

the future of us_

Innovation Forum

The Balance
Edition

Forum Members Insight Report
by Ant Morse – Chair | Innovation Forum

Apr|26

adventa_

“*The value of an idea lies in the using of it*”
Thomas Edison

Executive Summary

As a forum, we have always prided ourselves on being ahead of the curve. While back in 2021 the Big 4 were still evangelising Cloud and Digital transformation strategies, we were deep into the world of Generative AI, Agents and Robotics, and while many raised an eyebrow in our direction at the time. I'm grateful you stayed with us, as the landscape then rapidly moved decisively in the direction we predicted, and the next wave of transformation is likely to make Gen AI look like a warm-up act, as we now head into a period of change many cannot yet quite imagine.

In this edition I'll start to expand on this next stage, which in brief will see us move from Large Language Models to Context-Aware World Models, where Spatial computing, Robotics, and AI Personal Assistants are no longer concepts or gadgets; they are instead the data feeds that will power more symbiotic AI, which no longer rely purely on our prompts, but instead also start to think for themselves, based on our environments and circumstances.

And while that last sentence might raise an eyebrow, it's possibly not as radical as it might sound, nor to be as clear an introduction to the often claimed AI General Intelligence, it's more simply a collision and progression of existing technology capabilities, and a commercial battle by big tech for our wallet spend.

But as I make a big claim here, as always, I hopefully bring common sense and logical scepticism, and if the last three and a half years have taught us anything, it's that progress and risk go hand in hand and finding the happy medium is as ever an incredible balance to find, so we've called this edition **The Balance Edition** for good reason.

That title also ever relevant for us, as leaders as while AI continues to advance at extraordinary speed, the reality of corporate adoption tells a more complex story, and from our work at Adventa we are spending more of our time pausing projects, closing down pilots, or having some candid conversations about Governance or ROI cases that just don't stack up.

Our current landscape and stage of AI progress needs very careful focus and action, and as a collective forum rather than shy away from this tension, we tackle it head on, from the real questions around value, data security, governance and cultural readiness, to the headlines of hype around AGI and the Magnificent 7 bubble debate. Living by our premise that the organisations that will win are not necessarily the fastest movers, but the most informed. And on that front, as ever, as a forum...I think we're all in the right place.

Very Best,

Ant

Ant Morse – Chair | Innovation Forum

adventa_



Q4 Report Content_

Page 4

Latest News from the world of tech & the AI & Tech Industry

Page 6

Horizon Scanning – LLMs to World Models

Page 8

Has Gen AI hit a bump in the road

Page 10

In Spatial Computing we explore the feeders of World Models

Page 12

We hear from Northern Accountants COO on their actual practical AI Apps

Page 14

AI Job Losses, real or simply a smokescreen?

Page 16

Crazy or Credible Experts' genuine views or Agendas

Page 5

Our AI Moment – Balancing the Reality v Rhetoric

Page 7

The Magnificent Bubble is it going to burst or just a lot of hot air

Page 9

AI Agents - The Great Corporate Hope

Page 11

Tech for good Forum Member – Robin Spinks shares his work with Waymo

Page 13

A check in on Societal and Economic drivers

Page 15

Social Media and Model Guard Rails hit the headlines

Page 17

Update on another busy period, with you Our Forum Members

Tech News_

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

Chip Wars Continues

As the AI bubble debate intensifies and big tech scrutinises every line of its balance sheet, the question of chip supply and dependency has moved from a background concern to a boardroom priority.

Few predicted quite how quickly the demand for AI compute would outpace the ability to supply it. Nvidia read the market brilliantly and found itself in a position few companies ever occupy - effectively indispensable. But indispensability at scale always attracts two things: competition and cost pressure. And it was only a matter of time before both arrived.

While many assumed the chip wars would play out as a US versus China battle, and that narrative is very much still live, the more interesting move came from closer to home. The deep-pocketed hyperscalers, Google and Amazon, began quietly developing their own AI chips, reducing their reliance on Nvidia and reclaiming control of both cost and capability within their own ecosystems. It's a classic vertical integration play, and a smart one.

Sovereign AI

Major economies are now treating AI as critical national infrastructure. France committed €109 billion to AI, the UK launched a Sovereign AI Unit, and European nations are forming compute alliances to build GPU clusters independent of US and Chinese tech ecosystems. Signal: Organisations with multi-geography operations must develop region-specific AI strategies as infrastructure increasingly fragments along geopolitical lines.

In other news

Vibe Coding named Collins Dictionary Word of the Year 2025, with 92% of developers now using Vibe tools daily

India introduced an AI Judge, integrating AI into legal rulings, helping to reduce case times, but raising hallucination risks

DeepSeek's R1 didn't just shake AI, it shook assumptions. The more efficient model sparked a \$600B Nvidia sell-off as we rethink the cost of intelligence.

Data centre power demand set to double by 2030; SMRs and space-based data centres are the only credible solutions

Data Centres in Space – the new Edge networks

When most people think of space, they picture Elon Musk sending humans to Mars. And while that vision makes for compelling viewing, catching rockets out of the sky is genuinely spectacular, the reality is that a crewed Mars mission remains a very long way off. Bone-crushing radiation, zero return infrastructure, and a few other significant engineering challenges see to that.

The far more immediate and practical story is what those rockets are more likely poised to carry data centres and edge computing satellites into space.

The economics are compelling. In space, the two biggest headaches of earth-bound data centres simply disappear. Solar power runs approximately six times more efficiently, and cooling, one of the industry's most water-hungry and costly problems, requires none at all. Add to that the frequently overlooked point that satellite AI infrastructure is considerably more efficient than ground-based cellular networks, with low latency that could give early adopters a genuine competitive edge, and the case starts to look less like science fiction and more like sound infrastructure strategy. *See our infrastructure and bubble analysis on page 7.*

Move to World Models

While we are often very critical of the claims that we are close to achieving AI General Intelligence, our horizon scanning alarms are ALL going off as we witness some technology collisions that will bring us closer to AI World models that no longer rely purely on our prompts but start to think with us... and also for us.

Could this be the lever of change that will not only impact on our lives, but reshape society and the entire global economy? It's certainly looking that way... yet under the radar of many. Learn more on Page 6



Our AI Moment : Balance of Reality 'v' Rhetoric

Over the past year, as a forum, we have examined AI through leadership decisions and lived organisational experience rather than headlines of hype, balancing opportunity with the twin pressures of **FOMO** (fear of missing out) and **FOMU** (fear of messing up). As we head into 2026, a clearer picture has emerged: while AI capabilities are advancing at extraordinary pace, the outcomes being realised by organisations remain far more mixed than the public narrative suggests. This reflects our work at Adventa where we pause as many projects as we progressed due to issues with ROI, governance, security or cultural readiness. Yet we remain resolutely convinced that the next stage of AI will bring the value and controls we need. But here, let's look at what we feel are the key factors for our continued close evaluation, while highlighting why measured, informed leadership now matters more than ever.

Business Case – Has anyone seen my ROI ?

Do we all have robust business cases for AI, are we investing in efficiencies, or simply to be 'doing' AI ? Both count and while the ROI rigour is there, we also need to step into some progressive pilots ready for future collisions and customer or employee demand, or in response to competitive planning,. We need a health balance and supportive Board in sync

Customer Expectation - We want more!

