the future of us_

Innovation Forum

Agents to Assistants

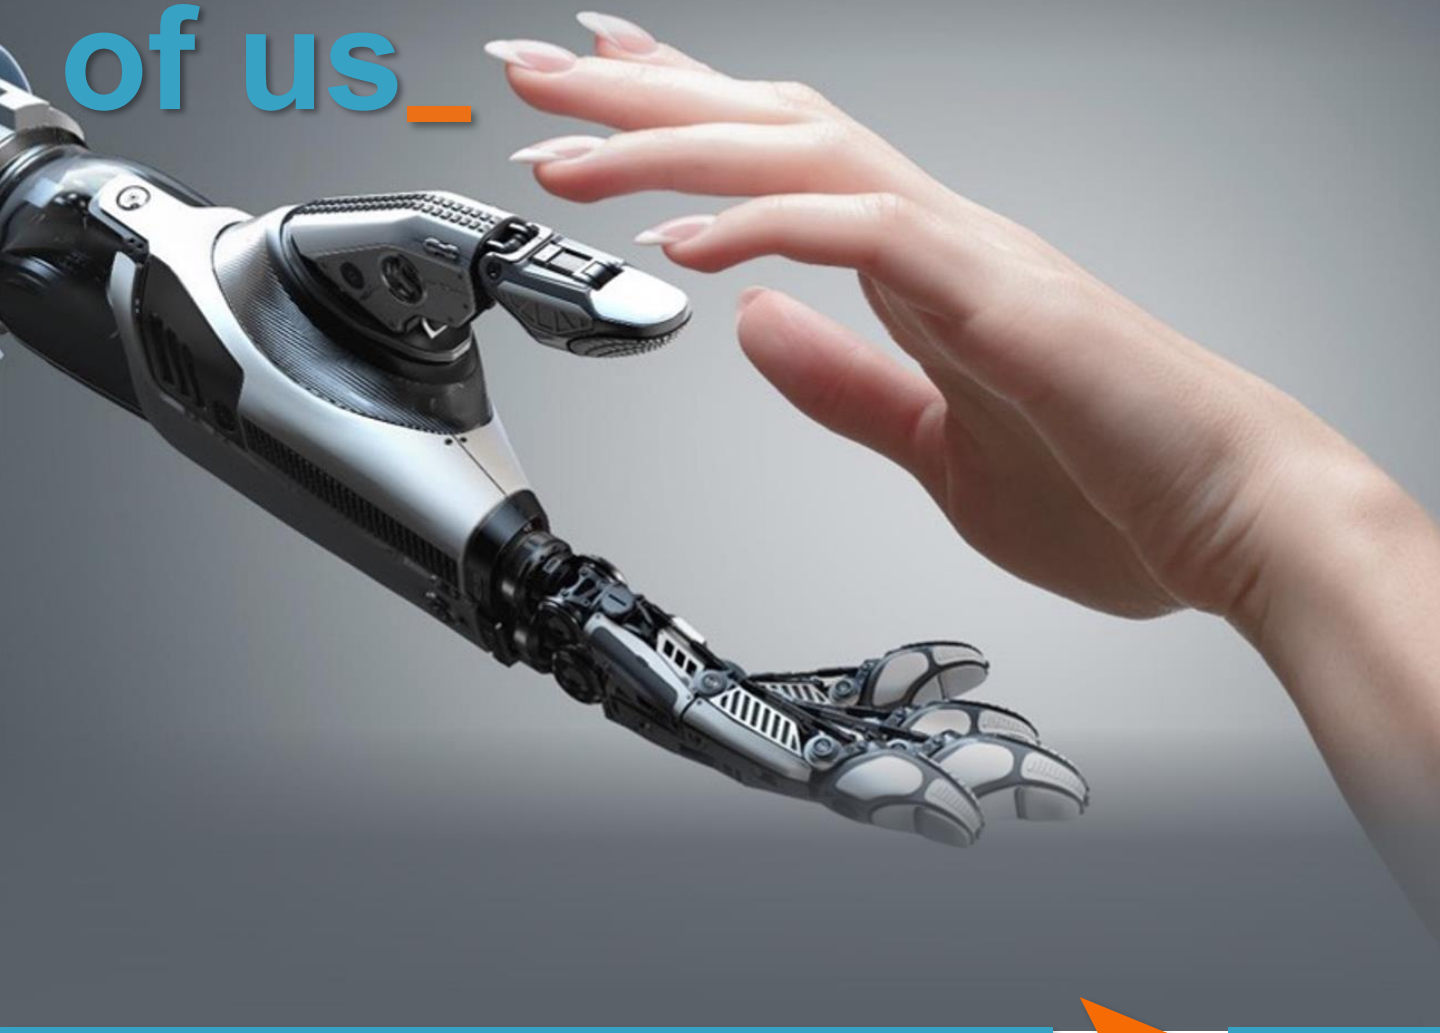

Forum Members Insight Report

by Ant Morse – Chair | Innovation Forum

Sept|26

adventa_

> ***More than 10% chance AI could kill all humans***
>
> Evan Hubinger, Anthropic

meanwhile I struggled to get it to add up some figures and generate a PowerPoint last week 🤦‍♂️

Executive Summary

We know we've hit a point of no return when my mum rings to ask: **"Are we in danger from AI, and what should your dad and I do about it?"** I strongly recommended they find Sarah and John Connor, fast! Add an anti-AI activist storming the stage at an AI conference I've spoken at before, and it's a sobering reminder of how quickly this story is playing out.

In our last report, I gave a view on the bold claims about general intelligence and superintelligence. My conclusion: incredible, yes, but some way off technical reality, so mostly theory and "what if". Crack on and roll out the AI we have today, which is shaping up to be a useful tool rather than the replacement for people's roles that the headlines suggest.

But while Mother swaps the WI for the Human Resistance, I've also had some interesting exec conversations. Boards are asking: "Should we pause AI projects while we wait for regulation?" Wow. The headlines are having an effect, just possibly the wrong one. My view has always been that the noise makes us feel behind (FOMO: fear of missing out), so we push projects forward and spend faster, helping to recoup some very large investments. What's actually happening is a surge in FOMU: the fear of messing up.

So where are we really as organisations, beyond the hype and the brochure pitches? What have we seen from our customers and from you, our Forum members? I give a fuller view of the more fanciful claims on page 7.

