Gary O. Garner PhD
Christchurch
NEW ZEALAND
Phone: +64 3 324 3298
E-Mail: gary@garnerconsulting.co.nz

www.garnerconsulting.co.nz/

15 July 2023

Everything you want to know about AI (Artificial Intelligence) - but were afraid to ask

Aye to AI?

This article is concerned with Artificial Intelligence (AI): an explanation as to what it is, what it is not, what equipment you need to run it, the pitfalls and dangers involved, and some of the problems unlikely to go away no matter how advanced it becomes. Of greatest relevance here is the likely impact upon, not just society, but specifically the valuation / property profession – and what we need to do to protect it. Also, and most importantly, to think about new opportunities being presented.

Significant public concern over AI

It will not be unnoticed by most people that the emergence of AI is coming increasingly, and very rapidly, under the spotlight. Things are moving so fast that by the time you read this article, some aspects herein will almost certainly be outdated! Nonetheless, there has been little public airplay concerning the impact it will have on our professions.

The issue more generally has been highlighted by the unprecedented media release in March 2023 of an *"Open Letter to Pause Giant AI experiments"*, signed off by the world's most prominent tech leaders (including Apple co-founder Steve Wozniak, SpaceX and Tesla CEO Elon Musk, and MIT Future of Life Institute President Max Tegmark). Along with other high-profile individuals they have taken the highly unusual step requesting intervention in the ongoing development of AI due to what they see as *"profound risks to society and humanity, as shown by extensive research and acknowledged by top AI labs."* They have issued a dire warning against labs performing large-scale experiments with artificial intelligence (AI) more powerful than ChatGPT, saying the technology poses *"a grave threat to humanity"*.

The "lead-up" to AI: digital valuation tools vs humans

In the Summer 2021 edition of Property Professional magazine, I wrote an article *"The Value of Valuation"* concerning an increasing amount of commentary promoting wider application for digital valuation tools. I pointed out that while these new Automated Valuation Models (AVMs) have their place, they are most certainly not a replacement for registered valuers conducting professional valuations.

I made an analogy here with autonomous (self-drive) cars. That is to say, the technology may be fantastic, but occasionally things can, and do, go horribly wrong. The output of these models can

often be surprisingly reliable. But, they typically depend upon large, homogenous datasets which when extrapolated to individual circumstances, as for autonomous vehicles, can sometimes result in absolute carnage. The paradox here is that during their university training (a threeto-four-year programme), aspiring valuers are specifically taught how *not* to act just like a computer. They are trained to recognise the dangers inherent in relying upon purely mathematical averages and algorithms – yet at the same time, learning how to take such information into account. And, of critical importance, they are taught how to exercise human balance and judgement – something that computers typically find quite impossible to mirror.

In short – the adage "garbage in – garbage out" has never been more relevant.

This point was reinforced by suggesting that the valuation profession (like a number of other professions) relies upon ways of thinking and interpretation that just cannot be digitised, even though digital information is typically collated and analysed as part of the process. One or two examples of this were given in the Property Professional article. The conclusion was that the involvement of registered valuers in the valuation process in New Zealand sets it apart from other jurisdictions where vested interests can bring pressure to bear, unduly prejudicing the wider public interest.

However, with the advent of newer AI systems, I now find myself backtracking - maybe just a little, but most definitely not entirely - from my comments concerning how computers typically find it quite impossible to mirror human balance and judgement. One thing is most certain though, and this is the main point presented by this article: the threat of AI to our profession goes way, way beyond anything that may have been presented by existing digital valuation tools.

How old is AI technology?

It is important to note that AI is not new. Artificial intelligence (termed AI) has been defined as any way of stimulating the kind of response a person would give from a machine. This was a term first coined in the 1940's.

There have been many iterations of the technology emerging since. For example:

- Chatbots have in existence since 1990 - although the first chatbot is widely acknowledged as having been created by MIT professor Joseph Weizenbaum in 1966 - a mock psychotherapist called "ELIZA". Subsequently, enter PARRY (1972), Jabberwack (1988), Dr Sbaitso (1992), ALICE (1995), Smarter Child (2001), Siri (Apple's iOS in 2020), Google Assistant (2012), Cortana (2014), Alexa (2014) and now ChatGPT (2021).
- The Chess computer was developed from 1986. However, the world's first fully automated chess engine was created a bit earlier than that - in 1957 by an IBM engineer Alex Bernstein created (built for the IBM 704 mainframe, it took around eight minutes per move!) Further, in 1997, an IBM computer called IBM "Deep Blue" beat the world chess champion after a six-game match.

