Titre: Title:	A responsible framework for applying artificial intelligence on medical images and signals at the point-of-care: the PACS-AI platform
Auteurs: Authors:	Pascal Thériault-Lauzier, Denis Corbin, Olivier Tastet, Élodie Labrecque Langlais, Bahareh Taji, Guson Kang, Aun-Yeong Chong, Derek So, An Tang, Judy Wawira Gichoya, Sarath Chandar Anbil Parthipan, Pierre-Luc Déziel, Julie G. Hussin, Samuel Kadoury, & Robert Avram
Date:	2024
Type:	Article de revue / Article
Référence: Citation:	Thériault-Lauzier, P., Corbin, D., Tastet, O., Langlais, É. L., Taji, B., Kang, G., Chong, A.-Y., So, D., Tang, A., Gichoya, J. W., Anbil Parthipan, S. C., Déziel, P.-L., Hussin, J. G., Kadoury, S., & Avram, R. (2024). A responsible framework for applying artificial intelligence on medical images and signals at the point-of-care: the PACS-AI platform. Canadian Journal of Cardiology, 025 (39 pages). https://doi.org/10.1016/j.cjca.2024.05.025

Titre de la revue: Journal Title:	Canadian Journal of Cardiology
Maison d'édition: Publisher:	Elsevier
URL officiel: Official URL:	https://doi.org/10.1016/j.cjca.2024.05.025

Canadian Journal of Cardiology 40 (2024) 1828–1840

Review

A Responsible Framework for Applying Artificial Intelligence on Medical Images and Signals at the Point of Care: The PACS-AI Platform

Pascal Theriault-Lauzier, MD, PhD,[a,b] Denis Corbin, MSc,[c] Olivier Tastet, MSc,[c] Elodie Labrecque Langlais, MSc,[c] Bahareh Taji, PhD,[b] Guson Kang, MD,[a] Aun-Yeong Chong, MD,[b] Derek So, MD, MSc,[b] An Tang, MD, MSc,[d] Judy Wawira Gichoya, MD, MS,[e] Sarath Chandar, PhD,[g] Pierre-Luc Déziel, JD,[h] Julie G. Hussin, PhD,[c,g,h] Samuel Kadoury, PhD,[d,i] and Robert Avram, MD, MSc[c,f]

[a] *Division of Cardiovascular Medicine, Stanford School of Medicine, Palo Alto, California, USA*

[b] *Division of Cardiology, University of Ottawa Heart Institute, Ottawa, Ontario, Canada*

[c] *Montréal Heart Institute, Montréal, Québec, Canada*

[d] *Department of Radiology, Radiation Oncology and Nuclear Medicine, Université de Montréal, Montréal, Québec, Canada*

[e] *Department of Radiology and Imaging Sciences, Emory University School of Medicine, Atlanta, Georgia, USA*

[f] *Department of Medicine, Université de Montréal, Montréal, Québec, Canada*

[g] *Mila-Québec AI Institute, Montréal, Québec, Canada*

[h] *Faculty of Law, Université Laval, Québec, Québec, Canada*

[i] *Polytechnique Montréal, Montréal, Québec, Canada*

See editorial by Avram, et al., pages 1769-1773 of this issue.

ABSTRACT

The potential of artificial intelligence (AI) in medicine lies in its ability to enhance clinicians' capacity to analyse medical images, thereby improving diagnostic precision and accuracy and thus enhancing current tests. However, the integration of AI within health care is fraught with difficulties. Heterogeneity among health care system applications, reliance on proprietary closed-source software, and rising cybersecurity threats pose significant challenges. Moreover, before their deployment in clinical settings, AI models must demonstrate their effectiveness

RÉSUMÉ

Le potentiel de l'intelligence artificielle (IA) en médecine réside dans sa capacité à améliorer la capacité des cliniciens à analyser les images médicales, améliorant la précision et l'exactitude diagnostiques et renforçant ainsi les tests actuels. Cependant, l'intégration de l'IA dans les soins de santé est semée de difficultés. L'hétérogénéité parmi les applications des systèmes de soins de santé, la dépendance à l'égard de logiciels propriétaires à code source fermé et les menaces croissantes en matière de cybersécurité posent des défis importants.

The use of artificial intelligence (AI) in health care has sparked a wave of excitement for the analysis and interpretation of medical images and signals.[1,2] Simultaneously, there has been an explosive growth in the use of radiologic tests globally, with more than 5 billion diagnostic examinations performed each year.[3] This surge presents an opportunity for innovation in the health care industry by deploying real-world clinically actionable AI solutions. From 2021 to 2023, more than 10,000 scientific papers related to AI in health care were indexed annually on PubMed.[4] By the end of 2023, 692 AI-enabled medical devices had been approved by the US Food and Drug Administration (FDA).[5] AI systems have been shown to augment clinician awareness, screening, or diagnostic tasks,[6] predict clinically relevant events such as patient deterioration,[7] and augment preexisting tests by allowing the detection of diseases that human experts are not trained to identify.[2]

However, despite the enthusiasm over AI in health care and its stunning advances, a major challenge remains the

Received for publication February 19, 2024. Accepted May 26, 2024.

Corresponding author: Dr Robert Avram, Division of Cardiology, Department of Medicine, Montréal Heart Institute, 5000 Rue Belanger, Montréal, Québec H1T 1C8, Canada.

E-mail: robert.avram.md@gmail.com

See page 1839 for disclosure information.

https://doi.org/10.1016/j.cjca.2024.05.025

across a wide range of scenarios and must be validated by prospective studies, but doing so requires testing in an environment mirroring the clinical workflow, which is difficult to achieve without dedicated software. Finally, the use of AI techniques in health care raises significant legal and ethical issues, such as the protection of patient privacy, the prevention of bias, and the monitoring of the device's safety and effectiveness for regulatory compliance. This review describes challenges to AI integration in health care and provides guidelines on how to move forward. We describe an open-source solution that we developed that integrates AI models into the Picture Archives Communication System (PACS), called PACS-AI. This approach aims to increase the evaluation of AI models by facilitating their integration and validation with existing medical imaging databases. PACS-AI may overcome many current barriers to AI deployment and offer a pathway toward responsible, fair, and effective deployment of AI models in health care. In addition, we propose a list of criteria and guidelines that AI researchers should adopt when publishing a medical AI model to enhance standardisation and reproducibility.

En outre, avant d'être déployés en milieu clinique, les modèles d'IA doivent démontrer leur efficacité dans un large éventail de scénarios et être validés par des études prospectives, mais pour ce faire, il faut les tester dans un environnement reflétant le flux de travail clinique, ce qui est difficile à réaliser sans logiciel dédié. Enfin, l'utilisation de techniques d'IA dans les soins de santé soulève d'importantes questions juridiques et éthiques, telles que la protection de la vie privée des patients, la prévention des biais et le contrôle de la sécurité et de l'efficacité des dispositifs pour assurer la conformité aux réglementations. Cette étude décrit les défis posés par l'intégration de l'IA dans les soins de santé et fournit des lignes directrices sur la manière d'aller de l'avant. Nous décrivons une solution en source ouverte que nous avons développée et qui intègre des modèles d'IA dans le système d'archivage et de transmission d'images (PACS, de l'anglais 'Picture Archives Communication System'), appelé PACS-AI. Cette approche vise à améliorer l'évaluation des modèles d'IA en facilitant leur intégration et leur validation dans les bases de données d'imagerie médicale existantes. Le PACS-AI peut surmonter de nombreux obstacles au déploiement de l'IA actuels et offrir une voie vers un déploiement responsable, équitable et efficace des modèles d'IA dans les soins de santé. En outre, nous proposons une liste de critères et de lignes directrices que les chercheurs en IA devraient adopter lors de la publication d'un modèle d'IA médicale afin d'améliorer la normalisation et la reproductibilité.

translation of these AI-derived models, trained *in silico*, into clinical practice in an effective and responsible fashion.[8] Very few models developed by academia are integrated in the clinical environment such that they can be validated and integrated in the clinical workflow.[9] A limited number of models are commercially available and integrated within medical devices; these are often single-vendor, single-algorithm systems. In the future, this could result in a fragmented environment with dozens of independent platforms. Furthermore, open-source algorithms,[10] which could be valuable to test and adapt locally, are particularly difficult to implement in clinical environments owing to a lack of support from third-party vendors. In this article, we explore the many challenges facing the translation to clinical practice of AI algorithms in cardiology and medical imaging. We further describe a path forward using existing technology and conclude by describing Picture Archives Communication System (PACS)-AI, a new academia-led open-source platform for integrating AI in medical imaging.

