Artificial intelligence foundation models in healthcare: A Malaysian perspective

Shen-Han Lee, PhD[1], Irfan Mohamad, MMed (ORL-HNS)[1,2]

[1]Department of Otorhinolaryngology-Head & Neck Surgery, School of Medical Sciences, Universiti Sains Malaysia, Health Campus, 16150 Kubang Kerian, Kelantan, Malaysia, [2]Hospital Pakar Universiti Sains Malaysia, Universiti Sains Malaysia, Health Campus, 16150 Kubang Kerian, Kelantan, Malaysia

ABSTRACT
One of the most revolutionary breakthroughs in modern artificial intelligence research over the past decade has been the introduction of foundation models – deep learning models trained on extensive datasets that can be adapted to tackle a wide range of downstream tasks. The rise of foundation models has significantly accelerated the adoption of AI in healthcare where there is increasing digitalization, enabling the integration of medical imaging, clinical notes, and genomic data to provide a more holistic understanding of patient health and supporting personalized interventions. In this Editorial, we will explore how foundation models are catalysing insights in precision health and offer our perspective of how foundation models can be integrated into the Malaysian healthcare system. In addition, we will also highlight some of the issues concerning governance, ethical, regulatory and policy challenges of implementing these foundation models in the Malaysia.

KEYWORDS:
Artificial intelligence; Foundation models; Large language models; Clinician; Healthcare system; Malaysia

INTRODUCTION
One of the most revolutionary breakthroughs in modern artificial intelligence (AI) research are foundation models including large language models which have been in the limelight since the release of Chat Generative Pretrained Transformer (ChatGPT) by OpenAI in November 2022.[1] This watershed moment heralded the era of generative artificial intelligence in the public domain with its widespread use across all fields including medicine. In this Editorial, we will explore how foundation models are catalysing insights in precision health and offer our perspective of how foundation models can be integrated into the Malaysian healthcare system. In addition, we will also highlight some of the issues concerning governance, ethical, regulatory and policy challenges of implementing these foundation models in Malaysia.

1. What are Foundation Models?
AI, broadly defined, is the use of technologies to build computational systems capable of mimicking human cognitive function and intelligence. Generative AI, a subset of AI, are deep learning models that can be trained to create new data. The term "foundation model", first coined by the Stanford Institute for Human-Centered Artificial Intelligence in 2021, refers to a form of generative AI, defined as large-scale deep learning models with broad capabilities that are trained on extensive datasets in an unsupervised fashion and can be adapted (e.g. fine-tuned) to tackle a wide range of downstream tasks.[2] In essence, these generative AI models are flexible and broadly applicable across various domains.

The modern evolution of foundation models began with the milestone introduction of the transformer architecture in 2017, a design that uses self-attention to process sequences in parallel and capture complex data dependencies more efficiently than previous deep learning models, such as recurrent neural networks.[3] Transformers emulate human attention by learning to focus on key aspects of input data, thereby improving prediction accuracy across varied applications such as natural language processing and computer vision. Today, foundation models are built on generative deep neural networks using the transformer architecture and leverage the principles of masked language modelling to understand the data structure, and transfer learning to reuse this knowledge to boost their performance of a related task. In simple terms, these large-scale deep learning models are trained on very large amounts of data without being given explicit instructions (unsupervised learning), allowing them to learn general patterns of these data on their own. Once trained, these models can be further adapted (fine-tuned) to perform tackle a wide range of downstream tasks, rather than being built only for a single purpose. Well-known examples of foundation models include Generative Pre-trained Transformer (GPT),[4] Bidirectional Encoder Representations from Transformers (BERT),[5] Gemini,[6] Large Language Model Meta AI (LLaMA),[7] and DeepSeek.[8] Essentially, the original model serves as a base (hence the term "foundation") on which other functions can be built upon, in contrast to other AI systems specifically trained for a particular use. A glossary of AI terminology is listed in Table I.

