Revenue
Innovations

A B2B LEADER'S GUIDE:

STRATEGIC AI

FOCUSING ON PROBLEMS TO SOLVE RATHER THAN HYPE

"

**A REVIEW OF THE ISSUES
IMPACTING EXECUTIVES**

Authored by Ben Scoones

A CREATION OF REVENUE INNOVATIONS

www.revenueinnovations.com

The Certainty of Change

When OpenAI made ChatGPT-3 available to the public in November of 2022, with its astounding capacity to understand and interpret natural language prompts and produce often coherent and accurate responses, it became clear to many professionals - technical and non-technical alike - the colossal impact AI may soon have on their work, their organizations, and their industries

Excitement & Fear

While exciting and thought provoking for many, for some this development also elicits fear and uncertainty regarding the future of their professions and their place within them. This is not a concern only for individuals, but also for businesses - as industries change, the need to adapt to the new world of AI is also at the forefront of many leaders' minds. Of course such a thought is only natural when on the horizon of great upheaval, but the present is not unique for this. Technology has been developing for centuries, and in the past there too has been technological change so significant whole industries have been warped - consider the development of machinery in the industrial revolution, or the widespread use of computers in the 70s and 80s. Note, though, that we can look back on these examples and see how beneficial they have been for industry and the economy as a whole, allowing fewer to do more with less. Indeed, it is because of these historic developments that so much modern industry will be affected by AI, having been pivotal in facilitating a more knowledge-based economy.

AI, therefore, should be the cause of great optimism within the professional world, bringing with it not only great economic growth but also the opportunity for human work to be made more meaningful than ever.
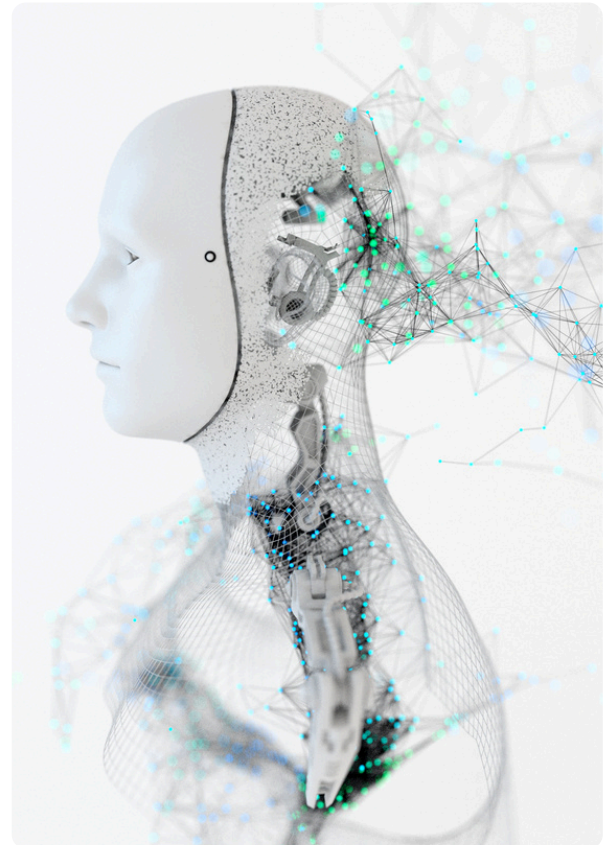
Take the Right Actions

With its rapid development and similarly rapid adoption, now is the time to act. In this paper, we will prepare you for doing just that. We will present a framework for how you can understand, prepare for, and adopt AI within your organization, reviewing its benefits and risks, to ensure that your organization not only survives, but thrives by its use.

The Nature of AI: Simple at Its Heart

The concept of artificial intelligence has been in existence for several decades and, while several definitions exist, it is, at its heart, a very simple concept: a machine which can perform similar cognitive functions to a human, whether that's mathematical calculations, playing chess, differentiating between cats and dogs, or even having a conversation - in other words, an "intelligent" machine. Of course, in this list there are some tasks which are harder than others. For many decades now (and several centuries in the form of calculators), computers have been able to perform mathematical calculations with unmatched accuracy and rapidity. The higher order problem of chess proved a harder challenge, but IBM's supercomputer Deep Blue famously defeated renowned chess player Gary Kasparov in 1997, almost three decades ago. The sophistication of machines for image classification has developed greatly over the last decade, with models becoming broader and more accurate, even being applied within fields such as healthcare for classification of medical images. Now, since the first public release of ChatGPT-3 in 2022, machines are capable of conducting conversations that are not only grammatically correct but, apparently, sufficiently sophisticated to convince some into believing the machine is sentient (Klimek, 2022). **Development of AI has indeed been dramatic, but this most recent advancement is importantly distinct to those that came before it.** First, however, we must understand more about the nature of such machines.

