Kidney360

Multimodal Artificial Intelligence in Medicine

--Manuscript Draft--

Manuscript Number:	K360-2024-000329R1
Full Title:	Multimodal Artificial Intelligence in Medicine
Short Title:	Multimodal AI in Medicine
Article Type:	Review - Invited
Section/Category:	Acute Kidney Injury and ICU Nephrology
Corresponding Author:	Conor Judge, BMBS BENG PHD University of Galway CO. GALWAY, Co. Galway IRELAND
Corresponding Author E-Mail:	conorjudge@gmail.com
Other Authors:	Finn Krewer, PhD
	Martin J. O'Donnell, PhD
	Lisa Kiely, BMBS
	Donal Sexton, PhD
	Graham W. Taylor, PhD
	Joshua August Skorburg, PhD
	Bryan Tripp, PhD
Order of Authors (with Contributor Roles):	Conor S. Judge, BMBS BEng PhD (Conceptualization; Visualization; Writing – original draft; Writing – review & editing)
	Finn Krewer, PhD (Writing – original draft; Writing – review & editing)
	Martin J. O'Donnell, PhD (Writing – review & editing)
	Lisa Kiely, BMBS (Writing – review & editing)
	Donal Sexton, PhD (Writing – review & editing)
	Graham W. Taylor, PhD (Writing – original draft; Writing – review & editing)
	Joshua August Skorburg, PhD (Writing – original draft; Writing – review & editing)
	Bryan Tripp, PhD (Conceptualization; Supervision; Writing – original draft; Writing – review & editing)
Manuscript Classifications:	10: Acute Kidney Injury
Abstract:	Traditional medical Artificial Intelligence models, approved for clinical use, restrict themselves to single-modal data e.g. images only, limiting their applicability in the complex, multimodal environment of medical diagnosis and treatment. Multimodal Transformer Models in healthcare can effectively process and interpret diverse data forms such as text, images, and structured data. They have demonstrated impressive performance on standard benchmarks like USLME question banks and continue to improve with scale. However, the adoption of these advanced AI models is not without challenges. While multimodal deep learning models like Transformers offer promising advancements in healthcare, their integration requires careful consideration of the accompanying ethical and environmental challenges.
Funding Information:	Health Research Board (CSF-2023-016) — Conor S. Judge
Additional Information:	
Question	**Response**
Is this a Basic Science or Clinical Science	Clinical

topic?	
Section/Category: Select the Section or Category related to your manuscript from the drop-down menu below.	Acute Kidney Injury and ICU Nephrology
Study Group: Does your paper include study group(s)? If yes, please provide a list of study group(s) and members that have contributed to or participated in the submitted work in some way. This list may contain either a collaboration of individuals (e.g., investigators) and/or the name of an organization (e.g., a laboratory, educational institution, corporation, or department) and its members	No

OPEN

Kidney360 Publish Ahead of Print

DOI: 10.34067/KID.0000000000000556

Multimodal Artificial Intelligence in Medicine

Authors

Conor S. Judge (a,b)

Finn Krewer (a)

Martin J. O'Donnell (a)

Lisa Kiely (a)

Donal Sexton (c)

Graham W. Taylor (d,e)

Joshua August Skorburg (d)

Bryan Tripp (f)

Author department and institution affiliations

(a) HRB-Clinical Research Facility, University of Galway, Galway, Ireland
(b) Insight Data Analytics, University of Galway, Galway, Ireland
(c) Trinity College Dublin, Ireland
(d) University of Guelph, Guelph, Ontario, Canada
(e) Vector Institute, Toronto, Canada
(f) University of Waterloo, Waterloo, Ontario, Canada

Correspondence to

Conor Judge: conor.judge@universityofgalway.ie

Abstract

Traditional medical Artificial Intelligence models, approved for clinical use, restrict themselves to single-modal data e.g. images only, limiting their applicability in the complex, multimodal environment of medical diagnosis and treatment. Multimodal Transformer Models in healthcare can effectively process and interpret diverse data forms such as text, images, and structured data. They have demonstrated impressive performance on standard benchmarks like USLME question banks and continue to improve with scale. However, the adoption of these advanced AI models is not without challenges. While multimodal deep learning models like Transformers offer promising advancements in healthcare, their integration requires careful consideration of the accompanying ethical and environmental challenges.

Introduction

The collection and interpretation of multimodal data, including images (e.g., x-rays), unstructured text (e.g., clinical notes), and structured data (e.g., lab results), has been a critical component of caring for patients (1,2). Healthcare workers collect and combine these data, making informed diagnoses and treatment decisions, often using a semi-quantitative approach. While seasoned clinicians understand the inherently multimodal nature of diagnostic and therapeutic medicine, most current medical AI models, approved for clinical use, are not multimodal (3), employing AI models that are specialized for one specific task and usually use one type of clinical data (4). Our article focuses on recent developments in multimodal deep learning applications in healthcare and discusses some of the risks and potential opportunities of this very fast-paced technological revolution.

Scenario 1 – Multimodal Human Intelligence (Current Clinical Practice)

Mary is a 70-year-old female who attends the ER with a persistent cough with green-colored sputum and fever. She has a background of mild heart failure (Ejection fraction 45%), which has been stable for the past year. She has continued her routine medications, which are ramipril, frusemide, aspirin, and occasional non-steroid anti-inflammatories. In the emergency department, routine blood tests revealed raised CRP (96) with normal renal and liver profile. Her CURB-65 score was 3 and she was admitted to hospital for community-acquired pneumonia. All pre-admission medications are continued, as there is no specific concern by the admitting team. By Day-2 of admission, serum urea and creatinine become elevated and she was diagnosed with an acute kidney injury. The team discontinued all medications with adverse renal effects, including ramipril, frusemide, and non-steroidal anti-inflammatories. On Day 5 of admission, Mary has an exacerbation of heart failure, which extends her hospital admission by a further 8 days as it is also associated with an episode of delirium.

Scenario 2 – Assistance of Single Modal AI

In Scenario 2, the acute hospital has recently introduced an Acute Kidney Injury (AKI) alert system that takes data from the electronic health record (laboratory results, diagnosis codes, vital signs, and medication records) and generates a predicted probability of developing acute kidney injury within 48 hours of admission (Figure 1)(5). The AKI alert system employs a deep learning architecture called a recurrent neural network. The AKI alert system indicates a high probability of Mary developing an AKI and recommends holding all nephrotoxic medications (ramipril, frusemide, and non-steroidal anti-inflammatories), despite normal renal labs. This model does not have any information about her volume status (hypovolemia, euvolemia, or hypervolemia). While Mary does not develop an AKI, she does have an exacerbation of acute heart failure with prolonged admission associated with delirium.

Scenario 3 – Assistance of Multimodal AI

In Scenario 3, the same acute hospital has recently introduced a multimodal AKI alert system that takes in multiple types of data (clinical notes, laboratory data, imaging request details, medical images, imaging reports, clinical measurements, and prescription records) (Figure 1). The multimodal AKI alert system is built using a deep learning architecture called a ***Transformer Model*** (See Table 1). In addition to the predictive probability of AKI, the AI-system also considers findings on chest X-Ray, which reveals subtle features of mild pulmonary edema, along with a pneumonic infiltrate. In this scenario, the AI-algorithm recommends holding ramipril and non-steroidal anti-inflammatories, but to continue diuretic therapy. In this management scenario, the patient recovers from community-acquired pneumonia and discharged on Day 3 of admission, and does not experience AKI, exacerbation of heart failure or delirium.

