Wonca
World family doctors. Caring for people.

GUIDANCE

on the ethical use of artificial intelligence in primary care

A document developed by the World Organization of Family Doctors



30 questions to help family doctors review an AI tool

Guidance document on the ethical use of artificial intelligence in primary care

Authors: Tania Moerenhout[1], Avni Patel[2], Andrew Papanikitas[3], Issam Shaarani[4], Keith Thompson[5], and the WONCA Working Party on Ethics and Professionalism.

[1] Department of Bioethics, University of Otago, Dunedin, Aotearoa New Zealand

[2] Tanglin Health Clinic, Ministry of Health Malaysia, Kuala Lumpur, Malaysia

[3] Nuffield Department of Primary Care Health Sciences, University of Oxford, United Kingdom

[4] Faculty of Medicine, Beirut Arab University, Beirut, Lebanon

[5] Department of Family Medicine, Schulich School of Medicine & Dentistry, Western University, London, Ontario, Canada

This guidance has been created by the WONCA (World Organization of Family Doctors) Working Party on Ethics and Professionalism, with the support of the WONCA Working Party on Digital Health and was endorsed by the WONCA Executive Committee on 29 January 2026.

SUGGESTED CITATION

Moerenhout T, Patel A, Papanikitas A, Shaarani I, Thompson K, WONCA Working Party on Ethics and Professionalism. *Guidance document on the ethical use of artificial intelligence in primary care*. World Organization of Family Doctors (WONCA); 2026. Available from: https://www.globalfamilydoctor.com/Resources/ethical-ai-primary-care.aspx

Version 1.0 · October 2026

© World Organization of Family Doctors (WONCA), 2026

CONTENTS

At a glance 4

Executive summary 5

Introduction 7

Definitions: artificial intelligence, machine learning, and large language models 8

Opportunities in primary care 10

Ethical challenges of AI use in primary care 12

Ethical values and principles that (should) guide AI use in primary care 17

Practical guidance for the ethical use of AI tools in daily practice 21

Table 1: Questions for the review process of AI tools in primary care 25

Next steps: roundtable dialogues for ethical AI in primary care 27

Reference list 28

About WONCA 31

AT A GLANCE

The guidance makes **11 recommendations**: 2 fundamental requirements and 9 practical considerations, supported by **30 questions** for reviewing an AI tool.

2 FUNDAMENTAL REQUIREMENTS

1 · VALIDATION AND REGULATION

Family doctors should only use AI tools that have been validated and that meet regulatory standards specific to clinical use in primary care.

2 · COLLABORATION

The implementation of AI in primary care is too complex to be delegated to individual practitioners.

⚠ The default position should be not to use generative AI applications in primary care unless they have undergone sufficient scrutiny and meet ethical and clinical requirements (precautionary principle).

↓ Most AI tools have not been thoroughly validated, and robust AI regulatory processes are lacking in most countries. For promising tools that have some evidence but have not been fully validated or are unregulated, family doctors can use the 9 considerations below.

9 PRACTICAL CONSIDERATIONS

1 Evidence	**2** Risk of harm	**3** Data review
4 Co-design	**5** Patient information and consent	**6** Inclusion and equity
7 Sustainability and planetary health	**8** Integration and impact on workflow	**9** Ongoing monitoring

A printable checklist of the 30 questions is available at www.globalfamilydoctor.com/Resources/ethical-ai-primary-care.aspx.

EXECUTIVE SUMMARY

This guidance has been developed to support family doctors who are considering the use of artificial intelligence (AI) in primary care. Its aim is threefold. First, it provides an overview of the specific opportunities and ethical challenges presented by AI in primary care. Second, it articulates which values and principles should guide the use of AI in primary care. This includes two strands: ethical values related to AI, and primary care principles that AI tools should support. These values and principles are not only relevant to clinicians but also to designers, developers, and health managers. Third, this guidance document offers more practical ethical guidance for family doctors who wish to incorporate AI tools into their practice. Here, 11 recommendations are provided, aiming to be both aspirational and practical. The first 2 recommendations are *a priori*, focusing on validation and collaboration. They specify that, under ideal circumstances, family doctors should only use AI tools that have been validated and meet regulatory standards for clinical use in primary care, and that the implementation of AI in primary care is too complex to be delegated to individual practitioners. However, we also acknowledge that these recommendations may be highly restrictive, as, at the time of writing, such regulatory standards and clear decision-making processes have not yet been established in most countries. Following this, we formulate 9 recommendations for family doctors to use if they have access to promising AI tools with some evidence supporting them, but that have not been fully validated or are unregulated. Ideally, such a review should not fall on the shoulders of the individual family doctor, as specified in the recommendation on collaboration above. However, family doctors need a process in place should they wish to start using AI tools that have not undergone rigorous validation or regulation.



Tania Moerenhout presents the WONCA Working Party on Ethics and Professionalism's draft guidance on the ethical use of AI in primary care at the 25th WONCA World Conference, Lisbon, September 2025.

Family doctors should consider:

1 Evidence: Tools should be supported by independent evidence and tested in real-world settings.

2 Risk of harm: Potential harms should be assessed across multiple domains, including patient, clinician and system-level impacts.

3 Data review: The training data and ongoing data collection should be reviewed to ensure the AI tool is relevant to the local context and population and meets privacy and confidentiality requirements.

4 Co-design: Tools should be developed with meaningful input from clinicians, patients, and communities.

5 Patient information and consent: Use of AI tools requires appropriate patient disclosure, transparent information, and, where necessary, explicit consent.

6 Inclusion and equity: Implementation should avoid reinforcing disparities or excluding vulnerable groups.

7 Sustainability and planetary health: AI tools may have a substantial impact on carbon emissions, water use, and overall planetary health, which should be weighed against potential benefits.

8 Integration and impact on workflow: AI tools should align with clinical workflows and support rather than disrupt team-based care.

9 Ongoing monitoring: Use of AI tools should be continually monitored for performance, safety, and unintended consequences.

Alongside these recommendations, the guidance also includes a list of 30 questions to help family doctors conduct this review. We are acutely aware that any guidance on AI risks being outdated as soon as it is published. Nevertheless, we believe that the underlying values and principles remain stable over time, and that incorporating them into technology design and implementation is vital if we want AI tools to succeed and support family doctors in their work, rather than undermine them.

INTRODUCTION

The past few years have seen an unprecedented boom in the development and deployment of artificial intelligence (AI) applications, including in healthcare. Much of this momentum is driven by advances in large language models (LLMs) and so-called 'generative' AI. This incoming wave of AI technology brings tremendous opportunities for primary care, but there are also potential harms which generate ethical questions and challenges.

We consider that primary healthcare has many features that make adoption of AI applications likely or desirable. Practitioners often work independently yet require instant access to clinical knowledge and decision-support tools that could be provided by AI technology. However, AI development requires that large volumes of diverse, high-quality healthcare data be collected, stored and managed. This data collection needs to occur at both the individual and population level. Patient empowerment, self-care, and administrative efficiency may benefit greatly from AI support.

While the potential benefits are substantial, AI tools can cause harm if introduced without appropriate evidence, safeguards, or ethical oversight. In response to this, the Working Party on Ethics and Professionalism of the World Organization of Family Doctors (WONCA) developed this guidance document to guide the ethical use of AI in primary care globally.

This guidance document:

1. Defines AI, machine learning, and LLMs, in the healthcare context;
2. Outlines potential opportunities for primary care;
3. Explores ethical challenges that may arise with AI use in primary care;
4. Articulates guiding ethical principles and primary care values for ethical AI use;
5. Provides practical guidance for clinicians thinking of using AI in their daily practice;
6. Considers the next steps, starting with roundtable dialogues.

The document is built on existing literature and guidelines, but is not exhaustive. Rather, it has a specific focus on ethical considerations of AI use in primary care. The first part is situated at the systemic level, identifying opportunities and challenges, values and principles, while the last part is aimed at the individual practitioner using AI tools. In this sense, this guidance seeks to include both conceptual frameworks and actionable advice.

