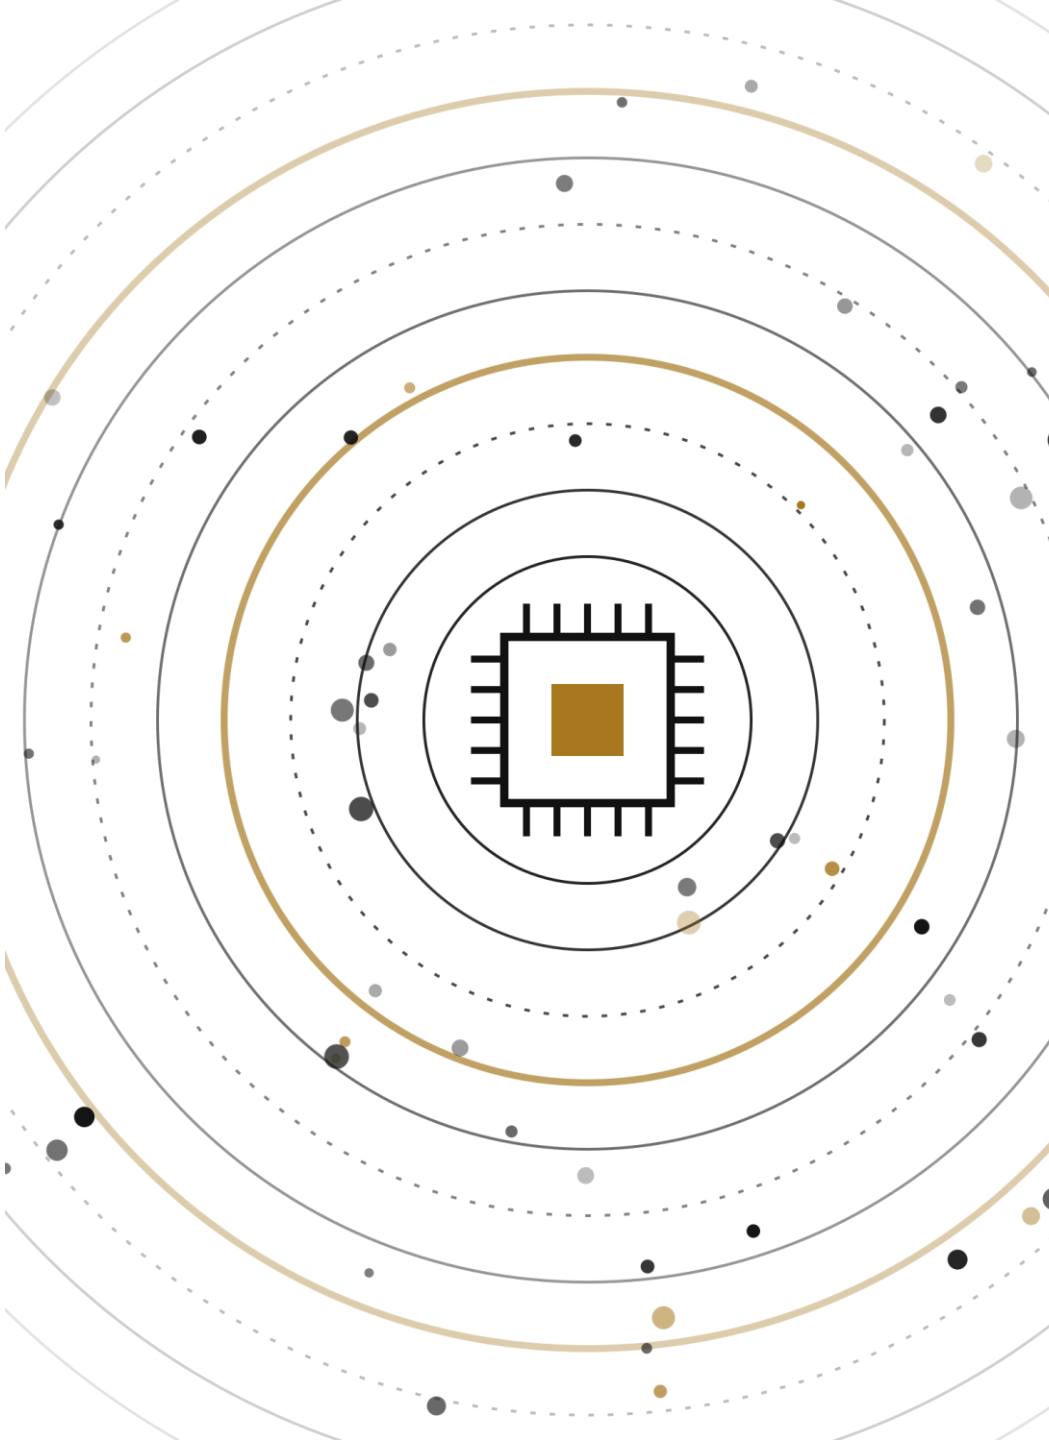
MODULE 5 · CRITICAL THINKING

The Confident Machine

AI Limitations, Hallucinations and Bias

Dr Olusola Ayoola
Founder and CEO, Robotics and Artificial Intelligence Nigeria (RAIN)

AI Training for Academic Staff · College of Medicine, University of Ibadan · 16 October 2026



Thank you.

To the Provost and the Principal Officers of the College of Medicine, University of Ibadan, for this invitation.

Even if you don't use AI, AI will use you.

So we must be wary. And we must be aware.

A partner we cannot deny



Education



Research



Practice



Solutions

Every profession must agree its own standards for AI.

Why we built it

Easier

lives

Faster

progress

Leapfrog

