
People's AI Stewardship Summit

London, 14th July, 2025



Scribe Capture Document

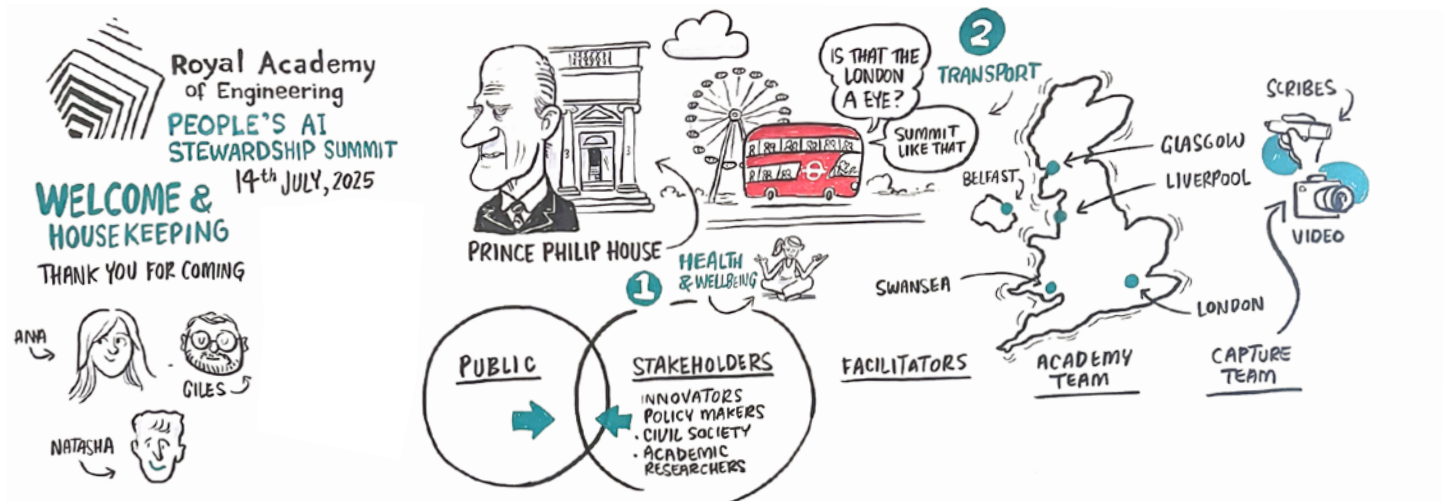
Captured by

We are Cognitive

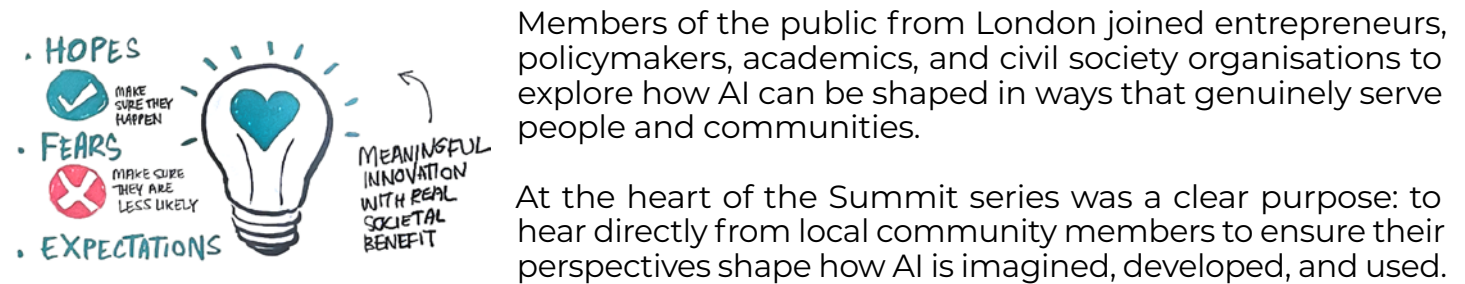


**Royal Academy
of Engineering**

Introduction




On 14th July 2025, the Royal Academy of Engineering (or “the Academy”) hosted the final event in its People’s AI Stewardship Summit series. This ended a UK-wide journey via Belfast, Glasgow, Liverpool, and Swansea.



The London event focused on two themes with relevance to the city: health and transportation.

About the Academy

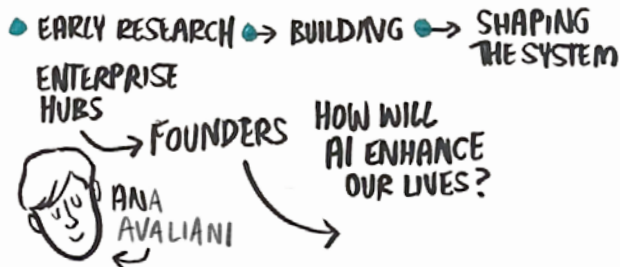
The day began with a welcome from Dr Natasha McCarthy, Associate Director of Policy at the Academy. She introduced the Academy as a fellowship of some of the UK’s most accomplished engineers, including people behind the chips in most smartphones, major infrastructure projects, medical breakthroughs, and creative feats of design. The Academy also supports research and advises the government on engineering and technology policy.



A Note from the Enterprise Hub

Ana Avaliani, Director of Enterprise, offered a brief welcome on behalf of the Enterprise Hub.

