AI FOR HEALTH: CONVERTING MOMENTUM INTO MUSCLE

HEAL Advisory Thinktank

February 2026

Leadership Council Report


OVATE
OVATE

Contents

Foreword

By Minister for Health and Coordinating Minister for Social Policies

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.

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.
Responsible adoption depends on proportionate, clinically anchored oversight that protects trust without stifling progress. The goal is not to eliminate risk. It is to manage it intelligently, with transparency, monitoring, and the ability to adapt.

6 Grow confidence across the workforce, not only capability in specialist teams.
AI fluency cannot sit only with enthusiasts. Health systems need practical literacy, clear guardrails, psychological safety, and visible support so that staff can be curious without being careless, and cautious without being stuck.

7 Build evaluation into deployment, not after it.
If learning is optional, it will be crowded out. Systems need embedded evaluation and "signals" that inform decisions in real time: what is working, for whom, under what conditions, and with what trade-offs.

8 Measure what matters, not only what is easy.
Leaders should track not just ROI and throughput, but also staff experience, cognitive burden, time to care, trust, safety signals, and whether value is being fairly distributed across teams and populations.

9 Partnerships must support stewardship, not dependency.
As AI becomes more embedded, relationships with vendors, academia, and regulators need to evolve beyond procurement towards shared accountability for outcomes, transparency, and continuous improvement.

10 The critical move is coherence: collaboration without uniformity.
Scaling does not always mean copying a single model everywhere. It means enabling local adaptation within shared guardrails, so that progress remains safe, comparable, and learnable across the system.

The leadership heart of the report: Turning momentum into muscle

Taken together, these findings point to a leadership challenge as much as a technical one. The work is to convert momentum into muscle: building the strategic capability to adopt AI responsibly and repeatedly, at scale, over time.

That places leadership at the centre - not as the "owner" of every initiative, but as the conductor: shaping alignment around what matters, building reusable foundations, governing proportionately, and making learning cumulative rather than accidental.

That requires an operating stance suited to this era: ambidexterity. Leaders must orchestrate guardrails, platforms, standards, evaluation, and safety, while enabling participation through frontline experimentation, local adaptation, ownership, and faster learning. When this balance is struck, AI adoption becomes less like a sequence of disconnected projects and more like a compounding system learning cycle: each deployment strengthens the system's readiness for the next.

This is also where the report's broader proposition comes into view. Many of these ideas, such as learning systems, adaptive governance, and relational approaches to change, are not new. What feels new is the impetus. AI's pace and accessibility have moved these principles from theory into operational necessity. Health systems are being pulled towards a model in which intelligence, learning, and decision support are more distributed across the workforce and care settings, rather than concentrated at the centre.

In this emerging era of distributed intelligence, leadership is less about control and more about intentionality: setting direction, building enabling foundations, creating the conditions for learning, and holding the paradoxes of speed and safety, innovation and assurance, ambition and trust. The technologies may be new, but the central task is familiar and enduring: building systems that can learn their way to better care.

The opportunity ahead is therefore substantial: to relieve strain, improve reliability, and create space for the relational work of care – so that humanity is not the casualty of modernisation, but its outcome.

Chapter 1 - From Possibility to Purpose

Strategic Context and The Promise of AI

Health systems around the world are facing a confluence of pressure and possibility. Ageing populations, the rising burden of chronic disease, and workforce shortages are converging with new digital capabilities that promise to fundamentally change how care is delivered. The arrival of artificial intelligence (AI), particularly generative and agentic models, has added a new layer of both optimism and urgency.

Across health organisations, there is an abundance of ideas, energy, and ambition – driven by the scale of the challenges systems face. Yet much of this activity is not yet translating into the kind of fundamental, transformational change required for sustained impact. There are pilots everywhere, but few pathways to scale; numerous breakthroughs, but bottlenecks in adoption. The result is progress that is visible in pockets rather than pervasive across systems.

We are witnessing a familiar pattern: a flurry of experimentation without commensurate transformation. As George Westerman of MIT cautions, ***"When digital transformation is done right, it's like a caterpillar turning into a butterfly. But when done wrong, all you have is a really fast caterpillar."*** In today's health systems, leaders are discovering that to meet rising demand, rising costs, and rising expectations, speed alone is not enough. Transformation requires changing form, not merely accelerating the status quo.

Against this backdrop of opportunity and challenge, a new wave of generative AI has catalysed fresh energy across health systems – marking a pivotal moment, not just in technology, but in how people imagine and engage in the process of care. This is why AI feels like an inflection point – because it reaches into the cognitive and relational fabric of healthcare: how knowledge is used, decisions are made, and interactions occur. The opportunity is not simply to automate tasks, but to redesign workflows, personalise prevention and care, and restore capacity to systems under strain.

What the Promise Looks Like - Beyond Hype, Towards Augmentation

Across the world, health systems are experiencing a moment of intense acceleration. Artificial intelligence (AI) has moved from the margins of digital strategy to the centre of every conversation. New tools appear almost daily. Teams that once waited for central pilots are now experimenting on their own. For many leaders, the challenge is no longer convincing people to try AI – it is channelling that enthusiasm into meaningful, measurable transformation.

The excitement is justified. AI can help us work smarter, see patterns faster, and ease the administrative and cognitive burdens that sap time and energy from care. But this new wave also exposes a familiar pattern: early breakthroughs followed by uneven adoption. Many health systems now find themselves with pockets of excellence rather than a coherent ecosystem. The potential is extraordinary, but the reality often stops short of transformation.

These challenges are not new. The barriers are rarely technological. They are systemic: the result of fragmented intent, uneven data foundations, and governance models still adapting to the pace of change. This chapter explores how the promise of AI can be turned from early potential into sustained, system-wide impact.

How We Got Here – From Early Promise to Partial Realisation

AI in healthcare did not begin with ChatGPT. The story stretches back decades – from expert systems that could diagnose based on rule sets, to statistical risk calculators and imaging algorithms trained on vast datasets. Each generation offered glimpses of the future. Yet most remained trapped in pilots, constrained by integration challenges, regulatory complexity, and the characteristic cautiousness of clinical change.

What changed with generative AI was accessibility. Suddenly, anyone with a digital device could test, iterate, and experience the power of these systems firsthand. Large language models (LLMs) and AI assistants have made AI tangible. Their fluency with language – the universal medium of healthcare – created an inflection point. Conversations that once required data scientists are now happening in corridors, clinics, and call rooms. Curiosity became capability. As one of the Innovate speakers put it,

"Generative AI has brought innovation into everyone's hands."

This has brought enormous energy but also risk. Productivity gains are widely self-reported, yet impact is uneven and often hard to verify. Enthusiasm can outpace understanding. What NHG Health and the HEAL Advisory Thinktank (HAT) have found is that this moment calls not for restraint, but for intentional direction – ensuring experimentation builds organisational muscle, not fragmentation.

AI's potential spans clinical care, operations, and population health. In clinical settings, it can help surface risk earlier, reduce diagnostic delay, and improve safety through detecting subtle signs of risk earlier and enhancing decision support. In operational settings, it can reduce administrative load, increase reliability, and support smoother flow through forecasting, triage support, rostering optimisation, and intelligent coordination. At system level, it can strengthen

prevention and planning by helping anticipate population needs, forecast pressure on services, and support more proactive and coordinated approaches to care.

Spotlight: The Generative Moment – A Turning Point in Adoption

One of the strongest themes to emerge from discussions across the HEAL Advisory Thinktank (HAT) and the CHI Innovate conference was the sheer pace and reach of the new wave of generative AI. The release of large language models (LLMs) has accelerated interest in artificial intelligence to levels unseen in previous decades. What began as a technical breakthrough has rapidly become a cultural and organisational phenomenon, transforming how people think about work, creativity, and knowledge itself.

This is not simply another phase of digital innovation. The accessibility of generative tools – intuitive, conversational, and immediately useful – has brought AI directly into the hands of the workforce. Staff across clinical, administrative, and operational domains are experimenting with ways to summarise, draft, and analyse more efficiently. Self-reported productivity gains are common, but perhaps more significant is the shift in sentiment: curiosity and confidence in AI have risen sharply.

This moment has prompted a deliberate pivot in approach. Systems like NHG Health are recognising the need for ambidexterity – combining the discipline of orchestrated, centrally governed deployment with the energy and learning that come from participatory, distributed adoption. Making generative tools widely available, while investing in responsible use and new skills such as prompt engineering, has proven a powerful way to accelerate confidence and capability.

At the same time, this new accessibility brings risk – of over-reliance, variable quality, and potential deskilling. Ensuring that empowerment is matched by proportionate safeguards and continuous learning has therefore become a defining challenge for health systems seeking to convert momentum into sustainable muscle.

The next wave is already visible. As large language models evolve into more autonomous, agentic systems capable of performing multi-step tasks, initiating actions, and collaborating with humans, the implications for healthcare work will be profound. Roles will need to adapt, workflows will be redesigned, and new forms of oversight and accountability will be required. The distinction between user and system is beginning to blur.

For health systems, this means investing not only in digital literacy but in workforce transformation – rethinking how expertise, judgement, and empathy are combined with machine capability. The organisations that thrive will be those that treat this disruption not as a threat to professional identity but as an opportunity to redefine it.

Three Levels of Promise

Clinical Promise	*Operational Promise*	*System Promise*

AI can enhance safety and judgement by processing information at a speed and scale no human can match. Imaging triage, early deterioration alerts, or personalised treatment suggestions can bring signal to the noise. The best tools complement intuition rather than override it – guiding attention, highlighting risk, and freeing clinicians to focus on the moments that require human presence and compassion.

Example: AI-enabled fall risk assessment

In partnership with National University Health System, NHG Health is adopting artificial intelligence to strengthen fall risk assessment by continuously analysing patient data and generating timely, evidence-based risk scores. This automation replaces manual checks that nurses currently perform at admission and whenever a patient's condition changes, a process vulnerable to missed changes in clinical status.

An AI-enabled system automates routine screening, monitors patients in real time, and alerts staff only when risk meaningfully increases. By processing data from multiple sources including vital signs, medication records, and mobility patterns, the system detects subtle changes that would otherwise go unnoticed between scheduled assessments.

This approach reduces nursing workload, supports consistent adherence to fall-prevention protocols, and enables proactive intervention. The potential impact includes fewer fall incidents and their associated complications, costs, and extended hospital stays.

Example: Radiology co-pilots in primary care (NHG Polyclinics)

NHG Polyclinics (NHGP) has begun embedding AI into everyday radiology workflows to improve speed, consistency, and diagnostic confidence across high-volume primary care settings. Rather than replacing clinical judgement, these tools act as a "second set of eyes" – helping clinicians and radiologists prioritise risk, reduce missed findings, and focus attention where it matters most.