Challenges to AI Implementation

Heterogeneous infrastructure

The clinical implementation of AI is strongly linked to the health care information technology (IT) infrastructure. Each health care institution has the flexibility to adopt software solutions from a wide array of publishers and device manufacturers. There are undeniable advantages to the flexibility afforded by such heterogeneous environments, such as the precise matching between implemented software solutions and clinical needs of patients treated at a given institution. However, this flexibility comes at a cost. The data formats generated by current health care platforms are often incompatible and lack standardisation.[11] This affects both the development and the deployment of AI models. Algorithms need data to be in a format similar to the one it used for training to permit analysis. However, different manufacturers may choose different acquisition parameters (eg, image resolution, pixel value normalisation, frames per second), apply different post-processing approaches (eg, reconstruction kernels, slice thickness reconstruction), or include unique metadata (eg, age or sex). Such heterogeneity in data formats can hinder the algorithm's ability to accurately interpret and analyse data, leading to inefficiencies or inaccuracies in its application. In a recent review of external validation studies published on radiologic image—based diagnosis, more than 80% of algorithms tested had a worse performance when applied to external data sets.[12] This performance degradation is partially explained by heterogeneous data formats. It limits the scalability and generalisability of AI solutions, because models may need to be retrained or significantly adjusted for each unique IT environment. Furthermore, diverse user interfaces and workflows of different health care IT systems can pose challenges in integrating AI tools in an intuitive and efficient manner for health professionals. If an AI solution does not seamlessly integrate with the existing workflow, it may face resistance by health care providers to adapt it, and they may lose confidence in its output.[13,14]

Proprietary closed-source software

Most software used in clinical medicine and cardiology are closed source and vendor specific. Closed-source software, also known as proprietary, refers to computer software of which the source code is not freely available or accessible to the public. For closed-source software, the rights to the source code are held exclusively by its creators, who control and restrict its use, modification, and distribution. Proprietary software typically does not allow users to modify or customize

the code to meet specific needs. This rigidity can be a major impediment in clinical settings, where AI models need to be tailored to suit unique medical practices, patient demographics, or specific types of data. AI models require access to large amounts of data for their development and validation. Although closed-source systems are typically well maintained and offer support, they may restrict access to data or make data extraction challenging; this in turn limits the ability to train and refine AI models effectively. This also means that AI models developed in academic settings cannot be directly integrated into such software, because AI researchers cannot modify the codebase of the software (eg, electronic medical record or image viewer) without explicit permission from the software provider and often involving hefty costs and prolonged negotiation processes. Closed-source environments can stifle innovation by limiting collaboration and sharing of knowledge and tools within the broader medical and scientific community. Open-source platforms, in contrast, can foster collaborative development and rapid advancement of AI technologies. Closed-source AI models have obvious financial advantages for their owners who may collect licencing fees and maintain a competitive advantage by keeping the details of the model architecture private. The above discussion is limited to narrow AI models, ie, relatively small models targeted at few specialised tasks. A distinction must be made when discussing much larger general AI models, which may pose risks if they become controlled by a user with nefarious intents. A review of these risks is outside the scope of the present article.

Cybersecurity risk

Health care organisations are increasingly the target of malicious attacks. Ransomware, a type of attack that prevents users from accessing their electronic systems by demanding a ransom to restore access, has grown in frequency and sophistication. From 2016 to 2021, 374 ransomware attacks have exposed the protected health information (PHI) of 42 million patients in the United States.[15] This not only represents a major financial burden, but also may result in unacceptable compromise to patient care. AI algorithms in health care often require access to large volumes of sensitive patient PHI. Any vulnerability in the AI system can lead to data breaches, compromising patient privacy and violating local (provinces, states), federal (eg, US Health Insurance Portability and Accountability Act), and international (eg, General Data Protection Regulation) regulations. The fear of such breaches makes health care providers cautious of implementing untested software packages. Addressing cybersecurity concerns requires significant investment in security infrastructure, regular audits, and staff training.[16] Allocation of resources can be a burden, especially for academic researchers looking to validate or implement AI models without access to a team of software developers with cybersecurity expertise. Cybersecurity risks and regulations represent a barrier to deployment of AI systems in a clinical setting.

Scientific evidence

Most of the early medical AI studies were validated by comparison with expert performance. However, those studies were largely retrospective and relied on historically labelled data sets for training and testing models.[17] The performance of the algorithms was often worse on data stemming from new sources as opposed to the development data set used for training.[12] It is the responsibility of developers and to regulators to ensure that models have been adequately validated before they are used clinically. For AI systems to be trusted by clinicians, prospective validation studies are needed to understand their performance in real-world situations, as well as illustrate how these algorithms are integrated in the busy clinical workflow. Furthermore, many medical AI manuscripts are published on preprint servers, thus foregoing the peer review process, and many studies do not conform with clinical research reporting guidelines.[18] High-quality peer-reviewed studies, such as multicentre randomised controlled trials with meaningful end points, are critical to the acceptance of AI models in clinical practice, as well as to demonstrate that the benefits of the AI system outweigh the risks.

Bias and data set shift

The issue of generalisability is related to that of bias, which is covered in a separate article in this issue of the *Canadian Journal of Cardiology*.[19] AI systems trained on historical data sets may inadvertently incorporate societal biases, amplify their adverse impacts on minority groups, and lead to discriminatory practices.[20] It is crucial for researchers to validate the performance of their systems in varied population subgroups including geographic location, sex, age, race, ethnic background, and socioeconomic status.[17] Furthermore, the population from which the model was derived may evolve over time and differ in important ways from the population on which the model is applied. This phenomenon, known as data set shift, can be caused by a variety of factors, including variation in data acquisition devices, changes in disease incidence, seasonal variability, evolution of new disease, changes in patient behaviour, migrations, and clinician behaviour.[21] Various organisations have published guidelines on "good machine learning practices," which includes guidelines for real-world performance monitoring and retraining of the model based on evolution of the target population.[22-25] However, retraining of AI systems can be problematic as it may result in worse performance despite a newer model version. A primary concern is the introduction of new biases or errors, especially if the retraining data is not representative or contains inaccuracies. Monitoring model performance after deployment must be performed to detect performance degradation and make corrections. Such postmarket surveillance is included in regulations.[22] Moreover, explainability of AI models plays an important role in addressing these challenges. Many AI methods are "black boxes," which means that humans cannot easily explain how decisions are made; this can cause problems for clinicians who need to understand and trust the suggestions made by AI systems.[26] By understanding how AI models arrive at their conclusions, developers and users can more effectively identify and mitigate biases and recognise when data shifts occur, although studies looking at providing explainable predictions to physicians have had mixed results.[27,28] This level of transparency is necessary for trust and efficacy in clinical applications.[25]

Regulation

Regulating AI in medicine involves a delicate balance. On the positive side, regulation ensures that AI systems are safe,

effective, and reliable, all of which fosters public trust and encourages standardised high-quality care. In addition, it helps to protect sensitive patient PHI and to mitigate legal as well as ethical concerns such as bias and discrimination.[29] However, regulation also poses some risks. It can delay innovation, making it harder for new and potentially beneficial technologies to reach the market. Increased regulatory processes can also add significant costs, affecting smaller innovators and many academic institutions. The risk of regulatory capture (or agency capture), where regulatory bodies prioritise the interests of the industry over public health, can lead to compromised regulations that do not adequately protect patients or may hinder competition and innovation by favouring established players.[30] Effective regulation should therefore aim to protect public health, promote health equity, and foster innovation while avoiding biases toward particular industry interests.[31]

Cost

The development of AI models in health care has far outpaced health economics evaluation of them. A recent systematic review[32] found that such formal health economics studies are limited in number and quality. They conclude that this may impede the clinical deployment of AI models in clinical settings. Although medical AI models may lead to time and cost savings due to improved efficiency, those can be outweighed by implementation costs. Therefore, optimising costs is of paramount importance to the deployment of clinical AI models. Specifically, developing a bespoke software platform for each new AI model in medicine is prohibitive in practice. The development process, including designing, coding, testing, and integrating a suite of AI models into existing health care IT systems may result in an economically nonviable solution without a preexisting overarching platform. Maintenance and updates, as well as ensuring data privacy and security, further increase expenses. These factors make it challenging for academics to validate their AI models clinically and translate their work into a clinical environment. Consequently, only a small proportion of AI innovations that are published end up being used by clinicians.

Solutions and Path Forward

In this section, we explore technologic, regulatory, and administrative solutions to implement AI algorithms responsibly in a clinical setting (Fig. 1).

