

Patient Self-Administered Screening for Cardiovascular Disease Using Artificial Intelligence in the Home

*Patrik Bächtiger, Mihir A. Kelshiker, Marie E. G. Moe,
Daniel B. Kramer, and Nicholas S. Peters*

I INTRODUCTION

The United Kingdom (UK) National Health Service (NHS) is funding technologies for home-based diagnosis that draw on artificial intelligence (AI).¹ Broadly defined, AI is the ability of computer algorithms to interpret data at human or super-human levels of performance.² One compelling use case involves patient-recorded cardiac waveforms that are interpreted in real time by AI to predict the presence of common, clinically actionable cardiovascular diseases. In this case, both electrocardiograms (ECGs) and phonocardiograms (heart sounds) are recorded by a handheld device applied by the patient in a self-administered smart stethoscope examination, communicating waveforms to the cloud via smartphone for subsequent AI interpretation – principally known as AI-ECG. Validation studies suggest the accuracy of this technology approaches or exceeds many established national screening programs for other diseases.³ More broadly, the combination of a new device (a modified handheld stethoscope), novel AI algorithms, and communication via smartphone coalesce into a distinct clinical care pathway that may become increasingly prevalent across multiple disease areas.

However, the deployment of a home-based screening program combining hardware, AI, and a cloud-based digital platform for administration – all anchored in patient self-administration – raises distinct ethical challenges for safe, effective, and trustworthy implementation. This chapter approaches these concerns in five parts. First, we briefly outline the organizational structure of the NHS and associated regulatory bodies responsible for evaluating the safety of medical technology. Second,

¹ United Kingdom Government Department of Health and Social Care, *Health Secretary Announces £250 Million Investment in Artificial Intelligence*, Gov.UK (August 8, 2019), www.gov.uk/government/news/health-secretary-announces-250-million-investment-in-artificial-intelligence.

² Patrik Bächtiger, et al., *Artificial Intelligence, Data Sensors and Interconnectivity: Future Opportunities for Heart Failure*, *Cardiac Failure Rev.* 6 (2020).

³ Patrik Bächtiger, et al., *Point-of-Care Screening for Heart Failure with Reduced Ejection Fraction Using Artificial Intelligence during ECG-Enabled Stethoscope Examination in London, UK: A Prospective, Observational, Multicentre Study*, 4 *Lancet Digit. Health* 117, 117–25 (2022).

we highlight NHS plans to prioritize digital health and the specific role of AI in advancing this goal with a focus on cardiovascular disease. Third, we review the clinical imperative for early diagnosis of heart failure in community settings, and the established clinical evidence supporting the use of a novel AI-ECG-based tool to do so. Fourth, we examine the ethical concerns with the AI-ECG diagnostic pathway according to considerations of equity, agency, and data rights across key stakeholders. Finally, we propose a multi-agency strategy anchored in a purposefully centralized view of this novel diagnostic pathway – with the goal of preserving and promoting trust, patient engagement, and public health.

II THE UK NATIONAL HEALTH SERVICE AND RESPONSIBLE AGENCIES

For the purposes of this chapter, we focus on England, where NHS England is the responsible central government entity for the delivery of health care (Scotland, Wales, and Northern Ireland run devolved versions of the NHS). The increasing societal and political pressure to modernize the NHS has led to the formation of agencies tasked with this specific mandate, each of which plays a key role in evaluating and deploying the technology at issue in this chapter. Within NHS England, the NHSX was established with the aim of setting national NHS policy and developing best practices across technology, digital innovation, and data, including data sharing and transparency. Closely related, NHS Digital is the national provider of information, data, and IT systems for commissioners, analysts, and clinicians in health and social care in England. From a regulatory perspective, the Medicines and Healthcare products Regulatory Agency (MHRA) is responsible for ensuring that medicines and medical devices (including software) work and are acceptably safe for market entry within the scope of their labelled indications. Post Brexit, the UK's underlying risk-based classification system remains similar to that of its international counterparts, categorizing risk into three incremental classes determined by the intended use of the product. In practice, most diagnostic technology (including ECG machines, stethoscopes, and similar) would be considered relatively low-risk devices (class I/II) compared with invasive, implantable, or explicitly life-sustaining technologies (class III). One implication of this risk tiering is that, unlike a new implanted cardiac device, such as a novel pacemaker or coronary stent, the market entry of diagnostic technology (including AI-ECGs) would not be predicated on having demonstrated their safety and effectiveness through, for example, a large trial with hard clinical endpoints.