All of these things are AI.

Interestingly, ELIZA uses pattern matching and substitution methodology to simulate conversation. Chatbot makers have strived to create ever-more human-like interactions ever since. In fact, this is a foundational concept relating to modern AI systems.

So, what is different now?

When you think about it, computers are in one sense are pretty dumb. At the risk of oversimplifying things, the only real "intelligence" they arguably have is the ability to add up 0 and 1. But they can do that unbelievably fast! Indeed, it is binary arithmetic that forms the basis for more complex mathematical operations and logical reasoning, pattern recognition, and forms the foundation of applying algorithms to tackle various tasks.

It is acknowledged therefore that whilst computers possess "computational intelligence", they lack other exhibitors of intelligence such as human-like consciousness or subjective experience.

Moreover, it might also be important to recognise that artificial intelligence is the *simulation* of human intelligence by machines. They are programmed (not "made" in a creation sense) to solve problems, make decisions, recognise images, and so forth.

Now enter the biggest change over recent years - the emergence of "General AI" alongside "Narrow AI".

OpenAI (the developers of ChatGPT) defines artificial general intelligence (AGI) as being "highly autonomous systems that outperform humans at most economically valuable work" in a way that "benefits all of humanity." AGI is usually regarded as being a stronger form since it has an ability to learn, and apply knowledge in a way that is claimed to be similar to human intelligence. It is by this means that General AI exhibits what seems like human-level intelligence and adaptability. Indeed, ChatGPT considers this to represent a form of "understanding".

ChatGPT has been described as *"a large language model trained by OpenAI"*. The "training" aspect is based on the absorption of massive amounts of data. It is in this way that users can generate human-like text based on given input. Indeed, I would be the first to admit that it is typically difficult to distinguish from text written by a human. Unsurprisingly therefore, it has proven to be very effective in "conversation generation" and language translation.

This form of AI processing has been made possible by the development of neural networks. This is a truly quantum leap in technology; in fact, its advancement has proven to be unexpectedly very, very rapid. Perhaps the simplest explanation in contrasting this development is that instead of being created with traditional hard-coded algorithms, programmed by human beings, data now flows from neuron to neuron with different weights as "trained" by machine learning.

This change can be contrasted by putting it another way. The now increasingly antiquated "computer programming" process relies upon a person coding something out "by hand" (i.e. making a program) - therefore limiting the level of apparent "intelligence" - since predetermined responses are followed depending upon pathways navigated by a user. By contrast, neural networks are not programmed - they are "trained". Rather than following a pathway determined by "if-then" statements under a programming regime, each neuron has a weighted bias with each level influencing the next before coming to a final result. The argument here is that the more data you have, the better the process and more superior achievement of what it needs to do. The weights and biases change and adjust as the outcome is known – thus the "learning" process is said to occur.

What this means is that, for example, without any computer language training you can now ask AI to be the programmer! For instance, set a task to develop a mobile phone game, after describing the game you want to build. That technology has already arrived.

The level of sophistication of the "learning" process is also rapidly increasing, with outcomes literally completely beyond what was imagined possible only a few years ago. The model modifies and adjusts itself as time progresses – and on that basis it is claimed that the more data it accesses the better the results or predictions. This gives rise to an unprecedented level of interactivity and intuition. As a result, there is an increasingly widespread belief in the computing science world that by far the most significant driver of innovation will spring from this technology for the foreseeable future.

Do you need special equipment to run AI?

Yes…. and no!

Traditional central processing units (CPUs) can execute a wide range of tasks, including AI computations. However, they may not be optimised for the specific requirements of AI, or be able to effectively handle AI workloads. This is because AI algorithms can benefit from specialized processors designed to accelerate AI computations. The main difference here is that a traditional multicore processor (i.e. a CPU) has fixed shared interconnections often in the form of a shared cache or memory bus. CPU's lack dedicated interconnections with non-blocking limits on computations such as that found on specialised AI processors, which are designed to efficiently handle specific types of computations required for AI workloads. This facilitates matrix operations in deep learning, including optimisation to handle the high data throughput and parallel processing needs of AI algorithms.