Transformer AI Models (Origins)

To appreciate the anticipated role of Transformer Models in clinical care, it is useful to consider their origins. A ***Transformer*** is an Artificial Intelligence deep learning model ***Architecture*** (See Table 1) that is predominately used to process information provided in a sequence, and they were first employed for translating sentences from one language to another (6). In addition to translating individual words, a Transformer Model captures the context of each word with interpretive relevance to all other words in the sentence, a process known as ***Attention*** (See Table 1). The Transformer architecture is the foundation of complex Large Language Models (LLMs), which employ extensive and interconnected Transformer Models in the interpretation of information from diverse sources. Following interpretation, Transformer Models can provide summary interpretation to the human user. Such models are referred to as Generative LLMs (e.g. Generative Pretrained Transformer (GPT) family of models). Generative LLMs are trained on extensive amounts of general text data (internet web pages, books, and Wikipedia), allowing them to generate human-like text and complete complex language tasks (7–9), resulting in the ability to provide detailed answers to specific questions.

Transformers/LLMs Go to Medical School: As the sophistication of transformer models to acquire and interpret information grew, an obvious target for deployment was healthcare delivery. We might consider the evolution of LLMs to the development and training of a medical doctor. An initial step, therefore, is the acquisition and assimilation of medical information, suitable to correctly answer knowledge-based medical questions. Several LLMs have undertaken training (studying) in large datasets of medical information and then tested their knowledge in conventional medical examinations. For example, BioMedLM, trained on biomedical literature from PubMed, achieved a score of 50.3% in medical question answering on the MedQA benchmark which is a USMLE-style question bank (December 2022) (10). Med-Palm 2 is closed source medical transformer model that was trained on general data and then fine-tuned on medical data (11). Med-Palm 2 has reached a score of 86.5% on the MedQA benchmark (June 2023) (11) and GPT-4, also closed source, has surpassed 90% (12). Open source models are quickly reaching this standard also, Meditron-70B is a smaller open source model that achieved a score of 70.2% on MedQA (November 2023) (13). Currently, LLMs can successfully complete knowledge-based assessments of medical knowledge.

Transformers/LLMs Can Be Bad Students: LLMs often produce incorrect “hallucinated” text such as fabricating information in responses, a behavior not uncommon among medical students (14). The Galactica LLM by Meta (15) was trained on high-quality sources including medical textbooks and journal articles, and provided an interactive interface. However, users discovered that it produced false output with an authoritative tone and convincing detail, such as well-written journal articles on the benefits of eating crushed glass (16). Such tendencies arise from the way LLMs learn to model language (17). Training generative transformer models mostly consists of learning to predict the next word in human-written texts (autoregressive training). For example, the words “presented with urinary frequency and 48 hours of fever and rigors. He was hypotensive and” are likely to be followed by “tachycardic”. To some extent, LLMs encode this underlying reality (18).

LLMs may output fabrications, such as inserting a plausible body mass index (BMI) into a generated medical note despite lacking access to the patient’s weight (19). As such, the LLM may provide a ‘best guess’ rather than accurate information, but not reflect uncertainty in the response. To improve accuracy of responses, a process of supervised fine-tuning involves a labelled dataset of prompts and desired responses written by humans to fine-tune the responses that the model gives (20). Alignment training of GPT models uses reinforcement learning with reward-based human feedback, based on human-interpreted quality of output (21). Such methods have greatly improved GPT models, but confabulation persists. A recent study reporting the use of ChatGPT in an Intensive Care Unit setting found that the answers provided by ChatGPT were often erroneous or dated, and the output would often contain incorrect factual pieces of information like dates and doses of medications (22). A further challenge is that LLMs are unable, currently, to distinguish which outputs are based on factual information and which are based on hallucinations.

Transformers work effectively with any modality of data – Medical Data Interpretation

Transformers like other neural network architectures require numeric input. To enable a Transformer to read text, the words must be converted into a list of numbers called a ***Vector*** through a process called ***Encoding*** (See Table 1). Beyond that first step, nothing about the operation of the transformer model is specific to language. The same network structures and training methods work well with entirely different modalities of data, such as sounds and images. Vision Transformer (ViT) models, like MedSAM (Medical Segment Anything Model) achieved state-of-the-art performance on many medical image segmentation tasks over pre-existing specialized non-transformer models (23). Transformers also excel in audio processing tasks, such as classifying sounds (24) and transcribing speech (25–28). Recently, transformers have been applied to bio signals as well. For example, strong performance has been reported in the classification of arrhythmias and epileptiform activity, with transformers that were trained on large unlabeled ECG and EEG datasets and fine-tuned on smaller classification datasets (29). Where direct comparisons have been made, transformers often outperform other kinds of networks such as convolutional neural networks if they are first trained on large volumes of data (30) or on tasks analogous to next-word prediction (25,29,30). Transformers need large volumes of unlabeled data (e.g. images only), after which they can quickly learn a particular task with limited labeled data (e.g. image-category pairs).

Multimodal Transformers – Interactive Medical Cases

Transformers are often trained to process multiple modalities. The DALL-E image generation model was trained on hybrid sequences consisting of image descriptions followed by image patches (31). Given only an image description, it could predict, patch-by-patch, an image likely to match the description.

Multimodal transformers like GPT-4 and Gemini can describe and reason about images. The company Be My Eyes is using this technology to develop a smartphone app for visually impaired people that can describe and answer questions about their visual surroundings (32). At the time of publication (July 2024), the image-related capabilities of GPT-4 Vision were evaluated on NEJM (88.7%) and JAMA (73.3%) image challenges, consistently outperforming human readers (33).. A technical paper released by OpenAI tested the ability of GPT-4V to give medical advice (34) reported serious interpretation errors for medical imaging including misdiagnosing the lateralization in medical images.

Med-PaLM Multimodal (Med-PaLM-M) is a large milestone towards a generalist biomedical AI system. It encodes and interprets multimodal biomedical data including clinical language, imaging, and genomics with the same set of model weights (35). It required a new testing benchmark (MultiMedBench) to be curated and includes multimodal medical tasks such as Radiology and Pathology visual question answering and dermatology and mammography image classification for which Med-PaLM-M reached state-of-the-art performance.

In Radiology, ELIXR is an approach that combines large language models with vision encoders and was trained on the MIMIC-CXR dataset that contains pairs of chest x-ray images and chest x-ray text reports (36). ELIXR demonstrated state-of-the-art classification performance comparable to other supervised learning methods on complex visual questions.

HeLM (Health LLM) combines high dimensional clinical information with text prompts to estimate disease risk (37). This approach, based on the PaLM-E model, trains an encoder which maps high dimensional data into the LLM's token embedding space. Using text and encoded tabular data from the UK Biobank HeLM predicts all-cause mortality and hypertension with superior accuracy to traditional machine learning algorithms using the same dataset. The HeLM model demonstrates improved accuracy in predicting asthma phenotype when combining spirometer data with text prompts and encoded tabular data over text prompt and tabular data only.