Once a medical device receives regulatory authorization from the MHRA, the UK takes additional steps to determine whether and what the NHS should pay for it. The National Institute for Health and Clinical Excellence (NICE) evaluates the clinical efficacy and cost-effectiveness of drugs, health technologies, and clinical practices for the NHS. Rather than negotiating prices, NICE makes recommendations for system-wide funding and, therefore, deployment, principally based on

using tools such as quality-adjusted life years. In response to the increasing number and complexity of digital health technologies, NICE partnered with NHS England to develop standards that aim to ensure that new digital health technologies are clinically effective and offer economic value. The subsequent evidence standards framework for digital health technologies aims to inform stakeholders by exacting appropriate evidence, and to be dynamic and value-driven, with a focus on offering maximal value to patients.⁴

Considering the role of the regulatory bodies above, as applied to a novel AI-ECG device, we observe the following: Manufacturers seeking marketing authority for new digital health tools primarily focused on the diagnosis rather than treatment of a specific condition (like heart failure), must meet the safety and effectiveness standards of the MHRA – but those standards do not necessarily (or likely) require a dedicated clinical trial illustrating real-world clinical value. By contrast, convincing the NHS to pay for the new technology may require more comprehensive evidence sufficient to sway NICE, which is empowered to take a more holistic view of the costs and potential benefits of novel health tools. The advancement of this evidence generation for digital health tools is increasingly tasked to NHS sub-agencies. All of this aims to align with the NHS Long Term Plan, which defines the key challenges and sets an ambitious vision for the next ten years of health care in the UK.⁵ AI is singled-out as a key driver for digital transformation. Specifically, the “use of decision support and AI to help clinicians in applying best practice, eliminate unwarranted variation across the whole pathway of care, and support patients in managing their health and condition.” Here we already note implicit ethical principles: Reducing unjustified variability in care (as a consideration of justice) and promoting patient autonomy by disseminating diagnostic capabilities that otherwise may be accessible only behind layers of clinical or administrative gatekeeping. Focusing on the specific imperative of heart failure, this chapter discusses whether either of these or other ethical targets are, on balance, advanced by AI-ECG. To do this, we first outline the relevant clinical and technological background below.

III SCREENING FOR HEART FAILURE WITH AI-ECG

The symptomatic burden and mortality risks of heart failure – where the heart is no longer able to effectively pump blood to meet the body’s needs under normal pressures – remain worse than those of many common, serious cancers. Among all chronic conditions, heart failure has the greatest impact on quality of life and costs the NHS over £625 million per year – 4 percent of its annual budget.⁶ The NHS

⁴ National Institute for Health and Care Excellence, *Evidence Standards Framework for Digital Health Technologies* (2018), www.nice.org.uk/corporate/ecd7.

⁵ NHS England, *The NHS Long Term Plan* (2019), www.longtermplan.nhs.uk/.

⁶ Nathalie Conrad, et al., Temporal Trends and Patterns in Heart Failure Incidence: A Population-Based Study of 4 Million Individuals, 391 *The Lancet* 572, 572–80 (2018).

Long Term Plan emphasizes that “80% of heart failure is currently diagnosed in hospital, despite 40% of patients having symptoms that should have triggered an earlier assessment.” Subsequently, the Plan advocates for “using a proactive population health approach focused on ... earlier detection and intervention to treat undiagnosed disorders.”⁷ While the exact combination of data will vary by context, a clinical diagnosis of heart failure may include the integration of patients’ symptoms, physical exams (including traditional stethoscope auscultation of the heart and lungs), and various cardiac investigations, including blood tests and imaging. Individually, compared with a clinical diagnosis gold standard, the test characteristics of each modality vary widely, with sensitivity generally higher than specificity.