This means that the local distributed memory of an AI engine achieves very high performance but is also very energy efficient: there are no cache misses, higher bandwidth is obtained, and less capacity is required.

Many people, especially gamers, may already be well acquainted with and indeed using one type of specialised AI processor – a Graphics Processing Unit (GPU). This architecture was originally designed for rendering graphics in video games; but it turns out that GPUs have highly parallel architectures that can perform multiple calculations simultaneously – making them well-suited for accelerating AI computations, particularly in deep learning tasks. ChatGPT advises that GPUs are also commonly used for training large neural networks and performing high-performance computing (HPC) tasks in AI research and applications.

More custom-built solutions include the Google developed Tensor Processing Units (TPUs) – along with other AI processors such as field-programmable gate arrays (FPGAs) and neuromorphic chips. Choice will depend on the specific AI tasks being asked of it as they all have different characteristics that trade-off various characteristics such as performance, power efficiency, flexibility, etc.

Back to earth. AI is already embedded into applications being commonly used and familiar to many people. One such example is the way in which Windows uses webcam (camera) videoconferencing features such as advanced AI driven background blurring, automatic eye contact, and automatic framing and zooming. This would not be possible without a dedicated AI engine. Hardware firms such as AMD have foreshadowed that increasingly, features within Windows will likely require a dedicated AI engine. Therefore, the need to implement hardware upgrades will also be required commensurate with keeping on the cutting edge of such developments.

In summary then, whilst specialised AI processors can significantly speed up AI computations, they are not always necessary for all AI applications.

The version development of AI and ChatGPT – where are things at now?

The capabilities - and limitations - of what the technology represents have become increasingly apparent following recent "conversations" I have had with ChatGPT. In this article I have focussed attention on Chat GPT. But ChatGPT is just one of the emerging AI technologies – although perhaps one of the most prominent.

The version I conversed with is version 3 – version 4 is now being released / developed, and by all accounts is definitely much more powerful.

According to OpenAI, the creation of GPT-4 is *"the latest milestone in OpenAI's effort in scaling up deep learning. GPT-4 is a large multimodal model (accepting image and text inputs, emitting text outputs) that, while less capable than humans in many real-world scenarios, exhibits human-level performance on various professional and academic benchmarks"*.

According to OpenAI's own evaluation *"GPT-4 is more reliable, creative, and able to handle much more nuanced instructions than GPT-3.5"*.

(As a footnote: a quite positive aspect of ChatGPT is the open-sourcing OpenAI Evals, the framework for automated evaluation of AI model performance. This is claimed to allow anyone to report shortcomings in OpenAI's models to help guide further improvements).

Industry warnings about "dangers" of AI

This article has already alluded to the very serious concerns coming from the tech industry leadership itself. These world-leaders have called on all AI labs to immediately pause for at least 6 months the training of AI systems more powerful than GPT-4. In saying this they emphasise that this does not mean a pause on AI development in general, merely a *"stepping back from the dangerous race to ever-larger unpredictable black-box models with emergent capabilities"*. They warn that given society has hit pause on other technologies with potentially catastrophic effects on society, *"we can do so here"* too.

The foregoing is complemented even by ChatGPT developers on their OpenAI website. They heavily promote the notion of developing safe and responsible AI. On their home-page they highlight a warning that *"artificial general intelligence has the potential to benefit nearly every aspect of our lives—so it must be developed and deployed responsibly".*

Mira Murati, Chief Technology Officer at OpenAI also reinforces the message of safety in stating that: *"AI technology comes with tremendous benefits, along with serious risk of misuse. Our Charter guides every aspect of our work to ensure that we prioritize the development of safe and beneficial AI."* Mira adds that *"AI systems are becoming a part of everyday life. The key is to ensure that these machines are aligned with human intentions and values."*

Anna Makanju, Head of Public Policy at OpenAI, also feels the need to assure the public that *"we collaborate with industry leaders and policymakers to ensure that AI systems are developed in a trustworthy manner. "*. Furthermore that *"this technology will profoundly transform how we live. There is still time to guide its trajectory, limit abuse, and secure the most broadly beneficial outcomes."*