DEFINITIONS

artificial intelligence, machine learning, and large language models

Although a commonly agreed-upon definition is lacking, AI usually refers to intelligence displayed or simulated by code (i.e., algorithms) or machines, formulated as: "*Technologies with the ability to perform tasks that would otherwise require human intelligence*".[1,2] AI is not new: the term was coined at a Dartmouth College workshop in the summer of 1956 led by John McCarthy and centered around the idea of simulating human intelligence.[2] What intelligence means can vary: it may include all kinds of cognitive functions and tasks humans usually perform, such as learning, perception, planning, decision making, and problem solving.[2] This is why two types of AI can be discerned. N*arrow* (or weak) AI can only perform in a specific domain or fulfill a specific task, while *general* (or strong) AI is capable of carrying out any cognitive tasks that humans can do.[2,3] Although many of the Dartmouth workshop participants expected to see the creation of a machine as intelligent as a human being within a generation, general AI is not on the horizon yet.[2]

Another relevant difference is between symbolic or rule-centric AI, where the AI is developed along a predetermined decision tree or flowchart (meaning its reasoning is transparent), and the *machine learning or data-centric AI* approach, where the technology learns from data and examples without explicit programming. [2–4] *Machine learning* can be *supervised*, where the algorithm's learning is based on predetermined variables that are provided to the model, or *unsupervised*, where the AI finds patterns in datasets and creates categories on its own.[2] In health care, *reinforcement learning* is often used, a type of machine learning in which the algorithm receives feedback on its output, much like a trial-and-error or reward-and-punishment system.[2,5] For example, the algorithm may receive feedback on its assessment of images of skin lesions as cancerous or non-cancerous, or it may learn from a surgeon's physical motions to improve robotic-assisted surgery.[4]

Most of the current progress in AI has been made in a subfield of machine learning called *deep learning*. Deep learning is a multi-layered machine learning system that can accept multiple data types as input and progressively extracts features from data as it flows through the layers of the system.[4,5] The underlying structure is referred to as a *neural network*, consisting of several layers of 'neurons' or simple processing units.[2] *Large Language Models* (LLMs) are neural networks with a primary focus on language; they are trained on extensive text datasets and can comprehend and generate text across a wide range of topics.[6] This is why they are also called *generative AI*. One of the best-known examples is OpenAI's ChatGPT, a large language model designed as a generative AI chatbot.[7] Following its launch in November 2022, the platform reached the milestone of 100 million users in just two months, generating excitement over its potential in various fields, including healthcare.[7,8] As shown by a recent review, LLMs show promise in diagnostic support, patient notes and discharge letter writing, and in empowering patients to navigate the healthcare system, but raise concerns of risk of harm due to bias and 'hallucinations' (the generation of inaccurate but convincing information, sometimes also referred to as confabulation), privacy and security limitations, lack of validation and evaluation, and other ethical and legal issues.[9,10] The advent of multimodal LLMs (sometimes also called generative large language multimodal models,

GLLMMs, pronounced ‘Gollums’) allows the inclusion of not only textual content, but also medical imaging, electrophysiological data, sensory data, videos, audio, and omics data, resulting in multimodal learning and output.[11]

To summarise, AI consists of algorithms that analyze large datasets quickly.[1] As a result, the quality of AI tools will, to a large extent, be determined by the quality of the data, and data ethics play an important role in AI ethics. We will first consider the opportunities of AI in primary care, before turning to the challenges.

OPPORTUNITIES IN PRIMARY CARE

Globally, primary care is facing multiple challenges, and the hope is that AI may mitigate at least some of them. We identify 4 core challenges based on the literature and recent work done within WONCA[12]:

1. **Resource constraints** mean current needs cannot be met, and underfunding of the primary care sector risks compromising quality of care and accessibility
2. **Staff shortages** increase burnout levels in family doctors and lower retention rates
3. Multimorbidity and chronic diseases lead to increasingly **complex care** and patient management
4. **Global threats** such as climate change, war, and pandemics will put more strain on primary care and require adaptations and flexibility[13]

With these challenges in mind, AI offers opportunities in primary care in 6 key domains:

1. **Clinical documentation and administrative support:** LLMs can generate patient summaries based on the information available in the electronic health record and can produce patient or discharge notes based on written information or on the conversation with the patient.[9,14] For example, ambient AI scribes use a microphone to transcribe clinical encounters and organise the transcripts into clinical note templates that can be integrated into the electronic record.[15] Generative AI can also be used to automate routine tasks in healthcare, such as through intelligent scheduling systems where chatbots or voice assistants allow patients to book or cancel appointments.[14]
2. **Precision medicine and therapeutic support**: Generative AI may predict which treatments will work best for an individual patient based on their genetic information, pathology reports, imaging, lifestyle parameters and other data sources.[14] This may be used to enhance cancer treatment or to optimize drug dosing for epilepsy.[16]
3. **Risk prediction**: By processing large amounts of data from electronic health records, AI tools can predict who is at risk of certain illnesses, post-operative complications, and even death. For example, AI can help predict the risk of chronic kidney disease in patients and forecast kidney function decline.[16] Natural language processing may assist in identifying people with suicidal ideation, thereby contributing to suicide prevention.[17] ChatGPT has shown a robust performance in cardiovascular disease risk prediction, comparable to traditional methods.[18]
4. **Diagnostic support**: AI tools may improve diagnostic accuracy and speed, particularly in imaging, leading to applications in radiology, pathology, ophthalmology, dermatology, cardiology, and gastroenterology.[19] For example, deep learning algorithms for the diagnosis of diabetic retinopathy based on images have proven their validity in real-world settings, potentially leading to improved screening.[19,20] Also, AI-human partnerships have the potential to improve diagnostic accuracy in skin cancer diagnosis, showing greater benefit among non-dermatologists.[21]
5. **Self-management and patient empowerment:** AI could assist in self-care through chatbots, self-monitoring and risk prediction tools, and through technologies to support people living with disabilities.[5] Integrating AI-augmented monitoring technologies and other applications into nursing care may streamline processes and support nurses to provide safer and more effective healthcare

services.[22] Chronic conditions could be better managed with virtual coaching, giving individualized advice based on medical literature and patient data.[19] Generative AI can provide patient education content at different reading levels and in multiple languages, helping to improve health literacy.[14] This may prove specifically invaluable in settings where mass migration means that there are inadequate human healthcare services or migrant populations do not speak the same language as healthcare providers. LLMs can empower patients to navigate the healthcare system.[9]

6. **Education for providers**: AI offers opportunities to enhance education and training for primary care providers. Generative AI tools show promise in delivering scalable, simulation-based education with personalised feedback, potentially improving access to quality training in diverse settings.[23] For example, AI-generated virtual patients can be used in simulations for both medical students and more senior clinicians, providing immediate feedback and even integrating realistic high-fidelity videos for training.[24]

These opportunities are closely aligned with potential benefits that have previously been identified, such as by the Royal Australian College of General Practitioners (RACGP) in their position statement on AI in primary care.[3] They also focus on reducing the administrative burden, filling service gaps in remote locations, reducing risks to patient safety, personalising treatment, increasing patients' access to care, and improving diagnostic accuracy and efficiency. Similarly, chart review and documentation, practice management, risk prediction and intervention, diagnostics, remote patient monitoring, and digital health coaching have been identified as application areas for AI in primary care in the United States.[25] Another area of use which we have not included in the above overview is AI in research and innovation.

Some general practitioners (GPs) are already using AI-based transcription software to support note writing during consultations. Several electronic health record systems now integrate AI into alerts and pop-ups, presenting family doctors* with information on medication interactions or suggesting treatment options in real-time. And many patients have started using LLMs to look up medical information and as a tool for self-management. With LLMs being freely and easily accessible, this may lower the threshold for patient engagement and education and improve equity.

* Globally, different terms are used to indicate doctors working in primary care: general practitioners, family doctors, and primary care clinicians or primary care physicians are the most common. Throughout this guidance document, we will adhere to the term 'family doctors'.