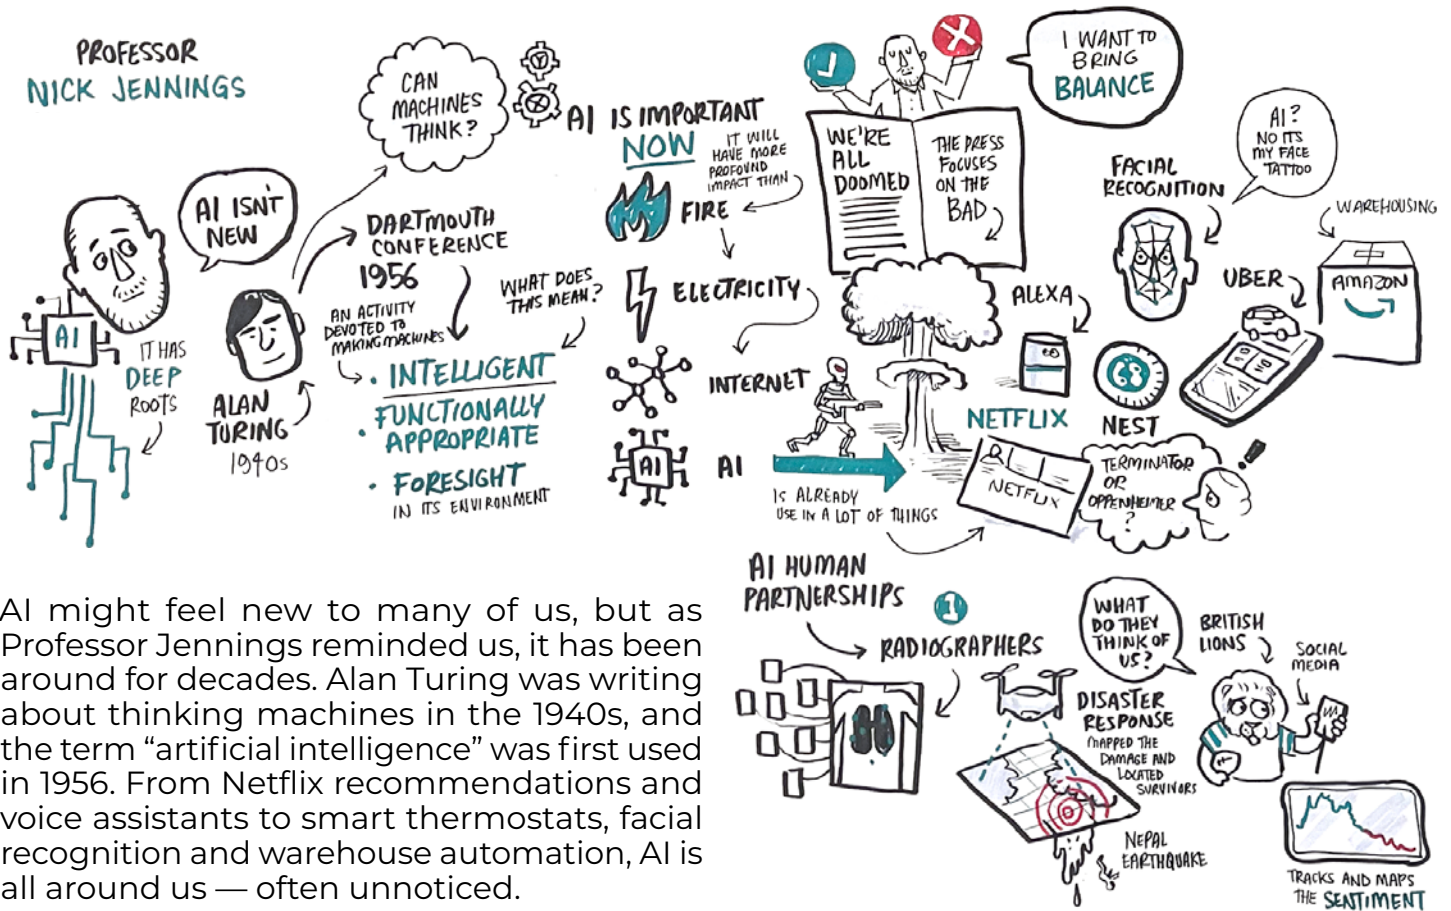
She introduced the Hub as a space that backs some of the UK’s best and brightest engineering entrepreneurs with funding, mentoring and expert support, helping turn ideas into real-world solutions.



"These summits are a fantastic opportunity for our founders to engage with the end users of the technology."

Setting the Scene: On Artificial Intelligence

To open the day's discussions, we heard from Professor Nick Jennings, Vice Chancellor of Loughborough University, a leading expert in AI, autonomous systems, and cybersecurity.



AI might feel new to many of us, but as Professor Jennings reminded us, it has been around for decades. Alan Turing was writing about thinking machines in the 1940s, and the term "artificial intelligence" was first used in 1956. From Netflix recommendations and voice assistants to smart thermostats, facial recognition and warehouse automation, AI is all around us — often unnoticed.

Professor Jennings encouraged us to move beyond the headlines. While concerns about jobs and misuse are real, AI also holds enormous potential for public benefit. He shared examples from his own work, including using AI to support disaster response in Nepal and helping the British & Irish Lions rugby team track and respond to social media in real time.

He left us with three takeaways:

- AI lets us do things we couldn't do before, for good or bad
- It is already central to solving complex problems
- Its greatest strength lies in collaboration with people



Q&A Highlights



The discussion that followed surfaced thoughtful reflections on how AI is both represented and governed.

One participant drew attention to the white humanoid robots shown in the presentation slides, noting how common and potentially misleading this imagery can be. Professor Jennings explained that the images were generated using ChatGPT and included deliberately to prompt reflection on the kinds of assumptions that often shape how AI is portrayed and perceived.

Another contributor, speaking from a background in AI regulation, highlighted that the risks of AI go far beyond bias or privacy. From mental health to safety, the harms are already real.

"Anything a human can get wrong, an AI can get wrong faster and at scale."

Professor Jennings agreed, adding that the idea of removing bias entirely is a myth. All systems reflect bias, just as all humans do, which is why critical thinking remains essential.