The most widely recognized use of foundation models in healthcare is in natural language processing (NLP) through the use of large language models (LLMs). These models are trained on vast amounts of text data and can be adapted (fine-tuned) for specific tasks such as summarizing medical documents, answering questions, and organizing or classifying information.[9] The same approach has also been

Corresponding Author: Irfan Mohamad
Email: irfankb@usm.my

Table I: Glossary of AI terms

Term	Definition
Artificial Intelligence (AI)	The use of technologies to build machines and computers capable of mimicking human cognitive functions and intelligence
Generative AI	AI systems that are designed to create new content, such as text, audio, or images in response to user prompts, by learning patterns from existing data
Machine Learning	A sub-field of AI that uses statistical algorithms to learn and make predictions based on input data without explicit programming, thereby "learning from experience."
Deep Learning	A variant of machine learning that uses neural networks to extract higher features of unstructured data such as images, video, and text and provide a suitable output such as prediction or evidence for a decision. These neural networks simulate the complex decision-making abilities of the human brain.
Neural Network	An AI system inspired by the human brain that consists of a large set of interconnected computational units (akin to neurons). Neural networks can have hundreds of layers of these neurons, with each layer playing a role in solving problems related to complex tasks.
Natural Language Processing	A subfield of artificial intelligence that facilitates automated analysis of language using machine learning, enabling computers to understand and communicate with human language.
Foundation Model	Large-scale deep learning models with broad capabilities that is trained on extensive datasets in an unsupervised fashion and can be adapted (e.g. fine-tuned) to tackle a wide range of downstream tasks
Transformer	A deep learning architecture that processes sequential data using a mechanism called parallel attention that emulates human attention, learning to focus on key aspects of input data and capture their long-range dependencies more effectively than previous deep learning models. In essence, the model dynamically weighs the significance of different parts of the input data and process them in parallel, making them highly scalable to process large datasets efficiently.
Large Language Model	A deep learning model that is trained on text to perform natural language processing tasks.
Transfer Learning	A machine learning approach in which a model trained on one task is adapted to perform a different but related task, using less data and training time.
Masked Language Modelling	A form of learning where certain tokens in a sequence (e.g. words in a sentence) is masked and hidden, challenging the models to predict these tokens and thereby develop an understanding of language structure.
Token	A basic unit of a sequence, such as a word, amino acid, or gene, that serves as the fundamental input unit for a foundation model.
Embedding	A numerical representation of real-world objects (e.g. words) as a vector in a high dimensional space that machine learning and AI systems use to understand complex knowledge domains
Pre-training	The first step of transfer learning, in which a machine learning model is trained on a large dataset to learn general features.
Fine-tuning	The second step of transfer learning in which a pre-trained model is adapted to a specific task by training it on a smaller and task-specific dataset.
Zero-shot	The ability of AI models to perform a task without needing any specific training on that task.

extended beyond text to other types of data, including medical image pixels,[10] audio signal waveforms[11] and even biological sequences such as protein amino acid sequences.[12]

2. When Foundation Models meets Digital Health and Big Data

The rise of foundation models has dovetailed with the big data and digital health revolution to transform how patient data can be used to derive insights that help clinicians make better decisions and deliver more personalized healthcare, ranging from disease diagnosis and treatment to genomics and drug discovery. One of the most notable applications of foundation models in healthcare is in medical imaging where vision-based foundation models have been used to analyse medical images, in particular, from three areas of medicine: radiology,[13] pathology,[14] and ophthalmology.[15] The field is rapidly evolving field with the development of various vision language models trained on multi-centre CT scan images,[16] retinal images,[17] and pathology slides[18] that are capable of outperforming other baseline deep learning models on the same tasks.

In the case of electronic medical records (EMRs), foundation models have been used to predict patient outcomes, and identify disease risks and trajectory despite the complexity and heterogeneity of EMR datasets.[19–22] In public health, foundation models been deployed to analyze large-scale public health and geospatial data for disease surveillance and disease outbreak prediction.[23,24] In precision medicine, foundation models have been used to learn complex patterns in genomic datasets,[25] identify genetic markers for disease[26] and predict treatment responses at the single cell level.[27] In drug discovery, foundation models can accelerate drug discovery by predicting drug-target interactions, design drug molecules, and drug repurposing.[28–30]

3. How can we implement AI foundation models in Malaysian Healthcare?

Recently, the Malaysian Ministry of Health has announced plans to scale cloud-based clinic management systems to more than 2000 government primary care clinics nationwide and expand the reach of total health information system to more government hospitals.[31] The expansion of this digital health infrastructure is in line with the Malaysia's aspiration of implementing a national EMR system to achieve the 'One Patient, One Record' vision.[31]

In light of this national electronic health record initiative, we propose that foundation models, including but not limited to large language models, should be built into this cloud-based framework and national EHR. These models will enable active processing of high dimensional, heterogeneous time-stamped health data from these EHRs. Once trained, these foundation models can generate multiple simulated future patient health timelines and estimate the probability of clinical events occurring within those trajectories (e.g. length of stay, adverse events, hospital readmission, ICU admission). These probabilities may serve as dynamic risk estimates and serve as an early warning system, and as new patient data becomes available, the risk estimates are automatically updated.