Sure, customers want improved services and better experiences, and with their wider awareness of AI's capabilities, the pressure is growing to provide more self-service proactive solutions and to free up agents to be available to 'speak to a human' when it counts. Also as AI to AI overlaps into AI Assistants, AI might make the choice of who to use, based on service, before we even get the enquiry...Service will become the first priority in the near future.

Cultural Evolution – new future of work

Tech takes the headlines but as we've explored many time, we are also living through a time of considerable societal and cultural evolution. Peoples' expectations, wants and desires of what a fulfilling career looks like have changed.

And with the rise of Sickfluencers and NEETs now hitting 1m in the UK. The talent funnel already looks a bit different. And this isn't simply driven by generational attitudes it's a far wider societal and economic evolutions, as the gap between taxes collected £450bn (up 3%) versus benefits paid out £334bn (up 5%) closes. A troubling trend, and our research suggests we are only at the very start. So how might AI help us make work more attractive or very harshly provide the backfill

Where Next – Horizon Scanning

And while our focus is on AI, many are missing some of the wider levers of transformation such as spatial computing and Robotics, with a degree of comfort that they still feel 'a little bit gadgetry' and one we can kick into the long grass, blissfully unaware that the focus on these technologies to capture our attention and usage of their 'heavily invested' AI solutions will inevitably push us towards the next stage of AI not based on technology competitions, but more so on commercial edge and in some cases survival. See more on page 6

Our Board Summary

AI doesn't sit neatly in any one place on the org chart, and that's part of the challenge. It originates in tech but its impact runs far deeper, reshaping what we can do operationally, how our customers experience us, and what they'll expect next. The Board naturally looks to the CIO or CTO to lead, but the truth is this cuts across every function, from customer experience to work force planning to commercial strategy.

As a forum, we track all of these threads, but are firms all joining the dots internally? Do we have a person/team, dedicated to sitting across these intersecting forces? not just the tech, but the cultural, commercial and competitive shifts it brings? Because the organisations that treat AI as a standalone IT initiative will find themselves playing catch-up to those that recognised it for what it is: a boardroom-wide conversation about the future shape of the business.

Headlines of Hype – Crazy or Credible

As we consider if we are learning enough, preparing, piloting and deploying enough AI, rarely does a day go by without an aging expert or tech leader or visionary tell us AI General Intelligence is close the end of jobs and mankind is neigh.

And while we can't dismiss these incredible minds, it's not particularly helpful while we're still trying to get John in accounts to stop clicking links in email and Sue in Sales sharing half our company data with Chat GPT. And ultimately the technology to reach AGI is not even close to the Gen AI solutions we have today, causing an unnecessary distraction to our current task in hand. Safe, sensible, useful and efficient AI. See Page 16 for a closer review

Is it all just a Bubble

We've explored and tracked the performance of the Mag 7 for some time and called out overly high valuations way before Bubbles were mentioned, but for us unlike the DotCom bubble we didn't take an extra holiday because the market double the valuation of LastMinute.com... but ask yourself this are 'you paying for an AI solutions or would you if the free versions disappeared?

Actual Progress – Doughnut AI

when is AI enough? Today's systems are already capable of delivering meaningful efficiency gains, decision support and productivity improvements across most organisations, yet the industry narrative insists on relentless acceleration at any cost. Rather than assuming progress must always mean bigger models, more compute and higher energy demand, leaders should pause and ask whether existing capability already sits in the "useful zone" creating economic value without unnecessary environmental or financial overshoot. The challenge is not whether AI can keep advancing, but whether it needs to in every context. A more responsible AI future is one that deliberately defines pause points, distinguishing between progress that genuinely improves outcomes and progress that simply burns capital and carbon for marginal returns.

Regulation & Security

AI brings inevitable security risk and increases the threat landscape, and we have to balance the safe option to do less with the pressure to do more from multiple angles. And while we operate in a time of no official UK AI guardrails we know they will inevitably follow, so how do we impose our own benchmark today, address risk already here and prepare for future regulation, between the FOMU (fear of messing up) as we await an inevitable newsworthy disaster

Horizon Scanning — LLM's to World Models

'What if you're wrong?' ...were the words from a renowned publishing house, which caused a few sleepless nights, when they advised me that writing a book predicting the future was a career limiting project, as surely no one can accurately predict the future. 3 months later I self-published my book, which enjoyed its 3rd birthday in Dec and successfully 'did' predict the future in terms of Gen AI, Agents, Spatial computing and a number of societal and economic transformations. Over half of the predictions are now an everyday part of our livesand the other half?...they are now progressing at pace and highly likely to appear very soon.

This isn't to brag (I also backed Nike 5 years ago to be the leader in smart clothing, but it does raise an important question: *how did I get so much right...* and why didn't I buy more Nvidia shares?

The answer, at least on the transformation front, lies in our research methodology called **Collision Sequencing**. It's a research process this forum has used many times over the years, helping us spot early signals and focus on the right topics at the right time. Spotting Gen AI, Agents but also rise in unemployment and cost of living challenges from this methodology. .

So, the question for today is simple: **what are the next 'collisions' ahead of us?** Does the process still hold true, or did I get lucky or is this the real 'where next' in the Future Of Us

LLM's to World Models

Today's AI is undeniably impressive, but as we are now finding its limits and reality that in its current guise, it's primarily a tool for human amplification, rather than replacement, with us directing it by prompts.. some of which are so detailed, we might be better just doing the work ourselves. So, with mass of investments and all out brutal competition between big tech, how will this game changing tech now change the game again?

Well, the adventa collision sequencing alarms are going off, and calling out what we feel is going to be possibly the biggest progression since Gen AI itself, and result in one of the key outputs needed to reach Contextually Aware World Models.

World Models enable AI systems to build an internal representation of their environment, learning how the world works, how objects behave, how actions lead to outcomes. Instead of relying on our detailed prompts or static knowledge, AI can begin to infer context dynamically: understanding local conditions, physical constraints, and cause-and-effect relationships without being explicitly told. This is where the next phase of AI acceleration begins. World Models will be trained not just on text, but through interaction with the physical and spatial world, learning from spatial computing hardware, wearables such as glasses, robotics, sensors, and real-time environment data and wider IoT devices. By observing, moving, and interacting, AI starts to close the gap between knowing about the world and understanding it, and reducing the level of prompting we need to provide.

Three of the key factors needed to deliver world models have now progressed from the reels of sci-fi movies, into now-fi realities, and we need to see them not as gadgets, but as the 'AI feeds' they will soon become.

Spatial Computing

While many dismissed Glass as a gadget, but its far from, as it holds two key focuses (excuse the pun) beyond novelty. Firstly, it's a very sticky access point to AI. I use Meta Glasses and I can only speak to Meta on them. Secondly, and my point here these solutions will soon provide a world model feeding system, with some subtle but great test example appearing, where users open the camera and ask about their location, such as in sight loss where our forum member Pete Osborne used them on our last call instead of a usual screen reader. Another example being that Tesla knows where all the Potholes due to their AI Lidar camera reading the roads. In the future when you ask AI a question the first thing it will do is check where you are and if that info can help it provide a better results

Robotics

Similarly to Glass solutions, they will become a forced access point to the manufacturer's AI, and imagine if you have a Robot in your home in the future having yet another AI voice assistant will be overkill so we will likely use the one most useful to us, and the one doing the cooking and cleaning will take some beating. This is already happening with Alexa, it's just not that useful overall, but its no mistake. Also, like Tesla cars Robots will learn from each other and communicate with each other, building a mass awareness of the world. Bots will also speak to each other and share awareness locally and globally, learning from the lives they support. Gadget for today but key feeds for world models tomorrow.