Transformers improve reliably with model and data scale – Clinical Experience

While there have been important technical advances in recent years, e.g. to allow LLMs to carry on longer conversations (38–40), today's transformers are structured in much the same way as older transformers, except that they are larger in every dimension. By some measures, the performance of LLMs improves predictably with scale (9,41–43). Other abilities like mathematical operations seem to emerge suddenly at certain model scales (44,45). The sudden emergence of a capability might say more about the task than the model, in that the success measure may be a nonlinear function of a deeper capacity that emerges more gradually (46).

Transformers are much less sophisticated than the human brain, so one might suspect that they will not surpass humans without significant changes. However, they have advantages of scale. Some transformers have detailed working memory for hundreds of pages of text, whereas human working memory begins to struggle with any more than a phone number. PaLM (47) was trained on 780 billion "tokens" (roughly 585B words), corresponding to roughly ten thousand years of full-time reading for a human. Later stages of LLM training currently rely on humans to rank their responses. However, LLMs seem to be even better at critiquing than producing answers (19), suggesting a potential path to indefinite improvement. The most imminent limiting factor is that there is only so much high-quality text available to consume. Recent transformers have already been trained with perhaps one-tenth of the total high-quality text that exists (48).

Future Directions and Risks for Multimodal AI in healthcare

Another set of considerations are whether current ethical and legal frameworks are adequate to address the challenges raised by multimodal AI in healthcare. The last decade has seen a burgeoning literature develop on the ethics of AI in healthcare generally. Issues such as data ownership, informed consent, bias, privacy, responsibility gaps, opacity, digital divides, environmental impact, have been extensively studied (49). Rapid LLM development and lag time in academic publishing result in few studies on LLM ethics in healthcare (50–54). High-level ethical principles often lack actionable guidance in specific medical AI contexts. Thus, we think that recent approaches which focus on stakeholder engagement and context-specific analysis of potential harms (55–57) strike the right balance between realizing the transformative potential of AI, while avoiding ethical pitfalls.

Importantly, large models could be detrimental to health, even if they benefit healthcare. LLMs use large amounts of electricity and could drive significant CO_2 emissions, exacerbating potential health impacts of climate change (58,59). Large model training is carbon-intensive, e.g. training GPT-3 produced about 552 metric tons of CO_2-equivalent emissions. But most emissions tend to come from use of the model after training (60). Future CO_2 emissions depend on model sizes, hardware, the source of electricity used, and algorithm improvements (60–63), but there is a risk of substantial impact. For instance, processing each health data item once, a model like PaLM (47) requires as many floating point operations as it has parameters (540B). Health data may currently be growing by about $1.5\text{x}10^{21}$ bytes per year (64). Assuming for simplicity that much health data consists of 16-bit images (30), this would correspond to roughly $3\text{x}10^{18}$ tokens per year. If this processing were done on typical current hardware such as Nvidia A100 GPUs, using power from the eastern US, which is moderately carbon-intensive (61), this would result in about $3.5\text{x}10^{10}$ metric tons of CO_2 emissions per year (47), roughly 10% of the current total for healthcare (65,66).

Environmental impacts are just one of several categories of AI risk (63). Large models may also pose various extreme risks that are difficult to mitigate (67). For example, it is hard to rule out the possibility that a future model could have ill intent toward humans (perhaps due to a malicious developer, or as an unexpected side-effect) that it hides until it is widely deployed (67). Additionally, it is important to avoid a scenario in which healthcare's legitimate needs and large market motivate AI developers to develop capabilities that ultimately prove detrimental to health overall. The weighing of risks and benefits to health should be performed by healthcare providers, policy makers, and independent researchers, rather than by companies that stand to profit more from one outcome than another.

Conclusion

In conclusion, the rapid evolution of large language models and their multimodal counterparts presents both promising opportunities and formidable challenges for healthcare. The expanding capabilities of these AI models, exemplified by GPT-4 and Med PaLM 2, demonstrate remarkable advancements in tasks ranging from medical question answering to medical image reporting. The integration of these models in areas like radiology and disease prediction has shown potential for significant improvements in clinical outcomes. However, this technological progress is not without its ethical, legal, and environmental ramifications. Furthermore, the considerable environmental impact of these models, particularly in terms of carbon emissions, raises concerns about their sustainability and the indirect health consequences of climate change.

The future of healthcare will undoubtedly be shaped by multimodal AI, and it is our collective responsibility to steer this course with a commitment to the betterment of human health.

Acknowledgements

The corresponding author attests that all listed authors meet authorship criteria and that no others meeting the criteria have been omitted.

Table 1 – Artificial Intelligence Concepts and Definitions

Concept	Description
Architectures	Artificial Intelligence systems can be broken down into architectures, model families, models, and model instances. Common AI architectures include Convolution Neural Networks (CNNs) in image or pattern recognition tasks and Transformers in sequence-to-sequence prediction and classification tasks. Various transformer model families have been developed such as GPT and Flan-T5. Within model families models are differentiated by their size or model improvements such as flan-t5-small, flan-t5-base and flan-t5-large or GPT-2 and GPT-4. Model instances refer to a specific set of weights trained in a model, often for a specific problem domain such as Llama-2-13B-Instruct.
Vectors	Vectors are essential components of transformers. A vector is a list of numbers. In natural language processing, words are converted into vectors to be processed by a machine learning model. These vectors are optimized to represent the semantic and syntactic characteristics of the words, allowing the model to understand and manipulate language. High-dimensional (long) vectors can capture more nuanced relationships and meanings than lower-dimensional ones.
Transformers	In a text-processing application such as a chatbot, the input to a transformer is a sequence of words. The transformer converts each word into a vector, i.e. an ordered list of numbers. Once a word has been converted into a vector, the transformer repeatedly updates each vector with information from other vectors, corresponding to words that appear earlier in the sequence. It does this selectively using an attention mechanism (described below). This adds context to the current word. For example, at the end of the sequence, "... confirms the right ventricle is enlarged", the attention mechanism might add together the vectors for "right", "ventricle", and "enlarged" to produce a new vector that means roughly "enlarged right ventricle". Large transformers repeatedly combine and transform the inputs, ultimately producing a prediction about which word should come next. Mathematically, these combinations and transformations rely on matrix multiplications.
Attention	The attention mechanism is based on the idea of key-value pairs, which are widely used in information systems. A hospital information system might use key-value pairs to keep track of which patient is in which bed, e.g. the bed "Red-North-12" (a key) may have the patient with ID "12345" (a value). If a user provides a bed identifier (a "query") that matches one of the keys, the system will provide the corresponding value (the patient ID). Transformer attention uses partial matches between queries and keys. This is roughly analogous to querying the bed system with "Red-North" and getting back a list of all the patients in that unit. In a transformer, the queries, keys, and values are vectors. The transformer creates a query vector by multiplying the current word's vector by a matrix (the matrix elements are learned network parameters). Similarly, it creates the keys and values by multiplying the previous words' vectors by different matrices. It calculates partial matches according to the similarity between the query vector and all the key vectors. The transformer then combines the value vectors that correspond to any partially matching keys – not as a list, but as a weighted sum. This is how previous sequence elements are combined to add context and other nuance to the current word vector.
Encoding	Encoding in the context of transformers refers to the process of converting input data (like text, images, sound) into a numerical format that can be processed by the model. In NLP, this often involves converting words (or part of words) into vectors. Encoding is a critical step in preparing data for complex tasks such as medical question answering, report generation, and medical summarisation.