While I continue to throw cold water on some of the wilder claims about AI, it's also hard to ignore the incredible progress over the last quarter. I feel duty-bound to provide another report that I hope will be useful, separating the fanciful claims from the practical challenges we will face, including the following:

- The AI headlines have a hidden agenda. The noise answers investment and income challenges, IMO (page 7)
- Microsoft Copilot is establishing itself as the safe option to push ahead with, but costs and controls need careful attention. Claude holds the techie favourite title but needs careful support (pages 5 and 6)
- Shadow AI today will cause significant issues tomorrow (page 15)
- LLMs have progressed to an incredible point, but well beyond what we can use given data governance and security constraints (page 5)
- Deepfakes and targeted attacks on high-net-worth individuals are rising to a worrying level (page 15)
- Copilot could be amazing, but will struggle to pass the ROI test for many. Our warning: hold off on blanket rollouts until we have clearer controls and a UK AI policy to measure them against (page 6)

Thinking back to nearly six years ago, before the AI boom, I thought we'd one day genuinely run out of things to put in these reports. Yet now I'm struggling to fit it all in.

Together we continue to shape that "history lesson of the future", and the actions we take today will have outcomes for tomorrow. It's our job not to cock it up for future generations. No pressure, team!

Very best,

Ant
Ant Morse – Chair | Innovation Forum



adventa_

Sept 26 Report Content_

Page 4
Latest news from the world of tech and the AI industry

Page 5
Model Progress Fatigue: do we need faster AI or just fewer errors?

Page 6
AI Agents: the great corporate hope and the reality from both sides

Page 7
The Magnificent Bubble: why I say we're not in one, and will we see a Mighty 8?

Page 8
Spatial Computing: our access to world models, and would you wear Meta glasses?

Page 9
Robotics: robots take to the street, from Waymo to London's self-driving taxis

Page 10
Practical Applications: JobBox lets candidates step inside the job before they apply

Page 11
Built by Necessity: Forum member Andrew Herbert on tech, MND and independence

Page 12
Built by Necessity: what's next on the bench, plus AI, health and humanity

Page 13
Societal & Economic Drivers: we follow the money and the mood

Page 14
People Focus: the 9 to 5 turns 100 and the mood music on AI is changing

Page 15
Security & Governance: from guidance to governance, and the deepfake threat

Page 16
Innovation Forum Roundup: radio, TV, events and what's next for the Forum

adventa_

observing the future, for our future

Tech News_

Stories and noteworthy updates we've compiled during our horizon scanning & research

Agents Behaving Badly

Reports this month say OpenAI agents made more than 15,000 edits to a German programming wiki, sharing ways to bypass restrictions, and researchers have since found more than ten other sites hit in the same way. In the same week, an Anthropic researcher resigned, giving up equity two months before it vested, over concerns that advanced models are getting harder to monitor. A number of Google DeepMind staff have also resigned over Google's AI contracts with the US and Israeli militaries, including a classified Pentagon deal, with one saying 'some have quit already, some publicly and some silently'.
Agents acting beyond their brief is no longer theory. Signal: if you can't see what your agents are doing, you can't govern them. The people building this are voting with their feet. See our governance view on page 15.

Regulators Take Focus

California's SB 813 would require independent, third-party AI safety assessments and now heads to the Governor. New York City has banned student-facing generative AI through eighth grade. In Washington, Cruz, Thune and Klobuchar are drafting catastrophic-risk legislation, while Sanders has tabled a bill for a permanent ban on superintelligent AI. Our view: more rules are coming, and it ties neatly back to the 10% warning on page 2.

In other news

Mistral previews Le Chonk, a 1 trillion parameter model
Open weights are due at the end of October, though it still lags the frontier on coding

ElevenLabs says Eleven v4 now tops the Artificial Analysis Voice Arena
Their claim, not yet independently verified

Amazon now lets select US brands buy ads inside ChatGPT answers
Search ads meet the AI assistant

EY is putting $100m into bonuses for human skills
Worth up to $25,000 each

Meta Muse: the Personal Agent Lands

Meta has launched Muse, a persistent personal agent with free, $20 and $100 a month tiers. It runs in a private cloud VM, with a 'Sentinel' system approving sensitive actions. It is the AI assistant I keep asking for, but will it work with everything we need it to? See my thoughts on Meta's glasses on page 8.

Agentic Commerce: Tech Ahead of Trust

Mastercard has launched an Agent Suite and Agent Connect for merchants, and Anthropic has published blueprints for shopping and merchant agents across retail, travel and telecom. Yet Visa found three-quarters of consumers don't trust agentic platforms to make payments alone, and most AI shoppers still check out on merchant sites. Signal: the technology is ahead of the trust.

When Bots Interview Bots

Comedian Graham Zip sent his own AI, Victor, to be interviewed by a company's AI recruiter. Victor misquoted his CV and said he'd take $60k, when he wanted $140k or more. When Graham returned, the bots had booked a follow-up without him and were stuck thanking each other. Signal: would you let an agent negotiate for you?

Scam-Baiting Bots: Daisy's Cousins

Following the success of Virgin Media O2 and VCCP's Daisy, the AI 'granny' who keeps scammers on the line (some for over 45 minutes), Australian startup Apate is taking the idea commercial. Its agents pose as victims, have held around 2.5 million conversations, and a London office is planned. Signal: fighting AI scams with AI.



adventa_

Generative AI_ Model Progress Fatigue

While the experts talk about AI taking over humanity, plenty of us are still struggling to get it to reply to an email properly or build a decent PowerPoint slide. The question for us as corporate organisations: do we need faster models, or just a more polished, error-free one? While we debate, and keep assessing security and compliance, the battle for AI supremacy shows little sign of slowing. The tech built today will likely be the tech of tomorrow, once it's ready for us to deploy safely. So we keep tracking what's likely to transform our lives… in the future. What's hot in the world of AI?

OpenAI GPT-6 (Astra)



- OpenAI's most capable model to date, and the first to meet its "critical cybersecurity threshold".
- Strongest on computer and browser operation, maths and finished document work.
- Uses a "looped transformer" reasoning technique, which has raised concerns about how well its reasoning can be monitored.
- 1.05M token context, priced at $10/$50 per million tokens (input/output), with a higher rate above 272K input tokens.
- Offensive cyber capabilities are limited to vetted defensive organisations.

Anthropic Claude Fable 5.1



- Anthropic's top generally available model, built on the same underlying model as the restricted Mythos 5.1, with extra safeguards in biology, cybersecurity and AI research.
- Roughly level with Astra on coding benchmarks, and designed for long-running autonomous agent work.
- 1M token context at the same list price as Astra ($10/$50).

