FUTURE OF AI IN

"We're building a world where doctors spend more time looking their patients in the eye," says **LUCY ORR-EWING,** Head of Policy at the Coalition for Health AI.

After receiving a degree in classics from the University of Bristol, Lucy Orr-Ewing forged a deep level of expertise at the intersection of technology and healthcare, among other positions serving as Chief of Staff for Tech Policy at the UK's National Health Service. From there, she was selected for the prestigious Harkness Fellowship, the reciprocal Rhodes Scholarship, and joined Stanford Medicine's faculty researcher, thus launching her presence in the US. Today she is the Policy Lead at the Coalition for Health AI (CHAI). In an interview with Brunswick's Kirsty Cameron, Orr-Ewing spoke about tech optimism, system friction and why AI in health is really about people.

How did you end up working at the intersection of healthcare and AI?

I studied classics at university and it made me fascinated with deciphering people and what they mean, not just what they say. That translated really naturally into healthcare, because the biggest challenges aren't technical. They're people problems: trust, communication, behavior.

I started working in the NHS during COVID. You could clearly see demand rising and clinician capacity falling and the only way to keep up was through tech. That's what pulled me into health AI. I knew that technology wasn't optional, it was the only way forward.

What was your experience like working inside the NHS during the pandemic?

It was intense, but incredibly powerful. I was lucky to work closely with brilliant people, including Catherine Pollard and Tim Ferris, who led tech and transformation for the NHS at the time. It felt like a total war room. We had this moment of urgency where things moved fast. But now, with budget cuts and uncertainty, much of that progress is already being undone.

How would you compare the innovation environment in the UK vs. the US?

During my fellowship at Stanford, I came to see what a moonshot factory the campus facilitates. I made it my business to track down researchers working on huge, visionary ideas. One Stanford professor was 3D-printing human hearts—with $26.3 million in US government funding. He was a Brit, who actually had pitched the same idea to Cambridge and was offered only £40,000. That really showed me how different the pace and scale of innovation is between the UK and US systems. The "anything is possible" legacy of the American Dream really permeates the American appetite and expediency for innovation.

The loop between problem and solution is much tighter in the US. Here, clinicians and industry are in the same clusters, both virtually and geographically, working on solutions to some of the most pressing healthcare challenges. In the UK, the public and private sectors are often speaking different languages and struggle to move in tandem on solving these challenges. The NHS has amazing longitudinal data and relatively high public trust, good patient outcomes and an incredibly talented workforce, but innovation integration is hard. Every trust has thousands of different systems, workflows and financial constraints. Innovation has an impossible pathway to scale.

Let's talk about your work at CHAI. What is the Coalition for Health AI?

CHAI is a nonprofit coalition founded by the health and technology ecosystem to advance responsible AI. We began with incredible support from major US health systems like Stanford Health, Mayo, Duke and Mercy, alongside major tech companies like Microsoft, Google and Amazon. Importantly, we ensured to bring smaller startups to the table as well, including Aidoc, Abridge, Nabla, Ambience and many more. CHAI also works with patient groups, medical societies and policymakers. We intend to

PHOTOGRAPH: COURTESY OF COALITION FOR HEALTH AI

HEALTHCARE

act as the connective tissue between government, providers and industry, building consensus and governance frameworks to help the private sector move faster, the motto of the tech industry, and do no harm, the hippocratic oath of healthcare.

What are some of the tools CHAI is building?

CHAI keeps very busy, building many solutions and tools to help progress the use of AI in healthcare. One major tool is the applied model card—basically a nutrition label for AI. It gives health systems the key facts about an AI model: how it performs, where it works well, where it doesn't. Does it work on women? What about minority populations? The goal is transparency and trust, so doctors know what they're using, and don't feel like they're forced to use something they don't understand. We're also building the first ever model card registry, so health systems can track which AI tools they're using and whether they're still working safely over time. The difference between a scalpel and an algorithm is that the scalpel won't drift to hallucinate or degrade in quality once deployed. That's where CHAI comes into the picture.

What's the biggest risk you see in the adoption of AI in health?

Ironically, I think the biggest risk is not having guardrails. If you remove regulation, doctors and hospitals won't use the tools. They need to know something's been tested, validated and governed. Otherwise, they'll default to risk aversion, which is understandable, given the stakes.

What excites you most right now?

Honestly? The speed. We've had great movement with US government agencies in pushing forward AI priorities. Health agencies like the Centers for Medicare and Medicaid Services, the FDA and the Office of the National Coordinator for Health IT are working together in a way I've never seen before. It's exciting. And internationally, we're expanding CHAI's work to Singapore, Europe, India, Taiwan, South Korea and beyond—building a shared infrastructure for trustworthy AI.

And what keeps you up at night?

I do worry about AGI, the idea of AI evolving into something superintelligent and ungovernable. It sounds sci-fi, but timelines are shrinking. That said, most of the real risk today is about not building intentionally. If we don't design AI to serve underserved communities, for example, they'll be left out entirely.

Any final thoughts on trust and adoption?

We can't rely on tech alone. Trust lives in relationships—in how doctors talk to patients, in how transparent we are about tools, in whether people feel heard. AI should enable that, not replace it.

At CHAI, I'm proud to be working with a team to build a world where doctors spend more time looking their patients in the eye, because tech-enabled tools free them up to focus more on live patient interactions. ◆

KIRSTY CAMERON is an Associate and Digital Specialist in Brunswick's San Francisco office.