Similar to most chronic diseases in high-income countries, the burden of heart failure is greatest in those who are most deprived and tends to have an earlier age of onset in minority ethnic groups, who experience worse outcomes.⁸ Therefore, heart failure presents a particularly attractive target for disseminated technology with the potential to speed up diagnosis and direct patients toward proven therapies, particularly if this mitigates the social determinants of health driving observed disparities in care. Given the epidemiology of the problem and the imperative for practical screening, a tool supporting the community-based diagnosis of heart failure has the potential to be both clinically impactful and economically attractive. The myriad diagnostics applicable to heart failure described, however, variously require phlebotomy, specialty imaging, and clinical interpretation to tie together signs and symptoms into a clinical syndrome. AI-supported diagnosis may overcome these limitations.

The near ubiquity of ECGs in well-phenotyped cardiology cohorts supports the training and testing of AI algorithms among tens of thousands of patients. This has resulted in both clinical and, increasingly, consumer-facing applications where AI can interrogate ECGs and accurately identify the presence, for example, of heart rhythm disturbances. Building on an established background suggesting that the ECG can serve as an accurate digital biomarker for the stages of heart failure, a recent advance in AI has unlocked the super-human capability to detect heart failure from a single-lead ECG alone.⁹

The emergence of ECG-enabled stethoscopes, capable of recording single-lead ECGs during contact for routine auscultation (listening), highlighted an opportunity to apply AI-ECG to point-of-care screening. The Eko DUO (Eko Health, Oakland, CA, US) is one example of such an ECG-enabled stethoscope (see Figure 5.1). Detaching the tubing leaves a small cell phone-sized device embedded with sensors (electrodes and microphone) for recording both ECGs

⁷ NHS England, *supra* note 5.

⁸ Claire A Lawson, et al., Risk Factors for Heart Failure: 20-year Population-Based Trends by Sex, Socioeconomic Status, and Ethnicity, *13 Circulation: Heart Failure* (2020).

⁹ Patrik Bächtiger, et al., *supra* note 3.




FIGURE 5.1 Left to right: Eko DUO smart stethoscope; patient-facing “bell” of stethoscope labelled with sensors; data flow between Eko DUO, user’s smartphone, and cloud for the application of AI

and phonocardiograms (heart sounds). Connectivity via Bluetooth allows the subsequent live streaming of both ECG and phonocardiographic waveforms to a user’s smartphone and the corresponding Eko app. Waveforms can be recorded and transmitted to cloud-based infrastructure, allowing them to be analyzed by cloud-based AI algorithms, such as AI-ECG.

While the current programmatic focus is on identifying community heart failure diagnoses, AI can, in theory, also be applied to ECG and phonocardiographic waveforms to identify the presence of two additional public health priorities: Atrial fibrillation, a common irregular heart rhythm, and valvular heart disease, typified by the presence of heart murmurs. Therefore, taken in combination, a fifteen-second examination with an ECG-enabled smart stethoscope may offer a three-in-one screening test for substantial drivers of cardiovascular morbidity and mortality, and systemically important health care costs.

The authors are currently embarking on the first stage of deploying such a screening pathway, anchored in primary care, given the high rates of undiagnosed heart failure and further cardiovascular disease, including atrial fibrillation and valvular disease, in communities across England.¹⁰ The early stages of this pathway involve using NHS general practitioner electronic health records and applying search logic to identify those at risk for heart failure (e.g., risk factors such as hypertension, diabetes, previous myocardial infarction). Patients who consent are mailed a small parcel containing an ECG-enabled stethoscope (Eko DUO) and a simple instruction leaflet on how to perform and transmit a self-recording. Patients are encouraged to download the corresponding Eko App to their own phones (those who are unable to are sent a phone with the app preinstalled as part of the package). Patients whose data, as interpreted by AI, suggests the presence of heart failure, atrial fibrillation, or valvular heart disease are invited for further investigation in line with established NICE clinical pathways.

¹⁰ Michael Soljak, et al., Variations in Cardiovascular Disease Under-Diagnosis in England: National Cross-Sectional Spatial Analysis, 11 *BMC Cardiovascular Disorders* 1, 1–12 (2011).