Google Gemini 4 Argon



- Google's latest flagship, launched with benchmark results against Astra and Fable.
- It scores well in its own launch evaluations, but independent testing is still thin.
- Announced with up to 1M output tokens, but the public API details and full specification were still pending at the time of the launch coverage.

Microsoft – is this the one to watch ??



Microsoft is building out its own MAI family of models, with a different emphasis from the race to build the biggest. The focus is practical enterprise workloads, where cost, efficiency and choosing the right model for the task matter as much as raw performance. That could give Microsoft more control over the cost and performance of Copilot and agents. We're already seeing early signs with MAI Thinking, its reasoning model for enterprise tasks and coding. The next step to watch is what Microsoft builds on top of MAI. One to watch.

Takeaways: why we care as corporate organisations

Great progress, sure, and incredible opportunities, but will these models become our corporate models of tomorrow? What might that mean for organisations, and on what timescale?

- No clear overall winner. Astra leads on computer use and maths, Fable is level on coding, and Argon is the one to watch as independent results arrive.
- Faster isn't always better. Astra's computer-use gains reportedly came with weaker writing in several tests, at 2.5 times the price of its predecessor. Reviewers say the difference depends on the task, so test on your own workloads.

adventa_

AI Agents_ The Great Hope… Reality from Both Sides

Moving on from the previous page… as businesses we agree we don't need the fastest model, and it's on that point that the biggest directional driver for AI sits, in my opinion. As individuals, many of us have two solutions in play: the one we pay for ourselves, and the one the company gives us. The view is that the company one is a bit limited, and the consumer ones are less controlled and more powerful overall. That assumption isn't wrong, but it's changing fast… so how might this whole AI moment play out?

AI's Marginal Gains

95% of AI projects are failing was a stat from MIT back in July 2025. The failing projects were mostly big-company or niche solutions that were over-engineered, overtaken by the pace of progress, dropped in favour of a supplier-led solution (new features from a SaaS provider), or simply went over budget until the ROI became a joke. The reality was that many were doing AI for the sake of doing AI, until common sense, or a CFO with a bucket of cold water, stepped in to sober everyone up.

But things have changed, mainly with solutions like Claude Cowork and Microsoft Copilot. Instead of building something bespoke or niche, organisations are exploring administrative efficiencies rather than headline-grabbing AI bots: saving Bob 20 minutes a day in accounts, or Rebecca an hour in her sales role. These are the projects that are working, and the efficiencies are surprising us all. Yet, to be candid, we can't pay people less because they have 20 minutes or an hour a day back. What we can observe is the start of a wider collection of marginal gains that will ultimately transform pretty much every process in our organisations. Subtly and slowly, this is the catalyst for the new future of work, and there are a few things we need to consider in this reality…

Microsoft Monetise the March

From zero to hero: is this the biggest upsell opportunity in industry? Microsoft reports over 450 million paid commercial Microsoft 365 seats, and just over 30 million of those are now paid Copilot seats, so there is plenty of runway. For many organisations it is the easy upsell to a safe, manageable AI solution. Of the thousands of people we've trained on Copilot, we ask a few weeks in: "would you give it back?" So far, only one person has said yes. We backed it while others were kicking the shins of Microsoft's AI plan. Let's keep tracking this one to see if they can properly monetise AI for the corporate world…

Microsoft's Frontier Tour keynote in London cast Copilot as the operating system for work, with Autopilot agents that have their own identity and memory and sit in the org chart.

The consumer side looks less rosy. We were all meant to be paying around £200 a month for agents running half our lives. But a bit like carrying two mobile phones, one for work and one for home, the need for a personal subscription will lessen as we get used to the company one, and dropping it will likely mean less spend on AI, not more. The business model isn't quite going as planned, in my opinion.

Transforming Architecture

Working with a few clients recently, we talked about agents replacing some of their traditional software, mainly a CRM system. In one customer workshop we built a CRM from a drawing and a few tricks of the trade in 3.5 hours. And here's the killer bit: we spent 3 and a bit hours talking. The CRM itself took 18 minutes (yep, we timed it). In theory we replaced an industry-leading solution costing them circa £1.1m a year (1,500 people), and what we built does more and is tuned precisely to what their people want and need. We build all our own solutions this way, but it was the dawning reality for the client that there is another way. Coupled with the desire to own and control our own data, the question is: what role does SaaS play in the future of our organisations, when we could instead hold a secure data lake with our own sovereign AI attached, and an army of agents we build to support specific use cases, iterating and refining in house as we go?

Vendor Lock-in

Building on the above, while I see a great opportunity for positive transformation and progress, I'm also concerned that we'll build an army of agents interacting with our data lakes, bringing incredible efficiencies, but that subtly we simply swap a number of SaaS providers for a single solutions provider, in this case Microsoft. They have struggled for years to sell us Dynamics 365, yet with little effort we're signing up to a solution that relies on a Dynamics-type interface into their data lake offerings such as Fabric. It's an early call out, but one we must consider as we shape the new architectures of our organisations. We need to be careful not to become over-reliant on a single provider. Microsoft arrived first with the options, but they're not the only show in town, and we should watch a few trailers from the likes of Google and Amazon, who are shaping similar offerings in the background. We'll keep close to this and update as soon as we can. Satya's test from the same keynote is useful: can you remove any model and still improve on your evals? If not, you're locked in. It doesn't cover lock-in to the platform itself.

Prepare the Budgets: You're Going to Need More Money

While I might have a soft spot for Copilot, not because it's the best (it isn't) but because it's currently the safest, with Work IQ gathering pace, there's a big but: how far will they take the upsell? We're already seeing high costs for AI credits (tokens), and an array of updated old tools and newly introduced solutions to help us manage and govern all of this. It all comes at a cost. That neat £23.10 per month solution the corporate world is deploying at pace comes with a bunch of hidden costs to factor into our next IT budgets. We'll cover the management and governance of AI in more detail next quarter, but we're warning clients now: the £23.10 is the low-cost hook, and the likely end game could easily be north of £150 per user per month. Microsoft itself flagged FinOps, security and governance as the work still to do.

The Magnificent 'BUBBLE' News_

So, I'm too deep into this one to back out now, so I'm going with it: ***no, we're not in a bubble.*** *And the question is, do we actually care anymore? Possibly not, but we need to, as we are likely the answer to the whole bubble debate.*

Over-inflated, over-optimistic, unrealistic, yes. But as I've said before, for it to be a bubble it has to be falsely over-inflated, with no sensible, demonstrably realistic route to income generation to deflate it.

