


**There's an AI for That.  
And that's Exactly the Problem**

---

**A/Prof. Adam CHEE**



**Adam CHEE**  
Practitioner Professor  
in AI-Enabled Health  
Transformation

## Primary Appointments

- Director, AI Development Office  
Singapore General Hospital
- Adj. Assoc Professor (Primary)  
SSH School of Public Health  
National University of Singapore
- Adj. Assoc Professor (Joint)  
YYL School of Medicine  
National University of Singapore

## (some) Other Academic Appointments


- Honorary Professor (Health Systems  
Transformation), Swansea University
- Visiting Professor (Medical Informatics),  
Taipei Medical University
- Adjunct Faculty (Health Technology  
Management), University of Hong Kong

## (some) Other Appointments

- Co-Director of Datacise Open  
Learning, UK
- Expert (Digital Health), World Health  
Organisation
- Expert (Health Data), Asia-Pacific  
Economic Cooperation (APEC)
- Chair, Health Level Seven (HL7)  
Singapore
- Chair, Technical Committee for  
Healthcare & Health Informatics,  
Singapore Standards Council,  
Enterprise Singapore


# There's an AI for that

- We've all heard this before: **'There's an AI for that'**
- And increasingly... that's where the problem begins
- Fact is, most AI fail shortly after deployment
- And when that happens, we often blame the model
- But in reality, the failure starts much earlier (and it's usually nothing to do with the model)



# There's an AI for that (cont.)

- In many cases, we've already decided on the solution...
- before we've actually understood the problem.



# The lenses I'm wearing today

- Vendors focus on what's possible
  - End users focus on what's usable
- IMPORTANT**
- But when you're running an AI Development Office in the largest public hospital in Singapore, you have a different problem
    - You don't have unlimited time, data, or people
    - Every project your office take on means something else doesn't get done
  - The interesting reality is, **“everything is important, everything is urgent”**

# We're measuring the wrong things

- When we say “AI solution,” do we mean:
  - Prediction?
  - Automation?
  - Decision support?
  - Workflow redesign?
- These are completely different things. Different **data**. Different **risks**. Different **responsibilities**.”
- But we often use one word - ‘AI’, to describe all of them.
- Pilot success ≠ Real success
- Model performance ≠ Clinical impact
- Deployment ≠ Adoption
- Usage ≠ Behaviour change
- We’ve become very good at making AI look successful... without actually changing anything that matters
- (Especially if we are not responsible for it’s success - **Model vs Product vs Solution**)

# We're measuring the wrong things (cont.)

## **Pilot ≠ Real-world success**

- Controlled conditions don't reflect real clinical complexity.
- E.g. A sepsis model performs well in a curated pilot ward but breaks down when deployed across different departments with varied workflows.

## **Model performance ≠ Clinical impact**

- High accuracy doesn't guarantee better decisions or outcomes.
- E.g. A model predicts deterioration with 95% AUC, but clinicians ignore it because it doesn't fit their decision-making process.

# We're measuring the wrong things (cont.)

## **Deployment ≠ Adoption**

- Making a tool available doesn't mean it's actually used.
- E.g. An AI dashboard is integrated into the system, but clinicians rarely open it during routine care.

## **Usage metrics ≠ Behavior change**

- Interaction with a system doesn't mean practice has improved.
- E.g. High click rates on AI alerts, but no reduction in adverse events or change in treatment decisions.