ETHICAL CHALLENGES

of AI use in primary care

Alongside these potential benefits, we need to acknowledge that the use of AI also carries risks and challenges for individual healthcare professionals, patients, the profession, society, and our planet.

A TECHNOLOGY-RELATED CHALLENGES: design, development, and regulation

- **Overhyping:** many AI tools are overpromising and underdelivering. Overhyping becomes a problem if unrealistic expectations are created and problems not addressed. This could be a reflection of the disconnect between the values and priorities of the technology industry and of healthcare:
    i. Medicine operates on the *'Do no harm' principle*, while risk-taking is very common and often even required in the technology industry.
    ii. Medicine adheres to *privacy and confidentiality*, while data access and sharing are fundamental to AI design and development.
    iii. General *bioethics principles such as beneficence, autonomy, justice and the avoidance of harm* are prioritized in medicine, whereas industry is predominantly driven by profit-making, sales figures being taken as a surrogate marker of public good.
    iv. New healthcare innovations tend to be *time-consuming due to regulatory and other requirements* (for example, given the time it takes to bring new pharmaceuticals to the market). For the technology industry, being the first on the market is often of crucial importance. The need to comply with regulation and demonstrate benefit may not align well with fast-paced AI development and the economic goal of profitability.
    v. In healthcare, having a clear understanding of the potential challenges of a tool is crucial, meaning *realistic expectations* are required. For technological innovation, selling the products is prioritized, which may result in high expectations and unrealistic goals, with very little attention going to the management of potential weaknesses and downsides.
    vi. Healthcare in general and primary care specifically are concerned with providing care to all patients, including *vulnerable patient populations and minorities*. The technology industry does not always include vulnerable patients in its target audience, resulting in products that sometimes only reach the so-called "worried well" or early adopters, and larger populations at a later stage. However, they may never reach those people who could potentially benefit from the applications the most.

All of the above results in AI applications being brought onto the market at an early stage in their development, when robust evidence is still lacking.[19,26] This phenomenon is also described as "technical debt".[27] Dealing with so-called 'minimally viable products' contributes to the ethical challenges Family doctors encounter in the implementation of AI tools.

HEALTHCARE AND THE TECHNOLOGY INDUSTRY

	MEDICINE AND HEALTHCARE	TECHNOLOGY INDUSTRY
Harm and risk	Operates on the 'do no harm' principle	Risk-taking is very common and often even required
Data	Adheres to privacy and confidentiality	Data access and sharing are fundamental to AI design and development
What drives decisions	Bioethics principles: beneficence, autonomy, justice and the avoidance of harm	Predominantly profit-making, with sales figures taken as a surrogate marker of public good
Pace	Innovation is time-consuming because of regulatory and other requirements	Being first on the market is often of crucial importance
Expectations	A clear understanding of a tool's potential challenges; realistic expectations	Selling is prioritised; high expectations, with little attention to weaknesses and downsides
Who is reached	Care for all patients, including vulnerable populations and minorities	Products may reach the "worried well" and early adopters first, and may never reach those who could benefit most

- **Lack of regulation**: due to the high pace of developments in AI health technology, regulatory bodies are struggling to keep pace with new technologies and to adjust their methods for assessment and evaluation. Most countries lack clear legal and ethical processes for evaluating AI technology. If AI tools guide clinical decisions, they may be labelled as software as a medical device (SAMD), incurring stricter rules and requirements for their use, including validation. Without regulatory clarity, some AI vendors may self-designate their tools as non-medical devices to circumvent regulatory scrutiny intended for SAMD.
- **Inequities and bias in data sets**: As discussed above, AI technologies rely on large datasets for training and further development. The quality of the data the AI system builds on will determine the quality of the output, captured by the phrase 'garbage in, garbage out'. Most AI tools are built on datasets that predominantly originate from resource-rich countries, lacking representation of other nations and populations. Such suboptimal representation in datasets introduces bias, potentially perpetuating inequities through new AI technology. Bias may be inherent in the lifecycle of the production of the algorithm. Bias may creep into the algorithm during conceptualization, design, development and validation of the models being produced. There is also concern that data sets used to create AI models and their training can lose accuracy over time, a concept called data drift.[29] This has serious implications for potential errors within healthcare systems when AI models are unintentionally applied to populations outside of the training sets or when vendors fail to update the algorithms on a regular basis.

- **Black box phenomenon:** Machine learning or data-centric AI methods intrinsically lack explainability. This means that the reasoning behind a certain output or conclusion is not known, and in some cases, cannot be known, which means some types of AI operate as a black box. This lack of explainability is problematic in primary health care as it impacts consent processes, especially when diagnostic or therapeutic AI tools provide suggestions that seem illogical. This also makes it more challenging to identify and correct hallucinations that occur in generative AI applications.
- **Lack of co-design:** Most stakeholders involved in AI design and use will agree that ideally, AI design ought to include stakeholder engagement from early on in the process. However, robust methods for incorporating stakeholder involvement in the design process are still mostly lacking, and this is particularly true for ethical considerations. Value-sensitive design and embedded ethics methods show promise but have not yet been widely implemented in AI design.

B PROVIDER-RELATED CHALLENGES

- **Loss of skills, procedural medicine, dehumanisation**: Concerns have been raised that the increasing use of AI tools in primary care, specifically for diagnostic and therapeutic support, may negatively affect diagnostic skills and clinical reasoning in family doctors. This is closely related to the concern that AI implementation will require adaptations in the way we practice medicine, leading to more so-called 'procedural medicine', where clinical reasoning and actions operate within narrow guidelines, flowcharts, and decision trees. This may lead to a loss of richness and depth in clinical reasoning, a juxtaposition with narrative-based medicine, and might risk a dehumanisation of the primary care profession.[30]
- **Automation bias - overreliance on AI**: Closely related to the skill loss is the potential overreliance on AI tools. In particular, the use of generative AI introduces hallucinations as an inherent problem, especially with early generative AI models and open-access, free versions, necessitating careful review of their outputs. Counterintuitively, increased AI accuracy levels may lead to increased trust in its performance, which in turn increases the risk of missing errors or over-relying on the tools. This is referred to as *automation bias* in the literature.
- **Alert fatigue:** The use of AI tools may lead to a further inflation of alerts popping up in electronic health records and patient portal systems. Alert fatigue is not a new phenomenon and has been well-examined in the context of electronic health records and other areas of medicine. Although the goal for AI is to reduce medical errors, redundant and erroneous alerts may have the opposite effect, leading clinicians to dismiss or overlook important notifications. Over time, this can erode trust in digital systems and contribute to decision fatigue, a state of mental exhaustion from constant information processing and limited cognitive recovery, which can impair clinical judgment and increase the risk of errors.[31] A growing body of evidence highlights that health information technology-related stressors, such as insufficient time for documentation, after-hours EHR use, and poor system usability, are significantly associated with clinician burnout.[32] Without careful design and integration, AI risks becoming another contributor to the technological overload clinicians already face.

- **Responsibility and liability**: Given that some AI tools tend to operate more independently than any other technology, almost as a third agent in the room next to the family doctor and patient, questions arise regarding responsibility and liability, mainly in case of errors or hallucinations. Currently, responsibility falls mainly or entirely on the clinician. However, it could be imagined that the AI tool, or rather the industry partner behind it, will come to share some of this responsibility. How this shared responsibility could work, both from an ethical and legal perspective, is still unclear.
- **Hidden work**: While the focus is most often on how AI tools will help clinicians save time, reduce the administrative burden, and improve quality of care, there is a lack of consideration of the hidden work potentially caused by AI integration. LLMs are prone to errors and bias, including but not limited to so-called hallucinations. Family doctors will need to filter out these errors before integrating any AI-generated output. The need to regularly review and verify AI outputs can contribute to cognitive overload and decision fatigue among clinicians. When family doctors use AI tools that have only recently been released, they may need to address teething issues in the product, and in a sense, become part of testing the product as it becomes available to a wider audience. Lastly, incorporating AI tools into the daily clinical practice may require adjusting existing processes and workflows, which in itself will require effort and time. Acknowledging hidden work caused by AI integration into clinical practice is both a clinical and ethical prerequisite for its successful application.