Large language models can be trained to learn patterns from EHRs, such as clinic notes and investigation results. Furthermore, unstructured text in EHRs such a radiology reports offer a wealth of real-world data, and large language models offer new ways of extracting clinical insights from these large sets of unstructured medical text. In practice, these clinical data can be 'tokenized' and foundation models can then be trained to predict what information is likely to come next based on past data. Once trained, these models can be used to predict patient outcomes such as risk of mortality during or after hospitalization, the need for ICU admission, hospital readmission, and prolonged hospital stays.[19] Such models could also potentially help to improve patient risk stratification, suggest targeted interventions, and allow hospitals to better predict resource requirements.[32] In the future, these models could also help us to better understand the aetiology and progression of disease.

4. Building Trustworthy and Appropriate Foundation Models that Conform to Standards

As the worldwide AI community move towards more advanced foundation models such as multimodal systems that can reason across both images and text, there are several challenges that the various stakeholders in Malaysia need to work together to address.

AI systems need to be trustworthy, appropriate, transparent, conform to standards, and be clear on the issue of liability – these are principles espoused by the UK National Health Service AI Lab and Health Education England,[33] and more recently, the 2025 Malaysian Medical Council (MMC) Guideline on the Ethical Use of Artificial Intelligence in Medical Practice.[34] To this end, standards needs to be established with regards to data governance and the digital infrastructure should support safe and scalable deployment. Furthermore, clinical validation and performance monitoring should also evolve alongside technology.

However, there are complex challenges of integrating foundation models such as LLMs into existing clinical workflows. LLMs have a tendency to generate inaccurate or nonsensical information, commonly known as 'Hallucinations'. As such, who would take responsibility for medical errors, particularly in the context of these "hallucination" by LLMs? There have been reports on real-world harms to patients who relied on layperson-facing applications of foundation models for self-triage of transient ischemic attacks. Such cases raise concerns about whether these models can safely be integrated into healthcare systems. Ultimately, the MMC guidelines stipulates that the responsibility and liability fall on the registered medical practitioner, and therefore clinicians should be aware of the limitations of these foundation models to safeguard patient safety.[34]

5. Governance of AI Systems

The variety of public and private healthcare institutional structures in Malaysia means that there is fragmented ownership of health data, governance, procurement and data stewardship arrangements. This poses a governance challenge because there is no single authority that 'owns' responsibility for AI systems deployed across these silos. Furthermore, data used to train or fine-tune foundation models may come from multiple institutions with inconsistent data quality and oversight standards. Moreover, most hospitals in Malaysia lack AI governance committees and AI leads. As such, foundation models may only enter clinical workflows through research projects, isolated pilot initiatives, and informal use by clinicians and risk becoming a 'shadow infrastructure' outside proper institutional oversight with no audit trail, institutional responsibility or mechanism to learn from errors.

6. Data-sharing and Protecting Patient Privacy

Most existing foundational models are trained on datasets that are geographically and demographically 'narrow', limiting their global representation, generalizability and effectiveness particularly in under-represented regions and populations. Yet, there are concerns regarding data-sharing across institutions and borders, as well as the need to protect patient privacy. Therefore, robust data-sharing frameworks should be in place to address these concerns. To this end, we wish to highlight an example from the Global RETFound consortium, a global collaboration led by the National University of Singapore, Chinese University of Hong Kong, Moorfields Eye Hospital and University College London, that involves more than 100 study groups across 65 countries to create the world's first globally representative foundation model in medicine, trained on over 100 million fundoscopy images.[35] This consortium employs two robust, adaptable data-sharing frameworks – firstly, using only synthetic data generation through local fine-tuning of a foundation model, with only sharing of model weights, and no exchange of real data. As such, this ensures that the source data remains within their originating institutions, ensuring data sovereignty. The second data-sharing framework allows the transfer of anonymized real data within a secure cloud-based platform, governed by stringent protocols and legally-binding agreements. The group has released a comprehensive framework of their study protocol and data-sharing structure, allowing the use of these resources by other medical specialties and study groups to create similar initiatives within their own domains.[35]