Technologic considerations

Several technologic frameworks and concepts offer major opportunities for the development of a solution that is vendor agnostic yet integrates within the existing clinical workflow. The principle of interoperability is defined as the ability of 2 or more systems or components to exchange information and to use the information that has been exchanged.[11]

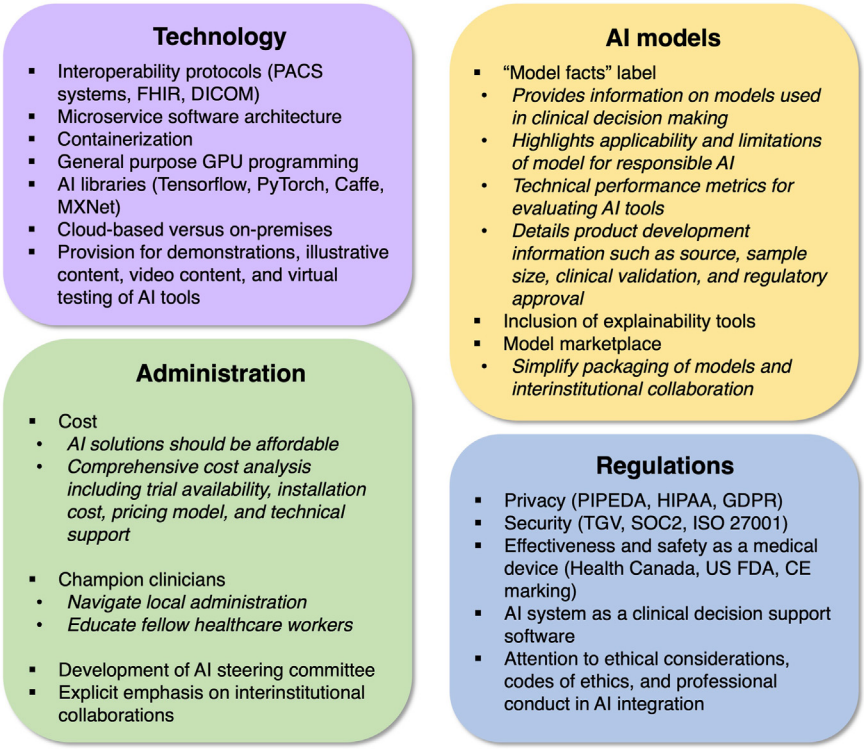

Figure 1. Illustration of the varied solutions that can be leveraged to solve the challenges to the clinical implementation of artificial intelligence models in clinical medicine and cardiology. AI, artificial intelligence; CE, Conformité Européenne; DICOM, Digital Imaging and Communications in Medicine; FDA, Food and Drug Administration; FHIR, Fast Healthcare Interoperability Resources; GDPR, General Data Protection Regulation; GPU, graphical processing unit; HIPAA, Health Insurance Portability and Accountability Act; ISO, International Organization for Standardization; PACS, Picture Archives Communication System; PIPEDA, Personal Information Protection and Electronic Documents Act; SOC2, Service Organization Control Type 2; TGV, Trousse Globale de Verification.

Interoperability is crucial for the development of digital health. Significant efforts have been invested in the development of standards for health care software. The Fast Healthcare Interoperability Resources (Health Level Seven International, Ann Arbor, MI) provide a set of rules and specifications for exchanging electronic health-related data in software solutions and is a rapidly growing standard. The Picture Archives Communication System (PACS) is a medical imaging technology to securely store, digitally transmit, and access medical images and related documentation. PACS systems eliminate the need to manually file, store, retrieve, or transport film jackets (traditional film-based imaging). PACS rely on the Digital Imaging and Communications in Medicine (DICOM; Medical Imaging and Technology Alliance, Arlington, VA), which is a standard protocol for the management and transmission of medical images and related data. DICOM ensures the interoperability of systems used to produce, store, display, process, send, retrieve, query, and print medical images, as well as to manage related workflow. Essentially, PACS uses the DICOM standard for handling, storing, printing, and transmitting information in medical imaging. The vast majority of hospitals worldwide use technology that is compatible with the DICOM standards.[33,34] The availability of PACS systems is nearly universal in Canada and in high-income countries. However, this is not the case in low- and middle-income countries[35]; however, web-based and cloud-based PACS are being implemented in low-resource settings. The application of AI models is thus enabled for medical images.

Recently, software engineering has embraced the microservice paradigm. This approach constructs an application as a collection of small independent services, each operating in its own process and interacting via simple lightweight methods.[36] These services, organised around specific business or mission functions, can be deployed independently through automated systems. This method simplifies building and maintenance of certain types of applications by breaking them down into smaller interworking pieces. Each component is developed separately, making the application a composite of the individual parts. This contrasts with the traditional monolithic approach, where an application is developed as a single indivisible unit. In addition, containers are a technology that has gained prominence in parallel with the microservice architecture. Containers encapsulate everything an application needs to run (including libraries, system tools, code, and runtime), ensuring that it behaves identically regardless of where it is deployed.[37] Containers and microservice architectures are closely related. Containers provide an ideal environment for deploying microservices because each container can host a separate microservice, ensuring isolation and reducing conflicts between services. This isolation simplifies development, testing, and deployment processes. Containers are lightweight, making them suitable for the dynamic and scalable nature of microservices. They offer a consistent environment from development through production, enhancing the efficiency and reliability of microservices deployment. This synergy between containers and microservices leads to more efficient, scalable, and maintainable applications.

One of the main factors leading to the AI revolution has been the use of hardware accelerators that can process large quantities of data in parallel. These hardware accelerators include graphical processing units (GPUs) that were initially developed for computer graphics processing but are very well suited to accelerate the training of neural networks. The use of low-level libraries for such general-purpose GPU programming is notoriously difficult; several machine learning and deep learning software libraries have been developed to make implementation of algorithms easier for AI practitioners. These include TensorFlow[38] (Google, Mountain View, CA), Caffe[39] (Berkeley AI Research, Berkeley, CA), PyTorch[40] (Meta, Menlo Park, CA), and MXNet[41] (Apache, Wakefield, MA). Each of these libraries can be easily used to create a containerized version of trained AI models, thus easing their integration in microarchitecture-based software applications. The Medical Open Network for Artificial Intelligence (MONAI) project has developed a set of standards based on best practices for packaging medical imaging AI models for research as well as clinical deployment.[42] MONAI is built as an extension of the PyTorch framework.

The physical location where AI models are executed is a related important consideration for medical applications. Two main options exist in this case, on-premises or cloud-based deployment.[43] Deploying AI models on the premises involves setting up the necessary infrastructure within a local environment (ie, within the hospital or hospital network). This approach offers full control over data and infrastructure, ensuring improved compliance with data privacy regulations. However, it requires significant up-front investment in hardware and infrastructure, it incurs recurring maintenance costs, and scalability can be limited. Conversely, cloud deployment uses external cloud services.[44] It offers scalability and flexibility, and typically lowers up-front costs but increases recurrent costs. Cloud services provide advanced computational resources that are easily upgradable to support additional AI models or additional centres. Major cloud infrastructure vendors provide platforms that offer streamlined hosting of AI models for training and predictions in the cloud; SageMaker (Amazon), Vertex AI (Google), and Azure AI (Microsoft) are examples of such services. However, this approach can raise concerns about data security and privacy, depending on the cloud service provider's policies, the location of the servers, and the nature of the data being processed.[45]

For medical applications, most providers consider that the safest option is to have PHI stored on premises exclusively and execute AI model predictions on local servers. However, as already noted, this increases infrastructure purchase and maintenance costs. The most flexible option is to deploy the entire application on cloud computing providers with encryption of information when saved to a hard drive (at rest) and while it is being transmitted from one service to another (in flight).[44] A third, hybrid, approach involves storing PHI on hospital servers while using cloud computing infrastructure to process AI inference requests. In this approach, the data is de-identified, such that no PHI leaves the hospital systems or, if PHI is transmitted, it is immediately deleted after use by the cloud system.

AI model considerations

Clinicians must understand an AI system's outputs and predictions to apply them responsibly to guide clinical

decisions. Recently, researchers have proposed the "model fact" label concept, aiming to provide clinicians with clear actionable information about machine learning models used in medical decision making.[46] Its format is designed to mirror product information for food, drugs, and devices. It highlights the importance of transparent communication of a model's capabilities, limitations, and appropriate usage contexts to prevent patient harm. The label is designed to be concise and includes sections such as model name, model version, data set characteristics during training and validation, performance in those data sets, indications for use, and contraindications. The model facts label seek to standardise the presentation of AI model information in clinical settings, emphasising the target population, input data type, training data location, and time and type of AI model used. This ensures that clinical end users can apply AI models, interpret the outputs, and integrate them appropriately in their decision making with full knowledge of their limitations.

Guidelines and checklists, including CLAIM (**C**heck**l**ist for **A**rtificial **I**ntelligence in **M**edical Imaging),[47] CONSORT-AI (**Con**solidated **S**tandards **o**f **R**eporting **T**rials—**A**rtificial **I**ntelligence),[48] DECIDE-AI (**D**evelopmental and **E**xploratory **C**linical **I**nvestigations of **De**cision Support Systems Driven by **A**rtificial **I**ntelligence),[49] and SPIRIT-AI (**S**tandard **P**rotocol **I**tems: **R**ecommendations for **I**nterventional **T**rials—**A**rtificial **I**ntelligence,[50] have been proposed to improve the quality of publications on AI models in medicine. High-quality peer-reviewed studies with meaningful end points are very important to the acceptance of AI models in clinical practice, as well as too demonstrating that the benefits of an AI system outweigh its risks. Such research requires the deployment of the model in the real-world clinical setting; yet the logistical and technical challenges involved in deploying such models in multiple institutions remain substantial.

Model prediction may vary substantially between institutions owing to the absence of uniform pre-processing routines, even when using the same data type and same model.[51] Furthermore, a mechanism for flagging accurate and inaccurate results would enable monitoring of the system's real-world performance. Post-deployment monitoring involves setting up systems to receive and incorporate feedback from health professionals and implementing continuous learning strategies to ensure regular updates to models.[52] Such post-market surveillance will become a standard requirement for responsible deployment of AI technology in health care.[22]

Regulatory environment

Several regulatory frameworks and certifications apply to the field of AI systems. Regulations can be grouped into 3 categories: privacy, security, and effectiveness/safety for clinical use. Several jurisdictions have developed mandatory privacy regulatory frameworks with which organisations handling user data must comply; these include the Health Insurance Portability and Accountability Act in the US, Personal Information Protection and Electronic Documents Act in Canada, and General Data Protection Regulation in the European Union. In Canada, provinces have enacted regulations regarding the management of PHI specifically.