This sets the scene for a novel population health intervention that draws on a technology-driven screening test, initiated in the patient's home, by the patient themselves. The current, hospital-centric approach to common and costly cardiovascular conditions combines clinical expertise and the available technologies to screen and unlock substantial clinical and health economic benefits through early diagnosis. Opportunities for more decentralized (outside of hospital), patient-activated screening with digital diagnostics will surely follow if AI-ECG proves tractable. Notably, here we have described what we believe to be among the earliest applications of "super-human" AI – accurately inferring the presence of heart failure from a single-lead ECG was previously thought impossible – with the potential for meeting a major unmet need through a clinical pathway that scales access to this potentially transformative diagnostic.

IV ETHICAL CONSIDERATIONS FOR SELF-ADMINISTERED CARDIOVASCULAR DISEASE SCREENING AT HOME

Having outlined the health policy and stakeholder landscape and specified how this relates to heart failure and AI-ECG, we can progress to discussing the unique ethical challenges posed by patient self-administration of this test in their own homes. Enthusiasm for such an approach to community, patient-driven cardiovascular screening is founded in not only clinical expediency, but also a recognition of the way in which this pathway may support normative public health goals, particularly around equity and patient empowerment. Despite these good-faith expectations, the deployment of such a home-based screening program combining hardware, AI, and a cloud-based digital platform for administration – all hinging on patient self-administration – raises distinct ethical challenges. In this section, we explore the ethical arguments in favor of the AI-ECG program, as well its potential pitfalls.

A Equity

One durable and compelling argument supporting AI-ECG arises from well-known disparities in cardiovascular disease and treatment. Cardiovascular disease follows a social gradient; this is particularly pronounced for heart failure diagnoses, where under-diagnosis in England is most frequent in the lowest-income areas. This tracks with language skills, a key social determinant of health related to a lower uptake of preventative health care and subsequently worse health outcomes. In England, nearly one million people (2 percent of the total population) lack basic English language skills. AI-ECG attenuates these disparities in several ways.

First, targeted screening based on risk factors (such as high blood pressure and diabetes) will, based on epidemiologic trends, necessarily and fruitfully support vulnerable patient groups for whom these conditions are more prevalent. These

same patients will also be less able to access traditional facility-based cardiac testing. AI-ECG overcomes these concerns for the patients most in need.

Second, AI-ECG explicitly transfers a key gatekeeping diagnostic screen away from clinicians: The cognitive biases of traditional bedside medicine. Cross-cultural challenges in subjective diagnosis and treatment escalation are well documented, including in heart failure across a spectrum of disease severity, ranging from outpatient symptoms ascertainment to referral for advanced cardiac therapies and even transplant.¹¹ AI-ECG overcomes the biases embedded in traditional heart failure screenings by simplifying a complex syndromic diagnosis into a positive or negative result that is programmatically entwined with subsequent specialist referral.

These supporting arguments grounded in reducing the disparities in access to cardiac care may be balanced by equally salient concerns. Even a charitable interpretation of the AI-ECG pathway assumes a relatively savvy, engaged, and motivated patient. The ability to mail the AI-ECG screening package widely to homes is just the first step in a series of necessary steps: Opening and setting up the screening kit, including the phone and ECG-enabled stethoscope, successfully activating the device, and recording a high-quality tracing that is then processed centrally without data loss. While the authors' early experience using this technology in various settings has been reassuring, it remains uncertain whether the established "digital divide" will complicate the equitable application of AI-ECG screening. Assuming equal (or even favorably targeted) access to the technology, are patients able to use it, and *do they want to?* The last point is critical: In the UK as well as the United States, trust in health care varies considerably and, (broadly speaking) in cardiovascular disease, tracks unfortunately and inversely with clinical need.

Indeed, one well-grounded reason for suspicion recalls another problem for the equity-driven enthusiasm for AI-ECG, which is the training and validation of the AI algorithms themselves. The "black box" nature of some forms of AI, where the reasons for model prediction cannot easily be inferred, has appropriately led to concerns over insidious algorithmic bias and subsequent reservations around deploying these tools for patient care.¹² Even low-tech heart failure screening confronts this same problem, as (for example) the most widely used biomarker for heart failure diagnosis has well-known performance variability according to age, sex, ethnicity, patient weight, renal function, and clinical comorbidities.¹³ Conversely, studies to date have suggested that AI-ECG for heart failure detection does not exhibit these

¹¹ Fouad Chouairi, et al., Evaluation of Racial and Ethnic Disparities in Cardiac Transplantation, 10 J. of the Am. Heart Ass'n (2021).