old barriers

Optimise

scarce resources

AI only knows what it is given



Text



Images



Audio



Patterns

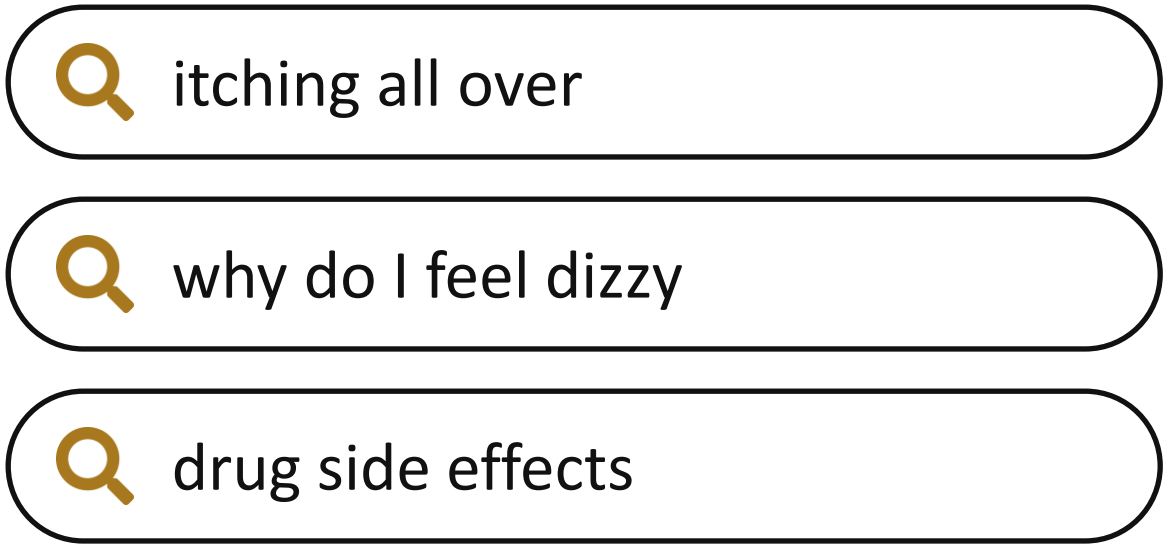
No intuition. No conscience. No lived experience.

It cannot reflect on its own. It must be fed.

It reads what you did not say

“I’m fine, doctor.”

What the patient said



What the search history said

Patterns are data too.

It learns from its failures

Predict → Fail → Look deeper → **Predict better**

50 : 50 → 60 : 40

Data it forgot to collect. Conditions it forgot to observe.

Speed is not depth

Seconds

AI reads more in seconds than we read in hours.

Reflection

We read, then we reflect, and bring a whole life to the page.

So what are the limitations?