Interactive Session: Mapping Hopes, Fears and Uncertainties

Next up, groups explored their hopes, fears, and uncertainties, capturing a wide range of insights, ambitions, and concerns.


Health and Wellbeing

Hopes

Faster, more accurate diagnoses came up repeatedly — from "fast and accurate scan reports" and "early identification of illness" to systems trained on more examples than any human could remember. Participants hoped AI could catch the notes that staff miss and help clinicians feel more confident in their decisions.

Many saw AI easing pressure on the NHS by handling admin and triage, reducing wait times and "making simple things fast."

Beyond clinical use, hopes included fairer resource allocation, data-informed decisions, and improved access to care. Some imagined ambitious possibilities — from "preventative health benefits from birth" to "finding cures for all diseases."



Uncertainties

Participants raised questions about how AI would function in real healthcare settings. Some wondered whether it would truly lower NHS costs or simply reduce staff, leading to “robot care homes” and “bedside robots.”

Accountability was a key concern; if something goes wrong, who is responsible?

Questions were also raised about whether AI could match “GP-level understanding,” and how it might affect staff training, junior roles, and diagnostic consistency.

Fears

Misdiagnosis was a common concern, along with the risk that clinicians might “blindly trust AI opinion” without scrutiny.




Some worried that AI, trained only on existing data, might limit medical research, offering no real innovation, just reinforcing what's already known. Others raised concerns about bias in training data, particularly relating to race, gender, sexuality, and weight.

There was discomfort around AI replacing human interaction, especially in mental health or end-of-life care.

Broader fears included poor data governance, vendor lock-in, and the growing power of private companies.

Transport and Infrastructure



Hopes

A major theme was efficiency. People wanted smoother, faster journeys, with AI helping to plan routes, reduce traffic, and improve scheduling by adapting in real time.

Some saw potential for AI to support smarter road planning, detect faults in vehicles, and enable more proactive maintenance of public transport systems. Hopes included self-driving trains, self-repairing infrastructure, and even “superflight” instant transport.

Uncertainties

A key concern was pace: would systems like driverless cars be rolled out too fast, without sufficient testing?


There were also unresolved questions about purpose and need. With more people working from home, how much transport infrastructure is necessary?

Would AI-supported systems be designed with the needs of older people, disabled travellers or women travelling alone in mind? Would rural areas be overlooked in favour of cities?

Fears

Driverless vehicles raised fears of accidents, loss of human oversight, and sabotage. Some worried AI could be misused in human trafficking or identity theft, or used to restrict personal freedom.

Surveillance was another fear, especially facial recognition, and digital IDs in public spaces, with little clarity on who's in control or how that data might be used.



One participant warned of “false optimism,” where AI is sold as a fix-all while infrastructure and services for people are quietly cut. At root, many feared that decisions would be made for efficiency or profit, not for public good.



Expectations

Next, participants were asked what government, industry, academia, and civil society should do to support their hopes for AI in health and address their concerns.

For Health and Wellbeing

Government was expected to act early, regulate clearly, and invest in public education and independent audits. Participants wanted funding not just for new tech, but for better care.

Industry was urged to prioritise safety, transparency, and accessibility.

Academia was seen as a leader in ethics and long-term thinking, with a role in producing accessible, applied research and disclosing funding sources.


Civil society was urged to advocate, question, and elevate underrepresented voices in decisions about AI in healthcare.

Across all sectors, participants emphasised the need for honesty, openness, and public education so people can have conversations about how AI is used in their care.

For Transport and Infrastructure

Government was expected to act fast, upskill itself, and resist corporate capture.

Industry was urged to design with ethics in mind, be transparent about their use of data, and avoid complicity in misuse.



Academia was expected to provide high-quality evidence and treat AI as a social as well as technical challenge.

Civil society was seen as a watchdog and bridge: surfacing local knowledge, challenging poor decisions, and keeping the public's voice in the loop.

There was support for independent regulators “with teeth” — capable of enforcing rules, rectifying harm, and holding both government and industry to account.

Visioning: Positive Futures for AI

Participants created posters sharing their visions for AI in health and transport.

Health and Wellbeing

Building Trust Through Education and Representation

"AI can only be as powerful as our data," raising the risk of failure if underrepresented communities are excluded. One poster imagined the foundation blocks needed for trustworthy AI in healthcare, including *education for young people and equitable research*.

FOUNDATIONAL BLOCK - CAN WE TRUST A SYSTEM THAT USES AI?

- 1. EDUCATION
- 2. EQUITABLE RESEARCH TO PREVENT BIAS

BIAS AWARENESS

Ending the Postcode Lottery


Many envisioned AI helping reduce health inequalities, enabling faster, more precise diagnoses and tackling delays. If built well, an AI-supported NHS could offer fairer outcomes for all, not just the already well-served.

SYSTEM

- END POSTCODE LOTTERY
- ACCESSIBLE TO ALL
- FASTER, MORE PRECISE
- END OF LIFE - SMART CARE HOMES

Tech which supports, not replaces humans

Participants imagined smart care homes that use AI for medication reminders or social prompts but keep human care at the centre.



Empowering Patients Through Knowledge

By providing patients with more knowledge, AI could help patients be more confident contributing to decisions about their care, rebalancing power between patients and doctors.

Public Good Over Private Gain

One group imagined AI surrounded by a "halo of governance," serving people and the planet, not just the powerful.


Healthcare as a Whole System

AI wasn't seen in isolation. Groups spoke about food, social care, and dementia support, and warned against narrow technical fixes. "It's not just the O-ring that crashed," one group reminded us, referencing past failures that had social, cultural, and political as well as technical causes.

Transport and Infrastructure

Human Oversight by Design

"Red buttons" were proposed — literal or symbolic safeguards to ensure people stay in control. Even with self-driving cars or pilotless planes, AI must support human judgment, not override it.