Agents to Assistants

And beyond the physical wearables and robotics, my big prediction back in 2015 is now finally starting to appear in the form of AI Personal Assistants. These will be small language models with unlimited memory, that act as agents on our devices interacting with our current apps on our behalf. We will program with our likes, dislikes, morals, ethics and preferences, and tweak over time.

They will move from useful, to impossible to live without, as they truly assist many aspects of our lives. The race is on to provide them, and the key thing is we will pay for them, and the possibilities are endless, so too are the subscription options from services that will support our health, to doing a quarter if not more of our day jobs for us, and aiding our lives in many simple but neat ways. This is the breakthrough for me that means we are not in an AI bubble. More a speculative early over inflation.

How might this play out...

Rapid progress is now BAU, and we'll see more spatial tech hit the market in the form of more glass tech, as Google partner with Samsung, and the Meta Display comes to the UK. We'll also see more personal AI devices, (Open AI's device might be 2027 but who knows) and expect to see the smart phone market fight back, with Apple possibly coming back fighting with their new Google powered Siri. AI Assisted Health will be a primary focus for AI Assistants, and the one we'll pay for, so I expect to see someone offer an intelligent Health agent very soon.

The question is, was this progress and focus always intended, in spatial to feed context and robots to learn the world around us, or is it just serendipitous progress, or incredible collective benefit.. Well weather we care or not, it's happening, on our watch and its likely to be fascinating... and who's going to win the race? Well, I think we've found our winner, but that's one to talk through over coffee, adding I've gone 'all in' on two firms...but unlike Nvidia this time I'm going big.

The Magnificent 'BUBBLE' News

Way before talk of a Bubble, as a forum we raised the question of the 'interesting' scaling valuations of some of the Mag 7 companies and wider, aware that all are at very different stages today, and their innovations of the past are no guarantees of surviving this current and next stage of change....such is the enormity of change, the only certainty is change.

So when the Press started to talk of a Bubble, we largely ignored it, feeling we had this one in hand, or certainly in our sights already, but not from an economic implosion aspect, but more around who we should partner with for our enterprise AI solutions.... think Blackberry v iPhone, or Lotus Notes v Microsoft Outlook.... This is not our first rodeo.

So, are we in a bubble ? And does it matter ? No! Yes! , well Sort of!it's complicated

No, in that unlike the Dot Com bubble the only input to inflate the bubble then was speculation of our spending, and as an example, speculating that an on-line travel company was worth £x millions, didn't make us suddenly take an extra holiday or two, it was falsely inflated and similarly in the subprime housing market crash of 2008, people didn't suddenly earn more to pay for mortgages they were given without due diligence of the lenders, in arrogance to the assumption of continuous growth and progress. The difference with AI is we 'are....and will ' pay for it., albeit gradually, and it will bring new solutions that provide services so useful to us (and wider society) that speculation will be followed by our hard cash to deflate the bubble... Now to put into perspective to reach the current valuations of some of these companies, it would need a miracle, or some serious spending and insane wage rise, but the spend here is not the cost of an extra holiday, or a stretching mortgage, it's more like a couple of cups of coffee a week, that collectively is far more likely to be spent, and by a vast proportion of society.


Yes, in that a number of these companies are massively over valued (most x35, some x100+) and their futures are far from certain; such is the scale of change and competition underway....reality is not all will survive. The point here is the future will not just a few neat services or solutions with a plethora of choice to consumers, if they build their solutions the way we think they will, they will result in very sticky offerings that see us wedded to their solutions and justifying far more of our spends (See further thoughts on page 5)

Also, in most cases the valuations are not based on results, its pure speculation. While Nvidia's earnings call drives valuations up, Open AI's should by the same see future valuations negative, yet they are poised to go all in on their future IPO. It's contradictory madness, but let's remember the markets incompetence pre 2008, they didn't need to be Michael Burry to see what was happening in hindsight. They all just self fulfilling bought new surfboards to ride the tide of speculation #bonusseason

Does it matter to us ? 'The Balance' Yes, for a few key reasons. **Firstly**, as firms we don't want to back the wrong horse, be it from the financial impact of having to swap suppliers further down the line, not to mention the upheaval. But **secondly**, and more critically, around the question how do we swap the memory & training these systems will build and refine over time? You can port your mobile phone, but can we port AI memory & training ? I'm sure port options will follow but that will take time.

On a personal note, our savings and investments are heavily reliant on the stability of these stocks, and while war chests and big backers mean we don't predict anyone going bust, the investment choices link to future innovation will be key and should be the focus on fund managers, not just the polished CEOs' developer day pitches.

Jensen Huang CEO of Nvidia describes AI as a 5-layer cake (see below) and we've added a crucial 6th something that many are missing as the layer needed to deflate the bubble ' Income ' so here's our take on the current bubble and forecasts



* SMR – Small Modular Reactors (mini nuclear power generators designed to power AI)

Advice :

Short contracts, not purely from a risk of future, but more to have the ability to move quickly should we need to due to advance tech capabilities and disruptive commercials. And also, lets all start raising the question of Migration of Memory and Training, to help bring more awareness to this, as we are way ahead of regulation again on this point

Generative AI — bump in the road? or are LLM's reaching a limit?

Trying to keep up with the latest from the world of Gen AI is proving a full-time and tough task, such is the pace of change and the constant introduction of new features, that we often have to re-write this section many times before publication.

But here's a contradictory thought.... While progress is moving ahead at lightning speed, can we already see the end of the road for GPT and LLM's? Madness, we hear you cry, as we share news of the billions invested, 11x company valuations, Bubbles, and the increasing battle for AI supremacy, surely, we're only at the very start of this incredible journey...? Right ???

While we make this claim with some mischief, we wanted to share a few points to explain our thinking and observations here...