Table 2 - Opportunities and Challenges of Multimodal AI applied to healthcare

Opportunities		Challenges	Solutions to Mitigate Challenges
Effectively process and interpret various data types such as text, images, and structured data.		Data Ownership and Privacy	Implement robust data governance frameworks and ensure informed patient consent.
Enhanced Diagnostic and Predictive Capabilities due to integration of additional data from diverse data types		Transparency	Create trust in medical AI systems through the application of rigorous randomized controlled trials (68) and involvement of relevant stakeholders (56,57).
Early Detection of Diseases		Increasing Inequity	Low- and Middle-Income Countries continue to be underrepresented in datasets and access to resources and compute. Continue working with communities to locate compute in those developing economies so that they can benefit from the AI advances.
Excellent Performance on Standard Benchmarks		Environmental Impact	Prioritize energy-efficient model architectures. Choose cloud services that use renewable energy sources.
Continuous Improvement with Scale		Hallucinations and Errors	Incorporate guardrails when developing and deploying models (69). This is an active research area (70,71), so it will be important for system implementers to remain up to date on developments.
Integration into Clinical Workflows		Integration into Clinical Workflows	Both an opportunity and a challenge. Requires alignment with existing practices, overcoming resistance to change among healthcare professionals, and especially training healthcare professionals how to use AI systems safely and effectively (72).

Figure 1 – Patient Scenarios without AI, with Single Modal AI and with Multimodal AI

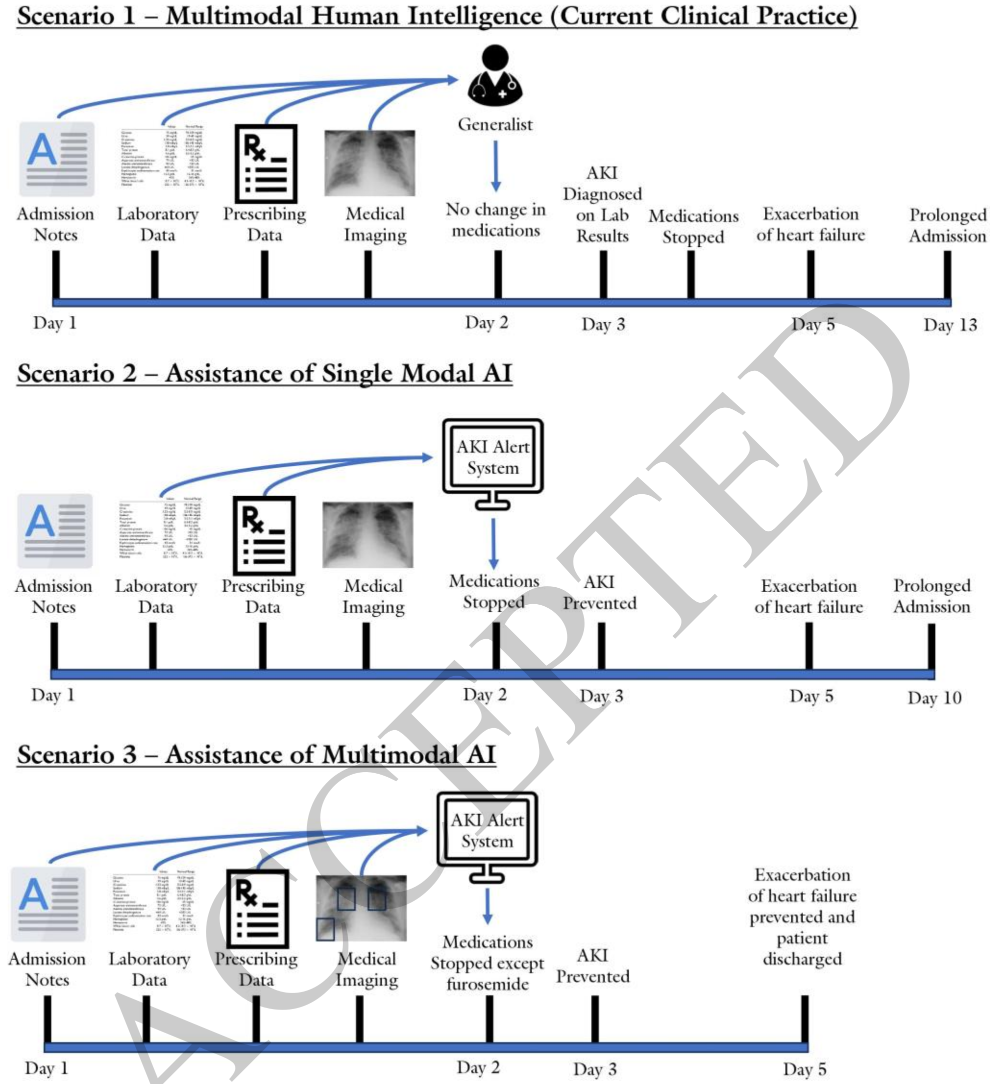

Figures 2 – A future multimodal Transformer based Acute Kidney Injury Alert System