One example is an AI chest X-ray triage tool (e.g., for TB screening and urgent abnormalities) that flags and prioritises higher-risk films for earlier re-view. It was progressively introduced at a single polyclinic site in March 2025 and has now been widely adopted. By structuring triage within the workflow - not as an "extra step" - it supports faster escalation for concerning cases while maintaining safety and consistency at scale.

In parallel, NHGP has explored AI-enabled decision support for earlier cancer detection in primary care. A skin cancer recognition and e-referral workflow ("C.A.R.E") helps streamline clearly benign presentations while fast-tracking suspicious lesions into appropriate referral pathways. Additional pilots – including fracture detection, mammogram second-reading support, and thyroid ultrasound assistance – are planned to further strengthen patient safety and reduce turnaround times as imaging demand grows.

Example: Ambient voice transcription to give time back to care

Since late 2024, NHG Health has been introducing ambient AI documentation tools to reduce the administrative load on frontline teams and help clinicians stay present with patients. These tools listen during consultations or sessions and generate a structured draft note that staff can quickly review, edit, and finalise, rather than documenting from scratch.

At Woodlands Hospital, Scribe was first piloted with Medical Social Workers (MSWs) in October 2024. For suitable sessions, it reduced documentation time by 41.8%, and has since been adopted by MSWs across the system. In early 2025, NoteBuddy was deployed for use by clinicians in NHG Polyclinics (NHGP) with a similar approach, converting consultations into concise, structured summaries to speed up documentation while maintaining quality.

Why does this matter? By making routine documentation lighter, voice-to-text tools reduce cognitive burden, improve the quality and consistency of notes for handover and continuity, and allow staff to focus more attention on the person in front of them, not the keyboard. Both Scribe and Note Buddy also support multilingual capture, which helps teams working across diverse language contexts.

In NHG Health and other HAT partners, the most immediate gains have emerged behind the scenes. Generative AI is helping to draft correspondence, summarise records, and coordinate care. Predictive models optimise rosters, anticipate bed pressures, and align diagnostics to demand. Productivity gains are real, but the greater value is reliability: smoother flow, fewer delays, and teams that spend more time on care and less on correction.

Example: Predictive workforce scheduling ("liquid nursing")

Tan Tock Seng Hospital of NHG Health was conferred the National Healthcare Improvement and Productivity (NHIP) 2025 Best Practice Medal (Workforce Transformation category) for their Liquid Nursing Initiative.

In NHG Health a "liquid" workforce model rethinks nursing rosters as a dynamic matching problem rather than a fixed monthly plan. Using real-time data to match staffing supply with patient demand, the model supports predictive scheduling and more modular roles so staffing can flex where care needs are shifting, without constantly rebuilding rosters from scratch. Instead of managers spending hours wrestling with rigid templates, the system helps generate workable options quickly, leaving human leads to apply judgement, resolve exceptions, and support team preferences.

In practice, this kind of approach has delivered sizeable operational and workforce benefits. Roster planning time fell by 83% (from ~90 minutes to ~15 minutes per workspace). Nursing attrition improved from 7.8% to 5.8%, exceeding target. And by reducing avoidable inefficiencies in allocation and coverage, the programme delivered savings equivalent to ~59.5 FTEs (about $4.46m annually). The wider point is not just cost: predictive workforce planning can reduce administrative burden, improve staff experience, and make service delivery more resilient when demand is volatile.

Example: LLMs in practice - from quick wins to system-scale impact

Across the more than fifty large-language-model (LLM) projects submitted to the CHI Innovate Conference from across NHG Health, a clear pattern emerged: some of the biggest gains come not from "headline" clinical AI, but from everyday friction points - translation, information retrieval, training, screening, and access. The examples below illustrate a range from small, targeted workflow accelerators to system-level digital front doors:

Making information clear and accessible
Frontline teams highlighted the time lost rewriting, simplifying, or translating patient-facing materials – often with long delays before publication. Tools such as PEanut and SG Med Translate now generate plain-language versions and multilingual translations (Mandarin, Malay, Tamil) within minutes, reducing turnaround from days to hours and achieving over 95% accuracy. The result is 179,000 hours saved and over $1 million in avoided cost, while enabling more consistent, equitable communication with diverse communities.

Supporting clinical decision-making at the point of need
Radiographers described the difficulty of locating the correct protocol in time-pressured settings, often interrupting senior colleagues. XRBuddy, a RAG*-enabled assistant, puts validated imaging protocols into a conversational interface accessible in under 30 seconds – reducing lookup time by nearly 90%. Early evaluation showed high scores for faithfulness and correctness, giving staff confidence and protecting patient safety through more consistent practice.
**retrieval-augmented generation*

Learning safely with virtual patients
Nurses needed scalable, realistic communication training without relying on resource-heavy simulation. A generative AI nurse–patient conversation chatbot now offers seven interactive scenarios with immediate, structured feedback. It has delivered up to 450% improvement in learner scores, strengthening confidence and preparedness for assessments while creating a psychologically safe environment to practise complex interactions.

Improving financial and discharge planning
Financial Care Officers previously spent hours manually screening thousands of patient records to identify potential bad debt risk. BEACON, an AI-supported risk-flagging system, automates this screening across more than 2,700 cases each month, saving 458 man-hours and enabling earlier interventions. Staff described being able to "shift from administrative sorting to meaningful counselling," improving both timeliness and quality of support.

A single front door to the system
Patients often struggle to know where to go or whom to ask. The enhanced NHG HealthBot, now supported by LLMs, provides natural, context-aware guidance across platforms — from appointments to service information. With WhatsApp penetration at over 80%, HealthBot is projected to reduce enquiry workload by 24% and avoid millions of dollars in call-centre costs by FY29, while providing more consistent access to timely information.

What These Examples Show
Together, these projects reflect a broader truth surfaced throughout the report. AI's value lies not in its novelty, but in its ability to solve everyday problems in ways that build confidence, coherence, and capacity. They illustrate how generative AI, when governed proportionately, can turn small moments of friction into moments of flow.

At population scale, AI can illuminate the invisible. It can forecast demand, segment risk, and test the likely impact of prevention strategies before they are implemented. But these insights matter only if they are embedded in decision cycles – procurement, planning, and performance – where they inform resource allocation and strategy rather than being confined to static reports and dashboards.

However, at system level, the promise of AI is more emergent – and more constrained – than in clinical or operational domains. While health systems hold vast quantities of data, much of it remains fragmented across settings, episodic in nature, and oriented toward documentation and reporting rather than prediction and planning.

This makes population-level AI fundamentally different. Its value does not lie primarily in automating existing processes, but in enabling a shift from retrospective analysis to anticipatory insight: forecasting demand, identifying emerging risk earlier, and supporting proactive approaches to prevention, service design, and resource allocation.

Internationally, real-world examples of population health AI at scale remain limited. Where progress has been made, it is often in narrowly defined use cases – such as risk stratification for specific conditions, demand forecasting for selected services, or scenario modelling for capacity planning – rather than fully integrated, system-wide approaches. The constraint is rarely algorithmic sophistication; it is the absence of longitudinal, linked, and trusted data foundations that allow insight to travel across time, settings, and decisions.

The significance of this should not be underestimated. If realised, AI-enabled population health capability could help systems move upstream – shifting effort from reacting to demand toward shaping it. But doing so requires deliberate investment in data integration, governance, and analytic capability, alongside clarity about how predictive insight is used in planning and commissioning decisions. In this sense, population health represents not the most mature application of AI today, but one of its most consequential horizons.

The Human Promise

HAT participants repeatedly observed that AI's deepest potential is not technological but human. When well used, it gives professionals back time, attention, and confidence. It can shorten meetings, reduce cognitive burden and burnout, and make multi-disciplinary teams work as one, as one participant said,

"Technology can't replace care, but it can give us more of it to give."

The shift to generative AI has shown that participation is power: when people are trusted to experiment safely, they learn faster and adopt more deeply. NHG Health's ambidextrous approach – combining orchestrated, centrally supported projects with distributed, participatory adoption – is a reflection of our learning on this. It recognises that innovation thrives at the edge, but alignment must come from the centre. Providing access to tools, pra-

ctical guardrails, and training in emerging skills such as prompt engineering has built confidence and trust. Yet leaders remain mindful of the risks: over-reliance, hallucination, and gradual deskilling. Managing this balance will define the next phase of workforce transformation.

Ambidextrous Leadership - Balancing Orchestration and Participation

A recurring theme across the HAT and Innovate discussions was the need to balance orchestrated and participatory approaches to adoption. Traditional technology rollouts rely on centralised control and linear decision-making. But in the era of generative AI, where tools evolve rapidly and use cases multiply daily, control alone cannot keep pace.

Health systems such as NHG Health and others internationally are adopting an ambidextrous approach - combining strategic orchestration with distributed participation.

Central governance defines guardrails, ethical principles, and shared infrastructure, while local teams are empowered to explore, adapt, and learn within those boundaries.

This duality creates both safety and speed, ensuring that innovation flourishes within a framework of responsibility. As one senior participant at Innovate summarised,

> *"If we try to control everything, we'll slow down. But if we try to do everything everywhere, we'll drown. The goal is disciplined freedom - learning together while staying aligned."*

This mindset shift – from command and control to coordination, from compliance to capability - is proving essential to make AI adoption both scalable and sustainable. It reflects a mature form of leadership: one that creates the conditions for learning, not just the rules for compliance.

Why The Promise Still Outruns Reality - Even as Activity Accelerates

Despite extraordinary investment and enthusiasm, most AI initiatives in healthcare have yet to achieve systemic impact. Learning from across our Thinktank participants identified a consistent pattern – barriers are rarely primarily technical, they are systemic, rooted in how health systems set intent, establish data foundations, govern risk, grow capability, and learn from real-world use. These are not failings. They are features of health systems trying to absorb a technology that is evolving faster than established governance and adoption processes can adapt. But the implication is clear: sustainable progress requires deliberate system-building, not more pilots. In other words, AI's binding constraint is not invention; it is implementation – the ability to embed, govern, evaluate, and sustain adoption in real-world practice.

A central theme is that activity is outpacing coherence. Pilots proliferate, but routes to sustained adoption remain unclear. Lessons are generated, but they do not reliably spread. Value appears in pockets rather than becoming part of routine practice. In parallel, many of the mechanisms health systems rely on to manage change - hierarchical decision pathways, process-heavy governance, and procurement designed for fixed products - struggle when the technology is adaptive, user-dependent, and requires continuous assurance.

These challenges intensify as innovation becomes increasingly distributed. Generative tools lower the friction of experimentation and place usable capability directly in the hands of frontline staff. That is a strength – but it also increases variability, amplifies risk, and widens the gap between early adopters and those who feel uncertain, excluded, or anxious about safety and accountability. Without shared foundations and proportionate oversight, momentum can become diluted.

