



## OPEN ACCESS

EDITED AND REVIEWED BY  
Marcello Pelillo,  
Ca' Foscari University of Venice, Italy

\*CORRESPONDENCE  
Sokratis Makrogiannis  
✉ smakrogiannis@desu.edu

RECEIVED 12 November 2025  
REVISED 12 November 2025  
ACCEPTED 17 November 2025  
PUBLISHED 26 November 2025

CITATION  
Makrogiannis S (2025) Editorial: Foundation  
models for healthcare: innovations in  
generative AI, computer vision, language  
models, and multimodal systems.  
*Front. Comput. Sci.* 7:1744581.  
doi: 10.3389/fcomp.2025.1744581

COPYRIGHT  
© 2025 Makrogiannis. This is an open-access  
article distributed under the terms of the  
[Creative Commons Attribution License \(CC  
BY\)](#). The use, distribution or reproduction in  
other forums is permitted, provided the  
original author(s) and the copyright owner(s)  
are credited and that the original publication  
in this journal is cited, in accordance with  
accepted academic practice. No use,  
distribution or reproduction is permitted  
which does not comply with these terms.

# Editorial: Foundation models for healthcare: innovations in generative AI, computer vision, language models, and multimodal systems

Sokratis Makrogiannis\*

Division of Physics, Engineering, Mathematics, and Computer Science, Delaware State University,  
Dover, DE, United States

## KEYWORDS

editorial, foundation models, multimodal systems, healthcare, deep learning

## Editorial on the Research Topic

[Foundation models for healthcare: innovations in generative AI, computer vision, language models, and multimodal systems](#)

Foundation models—large pre-trained vision, language, and multimodal systems—are reshaping how we approach medical data, from image segmentation to synthetic clinical data generation and multimodal fusion for diagnostics. The Research Topic Foundation Models for Healthcare collects an informative cross-section of recent advances that illuminate the promise and the obstacles of applying foundation models to real-world clinical problems. This editorial synthesizes the Research Topic's contributions, highlights emergent lessons, and outlines priorities for translating foundation models safely and effectively into healthcare practice.

## Key trends and advances

Three clear trends run through the Research Topic. First, foundation vision models can dramatically lower the barrier to effective image analysis. [Joas et al.](#) show that the Segment Anything Model (SAM), used zero-shot, provided excellent confluence estimates for mesenchymal stem cell cultures—in their setup outperforming fine-tuned specialist models and rendering exhaustive annotation unnecessary. This result is a striking demonstration that, for certain homogeneous imaging tasks, generalist foundation models can reduce annotation costs while delivering high performance.

Second, large language models are emerging as practical tools to generate realistic synthetic clinical data. [Barr et al.](#) used GPT-4o to generate perioperative tabular datasets and found that most parameters' distributions were statistically similar to an open real dataset, suggesting LLM-based synthetic data could alleviate privacy and access bottlenecks for secondary analyses and method development. However, synthetic realism does not automatically equate clinical utility or bias-free data; rigorous validation is still required.

Third, the Research Topic emphasizes multimodality and clinical end-goals. Several contributions—including multimodal MRI radiomics for HIFU prediction by [Wen et al.](#), and facial-gesture + paralinguistic systems for pain detection by [Gutierrez et al.](#)—illustrate that combining imaging, structured data, and behavioral signals can enable

clinically relevant predictions, but they also underscore challenges in annotation, generalization across centers, and clinically meaningful evaluation metrics.

## Recurring methodological and translational lessons

- Zero-shot/few-shot performance can be surprisingly strong—as Joas et al. show, foundation models can sometimes replace arduous labeling efforts for homogeneous tasks.
- Synthetic data is promising but must be validated beyond distributional similarity—the Barr et al. study demonstrates realistic distributional match, but downstream predictive value, bias propagation, and leakage risks require thorough testing and domain-expert scrutiny.
- Robustness to preprocessing and imaging pipelines matters—imaging pipeline studies (e.g., CT noise-reduction assessments) show that preprocessing choices can materially change inputs to AI systems; pipelines must be validated end-to-end.
- Multicenter generalization remains a major bottleneck—several works (radiomics multicenter studies, systematic reviews) highlight heterogeneity in acquisition and labeling leading to variable performance; building foundation models that generalize across institutions remains essential.
- Evaluation needs to reflect clinical utility—moving beyond conventional metrics (IoU, AUC) to outcomes-oriented, prospective, and human-in-the-loop evaluations is critical for translation.

## Ethical, safety, and regulatory considerations

Foundation models bring particular ethical and regulatory questions. Synthetic data generation can aid privacy but might still encode biases or create spurious correlations; LLMs can hallucinate structured records that appear realistic yet contain impossible combinations unless constrained and validated. Interpretability and auditability remain challenging with large opaque architectures; for clinical acceptance, transparency, provenance tracking, and failure-mode analysis must be standard. Lastly, regulatory pathways (e.g., approvals of software as a medical device) need case studies and benchmarks to consider foundation-model-specific risks.

## Roadmap—Priorities for the next 3 years

- Benchmarking and shared datasets: Create multicenter benchmarks that measure clinical relevance and generalization. The Research Topic's contributions provide initial seeds; more coordinated datasets are needed.
- Synthetic data governance: Define standards for synthetic health data, including leakage testing, bias audits, and downstream predictive validity checks.

- Lightweight foundation models and efficient deployment: Promote hybrid architectures and distilled models for hospitals with limited compute, inspired by MoNetViT (Triyono et al.) and efficient CNN/transformer work.
- Clinical validation pathways: Fund and run prospective trials and real-world deployments (not only retrospective benchmarks) to verify clinical value and safety.
- Explainability and human-in-the-loop design: Integrate clinicians in the loop and deploy explainability tools that matter for decision-making and error detection.

## Closing remarks

The Frontiers Research Topic brings together work that illustrates both the promise and the complexity of applying foundation models in healthcare. From zero-shot microscopy segmentation to LLM-driven synthetic data generation and multimodal prognostic systems, the field is moving rapidly. The path to clinical applications requires rigorous validation, improved evaluation frameworks, and multidisciplinary coordination among AI researchers, clinicians, ethicists, and regulators. The articles in this Research Topic are a valuable step forward and provide concrete starting points for the coordinated effort needed to translate foundation models into safe, equitable, and useful clinical tools.

## Author contributions

SM: Writing – original draft, Writing – review & editing.

## Funding

The author(s) declare that financial support was received for the research and/or publication of this article. The author acknowledges support by the National Institute of General Medical Sciences of the National Institutes of Health (NIH) awards #SC3GM113754, #1U54MD015959-01A1, and the National Science Foundation awards #2234871 and #2401835.

## Conflict of interest

The author declares that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

## Generative AI statement

The author(s) declare that no Gen AI was used in the creation of this manuscript.

Any alternative text (alt text) provided alongside figures in this article has been generated by Frontiers with the support of artificial intelligence and reasonable efforts have been made to ensure accuracy, including review by the authors wherever possible. If you identify any issues, please contact us.

## Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated

organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.