Monitoring, Security, and Independent Checks


AI could play a role in real-time monitoring of infrastructure, emergencies, even mass migration, but must be accountable and auditable.

Transport That Serves People and Planet

AI's potential to create greener, more affordable transport came through strongly. "No vanity space travel," one group wrote. The priority was wellbeing, not spectacle.

Reflections and Open Discussion

Following the poster presentations, participants moved into open discussion.



Some reflected on trust and the gap between what patients understand and what professionals or systems know: "We take a pill without knowing what's in it." What, then, do we need to know about AI?

Others raised concerns about self-diagnosis, unreliable information, and the spread of medical myths, with one participant citing collagen and horse tranquiliser as examples of misguided treatments propagated by AI.

The group debated whether people should be able to opt out of AI. Many supported the choice in principle, but opinions varied on practicality and ethics. Some saw it as essential to autonomy ("I believe in the right to self-determination") while others argued that the common good overrides individual self-interest. One person

quipped that the alternative is to rely on the wisdom and kindness of billionaires.

When asked whether they would trust AI to make an important medical diagnosis, participants were more divided. One person reflected on how their diagnosis had improved when seen by a doctor of the same ethnic background and questioned whether an AI system would recognise their context.

"AI won't know who I am. I don't want it to know who I am, either!"

Others saw value in AI as a screening tool or second opinion alongside human correction.

Expert Insights: AI in Practice

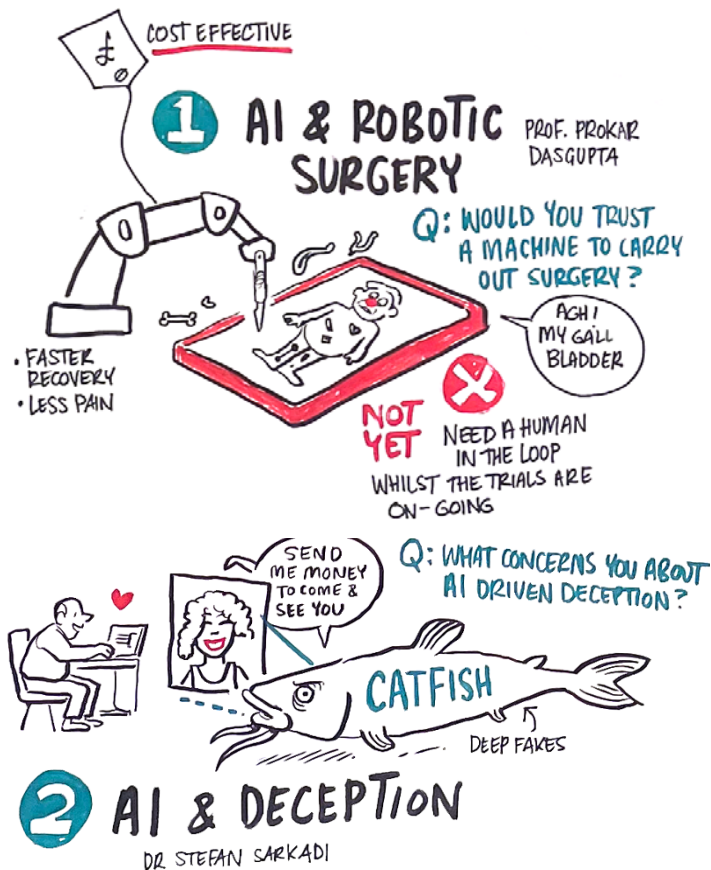
Participants had the chance to learn from experts working at the forefront of AI across different sectors. Each shared a real-world case and posed a provocation for discussion.

AI and Robotic Surgery

Expert: Professor Prokar Dasgupta, King's College London

Provocation:

Robotic surgery is increasingly common in cancer care, with faster recovery and better outcomes. AI could make it even more accurate, but would you trust a machine to operate on you autonomously?



Reflections:

Professor Dasgupta noted that fully autonomous surgery isn't science fiction: just days earlier, a machine had successfully removed a gallbladder from pigs. Still, most participants said, "not yet." They stressed the need for a human in the loop.

AI and Deception

Expert: Dr Stefan Sarkadi, King's College London

Provocation:

AI-driven deception is an urgent societal challenge. While not all forms are harmful (e.g. entertainment or therapeutic use), how should we manage the risks of malicious deception while preserving any benefits of pro-social applications?

Reflections:

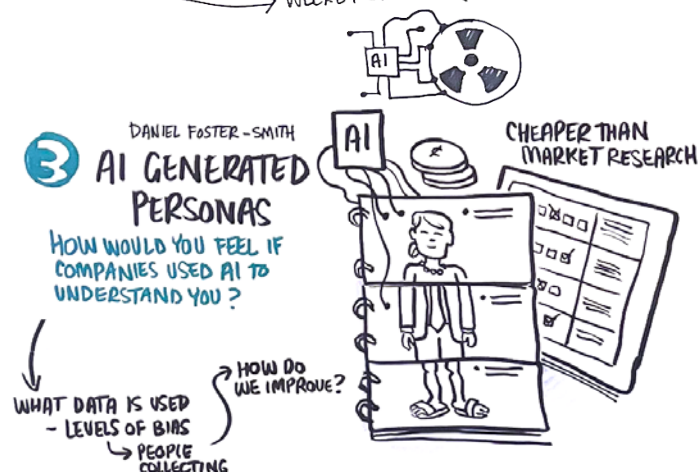
The group felt that deceptive AI is a tool shaped by human intent. While it can be used constructively—such as softening the delivery of serious medical news—they were most concerned about its potential for harm: flooding people with misleading content, fuelling disinformation, or manipulating public opinion.