¹² Matthew DeCamp & Jon C. Tilburt, Why We Cannot Trust Artificial Intelligence in Medicine, 1 Lancet Digit. Health 390 (2019).

¹³ Theresa A. McDonagh, et al., 2021 ESC Guidelines for the Diagnosis and Treatment of Acute and Chronic Heart Failure: Developed by the Task Force for the Diagnosis and Treatment of Acute and Chronic Heart Failure of the European Society of Cardiology (ESC) With the Special Contribution of the Heart Failure Association (HFA) of the ESC, 42 European Heart J. 3599, 3618 (2021).

biases. It may still be the case that biases do exist, but that they require further large-scale deployment to manifest themselves.

To address these concerns, we propose several programmatic features as essential and intentional for reinforcing the potential of wide-scale screening to promote equity. First, it is imperative for program managers to prominently collect self-identified race, ethnicity, and other socioeconomic data (e.g., language, education) from all participants at each level of outreach – screened, invited, agreed, successfully tested, identified as “positive,” referred for specialist evaluation, and downstream clinical results. Disproportionate representation at each level, and differential drop-out at each step, must be explored, but that can only begin with high-quality patient-level data to inform analyses and program refinement. This is an aspiration dependent on first resolving the outlined issues with trust. Trust in AI-ECG may be further buttressed in several ways, recognizing the resource limitations available for screening programs generally. One option may be providing accommodations for skeptical patients in a way that still provides suitable opportunities to participate through alternative means. This could simply involve having patients attend an in-person appointment during which the AI-ECG examination is performed on them by a health care professional.

The patient end-user needs to feel trust and confidence in using the technology. This can be achieved through user-centric design that prioritizes a simple protocol, to maximize uptake, with the requisite level of technical detail to ensure adequate recording quality (e.g., getting the right position). The accuracy of AI-ECG depends on these factors, in contrast with other point-of-care technologies where the acquisition of the “input” is less subject to variability (e.g., finger-prick blood drop tests).

The centralized administration of NHS screening programs by NHS England paired with NHS Digital’s repository data on the uptake of screening offers granular insights to anticipate and plan for regions and groups at risk of low uptake. We propose enshrining a dedicated data monitoring plan into the AI-ECG screening protocol, with prespecified targets for uptake and defined mitigation strategies – monitored in near real-time. This is made possible through the unique connectivity (for a screening technology) of the platform driving AI-ECG, with readily available up-to-date data flows for highlighting disparities in access. However, a more proactive approach to targeting individuals within a population with certain characteristics needs to be balanced against the risk of stigmatization, and, ultimately, potential loss of trust that may further worsen the cardiovascular outcomes seeking to be improved.

Lastly, equity concerns around algorithmic performance are necessarily empirical questions that will also benefit from patient-level data collection. We acknowledge that moving from research in the form of prospective validation studies to deployment for patient care requires judgment in the absence of consensus, within the NHS or more globally, around the minimum scrutiny for an acceptable level

(if any) of differential performance across – for starters – age, sex, and ethnicity. To avoid these potentially impactful innovations remaining in the domain of research, and to anticipate the wide-reaching implications of a deployment found to exhibit bias retrospectively, one possible solution would be to, by design, prospectively monitor for inconsistent test performance. Specifically, in the context of AI-ECG offering a binary yes or no screening test result for heart failure, it is important to measure the rate of false positive and false negative results. False positives can be measured through the AI-ECG technology platform linking directly into primary care EHR data. This allows positive AI-ECG results to be correlated with the outcomes of downstream gold-standard, definitive investigations for heart failure (e.g., echocardiography ultrasound scans). Such a prospective approach is less feasible for false negatives due to both the potentially longer time horizon for the disease to manifest and the uncertainty around whether AI-ECG truly missed the diagnosis. Instead, measuring the rate of false negatives may require a more expansive approach in the form of inviting a small sample of patients with negative AI screening tests for “quality control” next-step investigations. All of this risks adding complexity and, therefore, cost to a pathway seeking to simplify and save money. However, given this program’s position at the vanguard of AI deployments for health, a permissive approach balanced with checkpoints for sustained accuracy may help to blueprint best practices and build confidence for similar AI applications in additional disease areas.