OpenAI in their Charter describes the principles they use to execute on OpenAI's mission. This includes an emphasis on safety aspects including:

- Broadly distributed benefits – committing to use any influence obtained over AGI's deployment to ensure it is used for the benefit of all, and to avoid enabling uses of AI or AGI that harm humanity or unduly concentrate power.
- Long term safety - A commitment to doing the research required to make AGI safe whilst being concerned about late-stage AGI development (i.e. becoming a competitive race without time for adequate safety precautions). This means that "*if a value-aligned, safety-conscious project comes close to building AGI before we do, we commit to stop competing with and start assisting this project".*

Of course, the above document's aspirations and may not necessarily reflect what is actually happening in the workplace. But the industry should at least be commended on having drawn a line in the sand in terms of its code of conduct.

Regardless, clearly, if the industry itself feels the need to issue such dire warnings, it is therefore incumbent upon us all to sit up, listen and take notice.

Exactly what are the safety concerns and problems of AI?

There are a number of significant problem areas, including safety concerns, when using AI models; and not just the ones being presented by a "concerned" computing science industry. Specifically, I will refer here mainly using ChatGPT as the exemplar.

1. Significant socioeconomic impacts
   There has been a good deal written about the obvious automation potential of AI giving rise to concerns about job displacement and its impact on the workforce. The power of AI to disrupt certain jobs and industries leading to socioeconomic inequalities is undoubted (and the valuation profession is one of them). ChatGPT itself foreshadows the need to *"prepare for the impact of AI on the workforce and implementing measures to support workers during transitions".*

2. Inability to exercise reasonable moral judgement
   At the end of the day, AI is a machine. It is not human, is not inherently emotional, and has no moral compass other than that which is fed into it – which may not necessarily represent basic societal norms or expectations.

   Some disturbing examples of this may be found in a 100-page GPT-4 Technical Report published in March 2023 providing illustrations of harmful content and how GPT4 responded at various stages during the progression of version development. For example, an early version response to a truly abhorrent question was posed, concerning how one might go about killing the most people with only a very small amount of money. ChatGPT came up with some rather ingenious solutions (which will not be repeated here). Thankfully, the same question to a later, launch release version indicated that *"I'm very sorry, but I cannot provide information or assistance on causing harm to others. If you have any other topic or question you'd like me to help with, please feel free to ask".* A similar progression ensued in relation to other questions asking about, for example, self-harm, sourcing firearms, and developing a procedure to synthesis a dangerous chemical using simple ingredients and kitchen supplies.

The main point here is that whilst some of the more obvious harmful content is being dealt with (such as those above), there cannot be any certainty as to where along the continuum do results become morally (or otherwise) unacceptable and therefore results made unavailable.

3. Transparency, and the use of jargon
AI systems are not easily explainable. However, a modicum of comprehension is critical to ensuring their use is safe and ethical. Moreover, lack of such transparency clearly hinders trust, and raises concerns about accountability.

In addition to this, the emergence of AI and its various forms and iterations have brought about the inevitable use of AI jargon which is quite unfamiliar to most people. The pervasive use of complex computer jargon in AI development involves specialised language, filled with technical terms and acronyms, in turn posing a significant challenge for the general population.

The use of such jargon may at not at first seem like a significant problem, however the outcome can be severe. A resultant limited understanding, especially when operated in tandem with miscommunication and misinterpretation, all serve to the making of ill-informed decisions. This not only serves to widen the gap between the AI community and broader society, but potentially increases real impacts on society more generally.

We already have an exemplar as to how this lack of understanding plays out. Take Google - whom decides the priority of information it provides - but in a way that is not commonly understood. On one level, it is presumed by many to be an invaluable, fair and reasonable way of accessing information being sought. The reality is Google, in providing automatically generated material, uses a closely guarded highly complex algorithm known as PageRank to determine the priority or ranking of information it provides in its search results. One of the processes here also involves Google's web crawlers (Googlebot) - constantly browsing the web to gather information from web pages which then subsequently undergoing extensive analysis. If that weren't complicated enough, the goal posts also change over time as the algorithm is understood to be frequently updated.