C PATIENT AND POPULATION-RELATED CHALLENGES

- **Privacy and security breaches**: AI tools build on large, usually deidentified, datasets. We are increasingly seeing the interlinking of datasets at a population level to create even larger datasets for AI and research. The more data points are available on one individual, the higher the risk of reidentification becomes. Some have argued that all deidentified datasets should be treated as potentially identifiable, which changes the legal and ethical considerations for data use. Certain commercial AI tools permit the sharing of their de-identified datasets (or parts thereof) with third parties, which raises specific privacy concerns when health data are being shared. The risk of re-identification is real and needs to be managed appropriately.
- **Data and algorithmic bias**: When datasets are limited to specific patient populations, the output generated through AI tools may not be directly applicable to other patient populations. For example, AI tools that identify skin lesions have been shown to perform worse on dark skin than on white skin, because their training sets mainly include lesions on white skin.[33] Existing bias in healthcare is, by definition, inherent in any dataset used to train AI. This means that, unless we actively search for and eliminate such bias, AI-generated output risks perpetuating and even amplifying existing bias. This may be particularly problematic in diagnostic and therapeutic applications, but may even present itself in administrative tools such as AI scribes when they insert hallucinations in consultation notes reflecting bias.
- **Equity**: AI tools offer the potential to reduce digital and health inequities by supporting patient empowerment, but also through risk prediction, precision medicine, and diagnostic support. On the other hand, there is a risk that the use of AI may contribute to further widening existing inequities.[34] Underserved populations may not have access to AI tools and their potential benefits due to a lack of

hardware, software, internet access, and education.[35] Minorities may not be represented in datasets, resulting in AI output that is not well-adjusted to their needs. A risk of harm arises if AI tools become a replacement for traditional healthcare provision in underserved areas and populations: it may reduce the incentive to expand in-person healthcare access. Overreliance on AI tools, because there is no alternative option, may result in poor care despite the good intention of providing 'something' rather than 'nothing'.

- **Dehumanisation of care**: The increased use of AI in primary healthcare may lead to a decrease in human interactions. From an administrative perspective, we already see that some conversations are being replaced by the use of chatbots or other online communication systems. This may make scheduling more efficient, but it may also lead to patients having more detached experiences. Similarly, AI support for self-care and virtual coaching may contribute to a dehumanisation of care.
- **Lack of evidence:** A lack of robust evidence is not only a problem for AI tools, but is an issue for other health technologies as well, in particular self-monitoring apps.[26] Part of the problem is that AI tools may fall outside of medical device regulations (SAMD) if they do not directly involve diagnostic or therapeutic interventions (such as AI scribes). Moreover, the lack of prospective validation studies in real-life situations[19] means most of the existing evidence will have been generated under specific test conditions, which has limitations in understanding how the AI tool behaves in a real-life setting. This lack of robust evidence affects all types of users.
- **Misinformation and idle empowerment:** Generative AI is increasingly a potential source of misinformation, both deliberate and due to hallucinations. Patients may be exposed to harm or dismiss the validity of scientific evidence if the outputs are perceived as true or valid.[28] When patients turn to AI tools for information, idle empowerment may occur: they may seem or feel empowered by the tool, but it may be an idle empowerment if it is based on incorrect information and creates false beliefs.
- **Trust**: AI applications face challenges with both distrust and over-trusting. Some individuals and patients distrust healthcare organisations and institutions for valid reasons, such as past harm. This distrust may be transferred to AI tools as they become integrated into the healthcare system. On the other hand, patients may not have the knowledge, expertise, or skills to critically engage with AI agents and tools, and see them as equal to medical professional input.
- **Sustainability**: AI and LLM use significantly contribute to carbon emissions, and concerns have been raised about their substantial energy and water consumption.[36] However, some studies have argued that, in comparison to human labor, the environmental impact of LLMs is substantially lower.[37] LLMs also require human labor to run, which is often sourced in low-income countries, risking exploited labor including psychological harm (when training AI on violent, hateful, disturbing material etc).[38,39]

ETHICAL VALUES AND PRINCIPLES

that (should) guide AI use in primary care

Several ethical values and principles have been identified over the past decade that should guide the use of AI in general, and specifically in healthcare. Based on the current literature, we have selected **6 AI values** that we consider essential for AI use in primary care, but that also apply to other areas of AI application. Next to these values, we argue that **AI use in primary care should support primary care principles**†. To date, there are no universally agreed-upon primary care principles; however, recent work done by WONCA on core values in primary care has identified both common values and region-specific principles.[40] For the purpose of this document, we selected 5 primary care principles that are specifically relevant in the context of AI use. AI applications should ideally be trained with these principles to safeguard their outputs, or at a minimum support these values, to strengthen primary care and help it achieve its goals.

AI VALUES

1. **Patient safety and non-maleficence: do no harm principle.** Healthcare providers' first duty to the patient is to avoid causing harm, and the same principle should apply to AI. Similar to other interventions in health care, it is not always possible to avoid all harm. Understanding the specific risks of harm attributed to each AI tool is paramount. *The risk profile of any AI application used in primary care should be well understood and clearly communicated to all involved stakeholders*. Because many AI applications rely on continuous learning, robust monitoring systems should be in place to identify any potential harm early and, if possible, correct it. The technology industry, including project managers, designers, and developers, should take up responsibility for reliable risk identification processes and continued monitoring of any tools used in primary health.

   There are several different types of harm AI may cause: next to physical and mental harm to the individual, group harm and community harm due to bias or perpetuated discrimination should be considered. Harm may also occur through security breaches. Careful consideration of the risk of harm measured against expected benefits may result in a decision not to use a specific technology for a specific patient or patient group, or it may require adjustments to the standard use.

2. **Justice: equity and fairness.** Justice in terms of AI use in primary care has three distinct implications. First, unintended biases and discrimination should be prevented where possible, and processes should be in place to monitor, identify, and mitigate bias.[5,41] Careful attention should be paid to avoid exacerbating health disparities or widening the digital divide. Related to this, AI developers should take responsibility for ensuring that datasets are inclusive and diverse. Fair representation in datasets may require oversampling underrepresented groups.[5] Second, achieving

† The literature on primary care values uses different terms: core values, primary care principles, and often, these values and principles are incorporated in (core) competencies. For the purpose of this guidance document, we will use values and principles interchangeably; their definition, differences between them, and meaning within the field of (bio)ethics fall outside of the scope of this work.

equitable access requires that AI technology be widely available, including within primary care in low- and middle-income countries. From a health economics perspective, this raises familiar questions about financing and universal health coverage. Deploying AI tools usually requires significant upfront capital investment and ongoing costs. If AI-supported services are funded predominantly through out-of-pocket payments, patients who are better off will have preferential access, thereby widening existing inequalities. Ethical implementation of AI-supported services therefore requires transparent decisions about who pays for them. Furthermore, investment in AI should be aligned with national strategies to strengthen primary care, so that it does not undermine the availability and quality of basic, non-digital services through the opportunity costs of diverting scarce resources. Third, justice also applies to the benefits of AI to the workforce in terms of efficiency gains. The time released by AI must help clinicians with existing work rather than 1) generate more work (because of acceleration of clinical work e.g. by increased case-finding or errors that need time to rectify) or 2) result in the allocation of more 'complex cases' instead of allowing more time with existing patients. The implications of not doing this are a worse service to the public and an increased work burden on the primary care workforce. Given that AI technologies are marketed to the public as helpful to the workforce, this results in both an injustice to the workforce and a testimonial injustice: the workforce is, in fact, helped less while the public is told they are being helped more. Consequently, any serious discussion of the ethics of AI implementation must consider the workflow implications downstream from application.