B Agency

Another positive argument for AI-ECG screening aligns with trends in promoting agency, understood here as patient empowerment, particularly around the use of digital devices to measure, monitor, and manage one’s own health care – particularly in terms of cardiovascular disease. The enthusiastic commercial uptake of fitness wearables, for example, moved quickly past counting steps to incorporate heart rhythm monitoring.¹⁴ Testing of these distributed technologies has shown mixed results, with the yield of positive cases necessarily depending on the population at issue.¹⁵ Recalling the equity concerns above, the devices themselves may be more popular among younger and healthier patients, among whom true positive diagnoses may be uncommon. However, targeted and invited screening with AI-ECG may balance these concerns through enriching the population at risk by invitation.

¹⁴ David Duncker, et al., *Smart Wearables for Cardiac Monitoring – Real-World Use beyond Atrial Fibrillation*, 21 *Sensors* (2021).

¹⁵ Steven A Lubitz, et al., *Screening for Atrial Fibrillation in Older Adults at Primary Care Visits: VITAL-AF Randomized Controlled Trial*, 145 *Circulation* 946–54 (2022); Marco V Perez, et al., *Large-Scale Assessment of a Smartwatch to Identify Atrial Fibrillation*, 381 *New England Journal of Medicine* 1909–17 (2019).

Realistic concerns about agency extend beyond the previous warnings about digital literacy, access to reliable internet, and language barriers to ask more fundamental questions about whether patients actually *want* to assume this central role in their own health care. A key parallel here is the advent of mandates for shared decision-making in cardiovascular disease, particularly in the United States where federal law now requires selected Medicare beneficiaries considering certain cardiovascular procedures to incorporate “evidence based shared decision-making tools” in their treatment choices.¹⁶ However, patients may reasonably ask if screening with AI-ECG should necessarily shift the key role of test administration (literally) into their hands. Unlike the only other at-home national screening test in the UK – simply taking a stool sample for bowel cancer screening – self-application of AI-ECG requires the successful execution of several codependent steps. Here, even a relatively low failure rate may prove untenable for population-wide scaling, risking that this technology may remain in the physician’s office.

Putting such responsibility on patients could be argued to not only directly shift this responsibility away from clinicians, but also dilute learning opportunities. While subtle, shifting the cognitive work of integrating complex signs and symptoms into a syndromic diagnosis like heart failure may have unwelcome implications for clinicians’ diagnostic skills. We emphasize that this is not just whimsical nostalgia for a more paternalistic time in medicine, but a genuine worry about reductionism in algorithmic diagnosis that oversimplifies complex constellations of findings into simple yes or no diagnoses (AI-ECG, strictly speaking, only flags a risk of heart failure, which is not clinically equivalent to a diagnosis). Resolving these tensions may be possible through seeing the educational opportunity and wider clinical application of the hardware enabling AI-ECG.

Careful metrics, as described previously, will allow concerns about agency to be considered empirically, at least within the categories of patient data collected. If, for example, the utilization of AI-ECG varies sharply according to age, race, ethnicity, or language fluency, this would merit investigation specifically interrogating whether this variability rests in part on patient preferences for taking on this task rather than an inability to do so. At the same time, early patient experiences with AI-ECG in real-world settings may provide opportunities for patient feedback regarding whether this specific device, or the larger role being asked of them in their own care, is perceived as an appropriate assignment of responsibility or an imposition. If, for example, patients experience this shifting of cardiovascular screening out of the office as an inappropriate deferral of care out of traditional settings, this may suggest the need for either refining the pathway (still using the device, but perhaps keeping it in a clinical setting) or more extensive community engagement and education to ensure stakeholder agreement on roles, rights, and responsibilities.