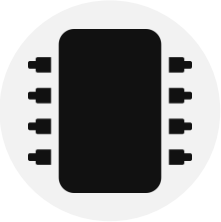
01 Processing

02 Time and place

03 Hallucination

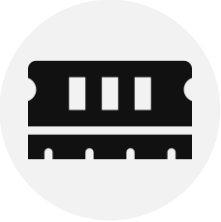
04 Bias

01 Processing has a ceiling



Speed is bound by hardware

CPU or GPU: some machines are far faster than others



Memory has a limit

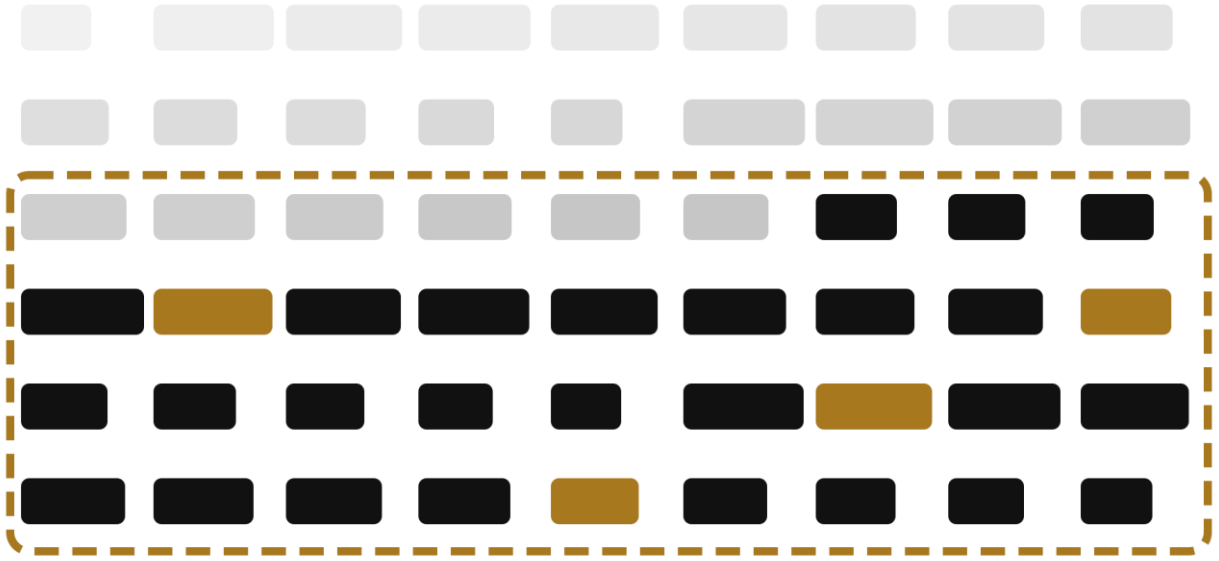
Some data is simply bigger than the machine can hold



Past the limit, it forgets

And it may not tell you that it has forgotten

It forgets the beginning



The dashed frame is what it can still see: **the context window.**

- Long chats push early facts out
- An honest AI asks you to repeat
- A silent AI builds on the gaps

Definition: **An AI LLM Context Window** is the working memory of a large language model (LLM), measuring the maximum amount of text (in tokens) the model can process and remember at any single time.

02 Knowledge has borders

Fever in Ibadan

Malaria is a leading suspect from the start.

Fever in Boston

Malaria is rare, and seen mostly in returning travellers.

A model trained elsewhere thinks like elsewhere.

Your patients are not the data



Haemoglobin C is West African, yet we rarely test for it.

Abroad, every newborn is screened, so C is counted. AI knows only what was measured and recorded.

Humans decide which borders apply



A Nigerian court is not bound by judgments made under another country's laws.

- Medicine has borders too
- AI must be told which data applies
- Each field must agree its limits

Eager to please.

When it does not know, it may **perform knowing**, until you correct it.

Not a theory. It has happened.

47%

of 115 references a certain AI LLM platform gave for medical papers were fabricated. Only 7% were real and accurate.