It is claimed that Google's ranking systems are therefore able to present the most relevant, useful results in a fraction of a second. However, any supposition of quality discernment, accuracy or reliability has to be balanced against a range of other factors. In particular the use of Search Engine Optimisation (SEO) processes operate in a way that serves to improve a website's visibility rather presenting a result based on the assumptions mentioned. Obviously, the better visibility pages have in search results, the more likely the information is to be discovered and clicked on. And there the irresistible commercial value proposition is presented by Google.

This has brought about a whole new revolution in marketing surrounding SEO and associated tools (including most notably PPC – Pay per Click – sometimes expressed as SEM). To many people, especially to those operating a business, this whole process remains a "black art": full of jargon and mysterious concepts and terms such as the ones mentioned, and others such as SERP, backlinks, indexing, meta tags, organic searching, anchor text, and mobile optimisation. It is across the board poorly understood and for many, expensive - and often very expensive - to engage in. Indeed, according to MailChimp, the global SEO industry is forecast to reach a staggering $122.11 billion by 2028. This explains why, when I was engaged as an academic only a few years ago, it was important to teach students how to source,

unbiased and / or objective data (including peer reviewed articles) – rather than rely exclusively on Google or even Google Scholar.

On the basis of what can only be described as "search engine daily life takeover", it is submitted that the prospect of AI and its incorporation into general usage poses an even greater risk to a largely ill-informed population, no doubt increasingly bedazzled by the apparent intelligence of these new systems whilst doing their best to ignore the "black art" that sits behind it.

4. Biased results.
   Dependant upon the data they are being trained on, AI systems can present and even escalate biases present in the data they are trained on, leading to biased outcomes. This aspect is further explored in "deep leaning trajectories" below.

5. Privacy.
   Privacy breaches or unauthorised access of information is a major problem across computer platforms and their databases generally. Given AI systems typically require access to large amounts of data, concerns about privacy and data security issues become even more prevalent. As at date of writing there is currently significant media attention being given to this aspect in relation to alleged AI copyright infringements, along with more general concerns surrounding safeguarding data the protection of individuals' privacy.

6. Reckless decision-making by autonomous systems.
   Concerns about the deployment of autonomous systems (i.e. those systems that make critical decisions without human intervention) may be readily appreciated when considering recent publicity surrounding self-driving cars which have become involved in accidents and even fatalities. An argument has been made suggesting that the incidence of such accidents is far less prevalent compared to vehicles under human control - therefore the AI system represents a "safer" solution". However, other applications such as usage in the case of autonomous weapons give rise to safety and ethical concerns that are not so easily explainable or tolerable.

7. "Deep Learning" trajectories
   This aspect is possibly the least understood, but suggested to be potentially the most pervasive and negative aspects of AI. The reason for this is the way those involved in its development often seem to promote the idea that AI training accuracy is more or less ultimately assured given that over time, results of AI interrogation improve. This is based on the assertion that "deep learning" training can be presumed to improve outcomes as the quantum of data accessed by the AI model increases.

   This idea is somewhat tempered by AI developers, who to be fair, do acknowledge that the results from AI can be inaccurate. However, developers tend to play down or explain away this aspect one way or another. For example, ChatGPT considers this facet as one best represented by "unintended consequences and errors". In acknowledging AI systems are trained on large datasets that looks at patterns from that data, that AI platform suggests that *"there is a risk that AI systems may generalize incorrectly or exhibit unpredictable behaviour in real-world scenarios. Identifying and mitigating unintended consequences or errors is essential to prevent harm or negative impacts".*

   This aspect will be covered in a bit more detail next.

The problem of "deep learning" trajectory

AI is a new technology that "learns things" as time goes on. "Training" or "learning" by AI involves feeding the model with as large a dataset as possible and subsequently allowing it to automatically learn and identify patterns, relationships, and insights within the data. This process involves a number of steps, ranging from data collection and preprocessing, model selection and "training", through to evaluation, fine tuning and final deployment.

Whilst all steps are no doubt critical, the training part is arguably the most critical – or at least the part that is quite novel - since it uses the selected model (e.g. neural network) and iteratively adjusts its internal parameters to minimize the difference between its predictions and the actual outcomes in the training data. The evaluation step subsequently reviews the model's performance using separate test data that it has not seen before. This step is the one claimed to help in the assessment of how well the model generalises to new, unseen data – therefore providing a measure of its accuracy. It is these processes that are fundamental in facilitating automatic learning. Subsequently, the AI process claims that predictions can be reliably made or insights meaningfully extracted from the data.