Security certifications usually involve the analysis of the software design and its effectiveness in ensuring secure operation over time. In the US, the American Institute of Certified Public Accounts publishes the System and Organizational Controls (SOC) compliance standard, which is a set of voluntary auditing guidelines that help organisations prove the effectiveness of their information security program.[53] SOC is not specific to health care data and is voluntary in nature. In Québec, the Trousse Globale de Vérification (TGV) is a required certification that includes security, protection of personal information and performance of a given technology for use in the health care sector.[54] Both SOC and TGV involve intrusion testing by a third party.

The effectiveness of a given AI system at a particular clinically relevant task is regulated as a medical device by the FDA, Health Canada, and CE Mark in the European Union.[5,55] For example, in the US and Canada, devices are classified based on the risk level to a person's health. The FDA has 3 classes[56] (I = low; II = intermediate; and III = high), and Health Canada has 4 classes (I = minimal; II = low; III = intermediate; and IV = high). Both organisations have released guidance documents on clinical decision support software (CDSS), which includes several digital health and AI systems.[57,58] The document outlines 4 criteria based by which CDSS is considered a medical device under the FDA regulation. While under investigation and before commercialisation, medical devices are generally classified as "research use only," which does not require approval. Most proposed AI systems fall under this category until their use is pursued outside of clinical research for real-world diagnosis. In contrast, high-risk devices must undergo premarket approval, which is a rigourous process involving an independent demonstration of safety and effectiveness. Premarket approval often requires large-scale scientific studies of the effectiveness of a device before approval for diagnosis or treatment of a clinical condition. Depending on claims of effectiveness made by developers of AI systems and the associated risk to patients, the regulatory burden will vary substantially between different systems. Regulatory burden in terms of clinical efficacy will likely be isolated to the AI model itself, but privacy and security regulations also apply to the supporting infrastructure that will interface with the clinical data sources.

Administrative considerations

Owing to the extensive regulatory and usability requirements involved in the development of medical devices, the purchase or subscription cost of medical software is often significantly higher than in other industries. Canadian hospitals are public entities and as such are bound by public procurement laws.[59] To ensure a competitive and fair procurement process, any purchase above $25,000 requires a public tender process via a government portal.[60] This adds a layer of complexity and substantial delays to the rapid deployment of AI systems in health care. An ideal solution to deploy AI systems in Canadian hospitals should be inexpensive.

Because of the novelty of AI systems in health care, their benefits may not be understood fully by clinical end users within the next 5-10 years; expertise at the institution level may be lacking to fully evaluate an AI solution. We suggest that successful deployment of AI solutions can be accelerated by "champion clinicians," who understand the benefits of a

technology and can work closely with local administration, engineering departments, and IT departments to deploy a system within the local health care infrastructure. The role of champion clinicians also extends to educating other health professionals on AI model capabilities and limitations. Furthermore, hospitals may consider administrative entities, such as AI steering committees, to support research on AI systems and their deployment within hospital systems.

PACS-AI Platform

We propose the PACS-AI platform as a tool for the responsible deployment of AI models in the hospital setting for clinical and research purposes. The PACS-AI platform is a vendor-agnostic software that works as an interface between an existing clinical PACS system, where medical images and reports are stored, and AI models. The main objective of the platform is to enable automated, near real-time application of AI models on clinical images for use at the point of care. It offers a web application interface that clinicians can use to search for an imaging study stored on the hospital PACS and select a compatible AI model to be applied to the associated images. The application backend then fetches the relevant images, prepares the data accordingly, and calls for an AI inference to be performed (Table 1). Then the user is presented with the results in the web interface. Figure 2 summarises the features of the PACS-AI platform, and Figure 3 shows the simplified architecture of the application. We summarise the strategies used by the PACS-AI platform to solve the challenges facing the clinical integration of AI models in Table 2.

We aim to further develop the PACS-AI platform into an open-source, broadly available, vendor-agnostic solution for clinical validation and deployment of AI models targeting medical images at the point of care. The PACS-AI platform is supported by the Canadian Institute for Advanced Research (CIFAR) as part of its Solution Network on Integrated AI for Health Imaging.[61] Over the next 3 years, we plan to publicly release the source code of the application as an open source project.

The PACS-AI platform was designed using mostly open-source libraries and software packages. The application was designed using the Docker containerized microservices (Docker, Palo Alto, CA), and each component of the application is designed as a separate container, each of which communicates using a RESTful application programming interface (API). The API gateway is based on the Kong platform (Kong, San Francisco, CA) and directs the user's requests to appropriate microservices within the application. Authentication is handled by the Auth0 platform (Okta, San Franscisco, CA), which can connect with local institutions' user management services or can be set up independently.

Table 1. PACS-AI model publishing guidelines

1. Model and version identification: Detailed identification of the AI model, including version number.
2. Framework: Mention the deep learning framework used (eg, TensorFlow, PyTorch) and its version.
3. Model architecture: Specify the model architecture (eg, ResNet, U-Net, EfficientNet).
4. Input requirements:
   - Normalization: Detail any input normalization required, including mean and standard deviation values, as well as how these values were obtained.
   - Preprocessing steps: Outline any preprocessing steps required before inferencing (resizing, cropping, etc).
   - Size: Provide the input size expected by the model (eg, 224 × 224 pixels).
   - Frames per second: If applicable, specify how many frames per second the model can process.
   - Data format: State the required data format for input (eg, DICOM, JPEG, PNG). For videos, define the expected codec and container (eg, MP4—H.264).
5. Output specifications*:
   - Class labels: List the class labels the model can predict (eg, normal, abnormal, specific conditions).
   - Confidence scores: State if the model output includes confidence scores.
   - Segmentation maps: For segmentation models, include information about the output format for segmentation maps, such as DICE score.
6. Clinical variables*: List and metadata of additional clinical variables that the model considers.
7. Performance metrics*: Detail the model's performance metrics (eg, accuracy, AUC, precision, recall). Include subgroup performance and bias evaluation.
8. Training data set*: Describe the data set used to train the model, including size, source, demographics, and any relevant characteristics (eg, modalities, number of classes).
9. Validation and testing*:
   - Validation data set: Information on the validation set used.
   - Test results: Summarize the model's performance on a test set, including metrics and data set characteristics.
10. Regulatory approval status*: Status of regulatory approvals.
11. Evidence of clinical efficacy*: Summary of studies demonstrating clinical utility.
12. Integration and Verification Processes*: Overview of integration into health care workflows and verification processes.
13. Monitoring and performance*: Strategy for continuous monitoring of model performance after deployment.
14. Usage instructions*: Provide clear instructions on how to use the model, including example code snippets if necessary.
15. Licensing*: Information on the model's licensing and any restrictions on its use.
16. Author information*: Provide contact information or profiles of the model's authors.
17. Economic consideration*: Cost-benefit analysis or economic impact assessment.
18. References and publications*: Relevant references and publications supporting the model's development and efficacy.
19. Dependencies: List any specific library or package dependency as well as their versions required to execute the model.
20. Hardware requirements: Detail the minimum and recommended hardware requirements (eg, GPU/CPU specs, RAM).
21. Model checkpoints: Direct links to download model weights/checkpoints.
22. Sample code/notebooks: Provide sample code or Jupyter notebooks demonstrating the model in action.
23. Support and updates: Information on how users can get support and whether there will be updates to the model.

AI, artificial intelligence; AUC, area under the receiver operating characteristic curve; CPU, central processing unit; DICOM, Digital Imaging and Communications in Medicine; GPU, graphical processing unit; JPEG, Joint Photographic Experts Group; MP4, Moving Picture Experts Group 4 Part 14; PACS, Picture Archives Communication System; PNG, Portable Network Graphics; RAM, random access memory.