Bhattacharyya et al., Cureus, 2023

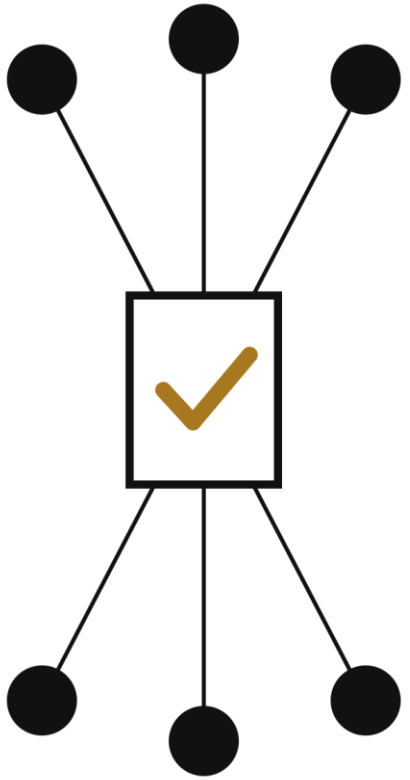
6

court cases invented by that AI LLM platform and filed by New York lawyers, who were then fined by the judge.

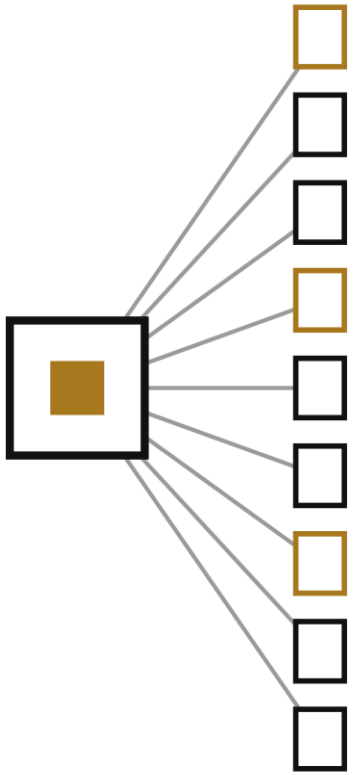
Mata v. Avianca, US District Court, 2023

Nobody appointed AI as judge

- Experts agree what is upheld
- AI weighs every source alike
- Search is not judgement

