

Artificial Intelligence

This blog collection focuses on artificial intelligence.

Is it time to start worrying about Artificial General Intelligence?




Image by neo tam from Pixabay

Concerns about the risks associated with run-away artificial intelligence superseding human intelligent (“The Singularity”) have been around for many years, with scientists as eminent as Stephen Hawkins expressing alarm. At SAMI we lean to the more technology-sceptic, in the sense that radical societal shift occurs far less frequently because of technological advance than is often suggested – flying taxis anyone? Technological change is often over-hyped in the short term, yet under-estimated on a longer timescale, being the accumulated results of successive incremental changes.

But recent advances in AI, such as large language models (LLMs), suggest that it might be time to re-assess. Clearly LLMs are nowhere near intelligent. But they do represent another step forward that could indicate more significant developments to come, even if as yet there



is no evidence that artificial general intelligence (AGI) is actually feasible. Only this week, OpenAI announced the release of GPT-4. Microsoft announced its search engine, Bing, would use GPT-4 customised for search. And Google have just released generative AI across its whole Workspace suite of tools.

Many experts however expect AGI before the end of the century. In the 2022 Expert Survey on Progress in AI, conducted with 738 experts, they estimated there was a 50% chance that high-level machine intelligence will occur by 2059. In older surveys (2012), 10% thought it would happen by 2022. So some scepticism is clearly warranted. Part of the argument for AGI relies on the exponential growth in computing power, set against the static capabilities of the human brain.

There is a problem in the terminology – what do we mean by “intelligence” anyway?

Historically, the definition of “intelligence” has a white-supremacist background, with the Stanford-Binet IQ tests being used to demonstrate the inferiority of non-white races. Howard Gardner’s theory of multiple intelligences extends to eight different types of intelligences consisting of: Linguistic, Logical/Mathematical, Spatial, Bodily-Kinesthetic, Musical, Interpersonal, Intrapersonal, and Naturalist. There is also Emotional Intelligence – the ability to perceive, control, and evaluate emotions – as popularized in a 1995 book by Daniel Goleman.

But let’s put aside our scepticism about the feasibility of AGI and consider the scenario where it looks as if it will soon come about. The singularity hypothesis, posits that a self-upgrading system will eventually enter a “runaway reaction” of self-improvement cycles, resulting in a powerful superintelligence that qualitatively far surpasses all human intelligence. At that point it becomes impossible to know for sure that human beings’ interests will be to the fore, and a wild range of dystopian stories have been written about the consequences.

Even AI developers, such as Sam Altman of OpenAI, recognise the problem and propose caution. Some are wary of his statement and want more immediate, co-ordinated action. Certainly it seems that waiting until AGI happens and then trying to take action is risky. So there is a groundswell of opinion that is pushing for an “anticipatory governance study” to get the conditions for AGI right before it is created.

The Millennium Project, a global think tank, is one such group. They point out that although there are international groups, such as the Global AI Ethics Institute, looking to get agreement about some of the difficult issues with today’s AI, there is not an equivalent for AGI. Without such agreement, the AI community is left to self-regulation, which in an era of competing power-blocks could result in an uncontrolled race with uncertain consequences. It’s true that international consensus is holding around major ethical issues in gene editing (just about), but leaving such a key issue to informal agreements also seems a poor answer.



SAMI Consulting

robust decisions in uncertain times

Getting global consensus in today's challenging geo-political environment would not be easy, even if technical conditions for the safe launch of AGI could be developed. Developing such a consensus on governance could easily take 10 years or more – by which time, what new advances in AI will have happened?

It would be easy to bury our heads and ignore the issue, to pray that AGI is actually not feasible, to leave it a while to see how things develop, or simply to say it's all too difficult and give in. But the whole point of foresight thinking is to explore the “what-if's” and to put in place plans that are activated at the right time. Is now the right time?

Written by Huw Williams, SAMI Principal

Published 17 March 2023



SAMI Consulting

robust decisions in uncertain times

Flying taxis crash into reality



Image by fajaws from Pixabay

In “The Fifth Element”, Korben Dallas is visited by a flying street food stall at his apartment window. Only one of any number of flying cars – including the one Bruce Willis is driving when Leeloo lands on it – the whole image of the cityscape is of high-rise buildings, and equally high-rise transport. From “Blade Runner”’s police ‘spinners’ to “The Jetsons”, airborne personal transport is a trope of the technological prowess of the future.