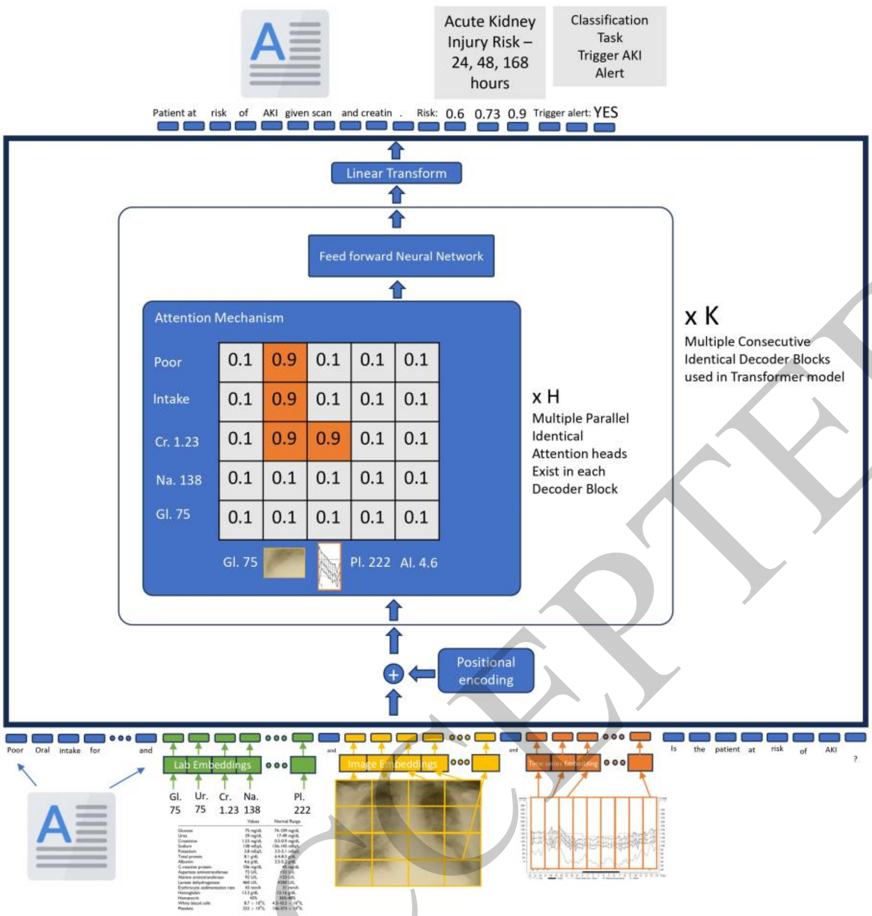

Figure 2 Legend

This figure depicts a multimodal Transformer-based system designed to predict acute kidney injury (AKI) risk by integrating text, structured data (e.g. lab tests), images (e.g. chest x-ray), and time series data. These are converted into numerical vectors through specific embedding layers. Positional encoding adds order information to these vectors, which are processed by the attention mechanism. The heatmap shows how the model focuses on relevant data parts, using keys, queries, and values to compute attention scores. Transformed vectors pass through a feed-forward neural network for pattern learning. The architecture includes multiple identical decoder blocks for iterative refinement. At the top, the classification task predicts AKI risk at 24, 48, and 168 hours, triggering alerts if risk thresholds are exceeded.

References

1. Rajpurkar P, Chen E, Banerjee O, Topol EJ. AI in health and medicine. Nat Med. 2022 Jan;28(1):31–8.
2. Sisson JC, Schoomaker EB, Ross JC. Clinical Decision Analysis: The Hazard of Using Additional Data. JAMA. 1976 Sep 13;236(11):1259–63.
3. Artificial Intelligence (AI) and Machine Learning (ML) in Medical Devices [Internet]. U.S. FOOD & DRUG ADMINISTRATION; 2020 [cited 2023 Sep 3]. Available from: https://www.fda.gov/media/142998/download
4. Meskó B, Görög M. A short guide for medical professionals in the era of artificial intelligence. NPJ Digit Med. 2020 Sep 24;3:126.
5. Tomašev N, Glorot X, Rae JW, Zielinski M, Askham H, Saraiva A, et al. A clinically applicable approach to continuous prediction of future acute kidney injury. Nature. 2019 Aug;572(7767):116–9.
6. Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, et al. Attention is All you Need. In: Advances in Neural Information Processing Systems [Internet]. Curran Associates, Inc.; 2017 [cited 2023 May 31]. Available from: https://proceedings.neurips.cc/paper_files/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html
7. Li Y, Wang H, Yerebakan HZ, Shinagawa Y, Luo Y. FHIR-GPT Enhances Health Interoperability with Large Language Models [Internet]. medRxiv; 2023 [cited 2023 Dec 14]. p. 2023.10.17.23297028. Available from: https://www.medrxiv.org/content/10.1101/2023.10.17.23297028v3
8. Zhao WX, Zhou K, Li J, Tang T, Wang X, Hou Y, et al. A Survey of Large Language Models [Internet]. arXiv; 2023 [cited 2023 Aug 1]. Available from: http://arxiv.org/abs/2303.18223
9. OpenAI. GPT-4 Technical Report [Internet]. arXiv; 2023 [cited 2023 May 31]. Available from: http://arxiv.org/abs/2303.08774
10. BioMedLM [Internet]. [cited 2023 Oct 17]. Available from: https://crfm.stanford.edu/2022/12/15/biomedlm.html
11. Singhal K, Tu T, Gottweis J, Sayres R, Wulczyn E, Hou L, et al. Towards Expert-Level Medical Question Answering with Large Language Models [Internet]. arXiv; 2023 [cited 2023 Oct 17]. Available from: http://arxiv.org/abs/2305.09617
12. Nori H, King N, McKinney SM, Carignan D, Horvitz E. Capabilities of GPT-4 on Medical Challenge Problems [Internet]. arXiv; 2023 [cited 2023 May 31]. Available from: http://arxiv.org/abs/2303.13375
13. Chen Z, Cano AH, Romanou A, Bonnet A, Matoba K, Salvi F, et al. MEDITRON-70B: Scaling Medical Pretraining for Large Language Models [Internet]. arXiv; 2023 [cited 2023 Dec 14]. Available from: http://arxiv.org/abs/2311.16079
14. Ji Z, Lee N, Frieske R, Yu T, Su D, Xu Y, et al. Survey of Hallucination in Natural Language Generation. ACM Comput Surv. 2023 Mar 3;55(12):248:1-248:38.
15. Taylor R, Kardas M, Cucurull G, Scialom T, Hartshorn A, Saravia E, et al. Galactica: A Large Language Model for Science [Internet]. arXiv; 2022 [cited 2023 Aug 1]. Available from: http://arxiv.org/abs/2211.09085
16. Greene T. 2022. Available from: https://twitter.com/mrgreene1977/status/1593278664161996801