¹⁶ Christopher E Knoepke, et al., Medicare Mandates for Shared Decision Making in Cardiovascular Device Placement, 12 *Circulation: Cardiovascular Quality and Outcomes* (2019).

C Data Rights

A central government, NHS-funded public screening program making use of patients' own smartphones necessarily raises important questions about data rights. Beyond the expected guardrails required by the General Data Protection Regulation (GDPR) and UK-specific health data legislation, AI-ECG introduces additional concerns. One is whether patient participants should be obligated to contribute their health data toward the continuous refinement of the AI-ECG algorithms themselves or instead be given opt-in or opt-out mechanisms of enrollment. We note that while employed in this context by a public agency, the intellectual property for AI-ECG is held by the device manufacturer. Thus, while patients may carry some expectation of potential future benefit from algorithmic refinement, the more obvious rewards accrue to private entities. Another potential opportunity, not lost on the authors as overseers for the nascent AI-ECG program, is the possibility that AI-ECG data linked to patients' EHR records might support entirely new diagnostic discovery beyond the core cardiovascular conditions at issue. Other conditions may similarly have subtle manifestations in ECG waveforms, phonocardiography, or their combination – invisible to humans but not AI – that could plausibly emerge from widespread use. Beyond just opt-in or opt-out permissions – known to be problematic for meaningful engagement with patient consent¹⁷ – what control should patients have around the use of their health data in this context? For example, the NHS now holds a rich variety of health data for each patient – including free text, imaging, and blood test results. Patients may be happy to offer some but not all of this data for application to their own health, with different decisions on stratifying what can be used for AI product development.

Lastly, AI-ECG will need to consider data security carefully, including the possibility, however remote, of malicious intent or motivated intruders entering the system. Health data can be monetized by cyber criminals. Cyber threat modelling should be performed by the device manufacturer early in the design phase to identify possible threats and their mitigants.¹⁸ Documentation provided about embedded data security features adds valuable information for patients that may have concerns about the protection of their personal data, and can help them to make informed decisions on using AI-ECG. Beyond privacy, threat modelling should also account for patient safety, such as from an intruder with access that allows the manipulation of code or data. For example, it could be possible to manipulate results to deprive selected populations of appropriate referrals for care. Sabotaging results or causing a denial-of-service situation by flooding the system with incorrect data might also

¹⁷ Susan A Speer & Elizabeth Stokoe, Ethics in Action: Consent-Gaining Interactions and Implications for Research Practice, 53 *British J. of Soc. Psych.* 54–73 (2014).

¹⁸ Medical Device Innovation Consortium, The MITRE Corporation, *Playbook for Threat Modelling Medical Devices* (2021), www.mitre.org/news-insights/publication/playbook-threat-modeling-medical-devices.

cause damage to the reputation of the system in such a way that patients and clinicians become wary of using it. Overall, anticipating these security and other data rights considerations beyond the relatively superficial means of user agreements remains an unmet challenge for AI-ECG.

V FINAL RECOMMENDATIONS

This chapter has outlined a novel clinical pathway to screen for cardiovascular disease using an at-home, patient self-administered AI technology that can provide a screening capability beyond human expertise. We set this against a backdrop of: (1) A diverse ecosystem of stakeholders impacted by and responsible for AI-ECG, spanning patients, NHS clinicians, NHS agencies, and the responsible regulatory and health economic bodies and (2) a health-policy landscape eager to progress the “use of decision support and AI” as part of a wider push to decentralize (i.e., modernize) care. To underscore the outlined considerations of equity, agency, and data rights, we propose two principal recommendations, framed against but generalizable beyond the pathway example of AI-ECG.