7. Evaluating Foundation Models in the Local Setting

Due to the black-box decision making of foundation models and real-world variability, it is difficult to use traditional clinical trials to evaluate the efficacy of these models. New standards for reporting clinical trial protocols and results to evaluate interventions with an AI component have been introduced.[36,37] In addition, there should be rigorous validation of these foundation models on local health data using retrospective data, protocol-driven trials for using these foundation models for clinical decision support systems, and randomized clinical trials using prospective data.[38]

CONCLUSION

In the new age of AI, the rise of foundation models has rapidly accelerated various advances in precision medicine and public health, and is driven by healthcare digitalisation and the big data revolution. In turn, the large amount of data generated can be used to train these foundation models to tackle a wide range of downstream tasks including prognostication, risk stratification, and predictive analytics. To implement foundation models such as large language models against the backdrop of a national electronic medical record in Malaysia, there should be a national AI in healthcare steering committee that spans healthcare institutions across the public and private sectors, with data-sharing and patient privacy governed by robust legal, ethical, and regulatory frameworks.

REFERENCES

1. Introducing ChatGPT | OpenAI. https://openai.com/index/chatgpt/ (accessed 21 Jan 2026).
2. Bommasani R, Hudson DA, Adeli E, Altman R, Arora S, von Arx S et al. On the Opportunities and Risks of Foundation Models. arXiv:210807258 2021.https://arxiv.org/abs/2108.07258v3 (accessed 29 Oct 2024).
3. Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN et al. Attention Is All You Need. Adv Neural Inf Process Syst 2017; 2017-December: 5999–6009.
4. OpenAI, Achiam J, Adler S, Agarwal S, Ahmad L, Akkaya I et al. GPT-4 Technical Report. arXiv:230308774 [csCL] 2023.
5. Devlin J, Chang MW, Lee K, Toutanova K. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. NAACL HLT 2019 - 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies - Proceedings of the Conference 2018; 1: 4171-86.
6. Gemini Team, Anil R, Borgeaud S, Alayrac J-B, Yu J, Soricut R et al. Gemini: A Family of Highly Capable Multimodal Models. arXiv:231211805 [csCL] 2023.
7. Touvron H, Lavril T, Izacard G, Martinet X, Lachaux M-A, Lacroix T et al. LLaMA: Open and Efficient Foundation Language Models. arXiv:230213971 [csCL] 2023.
8. Guo D, Yang D, Zhang H, Song J, Wang P, Zhu Q et al. DeepSeek-R1 incentivizes reasoning in LLMs through reinforcement learning. Nature 2025; 645: 633-8.
9. Raiaan MAK, Mukta MSH, Fatema K, Fahad NM, Sakib S, Mim MMJ et al. A Review on Large Language Models: Architectures, Applications, Taxonomies, Open Issues and Challenges. IEEE Access 2024; 12: 26839-74.
10. Ramesh A, Pavlov M, Goh G, Gray S, Voss C, Radford A et al. Zero-Shot Text-to-Image Generation. arXiv:210212092 2021.
11. Copet J, Kreuk F, Gat I, Remez T, Kant D, Synnaeve G et al. Simple and Controllable Music Generation. arXiv:230605284 2023.
12. Lin Z, Akin H, Rao R, Hie B, Zhu Z, Lu W et al. Evolutionary-scale prediction of atomic-level protein structure with a language model. Science 2023; 379: 1123-30.
13. Paschali M, Chen Z, Blankemeier L, Varma M, Youssef A, Bluethgen C et al. Foundation Models in Radiology: What, How, Why, and Why Not. Radiology 2025; 314.
14. Veremis B, Chen S, Campanella G. An Introduction to Pathology Foundation Models. Head Neck Pathol 2026 20:1 2026; 20: 9.
15. Chia MA, Antaki F, Zhou Y, Turner AW, Lee AY, Keane PA. Foundation models in ophthalmology. Br J Ophthalmol 2024; 108: 1341-8.
16. Gao Z, Zhang G, Liang H, Liu J, Ma L, Wang T et al. A lung CT vision foundation model facilitating disease diagnosis and medical imaging. Nat Commun 2025; 17: 35.