The encouraging news is that these patterns are addressable – but not individually. They are interdependent and demand a systemic response. Treating them as separate problems often produces local fixes that do not add up to system change. What is needed is an orchestrated approach: building the muscle of adoption, governance, evaluation, and learning, so that every deployment strengthens the system's capacity rather than adding to the noise.

8 Recurring Patterns that Stall System Progress

01 The Experimentation Trap

Pilots proliferate but remain disconnected from organisational priorities, generating activity without cumulative learning.

02 Fragmented Intent

Teams and institutions pursue divergent AI projects, diluting effort and confusing accountability.

03 Legacy Data Foundations

Systems built for documentation, not intelligence, lack the interoperability and quality to support the AI lifecycle.

04 Transactional Procurement Mindsets

Governance and contracting models designed for fixed products struggle to manage adaptive, learning systems.

05 Skills and Confidence Gaps

Technical fluency and AI literacy are uneven, slowing adoption and eroding trust.

06 Governance Overload

Risk aversion and compliance-heavy processes can stifle safe experimentation rather than enable it.

07 Short-Term Pressure

Leaders are pushed to show results quickly, but genuine transformation demands time, learning, and behavioural change.

08 Uneven Learning Cycles

Evaluation remains episodic rather than embedded, limiting the feedback required for improvement and scale.

Countermoves - How Systems Turn Activity into Impact

The counter to fragmentation is coherence. The practical countermoves that surfaced across our discussions are less about adding new tools and more about changing the system conditions around adoption – turning pilots into pathways, and isolated learning into shared capability.

First, systems need to anchor AI in strategic purpose. Not because technology is unimportant, but because it is abundant. Purpose creates focus: it clarifies which use cases matter most, where to invest, and what 'good' looks like in practice. That focus then enables hard choices – stopping work that cannot scale, consolidating overlapping efforts, and building a portfolio that compounds rather than scatters.

Second, systems need foundations that allow learning to travel. Interoperability, data quality, deployment pipelines, monitoring, and transparency are not technical 'nice-to-haves'. They are the adoption spine that makes progress scalable. Without them, each pilot becomes bespoke. With them, adoption becomes faster, safer, and more comparable across settings, and learning can diffuse across the system instead of staying trapped in local teams.

Third, governance has to evolve from static control to adaptive oversight. Oversight calibrated to risk - anchored in clinical practice and close enough to real workflows – creates both safety and speed. Done well, it makes responsible experimentation easier, not harder. It also supports a culture where staff can be curious without being cavalier, and exercise caution constructively. The principle is to make governance an engine of learning, not a barrier to it – as identified by an Innovate panellist "Governance should fuel confidence, not restrain it."

Finally, leaders must treat adoption as capability-building. Each implementation should leave the system stronger: with better skills, better methods, better infrastructure, and a clearer understanding of what works, for whom, under what conditions, and with what trade-offs. This is how activity becomes impact - and how impact becomes something that can endure.

8 Countermoves that Build Coherence

01 Anchor in Strategic Purpose

Begin with the health outcomes and system priorities AI should advance, not the technology itself. Clarity of purpose is the foundation of coherence.

02 Create Conditions for Learning, Not Control

Replace one-off pilots with an architecture for continuous experimentation and feedback, where every deployment strengthens capability and maturity.

03 Invest in Platforms and Data Foundations

Build interoperable, privacy-preserving infrastructure that supports the AI lifecycle from design to deployment and monitoring, enabling scale with confidence.

04 Embed Proportionate Governance

Establish adaptive oversight that differentiates between low and high-risk applications, rewarding responsibility rather than punishing innovation.

05 Grow People and Confidence

Develop the technical, clinical, and leadership skills needed to use AI wisely and well. Capability and confidence must grow together.

06 Measure What Matters

Evaluate not just performance metrics but learning, trust, and the redistribution of value across the system. Progress is as much about culture as outcomes.

07 Build Partnerships for Scale

Work with regulators, academia, and industry through shared sandboxes and responsible innovation frameworks that align incentives and expectations.

08 Lead with Intentionality

See adoption not as a sprint but as an evolving journey – one that builds the strategic muscle for sustained transformation.

What unites these countermoves is not technology, but intentionality. They represent a shift from opportunism to orchestration – from isolated excellence to collective competence.

At the heart of this shift is the idea of systems learning. Every deployment, successful or not, is a data point for improvement. Every lesson should flow – not just throughout institutions, but across systems, and even across nations.

This was the spirit that animated the HAT and Innovate gatherings: that progress in AI for health is a shared enterprise. The systems that will thrive are those that view adoption not as a race to deploy, but as a process of learning their way to scale.

Taken together, the golden threads here are straightforward: direction over speed; platforms over point solutions; capability over novelty; proportionate oversight over blanket control; and learning as the default mode of progress. These are the habits that convert momentum into muscle.

Golden Threads – What This Means for System Leaders

Direction over speed
Fewer, better bets aligned to system goals create coherence and focus.

Platform over point solutions
Connected infrastructure makes scaling possible and sustainable.

Capability over novelty
Long-term value depends on people, not just products.

Proportionate governance
Adaptive oversight enables confidence and continuity.

Continuous learning
Embedding feedback and reflection makes progress cumulative, not episodic.

These golden threads reflect the central message of this report: progress in AI depends less on invention than on intention, and from the leadership capacity to align people, purpose, and process around shared learning.

With this strategic context in view - the promise, the constraints, and the countermoves – the rest of the report turns to the practical leadership challenge: how to hold the paradoxes that come with AI, how to evolve responsible governance into a discipline of stewardship, and how to build the enabling conditions that turn experimentation into coherent, sustainable transformation.

From Hype to Muscle

The current wave of generative AI has brought renewed excitement and experimentation to healthcare. To convert this momentum into muscle the task of leadership is not to dampen enthusiasm but to channel it intentionally, using each pilot, platform, and partnership as a building block in a coherent architecture for impact. Transformation will not happen through isolated sparks of innovation, but through the patient work of designing systems that allow those sparks to ignite – providing the fuel (capability), oxygen (trust), and discipline (governance) – to create the relevance and draw into the system, to turn innovation into sustained, impactful change.

Ultimately, transformation at scale is not about having the most technology; it is about having the deepest capability for learning, collaboration, and trust. This is what turns motion into movement – and movement into meaningful, measurable progress.

Looking Forward – The Next Wave

The horizon is already shifting. Agentic AI – systems that can plan, call tools, and complete tasks collaboratively – will test the boundaries of current workflows and oversight. These advances will not just change how we use technology; they will reshape work itself. Preparing the workforce means more than literacy; it means designing new forms of collaboration between humans and machines, ensuring skills evolve as roles do.

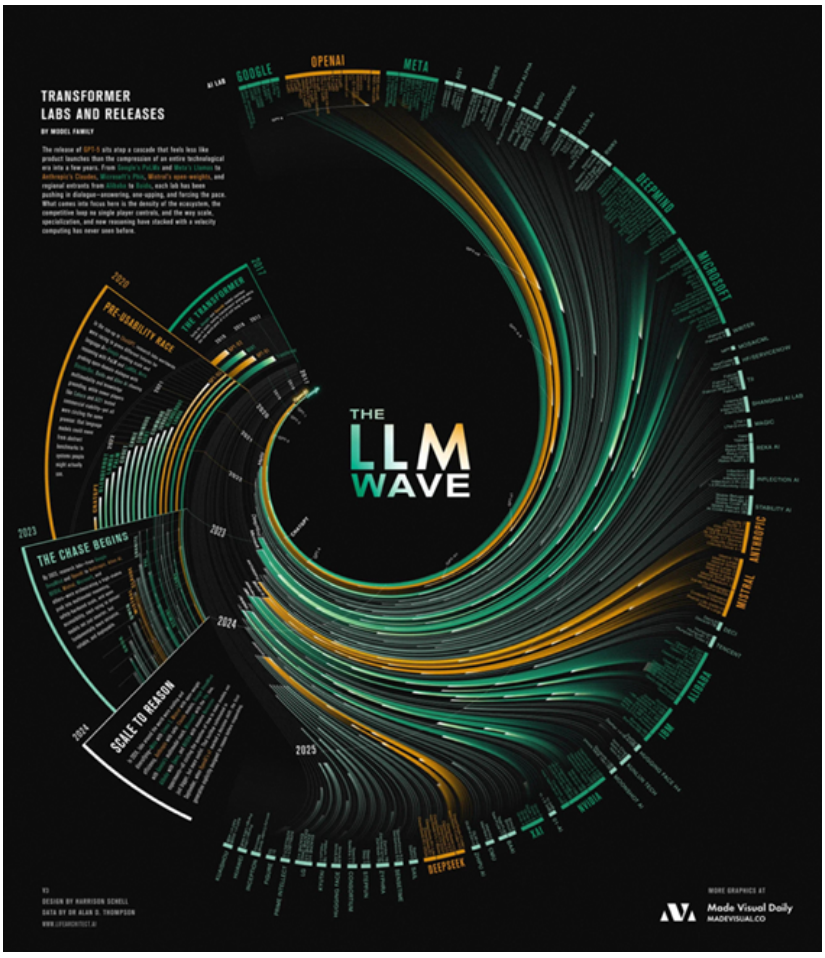

Graphic from Made Visual Daily [1]

The promise of AI in health remains profound: to restore capacity, insight, and humanity to strained systems. But realising it requires moving from enthusiasm to intention – from momentum to muscle. That is the work of leadership, and the focus of what follows. Every new capability brings its own tensions. As health systems accelerate change they also encounter the paradoxes that accompany rapid progress – where ambition can outpace alignment.

In Summary

The generative moment has expanded both possibility and participation. The next question for system leaders is how to translate this energy into purpose – turning exploration into evidence and exploration into lasting impact. AI's promise will not be realised through enthusiasm alone. It requires coordination, courage, and clarity of purpose.

The journey from momentum to muscle begins with understanding where impact stalls – and designing systems that can keep learning their way forward.

As health systems navigate this next phase of transformation, the challenge is not to 'do more' AI, but to become more adaptive as a system – able to respond to the rapidly changing opportunities that AI brings. Every new capability brings its own tensions.

[1] Harrison (2025). "8 Years of LLM Deployment Visualized + Methodology." Made Visual Daily. https://www.madevisual.co/8-years-of-llm-deployment-visualized-methodology/.

A Note on Leadership in the Age of Agentic AI

AI is redefining leadership as much as it is redefining work. As systems move from pilots to pervasive deployment, leaders are being called to steward transformation rather than simply sponsor it. The rise of large language models and agentic AI has created new dynamics – where individuals experiment before institutions regulate, and where confidence and curiosity often precede control.