17. McKenna N, Li T, Cheng L, Hosseini MJ, Johnson M, Steedman M. Sources of Hallucination by Large Language Models on Inference Tasks [Internet]. arXiv; 2023 [cited 2023 Jun 1]. Available from: http://arxiv.org/abs/2305.14552
18. Abdou M, Kulmizev A, Hershcovich D, Frank S, Pavlick E, Søgaard A. Can Language Models Encode Perceptual Structure Without Grounding? A Case Study in Color [Internet]. arXiv; 2021 [cited 2023 May 30]. Available from: http://arxiv.org/abs/2109.06129
19. Lee P, Goldberg C, Kohane I. The AI Revolution in Medicine: GPT-4 and Beyond. 1st edition. Hoboken: Pearson; 2023. 304 p.
20. Tay Y, Dehghani M, Rao J, Fedus W, Abnar S, Chung HW, et al. Scale Efficiently: Insights from Pre-training and Fine-tuning Transformers [Internet]. arXiv; 2022 [cited 2024 Jan 2]. Available from: http://arxiv.org/abs/2109.10686
21. Ouyang L, Wu J, Jiang X, Almeida D, Wainwright CL, Mishkin P, et al. Training language models to follow instructions with human feedback [Internet]. arXiv; 2022 [cited 2023 Dec 14]. Available from: http://arxiv.org/abs/2203.02155
22. Salvagno M, Taccone FS, Gerli AG. Can artificial intelligence help for scientific writing? Critical Care. 2023 Feb 25;27(1):75.
23. Ma J, He Y, Li F, Han L, You C, Wang B. Segment Anything in Medical Images [Internet]. arXiv; 2023 [cited 2023 Dec 1]. Available from: http://arxiv.org/abs/2304.12306
24. Gong Y, Chung YA, Glass J. AST: Audio Spectrogram Transformer [Internet]. arXiv; 2021 [cited 2023 May 30]. Available from: http://arxiv.org/abs/2104.01778
25. Enarvi S, Amoia M, Del-Agua Teba M, Delaney B, Diehl F, Hahn S, et al. Generating Medical Reports from Patient-Doctor Conversations Using Sequence-to-Sequence Models. In: Proceedings of the First Workshop on Natural Language Processing for Medical Conversations [Internet]. Online: Association for Computational Linguistics; 2020 [cited 2023 Jun 1]. p. 22–30. Available from: https://aclanthology.org/2020.nlpmc-1.4
26. Soltau H, Wang M, Shafran I, Shafey LE. Understanding Medical Conversations: Rich Transcription, Confidence Scores & Information Extraction [Internet]. arXiv; 2021 [cited 2023 Jun 1]. Available from: http://arxiv.org/abs/2104.02219
27. Radford A, Kim JW, Xu T, Brockman G, McLeavey C, Sutskever I. Robust Speech Recognition via Large-Scale Weak Supervision [Internet]. arXiv; 2022 [cited 2023 May 30]. Available from: http://arxiv.org/abs/2212.04356
28. Chen S, Wang C, Chen Z, Wu Y, Liu S, Chen Z, et al. WavLM: Large-Scale Self-Supervised Pre-Training for Full Stack Speech Processing. IEEE Journal of Selected Topics in Signal Processing. 2022 Oct;16(6):1505–18.
29. Yang C, Westover MB, Sun J. BIOT: Cross-data Biosignal Learning in the Wild [Internet]. arXiv; 2023 [cited 2023 May 30]. Available from: http://arxiv.org/abs/2305.10351
30. Dosovitskiy A, Beyer L, Kolesnikov A, Weissenborn D, Zhai X, Unterthiner T, et al. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale [Internet]. arXiv; 2021 [cited 2023 May 30]. Available from: http://arxiv.org/abs/2010.11929
31. Ramesh A, Pavlov M, Goh G, Gray S, Voss C, Radford A, et al. Zero-Shot Text-to-Image Generation. In: Proceedings of the 38th International Conference on Machine Learning [Internet]. PMLR; 2021 [cited 2023 May 30]. p. 8821–31. Available from: https://proceedings.mlr.press/v139/ramesh21a.html
32. OpenAI. Be My Eyes uses GPT-4 to transform visual accessibility [Internet]. 2023. Available from: https://openai.com/customer-stories/be-my-eyes
33. Han T, Adams LC, Bressem K, Busch F, Huck L, Nebelung S, et al. Comparative Analysis of GPT-4Vision, GPT-4 and Open Source LLMs in Clinical Diagnostic Accuracy: A Benchmark Against Human Expertise [Internet]. medRxiv; 2023 [cited 2024 Jul 16]. p.

2023.11.03.23297957. Available from: https://www.medrxiv.org/content/10.1101/2023.11.03.23297957v2
34. GPT-4V(ision) system card [Internet]. [cited 2023 Nov 21]. Available from: https://openai.com/research/gpt-4v-system-card
35. Tu T, Azizi S, Driess D, Schaekermann M, Amin M, Chang PC, et al. Towards Generalist Biomedical AI [Internet]. arXiv; 2023 [cited 2023 Aug 22]. Available from: http://arxiv.org/abs/2307.14334
36. Xu S, Yang L, Kelly C, Sieniek M, Kohlberger T, Ma M, et al. ELIXR: Towards a general purpose X-ray artificial intelligence system through alignment of large language models and radiology vision encoders.
37. Belyaeva A, Cosentino J, Hormozdiari F, Eswaran K, Shetty S, Corrado G, et al. Multimodal LLMs for health grounded in individual-specific data [Internet]. arXiv; 2023 [cited 2023 Nov 21]. Available from: http://arxiv.org/abs/2307.09018
38. Min B, Ross H, Sulem E, Veyseh APB, Nguyen TH, Sainz O, et al. Recent Advances in Natural Language Processing via Large Pre-trained Language Models: A Survey. ACM Comput Surv. 2023 Sep 14;56(2):30:1-30:40.
39. Dao T, Fu DY, Ermon S, Rudra A, Ré C. FlashAttention: Fast and Memory-Efficient Exact Attention with IO-Awareness [Internet]. arXiv; 2022 [cited 2023 Dec 1]. Available from: http://arxiv.org/abs/2205.14135
40. Beltagy I, Peters ME, Cohan A. Longformer: The Long-Document Transformer [Internet]. arXiv; 2020 [cited 2023 Dec 1]. Available from: http://arxiv.org/abs/2004.05150
41. Fedus W, Zoph B, Shazeer N. Switch transformers: scaling to trillion parameter models with simple and efficient sparsity. J Mach Learn Res. 2022 Jan 1;23(1):120:5232-120:5270.
42. Kaplan J, McCandlish S, Henighan T, Brown TB, Chess B, Child R, et al. Scaling Laws for Neural Language Models [Internet]. arXiv; 2020 [cited 2023 May 31]. Available from: http://arxiv.org/abs/2001.08361
43. Hoffmann J, Borgeaud S, Mensch A, Buchatskaya E, Cai T, Rutherford E, et al. Training Compute-Optimal Large Language Models [Internet]. arXiv; 2022 [cited 2023 May 31]. Available from: http://arxiv.org/abs/2203.15556
44. Srivastava A, Rastogi A, Rao A, Shoeb AAM, Abid A, Fisch A, et al. Beyond the Imitation Game: Quantifying and extrapolating the capabilities of language models [Internet]. arXiv; 2022 [cited 2023 May 31]. Available from: http://arxiv.org/abs/2206.04615
45. Wei J, Tay Y, Bommasani R, Raffel C, Zoph B, Borgeaud S, et al. Emergent Abilities of Large Language Models [Internet]. arXiv; 2022 [cited 2023 May 31]. Available from: http://arxiv.org/abs/2206.07682
46. Schaeffer R, Miranda B, Koyejo S. Are Emergent Abilities of Large Language Models a Mirage? [Internet]. arXiv; 2023 [cited 2023 May 31]. Available from: http://arxiv.org/abs/2304.15004
47. Chowdhery A, Narang S, Devlin J, Bosma M, Mishra G, Roberts A, et al. PaLM: Scaling Language Modeling with Pathways [Internet]. arXiv; 2022 [cited 2023 May 31]. Available from: http://arxiv.org/abs/2204.02311
48. Villalobos P, Sevilla J, Heim L, Besiroglu T, Hobbhahn M, Ho A. Will we run out of data? An analysis of the limits of scaling datasets in Machine Learning [Internet]. arXiv; 2022 [cited 2023 May 31]. Available from: http://arxiv.org/abs/2211.04325
49. Morley J, Machado CCV, Burr C, Cowls J, Joshi I, Taddeo M, et al. The ethics of AI in health care: A mapping review. Social Science & Medicine. 2020 Sep 1;260:113172.

