Time Decision-Making:**
  - Combining text, image, and bio-signal inputs through end-to-end multimodal frameworks such as transformers and cross-modal embeddings.
- **Artificial General Intelligence (AGI):**
  - Has the potential to revolutionize patient care by integrating advanced models that can understand and analyze complex clinical.

# Explainability and Trustworthiness

- **Explainability:**
  - Moving beyond black-box models by using techniques like SHAP values, saliency maps, and attention visualization.
- **Validation Frameworks:**
  - Ensuring safety and reliability in real-world applications through clinical trials and regulatory compliance.
- **Security:**
  - Incorporate robust FMs before deploying them into clinical settings so models are prone to adversarial attacks.
- **Ethical AI:**
  - Embed fairness and accountability mechanisms to build clinician and patient trust in AI systems.

# Ethical and Legal Considerations

- **Privacy-Preserving Models:**
  - Implementing secure federated learning to maintain patient confidentiality while enabling collaborative research.
- **Liability Concerns:**
  - Defining clear accountability frameworks and error-handling protocols for AI-assisted healthcare decisions.
- **Sustainability:**
  - Reducing the costs of data collection and processing, model training, and inference will stimulate the commercial advantages of HFMs and improve sustainability.

# Infrastructure Challenges

- **Efficient Training:**
  - Explore sparsity and pruning techniques to reduce computational costs without sacrificing accuracy.
- **Edge Computing:**
  - Deploy lightweight FMs in resource-limited environments, such as rural clinics.
- **Sustainability:**
  - Reduce environmental impact through energy-efficient training methods and reusable models.

# Collaborating for the Future

- **Industry-Academia Partnerships:**
  - Innovate by combining cutting-edge research with practical healthcare applications.
- **Open-Source Initiatives:**
  - Democratize AI development by sharing pre-trained models and datasets while ensuring privacy.
- **Interdisciplinary Teams:**
  - Build bridges between data scientists, clinicians, and ethicists for comprehensive AI solutions.

# Conclusion

- **Transformative Potential:**
  - Foundation models are revolutionizing healthcare, enabling improved diagnostics, and efficient resource use.
- **Overcoming Challenges:**
  - Addressing data scarcity, bias, interpretability, and computational demands is critical for widespread adoption.
- **Collaborative Future:**
  - Interdisciplinary efforts between researchers, clinicians, and policymakers are essential for sustainable progress.
- **Vision Forward:**
  - The integration of advanced AI technologies promises a future of equitable, accessible, and impactful healthcare solutions for simple and complex scenarios.

# Discussion

Joseph Moretto (jam5cq)

# Discussion

- In what situation would using foundation models not always be ideal in healthcare?
    - Highly specific tasks, lack of data
    - Task is very sensitive to bias
    - Extreme accuracy is needed
    - Deployment resource limitations

# Discussion

- How can privacy be ensured with foundation models in healthcare?
  - Data modification with the goal of privacy
  - Differential privacy to prevent inference attacks
  - LLM privacy agent for verification

# Discussion

- How will the adoption of AI in healthcare affect the workforce?
  - Training may be needed to use AI properly
  - Importance of knowing the AI's limitations and verifying
  - More productive

# Discussion

- Do you think AGI will help solve limitations of FMs?