Leadership in this context means creating permission and protection in equal measure. It requires holding space for uncertainty while maintaining clarity of direction – offering both the psychological safety to explore and the discipline to learn.

Effective leaders are shifting from commanding projects to orchestrating learning ecosystems, ensuring that every deployment, success, or failure feeds back into the organisation's collective intelligence.

In an age where AI can act, not just assist, leaders must remain the moral and strategic anchor.

They will need to guide teams through the paradox of increasing automation and deepening humanity, ensuring that technological capability amplifies purpose rather than replaces it. This is the art of leading transformation in the age of agentic AI – anchored in values, attuned to risk, and oriented always toward the human good.

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

Chapter 2 - Seeing the Wood for the Trees

Navigating the Paradoxes of Progress

Across health systems, enthusiasm for AI is palpable. Every week brings stories of new pilots, proofs of concept, and productivity gains – pockets of innovation emerging across the healthcare landscape. Yet despite all the energy, transformation struggles to take hold. The challenge is not failure but abundant activity without focus – plenty of effort, but little alignment.

Progress depends on a delicate balance of conditions – the capability that fuels it, trust that sustains it, and discipline that gives it form and direction. When any one of these is missing, activity disperses rather than deepens.

Participants at HAT and Innovate did not describe systems that had stalled, but ones that were alive – full of experimentation, curiosity, and ambition – yet still struggling to connect learning and convert it into consistent practice. Impact was fragmented, evidence inconsistent, and adoption patchy. This is not dysfunction, but a sign of systems in motion. The challenge is to channel this movement into muscle.

At the same time, leaders recognised the weight of expectation surrounding AI. The hype cycle has created a demand for rapid, visible returns – an impatience for results that can suffocate early innovation. Health systems are being asked to deliver transformation before the foundations of change have settled.

This tension is familiar to anyone leading large-scale transformation: how to balance accountability with incubation, expectation with exploration. Innovation needs time, trust, and space to mature – but those conditions are often constrained by short-term performance pressures and fiscal cycles.

The art of leadership lies in holding this paradox – giving new ideas oxygen without letting them overheat, ensuring that evaluation and expectation act as enablers of learning, rather than stifle it.

This is not failure, but part of how complex systems learn and adapt. As Peter Senge reminds us in The Fifth Discipline [2],

"the gap between vision and current reality is a source of creative tension – not frustration but energy for change."

In health systems, this tension is not something to eliminate but to hold, work with, and learn from. Progress in AI will always generate paradoxes: competing imperatives that must be balanced rather than resolved. The task of leadership is not to extinguish innovation but to shape the conditions under which it can coalesce into meaningful transformation.

The Paradoxes of Progress

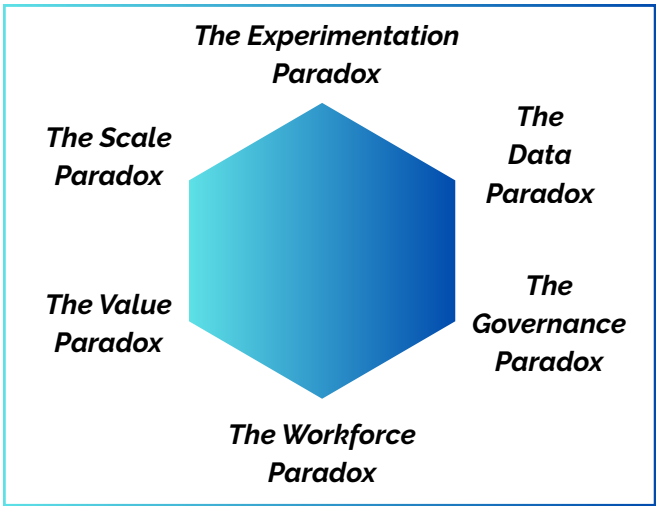

01 The Experimentation Paradox - Energy without Integration

Every health system is now an AI laboratory. The appetite to explore new tools, test new ideas, and experiment with generative models is strong. This creative energy is essential, but it also creates fragmentation: isolated pilots, parallel proofs of concept, and little shared learning.

Systems have learned how to start projects, but not yet how to sustain or compound their learning. When experimentation happens without intentional learning cycles or structured evaluation, enthusiasm becomes entropy. The challenge is not to stop experimenting – but to build the muscle of learning so that each experiment adds to collective capability.

"We're good at starting"
"but not yet at learning across." - one HAT participant shared.

02 The Data Paradox – Abundance Without Alignment

Most health systems sit on vast reservoirs of data, yet they remain stranded in silos – structured in ways that reflect history more than utility. More data does not automatically mean more intelligence. Without common standards, shared purpose, and clear custodian-

[2] Senge, P. (1990). The Fifth Discipline: The Art & Practice of the Learning Organization. Doubleday.

ship, information becomes noise. The paradox is that systems overflowing with information often struggle to use it – as Coleridge captures so beautifully in The Rime Of the Ancient Mariner (1798), *"water, water, everywhere, nor any drop to drink."*

Without common standards, trusted sharing frameworks, and the ability to connect and interpret data meaningfully, information becomes volume without value. Leaders at HAT described this as the difference between having data and using data well. Data readiness is not about scale or storage – it is about coherence. Real progress depends as much on governance and culture as on technology.

HAT participants often described this as moving from ***"data access"*** to ***"data coherence"***.

03 The Governance Paradox – Control Without Confidence

Traditional governance frameworks are designed for stability. They are procedural, sequential, and centralised – effective for static technologies, but optimised for predictability rather than continuous evolution, which is required for AI. As HAT's Responsible AI discussions highlighted, when governance tries to control innovation too tightly, it risks stifling the very learning that ensures safety.

The paradox is that excessive control can erode confidence. Responsible AI requires proportionate, risk-based governance that learns as it regulates. Stewardship must replace command and control. The shift is from compliance to confidence, from restriction to responsibility shared at every level.

"The goal," one speaker at Innovate observed,
"isn't perfect control – it's trusted adaptability."

04 The Workforce Paradox – Enthusiasm Without Enablement

Generative AI has transformed the workforce landscape almost overnight. Staff across clinical and administrative roles are experimenting with chatbots, LLM-based tools, and workflow automations. The result is a bottom-up surge of creativity – a distributed wave of adoption far outpacing the capacity of systems to support it.

This is progress of a kind: curiosity and confidence are building. Yet without structured enablement, self-directed innovation risks inequity, inconsistency, and even deskilling. The paradox is that enthusiasm, without coordinated capability-building, can widen the gap between pioneers and those left behind.

"Everyone's trying something," one leader reflected,
"and that's good – but they need a safe runway, not a warning sign."

05 The Scale Paradox – Diffusion Without Depth

Every system wants to scale what works. Yet scale, if misunderstood, can become the enemy of context. Replicating a solution without adaptation can erode relevance and trust.

The paradox is that standardisation and localisation must evolve together. Scaling responsibly means designing for difference – creating frameworks that enable shared learning while respecting local variation. True scale is not replication but diffusion with depth. As one HAT participant put it:

"Scaling means learning faster, not copying faster"

06 The Value Paradox – Progress Without Proof

Even when AI tools improve productivity or clinical processes, the evidence of impact often remains elusive. Systems seek measurable return on investment (ROI), yet measurement frameworks lag behind reality.

Leaders are often caught between pressure to demonstrate outcomes and the need to let innovation mature. The paradox is that systems demand evidence before scaling, but scaling is often necessary to generate evidence. Developing new forms of evaluation – blending quantitative data with qualitative learning – is therefore essential. The goal is to move from proof-of-concept to proof-of-value.

A deeper challenge lies in the expectations placed on innovation itself. Health systems operate under intense fiscal, political, and social pressure to show tangible results. The impulse to prove impact early is understandable – leaders must justify investment and maintain public trust. But when this expectation becomes excessive, it risks suffocating the very conditions that enable transformation. True innovation needs time, trust, and tolerance to learn and adapt. The art is to balance accountability with incubation – giving new ideas oxygen without letting them overheat under scrutiny.

Pitfalls and Progress

Holding these paradoxes is not easy. Under pressure for quick wins or certainty, systems can slip into familiar pitfalls that convert tension into traps. HAT discussions revealed **five recurring traps:**

Mistaking governance for control

Prioritising procedure over trust.

Confusing experimentation with progress

Mistaking activity for learning.

Treating data as a technology problem

Ignoring the cultural and strategic dimensions.

Defining scale as replication

Neglecting local adaptation and feedback.

Expecting impact before readiness

Compressing the time needed for innovation to mature.

Systems understandably want to show results and maintain momentum, yet early-stage change often follows a slow, uneven trajectory. When success is defined too narrowly or too soon, learning becomes an exercise in validation rather than discovery.

Avoiding these pitfalls requires leaders to work with paradox rather than against it. From the collective learning of HAT and Innovate, **five practices** stand out as enablers of progress:

Name the tension

Make paradox visible so it can be navigated, not denied

Balance orchestration and participation

Align distributed innovation with coherent direction.

Design for learning

Treat every deployment as a live experiment with embedded feedback cycles.

Anchor governance in trust

Move from control to stewardship, applying proportionality and shared accountability.

Build adaptive capability

Strengthen data, digital, and human competencies to sustain agility and optionality as technology evolves.

In every case, progress came not from resolution but from reframing – turning each paradox into a platform for collective learning.

During the HAT sessions, leaders described AI transformation not as a technical project but as "stewardship in motion".

"Stewardship in motion," as one participant reflected,
"means knowing when to hold back – not just when to push forward."

Like walking a transformation tightrope, the challenge is to balance the need for speed with the duty of care – to create space for experimentation while upholding responsibility. What emerged was a vision of leadership grounded in learning: collective, transparent, and humble.

"Our job isn't to predict the future of AI – it's to build systems capable of learning their way there."

Closing Reflection

Paradox is not the opposite of progress; it is the evidence of it. As systems experiment, learn, and adapt, contradictions surface – between speed and safety, centralisation and autonomy, experimentation and evaluation. These tensions are not signs of failure but indicators of growth.

The art of leadership lies in staying with the discomfort long enough for new patterns to emerge. In AI for health, progress will not be linear. But if systems can hold the tension between innovation and integrity, they can turn sparks into sustained flames – lighting the path toward transformational impact.

Addressing these paradoxes begins with responsibility. As we will see, to sustain trust and momentum systems must embed ethics, proportionality, and shared accountability into every layer of innovation.

Chapter 3 - From Governance to Stewardship

Responsible AI in Practice

As the adoption of artificial intelligence accelerates, so too does recognition that traditional models of oversight and accountability no longer suffice. Health systems have long relied on governance as a mechanism for safety – linear processes that authorise, monitor, and audit.

