

The AI Adoption Challenge in Medicine: A Never-ending Quest for Clinical Progress

Text and photo by Dr Charlene Liew



FEATURE

Not long ago, I had the privilege of visiting the Prado Museum in Madrid, Spain where I stumbled upon an ethereal, glowing masterpiece hanging on a massive canvas nearly twice my height. Almost hidden in the shadows of a dark corner of the museum, a solitary protagonist bathed in golden light, Sisyphus was hauling a boulder up the side of a mountain. I wondered why the museum would seemingly conceal such a spectacular work of art, one among the finest of all paintings I have ever seen. After all, the best masterpieces are typically placed in large viewing galleries, positioned front and centre; yet no benches were nearby for one to sit and contemplate this work by the master painter Titian. Perhaps we live in a world where Sisyphus metaphors are verboten, banished from sight. Perhaps we have become comfortable sitting where we are.

The impact of technology is often unanticipated

Roy Amara was an American scientist, a futurist and the president of the Institute for the Future, a non-profit think tank based out of Palo Alto, California – a suburb better known as the birthplace of technology companies such as Hewlett-Packard, Google, Apple, Facebook, Tesla and PayPal. He is credited for coining Amara's Law which states: "We tend to overestimate the effect of a technology in the short run and underestimate the effect in the long run."

As a child of the 1980s, I grew up watching cartoons like *The Jetsons* on network television and witnessed firsthand the 1990s' technological transformation of every conceivable aspect of life. My first encounter with the World Wide Web occurred at a computer fair where an innocuous-looking personal computer at a booth was wired up to the 28.8 kilobytes per second (Kbps) dial-up Internet, and I typed in a Netscape search bar to bring up *The X-Files*' fan page, which took

about a minute to load. I was instantly enchanted, swallowed into the rabbit hole, irreversibly ensconced in a world of unlimited possibilities – one where having almost instantaneous access to information at my literal fingertips had become reality.

This luxury was not immediately affordable, as I came to learn later. And as luck would have it, I had to wait for more than three years before my family had our own Internet connection, blazing at speeds of 56.6 Kbps. Life (and our phone bills) as we knew it was never the same again. The pace of technological change has been charging on ever since, along with the Internet boom and bust, and the dot-com bubble bursting less than half a decade later in March 2000.

The Cambrian explosion

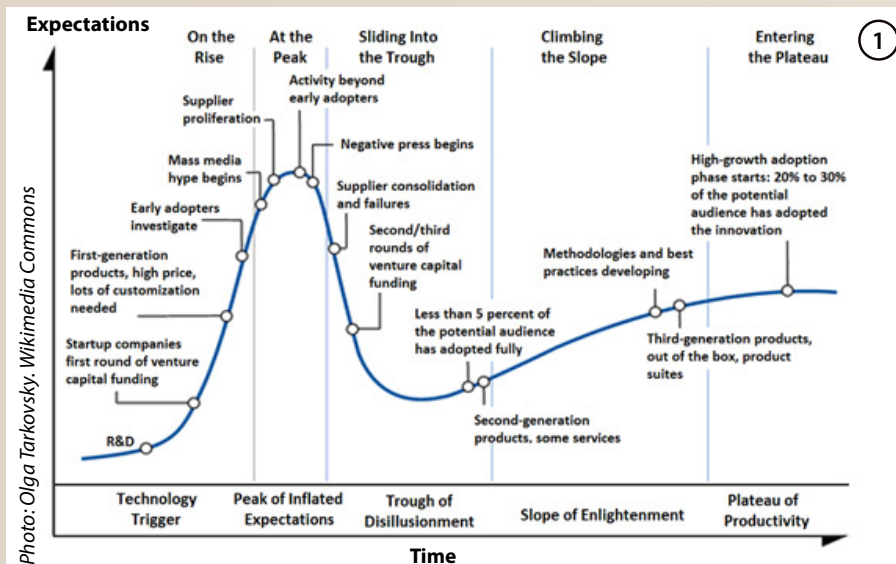
Roy Amara never lived to see the rise of electric vehicles (EVs) or artificial intelligence (AI) but he would recognise his law being played out in each of the EV or AI storylines, or "hype cycles" as the technology consulting firm Gartner likes to call these.

According to Gartner, a technological product cycle involves several stages, beginning with a "technology trigger".

The exact timeline for the current resurgence of AI is difficult to pinpoint, but my personal inclination is to bookmark the year 2005, the year Geoffrey Hinton achieved his mathematical breakthrough that allowed unsupervised training of deep neural networks. Before Hinton's eureka moment, the hidden layers of neural networks involved human supervision, which limited AI models' performance. Shortly after, Geoffrey Hinton and his collaborators coined the term "deep learning" and introduced the use of graphical processing units as engines to drive deep learning computing processes.