- The big solutions landed 3 years ago, and sure we wouldn't want to give them back, but has your firm given you a day back as AI is taking on more of your task?... Nope, its neat solutions but we're still very much needed and working hard.
- Hallucination rates have fallen dramatically, but it's still not and unlikely to ever be 100% accurate, meaning, while it's doing more heavy lifting, we're still needed in this loop to prompt it, or to mark its homework, especially in a corporate world.
- Building on the above, while 2026 has now seen 150,000+ tech layoffs (Meta 8,000, Microsoft 7%, Oracle 30,000), look closer and AI is largely the smokescreen for wider economic and trading pressures — these cuts fund AI infrastructure spend, not replace roles AI has actually taken. (see our thoughts on page 14)
- We're still struggling to find the default corporate fit. Many early 'self-build' AI solutions are already out of date/features and while Co-Pilot feels like the safe / right fit, two years in it has just 15 million paid seats — only 3.3% of its addressable market — and our people are still cutting and pasting from their own ChatGPT, Gemini or Claude apps, creating the world's biggest Shadow IT case study, and a possible future security headache.
- And the controversial point. In our opinion, GEN AI models are never going to move to true intelligence (AGI) regardless of how much investment, data, or compute power is thrown at them. Yes, GPT-5.5 and Claude Opus 4.7 plan, use tools and check their own work, impressive, but still sophisticated pattern-matching, not awareness or genuine reasoning. They will form part of the AGI tech stack, but it requires a fundamentally different type of AI which is still very much a work in progress, yet many still see it as a natural evolution. (see our thoughts on the middle ground on page 16)

So, as we pause to bring balance to this report, we'll leave the incredible minds in AI to continue to build true Artificial intelligence, and we move this section from ground-breaking innovations to more of a comparison Website update.... But that said, some of these updates will still impress and wow us, as possible the most important innovation of our time, now beds into its second phase of progress..... [So, what's new](#)

OpenAI (ChatGPT / GPT)

GPT-5.5 (April) is OpenAI's biggest leap yet a fully agentic model that plans, uses tools, checks its own work and keeps going across complex multi-step tasks without hand-holding. Strengths are clearest in coding, computer use, & knowledge work. GPT-5.5 Pro is now rolling out to paid plans. ChatGPT Images 2.0 also launched late April, claiming the top spot on every image leaderboard by a wide margin. Ads are now appearing for free-tier users in select markets, confirming the monetisation pressure is real.

Signal: The shift from 'chat' to 'do' is complete. GPT-5.5 doesn't just answer it acts autonomously. The ad rollout signals the gap between world-leading capability and a sustainable business model is still very much open.

Google (Gemini)

Gemini 3.1 Pro posted leading scores on 13 of 16 benchmarks, more than double its predecessor on ARC-AGI-2. Personal Intelligence now integrates with Gmail, Drive and Workspace apps for genuinely personalised responses. The Gemini app is now native on macOS. Deep Research Max is the market's most capable autonomous research agent. Flash Live adds real-time voice, now available across 24 languages in Google Vids.

Signal: Google's integration advantage is now its defining weapon — AI that already knows your calendar, email and docs without setup. The generational benchmark leap in 3.1 Pro suggests they are no longer chasing. The threat to Microsoft Office is becoming very real.

Anthropic (Claude)

Claude Opus 4.7 (April 2026) is the new flagship, stronger on coding, higher-resolution vision, built for long-running agentic tasks. Claude Design (Anthropic Labs) enables collaborative creation of visual outputs. Cowork is now generally available on Mac and Windows with enterprise controls. MCP has joined the Linux Foundation's Agentic AI Foundation as the open standard for AI tool integration.

Signal: Still the gold standard for coding and compliance-focused enterprise use. MCP's emergence as an industry standard is a quiet but significant power move, whoever owns the integration layer, owns the workflow.

Perplexity AI

Perplexity has pivoted hard from search engine to agentic platform and it's working. Revenues jumped 50% in a month, with ARR now topping \$450M. 'Perplexity Computer' (Feb) is a full computer-use agent orchestrating 19 different AI models. The API is now a full-stack agentic platform with Agent, Search, and Embeddings APIs. Integration with Snapchat now reaches nearly 1 billion users. A lawsuit alleging data-sharing with Meta and Google a concern.

Signal: Playing brilliantly in position. The pivot to agentic tooling and usage-based pricing is pulling serious revenue. The data privacy lawsuit is worth watching — citation-first trust is their brand, and any dent to it matters.

xAI (Grok)

Grok 4.3 Beta adds native document, spreadsheet and PowerPoint generation, and tighter integration with Grok Computer for autonomous desktop. The SuperGrok Heavy 16-agent system and 2M token context window remain key differentiators. A \$20B Series E has been raised. Grok 5 (6 trillion parameters, trained on the 1GW Colossus 2 supercluster) remains the most anticipated release.

Signal: Real-time X data and multi-agent architecture give Grok a genuinely different edge for trend analysis and live intelligence. Grok 5 could be a step-change, but the content moderation controversies are a recurring reputational risk that may limit enterprise adoption.

Meta (Muse Spark / Superintelligence Labs)

The biggest Meta story of April is not Llama, it's Muse Spark. The first model from newly formed Meta Superintelligence Labs (led by Alexandr Wang, backed by a \$14.3B Scale AI investment), Muse Spark is natively multimodal, features a 'Contemplating' parallel reasoning mode, and now powers Meta AI across all platforms. In a significant strategic pivot, it is closed source. Capital expenditure guidance of \$115–135B for 2026 signals this is no longer a side project.

Signal: Meta has torn up the Llama open-source playbook and gone all-in on proprietary AI. With 3 billion+ users across WhatsApp, Instagram and Facebook as the distribution layer, Muse Spark has the largest potential reach of any AI model ever launched. Watch this space.

AI Agents — The Great Corporate Hope...

Many of us are testing AI Agents to deliver the efficiencies the headlines promise, yet most projects remain paused or limited by uncertainty, errors, and security concerns. While we often take a measured 'fast follower' approach, the industry is accelerating at pace, big tech sees Agents as the route to monetise its AI investments, and for us they hold the promise of far greater autonomy and a real step toward true AI Assistants. For me this is the real opportunity for the future of work.

So, where are those signals and early services, set to change our future of work

While Agents were always a natural focus for Big Tech, their cutthroat competition on Gen AI has led them to focus on Agents market a little quicker than I think even they expected, to monetise their mass investments. This is causing huge disruption in the Agent AI sector, as big tech look to claim a piece of the action, and in doing so will make it super easy for us to use.. While the updates are making the press I still think they are some way off full disintermediation... but it's started and it's quickly bringing issues in security and governance we need to tackle now... or to be ready with the mop bucket very quickly.

Microsoft

The one to watch, as the 'safe solution' and switch to Partnering rather than building the right choice.

- Copilot Cowork - autonomous multi-step task execution across the entire Office ecosystem, and master stroke partnership with Anthropic to access Claude CoWork.
- Agent 365 went live 1st May at \$15/user/month a dedicated enterprise agent control plane for building, deploying and governing AI agents across the Office suite.
- Copilot Studio 2026 Wave 1 adds MCP connector support, a Shadow AI governance pane, and 'Hey Copilot' wake word for hands-free access.

Signal: The Claude integration is the boldest 'and best' move yet, signalling Microsoft will back whichever model wins. With Agent 365 now live and 41 Copilot Studio updates in April alone, the velocity is relentless. Don't bet against them... My prediction is they will move from 'behind' 'underdog' to leading safe Agents for work opportunity.

OpenAI

The company that started it all with ChatGPT, now fighting to stay relevant as the giants they awakened come for their lunch. Still innovating, but the moat is narrowing fast.

- Atlas browser with Agent Mode for multi-step autonomous task execution now integrated natively into ChatGPT. Operator evolved from standalone product into ChatGPT's core agent experience.
- AgentKit developer framework with Agent Builder (visual canvas), Connector Registry, and ChatKit, enabling enterprises to rapidly build and deploy multi-step automation workflows.

Signal: While they disrupted big tech in 2021, big tech caught up fast and they have war chests to play the long game.

Anthropic's enterprise revenue overtook them, Microsoft integrated Claude, and Google's platform play is formidable. GPT-5.5 is strong, but the agent war is far from won.