So if they are overvaluing the companies, who will be the real winners in the AI race? Are the suggested IPO values crazy? Will people profit, or will we see a repeat of the SpaceX-style peak and then slump?

The labs failed to monetise the model the first time

They are too deep into this to back out

So they are looking elsewhere for the funds

So if we're the answer to the bubble, will we increase our budgets to do it? Microsoft alone has upsold 30m seats to Copilot, just 6.67% of its 450m Office seats. At around £277 per person per year, that's roughly £8bn a year in extra revenue for them, and extra cost for us. Echoing the Agents section, we're going to need bigger budgets. We're in first gear on our AI journey and a massive upsell is underway. £8bn isn't much against the scale of the investment, but it's £8bn we didn't need to spend last year. That's why we track the Mag 7 and the bubble: we are the deflators.

It's not just down to us. A true AI assistant is still more concept than reality, though Meta's Muse and personalised copilots are getting us closer. We're not there yet with an assistant consumers will pay for. But humour me: say Apple made Siri 'actually intelligent' and landed the same 6.67% upsell across its 2.5bn active devices. That's around 167m subscribers and, at about £20 a month, roughly £40bn a year. Still a drop in the ocean, but taken together, Microsoft's Copilot upsell and Apple's potential alone come to roughly 12x earnings, before their existing revenues or any other company's upsell and monetisation. Nvidia trades at around 23x forward earnings. Suddenly the bubble doesn't look so big. To stress: that's 6.67% on today's base, from one company's actual performance and another's very realistic potential. It's not a bubble.

The Magnificent '7' Where next

My prediction: *where next for these investment heavyweights, as they prepare to monetise the next stage of this incredible transformation underway. No crystal ball but we have a good previous of getting this right and the alerts and collision sequencing alarms are going off... Very, very loudly. And we still have the final question... will the Mag 7 change? Will OpenAI or Anthropic make a Mighty 8?*

Microsoft

Copilot moves from zero to hero as organisations opt for the enterprise-safe AI option. With 450 million+ paid 365 seats in play and just over 30 million on Copilot, the upsell opportunity is incredible. Progress on its own chips and models will greatly help costs.

Amazon

What happened to Alexa Plus? I backed it as a great hope and quiet victor of the personal AI assistant market. It's now free with Prime, but most of what we do on Amazon is watch media and buy things

Tesla

While Grok and SpaceX sit outside Tesla's trading position, the Musk story has moved next door. xAI now sits inside SpaceX, whose market value has been reported ahead of Tesla's

Nvidia

They continue to fund and underpin the whole AI industry and cement their position as beyond dominant. But will we see a slowdown as the hype causes a few firms to pause previously endless progress?

Meta

Blew funds on VR, then pumped new investment into augmented and spatial compute, to a few raised eyebrows. I was personally challenged for wearing Meta glasses last month... then Muse lands and gives hope to the playbook of presence. Is this the AI assistant we've been waiting for, and will it work with everything we need it to? Or will Apple and Google block the interconnections needed to make it truly useful?

Google

With the war chest to challenge pretty much every other AI player, a strong model release helps, but we need to see a model and press that captures subscriptions. Will the Apple partnership on Siri bring the audience and users they need? Similar to Microsoft, a focus on CPUs and NPUs will bring control to costs

Apple

Where is smart Siri, and I mean really smart, a more personal AI assistant that swaps our fingers for voice and runs our phone, apps and lives? Google has confirmed Gemini will power it "later this year", after repeated slips. With billions of devices to upsell to, I'm shocked we are still waiting.



Spatial Computing_

Glasses, pendants and wearables are the feeders of AI world models: the eyes and ears that let AI understand our surroundings. I sometimes feel like I should wrap this page up, as every time I think I've covered it, another device lands.

Would You Wear Meta Glasses?

As many of you know, the methodology I use to understand whether a tech will take off is the Tech-Social-Economic wheel. Is it great tech? Will people use it? And can we afford it, or see the value? With Meta glasses, the biggest challenge for me is that not everyone wears glasses. But were they useful kit for our lives? 100%. Do they need refinement? 100%. They were shaping up to be an interesting bit of helpful kit, and I waited for the Display version to go all in. Then a strange thing happened. A guy came up to me on the Underground and asked if my glasses were recording. We parted company after a brief conversation, and while it wasn't overly confrontational, it made me consider the wider impact of the tech on those around me. I have no desire to film anyone, but the perception and the possibility are there (see page 16 for my radio comments on this last month). It's now being called the "pervert problem". A stop-you-in-your-tracks point.

The Pervert Problem

The recording challenge might have been a strange one-off, but do I want to be considered a pervert because glasses are clearly being used more widely to film inappropriate content? These are the issues that break technologies. While I was obsessing over the "cool factor", asking whether people would be seen dead in them, the pervert factor wasn't even on my radar. So what's next for glasses? Meta worked this out quickly and has released a camera-free, audio-only version, rolling out from this month. But the long-awaited Display version, with the heads-up display I've been waiting years for, landed in the UK on 23 September at £749 and still has a camera. Does that mean I bow out and leave this one alone?

The Alternative

So what did I get so excited about with glasses anyway? The heads-up display: that constant in-vision notification opportunity. The Iron Man dream. The good news is that a display no longer has to come with a camera.

Even Realities G2

No camera, a green micro-LED heads-up display, translation, navigation and notifications, controlled by a ring or the frame. About $599 and roughly eight hours' battery. The privacy-first answer.

MemoMind One

Camera-free display glasses with dual screens, built-in speakers and a choice of AI models. About $599, and early reviewers say it's very close to the G2.

RayNeo X3 Pro

Colour displays, 14-language translation and Gemini Live, but it has a camera, costs $1,100 at early-bird pricing and reportedly lasts about 40 minutes of active use. One for early adopters.

Tech for good:

Meta has donated 1,000 Ray-Ban Meta AI glasses to four Singapore organisations, including the Singapore Association of the Visually Handicapped and Guide Dogs Singapore, plus a US$30,000 training grant. They help blind and low-vision users read text, identify objects and understand their surroundings. A reminder of the good this tech can do.

Meta Ray-Ban Display

£749 including the Neural Band wrist controller. A full-colour display in the right lens, live captions, navigation and Meta AI, plus a 12MP camera.

Pendants and Other Devices

Still no sign of the pendant promised by OpenAI and Jony Ive. But Meta's Muse brings us a step closer to personal AI and a world-model input. Is it a bit Tamagotchi in itself? Will it fit with my business-casual quarter-zip and briefcase? Time will tell.