AI-Generated Personas

Expert: Daniel Foster-Smith, Persona Design

Provocation:

Many companies spend a lot of time and money talking to customers before launching new products, while smaller businesses often can't afford this research.



How would you feel if companies used an AI tool to create realistic profiles using their existing customer data?

Reflections:

The group focused on the risks of bias, not just in the data used, but also in how it's collected and who collects it. They questioned what kind of data would be fed into such systems, and how it could be accounted for.

Medical Imaging and Triage

Expert: Professor Ben Glocker, Imperial College London

Provocation:

Professor Ben Glocker asked whether people would accept fully automated AI triage systems that assess scans, decide next steps, and refer to specialists, especially if they reduced waiting times. Would people want to opt out?

Reflections:

The group focused on trust. People are used to trusting humans, not machines—like trusting product reviews rather than sellers. There was discussion about whether AI could work behind the scenes while a doctor remains the face of care. Some questioned whether we hold AI to higher standards than human clinicians, and whether we even know how accurate human judgement is.

AI in Mental Health Services


Expert: Niamh Roberts (South London and Maudsley NHS Foundation Trust)

Provocation:


Niamh asked how people would feel about a tool that could read a patient's mental health history and suggest treatments. What should it consider, and how should it be used?

Reflections:

The group valued the idea of AI identifying patterns across fragmented systems, such as GP and local authority data, especially when patients see different clinicians over time. But recognised that this posed significant ethical and privacy challenges in implementation.



- TIME TO CARE
- HOW DO YOU SPOT PATTERNS?
- GP DATA, LOCAL AUTHORITY DATA
- AMPLIFYING BIAS



The future of AI in London: Two Questions

Participants split into groups to explore two questions focused on London's AI future.

1. How can we ensure that the benefits from AI developed in London and using local resident data have local benefit?

The group reflected on the scale of data collected across London, from street CCTV to hospital and transport infrastructure systems, and questioned whether this wealth of data is delivering enough value for local communities. They imagined using real-time traffic data to suggest safer routes or analysing environmental data to track pests. Another welcome idea was AI systems detecting rail faults or trees that could fall on tracks before they cause travel disruptions.

1 AI AND LOCAL BENEFIT



2. Should AI use be identified or labelled — and if so, how?

There was strong support for labelling AI use, particularly where it shapes decisions or affects consumers. But the group acknowledged this isn't simple. Should we label all uses, or only when AI is central to the task? It depends on what the AI is doing. And how would this work in practice when many technology providers are global and not incentivised to be transparent?

2 LABELLING AI



Final Activity: What Matters Most?

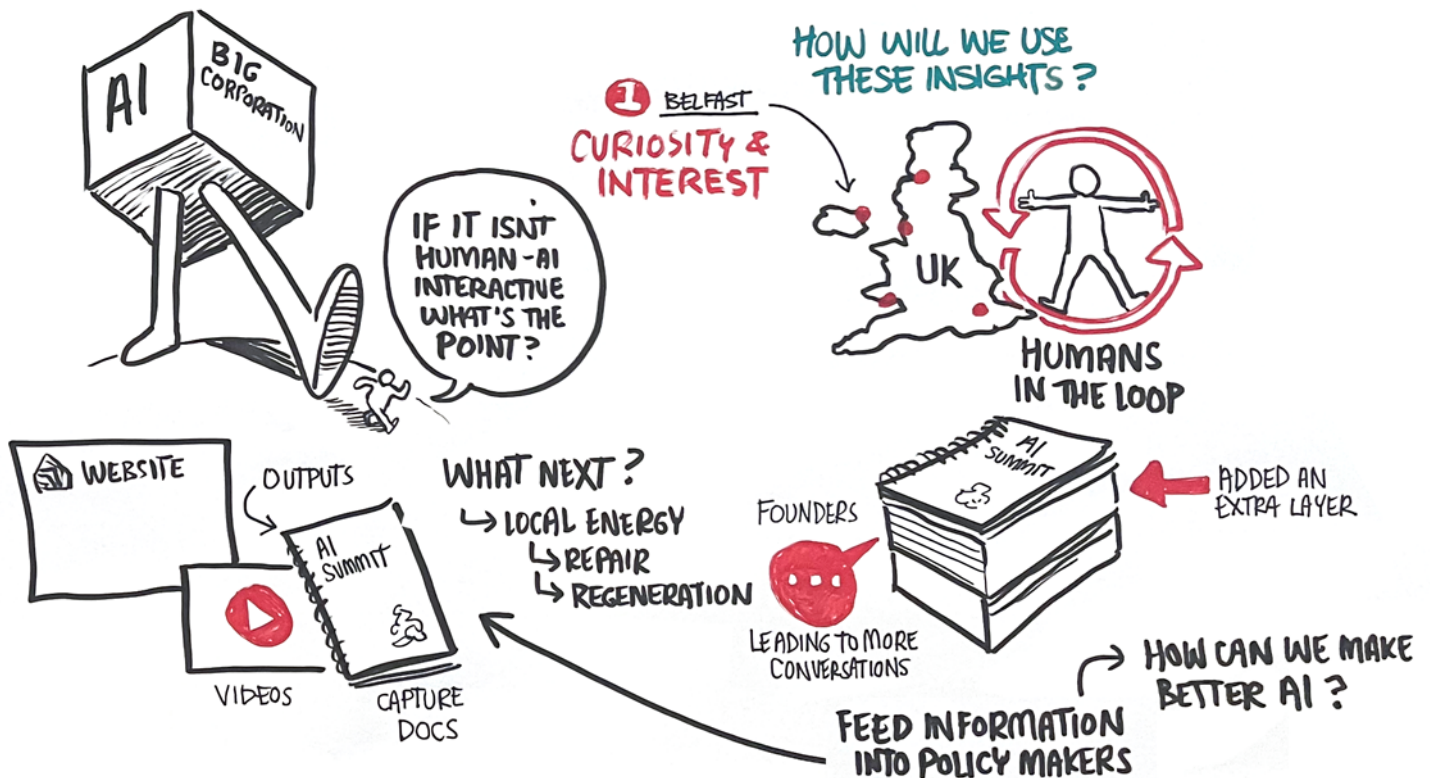
To close the day, the Academy team had generated 8 statements on posters that reflected common sentiments and responses heard across the previous 4 summits. Each participant was given five sticky stars to vote for the statements they felt were most important or exciting about AI to them.