50. Minssen T, Vayena E, Cohen IG. The Challenges for Regulating Medical Use of ChatGPT and Other Large Language Models. JAMA. 2023 Jul 25;330(4):315–6.
51. Harrer S. Attention is not all you need: the complicated case of ethically using large language models in healthcare and medicine. eBioMedicine [Internet]. 2023 Apr 1 [cited 2023 Nov 8];90. Available from: https://www.thelancet.com/journals/ebiom/article/PIIS2352-3964(23)00077-4/fulltext?ref=dedataverbinders.nl
52. Sallam M. ChatGPT Utility in Healthcare Education, Research, and Practice: Systematic Review on the Promising Perspectives and Valid Concerns. Healthcare. 2023 Jan;11(6):887.
53. Meskó B, Topol EJ. The imperative for regulatory oversight of large language models (or generative AI) in healthcare. npj Digit Med. 2023 Jul 6;6(1):1–6.
54. Li H, Moon JT, Purkayastha S, Celi LA, Trivedi H, Gichoya JW. Ethics of large language models in medicine and medical research. The Lancet Digital Health. 2023 Jun 1;5(6):e333–5.
55. Raji ID, Gebru T, Mitchell M, Buolamwini J, Lee J, Denton E. Saving Face: Investigating the Ethical Concerns of Facial Recognition Auditing. In: Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society [Internet]. New York, NY, USA: Association for Computing Machinery; 2020 [cited 2023 Nov 21]. p. 145–51. (AIES '20). Available from: https://dl.acm.org/doi/10.1145/3375627.3375820
56. Metcalf J, Moss E, Watkins EA, Singh R, Elish MC. Algorithmic Impact Assessments and Accountability: The Co-construction of Impacts. In: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency [Internet]. New York, NY, USA: Association for Computing Machinery; 2021 [cited 2023 Nov 21]. p. 735–46. (FAccT '21). Available from: https://dl.acm.org/doi/10.1145/3442188.3445935
57. Birhane A, Ruane E, Laurent T, S. Brown M, Flowers J, Ventresque A, et al. The Forgotten Margins of AI Ethics. In: Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency [Internet]. New York, NY, USA: Association for Computing Machinery; 2022 [cited 2023 Nov 21]. p. 948–58. (FAccT '22). Available from: https://dl.acm.org/doi/10.1145/3531146.3533157
58. Haines A, Ebi KL, Smith KR, Woodward A. Health risks of climate change: act now or pay later. The Lancet. 2014 Sep 20;384(9948):1073–5.
59. Ebi KL, Boyer C, Ogden N, Paz S, Berry P, Campbell-Lendrum D, et al. Burning embers: synthesis of the health risks of climate change. Environ Res Lett. 2021 Mar;16(4):044042.
60. Patterson D, Gonzalez J, Le Q, Liang C, Munguia LM, Rothchild D, et al. Carbon Emissions and Large Neural Network Training [Internet]. arXiv; 2021 [cited 2023 Jul 25]. Available from: http://arxiv.org/abs/2104.10350
61. Lacoste A, Luccioni A, Schmidt V, Dandres T. Quantifying the Carbon Emissions of Machine Learning [Internet]. arXiv; 2019 [cited 2023 Jul 25]. Available from: http://arxiv.org/abs/1910.09700
62. Strubell E, Ganesh A, McCallum A. Energy and Policy Considerations for Deep Learning in NLP [Internet]. arXiv; 2019 [cited 2023 Jul 25]. Available from: http://arxiv.org/abs/1906.02243
63. Weidinger L, Uesato J, Rauh M, Griffin C, Huang PS, Mellor J, et al. Taxonomy of Risks posed by Language Models. In: Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency [Internet]. New York, NY, USA: Association for Computing Machinery; 2022 [cited 2023 Aug 11]. p. 214–29. (FAccT '22). Available from: https://dl.acm.org/doi/10.1145/3531146.3533088

64. Reinsel D, Gantz J, Rydning J. The digitization of the world from edge to core [Internet]. International Data Corporation; Report No.: #US44413318. Available from: https://www.seagate.com/files/www-content/our-story/trends/files/idc-seagate-dataage-whitepaper.pdf
65. Karliner J, Slotterback S, Boyd R, Ashby B, Steele K, Wang J. Health care's climate footprint: the health sector contribution and opportunities for action. European Journal of Public Health. 2020 Sep 1;30(Supplement_5):ckaa165.843.
66. International Energy Agency. CO2 Emissions in 2022 [Internet]. Available from: https://www.iea.org/reports/co2-emissions-in-2022
67. Shevlane T, Farquhar S, Garfinkel B, Phuong M, Whittlestone J, Leung J, et al. Model evaluation for extreme risks [Internet]. arXiv; 2023 [cited 2023 Aug 11]. Available from: http://arxiv.org/abs/2305.15324
68. Byrne DW. Artificial intelligence for improved patient outcomes: principles for moving forward with rigorous science. Philadelphia: Wolters Kluwer; 2024.
69. Dong Y, Mu R, Jin G, Qi Y, Hu J, Zhao X, et al. Building Guardrails for Large Language Models [Internet]. arXiv; 2024 [cited 2024 Jul 24]. Available from: http://arxiv.org/abs/2402.01822
70. Farquhar S, Kossen J, Kuhn L, Gal Y. Detecting hallucinations in large language models using semantic entropy. Nature. 2024 Jun;630(8017):625–30.
71. McAleese N, Pokorny M, Uribe JFC. LLM Critics Help Catch LLM Bugs.
72. Elish MC, Watkins EA. Repairing Innovation: A Study of Integrating AI in Clinical Care.

ASN Journal Disclosure Form

As per ASN journal policy, I have disclosed any financial relationships or commitments I have held in the past 36 months as included below. I have listed my Current Employer below to indicate there is a relationship requiring disclosure. If no relationship exists, my Current Employer is not listed.

C. Judge has nothing to disclose.

I understand that the information above will be published within the journal article, if accepted, and that failure to comply and/or to accurately and completely report the potential financial conflicts of interest could lead to the following: 1) Prior to publication, article rejection, or 2) Post-publication, sanctions ranging from, but not limited to, issuing a correction, reporting the inaccurate information to the authors' institution, banning authors from submitting work to ASN journals for varying lengths of time, and/or retraction of the published work.

Name: Conor S. Judge
Manuscript ID: K360-2024-000329R1
Manuscript Title: Multimodal Artificial Intelligence in Medicine
Date of Completion: July 30, 2024
Disclosure Updated Date: May 22, 2024

ASN Journal Disclosure Form

As per ASN journal policy, I have disclosed any financial relationships or commitments I have held in the past 36 months as included below. I have listed my Current Employer below to indicate there is a relationship requiring disclosure. If no relationship exists, my Current Employer is not listed.

L. Kiely has nothing to disclose.

I understand that the information above will be published within the journal article, if accepted, and that failure to comply and/or to accurately and completely report the potential financial conflicts of interest could lead to the following: 1) Prior to publication, article rejection, or 2) Post-publication, sanctions ranging from, but not limited to, issuing a correction, reporting the inaccurate information to the authors' institution, banning authors from submitting work to ASN journals for varying lengths of time, and/or retraction of the published work.