Yet AI, particularly with the emergence of generative and agentic models, operates in ways that are dynamic, adaptive, and context-specific. It learns, changes, and interacts in real time. This demands not simply new rules, but a new mindset – a shift from governance to stewardship, from asking **"Who is responsible?"** to exploring **"How is responsibility practiced?"**

> *"We can't govern innovation the way we govern risk.*
> *It has to breathe, but it still needs to be accountable."*
> - HAT participant.

This shift was a consistent thread across the HEAL Advisory Thinktank (HAT) and CHI Innovate discussions. Participants emphasised that the old governance paradigm – risk-averse, centralised, procedural – is ill-suited to the distributed and iterative nature of AI adoption. Responsibility must move with capability. It must be proportionate to risk, embedded in practice, and reinforced through a culture of learning and reflection.

Operationalising Responsibility

Proportionate Governance

The foundation of responsible AI practice is proportionality – the idea that oversight should match the level of risk, impact, and autonomy of the system in question. Not every tool requires a full ethical review or central approval; many benefit more from clear guidance, documentation, and locally empowered leadership. For higher-risk applications – those that directly influence diagnosis or treatment – scrutiny must be more formal, multidisciplinary, and transparent.

"The right question isn't 'Is it safe?' but 'What kind of safety do we need for this kind of tool?'"

- Innovate panellist.

This approach replaces blanket control with calibrated assurance, allowing innovation to progress while maintaining trust and accountability.

Devolved Accountability

Responsibility in AI-enabled systems cannot reside solely in oversight committees or data offices. It must be shared and situational. Clinicians remain responsible for how they use AI in decision-making; digital teams for ensuring model integrity and security; organisational leaders for cultivating ethical norms and competence.

"Accountability has to live where the work happens."

- HAT participant.

The principle is not diffusion but distribution – each actor accountable within their domain, connected by clear processes for escalation and shared learning. This distributed model makes responsibility visible, local, and lived rather than abstract.

Learning-driven Evaluation

AI deployment should be treated as the start of learning, not the end of testing. Embedding evaluation into real-world use transforms governance from retrospective policing to active improvement. Continuous monitoring, drift detection, and post-deployment surveillance allow systems to identify where performance degrades or bias emerges.

"Every model teaches us something - even the ones that fail." - Innovate speaker.

At the HAT sessions, participants described this as creating a living laboratory – each deployment contributing to collective intelligence about what works, for whom, and under what conditions.

Ethical Reflexivity

Traditional governance often focuses on compliance with rules rather than reflection on consequences. A culture of responsible AI requires structured opportunities for teams to pause, question, and learn. Regular forums – short ethics reviews, post-implementation reflections, and multidisciplinary discussions – help teams anticipate unintended impacts, interpret model behaviour, and respond to emerging risks.

> *"We've spent years building governance for compliance. Now we need governance for curiosity."* - HAT participant.

Reflexivity shifts ethics from static checklists to continuous dialogue, embedding values in the everyday rhythm of clinical and operational work.

Capability Building

Perhaps the most significant enabler of responsible practice is capability. As AI becomes embedded in clinical systems and everyday workflows, responsibility becomes everyone's job. This means moving beyond generic training toward applied competence – understanding data provenance, interpreting model output, and identifying when to question it.

At Innovate, several speakers stressed that this is not simply literacy – it is leadership. Building the muscle of responsibility means equipping staff with the confidence to use AI safely, challenge its output when appropriate, and escalate concerns without fear.

From Governance to Stewardship

If governance is about control, stewardship is about care. The role of system leaders is evolving - from approving technologies to cultivating trust, capability, and curiosity across the workforce.

> *"The task isn't to slow innovation, but to make sure it grows in the right direction."* - Innovate panellist.

In a world where AI systems learn, adapt, and act, leaders must ensure that organisations learn just as quickly. This means orchestrating feedback cycles, resourcing reflective spaces, and recognising that every AI deployment – whether successful or not – contributes to collective learning.

Stewardship also demands moral clarity. As agentic AI grows more autonomous, leaders must remain the moral and strategic anchor – ensuring that human judgment and empathy stay at the centre of healthcare. The challenge is to preserve human purpose even as technology becomes more capable of independent action. That requires humility, vigilance, and an unwavering focus on the public good.

Ultimately, responsible AI is not a framework to be implemented but a discipline to be practiced. It is the ongoing work of aligning innovation with integrity, balancing permission with protection, and transforming governance from a brake on progress into an engine of learning and trust.

Summary - Building the Muscle of Responsibility

	What It Looks Like in Practice	Purpose
Proportionate Governance	Risk-based review pathways, tiered assurance models.	Enable oversight that scales with risk.
Devolved Accountability	Clear role delineation and escalation protocols.	Make responsibility distributed but visible.
Learning-driven Evaluation	Continuous model monitoring, post-market feedback cycles.	Treat deployment as discovery.
Ethical Reflexivity	Regular reflective reviews and cross-disciplinary discussion.	Embed values in day-to-day decision-making.
Capability Building	Applied AI literacy, scenario-based training, leadership engagement.	Create confident, competent, and curious users.

AI Guiding Principles for Health Systems

Ethics and Fairness – Design and deploy AI with integrity – minimising bias through diverse data, transparent governance, and continuous evaluation to ensure equitable outcomes for all populations.

Transparency and Explainability – Ensure AI systems are understandable and traceable. Use clear documentation and explainable methods so that clinicians, patients, and the public can trust how decisions are made.

Accountability – Define clear roles, responsibilities, and escalation paths. Establish governance structures that align with existing regulatory and professional standards to ensure safe, responsible use.

Safety and Security – Embed safety-by-design principles. Continuously test, monitor, and secure AI systems to protect patients, prevent harm, and ensure continuity of care.

Privacy and Data Protection – Apply privacy-by-design principles, using robust safeguards such as encryption, data minimisation, and privacy-enhancing technologies to protect individual data rights.

Inclusivity and Accessibility – Develop AI that works for everyone – reflecting the diversity of the populations served and ensuring accessibility for people of all backgrounds and abilities.

Reliability and Robustness – Validate and monitor systems to ensure consistent, high-quality performance. Retrain and adapt models as data, practice, and context evolve.

Human-Centric Design – Keep humans in control. AI should augment, not replace, professional judgment – supporting better decisions, stronger relationships, and improved outcomes.

Sustainability – Pursue energy-efficient and environmentally responsible AI operations as part of a broader commitment to sustainable health systems.

Continuous Improvement and Innovation – Treat every deployment as a learning opportunity. Monitor outcomes, refine systems iteratively, and enable safe, evidence-based innovation.

Trustworthiness – Build and maintain trust through ethical conduct, open communication, and proactive management of misinformation and misuse.

Responsible AI is the foundation – but progress depends equally on the scaffolding around it. The following section explores how leaders can create the conditions that make responsible, intentional, and adaptive change possible. **In the end responsible AI is less about keeping the technology in check, and more about keeping humanity in charge.**

Chapter 4 - From Momentum to Muscle

Creating the Conditions for Intentional Change

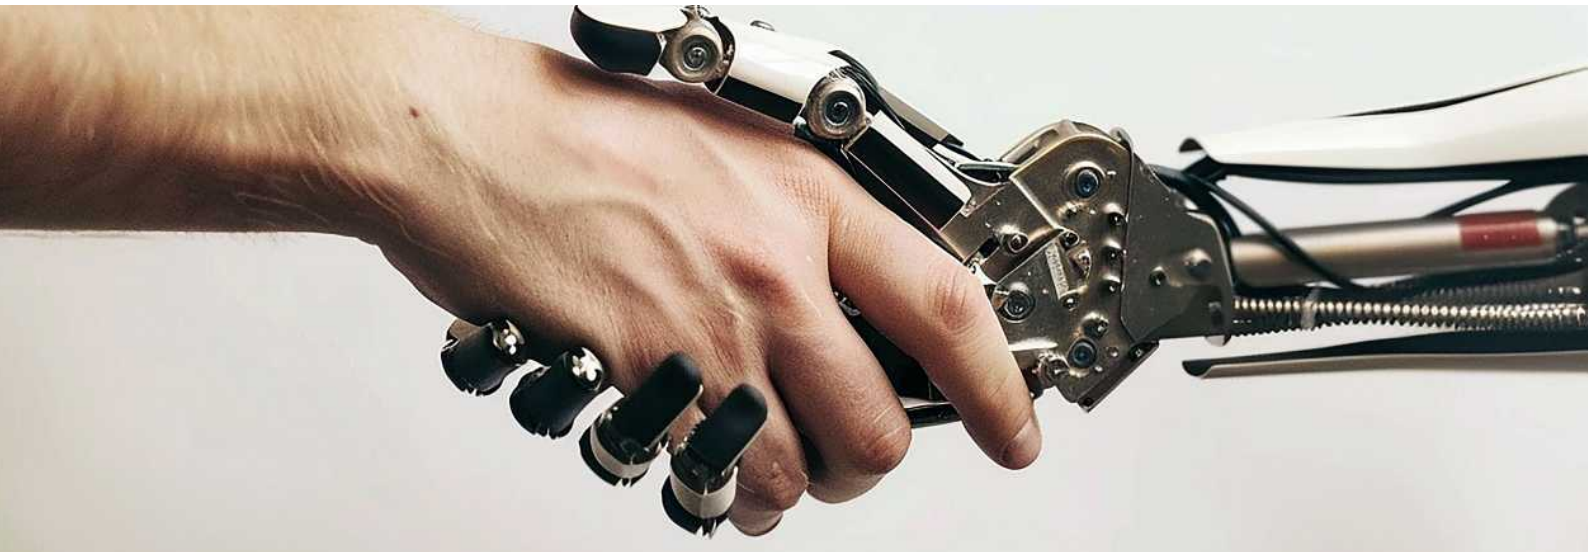

Having explored the paradoxes and pitfalls that often slow transformation, this chapter turns to what enables it – the conditions that make change intentional rather than incidental.

From Energy to Intentionality

When health systems first began experimenting with AI, the energy was unmistakable.

Everywhere, ideas were blossoming – clinicians prototyping new tools, data scientists eager to apply models, administrators hopeful for efficiency gains. That creative momentum was, and remains, a vital part of transformation.

But as one HEAL Advisory Thinktank (HAT) participant observed,

"There comes a point when you realise energy isn't the problem — it's direction."

The discussions across the HAT and CHI Innovate forums suggest we are at a clear inflection point. The focus has to shift from opportunistic experimentation to intentional system-building. Transformation at scale is not an act of chance or charisma. It is a discipline.

It emerges when health systems bring purpose, people, platforms, process, partnership, and performance into alignment – not as slogans, but as lived principles.

01 From Opportunistic to Systematic Adoption

Intentionality transforms innovation from enthusiasm into evidence. It is the discipline that turns activity into progress.