17. Zhou Y, Chia MA, Wagner SK, Ayhan MS, Williamson DJ, Struyven RR et al. A foundation model for generalizable disease detection from retinal images. Nature 2023; 622: 156-63.
18. Ding T, Wagner SJ, Song AH, Chen RJ, Lu MY, Zhang A et al. A multimodal whole-slide foundation model for pathology. Nat Med 2025; 31: 3749-61.
19. Renc P, Grzeszczyk MK, Oufattole N, Goode D, Jia Y, Bieganski S et al. Foundation model of electronic medical records for adaptive risk estimation. Gigascience 2025; 14.
20. Yang Z, Mitra A, Liu W, Berlowitz D, Yu H. TransformEHR: transformer-based encoder-decoder generative model to enhance prediction of disease outcomes using electronic health records. Nat Commun 2023; 14: 7857.
21. Shmatko A, Jung AW, Gaurav K, Brunak S, Mortensen LH, Birney E et al. Learning the natural history of human disease with generative transformers. Nature 2025; 647: 248–56.
22. Akagi Y, Seki T, Kawazoe Y, Takiguchi T, Ohe K. A Generative Foundation Model for Structured Patient Trajectory Data. AMIA Annual Symposium Proceedings 2025; 2024: 124.
23. Agarwal M, Sun M, Kamath C, Muslim A, Sarker P, Paul J et al. General Geospatial Inference with a Population Dynamics Foundation Model. arXiv:241107207 2024. https://arxiv.org/pdf/2411.07207.
24. Kalahasti S, Faucher B, Wang B, Ascione C, Carbajal R, Enault M et al. Foundation time series models for forecasting and policy evaluation in infectious disease epidemics. medRxiv 2025.
25. Feng H, Wu L, Zhao B, Huff C, Zhang J, Wu J et al. Benchmarking DNA foundation models for genomic and genetic tasks. Nat Commun 2025; 16: 10780.
26. Machado Reyes D, Burch M, Parida L, Bose A. A foundation model for learning genetic associations from brain imaging phenotypes. Bioinform Adv 2024; 5.
27. Wang Q, Pan Y, Zhou M, Tang Z, Wang Y, Wang G et al. scDrugMap: benchmarking large foundation models for drug response prediction. Nat Commun 2025; 17: 730.
28. Chevalier A, Ghosh S, Awasthi U, Watkins J, Bieniewska J, Mitrea N et al. TEDDY: A Family Of Foundation Models For Understanding Single Cell Biology. arXiv:250303485 2025.
29. Huang K, Chandak P, Wang Q, Havaldar S, Vaid A, Leskovec J et al. A foundation model for clinician-centered drug repurposing. Nat Med 2024; 30: 3601-13.
30. Suryanarayanan P, Qiu Y, Sethi S, Mahajan D, Li H, Yang Y et al. Multi-view biomedical foundation models for molecule-target and property prediction. arXiv:241019704 2024. doi:10.48550/arXiv.2410.19704.
31. Ahmad D. My Say: Modernising the health system, elevating service quality. The Edge Malaysia. https://theedgemalaysia.com/node/788798 (accessed 21 Jan 2026).
32. Rajpurkar P, Chen E, Banerjee O, Topol EJ. AI in health and medicine. Nat Med 2022; 28: 31-8.
33. NHS AI Lab & Health Education England. Understanding healthcare workers' confidence in AI. Report 1. 2022.https://digital-transformation.hee.nhs.uk/ binaries/content/assets/digital-transformation/dart-ed/understanding confidenceinai-may22.pdf (accessed 20 Jan 2026).
34. Malaysian Medical Council. MMC Guideline on the Ethical Use of Artificial Intelligence (AI) in Medical Practice. 2025.https://mmc.gov.my/wp-content/uploads/2025/09/MMC-Guideline-on-The-Ethical-Use-of-Artificial-Intelligence-AI-.pdf (accessed 20 Jan 2026).
35. Tham YC, Goh JHL, Nderitu P, Zhou Y, Ran AR, Srinivasan S et al. Building the world's first truly global medical foundation model. Nat Med 2025; 31: 3580-5.

36. Cruz Rivera S, Liu X, Chan AW, Denniston AK, Calvert MJ, Darzi A et al. Guidelines for clinical trial protocols for interventions involving artificial intelligence: the SPIRIT-AI extension. Nat Med 2020; 26: 1351-63.
37. Liu X, Cruz Rivera S, Moher D, Calvert MJ, Denniston AK, Chan AW et al. Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: the CONSORT-AI extension. Nat Med 2020; 26: 1364-74.
38. Yuan H. Toward real-world deployment of machine learning for health care: External validation, continual monitoring, and randomized clinical trials. Health Care Sci 2024; 3: 360-4.