3. **Beneficence: reciprocity and solidarity.** The use of AI in primary care should enhance the health of patients and communities, while supporting the role of family doctors in care provision. The benefits of AI should be returned to the individuals who contributed their data, with processes in place to ensure that AI technology benefits the communities that bear the burdens of technology development. Similarly, AI technologies should be made available to those communities who stand to gain the most from AI, with specific attention paid to including and reaching underserved populations. This is why reciprocity and solidarity are emphasised within the value of beneficence.
4. **Transparency**: We should be mindful that explainability, in particular with the use of generative AI, may be a lofty but unattainable goal given the ‘black box’ nature of these applications. However, understanding the exact processes through which AI reached its output may not be necessary, as long as we have access to robust evidence and a thorough understanding of its outputs through real-world tests. Rather than emphasising the need for explainability, AI use in primary care should be transparent in terms of understanding the datasets and training data, limitations and potential risks of harm, and having robust evidence for its effects, including through ongoing monitoring after deployment. In this sense, regulation around what type of evidence new AI applications for primary care need to provide through validation processes has to be put in place so that Family doctors and patients can trust the technology they are using. Next to this, transparency in communication is also paramount. For-profit technology providers tend to highlight the benefits of their products or even overpromise on what they can deliver. Transparent disclosure of information around the benefits, risks, and realistic expectations of AI technology for primary care is essential and should be tailored to different stakeholder groups. In particular, communication with patients should avoid technical jargon and be adjusted to their needs.

5. **Data governance and sovereignty**: privacy, confidentiality, autonomy. Traditionally, protecting patient privacy and confidentiality regarding data use requires individual informed consent. Because AI builds on extensive datasets, we acknowledge that valid informed consent will not always be attainable. Moreover, if not executed well, it will lose some of its value if it is treated as the mandatory signing off of terms and conditions, as is the case in other areas of technology use. However, protecting patients' privacy and autonomy remains an important moral goal in AI use in primary care. This can be achieved by putting data protection laws in place[5] alongside data sovereignty principles in regulatory frameworks. Independent data governance groups can ensure the proper application of such laws and principles. This means that some of the burden of informed consent is transferred from the individual to the regulatory level. For example, clearly understanding the risk of reidentification in a dataset may be very challenging to the individual patient, but it may be managed at a governance level. This also means that data sovereignty principles can be adapted to the local context and needs, for example, by incorporating Indigenous data sovereignty principles.[42,43]
6. **Sustainability**: Sustainability in terms of AI use in primary care relates to both sustainable systems and protection of the environment. Existing healthcare systems and processes will need to be adjusted to ensure the successful and sustainable integration of new AI tools. This may require substantially more time and effort than simply procuring AI technology, and this should be taken into account. This may also mean potential job losses as AI technology may come to replace certain roles in healthcare.[5] Careful thought should go toward understanding the impact of AI technology implementation on existing healthcare systems. Secondly, AI systems contribute significantly to carbon emissions and water and energy consumption. New AI technologies should be designed and implemented in a way that minimises their carbon footprint and maximises energy efficiency.[41]

PRIMARY CARE PRINCIPLES

Digital health has been identified by family doctors as a challenge to the core values in family medicine.[44] In addition to the above-listed AI values, any AI implementation in primary care should also support the core values or ethical principles central to primary care. These principles need to be protected from erosion due to the integration of AI technology. Rather, AI applications should be designed to support and strengthen these principles.

1. **Patient-centered, comprehensive, and holistic care**: primary care provides person-centered care not just focused on disease, illness, health promotion, and how the patient experiences these, but on the whole person, with a generalist approach to the person and their family.[45,46] Attention also goes to the social determinants of health and patients' life stories.[47] The patient-doctor relationship builds on a therapeutic alliance in which power and responsibility are shared.[46] This means AI tools need to take into account the patient and their context, and they need to be able to be tailored to individual patient needs, i.e., one-size-fits-all solutions are unlikely to support the person-centered care primary care seeks to provide. For example, any AI tools aimed at patient self-management need to weigh individual responsibility against social determinants of health, and this should be carefully managed from the design phase onwards.

2. **Continuity of care**: In primary care literature, continuity of care is usually defined as the "relationship between a single practitioner and a patient that extends beyond specific episodes of illness or disease".[48] Although continuity of care is mainly understood as building a relationship over time through repeated interactions with a designated provider, it may also mean providing a team approach to patients under the umbrella of continuity of responsibility[45,49], or informational continuity, understood as formally recording information and making it available to all involved in a patient's care.[50,51] There is evidence that continuity of care may contribute to reduced hospital admissions, improved quality of care, the prevention of chronic illnesses, and even reduced mortality.[50,52,53] Regarding AI implementation in primary care, this implies that we should be sceptical of tools that primarily focus on episodic care without supporting continuity. Certain AI technologies may erode the patient-provider relationship rather than support it by reducing opportunities for patient-provider interactions. Any AI tools should be carefully integrated within the patient-provider relationship, aiming to support rather than weaken continuity of care.
3. **Accessibility**: Although accessibility or 'first contact' was one of Barbara Starfield's 4C's or initial four defining features of primary care, it seems to have received less emphasis in more recent literature on the core values of family medicine.[44–47,54,55] There should be a low threshold to consultation in primary care, as this level of care serves as the first point of entry into the healthcare system. Part of this is captured under equity of care and a community orientation, aiming to minimise inequities in health service access and delivery and allocating time and effort to those patients who need it the most.[47] Improving accessibility is often touted as one of the benefits of digital health implementation. Often, there are clear benefits in terms of accessibility through convenience: AI scheduling systems can facilitate appointment management, patients can receive advice through an AI chatbot in their own home, and AI tools can support telehealth consultations. However, AI tools should also support equitable access with the aim of reducing inequities and providing inclusive care.
4. **Community orientation**: The focus of primary care should move beyond the individual patient into promoting health at the community level, integrating clinical care, public health, and social services.[46,47] It may be challenging for primary care clinicians to reach this goal as it requires attention to population health, advocacy, political, and social engagement, which often fall outside of the usual working hours.[47] AI may play an important role in supporting this principle. To date, AI applications in primary care have primarily focused on the individual patient-provider encounter. However, AI could support the analysis of population data, resulting in better proactive care and greater health equity.[56]
5. **Professionalism and life-long learning**: This principle acknowledges that sustained competence and performance over a career require a continuous process of personal and professional development.[45,46] Professionalism also requires the implementation and promotion of research, including the development and integration of clinical and other guidelines into primary care.[47] Several countries have set up primary-care-led research networks for this purpose.[45] There will be requirements to assess the needs and determine barriers to participation in continuous professional development programs to upskill clinicians in the use of AI in primary care.[57] AI tools should support both education and research in primary care and facilitate lifelong learning.

PRACTICAL GUIDANCE

for the ethical use of AI tools in daily practice

The following practical recommendations aim to promote the cautious, ethical, and context-sensitive adoption of AI tools in daily primary care practice. They are not intended to discourage innovation, but rather to ensure that AI tools fulfil their potential while preserving values central to primary care delivery. We recommend that clinicians critically appraise tools using the criteria below, with special caution around generative AI applications, AI tools that have not been validated, and/or those developed by for-profit companies that have not been specifically designed for use in primary healthcare.

IN THIS SECTION

2 fundamental requirements

9 considerations for promising AI tools

30 questions for reviewing an AI tool

We start with 2 fundamental requirements for AI use in primary care.

1 VALIDATION AND REGULATION

Family doctors should only use AI tools that have been validated and that meet regulatory standards specific to clinical use in primary care.

This is often not possible as robust validation processes and local and international regulations for AI may not have been established yet. If validation and regulation processes are unknown, we recommend turning to Colleges, primary care departments at (local) Universities, indemnity insurance providers, and government bodies to check if any regulation is in place, in particular for primary care, and if not, to advocate for its establishment. In the future, independent AI oversight bodies may provide a fitting solution to the problem of AI validation and regulation in primary care [58].

2 COLLABORATION

The implementation of AI in primary care is too complex to be delegated to individual practitioners.

Many family doctors work in isolated settings or small teams and may not have access to larger networks. However, implementing AI in primary care will require teamwork. Our advice is to unite, share experiences, seek guidance, and collaborate at the local, regional, and international levels. This may include mentorship and support from more experienced clinicians to others with limited experience across regions and countries. Again, colleges, universities, and other professional organisations could play a central role here, in providing education and training for family doctors.