Google

The full-stack AI powerhouse. From DeepMind's labs to Workspace's 3 billion users, nobody else has the reach, the models, and the platform to make agents truly ubiquitous.

- Workspace Studio enables no-code agent building from plain language prompts. AI Agent Marketplace now features 120+ pre-built agents from Accenture, Deloitte and PwC. Agent Development Kit (ADK) launched for custom builds.
- Agent-to-Agent (A2A) open protocol launched, enabling multi-vendor agent orchestration. 200+ models in the Model Garden now include Anthropic Claude. \$750M partner fund announced to accelerate the ecosystem.

Signal: Agentspace launched as a unified enterprise agent hub, combining search, knowledge and agents in one control plane. NotebookLM Business went live with team collaboration and is fast becoming the knowledge worker's secret weapon.

Anthropic

Our heroes of AI, and 'go to' for most in the know. Taking ethics and innovation into Claude Code and now building this success model into Cowork and Managed Agents.

- Claude Cowork went live in Apr with full enterprise features, RBAC, OpenTelemetry observability, and Zoom MCP integration. Now available with enterprise controls
- Claude Managed Agents launched alongside Cowork. autonomous multi-step agents with early adopters including Notion, Asana and Sentry. MCP has joined the Linux Foundation's Agentic AI Foundation.

Signal: Still the gold standard for coding and compliance-focused enterprise use. ARR sprinted from \$9B to \$30B while competitors scrambled. MCP as an industry standard is a quiet but significant power move — whoever owns the integration layer, owns the workflow.

Where is this heading

This is likely to be where the real AI battle is fought. Gen AI interesting but for us Agents will enable our organisations, providing the real power and opportunity to truly automate work. For me it's the step before AI Assistants and imagine if instead of opening outlook this morning, you simply asked your agent via voice over coffee how your day looks... that's how my day now starts, and it's changed how I work forever.... take your seats... the 'real' race for AI is now underway.

Summary & Advice

- This one needs our urgent attention. The race is on to monetise, bringing new incredible solutions at speed.
- Solutions are coming ahead of controls and most IT firms are out of their depth on this.
- Our People are trying them out, causing a massive Shadow AI challenge, but also expectations of their use will follow.
- But this needs very careful review and deployment and Updates to our AI Policies, and I suggest a new DPIA

As big fans of spatial computing, we've always been realistic that the technology for the last 20 years has been little more than neat gadgets, with a VERY slight grip on some VERY Specialist niche use cases, in health, science and training... But two things have happened recently that have changed the game completely. Moving this from novelty to necessity.

It now works...really works

I had to jump into XR for the first time in a while to use 3 screens and was shocked at how well it worked. I used MetaQuest3 and Mac remote desktop, and spent around 3 hours in there. The screens were clear, the processing fast and I got a lot done in total immersion (no distractions) The gauge of how good it is... I've been back in for a few hours every week since. I think we've reached a tipping point of progress and not sure I want to give mine back now...
Video here:

Death of the Metaverse

Meta dramatically pivoted away from the metaverse in early 2026, announcing the shutdown of Horizon Worlds VR access on March 17 (before reversing 48 hours later under public pressure), with Reality Labs posting \$6.02B in losses in Q4 2025 alone, pushing cumulative losses beyond \$80B. The consumer metaverse dream is effectively over. What Meta has salvaged is its XR hardware strategy, cheap, capable devices feeding spatial data into AI world models. The metaverse brand is dead; the spatial computing infrastructure is not.

Open AI Device

OpenAI is developing hardware in partnership with design legend Jony Ive, a screenless, voice-first AI wearable (codenamed "Sweetpea") described as "sitting by a lake rather than walking through Times Square." Court filings confirmed the first device won't ship before late February 2027, later than originally planned. Initial production targets 40-50 million units through Foxconn. Additional form factors including a pen device ("Gumdrop") are in development. The device represents OpenAI's bet that conversational AI will become the primary interface layer, bypassing the smartphone entirely.

Vision Pro 2.0

We said this is possibly the most incredible tech device we've ever used. Yet we also called it out as a flop... even before it launched... why? The price! At £3,499 it tipped into a different level of affordability, pushing it out of what we call the adoption reach. Apple has since released an M5-refreshed Vision Pro at £3,199, a £300 cut, but still far from mass-market. While Apple might be the masters of innovation, they missed the importance of collision sequencing and the collective power of adoption to drive a product. Meta is playing the other card with affordable devices and over-demand in the US for a far cheaper product. No dramatic low-cost version has been announced yet, but our money is still on a lighter version, powered by Siri, powered by Google to feed Gemini... Come on Apple... I'm banking on you. And we have the share buy ready to go...

Spatial will fuel AI World Models

One of the limits of AI today is that its only as good as the data in its models or the prompts we use. The bit missing its contextual awareness, it doesn't understand our individual environments.. It's a one fits all.

A by-product of spatial computing is that it naturally sees or engages with our immediate environment, such as glass uses its tech to take images or video of our locations, this new insight is now being used to feed back into AI models and build its training, which will allow AI to progress greatly to the next stage... more to follow on this but we're ahead of the game once again.

Google Glasses are back

Google is launching AI-powered smart glasses this year through partnerships with Samsung, Warby Parker, & Gentle Monster both audio-only and display-enabled models built on Android XR with Gemini AI integration.

Google return to Glass is strategic rather than experimental: the glasses provide a real-time visual feed of the user's environment that directly feeds AI world models.

Meta's Ray-Ban AI glasses grew 210% year-over-year in 2025 Google has taken note and wants in.

Gemini Powers Siri

Apple and Google have formalised a multi-year deal in which Google's Gemini models will power the next generation of Siri, launching in spring 2026 via iOS 26.4. Apple is investing approximately \$1 billion annually and will use Google's 1.2 trillion parameter model while maintaining its privacy-first approach through on-device processing. This is a landmark moment as Apple conceding it cannot compete at the frontier model level alone, while Google secures the world's most recognisable AI voice interface. The convergence of Google AI inside Apple hardware is in line with our earlier prediction of Apple focusing on 'best access' to AI, and a win win for both firms

Neuralink

Neuralink has grown to 12 trial participants demonstrating computer and device control via thought alone, with reported expansion beyond the US. The company announced plans for high-volume production in 2026, with a breakthrough in automated surgery that threads electrodes through the dura without removal, dramatically reducing invasiveness. A \$650M Series E round in mid 2025 is funding the manufacturing scale-up. Signal: As Neuralink moves from proof-of-concept to manufacturing scale, organisations should monitor this as both a future accessibility tool and a signal of where human-computer interfaces are heading next.





Robin's First Waymo Ride

Atlanta, Georgia

RoboTaxis soon to be a reality in the UK

Waymo, the autonomous vehicle company spun out of Google's AI research, is bringing its self-driving taxi service to the UK, marking its first step into Europe. Operating fully electric, driverless vehicles under government-backed trials, the launch signals growing confidence in both the technology and the UK's regulatory direction for autonomous transport. Beyond safety and efficiency, the move raises important questions about how self-driving services could reshape mobility, independence, and accessibility.

Forum member Robin Spinks, Head of Accessibility at RNIB is part of the project and brings us the inside track on this important milestone.

So, Robin over to you.. "What's does this mean for the UK, accessibility and the 'future of us' overall?"