Name: Lisa Kiely
Manuscript ID: K360-2024-000329R1
Manuscript Title: Multimodal Artificial Intelligence in Medicine
Date of Completion: July 24, 2024
Disclosure Updated Date: July 24, 2024

ASN Journal Disclosure Form

As per ASN journal policy, I have disclosed any financial relationships or commitments I have held in the past 36 months as included below. I have listed my Current Employer below to indicate there is a relationship requiring disclosure. If no relationship exists, my Current Employer is not listed.

F. Krewer reports the following:
Employer: 9th Impact; ICON plc

I understand that the information above will be published within the journal article, if accepted, and that failure to comply and/or to accurately and completely report the potential financial conflicts of interest could lead to the following: 1) Prior to publication, article rejection, or 2) Post-publication, sanctions ranging from, but not limited to, issuing a correction, reporting the inaccurate information to the authors' institution, banning authors from submitting work to ASN journals for varying lengths of time, and/or retraction of the published work.

Name: Finn Krewer
Manuscript ID: K360-2024-000329R1
Manuscript Title: Multimodal Artificial Intelligence in Medicine
Date of Completion: July 24, 2024
Disclosure Updated Date: May 22, 2024

ASN American Society of Nephrology JASN Journal of the American Society of Nephrology KIDNEY360 Accessing Our World From Every Angle CJASN Clinical Journal of the American Society of Nephrology

ASN Journal Disclosure Form

As per ASN journal policy, I have disclosed any financial relationships or commitments I have held in the past 36 months as included below. I have listed my Current Employer below to indicate there is a relationship requiring disclosure. If no relationship exists, my Current Employer is not listed.

M. O'donnell reports the following:
Employer: NUI Galway

I understand that the information above will be published within the journal article, if accepted, and that failure to comply and/or to accurately and completely report the potential financial conflicts of interest could lead to the following: 1) Prior to publication, article rejection, or 2) Post-publication, sanctions ranging from, but not limited to, issuing a correction, reporting the inaccurate information to the authors' institution, banning authors from submitting work to ASN journals for varying lengths of time, and/or retraction of the published work.

Name: Martin O'donnell
Manuscript ID: K360-2024-000329R1
Manuscript Title: Multimodal Artificial Intelligence in Medicine
Date of Completion: August 11, 2024
Disclosure Updated Date: August 11, 2024

ASN
American Society of Nephrology
JASN
Journal of the American Society of Nephrology
KIDNEY360
Accessing Our World From Every Angle
CJASN
Clinical Journal of the American Society of Nephrology

ASN Journal Disclosure Form

As per ASN journal policy, I have disclosed any financial relationships or commitments I have held in the past 36 months as included below. I have listed my Current Employer below to indicate there is a relationship requiring disclosure. If no relationship exists, my Current Employer is not listed.

D. Sexton reports the following:
Employer: School of Medicine, Trinity College Dublin, Ireland.; Consultancy: AstraZeneca ; Boehringer Ingelheim; Takeda, Vifor; Honoraria: Boehringer Ingelheim; Takeda; and Advisory or Leadership Role: Honorary Secretary of the Irish Nephrology Society which is a not for profit organisation, it is a registered charity.

I understand that the information above will be published within the journal article, if accepted, and that failure to comply and/or to accurately and completely report the potential financial conflicts of interest could lead to the following: 1) Prior to publication, article rejection, or 2) Post-publication, sanctions ranging from, but not limited to, issuing a correction, reporting the inaccurate information to the authors' institution, banning authors from submitting work to ASN journals for varying lengths of time, and/or retraction of the published work.

Name: Donal J. Sexton
Manuscript ID: K360-2024-000329R1
Manuscript Title: Multimodal Artificial Intelligence in Medicine
Date of Completion: July 24, 2024
Disclosure Updated Date: July 24, 2024

ASN American Society of Nephrology
JASN Journal of the American Society of Nephrology
KIDNEY360 Accessing Our World From Every Angle
CJASN Clinical Journal of the American Society of Nephrology

ASN Journal Disclosure Form

As per ASN journal policy, I have disclosed any financial relationships or commitments I have held in the past 36 months as included below. I have listed my Current Employer below to indicate there is a relationship requiring disclosure. If no relationship exists, my Current Employer is not listed.

J. Skorburg has nothing to disclose.

I understand that the information above will be published within the journal article, if accepted, and that failure to comply and/or to accurately and completely report the potential financial conflicts of interest could lead to the following: 1) Prior to publication, article rejection, or 2) Post-publication, sanctions ranging from, but not limited to, issuing a correction, reporting the inaccurate information to the authors' institution, banning authors from submitting work to ASN journals for varying lengths of time, and/or retraction of the published work.

Name: Joshua August Skorburg
Manuscript ID: K360-2024-000329R1
Manuscript Title: Multimodal Artificial Intelligence in Medicine
Date of Completion: July 24, 2024
Disclosure Updated Date: July 24, 2024

ASN American Society of Nephrology | JASN Journal of the American Society of Nephrology | KIDNEY360 Accessing Our World From Every Angle | CJASN Clinical Journal of the American Society of Nephrology

ASN Journal Disclosure Form

As per ASN journal policy, I have disclosed any financial relationships or commitments I have held in the past 36 months as included below. I have listed my Current Employer below to indicate there is a relationship requiring disclosure. If no relationship exists, my Current Employer is not listed.

G. Taylor reports the following:
Consultancy: CentML; Royal Bank of Canada; and Ownership Interest: CentML; DentiAI; LatentAI.

I understand that the information above will be published within the journal article, if accepted, and that failure to comply and/or to accurately and completely report the potential financial conflicts of interest could lead to the following: 1) Prior to publication, article rejection, or 2) Post-publication, sanctions ranging from, but not limited to, issuing a correction, reporting the inaccurate information to the authors' institution, banning authors from submitting work to ASN journals for varying lengths of time, and/or retraction of the published work.

Name: Graham William Taylor
Manuscript ID: K360-2024-000329R1
Manuscript Title: Multimodal Artificial Intelligence in Medicine
Date of Completion: August 14, 2024
Disclosure Updated Date: August 14, 2024

ASN Journal Disclosure Form

As per ASN journal policy, I have disclosed any financial relationships or commitments I have held in the past 36 months as included below. I have listed my Current Employer below to indicate there is a relationship requiring disclosure. If no relationship exists, my Current Employer is not listed.

B. Tripp reports the following:
Employer: University of Waterloo; Ownership Interest: 15165172 Canada Inc.; and Research Funding: Smile Digital Health.

I understand that the information above will be published within the journal article, if accepted, and that failure to comply and/or to accurately and completely report the potential financial conflicts of interest could lead to the following: 1) Prior to publication, article rejection, or 2) Post-publication, sanctions ranging from, but not limited to, issuing a correction, reporting the inaccurate information to the authors' institution, banning authors from submitting work to ASN journals for varying lengths of time, and/or retraction of the published work.

Name: Bryan Tripp
Manuscript ID: K360-2024-000329R1
Manuscript Title: Multimodal Artificial Intelligence in Medicine
Date of Completion: July 24, 2024
Disclosure Updated Date: July 24, 2024