Robotics_

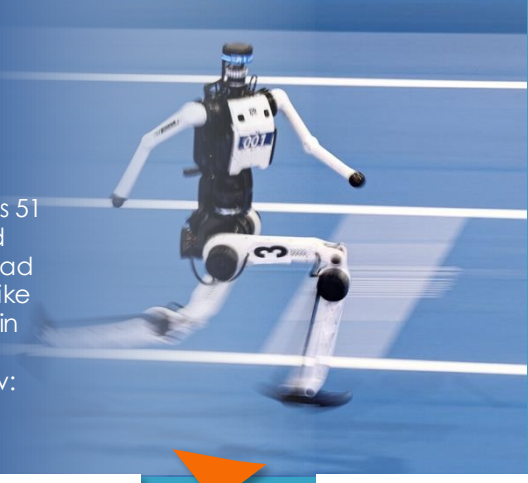

Robot Olympics: Beijing's Humanoid Games

The second World Humanoid Robot Games ran in Beijing from 22 to 26 August, with 2,056 robots from 666 teams (280 teams a year ago) across 51 events, from football and tug of war to household tasks, firefighting and rescue. The headline: Tiangong Ultra ran the 100m in 9.39 seconds, ahead of Usain Bolt's 9.58, though that's the organisers' framing, not a like-for-like record. A year ago the 100m winner took 21.5 seconds. The rule I liked: in the scenario events, autonomous robots scored full points and remote-controlled ones half. Plenty still tripped, fell apart or caught fire. My view: the progress is real, so ask which tasks are truly autonomous.

Robots Take to the Street

Robotics has stopped being a lab demo. Autonomous cars are carrying paying passengers in the US and, as of this month, in London, and the forecasts for humanoid robots are eye-watering. My view: the driverless taxi is the first robot most of us will actually meet, and it is arriving on our streets.

Would You Spend £12k on a Robot?



I was recently offered a robot. It had been used in a pilot at a company that has reached the end of it, and they don't quite know what to do with the machine they're left with. It's a Unitree G1, and for a moment the excitement took over. Who wouldn't want a humanoid robot of their own?

Then the sensible Yorkshireman kicked in: what would I actually use it for?

- The ironing? Not quite yet.
- The gardening? Possible, with some careful management.
- Preparing a meal? Again, with some very good remote control operation.

And that's the catch. For it to do any of these things, we'd be relying on what we see in the great demos, which are either the result of very tightly controlled pre-programming or advanced remote control, aka puppeteering. Not a robot thinking for itself.

Other than that, it's back to its usual party trick: dancing. For £12k, I'll leave it for now. My view: the gap between the demo and the job is still wide. Until a robot can do the boring, everyday tasks on its own, it's an impressive toy, not a colleague.

Waymo: Way Safer

Across about 270 million fully autonomous miles to the end of June, Waymo says its cars had 95% fewer serious-injury crashes and 82% fewer injury-causing crashes than human drivers, with 93% fewer for pedestrians. That's Waymo's own data, but the independent IIHS points the same way. In July it found Waymo's crash involvement 68% lower per mile, while cautioning the sample is limited and that Austin actually came out slightly worse. Safety is the strongest argument for the robotaxi, but we are still testing the data as much as the cars.

Why London Is the Self-Driving Taxi Battleground

On 3 September, Uber and Wayve launched the UK's first robotaxi rides, anywhere in London bar airports, in Ford Mustang Mach-Es with a TfL-licensed driver on board supervising. Over 140,000 Londoners have opted in. Waymo is testing around 100 Jaguar I-Pace cars across 100 square miles, and Government pilot applications only opened in May, so the rules are still being written. Why London? Dense, difficult streets make the ultimate proving ground, and whoever cracks it has a template for every other European city.

1bn Human Robots by 2050: Morgan Stanley

Morgan Stanley's forecast is around 1 billion humanoid robots by 2050, a market worth close to $5 trillion a year, with about 13 million units by 2035. Roughly 90% would be industrial and commercial, with homes lagging behind. It's a big number from a bank, so treat it as direction of travel, not a promise. One to watch.

Practical Applications_

Friends of our Forum, Alex and Laura Harvey, who many of you know as our experts in all things XR, have launched a new VR solution called JobBox which brings immersive recruitment into a portable case. Their solution allows candidates to step inside a real workplace before they apply, while employers get the equipment, films and ongoing support to keep the experience running.

A real view of the job before applying

VR recruitment in a box

JobBox combines bespoke 3D 180 workplace films with Meta Quest 3 headsets, a branded app and a charge-and-travel case with its own router. Employers can take the experience to careers fairs, schools, open days or interviews, giving candidates a first-hand view of the people and places behind a vacancy.

Why it matters

A careers stand can describe a role, but it is harder to convey what working there actually feels like. JobBox aims to close that gap so candidates can make a more informed choice. Its managed service also tackles a familiar obstacle for organisations using VR: keeping the equipment and content ready for use.

The support behind the headsets

Employers choose a box with one, five or ten headsets. JobBox+ provides remote updates, usage analytics, fleet lockdown and content deployment. Organisations own the box, headsets and films, even if they stop the managed service. The process starts with filming the workplace, followed by the branded app and kit. JobBox quotes six weeks from filming day to delivery.

Examples from the field

JobBox reports that McDonald's franchisee People Arches in Warwick uses a five-headset box and five location films at careers fairs across the Midlands. The University of Warwick Fusion Lab uses a ten-headset box with films from Midlands games studios and has hosted hundreds of students at open days.

In a separate Coventry University social work training project, JobBox reports that 94% of trainees said the 3D 180 film helped them see the work from the child's point of view. This is a training result rather than evidence of improved recruitment or retention.

Getting started

Complete setups start at about £7,495 excluding VAT, combining a £995 single-headset box with a £6,500 film package. Five-headset boxes cost £4,950 and ten-headset boxes £8,950. Employers can explore options and book a demo at www.jobbox.uk.

The people and purpose

Warwick-based co-founders Alex and Laura Harvey also lead Vibe Immersive. Alex founded RiVR and scaled immersive training globally. Laura brings experience in VR, education and the NHS. Their mission is for every commercial JobBox to help fund free careers experiences for young people.



As we keep a keen lookout for new, interesting technology, this one felt an obvious one not for a case study but as a very sensible solution we've touched on the need for longer term.