- **Human-AI interaction (52 votes):** A strong desire for AI to support, not replace, human decision-making.
- **Inclusive AI (42 votes):** Participants wanted AI that actively reduces inequalities.
- **Trustworthy AI (36 votes):** People wanted AI systems to explain how they reach conclusions and to be accountable in their workings.

Other popular priorities included transparency, public involvement in design, sustainability, education, and ensuring AI benefits the communities where it is developed. The posters and stars prompted some final reflections on common themes and topics that have flowed across all 5 summits, as well as the growing public awareness of AI and its implications for society.

Wrap-Up: Thank You

Thank you to all participants for being part of this final stop on the People's AI Stewardship Summit tour. Your time, energy, and thoughtfulness have shaped a rich and varied conversation — and it doesn't end here.

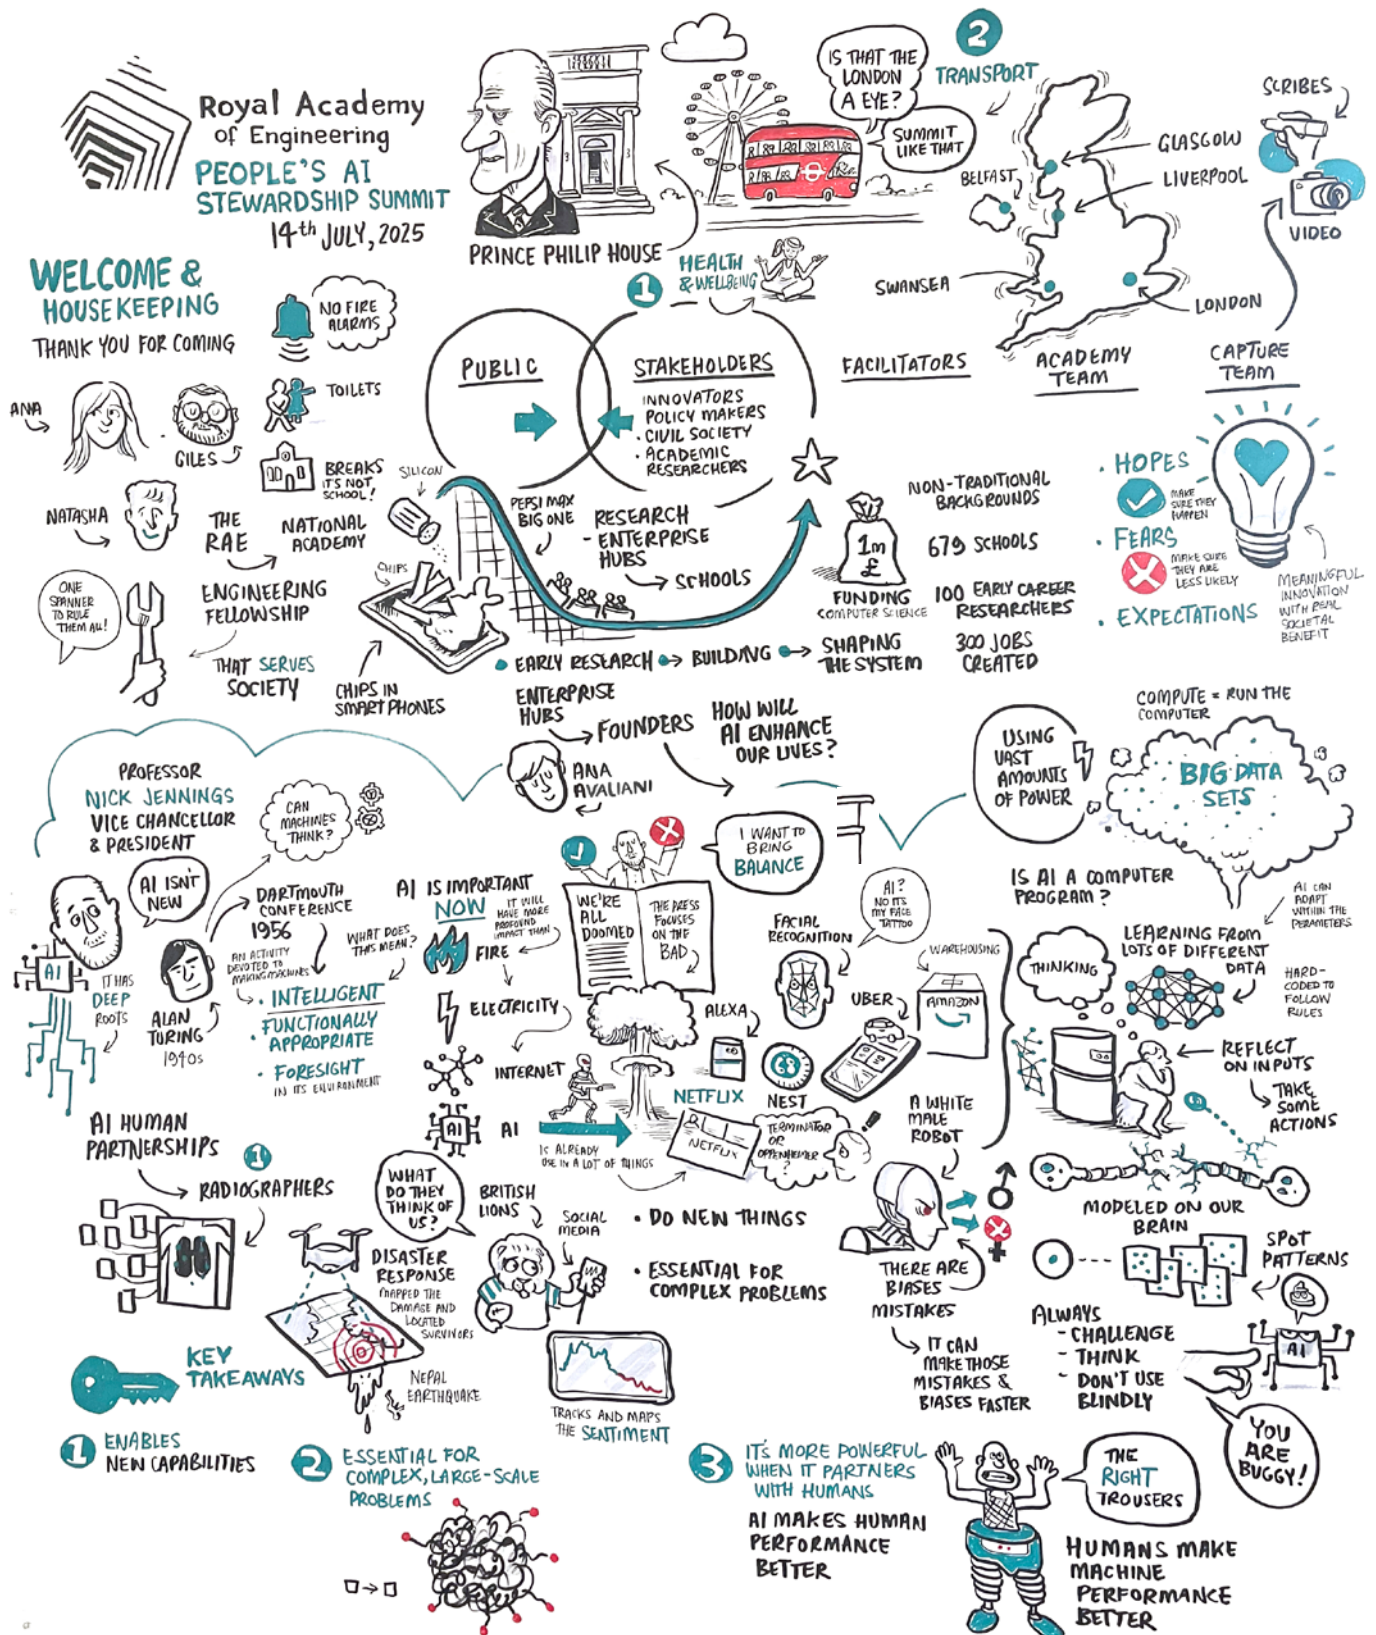


The perspectives which participants shared will directly inform the Academy's work with founders, engineers, policymakers, and the wider AI community. It will help us push for more transparent, inclusive, and people-centred approaches and ensure those who are building these systems hear what matters most to people from communities across the UK.