Our recent research indicates that mobility remains a huge challenge for many in our community, and more than half of blind and partially sighted people stated that they want to get out more.

It's critical that all new AV systems are designed inclusively so that every part of the service works for disabled people. This includes vehicles, infotainment services, and apps.

AV systems also need to take account of the fact that a blind person might not see where the vehicle is or even be aware of its presence as a quiet vehicle.

Autonomous vehicles can enable a new level of spontaneity and privacy when making journeys. You can't currently summon a taxi and then crank up your favourite playlist on the car audio system. You'll be able to do that in an AV. You might even enjoy a karaoke with your friends while on a journey or conduct a confidential meeting without the privacy concerns of a typical taxi.

Getting the systems to work well for everyone will be possible with genuine co design and deep community engagement and education. It might feel like a strange experience at first but once you get started, you'll quickly build confidence and appreciate the opportunity to enjoy a more private, flexible, and personal travel option.

Watch Robin's First Ride experience here - <https://www.youtube.com/watch?v=SSor67xxXhA>

R N I B

Robin, a founding member of our forum, is Head of Inclusive Design at RNIB, where he leads the organisation's inclusive design teams and drives innovative global partnerships, including many of us. With two decades across the disability sector, Robin brings deep insight into the needs of disabled people and is a recognised voice on smart tech, inclusion and accessibility.

As a forum, we gain real value from learning from each other, and Northern Accountants are doing some incredible work in this space. Adventa has been supporting their leadership team on the AI journey, helping shape an AI strategy aligned to their existing strategy and culture, identifying practical use cases, and proving real operational value, with the emphasis firmly on **AI assistance, not AI replacement**. Rather than share our own perspective, there's greater value in hearing it first-hand, so over to Ben Muncaster, COO at Northern Accountants, to share his experience and insights.

1. Why Explore AI Now?

My initial brief was that we needed an AI strategy and 'not to get left behind', so my first task was to understand what this really meant. Originally, the common consensus from within the business was to investigate and leverage AI solutions that were being bandied about from the various forums and 'thought leaders' but I never thought it was the right approach. Not that current off the shelf solutions didn't have a place, but more often than not, they were built around their processes, only part of the solution, or simply didn't live up to the hype. Ultimately, I wanted a strategy that would provide real benefits without compromising our data security. This led me to creating a cross functional team where we harnesses development capability, industry knowledge with lean Six Sigma methodologies and a delivery framework to ensure we had full control and oversight of the strategy. As with any new strategic decision, there has been a period of learning, but I feel from a standing start, that in 6 months we are now aligned for an even more effective period of delivery where the team will feel the true benefit of the work that has been put in

2. AI as Workforce Support, Not Replacement

We have a very strong people focus, and we engaged with our people early in the process, candidly sharing what AI is and isn't and how we planned to use it. Creating an internal working group and Improvement Champions is certainly something I would encourage. I know there is a fear amongst senior leaders that people are fearful for their jobs but I think have a progressive improvement culture and the sell to the team always being, we want to make their jobs easier and ultimately remove a lot of what they don't like doing. Also my focus has not necessarily been on headcount but really about freeing up capacity to aggressively out pace our competition by being so efficient in the way we operate that they are unable to compete with us. Yes there will be headcount savings from churn but right now it isn't the dominating factor

3. How did you Choose the Right Solutions

It's a big market and working with Ant, we were quickly able to whittled this down to the solutions that were practical, safe in terms of governance with respect of the data we utilise to service our clients, and within our initial automation budget. What I didn't want was to commit to the wrong solution but at the same time, ensure that we could deliver proof of concept projects to determine the route we would take. Looking closely at the solutions, from large providers, mid tier developers, through to agent platforms, and then comparing these to our requirements have provided us with the knowledge to pair the right solution to current/future projects.

4. What did the Trials Look Like in practice

We took the view to target small scale trials focusing on capability that we could reproduce easily for similar projects to be able to accelerate the returns on their successful completion. This was also an opportunity to test the effectiveness of the teams and the processes we had put in place to control them in terms of lessons learned. I think my main takeaway is that you have to manage expectations in those early stages as progress will be slow as you learn more about your teams, you settle on the processes and you stress test your due diligence on both security and ROI calculation.

5. Your Governance, Risk focus and Controls

Customer data and security were the non-negotiables, and we wanted to ensure that all risks were known, evaluated and controlled or we simply failed fast. We implemented a new process that focuses on strong initial assessments and due diligence with a number of check points and controls to move us safely from idea to deployment. Further to the above, this process of due diligence also (and crucially) is also carried out on all our business systems, partners and platforms we choose to operate with. We carry this process out in line with GDPR and ISO42001 frameworks and recommendations. All the time having a clear and documented audit trail of what we have done and how we have assessed each stage.

6. The Impact on Teams and Leadership

I think this is more of a question for the wider leadership community, but I do feel there is an understanding amongst progressive and bigger business that AI is an inevitable part of the future business landscape. The challenge now is to navigate how you implement this as a strategy in your business. In terms of what this needs to look like, my personal opinion is born from something that Ant said in the very early days that 'you don't need to know code to lead a AI strategy'. What I have learnt from working with so many talented people in this space, is that the need for a team with a wide range of skill sets to deliver this as an effective strategy is crucial. Having the ability to understand your processes, map and scope the future state, and control the delivery is absolutely crucial. I think this will shape future roles in business as success depends on a wider understanding of how all these different elements work alongside each other.

7. Your Advice to Other Firms

I can only really talk about this from my own perspective, but my advice would be to create a diverse team around you that complements you and each other. Some of these will be available to you internally but also look externally for the skills and influence you need. Spend your time understanding exactly what you want to accomplish and how, especially in terms of creating the culture in which this will succeed. Marry this alongside a strategy that is consistent, controlled and ambitious, and my belief is that you will achieve your goals. What would I do differently? Probably start off with fewer projects as managing multiple projects can slow down the turnaround time of success and understanding of where improvements can be made

Finally. One piece of advice for peers considering AI agents?

These are coming and will be highly effective within organisations. They are in effect digital apprentices so if you treat them as such - choose the right ones, spend time training them, audit and check their progress, then they will become integral to your businesses.

Thanks for sharing your insights, Ben, and we look forward to continuing to support Northern Accountants on your AI journey and raising the mischievous question.. 'will your competitors be able to keep up?'

Societal & Economic Drivers

While many expect a futurist's focus to be mainly around technology, while tech is often the key focus of our research, our methodology and process quickly kicks in to ask how will society use it? and as an economy how, and who, will pay for it? In practice, 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

ONS data to March 2026 shows CPIH inflation rising to 3.4%, up from 3.2% in February, driven by a sharp jump in motor fuel prices linked to the Middle East conflict. CPI hit 3.3%, with food inflation climbing to 3.7%. The Bank of England held the base rate at 3.75% in April, with one MPC member voting for a rise. GDP grew 0.5% in the three months to February, while unemployment fell to 4.9% from 5.2%, though vacancies dropped to their lowest since 2021. Payrolled employees fell by 65,000 on the year. The OBR has cut its 2026 growth forecast to 1.1%, down from 1.4%.