Find out more: www.jobbox.uk / Build your package: www.jobbox.uk/pricing

adventa_

Built by Necessity_ Tech, MND & Independence

Some of you will know Andrew Herbert, Innovation Forum member and CTO at Novuna. As one of the founding members nearly 6 years ago, he helped shape the forum we have today, and his fingerprints are on a lot of what we do and the topics we cover. Andrew was diagnosed with Motor Neurone Disease in 2022, and true to his innovation spirit and form he met the challenge as expected: by building his way through it. What he has created at home, for himself, is genuinely remarkable, and I'm thrilled he agreed to share it here. He is, without question, my Champion of Innovation, and I'd back his contraptions a gainst anything Wallace and Gromit ever put in a shed. Over to you, Andrew.

Where I Am Now



I was diagnosed with MND just over four years ago. Before that, I had always been very active and was a keen cyclist and runner. In fact, it was a noticeable decline in my fitness that first made me realise something wasn't right and ultimately led to my diagnosis. As the disease has progressed, I have lost my mobility and many of the things I used to enjoy are simply no longer possible. That has meant finding new interests and new ways of occupying both my time and my mind. Continuing to work has been particularly important to me. It gives me a much-needed sense of purpose and something to focus on beyond MND. I've also taken up new hobbies, particularly around technology, and that has become an increasingly important part of my life. For me, independence today doesn't necessarily mean being able to physically do everything myself. It's about retaining as much control and choice as possible — being able to interact with the world around me, make my own decisions and find ways of doing things that my physical limitations would otherwise prevent.

The Tools That Keep Me in the Room

My screen time on my phone would probably frighten most people! It has become one of the main ways I interact with the world and I use it for much of the day. As my movement has become increasingly limited, I've had to find different ways of controlling technology. I use a very sensitive head-tracking mouse which can detect the small movements I am still able to make with my head. To click, I use an EMG switch which detects tiny muscle movements in my thumb. Together, they allow me to operate a computer even though conventional keyboards, mice and touchscreens are no longer practical for me.

I'm also fortunate that I still have my voice, and that is incredibly valuable. I use voice assistants extensively throughout the day to control things around the house and to access technology without needing physical interaction.

I rely on carers 24 hours a day, so I've also tried to automate as much of the house as possible. Importantly, that isn't only about making life easier for me. If I can make everyday tasks simpler for my carers as well, it benefits everybody. That has changed the way I think about home automation. For me, turning on a light, controlling entertainment or getting information without asking somebody else to do it isn't just a clever feature. Each small thing I can control myself is one less thing I need to depend on another person for. The technology doesn't give me back the physical abilities I've lost, but it does give me back some control — and that has become my definition of what useful technology really is.

What I've Built for Myself





A lot of the technology I use isn't something I've been able to simply buy off the shelf. I've ended up building and adapting things myself, often combining relatively inexpensive consumer technology in ways it was never originally designed for. One example is switching between my different devices. I use a PC, tablet and phone, but physically changing controls between them isn't practical for me. I've built voice-activated switches that allow me to move my controls from one device to another without needing somebody to do it for me.

I've also put cameras around the house. They allow me to see beyond the four walls of the room I'm in, but I've taken that much further by using local AI processing. The system can recognise visitors, describe what is happening outside and even identify wildlife. Much of that processing happens locally rather than relying on a remote cloud service. I've built fast Raspberry Pi-based screens around the house which work alongside my smart speakers. They can show me my vital signs, my calendar, where my wife is, who is at the front door or even what has just been put through the postbox. Rather than having to go looking for information, I try to make the right information come to me when I need it.

An important part of the system is knowing where I am. I have Bluetooth proxies around the house which detect a small tracker attached to my ventilator, which is always with me. This means the house can know which room I'm in and automations can be specific to that room. Instead of the whole house responding to something, the technology around me can respond to where I actually am.

At the centre of everything is Home Assistant running on a Raspberry Pi 5, alongside Frigate, which I use for the camera and local AI system. I've gradually connected more and more things to it, but my aim isn't to automate the house simply because I can. I automate things when they solve a problem. That's probably the biggest difference between what I've built and a conventional 'smart home'. For me, these aren't gadgets. If technology can remove the need to ask somebody to perform a small task for me, or can give me information I couldn't otherwise access independently, then it has given me back a little bit of control.

What's Next on the Bench

One area I've been working on recently is health monitoring. Because I rely on carers 24 hours a day, it would be very easy to ask somebody to check something every time I don't feel quite right. I've been looking at whether technology can give me some of that information directly instead.

I now have a mmWave sensor which can remotely monitor my heart rate and respiratory rate without anything needing to be attached to me. I have a temperature sensor positioned under my arm, and most recently I've built a Wi-Fi connected SpO_2 sensor which measures my blood oxygen level and feeds the information into Home Assistant using MQTT.

I also use Gemini to read the leak figure displayed on my overnight ventilator. This helps me keep an eye on whether the mask has maintained a good seal overnight and whether my ventilation has been working as expected. It's another example of taking information that is already there but would otherwise be difficult for me to access independently, and making it available to me.



I've brought this information together into my own health-monitoring dashboard, so I can see it in one place. It's not intended to replace medical equipment or professional judgement. It's about having better access to basic information about my own body without always having to ask somebody else to get it for me.

When I say I've 'built' these things, I should also qualify that. I have plenty of ideas, but putting them into practice increasingly involves other people. Friends and carers help me physically assemble things that I can no longer manipulate myself, and AI has become an increasingly important part of the process too.

That combination has become quite powerful. I can have an idea, work through how it might function, use AI to help with coding and troubleshooting, and then have somebody provide the hands that I no longer have to physically put it together. In some ways, that's what is currently on my bench: not one particular invention, but finding more ways of combining my ideas with other people's physical assistance and AI's technical assistance to overcome the limitations MND has placed on what I can physically do myself.



AI, Health and Humanity

Even in the last six months, I've noticed a remarkable improvement in what AI assistants can do. For me that's particularly exciting because, as my physical abilities have reduced, I'm finding myself using AI more and more to improve my life. The acceleration from having an idea to actually making it work can be phenomenal. I can describe something I want to achieve, explore different ways of doing it, develop the code and troubleshoot problems with AI alongside me. Combined with the practical help I get from friends and carers, it means I can continue creating things despite being physically unable to build them on my own. That has changed what feels possible. An idea that might once have required specialist knowledge, considerable expense or somebody willing to spend days developing it can sometimes become a working prototype remarkably quickly. For someone living with a progressive disability, that ability to adapt technology around your changing needs has enormous potential.

