AI FOR HEALTH: CONVERTING MOMENTUM INTO MUSCLE

HEAL Advisory Thinktank

February 2026


NTRE FOR HEALTHCARE INNOVATION


MINISTRY OF HEALTH
OVATE

Contents

Foreword

Artificial intelligence (AI) is reshaping many aspects of our life, and healthcare is no exception. We are already seeing how AI can support clinicians, enhance operations, and unlock new insights from data to improve care. But as with every technological revolution, its impact will depend not only on what the technology can do, but on how it is applied, governed and integrated into real clinical practice.

AI, as its name implies, can never replace the human touch. Healthcare is, and must remain, a profoundly human endeavour. At its core is trust: between patient and clinician, between science and society, between government and people. AI can strengthen that trust when it is deployed responsibly, transparently, and with purpose. But if we pursue novelty for its own sake, or treat AI as a substitute rather than an aid to human judgement, we risk undermining the very foundations of care.

Singapore's ambition is not to build the most algorithms, but to become the place where AI in health is implemented well – safely, equitably, and at scale. What matters is not speed alone, but reliability, accountability, and real-world impact. Achieving this requires close collaboration across our public healthcare institutions, innovators, and global partners. It also requires leadership: the courage to experiment, the humility to learn from failure, and the resolve to put in place proportionate safeguards as capabilities grow.

This report, *AI for Health: Converting Momentum into Muscle*, reflects that spirit. It moves beyond aspiration to execution, showing that the hard work of innovation lies in building the muscle of adoption, evaluation, and continual learning. It reminds us that implementation is not the end of innovation; it is where innovation proves its value.

As health systems around the world contend with ageing populations and workforce pressures, AI offers an important means to augment human capability, not replace it. Its success, however, will depend on how we collectively design, deploy, and derive value from it. The lessons captured here – from the HEAL Advisory Thinktank and partners in Singapore and abroad – are therefore not only technical. They are lessons in leadership, trust, and shared purpose. May we be encouraged to innovate with purpose and care, and to ensure that technology serves patients, supports professionals and strengthens the healthcare system as a whole.

Mr Ong Ye Kung
Minister for Health and Coordinating Minister for Social Policies
Republic of Singapore

Introduction

By NHG Health
Group CEO

The integration of AI into health systems is unfolding against a backdrop of deep and intersecting shifts – demographic, technological, clinical, and societal. These shifts are not new, but they are intensifying. To understand the promise and limits of AI in this moment, we must situate it within these broader currents.

Globally, populations are ageing and living longer, though not always healthier, lives. The burden of chronic disease is rising. Health inequities persist. Systems remain under strain, with workforce shortages, increasing demand, and financial pressure. At the same time, care itself is changing – moving beyond hospital walls and embracing prevention, personalisation, and digital technology.

Against this backdrop, AI has rapidly moved from the lab into mainstream conversation, fuelled by advances in machine learning, natural language processing, and generative models. It now touches everything from imaging to administration, from population risk stratification to personal health coaching. Yet despite a proliferation of pilots and proof-of-concepts, relatively few AI tools have achieved sustained, system-wide adoption.

This report is motivated by that observation. AI holds immense potential to reshape how we deliver, organise, and experience health and care. But we are still early in understanding how to realise that potential in practice. The gap between promise and impact remains wide.

Our aim is to narrow that gap – not by simply describing what AI can do, but by examining how systems can adopt it responsibly, equitably, and at scale. This report brings together the reflections of health and technology leaders who are navigating these challenges daily. It is not a playbook, but a set of pragmatic insights and shared principles to support progress.

Joe Sim

Prof Joe Sim
Group Chief Executive
NHG Health

Thinktank participants gathered in Singapore during the week of CHI INNOVATE 2025.

Attribution Note

AI is advancing with extraordinary speed, and health systems everywhere are under pressure to show how they will keep pace. It is tempting to respond with a roadmap – a sequence of steps that promise predictable progress. But AI does not behave like a linear programme, and health systems do not transform through instruction manuals. Contexts differ, digital maturity varies, and the work of adoption is shaped as much by culture, capability, and trust as by technology. This report therefore does not offer a recipe. Instead, it sets out the key ingredients, conditions, and leadership practices that make responsible, scalable adoption possible – whatever the starting point.

The Health Empowered by AI Launchpad (HEAL) Advisory Thinktank (HAT) was commissioned by the CHI Leadership Council to provide strategic insight into how health systems can move beyond experimentation towards responsible, system-wide adoption of AI. Convened alongside the CHI Innovate Conference in July 2025, HAT was designed as a focused, reflective forum rather than a technical workshop. It brought together more than 60 senior health system leaders, clinicians, academics, and technology experts with direct experience of AI implementation to surface shared challenges, practical lessons, and leadership considerations.

This report synthesises the learning from those discussions, capturing the principles, practices, and provocations that emerged from a collaborative inquiry into how health systems can translate AI potential into sustained, ethically grounded impact at scale.