Inflation	GDP	Interest Rates	Unemployment
3.4% ↑ 3.2%	0.5% ↓ 0.3%	3.75% → 3.75%	4.9% ↓ 5.2%
CPIH to Mar	to Feb	to Apr	to Feb

Society

We are living through both the biggest technology advancement in history and one of the most significant societal transformations. While AI dominates the headlines, the quieter shift in how we live, work and spend is arguably just as impactful.

Key trends driving this transformation:

- Cost of living has moved beyond ownership for many young people – home ownership, car ownership and traditional retirement are now seen as unachievable without family wealth
- The 'Bank of Mum and Dad' is being squeezed further, reducing intergenerational wealth transfer
- Benefits are shifting from safety net to lifestyle choice for a growing number, as work struggles to justify its costs against rising living expenses
- City living is drawing young people in record numbers, yet making saving for a mortgage near impossible
- Banks are pivoting from mortgages to lifetime rental portfolios

The bottom line: within one to two generations, the wealth divide will widen to levels not seen since the early industrial revolution, with state reliance growing enormously, yet the UK state pension is funded purely from tax with no reserve fund. This is a trajectory we must monitor closely.

Cultural Impacts

Technology grabs the headlines. The economy shapes the outcomes. Society is the driver. But it is the collision of all three that is now rewriting the culture of our time, a new zeitgeist, a new norm, unfolding in real time.

Work, redefined: AI isn't just changing what we do, it's changing what work means. When routine tasks vanish, the social contract around employment, purpose and identity is being rewritten.

Trust as currency: In a world of deepfakes, synthetic media and algorithmic influence, trust is the scarcest resource. Organisations and individuals that can prove authenticity will hold the power.

The wealth gap widens: Those with access to AI tools accelerate. Those without fall further behind. This isn't a future risk, it's already the dividing line between who thrives and who gets left behind.

Culture follows capability: Every major technology shift has reshaped how societies think, communicate and value things. AI is no different, except this time, the speed of change is outpacing our ability to adapt culturally.

We are not watching the future unfold. We are living through a history lesson, written in real time.

If balance is needed anywhere in our current transformation, it's for our people, as headlines continue to talk of mass job cuts due to AI. Our provocation... is this real, or a smokescreen for wider pressures?



We're supporting significant AI projects that **could** without doubt automate masses of roles and transform organisations. But common sense quickly prevails, just because it '**could**' doesn't mean it '**should**', given AI's current capabilities, error rates, data security and governance. Not to mention the customer experience.

As we've covered many times, AI is still in its **infancy of impacts** and the risks to data, security, customer expectation and brand have resulted in many projects being paused or canned completely.

But the headlines say heaps of organisations are making significant cuts due to AI automation.. Right? Well, here's the challenge...are you making cuts due to AI? The potential is there, and many are pausing backfill, but the reality is that aside from a few obvious cases, AI-related cuts are very rare (outside of Big Tech). And if we look to Call Centre roles, the most obvious area for disruption, high natural churn means if you're having to make redundancies here, you must be doing something pretty incredible.. Or like our Klarna case study from the last report, playing an interesting strategy.

LLM's are not working in Contact Centres - Our research is highlighting that we're spending more time correcting AI than they are saving? Adventa



I could offer heaps of cases on AI impacts, but I'll let you answer this one yourself - how many job cuts has your organisation made due to AI? One for our next call. As a Forum we are 350+ leaders from over 130 Organisations. I trust our stats over any other on this one.

The real driver

Back on page 13 as always we share our views (at a political distance) on the macro economics shaping recruitment, along with wider trading pressures, social evolutions and new attitudes towards work, from young people considering alternative lifestyles given the reality of home ownership and pensions. Our prediction is these trends collectively will continue to be the real drivers for cuts, validated by ONS data showing employee payroll numbers fell by 33k in Nov and an estimated 43k in Dec.

Our own daily interactions confirm it's a tough market a sea of green 'open to work' profile pictures on LinkedIn, colleagues back on the market for a second or third time since their last long term. Its getting tougher out there.

Where is this heading?

Our assessment is there are two parts. Current AI is really souped-up RPA (robotic process automation) like automating factories in the 1800s, yes disruptive, yet inevitable progress brought balance overall. It's the next stage that will bring the real impact on work: AI assistants automating many of our daily tasks, where we'll still be involved but most of the leg work is done for us. It will be like every worker having a mini team of assistants.

Picture this : head to the golf course, set direction for your AI assistants, take a couple of 'check-in calls' to ensure work is on track, quick read-through, approval sent from the 19th hole... But is this realistic? As AI assistants bring huge efficiencies, will our leisure time increase, or will employers take it back and work us harder? We've debated this but struggle to land on a likely outcome. Society v The Machine... seconds out, ding ding, round 2.



In January, Grok hit the headlines. And so did my phone. I was called into a number of TV and radio interviews as the story broke, xAI's model making news for all the wrong reasons. But it was a question asked off-air that stuck with me: "If you were in charge, what would you actually do about Grok?" It's a question that deserves a straight answer. So here it is.

Let's Be Clear About What This Is

First, I don't believe Elon Musk set out to cause harm. What we're more likely looking at is innovative minds exploiting gaps in controls, and then some deeply troubled individuals using those gaps for malicious purposes. That distinction matters, but it doesn't diminish the severity of the outcome.

As a father, I want to say plainly: the harm done is not being fully acknowledged. A fine and a few days of bad press is not a proportionate response. Not even close.

And therein lies the core of the problem. Fines are not working. Negative press isn't working either, in fact, for a platform positioning itself as the unfiltered, rule-breaking alternative in the AI space, controversy is marketing. Many people who had never heard of Grok went looking for it precisely because of the press coverage intended to warn them away. There is no such thing as bad press when your brand is 'bad boy of AI'.

My Proposal – Block It

The answer, I believe, is right in front of us. Block it. Simple as that.

Block Grok in the UK until xAI addresses the clear and unresolved public safety concerns. And if they want to take their bat and ball home (which they might) let's take that stand. I genuinely believe they need us more than we need them, and our competitive landscape means UK users will find an alternative quickly.

I'll also get ahead of the free speech argument now, because someone always raises it. Imagine it's you, or your child in those images. There is no case for free speech here. Not in a civilised world. This isn't muzzling debate, it's applying the first rule of robotics, as Isaac Asimov gave us: do not cause harm to a human being. xAI is, indirectly, breaking its own foundational code.

Why this will work

Here's what gets overlooked in most of these conversations. AI is a race, a race to become the model people use every day. And that race is fuelled by us. Our usage, our posts, our interactions. In Grok's case specifically, it is being trained daily by user activity, and shaped by the vast data sets flowing from X (formerly Twitter), which provide the cultural and political Zeitgeist that defines what the model learns and becomes.

Switch that off, block access in the UK, and you slow progress significantly. The catch-up from losing that live data feed is slow, painful, and expensive beyond any fine currently on the table, and they would fall behind.

Secondly, and perhaps more importantly: we are creatures of habit. If Grok isn't available, we move on, and fast. And that is a far worse outcome for big tech than a regulatory penalty, because the real race isn't just to capture our usage, it's to hold onto it. Competitors will swoop in, learn our preferences, and make their services stickier. The window to recapture us closes quickly and rarely reopens.