Captured by
We are Cognitive
www.wearecognitive.com

 **Royal Academy
of Engineering**

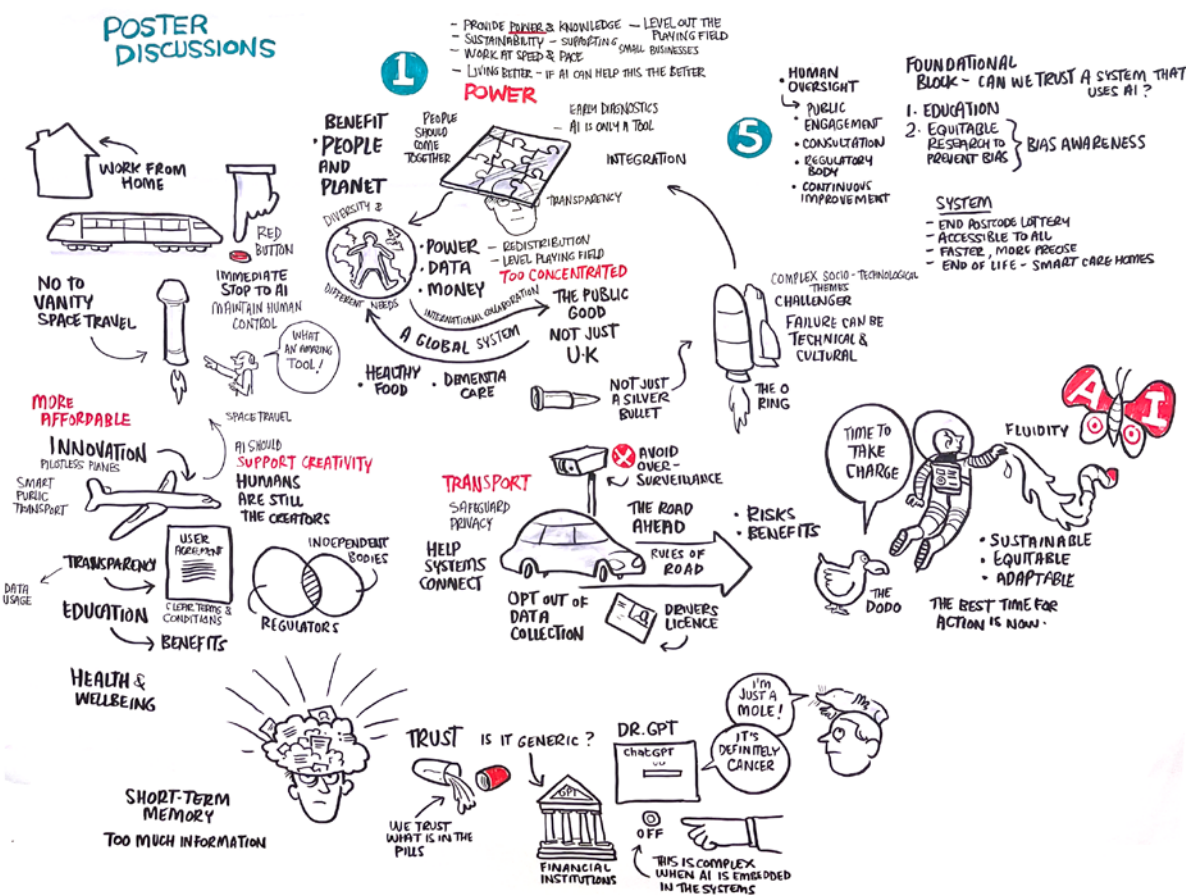
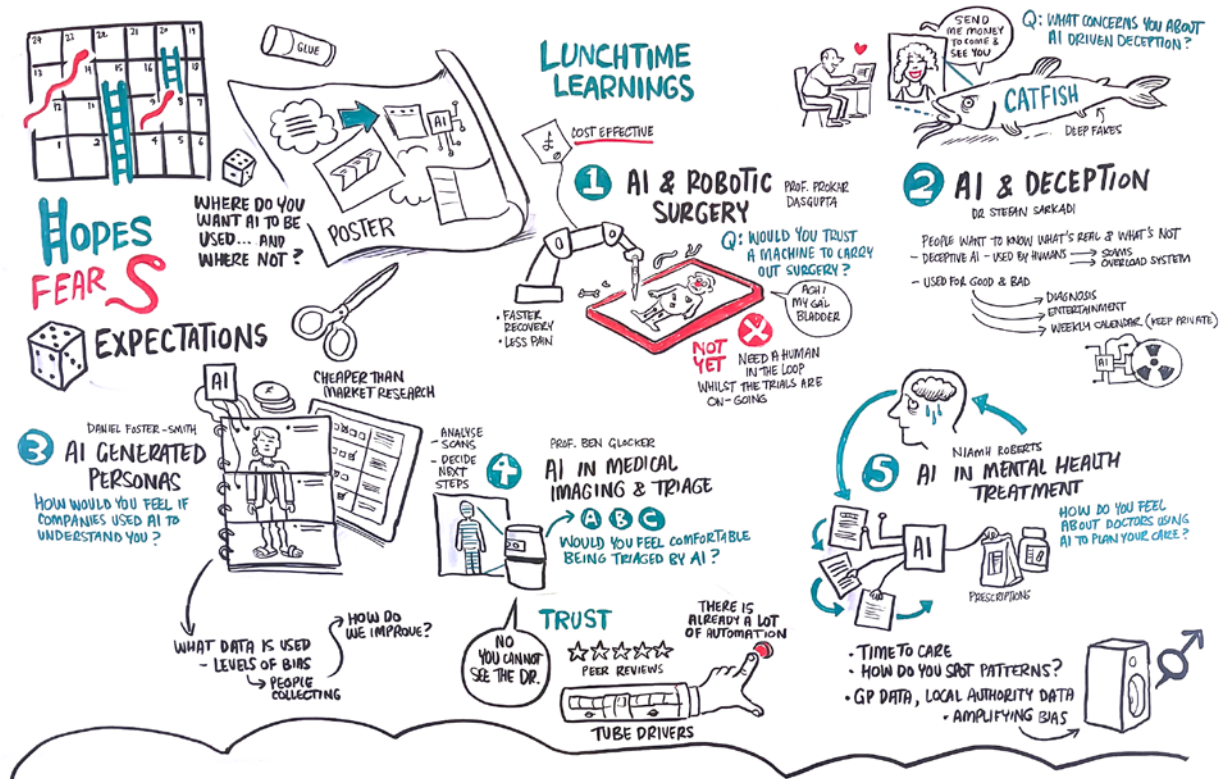
The Big Picture



Captured by:
We are Cognitive
 wearecognitive.com

The People's AI
Stewardship Summit
 July 14th 2025

The Big Picture



Captured by:
We are Cognitive
 wearecognitive.com

**The People's AI
 Stewardship Summit**
 July 14th 2025

The Big Picture

OPEN DISCUSSION

AI AND LOCAL BENEFIT

HOW CAN WE ENSURE THE BENEFITS FROM AI DEVELOPED IN LONDON & USING LOCAL RESIDENT DATA HAS LOCAL BENEFIT?

WHY DID YOU VOTE FOR THIS?

WHO SAYS IT'S FOR LOCAL BENEFIT?


HOW DID WE COMPARE TO THE REST OF UK?



HUMAN-AI INTERACTION

POWER BALANCE

REGULATORY



2 LABELLING AI

SHOULD AI USE BE IDENTIFIED OR LABELLED? WHY? WHY NOT?

CONSUMER CHOICE


HOW DO WE MEASURE?

WHAT WAS AI USED FOR?

SATIRISING: PRODUCTION? HOW IMPORTANT IS THIS TO PEOPLE?



HOW WILL WE USE THESE INSIGHT?



1 BELFAST
CURIOSITY & INTEREST

IF IT ISN'T HUMAN-AI INTERACTIVE WHAT'S THE POINT?




WHAT NEXT?

→ LOCAL ENERGY
→ REPAIR
→ REGENERATION

FOUNDERS

LEADING TO MORE CONVERSATIONS



FEED INFORMATION INTO POLICY MAKERS

→ HOW CAN WE MAKE BETTER AI?

Captured by:
We are Cognitive
wearecognitive.com

**The People's AI
Stewardship Summit**

July 14th 2025