Next, we acknowledge that most AI tools have not been thoroughly validated and that robust AI regulatory processes are lacking in most countries. What are family doctors to do in these situations, where they may have access to promising AI tools that have some evidence, but have not been fully validated or are unregulated? We propose the following reflection, along with prompts and questions (in Table 1), for family doctors to use to determine whether an AI tool is suitable for their local primary care setting. *The default position should be not to use generative AI applications in primary care unless they have undergone sufficient scrutiny and meet ethical and clinical requirements (precautionary principle).* Family doctors may not be able to answer all of the following questions when considering the use of an AI tool, or it may not be clear whether an AI tool meets a minimum standard for use in primary care based on these questions. In such cases, caution is recommended when implementing an AI tool.

1 EVIDENCE

Tools should be supported by independent evidence and tested in real-world clinical settings.

When considering the use of unvalidated AI tools, the first step is to review the available evidence on the tool, including early test results in trials and outcomes in real-world settings. Ideally, published studies would be available, but this is often not the case.

2 RISK OF HARM

Potential harms should be assessed across multiple domains, including patient, clinician and system-level impacts.

Family doctors should have a clear understanding of the potential harm of any technology or treatment used in primary care, and this applies to AI tools as well. This may involve a risk of harm to the patient, but also to the clinician or the healthcare system at the local or broader level.

3 DATA REVIEW

The training data and ongoing data collection should be reviewed to ensure the AI tool is relevant to the local context and population and meets privacy and confidentiality requirements.

Given that all AI tools build on datasets, it is useful to review these underlying data sources. This would require not only an understanding of the training data that was used to create the tool, but also of ongoing data collection after deployment. There is a tension between allowing data collection to improve AI tool performance, and data protection to ensure privacy, confidentiality, and data sovereignty principles in primary care.

4 CO-DESIGN

Tools should be developed with meaningful input from clinicians, patients, and communities.

When reviewing an AI tool for use in primary care, it is useful to examine whether clinicians, patients, and other relevant stakeholders have been involved in its design and development. Before implementing an AI tool, family doctors need to know what processes are in place to provide feedback to the AI provider.

5 PATIENT INFORMATION AND CONSENT

Use of AI tools requires appropriate patient disclosure, transparent information, and, where necessary, explicit consent.

Patients should receive adequate information about how AI tools are implemented into their care. In some cases, this will require explicit consent; in other cases, information provision will suffice. Family doctors need to carefully consider whether full informed consent would be required, and how information provision will be handled.

6 INCLUSION AND EQUITY

Implementation should avoid reinforcing disparities or excluding vulnerable groups.

Ideally, AI implementation contributes to closing inequity gaps and benefits those patients who need it the most. At a minimum, AI tools should not make existing inequities worse or create new ones. Before implementing an AI tool, family doctors should consider its potential impact on equity, particularly on underserved populations.

7 SUSTAINABILITY AND PLANETARY HEALTH

AI tools may have a substantial impact on carbon emissions, water use, and overall planetary health, which should be weighed against potential benefits.

Introducing AI tools in general practice may have a substantial impact on carbon emissions, water use, and overall planetary health. It is important for family doctors to understand these impacts and their consequences, and to weigh them against the potential benefits of AI use.

8 INTEGRATION AND IMPACT ON WORKFLOW

AI tools should align with clinical workflows and support rather than disrupt team-based care.

Integration of AI tools may impact existing clinical and administrative processes and overall workflow. This impact should be considered before implementation so family doctors can prepare accordingly. Complex clinical judgment and decision-making should not be delegated to unvalidated or unregulated AI tools.

9 ONGOING MONITORING

Use of AI tools should be continually monitored for performance, safety, and unintended consequences.

As stated above, family doctors should understand the processes in place to provide feedback on the AI tool; however, they should also consider establishing internal and external monitoring processes to track the tool's implementation.

Table 1: Questions for the review process of AI tools in primary care

CONSIDERATION	QUESTIONS FOR THE REVIEW
Evidence	**1.** Has the AI tool been through some form of validation, internally or externally? **2.** What evidence is available, and what is the quality of the available evidence? **3.** Is information available on the performance of the tool in real-world settings? **4.** Has the tool been tested specifically in primary care settings? **5.** What is known about the tool's reliability?
Risk of harm	**6.** What are the possible risks of using this tool for patients, and how can they be mitigated? **7.** Are any risks presented and communicated by the AI provider in a clear and transparent way? **8.** Could the tool negatively affect clinical judgement, well-being, and workload? **9.** Might it disrupt workflows or widen existing system inequities?
Data review	**10.** Was this tool trained on data from relevant settings? **11.** Is the tool applicable to the target population, or are adjustments needed? **12.** Are there risks of bias when using this tool in this patient population? **13.** Are privacy and confidentiality protected sufficiently? Is the tool secure?
Co-design	**14.** Were end users involved in the design and development process? **15.** Is the tool designed for use in general practice, and if not, do any adaptations need to be made for this use? **16.** Are there any processes in place for clinicians and patients to provide feedback on the AI tool?
Patient information and consent	**17.** Is explicit patient consent required for the use of this AI? If so, how will this be handled appropriately? **18.** How will the use of the tool be explained to patients in a transparent way? **19.** What are the consequences of opting out, and how will opt-out be handled?

CONSIDERATION	QUESTIONS FOR THE REVIEW
Inclusion and equity	**20.** Would certain patient populations encounter barriers in using this AI application, e.g., older adults, non-English speakers, rural populations? **21.** Could its use have any negative consequences for certain populations, and if yes, how could these be addressed? **22.** Is the AI tool culturally safe and appropriate?
Sustainability and planetary health	**23.** What is the impact of the tool on carbon emissions and water use? **24.** How do the potential benefits weigh up against the negative impact in terms of planetary health?
Integration	**25.** How will the integration of this AI tool impact workflow processes? **26.** Do any changes need to be made to existing processes? **27.** Could it lead to job loss, and if so, how would it be handled? **28.** How will this tool change clinical processes? Will it ease the workload or add new complexity or burdens?
Ongoing monitoring	**29.** Is a process in place to monitor the AI tool's performance and safety, log concerns, document possible harms, and refine use over time? **30.** If harm occurs, how will it be managed?

NEXT STEPS

roundtable dialogues for ethical AI in primary care

We propose holding future roundtables to bring together clinicians, patients, regulators, technology vendors, ethicists, and EDI (Electronic Data Interchange) experts to help formulate future roadmaps for the ethical development and eventual implementation of AI systems in healthcare. WONCA could take a leading role in organising these roundtables and ensuring primary care voices are heard and included in AI design and implementation. Three initial priorities may include:

1. **Equity, accessibility, and inclusive design**

   Despite consensus that co-design with patients and practitioners improves usability and acceptability[59], meaningful engagement is rare in AI development. Roundtables can explore co-design from the stages of needs assessment, model training, validation and deployment, thus ensuring the tools are reflective of the realities of primary care. To ensure equitable access to AI solutions and avoid perpetuating inequities, dialogue with stakeholders is essential to determine accessibility standards that address cultural, language, disability, connectivity, and digital literacy barriers. Particular attention to Indigenous populations and other underserved communities would ensure that AI benefits reach underserved groups, thereby reducing existing digital divides.[42]

2. **Data governance and consent models**

   A second suggested focus of the roundtables is a discussion around data governance and consent models. Maintaining individual consent and data control is becoming increasingly challenging as data is shared more widely for building AI tools and for secondary research. One option to address this is to examine the current state and future prospects of blockchain technology as a solution to ensure patient stewardship and control over healthcare data. This would enable patient agency, granular consent, and voluntary data sharing, accompanied by potential rewards or incentives to return benefits to contributing communities or individuals. This could be achieved within required privacy legislation and population-specific sovereignty frameworks.[60] The other option is to move away from individual consent models to a focus on collective ownership and data governance models. This would include ensuring roles for Indigenous Peoples and other underserved communities in data decision-making and protecting access to data as indicated by these communities to minimize harms and maximise benefits.[42]