On fines, I'm not saying abandon them. Go bigger. But redirect the majority of the penalty directly to those affected and the organisations supporting them. This creates something genuinely new: a societal shift, where legitimate citizen-led claims against big tech become viable and visible. Fight for the people, via the people. It's unconventional I know but new problems demand new thinking.

Could the World Join Us?

Australia has already shown appetite for this kind of stance. The EU has the regulatory architecture. If allies align, as they are doing on the under-16 social media ban, the pressure becomes existential rather than merely reputational for any platform that steps out of line.

And the competitive reality is this: if one platform exits the UK market, another will replace it within months, and the people will find an alternative for their all important (and valuable) attention. Trust me on this...the leverage is ours to use. The question is whether we have the political will to use it.

So, is blocking the answer to managing big tech? And should this be the conversation at Davos?

... I think it should.

Opinion_ Crazy or Credible



If **Doom Scrolling** social media was a sport, I'd certainly be a contender for its premier league. I can quickly move from a quick 'check in' to losing hours I will never get back, as my screen fills with funny pet videos, dance crazes, and ads for all manner of pointless gadgets I will never use. And in between these is a constant, growing and troubling splattering of aging experts offering a very bleak forecast for the future of AI. Be it the mass disruption of the jobs market and even the very future of mankind itself. And let's be clear these are big tech leaders, visionaries, ex google, ex Microsoft experts, who likely have far closer insights into the possible capabilities of AI than any of us. Which raises the questions. 'Are they credible and worth our attention?', or are they ex for a reason, and is their thinking a little too 'out there' for the big brands, or are their aims based on their own agendas set to influence our thinking and spending.... Or are they simply just now mad scientist level crazy?

I wrote a book about the future, and I would be a hypocrite to knock my peer's freedom of futurism, but quotes like this need fuller explanations

Humanity "needs to wake up" to the potentially catastrophic risks posed by powerful AI Systems **Dario Amodi**

The case for...

Some leading thinkers warn that AI's rapid progress could trigger **significant labour market upheaval and broader social strain**. Geoffrey Hinton, the "Godfather of AI," has cautioned that routine and even some knowledge-based work will increasingly be performed by AI systems, and that without policy responses like universal basic income, **wealth concentration and job displacement could intensify inequality**. He has also expressed concern that as AI approaches broader cognitive capabilities, humans could become **less essential in many roles**. Dario Amodi, CEO of Anthropic, has predicted scenarios in which AI could eliminate a sizeable portion of entry-level white-collar jobs and push unemployment higher if structural responses lag. Such warnings reflect broader anxieties that **economic disruption may outpace retraining and social safety nets**, potentially leading to societal unrest if transitions are not carefully managed.

The case against...

At the same time, experts grounded in current AI capabilities emphasise that **today's generative AI remains "narrow" and context-specific**, lacking the general reasoning and autonomy that define true artificial general intelligence (AGI). Andrew Ng, co-founder of Google Brain, has publicly called the notion that AI will imminently replace all jobs "*just ridiculous*", arguing that AGI has been overhyped and that, for the foreseeable future, **many tasks will remain beyond AI's reach**. Recent industry commentary also suggests that current tools are **augmenting rather than annihilating employment**, with CEOs of major tech firms highlighting AI's role in reshaping roles rather than eliminating them wholesale. Scholars studying AI advancement timelines generally find that transformative AGI systems truly rivaling human general cognition remains uncertain and distant, with significant technical obstacles ahead and estimates stretching decades into the future rather than months. In this view, the real challenge for society is not an overnight "AI takeover" but **responsibly integrating powerful tools while anticipating gradual labour and economic shifts**.

Our View

We sit somewhere between the two camps but lean towards the Against. As futurists we would be foolish to dismiss these very credible experts, but we're struggling to see the 'actual tech' that is going to drive the job disruption of wider tech capabilities. The reality is that Gen AI is a prediction machine and can never be 100% accurate, a simple fact that no amount of investment will ever fix. AGI is a different technology yet to fully materialise.

In balance in certain areas such as Customer Service, Data Entry, certain Admin roles, yes, I'm with them 100% and the wider aspect for me, is about evolution. We used to work in dangerous industrial environments to be replaced by machines, or we stood on production lines only to be replaced by robots, and while history tells us overall that disruption is inevitable, it also shows that outputs increased and in turn create more jobs than it replaced... I'd say it's simply our time, and we're masters of evolution. Rather than the catastrophes warned by these experts... The only real difference for me will be between 'the prepared and the unprepared'

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, work on some incredible projects, and attending some awesome events.

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



Festive'al of Innovation

Joining Angela (Head of Innovation) Martin (CIO) and the wider innovation team and their suppliers at their Xmas Festival of Innovation in Newcastle. A fantastic event and chance to explore some of the projects from the leaders of Innovation themselves. Thanks for having me along, always inspiring to spend time with you all.

Novuna Tech Bash 26



Thrilled to be back for another Novuna Tech bash in Leeds, providing an update on tech progressions since their last event. Another incredible event and fantastic chance to catch up with the team and some of our founding forum members Jo (Head of Software Engineering), Nick (CIO) and Andrew (CFO)

Cavell European Summit

Joining our good friends and Adventa partners Cavell at their European Summit in London. Taking along forum member Northern Accountant's Ben Muncaster (Director of Operations & Head of AI) taking to the main stage to share the progress of our AI Projects and Northern Accountants story of AI transformation. Cavell team, thanks for asking us along and Ben, thanks for joining me on stage to share your journey of AI transformation, which we're thrilled to be working with you on. And possible slightly biased.. But we have to say...it's pretty impressive.



Customer Innovation Forum – March 26

As many of you know I was a former exec at Virgin Media O2, for 29 years and the Founder and Chair of their own Innovation Forum, before leaving to set up Adventa. I'm very honoured to be asked to support them now as an independent on a number of projects and now as their official Resident Futurist, as we kick off a series of 3 session for 2026 focused on the topic of Connected Intelligence. Our March event homed in on Connected Intelligence and spatial compute supported by our friends at RiVR, before we explore the Future of Collaboration (June) and then Cybersecurity in the age of intelligent Networks (Nov).



Talk Radio TV

Joining Petrie Hosken a number of times recently on the topics of Social Media Addiction, Grok governance and the lack of LLM controls. My answer on both is simple, block UK access until we reach sensible measures. Too harsh? It would certainly prompt action as the biggest LLM feed.. trust me.



The Times

Joining Darryl Morris on Times Radio on the topic of Deepfakes. Asking the question how at risk are we from Deepfakes and how can we protect our identity as individuals, beyond the rich and famous brand protection cases. I predict a massive rise in corporate fraud from Deepfakes and will be providing further insight in our next report and advising on the steps we must take as corporate organisations.



A Partnership of
Friendship, Progress &
Support

A fitting close to this report is to recognise a partnership that means a great deal to me personally. Last year Adventa joined forces with Jo and Richard Groves founders of Active Digital, alongside Olympian Steve Backley, united by a shared mission to help customers navigate AI sensibly and efficiently. That partnership, obvious years ago, now goes from strength to strength, and it's wonderful to be working alongside long-standing partners I've known for 30 years and proudly consider friends.



Steve, Ant, Richard, Jo & Andy
at Apple HQ – back in 2016

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



adventa_

exploring the future, for our future

making the complex simple

and taking everyone on the journey with us

www.adventa.tech