Executive Summary

What this report is (and who it is for)

This report synthesises learning from more than 60 leaders across health systems, academia, and technology partners who are navigating these challenges in real time, brought together through HAT and the CHI Innovate 2025 Conference ('AI for All, AI for Change').

It is written primarily for system leaders – clinical, operational, digital, and policy leaders – who are not asking "What can AI do?" but "How do we make this work as system change?" It is deliberately pragmatic: not a playbook, but a set of shared patterns, paradoxes, and principles that can help leaders move from scattered experimentation to responsible adoption at scale.

To ground the discussion in operational reality, the report also includes tangible examples where AI is already making work lighter and care more reliable. For instance, ambient documentation tools can reduce documentation time substantially in suitable sessions and support more present, person-centred interactions. Likewise, "second set of eyes" radiology tools in primary care illustrate how AI can augment judgement and prioritise risk without displacing clinical responsibility. And several LLM-enabled service tools show how impact can range from micro-efficiencies to population-scale access – including a digital front door designed to reduce enquiry workload at scale.

The current context: extraordinary capability, genuine promise

Health systems around the world are facing a confluence of pressure and possibility. Ageing populations, rising chronic disease, and workforce shortages are colliding with rapid advances in artificial intelligence – particularly generative and agentic models – that can meaningfully reshape how care is delivered, organised, and experienced.

The promise is real, and it is wider than any single use case. AI can augment clinical judgement by bringing signal to the noise, improving consistency, and supporting safer decisions. It can reduce operational drag by taking on the "background work" of healthcare – documentation, coordination, forecasting, triage support, and information retrieval – so staff can focus more attention on the person in front of them. And at system level, it can strengthen prevention and planning by anticipating demand, segmenting need, and supporting more proactive, earlier intervention.

Crucially, the opportunity is not simply automation. It is augmentation: enhancing human judgement, reducing cognitive burden, and redesigning workflows so care becomes more reliable, more personalised, and more humane under strain. This is why AI feels like an inflection point: it reaches into the cognitive and relational fabric of healthcare – how knowledge is used, decisions are made, and interactions occur.

The complication: why momentum is not becoming system change

And yet, despite the hype and the surge of activity, a common experience is emerging internationally: most health systems remain caught between possibility and practice. Across countries and contexts, the pattern is recognisable: pilots proliferate, pockets of excellence emerge, but wider adoption and sustained impact remain stubbornly difficult.

To respond to that gap, NHG Health convened the HEAL Advisory Thinktank (HAT) alongside the CHI Innovate Conference, bringing together system leaders, clinicians, researchers, and technologists with direct experience in deploying and adopting AI in real-world care. Participants came with shared optimism about the opportunity, but also a shared determination to move beyond abstract debate towards practical insight: what does it actually take to adopt AI responsibly and at scale, and what shifts the dial from local innovation to system transformation?

The constraint is rarely a shortage of tools or ingenuity. It is the system work of adoption: making AI reliable in day-to-day practice, integrating it into workflows, governing it proportionately, evaluating it in the real world, and sustaining performance as context shifts. Put simply, the potential of AI is constrained less by innovation than by implementation, less by ideas than by execution.

AI raises the stakes in a way that feels different from earlier digital waves. It is more accessible, more iterative, and more entwined with judgement, knowledge work, and the lived experience of care. That means the familiar challenges of adoption still apply, but our theory of change has to evolve – this is as much about adapting the system as adopting the technology: moving from project delivery to capability-building, from one-off assurance to continuous oversight, and from linear change programmes to learning systems that can improve safely over time.

This isn't a story of failure. It is what it looks like when a fast-moving technology hits systems built for stability: procurement models designed for fixed products, governance designed for predictability, data environments designed for documentation, and change models designed for linear processes.

The practical question – for leaders everywhere – is what has to shift if we want AI's benefits to become real, reliable, and widespread.

Key findings: what shifts the dial from activity to impact

Across HAT and Innovate, a small set of practical lessons surfaced repeatedly. They are not slogans. They are the recurring design requirements of responsible adoption at scale that enable system leaders to become more intentional in their approach to AI adoption.

1 Start with purpose, not possibility.

The systems making progress are those that anchor AI in clear priorities, service needs, and outcomes that matter, rather than a growing catalogue of what is technically feasible. Focus is not a constraint on innovation. It is what allows learning, investment, and confidence to compound.

2 Treat adoption as capability-building, not delivery.

Every deployment should leave the system stronger: clearer methods, better data foundations, more confident teams, and more reusable pathways. Adoption is not a single "go-live". It is the system learning how to learn, repeatedly, under real conditions.

3 Replace accumulation with orchestration.

A growing portfolio is not a strategy. Without a scaling route, pilots turn into noise. Systems need a deliberate way to select, sequence, resource, and stop work, so that effort becomes cumulative rather than scattered.

4 Build the spine that makes scale possible, and practical.

Interoperability, data quality, deployment pipelines, monitoring, and secure environments are not technical nice-to-haves. They are what make responsible scaling feasible and reduce the repetitive effort that keeps teams stuck in bespoke builds.

5 Make governance a learning discipline, calibrated to risk.