*These values should also be a part of the model facts label.[46]

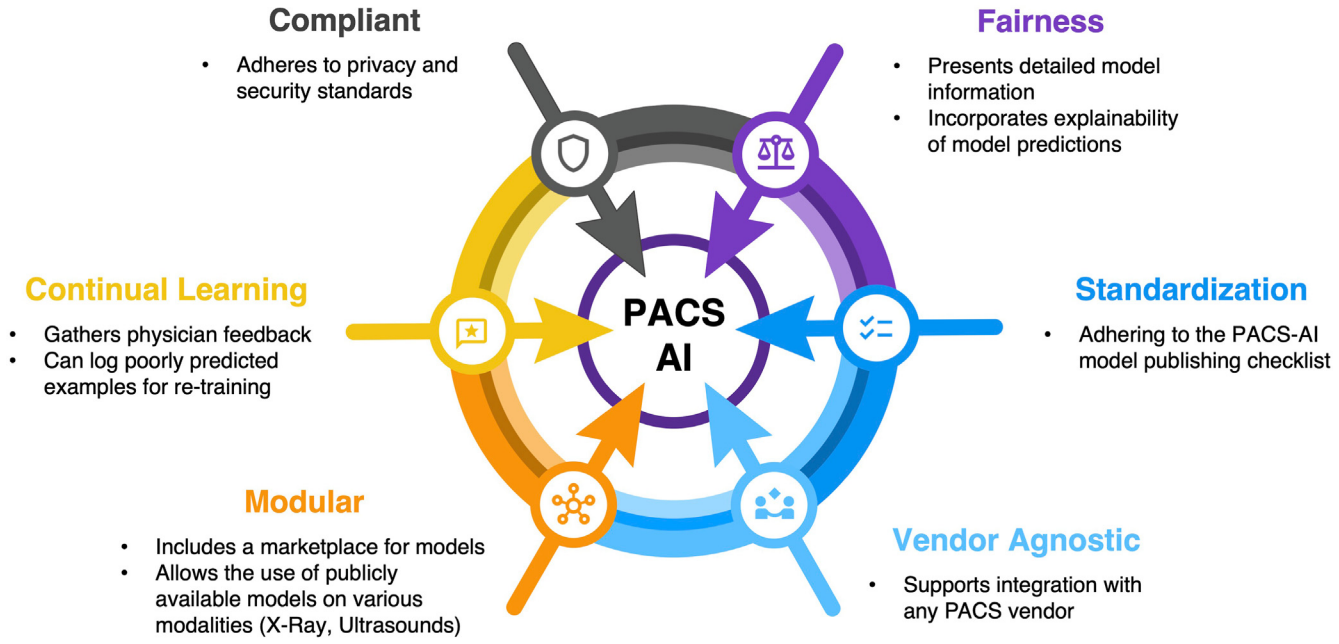

Figure 2. Features of the PACS-AI platform. AI, artificial intelligence; PACS, Picture Archives Communication System.

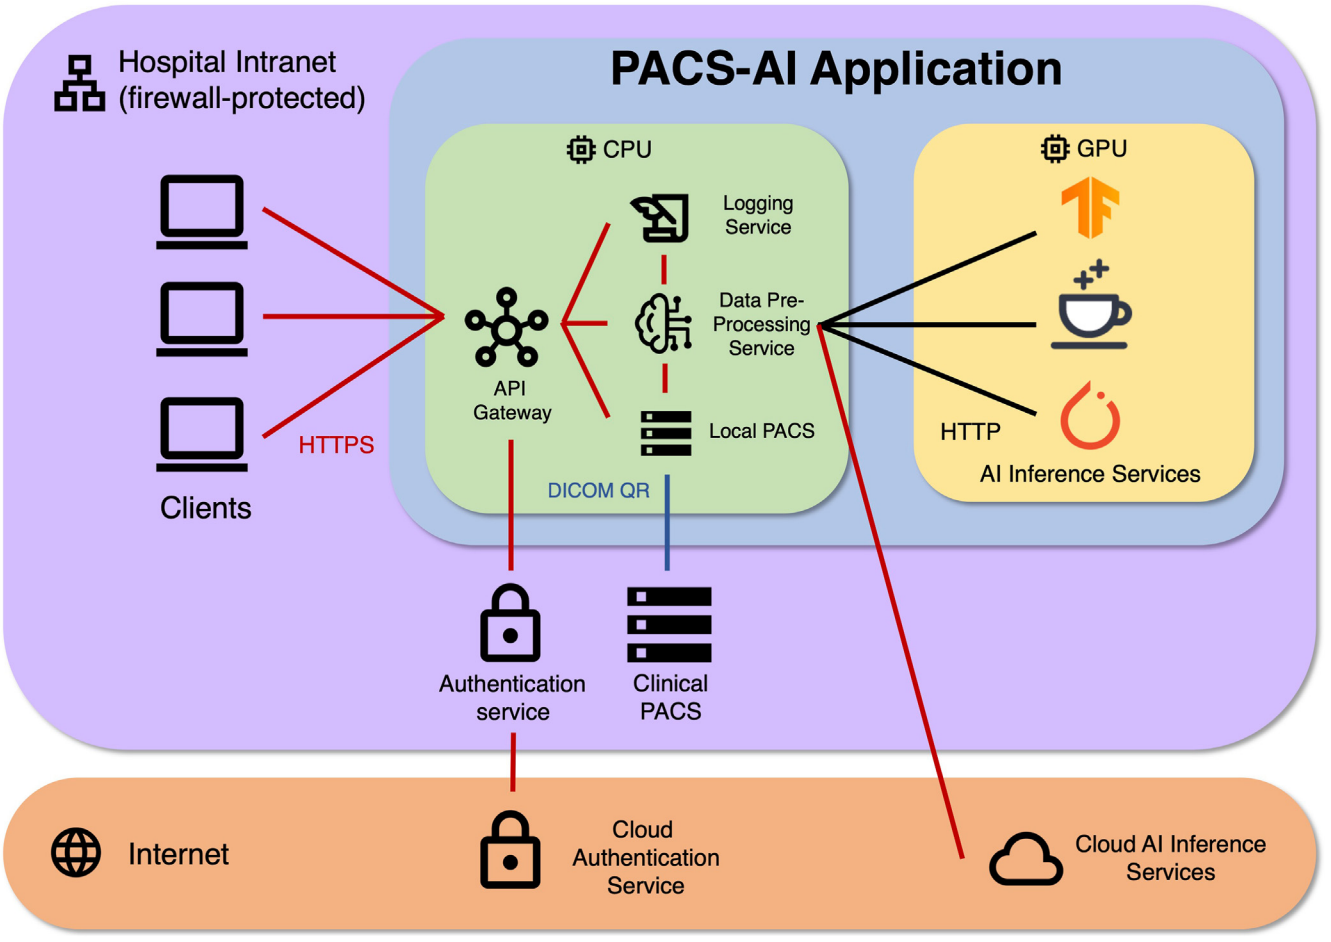

Figure 3. Architectural overview of the PACS-AI platform, demonstrating its integration with existing health care systems and its modular design for AI model deployment. AI, artificial intelligence; API, application programming interface; CPU, central processing unit; DICOM QR, Digital Imaging and Communications in Medicine query and retreive; GPU, graphical processing unit; HTTP, hypertext transfer protocol; PACS, Picture Archives Communication System.

Table 2. Strategies for overcoming challenges in AI clinical implementation via the PACS-AI platform

Challenges to AI clinical implementation	Solutions leveraged by PACS-AI platform
Variation in data format between vendors and institutions	Use standard pre-processing routines and interoperability standards (DICOM)
Cumbersome integration of multiple AI models into clinician workflow	Provide AI models within a single platform with the use of a multi-institution "marketplace"
Proprietary software platforms restrict access to data and integration of third-party AI models	Open-source software development philosophy
Cybersecurity risk and complex regulatory environment	Compliance with security standards including industry-standard auditing and encryption protocols
Lack of high-quality studies demonstrating clinical benefits of AI models	Provide an environment for the real-world multi-institution clinical validation and deployment of AI models
Application of AI models to patients under-represented in the training data set	Presentations of "model facts" label, which summarizes its capabilities, validation, and appropriate usage
Data set drift	Monitoring of model performance by integrating clinician feedback mechanism and automatically tracking model predictions over time
Lack of trust from clinicians	Display explanation or interpretation of model predictions
No standards for the publications of AI models in medicine	Proposal of the PACS-AI model publication checklist
High development cost of a bespoke software platform for each AI model	Open-source centralized hub integrating multiple models in a single platform

AI, artificial intelligence; DICOM, Digital Imaging and Communications in Medicine; PACS, Picture Archives Communication System.

Auth0 is the only closed-source library used in the project. Through Auth0, PACS-AI handles user permissions, enabling precise control over access and privileges within the application (eg, blinded users for clinical studies, admin users for audit management, and so on). Logging and auditing functionalities are managed using open-source projects Elasticsearch and Kibana (Elastic, Amsterdam, the Netherlands, and Mountain View, CA). The frontend web application was designed using the open-source Angular framework (Google, Mountain View, CA). Models trained using standard AI libraries such as PyTorch (Meta, Menlo Park, CA), and TensorFlow (Google, Mountain View, CA) are deployed using containerized inference servers to provide real-time predictions using TorchServe and TensorFlow Serving, respectively. To add new models into PACS-AI, the only prerequisite is the ability to package the desired model into a format compatible with the associated inference server. The interface with the PACS servers is handled by the Orthanc open-source DICOM server; DICOM files are cached locally for 24 hours after being used for an AI prediction and are then deleted. The platform uses the DICOM query/retrieve operations, which is a vendor-agnostic protocol for communicating with PACS servers. Pre-processing routines are written in Python using the FastAPI framework to format images after they are received from the local PACS server before passing them to the AI inference services. The PACS-AI platform currently uses a local, on-premises deployment strategy. Health care institutions, however, especially smaller primary- and secondary-care hospitals and clinics, may not possess the required GPU-powered infrastructure to take advantage of AI models using a fully on-premises deployment. We therefore plan to develop a hybrid architecture where PHI is stored on local servers but anonymized image data can be sent to cloud servers for AI model predictions. This would allow targeting a broader number of institutions while ensuring a uniform and transparent user experience.