3. **Regulatory harmonisation for AI Software as a Medical Device (SAMD) and Clinical Decision Support Systems (CDSS)**

   A blurring of boundaries currently exists between SAMD and CDSS in Generative AI applications, raising both regulatory and ethical concerns. While there is guidance on quality, safety, and risk management through standards such as Medical Device Single Audit Program (MDSAP), ISO 13485, and ISO 14971, there is little focus on ethical considerations for these tools and their application in primary care – particularly transparency of training data, equity of applicability across diverse

populations, and mitigation of algorithmic bias. Representation of primary care stakeholders within these governing bodies, as well as roundtable discussions on future standards and regulatory policy, could serve to integrate these ethical imperatives alongside technical compliance.[61–64]

REFERENCE LIST

1. Royal College of General Practitioners. *Artificial Intelligence and Primary Care*. London: RCGP; 2019. https://web.archive.org/web/20220615191943/https://www.rcgp.org.uk/-/media/Files/CIRC/CIRC-AI-REPORT.ashx
2. Coeckelbergh M. *AI Ethics*. MIT Press; 2020.
3. RACGP Practice Technology and Management. *Artificial Intelligence in Primary Care*. 2024. https://www.racgp.org.au/advocacy/position-statements/view-all-position-statements/clinical-and-practice-management/artificial-intelligence-in-primary-care
4. Esteva A, Robicquet A, Ramsundar B, et al. A guide to deep learning in healthcare. *Nat Med*. 2019;25(1):24-29. doi:10.1038/s41591-018-0316-z
5. World Health Organization. *Ethics and Governance of Artificial Intelligence for Health: WHO Guidance*. Geneva: World Health Organization; 2021. ISBN 9789240029200. https://www.who.int/publications/i/item/9789240029200
6. Li J, Dada A, Puladi B, Kleesiek J, Egger J. ChatGPT in healthcare: A taxonomy and systematic review. *Comput Methods Programs Biomed*. 2024;245:108013. doi:10.1016/j.cmpb.2024.108013
7. Raza MM, Venkatesh KP, Kvedar JC. Generative AI and large language models in health care: pathways to implementation. *npj Digit Med*. 2024;7(1). doi:10.1038/s41746-023-00988-4
8. Mesko B. The ChatGPT (Generative Artificial Intelligence) Revolution Has Made Artificial Intelligence Approachable for Medical Professionals. *J Med Internet Res*. 2023;25. doi:10.2196/48392
9. Park YJ, Pillai A, Deng J, et al. Assessing the research landscape and clinical utility of large language models: a scoping review. *BMC Med Inform Decis Mak*. 2024;24(1):1-14. doi:10.1186/s12911-024-02459-6
10. Wornow M, Xu Y, Thapa R, et al. The shaky foundations of large language models and foundation models for electronic health records. *npj Digit Med*. 2023;6(1):1-10. doi:10.1038/s41746-023-00879-8
11. AlSaad R, Abd-Alrazaq A, Boughorbel S, et al. Multimodal Large Language Models in Health Care: Applications, Challenges, and Future Outlook. *J Med Internet Res*. 2024;26. doi:10.2196/59505
12. Kueper JK, Terry A, Bahniwal R, Meredith L, Beleno R, Brown JB, et al. Connecting artificial intelligence and primary care challenges: findings from a multi stakeholder collaborative consultation. *BMJ Health & Care Informatics*. 2022;29:e100493. https://doi.org/10.1136/bmjhci-2021-100493
13. WONCA Working Party on the Environment. *The WONCA Working Party on the Environment Statement to Protect Health by Committing to Fossil Fuel Non-Proliferation and Climate Action*. 2023. https://www.globalfamilydoctor.com/site/DefaultSite/filesystem/documents/Groups/Environment/WONCA%20Position%20Statement%20on%20Fossil%20Fuels_final_20230220.pdf
14. Reddy S. Generative AI in healthcare: an implementation science informed translational path on application, integration and governance. *Implement Sci*. 2024;19(1):1-15. doi:10.1186/s13012-024-01357-9
15. Tierney AA, Gayre G, Hoberman B, et al. Ambient Artificial Intelligence Scribes to Alleviate the Burden of Clinical Documentation. *NEJM Catal*. 2024;5(3):1-15. doi:10.1056/cat.23.0404
16. Bouderhem R. Shaping the future of AI in healthcare through ethics and governance. *Humanit Soc Sci Commun*. 2024;11(1). doi:10.1057/s41599-024-02894-w
17. Arowosegbe A, Oyelade T. Application of Natural Language Processing (NLP) in Detecting and Preventing Suicide Ideation: A Systematic Review. *Int J Environ Res Public Health*. 2023;20(2). doi:10.3390/ijerph20021514
18. Han C, Kim DW, Kim S, et al. Evaluation of GPT-4 for 10-year cardiovascular risk prediction: Insights from the UK Biobank and KoGES data. *iScience*. 2024;27(2):109022. doi:10.1016/j.isci.2024.109022

19. Topol EJ. High-performance medicine: the convergence of human and artificial intelligence. *Nat Med*. 2019;25(1):44-56. doi:10.1038/s41591-018-0300-7

20. Romero-Aroca P, Verges-Puig R, De La Torre J, et al. Validation of a Deep Learning Algorithm for Diabetic Retinopathy. *Telemed e-Health*. 2020;26(8):1001-1009. doi:10.1089/tmj.2019.0137

21. Krakowski I, Kim J, Cai ZR, et al. Human-AI interaction in skin cancer diagnosis: a systematic review and meta-analysis. *npj Digit Med*. 2024;7(1):1-10. doi:10.1038/s41746-024-01031-w

22. Malla AM, Amin U. Scope of technology in health care, special focus on nursing. *J Integr Nurs*. 2023;5(4):300-310. doi:10.4103/jin.jin_92_23

23. Anthamatten A, Holt JE. Integrating Artificial Intelligence Into Virtual Simulations to Develop Entrustable Professional Activities. *J Nurse Pract*. 2024;20(9):105192. doi:10.1016/j.nurpra.2024.105192

24. Ning Y, Ong JCL, Cheng H, et al. How can artificial intelligence transform the training of medical students and physicians? *Lancet Digit Health*. 2025;7(10):100900. doi:10.1016/j.landig.2025.100900

25. Lin S. A Clinician's Guide to Artificial Intelligence (AI): Why and How Primary Care Should Lead the Health Care AI Revolution. *J Am Board Fam Med*. 2022;35(1):175. doi:10.3122/jabfm.2022.01.210226

26. Morley J, Laitila J, Ross JS, Schamroth J, Zhang J, Floridi L. An App a Day will (Probably Not) Keep the Doctor Away: An Evidence Audit of Health and Medical Apps Available on the Apple App Store. *Minds Mach*. 2025;35(1):1-30. doi:10.1007/s11023-025-09710-7

27. Petrozzino C. Who pays for ethical debt in AI? *AI Ethics*. 2021;1(3):205-208. doi:10.1007/s43681-020-00030-3

28. Monteith S, Glenn T, Geddes JR, Whybrow PC, Achtyes E, Bauer M. Artificial intelligence and increasing misinformation. *Br J Psychiatry*. 2024;224(2):33-35. doi:10.1192/bjp.2023.136

29. Kelly CJ, Karthikesalingam A, Suleyman M, Corrado G, King D. Key challenges for delivering clinical impact with artificial intelligence. *BMC Med*. 2019;17(1):1-9. doi:10.1186/s12916-019-1426-2

30. Ghenimi N, Govender R, Moodley K. The paradox of artificial intelligence (AI) and narrative-based medicine: challenges and potential for enhanced patient care. *AI Soc*. 2026;41(1):251-257. doi:10.1007/s00146-025-02418-3

31. Grignoli N, Manoni G, Gianini J, Schulz P, Gabutti L, Petrocchi S. Clinical decision fatigue: a systematic and scoping review with meta-synthesis. *Family Medicine and Community Health*. 2025;13:e003033. https://doi.org/10.1136/fmch-2024-003033