The first phase of AI adoption in many systems was, by necessity, opportunistic. It relied on inspired individuals and local pilots. That phase was essential: it built curiosity, revealed potential, and allowed early learning.

Yet it also created a pattern of fragmentation – a "thousand flowers" problem. Projects multiplied faster than they could be evaluated, integrated, or scaled. This proliferation taught a crucial lesson: without intentionality, innovation diffuses but doesn't deepen.

HAT participants reflected that the maturity of AI adoption is marked not by the number of pilots, but by the system's capacity to focus. As one Innovate panellist put it:

"We don't need more experiments — we need more experiments that matter."

Intentionality doesn't dampen innovation; it channels it. It means curating a portfolio of initiatives that align with system priorities, demonstrate real-world value, and are feasible to implement.

This is the pivot from exploration to execution – from the possible to the purposeful.

02 People with Purpose – The Heart of an Intentional System

Intentional change begins with intentional capability. People with purpose create alignment not through instruction, but through shared understanding.

At the centre of this shift are people – technology changes what is possible; people make it real.

Across HAT and Innovate, participants emphasised that capability is both the fuel and the limit of transformation. Systems must build AI fluency, not just literacy – moving beyond awareness to confident, critical, and creative use.

Fluency grows through doing. It requires environments where clinicians, analysts, and administrators can learn by applying AI safely, sharing insights, and seeing impact. These experiences create trust and capability far faster than theoretical training alone.

Leaders described this as "building the workforce muscle" – not by hiring more experts, but by cultivating adaptive expertise across disciplines. In the process, traditional hierarchies give way to networks of learning.

"AI is not replacing professionals," one HAT participant said.
"It's expanding what professionals can do – but only if they have the confidence to use it."

Leadership itself must evolve. The new generation of leaders must hold complexity lightly – balancing urgency with reflection, risk with responsibility. They must model curiosity, not certainty. These are not just technical leaders; they are meaning-makers who help staff see AI as a tool to fulfil their purpose, not obscure it.

03 Portfolios with Potential – Focusing Where Impact Meets Feasibility

Focus and discipline are not the opposite of innovation – they are what make innovation visible, viable, and valuable.

Intentionality also means choice. Health systems cannot – and should not – pursue every AI idea that emerges. The art of system maturity lies in prioritisation: focusing on what is both impactful and feasible, where complexity is manageable, and where lessons can compound.

This approach, refined through NHG Health's strategy development process following the HAT and Innovate sessions, defines a balanced AI portfolio – one that nurtures early wins while building future readiness. The criteria are pragmatic:

- **Clear alignment with system or service priorities**
- **Evidence of real-world demand or benefit**
- **Feasibility within existing infrastructure**
- **Capacity for evaluation and scale**

By focusing effort in this way, systems build credibility and confidence. Each successful deployment strengthens capability and coordination – creating a virtuous cycle that moves AI from peripheral to integral.

"We can't afford to chase novelty for its own sake," one HAT participant said. *"We need to show progress people can feel."*

Intentional portfolios also enable orchestration: aligning projects across domains to share learning and infrastructure. This reduces duplication and accelerates collective impact.

04 The Platform Ecosystem – Enabling Agility with Assurance

A platform ecosystem is not about building walls; it's about designing bridges — between experimentation and execution, between innovation and integrity.

Behind every credible AI programme lies a platform ecosystem that turns ideas into safe, scalable applications.

This is not merely infrastructure – it is an environment of interoperability, governance, and flexibility that allows systems to learn while they implement.

Example: Vendor-neutral AI deployment and monitoring - Newton's Tree

As health systems move beyond single pilots to running multiple AI tools safely at scale, a common bottleneck becomes the "plumbing" - integration, deployment, and ongoing assurance. This is where some systems are beginning to explore vendor-neutral platform layers that can standardise how AI solutions are onboarded, deployed, and monitored, without creating bespoke integrations for every tool.

One such platform we are exploring is Newton's Tree, which positions itself as an AI interoperability and deployment layer for hospitals: a capability that helps organisations connect to AI suppliers via a marketplace-style model, and manage deployments through a unified "mission control" view rather than bespoke integrations for every tool.

A key operational challenge is model drift - performance changing as patients, workflows, or equipment change. Through its Federated AI Monitoring Service (FAMOS), the platform is designed to monitor for shifts in data, patterns of use, and the outputs AI produces, supporting real-time risk management and continuous oversight. The approach also emphasises local hookup, where the AI is brought to the data with on-premise set-up, reducing the need to move sensitive data to a central server.

Why it matters: this kind of platform capability can reduce repeat integration effort, preserve vendor optionality, and make "responsible scaling" practical - because monitoring and assurance are built into the deployment pathway, not bolted on after the fact.

Example: Enriching clinical intelligence from messy data - Cogstack

Clinical data is often untidy: inconsistent standards, different documentation styles, and the reality that clinicians need expressive free-text notes all undermine data quality. As a result, valuable information can remain "locked" in narrative text, and teams end up doing significant manual curation before data is usable for clinical, research, or operational purposes.

One platform discussed in this context is CogStack – a clinical data enrichment suite designed to help health systems access and analyse information across electronic health records, clinician notes, and other digital sources, with deployments described across major UK organisations.

What this enables in practice:

- Faster, more reliable cohort identification and characterisation for research and trials through data enrichment and "deep phenotyping," supporting use cases like patient finding and population segmentation.
- Augmented clinical coding and registry workflows, positioning AI as an assistant to analysts and coders, designed to fit into existing processes rather than replace them.
- A more "reusable" data foundation - reducing one-off, manual extraction effort so that learning and insight can be generated more consistently across teams and over time.

Why this matters: this type of platform capability tackles a quiet but fundamental bottleneck in scaling AI – not the sophistication of models, but the usability of underlying clinical information. By making free text more computable and analysis-ready, systems can shorten the distance between documentation and decision-making, and reduce the manual drag that slows both operations and research.

At HAT, participants described an emerging architecture: vendor-neutral orchestration platforms that can test, validate, and deploy models safely; data environments that ensure quality and ethical use; and "AI marketplaces" that prevent vendor lock-in while accelerating experimentation.

The purpose of such platforms is not control – it is **optionality.** They allow systems to compare and deploy the best solutions available while maintaining oversight and trust.

Crucially, this ecosystem extends beyond technology. It encompasses the policies, standards, and agreements that make responsible AI possible. As one Innovate panellist put it:

"Trust is built into the architecture."

05 Processes for Learning and Scale – Turning Adoption into a Discipline

Processes should be predictable in structure, flexible in practice, and designed to make learning inevitable.

Every health system has processes. Few have processes that learn.

In many organisations, AI adoption has followed linear procurement pathways – specify, select, deploy, and hope for results. HAT participants agreed this model cannot keep pace with adaptive technologies.

Instead, AI implementation must become cyclical and learning-driven. Every deployment should create data for improvement, feedback for governance, and insight for scaling.

This is the essence of orchestrated and participatory adoption – the ambidextrous model NHG Health and others are developing. It combines central orchestration (to ensure safety, quality, and learning) with distributed participation (to harness frontline creativity and ownership).

"If governance is too rigid, innovation dies," one Innovate participant noted.
"If it's too loose, trust dies."

Implementation becomes a process of discovery; evaluation becomes part of daily practice. The result is a living system that evolves through use, not decree.

06 Partnerships – Co-creating Value and Progress

Partnerships are how trust travels – they turn boundaries into bridges and make progress a shared endeavour rather than a solitary race.

At HAT and Innovate alike, one principle was constant: transformation is relational before it is technical.

AI doesn't just connect systems – it connects people. Effective partnerships create the ecosystem of trust required for responsible progress.

- With vendors, systems move from contracting to co-creation – setting shared standards for safety, evaluation, and interoperability.
- With regulators, they evolve from compliance to collaboration – shaping frameworks that accelerate safe adoption.
- With academia and industry, they co-develop use cases grounded in real-world problems.
- With patients and communities, they build understanding and consent – recognising that public trust is the foundation of sustainability.

Partnerships also reinforce equity. By involving diverse voices early – from clinicians to communities – systems ensure AI reflects the people it serves.

As one HAT participant observed,

"True collaboration begins when we stop seeing vendors as suppliers and start seeing them as co-learners in the health ecosystem."

Collaboration, not competition, drives progress. The goal is not to outpace peers but to learn alongside them – sharing results, standards, and lessons so that improvement accelerates for everyone. Systems that cultivate open, trusted partnerships build momentum that no single organisation could achieve alone.

07 Performance – Measurement Matters

True performance is measured in progress, not perfection – in how systems learn, adapt, and grow more capable over time.

What we choose to measure shapes what we choose to value. In AI-enabled systems, that means widening the lens. Traditional indicators such as cost, throughput, and efficiency remain important, but they tell only a fraction of the story. The deeper signals of impact are often human: staff retention and recruitment, time returned to care, reductions in cognitive and administrative load, rising confidence, and stronger collaboration. These are the measures that show whether technology is helping people work better – not simply work faster.

Performance, in this sense, is not about policing; it is about understanding how improvement happens. Across HAT, participants called for more real-time, transparent, and human-centred performance intelligence – dashboards that track adoption, safety, trust, equity, and learning as actively as technical outputs. When performance helps teams see what is changing, why it is changing, and where additional support is needed, it becomes a tool for collective problem-solving rather than judgement.

Alignment of incentives also matters. Systems should reward learning and adaptation, not only delivery. Progress in AI adoption is rarely linear; it depends on iteration, feedback, and behavioural change. Performance frameworks that recognise this help build the psychological

safety required for responsible experimentation. As one HAT participant put it,

"Our goal isn't perfect models; it's better outcomes for people – that's the only performance measure that matters."

Crucially, performance should not become a contest between teams or institutions. Healthy comparison can inspire, but the real task is shared improvement – learning from variation rather than competing to erase it. Systems that surface and spread learning as readily as results build momentum that no league table can match.

A simple framework showing how systems can track progress across four dimensions of responsible AI adoption.

Health systems need indicators that show whether AI is helping the system learn and improve, not just whether a model is accurate. One way of doing this is through a balanced, capability-oriented indicator set that tracks four domains of progress:

Domain	What It Assess	Example Indicators
Adoption and Use in Practice	Is AI being meaningfully used?	• Proportion of eligible users adopting the tool • Patterns of real-world use (frequency, depth, setting) • Time-to-implementation for new tools
Workforce Experience and Capability	Is work getting better, not just faster?	• Time saved on cognitive or administrative tasks • Impact on workload and stress • AI literacy and confidence levels • Retention and role satisfaction in AI-enabled areas
Safety, Quality and Trust	Is the AI safe, reliable, and strengthening confidence?	• Stability of model performance over time • Alerts or deviations requiring follow-up • Staff trust, comfort, and perceived reliability • Equity indicators (e.g., variation in performance between groups)
Outcomes and Value	Are we improving care and system performance?	• Improvements in clinical or operational outcomes (e.g., change in turnaround time, capacity, or throughput) • Impact on clinical quality (e.g., diagnostic accuracy, avoided harm) • Time released to care / cognitive load reduction • Realised financial impact or cost avoidance • Resident and patient experience measures

Closing Reflection - The Architecture of Intentional Change

Intentional change is the quiet art of leadership – aligning energy and structure, courage and discipline, aspiration and accountability. It is how systems move from momentum to muscle, turning early sparks of enthusiasm into sustained, human-centred transformation.