Currently, the PACS-AI platform is used for research purposes only across several Canadian and American hospitals. Once models are validated and have regulatory clearance, PACS-AI will also become a solution for the deployment of AI models for clinical use. Health Canada and FDA clearance will be necessary for the platform to be used clinically in the US and Canada. We will also demonstrate compliance with local security and data privacy regulations.

To move beyond the 1-vendor, 1-model approach currently promoted by AI software companies, we imagine a "model marketplace." Multiple models compatible with a radiologic modality could be made available and presented as "apps," similar to those on a smartphone. A model may be available at a single institution, but also easily packaged and distributed externally. Such common framework for integrating AI models in practice and accessing them via a marketplace would enable the standardisation of models locally and across institutions.

Reporting guidelines have been proposed to guide research studies on AI algorithms, including CLAIM,[47] CONSORT-AI,[48] DECIDE-AI,[49] and SPIRIT-AI.[50] However, there are no accepted guidelines on how to publish an AI model itself. In Table 1, we propose the PACS-AI model publication checklist for information that should accompany published medical AI models. We think that such a checklist would ensure reproducibility and standardisation outside of the research or commercial development setting. The checklist was developed by the authors of this article, including professionals from the fields of cardiology, radiology, nuclear medicine, engineering, law, and bioethics. We hope to validate this initial version of the checklist during evaluation studies of the PACS-AI platform.

Repetitive, boilerplate code is needed to make a model compatible with a clinical data source, such as a system storing medical images. Tasks such as adjusting the image size or normalising pixel values may be necessary to meet the model's specifications. This task is made more complex by the range of data formats and devices used in hospitals. However, such boilerplate code can be standardised and reused, speeding up deployment of AI models. We suggest that pre-processing routines also should be published along with the model weights to facilitate integration. With PACS-AI, we aim to

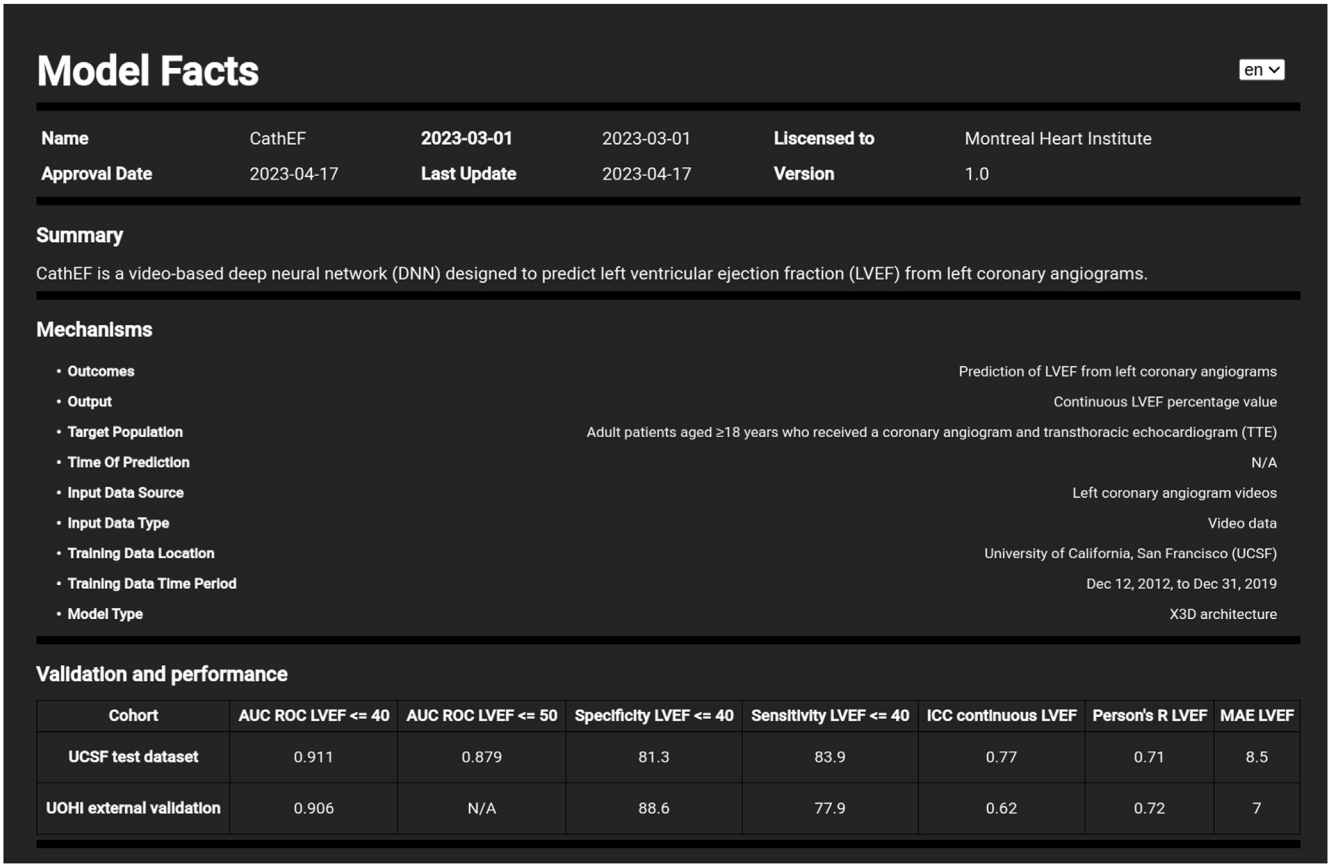

Model Facts

en

Name	CathEF	**2023-03-01**	2023-03-01	**Liscensed to**	Montreal Heart Institute
Approval Date	2023-04-17	**Last Update**	2023-04-17	**Version**	1.0

Summary

CathEF is a video-based deep neural network (DNN) designed to predict left ventricular ejection fraction (LVEF) from left coronary angiograms.

Mechanisms

- **Outcomes**: Prediction of LVEF from left coronary angiograms
- **Output**: Continuous LVEF percentage value
- **Target Population**: Adult patients aged ≥18 years who received a coronary angiogram and transthoracic echocardiogram (TTE)
- **Time Of Prediction**: N/A
- **Input Data Source**: Left coronary angiogram videos
- **Input Data Type**: Video data
- **Training Data Location**: University of California, San Francisco (UCSF)
- **Training Data Time Period**: Dec 12, 2012, to Dec 31, 2019
- **Model Type**: X3D architecture

Validation and performance

Cohort	AUC ROC LVEF <= 40	AUC ROC LVEF <= 50	Specificity LVEF <= 40	Sensitivity LVEF <= 40	ICC continuous LVEF	Person's R LVEF	MAE LVEF
UCSF test dataset	0.911	0.879	81.3	83.9	0.77	0.71	8.5
UOHI external validation	0.906	N/A	88.6	77.9	0.62	0.72	7

Figure 4. Example of a "model facts" label on the PACS-AI platform, providing essential information about an AI model's capabilities, validation, and appropriate usage. AI, artificial intelligence; PACS, Picture Archives Communication System.

handle the complex task of formatting data sources and presenting predictions to the end users. In radiologic applications, users should be able to easily search for a study, choose an AI model, apply it to the images, and view results on their screen in a timely manner. Each model implemented in the PACS-AI platform includes a model facts label which, as discussed earlier, promotes responsible use of AI (Fig. 4). It describes the model architecture, input data, output prediction, and the population of patients used in training and validation data sets. It can also be extended with confidence or uncertainty levels. The label also includes instructions for use and warnings.

The PACS-AI platform has been successfully used in a prospective validation clinical study of the LVEF Prediction During ACS Using AI Algorithm Applied on Coronary Angiogram Videos (CathEF)[62] model that recently completed enrollment at 2 separate institutions, the Montréal Heart Institute and the University of Ottawa Heart Institute (NCT05317286). The CathEF model estimates the left ventricular ejection fraction from left coronary angiograms. The integration in the PACS-AI platform enabled the real-time application of the model at the point of care, to obtain the left ventricular ejection fraction in patients undergoing urgent percutaneous coronary intervention for myocardial infarction (Fig. 5).

The PACS-AI platform and its deployment in Canadian hospitals is supported by a network of champion clinicians who understand the clinical use and limitations of specific AI models and can navigate the local administrative structure. These clinicians are crucial to the successful deployment of the platform in each health centre. With collaboration of champion clinicians and patient partners, our vision is to expand the reach of the PACS-AI platform across a multitude of health centres throughout Canada and internationally. By standardizing the approach to model deployment, releasing the code as open source and involving champion physicians in its deployment, PACS-AI addresses most challenges for implementation of AI for medical images in clinical practice (Table 2). Our goal is for PACS-AI to set a new standard in health care technology, becoming the central hub through which clinicians and scientists access and apply a spectrum of computer vision models to radiologic imagery. To enhance the platform's efficacy, we actively solicit and incorporate feedback from physicians, ensuring that the platform evolves to meet the dynamic needs of medical professionals.

We will specifically monitor the adoption rate, user satisfaction, accuracy of AI predictions presented on the platform, impact on health care delivery, and capacity of the platform to meet patient expectations regarding privacy. Specific assessment tools will be integrated to query users for feedback and present the information in a dashboard format.