Cats and the birth of computer vision

Two years later in 2007, one of the leading AI scientists, Dr Fei-Fei Li and her collaborators launched the ImageNet dataset, a large-scale visual database of 14 million images that became a critical resource for driving progress in deep learning for computer vision. By the summer of 2012, Google had linked 16,000 computer processors, connected them to the Internet and observed as the machines taught themselves how to identify cats by watching millions of YouTube videos. By 2017, Dr Li was the



vice president at Google and the chief scientist of AI at Google Cloud.

It was one of my career's highlights meeting Prof Fei-Fei Li at the Ministry of Health Office for Healthcare Transformation headquarters in 2021 for a closed-door meeting. She listened intently to our plans to build a national AI medical platform for imaging. Her eyes sparkled with inquisitiveness as she asked about our project, and we revelled in the abundant possibilities of harnessing clean, well-labelled datasets to train ever more robust AI models at a national level. Even more intriguing was the ability to deploy AI models at scale, unlocking the promise of AI in improving clinical outcomes.

In July 2023, AimSG, the AI medical imaging platform for Singapore public healthcare, was launched. AimSG allowed us to deploy commercially available AI models in our daily practice of radiology, where they are integrated into our picture archiving and communication system (PACS) viewing systems. We are still assessing its cost-effectiveness, although anecdotally, its value has been demonstrated to be able to save up to 70% of the time taken to create a radiology report in certain use cases. Not surprisingly, most of the productivity gains are realised by automating the drawing of borders (segmentation), a previously manual task, and by automating the generation of reports from simple clicks confirming AI-detected abnormalities. These models are “human-in-the-loop” systems, where expert input is required to generate an output.

The peak of inflated expectations and the trough of disillusionment

Ever the cautious pessimist, we often remind ourselves to temper our expectations of miraculous success in deploying AI models in clinical practice. Indeed, terms such as the technological “valley of death” which needs to be bridged, crossed or jumped over are often bandied about to warn those who bring new innovations into the real world.

Risk management becomes the overarching focus: careful planning;

scanning the marketplace; and challenging ourselves by questioning every possible assumption about how AI would work, work below expectations, fail overtly and potentially fail covertly without us realising until much later. All conceivable safeguards have been put in place, at the risk of raising the cost of implementation. We have over-invested in safety, for in failing early in healthcare, the penalty may be a scorched earth, a “never again” barren wasteland.

Are there any other pathways to frugal innovation? These other options exist surely, but we must first pass them through a sieve built upon the threads of ethical principles – *Primum non nocere*: first, we must do no harm. Connecting public healthcare information systems to AI models requires enterprise-level, defence-grade security, private cloud, and high-availability systems, all of which incur a hefty price tag.

Bringing along everyone for the ride

Where will we be in the next 20 years? A quote attributed to many famous individuals, from quantum physicist Niels Bohr to legendary baseball catcher Yogi Berra, imparts: “It is difficult to make predictions, especially about the future.” And yet, our founding fathers took it upon themselves to imprint upon our national mindset that it is possible to achieve a precise prophecy by engineering the future precisely.

By its very nature, the adoption of technology in healthcare involves everybody. It has never been a difficult task to get a few doctors to adopt a new technology; these early adopters eagerly set upon the job of getting busy trying new things and innovating on a small scale. The challenge comes from getting

everyone involved, from the junior staff and trainees to the most senior clinicians, to adopt new technology that has been inconspicuously named “artificial” into their daily working lives. These workflows include the participation of para-medical staff, healthcare assistants, nurses and other healthcare professionals who may be equally discombobulated by the notion of the ever-pervasive AI.

We partly have GPT-4 to thank for that. Generative language transformers entered our lives just as commercial medical-grade AI models started producing noticeable excitement among clinical users. As of 2022, there were over 500 AI models licensed by the United States Food and Drug Administration available commercially for healthcare. Large language models (LLMs) fine-tuned for healthcare use will catalyse the development of hundreds more commercially available AI models. Many of these LLMs will address digital data-entry and data-retrieval fatigue at the point of care. Ambient AI will act as invisible scribes, sitting in on patient-doctor consultations to transcribe notes into the electronic medical record (EMR), while contextual voice-activated searches for relevant EMR notes will occur through AI assistants as casually as we prompt our smart devices to show us the way to the nearest petrol station.

For those of us long-suffering techno-optimists bemoaning the lost dream of the age of self-flying cars like *The Jetsons*, this change cannot come about soon enough. If this is what it means to follow in the footsteps of Sisypheus, then count me in. ♦

Legend

1. The Gartner hype cycle
2. The launch of AimSG on 27 July 2023, with Synapse partners

Dr Liew is the deputy chief medical informatics officer at Changi General Hospital. She co-founded the Artificial Intelligence and Informatics section of the Singapore Radiological Society. She hopes she has made a prescient choice to be at the digital coalface that she may help to shape the progress of medicine's future.