# Solution first. Problem later.

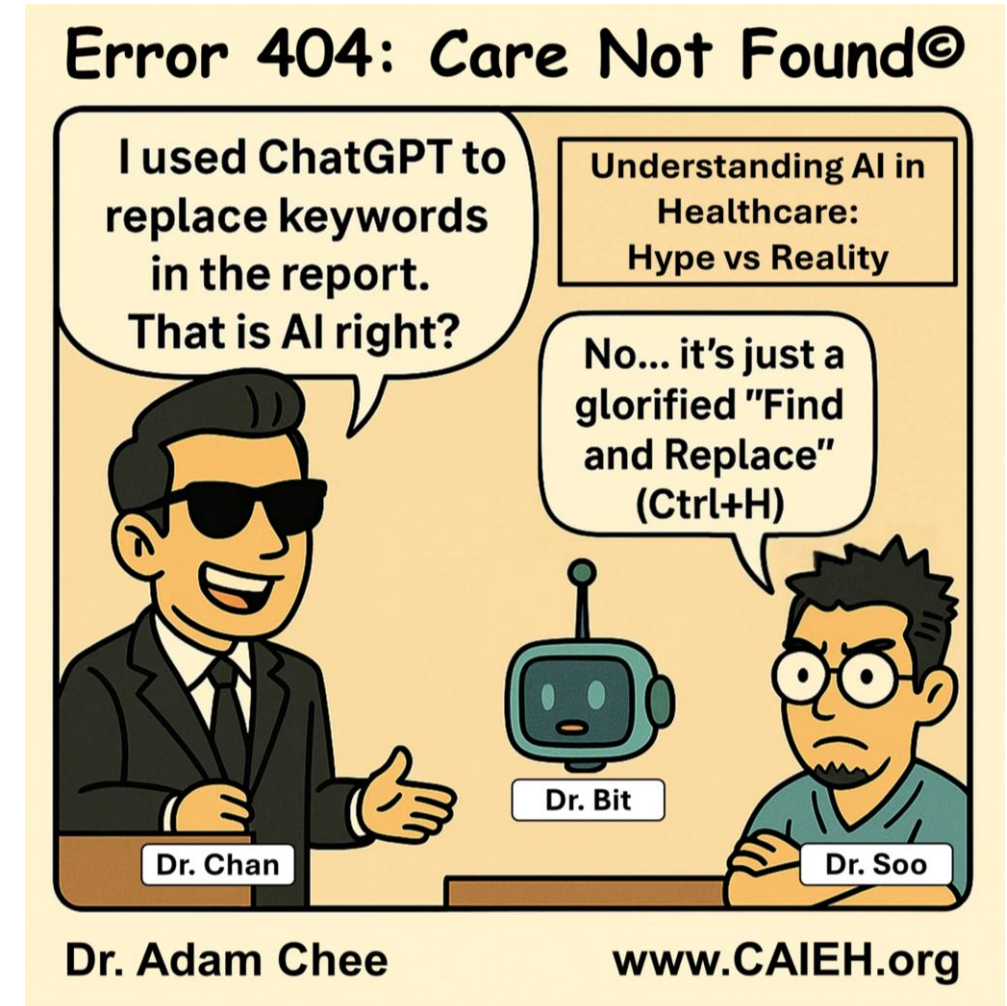
- AI is powerful.  
But that doesn't mean it belongs everywhere.
- The real question is not 'Can we build it?'
- It's 'Should we?'"
- We need to stop deciding on AI as the solution, before we've even understood the problem.
- AI will scale whatever we point it at.  
(assuming you have data)
- The question is whether we're pointing at the right thing.

# The common trap

- People ask for “an AI solution” before they can clearly answer:
  - Who is this for?
  - What unmet need are we solving?
  - What decision, behaviour, or workflow should change?
- Before development, we need clarity on:
  - target users and beneficiaries
  - real requirements, not vague feature requests
  - available data and whether it is good enough
  - clinical workflow fit
  - validation in the real setting
  - integration into existing Clinical Information Systems
  - monitoring after deployment

# Where things Break


- Vague requests (“build an AI for this”)
- Unclear requirements
  - What does success look like?
- Unknown data quality (if any data)
- No workflow integration
- No system integration (UX & UI)
  - Not to mention cost of integration, deployment & maintenance (yearly) etc.
- Building the wrong thing well... is still the wrong thing.



# The hard part isn't technical

- Not every request should become a project
- Leadership must support prioritization decisions
- Triage requires discipline, not popularity
- Saying “no” protects impact
- **Every ‘yes’ comes at the cost of something else**
- **If everything is a priority, nothing is.**
- AI will scale whatever we point it at
- Culture determines what we choose to point at

**There may be AI for many things.  
But not everything needs an AI solution.**



*Thank  
you!*