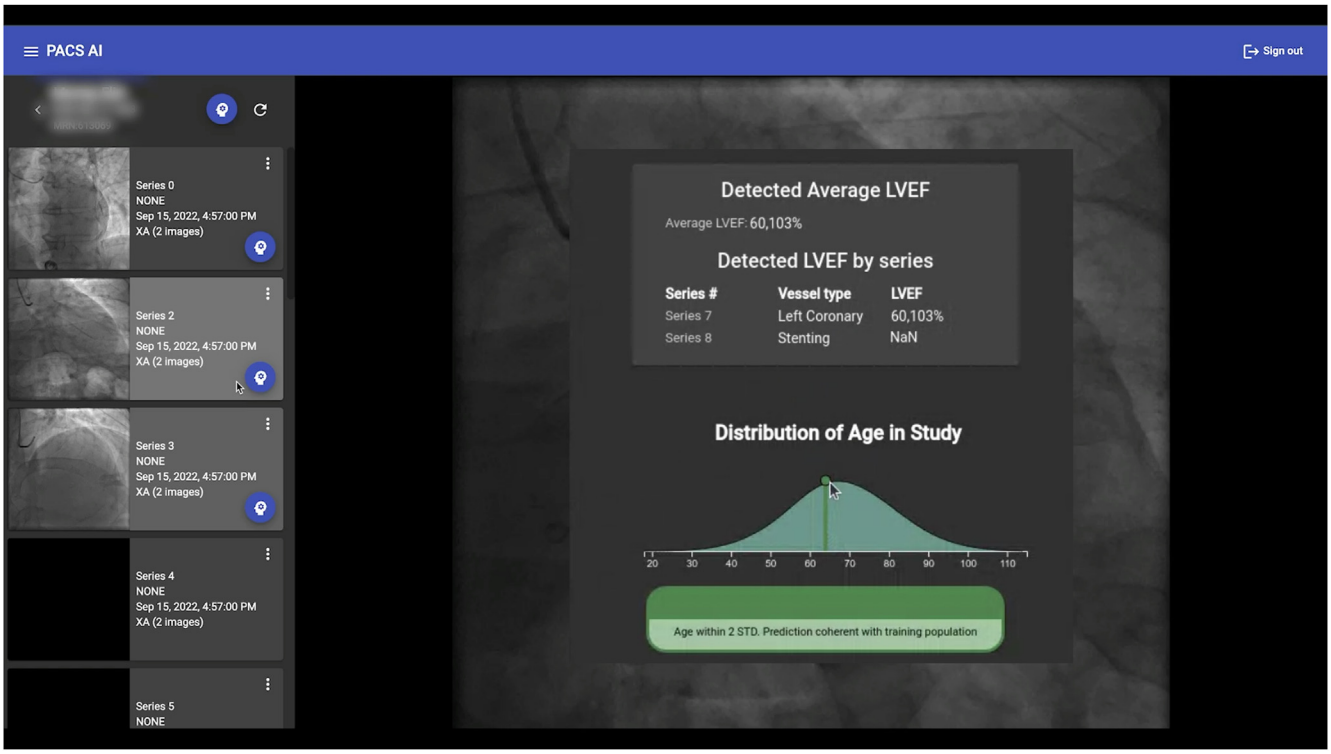

Figure 5. Example of the of the PACS-AI platform user interface in action, showcasing the real-time application of the LVEF Prediction During ACS Using AI Algorithm Applied on Coronary Angiogram Videos (CathEF) model in a clinical setting and its integration with patient care processes. The left menu allows the selection of an imaging series and displays an AI model button. Once a prediction is available, the results are displayed on screen. Information about how the target patient compares with those of the training data set is also provided. AI, artificial intelligence; PACS, Picture Archives Communication System.

Conclusion

This article highlights the need for effective and responsible deployment of AI models in cardiology and medical imaging. The challenges associated with IT infrastructure heterogeneity, closed-source software, cybersecurity risks, generalisability, bias, regulation, and cost have been critically examined. The proposed PACS-AI platform offers a promising path forward, addressing some of the major issues identified in this article. It integrates AI into existing clinical workflows, maintaining a focus on data security and regulatory compliance. With these features, it is our hope that the PACS-AI platform will become a major enabling technology for the responsible deployment of AI models in Canada and elsewhere. Future developments and wider adoption of such technology could significantly enhance diagnostic accuracy and efficiency in health care.

Ethics Statement

The authors confirm that ethical review is not applicable to this article as it is a review.

Patient Consent

The authors confirm that patient consent is not applicable to this article. This is a review article that did not enroll patients.

Funding Sources

This review and the development of the PACS-AI platform was funded by the Canadian Institute for Advanced Research Solution Network on Integrated AI for Health Imaging. Dr Tang is a Fonds de Recherche du Québec—Santé (FRQS) Senior Research Scholar (no. 298509). The funders played no role in study design, data collection, analysis and interpretation of data, or the writing of this manuscript. Dr Gichoya is a 2022 Robert Wood Johnson Foundation Harold Amos Medical Faculty Development Program and declares support from Radiological Society of North America Health Disparities grant (no. EIHD2204), Lacuna Fund (no. 67), Gordon and Betty Moore Foundation, and National Institutes of Health (National Institute of Biomedical Imaging and Bioengineering) Microbiology and Infectious Diseases Research Committee grant under contracts 75N92020C00008 and 75N92020C00021. Dr Hussin is an FRQS Junior 2 Research Fellow. Dr Avram is supported by the FRQS (grant no. 312758), the Montréal Heart Institute Research Centre, the Montréal Heart Institute Foundation, the Des Groseillers—Bérard Interventional Cardiology Research Chair, Natural Sciences and Engineering Research Council of Canada, the Institute for Data Valorisation, and the Fonds de Recherche du Québec—Nature et Technologies.

Disclosures

Dr Avram is a co-inventor in patent pending 63/208,406 (Method and System for Automated Analysis of Coronary Angiograms). The other authors have no conflicts of interest to disclose.

Editorial Disclaimer

Given his role as Associate Editor, Robert Avram had no involvement in the peer review of this article and has no access to information regarding its peer review.

Declaration of Generative AI and AI-Assisted Technologies in the Writing Process

During the preparation of this work, the authors used ChatGPT to improve the readability and language of the manuscript. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

References

1. Lauzier PT, Avram R, Dey D, et al. The evolving role of artificial intelligence in cardiac image analysis. Can J Cardiol 2022;38:214-24.
2. Langlais ÉL, Thériault-Lauzier P, Marquis-Gravel G, et al. Novel artificial intelligence applications in cardiology: current landscape, limitations, and the road to real-world applications. J Cardiovasc Trans Res 2023;16:513-25.
3. Ciarrapico AM, Ugenti R, Di Minco L, et al. Diagnostic imaging and spending review: extreme problems call for extreme measures. Radiol Med 2017;122:288-93.
4. National Library of Medicine. Available at: https://pubmed.ncbi.nlm.nih.gov/?term=artificial+intelligence. Accessed January 4, 2024.
5. US Food and Drug Administration: Artificial intelligence and machine learning (AI/ML)—enabled medical devices. Available at: https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-aiml-enabled-medical-devices. Accessed January 4, 2024.
6. Goto S, Mahara K, Beussink-Nelson L, et al. Artificial intelligence—enabled fully automated detection of cardiac amyloidosis using electrocardiograms and echocardiograms. Nat Commun 2021;12:2726.
7. Escobar GJ, Liu VX, Schuler A, et al. Automated identification of adults at risk for in-hospital clinical deterioration. N Engl J Med 2020;383:1951-60.
8. Emanuel EJ, Wachter RM. Artificial intelligence in health care: will the value match the hype? JAMA 2019;321:2281-2.
9. Khunte M, Chae A, Wang R, et al. Trends in clinical validation and usage of US Food and Drug Administration-cleared artificial intelligence algorithms for medical imaging. Clin Radiol 2023;78:123-9.
10. Ouyang D, He B, Ghorbani A, et al. Video-based AI for beat-to-beat assessment of cardiac function. Nature 2020;580(7802):252-6.
11. Lehne M, Sass J, Essenwanger A, Schepers J, Thun S. Why digital medicine depends on interoperability. NPJ Digit Med 2019;2:79.
12. Yu AC, Mohajer B, Eng J. External validation of deep learning algorithms for radiologic diagnosis: a systematic review. Radiol Artif Intell 2022;4:e210064.
13. Atasoy H, Greenwood BN, McCullough JS. The digitization of patient care: a review of the effects of electronic health records on health care quality and utilization. Annu Rev Public Health 2019;40:487-500.
14. Bronsoler A, Doyle J, van Reenen J. The impact of health information and communication technology on clinical quality, productivity, and workers. Annu Rev Econ 2022;14:23-46.
15. Neprash HT, McGlave CC, Cross DA, et al. Trends in ransomware attacks on US hospitals, clinics, and other health care delivery organizations 2016-2021. JAMA Health Forum 2022;3:e224873.
16. US Food and Drug Administration: Cybersecurity in medical devices: quality system considerations and content of premarket submissions. September 27, 2023. Available at: https://www.fda.gov/media/119933/download. Accessed April 1, 2024.
17. Kelly CJ, Karthikesalingam A, Suleyman M, Corrado G, King D. Key challenges for delivering clinical impact with artificial intelligence. BMC Med 2019;17:195.
18. Yusuf M, Atal I, Li J, et al. Reporting quality of studies using machine learning models for medical diagnosis: a systematic review. BMJ Open 2020;10:e034568.
19. Nolin Lapalme A, Corbin D, Tastet O, Avram R, Hussin JG, et al. Advancing fairness in cardiac care: strategies for mitigating bias in artificial intelligence models within cardiology. Can J Cardiol 2024;40:1906-20.
20. European Union Agency for Fundamental Rights: Bias in algorithms: artificial intelligence and discrimination. Available at: https://data.europa.eu/doi/10.2811/25847. Accessed February 7, 2024.
21. Finlayson SG, Subbaswamy A, Singh K, et al. The clinician and dataset shift in artificial intelligence. N Engl J Med 2021;385:283-6.
22. US Food and Drug Administration: Proposed regulatory framework for modifications to artificial intelligence/machine learning (AI/ML)—based software as a medical device (SaMD). Available at: https://www.fda.gov/files/medical%20devices/published/US-FDA-Artificial-Intelligence-and-Machine-Learning-Discussion-Paper.pdf. Accessed April 1, 2024.
23. Drukker K, Chen W, Gichoya J, et al. Toward fairness in artificial intelligence for medical image analysis: identification and mitigation of potential biases in the roadmap from data collection to model deployment. J Med Imag 2023;10:061104.
24. Medical Imaging and Data Resource Center: Medical imaging AI/machine learning bias awareness tool: an interactive decision tree. Available at: https://www.midrc.org/bias-awareness-tool-1. Accessed January 1, 2024.
25. Brady AP, Allen B, Chong J, et al. Developing, purchasing, implementing and monitoring AI tools in radiology: practical considerations. a multi-society statement from the ACR, CAR, ESR, RANZCR and RSNA. Radiol Artif Intell 2024;6:e230513.
26. Rajpurkar P, Lungren MP. The current and future state of ai interpretation of medical images. N Engl J Med 2023;388:1981-90.
27. Saporta A, Gui X, Agrawal A, et al. Benchmarking saliency methods for chest X-ray interpretation. Nat Mach Intell 2022;4:867-78.
28. Gaube S, Suresh H, Raue M, et al. Nontask expert physicians benefit from correct explainable AI advice when reviewing X-rays. Sci Rep 2023;13:1383.
29. World Health Organisation: Regulatory considerations on artificial intelligence for health. Available at: https://iris.who.int/bitstream/handle/10665/373421/9789240078871-eng.pdf. Accessed April 1, 2024.