First, we advocate for a multi-agency approach that balances permissive regulation and deployment – to align with the speed of AI innovation – against ethical and statutory obligations to safeguard public health. Bodies such as NHS England, the MHRA, and NICE each have unique responsibilities, but with cross-cutting implications. The clinical and health economic case for urgent innovation for unmet needs, such as AI-ECG for heart failure, is obvious and compelling. Agencies working sequentially delays translating such innovations into clinical practice, missing opportunities to avert substantial cardiovascular morbidity and mortality. Instead, the identification of a potentially transformative technology should trigger a multi-agency approach that works together and in parallel to support timely deployment within clinical pathways to positively impact patient care. This approach holds not only during initial deployment, but also as technology progresses. Here, we could consider the challenge of AI algorithms continually iterating (i.e., improving): For a given version of AI-ECG, the MHRA grants regulatory approval, NICE endorses procurement, and NHS England guides implementation. After evaluating a medical AI technology and deeming it safe and effective, should these agencies limit its authorization to market only the version of the algorithm that was submitted, or permit the marketing of an algorithm that can learn and adapt to new conditions?¹⁹ AI-ECG could continually iterate by learning from the ECG data accumulated during deployment, and also through continuing improvements in machine learning methodology and computational power. Cardiovascular data, including waveforms, imaging, blood, and physiological parameters, is generally high volume and repeatedly measured. This, therefore, offers a rich seam for taking advantage of

¹⁹ Boris Babic, et al., Algorithms on Regulatory Lockdown in Medicine, 6 *Science* 1202 (2019).

AI's defining strength to continually improve, unlike ordinary "medical devices." Parodying the ship of Theseus, at what point is the algorithm substantially different to the original, and what prospective validation, if any, is needed if the claims remain the same? Multi-agency collaboration can reach a consensus on such questions that avoids unfamiliarity with the lifecycle of AI disrupting delivery of care by reactively resetting when new (i.e., improved) versions arrive. For AI-ECG, this could involve the expensive and time-consuming repetition of high-volume patient recruitment to validation research studies. Encouragingly, in a potential move toward multi-agency collaboration, in 2022, NHS England commissioned NICE to lead a consultation for a digital health evidence standards framework that aims to better align with regulators.²⁰

Second, both to account for the ethical considerations outlined in this chapter and to balance any faster implementation of promising AI technologies, we recommend a centralized responsibility for NHS England to deploy and thoroughly evaluate programs such as AI-ECG. This chapter has covered some of the critical variables to measure that will be unique to using an AI technology for patient self-administered screening at home. Forming a comprehensive list would, again, be amenable to a multi-agency approach, where NHS England can draw on the playbook for already-monitoring existing national screening programs. An evaluation framework addressing the outlined considerations around equity, agency, and data rights should be considered not only an intrinsic but a mandatory part of the design, deployment, and ongoing surveillance of AI-ECG. The inherent connectivity and instant data flow of such technology offers, unlike screening programs to date, the opportunity for real-time monitoring and, therefore, prompt intervention, not only for clinical indications, but also for any disparities in uptake, execution, algorithm performance, or cybersecurity. Ultimately, this will not only bolster the NHS's position as a world leader in standards for patient safety, but also as an exemplar system for realizing effective AI-driven health care interventions.

Looking to the future for AI-ECG, translating the momentum for technological innovation in the NHS into patient benefit will require careful consideration of the outlined ethical pitfalls. This may, in the short term, establish best practices that build confidence for further applications. In the longer term, we see a convergence of commoditized AI algorithms for cardiovascular and wider disease, where increasingly sophisticated sensor technology may make future home-based screening a completely passive act. While moving toward such a reality could unlock major public health benefits, doing so will depend on bold early use cases, such as AI-ECG, that reveal unanticipated ethical challenges and allow them to be resolved. For

²⁰ National Institute for Health and Care Excellence, *Evidence Standards Framework (ESF) for Digital Health Technologies Update – Consultation* (2022), www.nice.org.uk/about/what-we-do/our-programmes/evidence-standards-framework-for-digital-health-technologies/esf-consultation.

now, the outlined policy recommendations can serve to underpin the stewardship of such novel diagnostic pathways in a way that preserves and promotes trust, patient engagement, and public health.

VI CONCLUSION

Patient self-administered screening for cardiovascular disease at home using an AI-powered technology offers substantial potential public health benefits, but also poses unique ethical challenges. We recommend a multi-agency approach to the lifecycle of implementing such AI technology, combined with a centrally overseen, mandatory prospective evaluation framework that monitors for equity, agency, and data rights. Assuming the responsibility to proactively address any observed neglect of these considerations instills trust as the foundation for the sustainable and impactful implementation of AI technologies for clinical application within patients' own homes.