Across HAT and Innovate, the message was clear: transformation cannot be improvised.

Systems that thrive are those that are intentional by design – clear on purpose, coherent in approach, and agile in execution.

They invest in capability not for its own sake, but to enable coherence.

They prioritise learning as the currency of progress.

And they balance innovation with integrity, ambition with accountability.

This is not a story of technology. It is a story of change – of how human intention shapes technological evolution.

The conditions for change are not accidental. They are built – intentionally, collaboratively, and with clear purpose in mind.

"The most transformative systems," one Innovate speaker concluded,
"are not those with the most tools, but those with the clearest intent."

Chapter 5 - Holding Tension, Leading with Intention

Leadership for an AI-Enabled Future

Artificial intelligence is reshaping the landscape of healthcare. Yet amid the noise – the hype cycles, vendor promises, and headlines – the quiet work of leadership has never mattered more. Across the HEAL Advisory Thinktank (HAT) and CHI Innovate conference, one message came through repeatedly: technology does not transform systems - people do. But for people to succeed, they need leaders who create clarity of purpose, coherence of effort, and confidence in change.

This chapter is a call to arms, built from that shared spirit. It offers provocations as mirrors, not maps: questions that help leaders move from enthusiasm to impact, and from the flickering flame of innovation to the steady burn of transformation. It also makes a simple claim: the story of AI in health is not primarily technological. It is profoundly human. It is about the systems we shape, the values we hold, and the judgement, compassion, and creativity that bring meaning to every capability we deploy.

From Momentum to Muscle - The Work of Leadership

In this report we have explored how health systems can turn the momentum of innovation into the muscle of sustained impact. The difference is intentionality.

Momentum is energy – the enthusiasm of early adopters, the excitement of new tools, the spark of discovery. Muscle is what enables systems to act with focus, coordination, and resilience. It is built through deliberate investment in people, processes, and platforms that can learn, adapt, and grow stronger over time.

Too often, health systems mistake activity for progress. The challenge is not a lack of innovation but a lack of integration. Pilots multiply, yet learning rarely spreads. Success remains local; evidence is partial; adoption inconsistent.

As Henry Ford famously observed, ***"If I had asked people what they wanted, they would have said faster horses."*** The point is not to accelerate the current way of working, but to rethink the system so that AI can be truly transformative. Leadership, in this context, is about changing form – turning diffuse experimentation into coherent, responsible adoption at scale.

The Leadership Provocations - Questions that Shape the System

The provocations below are deliberately simple. They are designed to be used in boardrooms, programme reviews, and leadership conversations as a discipline of reflection – to help systems learn, adapt, and lead responsibly in the age of AI.

Purpose and Focus

Are we leading with stewardship?

AI offers boundless possibility – to predict risk, personalise care, and optimise services. Yet the first question for any leader is not what AI can do, but which challenges warrant transformation.

- Does our use of AI align with outcomes that matter most for patients and residents – or merely what technology makes possible?
- Are we building a coherent portfolio tied to strategic priorities, or accumulating pilots without a route to scale?

Leadership at scale demands focus. It means saying "no" to protect coherence. It means treating choice as strategy – concentrating resources where they can create visible, replicable benefit. Stewardship is disciplined alignment: ensuring that every model deployed, every pilot launched, and every partnership formed advances a shared mission, with patient and public benefit at the centre.

Conditions for Change

Are we building capability, not just compliance?

AI adoption challenges long-standing habits of linear, compliance-driven management. Algorithms evolve, data changes, and context matters. Systems cannot manage adaptive technology with governance designed for predictable environments.

- Are we creating the conditions for responsible experimentation, or still trying to manage complexity with tools designed for predictability?
- Do our teams have "disciplined autonomy" - clear boundaries, permission to explore, and the support to learn?

The most effective leaders move from enforcing compliance to enabling capability. They invest in multidisciplinary teams, protected time for testing and reflection, and repeatable deployment pathways. They design governance that is clinically anchored and proportionate to risk – close enough to practice to keep pace with reality, and robust enough to sustain trust.

Trust and Confidence

Are we humanising change as capability spreads?

Generative AI has decentralised experimentation. Clinicians, administrators, and analysts are trying tools and adapting workflows in real time. This is exciting - and risky.

- Do our people have the confidence, literacy, and support to use AI responsibly?
- Do our systems have the maturity to learn from what happens next, without blame or overreaction?

Trust is the currency of adoption. Over-control can stifle innovation; unchecked enthusiasm can create harm. Leaders must model responsible optimism: encouraging curiosity while making safety visible. Building confidence also means cultivating care - psychological safety, clear communication, and space for dialogue where people can express concern, ask questions, and help shape what comes next. Resistance is often a form of care itself - a signal that people value what is at stake.

Ethics as Practice

Are we governing in a way that learns?

Governance frameworks and ethical principles remain essential, but they must evolve as technologies do.

- Are our governance processes proportionate, adaptive, and able to recalibrate as real-world evidence emerges?
- Does ethics live in documents - or in the everyday design and deployment choices we make?

Responsible AI is not about eliminating risk; it is about balancing it with purpose. Ethics cannot live in policy alone. It must live in practice: transparency, escalation routes, feedback loops, and continuous monitoring, with decision-making appropriately devolved to those closest to the data and the patient. When governance becomes a learning process rather than a barrier, it builds trust, autonomy, and responsible speed.

Partnership and Power

Are We Building an Ecosystem, or an Empire?

AI sits at the intersection of public purpose and private sector innovation. The relationship between health systems and technology companies must evolve beyond procurement towards shared accountability for outcomes.

- Are we engaging partners as suppliers, or as co-stewards of health outcomes?
- Do we have the technical literacy to negotiate as equals, preserve flexibility, and define standards that protect trust?

Progress comes through partnership and collective intelligence – across clinicians and technologists, policymakers and practitioners, regulators and innovators. This requires humility and courage: to make decisions under uncertainty, to share power, and to invite scrutiny. In a world of intelligent systems, one of the smartest leadership moves is to stay connected.

Learning and Measurement

Are we measuring what matters, when it matters?

Health systems face pressure to demonstrate return on investment. Yet transformation rarely follows quarterly cycles.

- Are our evaluation frameworks enabling learning, or prematurely demanding proof?
- Are we capturing the full value – including capability, confidence, data quality, workflow change, and trust - not only narrow performance metrics?

Early deployments often generate outcomes that are intangible but vital: new collaborations, stronger foundations, cultural readiness. These are not failures - they are prerequisites. Leaders must treat measurement as a learning act: a way to test assumptions, refine processes, and build credibility over time. Evaluation should not arrive at the end; it should be embedded from concept to scale so that every deployment - successful or not - generates insight that can travel.

Holding Paradox as Practice – Conviction Without Certainty

AI brings extraordinary uncertainty - technical, ethical, social, and professional. The instinct of many leaders is to respond with control and assurance. Yet the work ahead requires a more deliberate balance: conviction about direction, and curiosity about how the path unfolds.

Peter Senge described effective leadership in complex systems as "holding creative tension" – the space between current reality and future aspiration. The leaders who will shape the next decade of AI in health are those who can act with authority while staying open to learning: confident enough to decide, reflective enough to adapt.

Curiosity sustains innovation. Conviction sustains trust. The art of leadership lies in holding both – and knowing when to lean into each.

Why This Report Does Not Offer a Universal Roadmap

One recurring question in AI transformation is whether there is a definitive roadmap – a sequence of steps that systems can follow to adopt AI safely, effectively, and at scale. This report intentionally avoids presenting such a guide. The reality is that no two health systems begin from the same place. Digital maturity varies, data foundations differ, governance cultures evolve at different speeds, and the pressures facing each workforce are shaped by their own history and context.

A single, linear roadmap risks oversimplifying a complex journey. It implies predictability where the terrain is shifting. AI adoption is not a linear rollout but a developmental process – one that depends on learning, feedback, capability-building, and context-specific adaptation. What succeeds in one organisation may fail in another without the enabling conditions.

Rather than a fixed set of instructions, this report offers a set of core ingredients, shared principles, and practical considerations that help systems make sense of their own path. It highlights what matters most for responsible adoption; the capabilities and conditions that make scaling possible; the tensions leaders must hold and the choices they must make; the patterns observed across diverse contexts that can guide, not prescribe.

The aim is to provide a strategic scaffolding, not a universal blueprint – enabling leaders to design their own pathways, grounded in local realities but informed by collective insight.

AI transformation is less like following a map and more like navigating moving water: the currents shift, the depth varies, and judgment matters as much as direction. This report seeks to equip leaders with the tools, language, and perspectives to read their own environment and steer with intention – recognising that while every journey is different, the underlying work of learning, enabling, and leading is shared.

Closing Reflection: The Human Endeavour

Every generation of technology reshapes what we do; AI reshapes how we think. It prompts fundamental questions about judgement, trust, and purpose. As one Innovate speaker observed,

> *"AI will do what it is trained to do - but only humans can decide what matters."*

Keeping the human centre is not sentimental; it is strategic. Empathy, ethical reasoning, and contextual understanding are what allow systems to use AI wisely. Without them, efficiency risks outpacing equity, insights blur into intrusion, automation may erode agency - and the promise of AI risks undermining the trust it depends on.

If there is one defining insight from HAT and Innovate, it is this: **The true measure of progress is not how many models we deploy, but how well we learn.** Learning systems are willing to learn by doing - connecting policy to practice, and implementation to reflection. They understand that every deployment is an experiment, every result a lesson, and every setback an investment in capability.

The future of AI in health will not be decided by algorithms. It will be determined by people – together – who choose to use technology not as a substitute for humanity, but as an extension of it. That is the work ahead. And it is, unmistakably, a human endeavour.