30. Saltelli A, Dankel DJ, di Fiore M, Holland N, Pigeon M. Science, the endless frontier of regulatory capture. Futures 2022;135:102860.

31. Gallifant J, Nakayama LF, Gichoya JW, Pierce R, Celi LA. Equity should be fundamental to the emergence of innovation. PLoS Digit Health 2023;2:e0000224.

32. Voets MM, Veltman J, Slump CH, Siesling S, Koffijberg H. Systematic review of health economic evaluations focused on artificial intelligence in healthcare: the tortoise and the cheetah. Value Health 2022;25:340-9.

33. Morgan TA, Helibrun ME, Kahn CE. Reporting initiative of the Radiological Society of North America: progress and new directions. Radiology 2014;273:642-5.

34. Berkowitz SJ, Wei JL, Halabi S. Migrating to the modern PACS: challenges and opportunities. RadioGraphics 2018;38:1761-72.

35. Mollura DJ, Culp MP, Pollack E, et al. Artificial intelligence in low- and middle-income countries: innovating global health radiology. Radiology 2020;297:513-20.

36. Alshuqayran N, Ali N, Evans R. A systematic mapping study in microservice architecture. 2016 IEEE 9th International Conference on Service-Oriented Computing and Applications. Macau, China 2016:44-51. https://doi.org/10.1109/soca.2016.15.

37. Anderson C. Docker [software engineering]. IEEE Softw 2015;32. 102-c3.

38. Abadi M, Agarwal A, Barham P, et al. TensorFlow, large-scale machine learning on heterogeneous systems. https://doi.org/10.5281/zenodo.4724125.

39. Jia Y, Shelhamer E, Donahue J, et al. Caffe: convolutional architecture for fast feature embedding. June 20, 2014. https://doi.org/10.48550/arxiv.1408.5093. Accessed January 15, 2024.

40. Paszke A, Gross S, Massa F, et al. PyTorch: an imperative style, high-performance deep learning library. In: Wallach H, Larochelle H, Beygelzimer A, et al., eds. Advances in Neural Information Processing Systems 32. Curran Associates, 2019:8024-35. Available at: http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf.

41. Chen T, Li M, Li Y, et al. MXNet: a flexible and efficient machine learning library for heterogeneous distributed systems. December 3, 2015. https://doi.org/10.48550/arxiv.1512.01274. Accessed January 15, 2024.

42. MONAI Consortium. MONAI: medical open network for AI. https://doi.org/10.5281/zenodo.4323058.

43. Navale V, Bourne PE. Cloud computing applications for biomedical science: a perspective. PLoS Comput Biol 2018;14:e1006144.

44. Mehrtak M, SeyedAlinaghi S, MohsseniPour M, et al. Security challenges and solutions using healthcare cloud computing. J Med Life 2021;14:448-61.

45. Abba Ari AA, Ngangmo OK, Titouna C, et al. Enabling privacy and security in cloud of things: architecture, applications, security & privacy challenges. Appl Comput Inform 2024;20:119-41.

46. Sendak MP, Gao M, Brajer N, Balu S. Presenting machine learning model information to clinical end users with model facts labels. NPJ Digit Med 2020;3:41.

47. Mongan J, Moy L, Kahn CE. Checklist for Artificial Intelligence in Medical Imaging (CLAIM): a guide for authors and reviewers. Radiol Artif Intell 2020;2:e200029.

48. Liu X, Cruz Rivera S, Moher D, Calvert MJ, Denniston AK; SPIRIT-AI and CONSORT-AI Working Group. Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: the CONSORT-AI extension. Nat Med 2020;26:1364-74.

49. Vasey B, Nagendran M, Campbell B, et al. Reporting guideline for the early-stage clinical evaluation of decision support systems driven by artificial intelligence: DECIDE-AI. Nat Med 2022;28:924-33.

50. Cruz Rivera S, Liu X, Chan AW, et al. Guidelines for clinical trial protocols for interventions involving artificial intelligence: the SPIRIT-AI extension. Lancet Digit Health 2020;2:e549-60.

51. Felix EA, Lee SP. Systematic literature review of preprocessing techniques for imbalanced data. IET Softw 2019;13:479-96.

52. Kumari P, Chauhan J, Bozorgpour A, et al. Continual learning in medical image analysis: a comprehensive review of recent advancements and future prospects. https://doi.org/10.48550/arXiv.2312.17004. Accessed February 3, 2024.

53. Association of International Certified Public Accountants: 2018 SOC 2 description criteria (with revised implementation guidance—2022). Available at: https://www.aicpa-cima.com/resources/download/get-description-criteria-for-your-organizations-soc-2-r-report. Accessed January 1, 2024.

54. Santé et Services Sociaux—Québec: Technologies de l'information—internet. Available at: http://www.ti.msss.gouv.qc.ca/. Accessed January 1, 2024.

55. Government of Canada: Guide on the use of generative AI. Available at: https://www.canada.ca/en/government/system/digital-government/digital-government-innovations/responsible-use-ai/guide-use-generative-ai.html. Accessed April 1, 2024.

56. Jin J. FDA authorization of medical devices. JAMA 2014;311:435.

57. Clinical decision support software: guidance for industry and Food and Drug Administration staff. Available at: https://www.fda.gov/media/109618/download.

58. Health Canada: Predetermined change control plans for machine learning-enabled medical devices: guiding principles. Available at: https://www.canada.ca/en/health-canada/services/drugs-health-products/medical-devices/good-machine-learning-practice-medical-device-development/predetermined-change-control-plans-machine-learning-enabled-medical-devices.html. Accessed January 10, 2024.

59. Miller FA, Lehoux P, Peacock S, et al. How procurement judges the value of medical technologies: a review of healthcare tenders. Int J Technol Assess Health Care 2019;35:50-5.

60. Public Services and Procurement Canada: The procurement process. Available at: https://buyandsell.gc.ca/for-businesses/selling-to-the-government-of-canada/the-procurement-process. Accessed January 1, 2024.

61. Sandusky K. CIFAR announces launch of two AI for Health Solution Networks. Canadian Institute for Advanced Research, July 13, 2023. Available at: https://cifar.ca/cifarnews/2023/07/13/cifar-announces-launch-of-two-ai-for-health-solution-networks/. Accessed January 1, 2024.

62. Avram R, Barrios JP, Abreau S, et al. Automated assessment of cardiac systolic function from coronary angiograms with video-based artificial intelligence algorithms. JAMA Cardiol 2023;8:586-94.