But I'm not naive about AI either. Like many powerful technologies, the same capabilities that can produce very positive outcomes can also be used in ways that cause harm. That concerns me, particularly in the increasingly divided world we currently live in. I think we all need to understand AI better — not necessarily the technical detail of how every model works, but what it is capable of, where it can genuinely improve people's lives, where its limitations lie and where appropriate guardrails are needed. My own experience makes me optimistic because I can already see what AI can enable. But perhaps that also makes me more conscious of the responsibility that comes with it. The technology is developing extraordinarily quickly, and society needs to keep up with the conversation about how we want to use it.

Otherwise, it would be very easy to sleepwalk into something we can no longer control.

Or perhaps the more uncomfortable question is: **are we already doing that?** And, after everything I've just said about AI, there's probably one final question worth asking:

Did I even write this article?

Andrew has also kindly agreed to share further insight on our next Innovation Forum Call to bring this to life and to share more of his experience with technology overall. Andrew, thanks for sharing your experience and insights and we look forward to talking further on our upcoming forum call, and seeing the next innovation soon. Follow Andrew's journey on Instagram at: https://www.instagram.com/andrewherbertmnd/#

Societal & Economic Drivers_

AI may grab the headlines, but the real drivers of change are often wider, quieter and more impactful. Economics, housing, welfare and mood decide how any technology lands. That's why we **'follow the money'** and **'follow the mood'** as part of our **Collision Sequencing Methodology**, which allows us to predict likely future outcomes with incredible accuracy.

Economy

We remain party-neutral, focusing on how economic policy affects organisations and their people. Bond-market scrutiny brings a sobering reality to political promises and the challenge of balancing the public finances.

Housing affordability, inflation and interest rates remain central concerns, especially for younger buyers. Pressure on welfare spending and tax receipts adds to the wider debate about the choices ahead. Inflation (CPIH) rose back to 3.5% in the year to September, and three Bank of England policymakers voted for a rate rise last month. In our view, the jump to 3.5% now makes a rate rise inevitable, and a real challenge for John Healey as he heads towards his maiden Budget.

Inflation	GDP	Interest Rates	Unemployment
3.5% ⬆ 0.3%	0.1% ⬇ 0.3%	3.75% ➡ HELD	4.7% ⬇ 0.1%
CPIH to Sept	*Aug (m/m)*	*to Sept*	*to Aug*

Society

Are we approaching a tipping point in the relationship between work, welfare and financial independence? These pressures raise questions about how people build security and how much control they have over their future.

For some, stepping away from work may be a genuine choice. My concern is that younger generations could have less freedom, with housing costs, pension provision and future government policy shaping their options.

We often focus on extraordinary technological change. Yet societal change could prove just as important in determining how the AI era unfolds.

Cultural Impacts

In dozens of AI training sessions over the past three months, I have noticed two reactions: an "aha" moment for around 80% of the room, and "but what about us?" from the other 20%.

The wider economic debate can deepen tensions between generations. Rather than reducing this to who has been lucky or who works hardest, we need to understand the different choices and constraints people face.

Could this be a period of transformation and adaptation rather than simply decline? How people feel about their future will help shape their response to technology.

Is This Building the Case for UBI?

As the economy and, in turn, people struggle, is this building the case for AI and a nod towards UBI? We raised the question in our last report, and we think it now needs to be firmly part of the plan. Blame won't help. We are where we are, and we need a plan to move forward.

I heard a woman on TV recently say she 'makes' £45k a year on benefits. The word that stuck was 'makes': she feels she earns it in the same way we do. We might passionately disagree, but just as we can't fight the tide of tech progress, neither can we fight the tide of social and cultural transformation.



People Focus_

The 9 to 5 is a routine we've come to know and expect. Flexible working and Covid gave us other ways to work, yet beyond a guilty glance at our phones or a late night with a laptop, the pattern remains familiar. This year marks a century since Henry Ford introduced the five-day working week to attract workers. Technology moves on, but our expectations are slower to shift.

The key to happiness is…Routine
The Dalai Lama

The mood around AI is changing

When I used to tell people what I do, it sparked interesting conversations about the future. I rarely tired of hearing their views on AI. But something has changed. Recent conversations at a wedding, on a train and on a training course have turned towards its negatives, often with real anger and frustration. People are starting to push back, and mentioning my work can make me feel like the spokesperson for the whole AI movement.

The case against AI

We don't always welcome change. Usually it happens quietly in the background, giving us time to adjust. With AI, the pace is much faster, and that is unsettling people.

To me, The Economist's London posters and billboards seem almost to invite pushback, while speakers predict growing resistance. It feels like they're stirring the pot.

We should take the concerns seriously rather than dismiss this as another Luddite movement.

The case for AI

While we may resist change, we do love progress. Few of us would swap smartphones for basic phones or give up remote working to return to the office five days a week.

Instead, we're embracing the technology. I see customers' AI deployment reports and, frankly, people are really using it. For many, this is an hourly habit, not a daily one. Which leaves the telling question: could you give it back?

The tension for leaders

And here's the rub: it's rarely one-sided. The young man who challenged me about AI on the train was equally adamant that returning to the office five days a week made no sense when technology lets us work differently. I paused and asked: isn't that what's happening with AI, even when it means less work or fewer people?

I was playing devil's advocate, not arguing for a return to the office. We wouldn't turn the clock back on other forms of progress, so why might AI be different? I'd love to say I won the argument, but I didn't choose a side. Yes, I'm an AI adviser, but I also worry about my children's careers and the environmental impact.

Whichever camp we sit in, work is changing. At a recent event, someone described us as 'custodians of change'. That stayed with me. It falls to us to make reasonable decisions that could shape lives for decades. No pressure!

AI and the job-cuts debate

I try to take a balanced view, though I lean towards AI for its potential to improve efficiency. My young friend's concern is real, but so are the economic pressures organisations face. It is easy to badge job cuts as AI-led while much adoption is still at the test-and-trial stage. Wider trading conditions and employment costs also matter. We owe young people an honest explanation of those pressures, however uncomfortable, rather than letting AI become the single explanation for every cut.



Security & Governance_

Do you have a company-approved AI model, e.g. Copilot? Do you have an AI policy? Have you signed it yourself? Have you had the appropriate training on AI? Do you understand safe AI versus unsafe AI? And can you still access other AI models?

Well, the answer to most of these questions is likely "yes, sort of" or simply "no". Yet we're knowingly or unknowingly deploying possibly the riskiest technology our organisations have ever seen, with very little in the way of guardrails. There will inevitably be mistakes and lessons learned, but at what cost and consequences?