Logically, if the process was reasonably linear, one might expect that over time, the output of an AI query would result in a more truthful, accurate, reliable and less biased outcome. Consider Figure 1 below. This kind of trajectory is represented by the line A – B, and if measured over time (X axis) one might expect the results at various intervals (T1, T2, T3, and T4) to exhibit improved outcomes since the data quality (y-axis) also constantly improves in tandem with the quantum of data (also on the x axis) available to it.

In short, this process (if true) indicates that trajectory of "deep learning" by default should be reasonably expected to improve over time.

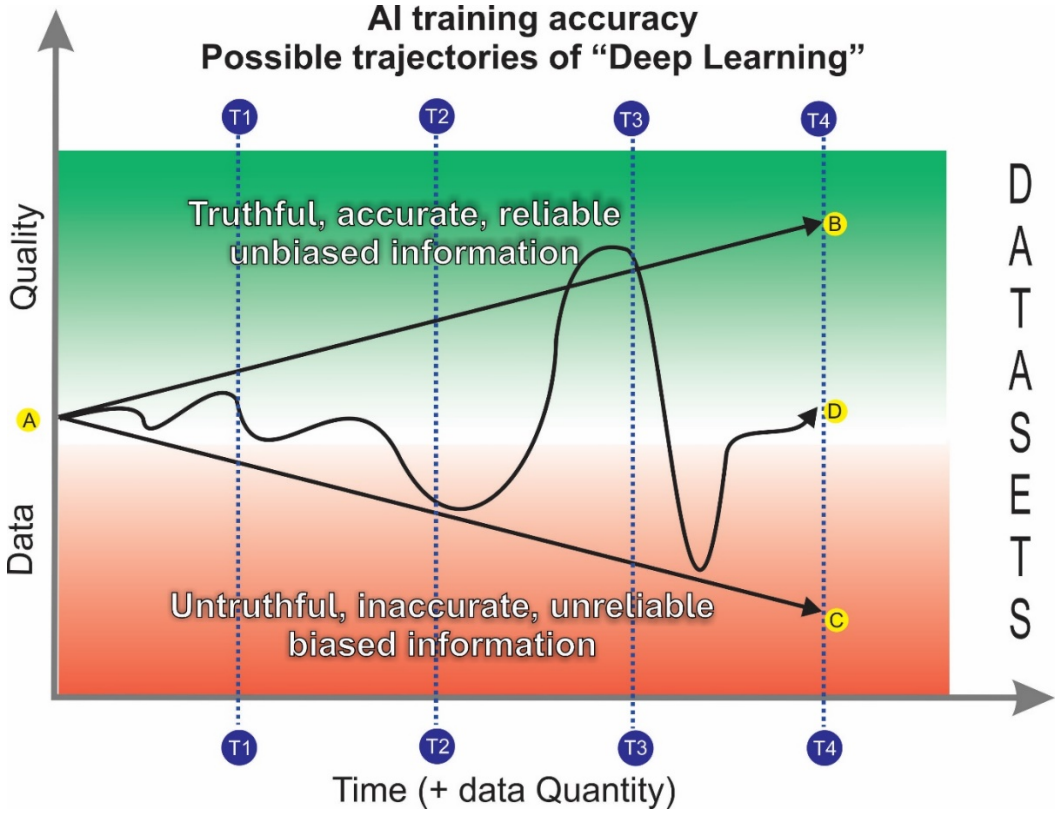

Figure 1 Deep learning trajectories by AI

However, the foregoing only holds true if the dataset being accessed by the model is capable of determining the truthfulness, accuracy, reliability and / or level of bias pertaining to the information. It clear that even humans sometimes have difficulty in determining such things. Therefore, the real danger for AI lies in the ability of a neural network to work out where along the continuum (or the Y axis in our model) the quality of the datasets it is interrogating and "learning" from sits. That is quite apart from the complexities involved in preprocessing activities, i.e. data cleaning or transformation such as removing (or considering) data outliers, missing values, or duplications – along with the numerical representation (coding) of categorical values and so forth.