As so often, the world tries to catch up with science fiction. Subaru’s “air mobility concept”, Xpeng Aeroht’s “moon rover for earth”, Pivotal’s Helix – to identify three ventures from the last couple of days only – all join the flying car revolution. In Guangzhou, the pilotless EHang has just been approved as a flying tour service. As The Guardian tells us “Your flying taxi is just around the corner”.

Sir Stephen Hillier, chair of the Civil Aviation Authority, says the travel industry is at an “inflection point” ahead of, to quote the Financial Times, “the “widespread” adoption of electric vertical take-off and landing aircraft (eVTOLs) designed to carry passengers on short trips.”

So we’re all going to get a flying car, the streets of London will be full of flying taxis, and we’ll finally be free of traffic jams forever.



No we're not. Not now, not soon, probably not for the medium-term future. Possibility is, as so often, outpacing ability. Here's some simple reasons why personal eVTOL is a non-starter.

Safety. Having countless vehicles flying around urban areas would inevitably lead to accidents, crashes, and loss of power, raining bits of metal (and passengers) on the people below. Even if the vehicles were autonomous, and controlled by AI, one system failure would prove the power of gravity. Wind shear, lightning strikes, let alone drivers intentionally causing harm, all mean that any vehicle would need extensive redundancy systems and fail safes.

Regulation. There are at present no regulations for "urban air mobility" worthy of the name because there is currently no urban air mobility. Simply adopting the rules currently in place for helicopters would invalidate the whole attraction of the eVTOL, of simple, easy, air transport. New traffic management systems would need to be put in place. Integrating air into urban transport modes would require oversight, approval and regulation from aviation regulators, transport ministries, transport authorities; flying taxis would require separate passenger regulations...

Noise. If you live anywhere near an urban centre, you know the sound of cars, airplanes, helicopters, motorbikes. Now imagine your neighbour had a four-rotor personal eVTOL aircraft to go to work in. And so does your neighbour on the other side. If they're lucky enough to have a house, they may 'park' in their garden, but if you're in an apartment block, they may land on the roof. Intrusive constant noise accompanies pervasive adoption of the technology.

Infrastructure. Landing and takeoff points close to people's homes would be required – these are flying cars, remember, and people want to park outside their home. Charging, maintenance, parking – a whole new infrastructure mirroring that of vehicles on land. Assuming that electric vehicles become the norm, adding flying vehicles to the mix puts even more pressure on the electricity grid.

And air traffic control is going to be a nightmare. The freedom of the air ignores the fact that currently, aircraft are incredibly tightly managed in the air and on the ground. There is no "freedom" when one has a set route, set height, closely detailed maps, and an air traffic controller instructing every change in course and speed.

Batteries are not yet good enough to provide real range, and until they stop exploding, there will be a real concern over putting current battery technology in vehicles travelling hundreds of yards above people's heads (see "safety" above).

Weather. Conventional aircraft can struggle with some weather patterns, and smaller aircraft or airports can be defeated by nothing more than fog. Small, easily blown about personal craft will be subject to high winds, or weather extremes.



An option could be to use eVTOL outside of urban areas. After all, that's where, for instance, gliders are used. Even there, though, the challenges of technology, regulation, training of drivers, infrastructure and noise still persist. And who would want a flying car if you can only fly from your field to someone else's?

It doesn't work. Putting lots of airborne metal in the skies above cities even with trained pilots, and adequate control is a constant risk. The core challenges around safety, regulation, infrastructure, noise, training and technology have no obvious solutions. If we need to improve urban and national travel, and we do, then high speed rail, networked electric vehicles, liveable cities, 15-minute neighbourhoods are all better options.

Despite the hype and excitement, flying cars and flying taxis are, absent massive changes in artificial intelligence-controlled traffic management, safety, and law, impractical and unnecessary. We have many more important technologies we need for our future. Let's focus on those.

Written by Jonathan Blanchard Smith, SAMI Fellow and Director

Published 27 October 2023



SAMI Consulting

robust decisions in uncertain times

Going public: Cyberevolution, SAFIRE, and scenarios



Image by Pete Linforth from Pixabay

Regular readers of this blog will remember our pride at the publication of the European Commission’s SAFIRE project. SAMI’s participation included the development of 44 discrete scenarios – four global scenarios, and four scenarios each for ten global regions. Produced for the Directorate-General of Research and Innovation, SAFIRE was a challenging and complex project “to support the design and implementation of EU R&I Framework Programmes with materials and procedures that improve the state of art in foresight use in EU policy-making.”