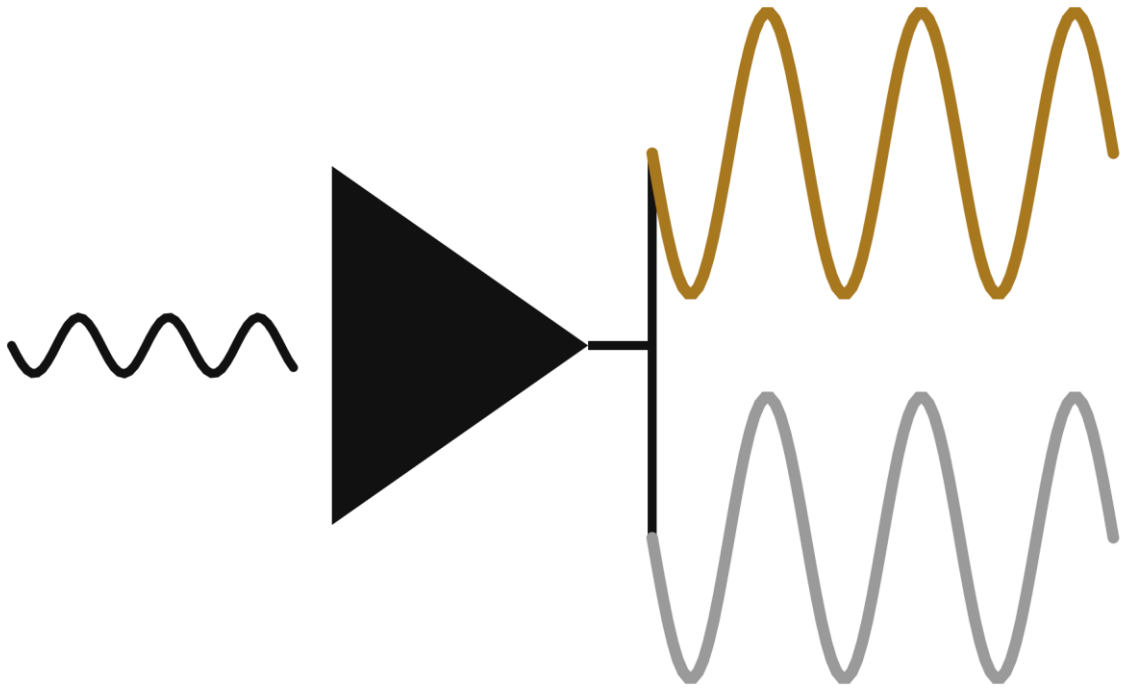


Peer review



AI

AI is an amplifier



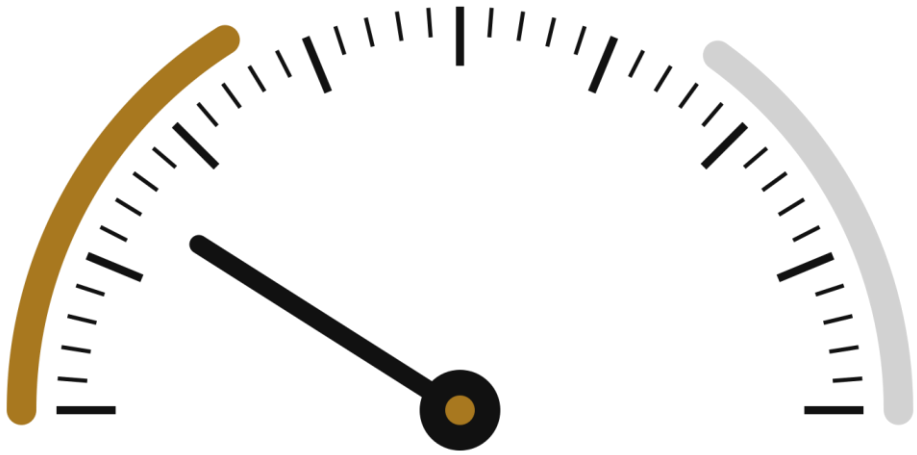
Excellence

Question. Challenge. Verify.

Laziness

Accept. Paste. Move on.

Turn down the temperature



"Your blood dey hot today."

- High: more invention
- Low: closer to the evidence
- Lower is safer, not safe

Definition: **AI LLM Temperature** is a configuration parameter that controls the randomness and creativity of text generated by large language models (LLMs). It does not change what the model knows, but rather how it chooses the next word or token from a probability distribution.

04 Bias in, bias out



Sources

Where did the data come from?



Integrity

Can those sources be trusted?

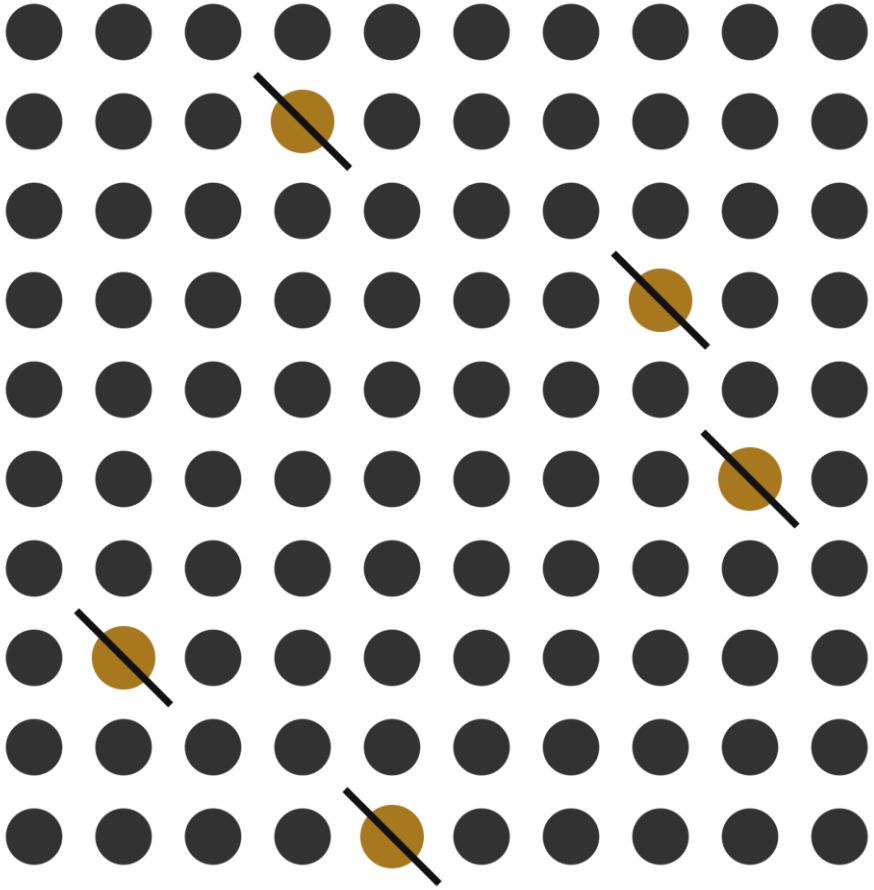


Balance

Who was counted, and who was left out?