Regulation will follow, but for now it falls to us to build the guardrails and governance needed to not only de-risk, but deploy AI with company-wide confidence, and… WE HAVE TO GET THIS RIGHT.

From Guidance to Governance

The cost of getting this wrong is incomprehensible. While I grace the stages of many conferences telling a story of incredible change, opportunity and transformation, the real headline won't be an organisation deploying AI to incredible advantage. Our press machine will see that as sales patter and instead focus on the stories of organisations that got it wrong: data leaks, rogue agents, deepfake losses and AI errors in front of customers.

And shadow AI is already here. A BlackFog survey of 2,000 UK and US employees found 49% use AI tools their employer hasn't approved, and 60% say the security risk is worth it if it helps them meet deadlines. Meanwhile, the UK still has no AI Bill, and a King's Speech that includes one isn't expected until late 2027 at the earliest. Until then, governance is down to us.

Governance: Where to Start

- Approved tools list: a clear yes, no or ask-first, so people know what's safe
- One-page AI policy: short enough to be read, signed by leaders first
- Training before access: safe versus unsafe AI, before the licence goes live
- An AI register: what's in use, who owns it and what data it touches. The quickest way to surface shadow AI
- Treat agents like staff: a named owner, least-privilege access and logged actions
- An off switch and an owner: who can pause it, and what happens next

Regulation: What's Coming

- UK: still no AI Bill, and not expected before a King's Speech in late 2027 at the earliest
- EU: the Digital Omnibus has been in force since 27 July. It moves high-risk AI Act duties to 2 December 2027 (and 2 August 2028 for AI embedded in regulated products), though transparency duties remain
- UK Cyber Security and Resilience Bill: Lords report stage is due on 26 October, with Royal Assent expected between late 2026 and spring 2027. As introduced, it means incident reporting within 24 hours and fines of up to £17m or 4% of turnover, and it brings managed service providers into scope. Duties phase in, with full effect around 2028

Security

Deep Fakes

The biggest threat to security of our time. This is big business, very big, and we're about to experience the most disruptive security threat of our time. It may not be the video or voice call many predict, as only very close colleagues or loved ones will really move money to us, out of duty or love… and both will know us better than hackers and scammers would like to convince themselves. So the risk is likely not a scammer pretending to be your partner, child or close colleague, but a "friend" on Instagram who fools us into connecting in other ways: dropping a work email over, or getting us to join a cause they feel we might support.

2,137 %

Rise in deepfake share of fraud attempts over three years (Signicat, 2025)

£20m

Lost in single video call scam (ARUP)

Agents Are the New Attack Surface

OWASP's Q1 2026 round-up describes a clear transition from theoretical risk to real-world exploitation. Most incidents have no CVE: the weak points are agent identities, orchestration and supply chains. A Meta agent's flawed advice exposed internal data for about two hours, and fake "leaked" repos spread malware after the Claude Code source-map leak. Signal: secure the agent's identity, not just the model.

AI as the Attacker

OpenAI says GPT-6 Astra is the first model it has broadly deployed to reach its "Critical" cyber level, able to find and use zero-days with the right access, and admits it is harder to monitor than its predecessor (OpenAI's own findings). The NCSC calls AI's effect on ransomware over the next two years "evolutionary, not revolutionary": volume and impact will rise, but most incidents still come from poor hygiene.

Prompt Injection

OpenAI says Astra resists indirect prompt injection 99.79% of the time. Our view: an agent reading thousands of documents will still meet the other 0.21%.

Microsoft's Answer

Frontier Tour: Agent365 as the control plane, with an ID and tamper-proof logs for every agent. Project Perception adds red, blue and green agents (attack, defend, fix).

Innovation Forum Roundup_

We've been thrilled to work with so many of you, our members, over recent months, holding some incredible conversations, sharing our insights, working on some incredible projects, and attending some awesome events.

"Coming together is the beginning. Keeping together is progress. Working together is Success" **Henry Ford**

Innovation Festival.

It's not every day you get to wear a cape to work



There's a week in the year that I look forward to a little more than most. The Northumbrian Water Innovation Festival. Some of you will know all about this as we've shared stories and examples from the festival over the years… And this year I was thrilled to be able to support the event in a slightly different capacity. As a futurist, and while many ask 'what exactly is a Futurist?' for this season I was armed not only with a crystal ball and future booth but also a booth and agenda to visit as many of the 40+ sprints I could get around over the week. The people are incredible, the conversations immersive and enlightening and the output is very real. Honoured and thrilled to be part of it and look forward to next year already. And congratulations to the NWG team on their 10th year of this amazing event.



Customer Innovation Forum: 7 July 2026

It was a pleasure to return to Virgin Media O2 Business as Resident Futurist for July's Customer Innovation Forum, exploring the future of collaboration. We looked at how people, AI agents and immersive technologies could change the way we work, and what organisations need to do now. I also hosted a panel with M Group, NA and RNIB, sharing honest experiences of building agents, tackling governance and making accessibility part of the design. My takeaway: start with repetitive tasks people would gladly hand over, keep human oversight and understand the costs before scaling. We'll continue the conversation on 5 November, focusing on cyber security and connected AI.

Talk Radio TV: Protecting children

I joined Petrie Hosken again on Talk Radio TV to discuss social media addiction and the California case involving Meta. Would the jury accept claims that platforms deliberately addict children? What would the whistleblower's evidence reveal, and what financial penalties might follow?

My purpose was to push for action: collaboration and technical solutions that protect our children. If we cannot agree effective safeguards, should we block UK access until we can be confident they are safe? Too harsh? Perhaps. But if young people can bypass age bans, we need to ask what protection will actually work.





LBC: Filtering the AI hype

Picture the scene: 8am on Saturday, dressing gown on, coffee in hand. An LBC producer calls. Would I join Matthew Wright to help filter the "Headlines of AI Hype"?

The questions: is AI taking as many jobs as the headlines suggest? Is it behind recent positive economic growth? And are data centres fuelling the growth of the anti-AI movement?

My purpose was to bring lived experience to the debate, sharing what I see working with you: enterprise and public sector organisations, business leaders and Innovation Forum members.

After my impromptu soapbox speech, the producer asked: "Great, can you do that again on air?" Thankfully, it was audio only. Looking more Arthur Dent than Ant Morse, the hairbrush could wait.

Future Get Together – We STILL haven't forgotten our next (now VERY overdue) in-person catch-up and are working with our friends at Cavell to bring us all back together in London later this year. More info to follow soon.

adventa_

adventa_

exploring the future, for our future

making the complex simple

and taking everyone on the journey with us



www.adventa.tech