32. Gardner RL, Cooper E, Haskell J, et al. Physician stress and burnout: the impact of health information technology. *J Am Med Informatics Assoc*. 2019;26(2):106-114. doi:10.1093/jamia/ocy145

33. Daneshjou R, Vodrahalli K, Novoa RA, et al. Disparities in dermatology AI performance on a diverse, curated clinical image set. *Sci Adv*. 2022;8(32):1-7. doi:10.1126/sciadv.abq6147

34. Moodley K. Artificial intelligence (AI) or augmented intelligence? How big data and AI are transforming healthcare: Challenges and opportunities. *S Afr Med J*. 2023;114(1):22-26. doi:10.7196/SAMJ.2024.v114i1.1631

35. Malpani R, Moodley K. Large multimodal models: boon or burden for low- and middle-income countries. *Int Health*. Published online 2025:1-3. doi:10.1093/inthealth/ihae073

36. Lucivero F. Big Data, Big Waste? A Reflection on the Environmental Sustainability of Big Data Initiatives. *Sci Eng Ethics*. 2020;26(2):1009-1030. doi:10.1007/s11948-019-00171-7

37. Ren S, Tomlinson B, Black RW, Torrance AW. Reconciling the contrasting narratives on the environmental impact of large language models. *Sci Rep*. 2024;14(1):1-8. doi:10.1038/s41598-024-76682-6

38. Perrigo B. OpenAI Used Kenyan Workers on Less Than $2 Per Hour to Make ChatGPT Less Toxic. *Time Magazine*. https://time.com/6247678/openai-chatgpt-kenya-workers/. Published January 18, 2023.

39. Dzieza J. AI Is a Lot of Work. *The Verge*. Published online June 2023. https://www.theverge.com/features/23764584/ai-artificial-intelligence-data-notation-labor-scale-surge-remotasks-openai-chatbots

40. Stavdal A, Sigurdsson JA, Goodyear-Smith F, eds. *Core Values in Family Medicine: Inspiring Global Change*. CRC Press, Taylor and Francis Group; 2026. doi:10.1201/9781003542353

41. Jobin A, Ienca M, Vayena E. The global landscape of AI ethics guidelines. *Nat Mach Intell*. 2019;1(9):389-399. doi:10.1038/s42256-019-0088-2

42. Carroll SR, Garba I, Figueroa-Rodríguez OL, et al. The CARE principles for indigenous data governance. *Data Sci J*. 2020;19(1):1-12. doi:10.5334/DSJ-2020-043

43. Brown PT, Wilson D, West K, et al. Māori Algorithmic Sovereignty: Idea, Principles, and Use. *Data Sci J*. 2024;23:1-16. doi:10.5334/dsj-2024-015

44. Arvidsson E, Švab I, Klemenc-Ketiš Z. Core Values of Family Medicine in Europe: Current State and Challenges. *Front Med*. 2021;8:646353. doi:10.3389/fmed.2021.646353

45. Hashim M. Principles of family medicine and general practice - defining the five core values of the specialty. *J Prim Health Care*. 2016;8(4):283-287. doi:10.1071/HC16006

46. Lawson HJO, Nortey DNN. Core Values of Family Physicians and General Practitioners in the African Context. *Front Med*. 2021;8:667144. doi:10.3389/fmed.2021.667144

47. WONCA Europe Core Values Working Group. *Core Values and Principles of General Practice/Family Medicine*.; 2023. https://www.woncaeurope.org/page/our-core-values

48. Haggerty JL, Reid RJ, Freeman GK, Starfield BH, Adair CE, McKendry R. Continuity of care: A multidisciplinary review. *Br Med J*. 2003;327(7425):1219-1221. doi:10.1136/bmj.327.7425.1219

49. McWhinney IR. Primary care: core values. Core values in a changing world. *BMJ*. 1998;316(7147):1807-1809. doi:10.1136/bmj.316.7147.1807

50. Baker R, Freeman GK, Haggerty JL, Bankart MJ, Nockels KH. Primary medical care continuity and patient mortality: a systematic review. *Br J Gen Pract*. 2020;70(698):e600-e611. doi:10.3399/bjgp20X712289

51. Guthrie B, Saultz JW, Freeman GK, Haggerty JL. Continuity of care matters. *BMJ*. 2008;337:a867. doi:10.1136/bmj.a867

52. Chan KS, Yuk-Fai Wan E, Chin WY, et al. Effects of continuity of care on health outcomes among patients with diabetes mellitus and/or hypertension: a systematic review. *BMC Fam Pract*. 2021;22(145). doi:10.1186/s12875-021-01493-x

53. Chau E, Rosella LC, Mondor L, Wodchis WP. Association between continuity of care and subsequent diagnosis of multimorbidity in Ontario, Canada from 2001–2015: A retrospective cohort study. *PLoS One*. 2021;16(3). doi:10.1371/JOURNAL.PONE.0245193

54. Starfield BH. *Primary Care: Concept, Evaluation, and Policy*. Oxford University Press; 1992.

55. Kolesnyk P, Bayen S, Shushman I, et al. Identification and Ranking of Core Values in Family Medicine: A Mixed Methods Study From Ukraine. *Front Med*. 2021;8(March):1-9. doi:10.3389/fmed.2021.646276

56. Basu S, Bermudez-Canete P, Hall TC, Rajpurkar P. Optimizing AI solutions for population health in primary care. *npj Digit Med*. 2025;8(1):1-3. doi:10.1038/s41746-025-01864-z

57. Gavarkovs AG, Kueper J, Arntfield R, Myslik F, Thompson K, Mccauley W. Assessing Physician Motivation to Engage in Continuing Professional Development on Artificial Intelligence. *J Contin Educ Health Prof*. 2025;45(4):287-292. doi:10.1097/CEH.0000000000000594

58. Bockting CL, van Dis EAM, van Rooij R, Zuidema W, Bollen J. Living guidelines for generative AI —why scientists must oversee its use. *Nature*. 2023;622(7984):693-696. doi:10.1038/d41586-023-03266-1

59. Sumner J, Ng CWT, Teo KEL, Peh ALT, Lim YW. Co-designing care for multimorbidity: a systematic review. *BMC Med*. 2024;22(1). doi:10.1186/s12916-024-03263-9

60. Kuo TT, Kim HE, Ohno-Machado L. Blockchain distributed ledger technologies for biomedical and health care applications. *J Am Med Informatics Assoc*. 2017;24(6):1211-1220. doi:10.1093/jamia/ocx068

61. IMDRF SaMD Working Group. Software as a Medical Device (SaMD): Key Definitions. *Int Med Device Regul Forum*. Published online 2013:1-9. http://www.imdrf.org/docs/imdrf/final/technical/imdrf-tech-131209-samd-key-definitions-140901.pdf

62. International Organization for Standardization. *ISO 13485:2016 Medical Devices—Quality Management Systems—Requirements for Regulatory Purposes*.; 2016. https://www.iso.org/standard/59752.html

63. Medical Device Single Audit Program. What is MDSAP? Accessed October 16, 2026. https://www.mdsap.global/about/what-mdsap

64. International Organization for Standardization. *ISO 14971:2019 Medical Devices—Application of Risk Management to Medical Devices*. 3rd ed. International Organization for Standardization; 2019. https://www.iso.org/standard/72704.html

ABOUT WONCA

WONCA, the World Organization of Family Doctors, is a not-for-profit organization founded in 1972. Through its 133 Member Organizations in 140 countries and territories, it represents about 500,000 family doctors who serve more than 90 per cent of the world's population across seven regions. WONCA works to strengthen family medicine and primary care worldwide.

WONCA WORKING PARTY ON ETHICS AND PROFESSIONALISM

The WONCA Working Party on Ethics and Professionalism is devoted to encouraging the maintenance of the highest ethical standards in practice, education and research in primary care. It does this by promoting widespread dialogue within the profession, initiating and supporting research, and disseminating information throughout the WONCA community, the international academic community, healthcare organisations, and government agencies.

Wonca
World family doctors. Caring for people.

wonca.com