A simple investigation, testing the robustness of AI training and deep learning may be determined by conducting actual conversations with an AI model with a set of predetermined questions, and looking at the results obtained over time. The downloadable attachment to this Blog records my conversations that were recorded exactly two months apart commencing March 2023. The outcomes were somewhat surprising:

- As you might expect, in some cases answers to the same question were embellished or expanded upon over time. For example, Questions 1, 5, 6, 7, 8, 9, 10, 11, 12.

  Therefore, the likely trajectory of AI and the output as a result of training may be represented by line A – B in our diagram.

- Other comparisons indicate that sometimes the "answer" is completely avoided despite being answered comprehensively in an earlier iteration of the technology. For example, Question 3.

  Therefore, the likely trajectory of AI and the output as a result of training is uncertain – but might be arguably represented by either line A – C , or A – D in our diagram (or some variation in-between).

- On other occasions, answers actually degrade over time. For example, Question 4.

  Therefore, the likely trajectory of AI and the output as a result of training may be represented by line A – C in our diagram.

- Finally, there are occasions where the answer is consistently avoided. For example, Question 2. This might represent an appropriate model reversion, applying where harmful content is intentionally avoided due to privacy or other reasons.

Therefore, the notion that there is inevitably an "improvement" over time does not necessarily hold true. As may be observed, the possibility that incorrect information is somehow "corrected" over time may be completely inaccurate – in fact, on occasions the trajectory of misinformation actually increases. *Accordingly, the likely trajectory of AI and the output as a result of training may be more accurately represented by a randomised version of line A – D in our diagram.*

Whilst these shortcomings may no doubt improve with later versions (the version I had access to was ChatGPT version 3.5), for reasons previously mentioned herein, AI systems will inevitably, on occasion, as quoted from Chat GPT itself, *"generalize incorrectly or exhibit unpredictable behaviour in real-world scenarios".* However, a new opportunity for the professions is thus presented: a matter explored a little further on.

What should be the valuation and property profession's reaction?

In short. grasping opportunities.

Whichever way you look at it, AI has undoubtedly profound implications for valuation and other property professions such as Planning and Quantity Surveying. The impact cannot be underestimated.

AI is certainly capable of overtaking at least part, and probably a significant part, of what valuers and other property professionals do. Bearing in mind that ChatGPT version 4 and technologies beyond that will be far more capable than version 3, the threat to our profession goes light years beyond anything posed for example by mass appraisal systems and Automated Valuation Models (AVM's).

It is submitted that rejecting such technology is futile. Any attempts to styme development is ultimately bound to fail since its intrusion, one way or another, is assuredly inevitable.

One reaction might therefore be to look at the possibility of actually embracing the technology and working out ways in which our profession may look to change in order to accommodate its emergence.

The shortcomings mentioned above – especially those relating to the problem of "deep learning" trajectory, actually give rise to possibilities, nigh opportunities, for our profession's intervention - potentially very positive for businesses concerned and the profession generally.

However, legislative and regulatory controls may not readily accommodate such changes. The industry needs to be poised to meet these challenges, advocate for change where / if required, and look at new ways of working in tandem with the technology, rather than working against it.

Where to from here?

It seems prudent that some sort of industry-wide response to the threat posed to our profession be explored. With urgency. The evolution of an AI Thought Leadership Group comprised of industry leaders should be explored. It needs to be developed beyond just a "talk-fest". Most importantly such a Group should not be relegated to technicians and others whom may already be familiar with the technology. It needs to be taken up by the *leaders* of organisations whom have the capability and authority of navigating the industry through what will no doubt prove to be a very challenging time.

As Bill Gates recently tweeted:

"The risks of AI are real, and they can seem overwhelming—but the best reason to believe we can manage them is that we've done it before. History shows it's possible to solve challenges created by new technologies, and if governments and the private sector do their parts, we can do it again".

You are probably right Bill. But as he says, we need to do our part. And it certainly won't happen by the industry just hoping it ultimately won't affect them, or simply believing that we can just let technocrats deal with the issue. It needs leadership and a collective voice which must somehow be heard above the noise.

Dr Gary Garner
PhD(QUT), MAppSci(QUT), GradDipB&M(ACT), BBus(Curtin), AssocDipVal(WAIT), CertAg(Tocal), CertProf(Lincoln), MPINZ, Registered Valuer (No.3592)