Next week, I shall be in Frankfurt, at the Cyberevolution conference. Focussing on the junction of AI, digitalization and cybersecurity, Cyberevolution brings together industry leaders, chief information security officers, and cybersecurity professionals in a range of plenary and workshop sessions to understand cybersecurity in an AI powered digital world.

That world exists, of course, in the future – which is why I shall be there. I’ll be running a few workshops, taking part in a few panels, and trying to help participants explore the threats – and opportunities – that the next ten years might bring.



And this is where SAFIRE comes in. We shall be using the SAFIRE global scenarios as our starting point. The scenarios are highly credible: robust, research and innovation focussed, have been through a rigorous round of review and editing within the Directorate-General, and are both published and third party endorsed. And they flex to accommodate challenge – such as cybersecurity. Or, as SAMI did recently, to form the basis for a professional association’s strategy workshops. Or, as the SIP project is currently doing, in understanding ways to achieve decarbonisation and Net Zero.

This flexibility got me thinking. There is a vast potential to be unlocked in using publicly available scenario sets.

The essence of futures thinking is not wedded to any single domain but is underpinned by methodologies that can be applied across the spectrum. Whether in technology, education, or urban development, the same techniques such as trend analysis, futures wheels, and implication matrices remain equally relevant and potent. **It is not the subject matter that dictates the efficacy of futures thinking but rather the sophistication and application of these techniques.**

Good public scenario sets, such as those sometimes produced by think tanks, governmental bodies, and research institutions, can be a goldmine for understanding possible futures. These scenarios often provide well-researched, diverse sets of potential futures, each reflecting different variables and outcomes. By leveraging these resources, we can construct, deconstruct, and reconstruct possible futures with greater ease and faster than starting from scratch.

While public scenarios are not bespoke to a particular project, they offer a starting point that can be adapted and personalized. By applying creativity and insight, we can customise these scenarios, aligning them with the specific strategic interests of an organization or community.

Further, incorporating a wide range of viewpoints is crucial in futures thinking. Publicly available scenarios inherently encourage this diversity by representing a variety of perspectives and assumptions. This multiplicity ensures that futures work is not myopic but rather integrative and expansive, accounting for different values, beliefs, and worldviews.

There are obvious disadvantages, which include:

- No scenario set is going to be as good as one crafted specifically for a project, using teams of people who know their subject, facilitated by futurists who know theirs.
- Some scenario sets are of higher quality than others, and it can sometimes be perishingly difficult to work out the difference. Rigorous methodology, substantial horizon scanning and data, underpinned by serious experience, are all key. SAFIRE had over 650 discrete data points, numerous workshops, expert authors and expert review, for instance.



- It's important to respect the rights of the scenario's creators, incidentally. SAFIRE is published under Creative Commons International 4.0. Randomly grabbing scenario sets off the internet without guaranteeing you have the rights to them is a risky business – and rightly so. The best ones represented many thousands of person-hours work, and those rights should be protected.
- Whilst I am currently preparing to speak at an AI conference and a therefore somewhat biased, beware of the dangers of AI. I can see a time soon – or even maybe now – when AI is extensively used in the production of scenario sets. The essential elements of workshopping, and its associated “aha” moments, serendipity and group understanding, are so far difficult to package into silicon.
- Scenario sets age. That's almost the most important of the disadvantages, in fact – if you're working with a published set that doesn't fit into your time horizon, or fails to accommodate the history between its publication and now, you might as well not use them at all. SAFIRE was written during the pandemic, for instance – it understands and includes it – which almost no scenario written prior to 2020 could encompass.

For speed, budget, and diversity; for flexibility; and for the security of knowing one is working with credible, third-party endorsed output, publicly available scenario sets can be a useful base for further work. The potency of futures thinking does not rest on the domain of application but thrives on the robustness of its methodologies (and we shall have more to say on this particular topic in future blog posts). Public scenario sets can offer rich, varied, and accessible insights into the possible paths that lie ahead. Have a look at SAFIRE – or the results from the work at Cyberevolution and its Cybersecurity Council – to see how.

Note

Cyberevolution takes place at Kap Europa, Frankfurt, 14-16 November 2023. Physical and remote attendance is possible. Booking is [here](#).

Written by Jonathan Blanchard Smith, SAMI Fellow and Director

Published 10 November 2023