Meet The Team

HEAL Advisory Thinktank

Abdullahi Sheriff
Managing Director, Singapore, Malaysia and Brunei, MSD

Cherdchai Nopmaneejumruslers
Vice President for Innovation Technology and Digital Transformation
Mahidol University

Cora Han
Chief Health Data Officer
University of California Health

Eleni Dimokidis
Head of Healthcare Technology for Asia Pacific and Japan, Amazon Web Services

Eugene Fidelis Soh
Chair of Centre for Healthcare Innovation (CHI) Leadership Council

Goh Kim Huat
Professor of Information System, Nanyang Business School

Haris Shuaib
Director
Newton's Tree

James Teo
Professor of Neurology
King's College Hospital
Chief Medical Officer
Artificial Intelligence Centre for Value Based Healthcare

Jennifer Anne Cleland
Professor of Medical Education Research and Scholarship Unit (MERSU)
Nanyang Technological University

Jonty Heaversedge
Chair of HEAL Advisory Thinktank

Junaid Bajwa
Senior Partner, Head of UK, and Science Partner Pioneering Intelligence, Flagship Pioneering

Kavitha Palaniappan
Project Lead for Health Services Regulation
Centre of Regulatory Excellence
Duke NUS

Sujoy Kar
Chief Medical Information Officer & VP
Apollo Hospitals

Tom Kirby
Chief Financial Officer
Coalition for Health AI

HAT Workshop Contributors

Adison Goh
Data Scientist (Digital Innovation Office)
Centre for Healthcare Innovation (CHI)

Andy Ta
Chief Data Officer
Synapxe

Aung Myint Oo
Clinical Lead (Digital Innovation Office)
CHI

Beverly Chew
Management Executive (Digital Innovation Office)
CHI

Cheong Yu Chye
Deputy Chief Data Officer
NHG Health

Chew Zonghui
Senior Consultnat, General Medicine
Tan Tock Seng Hospital (TTSH)

Chiang Kar Wai
Assistant Director (Digital Innovation Office)
CHI

Chua Pei Xuan
Assistant Director (Digital Innovation Office)
CHI

Chua Ying-Hong
Director (HealthTech Policy)
Ministry of Health (MOH)

Colin Lim
Chief Information Office
MOH

Daniel Fung
CEO
Institute of Mental Health (IMH)

Daniel Ting
Director (AI Office)
SingHealth

Danny Lee
Manager, Group Digital Health, NHG Health

Darion Chong
Chief Risk, Compliance & Corporate Services
NHG Health

Elizabeth Pan
Assistant Director (Digital Innovation Office)
CHI

Eric Wong
Group Chief Digital Health Officer
NHG Health

Eugene Fidelis Soh
Deputy Group CEO (Population Health)
NHG Health

Francis Goh

Assistant Director
(L&OD Office)
CHI

Furqan Mohammad Shaheryar

Principal Data Scientist
(Digital Innovation Office)
CHI

Gerald Goh

Senior Lead Analyst,
Group Digital Health,
NHG Health

Guangyi Zhuang

Deputy Director
(Digital Health)
Health Sciences Authority (HSA)

Iris Tan Siu Kwan

Assistant Director
AI Office
SingHealth

Ishrat Khan

Assistant Manager
(Digital Innovation Office)
CHI

Jace Han

Business Analyst
(Digital Innovation Office)
CHI

Jason Cheah

CEO
Woodlands Health Campus

Joe Sim

CEO
NHG Health

John Abisheganaden

Chief, HSOR
NHG Health

Jonty Heaversedge

Director (Digital Innovation Office)
CHI

Justin Yeo

Sr Executive (Digital Innovation Office)
CHI

Karen Ng

CEO
NHG Polyclinics

Kelvin Li

Consultant
TTSH

Khoo Swee Lan

Group Chief Human Resource Officer
NHG Health

Low Shiong Wen

Chairman Medical Board Ng teng Fong General Hospital (NTFGH)

Lim Xiao Wei

Sr Executive (Digital Innovation Office)
CHI

Lynette Ong

COO
TTSH

Marcus Tang

Deputy Director (Digital Innovation Office)
CHI

Matthew Chua

Former Director (Digital Innovation Office)
CHI

Michelle Lee

Deputy Director
(Partnerships Office)
CHI

Ng Yeuk Fan

Director, Corporate Development
Khoo Teck Puat Hospital & Yishun Health

Nicholas Nghai

Health Policy Analyst
(Regulatory Policy & Legislation)
MOH

Ong Jing Fang

Chief (Population Heath)
NHG Health

Palvannan S/O Kannapiran

Deputy Director, HSOR
NHG Health

Queenie Tan

Deputy Director, CEO Office-Chief Data Office
Synapxe

Raymond Chua

CEO
HSA

Raymond Lee

Augmentation Consultant
Synapxe

Ricardo Leite

CEO
HealthAI

Sandra Chua

Senior Assistant Director
(Regulatory Policy & Legislation)
MOH

Sucharita Hota

Sr Assistant Director, Group Medical Affairs, Medical Board Office
National University Health System (NUHS)

Tan Cher Heng

Group Chief Research and Innovation Officer
NHG Health

Tang Kang Choon

CEO
TTSH

Tan Ming Chu

Director, Technology Strategic Planning, Group Digital Health,
NHG Health

Tan Wei Hong

Sr Executive (Digital Innovation Office)
CHI

Tay Eng Sin

Sr Executive (Digital Innovation Office)
CHI

Terence Tang

Group Chief Clinical and Informatics Officer
NHG Health

Thomas Liew

Deputy Director (L&OD Office)
CHI

Tho Min Seng

Product Lead, Digital Innovation Office
CHI

Wong Hon Tym

Clinical Director
CHI

Wong Kirk Chuan

Group Chief Operating Officer
Woodlands Health Campus

Yong Keng Kwang

Group Chief Nurse and Chief Wellness Officer
NHG Health

Meet The Team

CHI Innovate Speakers

Abdullahi Sheriff

Managing Director, Singapore, Malaysia and Brunei, MSD

Amos Miller

CEO
Glidance

Annie Chang

Executive Director
Travel Agents and Tourist Guides Division
Singapore Tourism Board

April Chin

CEO
Resaro

Candy Gan

T1DM Patient Advocate & Lead
TypeOne.sg

Charlene Liew

Deputy Chief Medical Informatics Officer
Changi General Hospital

Cherdchai Nopmaneejumruslers

Vice President for Innovation Technology and Digital Transformation
Mahidol University

Chow Minyang

Consultant & Clinical Educator, TTSH

Cora Han

Chief Health Data Officer
University of California Health

Goh Kim Huat

Professor of Information System, Nanyang Business School

Habeebul Rahman

Asst Chairman Medical Board
TTSH

Haris Shuaib

CEO
Newton's Tree

Haslina Binte Hamzah

Deputy Director
Singapore National Eye Centre

James Teo

Professor of Neurology
King's College Hospital
Chief Medical Officer
Artificial Intelligence Centre for Value Based Healthcare

Jason Tamara Widjaja

Executive Director
AI Singapore Tech Center
MSD

Jennifer Anne Cleland

Professor of Medical Education Research and Scholarship Unit (MERSU)
Nanyang Technological University

Joanne Yoong

CEO
Research for Impact

Junaid Bajwa

Senior Partner, Head of UK, and Science Partner Pioneering Intelligence, Flagship Pioneering

Karen Wee

Executive Director
Lions Befriender Service Association (Singapore)

Lee Tze Ming

Co-founder
STUCK

Melissa Chew

Sr Principal Medical Social Worker
Woodlands Health

Ngiam Kee Yuan

Head of AI Office
NUHS

Ricardo Leite

CEO
HealthAI

Sally Lim

Directory of Nursing
IHH Healthcare Singapore

Sandy Ho

Assistant COO
National University Hospital (NUH)

Spencer Low

Head of Regional Sustainability, APAC
Google

Sujoy Kar

Chief Medical Information Officer & VP
Apollo Hospitals

Teo Yik Ying

VP (Global Health) & Dean, Saw Swee Hock School of Public Health, National University of Singapore

Tham Kok Leong

Partner
Allen & Gledhill LLP

Trishan Panch

Executive Director
Applications of AI in Healthcare
Harvard TH Chan School of Public Health

Meet The Team

Editorial Team

Jonty Heaversedge

Author (supported by ChatGPT!)

HEAL Advisory Thinktank (HAT)

Reviewers

Elizabeth Pan

Editorial Admin Team

Casper Ng

Editorial Admin Team

Germaine Tan and Shermaine Ong

Design Layout
Editorial Admin Team

About Centre for Healthcare Innovation

The Centre for Healthcare Innovation (CHI), hosted by NHG Health, serves as a transformation engine driving innovation across Singapore's healthcare system through a network of co-learning partners.

The network is advised by the **CHI Leadership Council,** an international advisory panel comprised of distinguished thought leaders in their respective fields, committed to pioneering systems innovation and driving change. They provide thought leadership to CHI and its partners in spearheading transformation across the health system, while fostering co-learning, collaborations, and facilitating innovation.

CHI innovates for future care, for future health, for future generations - by harnessing the power of our innovation cycle. This begins with redesigning care models to meet evolving needs, integrating smart technologies to enhance delivery, and reshaping workforce roles to build resilience and adaptability.

CHI's work spans multiple levels of impact: improving value at the point of care, enabling health and social change at the systems level, and cultivating a sustainable healthcare ecosystem for generations to come, driven by the goal of adding years of healthy life and shaping a healthier future for all.

Health Empowered by AI Launchpad
The Health Empowered by AI Launchpad (HEAL), hosted by the Centre for Healthcare Innovation (CHI), is NHG Health's dedicated capability for the safe, responsible, and scalable adoption of AI - translating innovation into real-world clinical and operational impact.

HEAL Advisory Thinktank
The HEAL Advisory Thinktank (HAT) was commissioned by the CHI Leadership Council and convened at the CHI Council Week during CHI INNOVATE 2025 to provide strategic insight into how health systems can move beyond experimentation towards responsible, system-wide adoption of AI.

For more information, visit www.chi.sg.

About NHG Health

NHG Health is a leading public healthcare provider in Singapore recognised for its quality clinical care and its commitment in enabling healthier lives through preventive health, innovative solutions and person-centred programmes tailored to every life stage. Our integrated health system, which spans primary care, hospitals and national specialty centres, includes Tan Tock Seng Hospital, Khoo Teck Puat Hospital, Woodlands Hospital, Yishun Community Hospital, NHG Polyclinics, the Institute of Mental Health, National Skin Centre and the National Centre for Infectious Diseases.

Together with academic and industry partners, we advance medical education, research and healthcare innovation in Singapore, addressing areas that are critical to Singapore's population needs. As the Regional Health Manager for Central and North Singapore, NHG Health partners general practices and health and social care agencies to ensure the physical, mental and social well-being of residents in the community.

Together, we are committed to building healthier and resilient communities, and Adding Years of Healthy Life to the people we serve.

More information at www.nhghealth.com.sg.

Follow us on Facebook | Instagram | LinkedIn | TikTok













Nelson MAndela



18 Jln Tan Tock Seng, Singapore 308443 E: chi@nhghealth.com.sg W: www.chi.sg