No data from your region means no knowledge of your region.

95% accurate. 0% useful.



100 patients. 5 are sick.

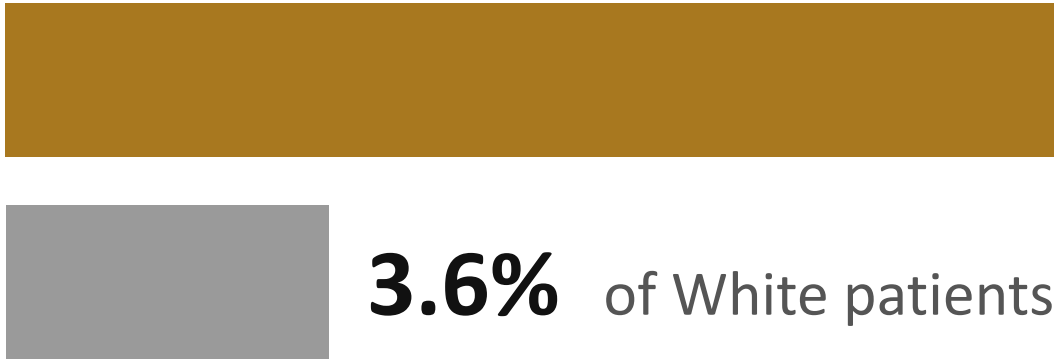
A model that calls everyone healthy scores 95%, and misses every sick patient.

Always ask:

How did it perform for each class?

Bias is already in the clinic

Same pulse oximeter reading of 92 to 96%. Dangerously low blood oxygen was missed in:



11.7% of Black patients

3.6% of White patients

If a device can be biased, so can the AI trained on its readings.

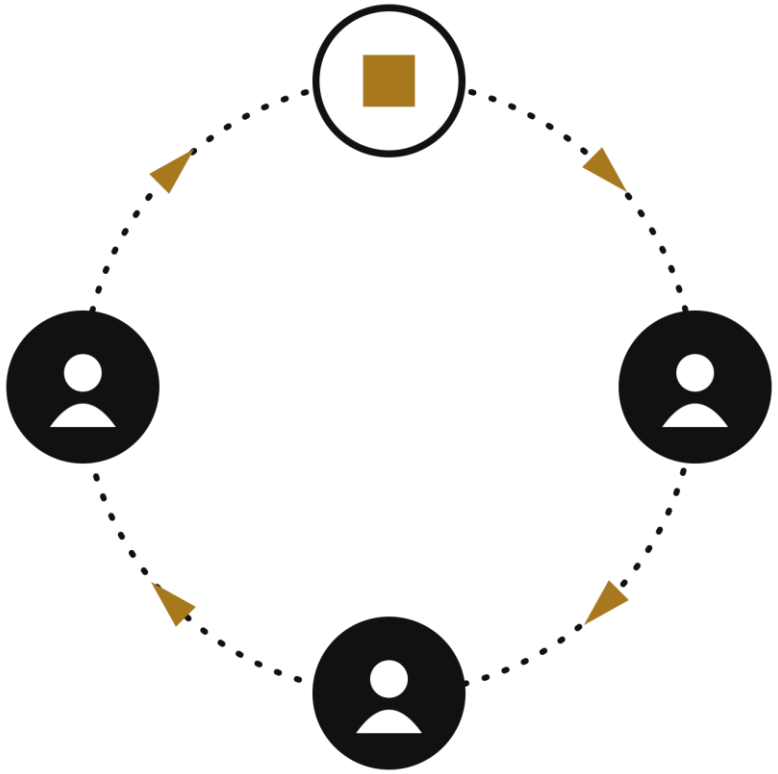
Sjoding et al., New England Journal of Medicine, 2020

Keep the human in the loop

Verify
every fact and every citation

Validate
against your own expertise

Train
before you trust



A knife in trained hands heals.

As AI advances and becomes unavoidable, professionals must learn how to

Relate. Rely. Respond.

With AI **On AI** **to AI**

Thank you.

Dr Olusola Ayoola · Robotics and Artificial Intelligence Nigeria (RAIN)