OpenAI's ChatGPT, and similar models such as Meta's LLaMA, Anthropic's Claud, and Google's Gemini, belong to a class of models known as large language models (LLMs), and LLMs belong to the class of AI called generative AI. **Generative AI itself is the concept of AI that is able to generate new content, whether that content is text, audio, images, and so on. LLMs are models that are built to understand language and therefore generate content in this first category and can also be interacted with through natural language.** LLMs are built by training the model on massive amounts of varied text-based data from which they learn the patterns in, and subsequently generate, language. Building an LLM therefore requires both a significant amount of data (more on this later) but also great skill and knowledge of machine learning (ML).

Illustration created by DALL-E

Democratizing AI: Foundation Models

While the barrier to building an LLM - or any other generative AI model, for that matter - is high, the barrier to using one can be very low. This is because m**odels like ChatGPT are foundation models, meaning they can be applied to a variety of use cases without need for further refinement or training.** Should a use case arise to which the model is poorly suited, the model can be fine-tuned by providing it with additional data from which to learn. In other words, the foundation model serves as the foundation for the new use-case driven model, hence the name. Foundation models therefore do a significant amount to democratize the use of generative AI: either the model itself is sufficient for your use case and so the barrier to entry is in learning how to utilize the model for the given purpose, i.e. prompt engineering; or it is in fine-tuning one of multiple available foundation models, in which case the technical challenge is greater in sourcing relevant data and tuning the model, but not so great as building one from scratch. Indeed, the ever growing number of fine-tuned models means it's quite possible a ready tuned model already exists for your given use case.

Increasingly, then, **the limitation on AI is not the technology itself,** but on two things: an organization's ability to adapt to the presence of the new technology and, relatedly, the additional data with which the model is provided. While the focus of many is on the former, the latter is of similar importance, and so we will discuss the essential value of data before shifting to the latter and primary focus of this paper.

Limitations of AI

A company's ability to adapt to new tech

The additional data the LLM is provided by the user

Data is Difference

At its most basic, data is difference, the distinction between what is and what could have been. It is that the coin toss came up heads, or that objects fall when dropped, or that consumers don't always make the optimal choice when given perfect information. What makes these things noteworthy is that each did not need to happen - there was the possibility that something different may have occurred instead. In some cases we are very aware of this, such as with a coin toss; when we toss a coin we usually do so precisely because there are two possible outcomes and one or the other can help make a decision or settle a bet. In others, the counterfactual is less present in our consciousness; the fact that

objects might not fall is something we rarely consider, but the laws of nature could have been different, just as they are on the moon, or in science fiction. The point is, the observation of some fact, some datum, informs us of how the world works, how people behave, and how events are related. Models use information in the same way.

Generative AI's ability to understand natural language, retrieve information, and generate content in such diverse contexts is truly impressive. **But, under the hood, what generative AI can do is, in some ways, no different to what the statistical and machine learning models of the past decades have done.** At bottom, these models **all learn relationships between variables and produce an output on the basis of what they've learnt.** Generative AI is widely applicable and its output non-numerical, unlike most earlier models, but how it's built is still based on that principle. As such, the performance of AI is, fundamentally and unchangeably, limited by the data - and specifically the information content present in the data - it has access to. The implications of this are significant for two reasons: unlike AI, data is not democratized; and not all data is of equal value.

What Makes Data Valuable

Ultimately, what makes data valuable to you may depend on a variety of factors, but some common ones include relevancy and novelty, as seen in this example, but also ease of use and scarcity.

| Relevancy | Novelty | Ease of Use | Scarcity |

Now, it is not quite true to say that all data is not democratized. In fact, a great deal of data is publicly available, and indeed chatGPT was trained, at least in part, on a large amount of publicly available data (OpenAI, n.d.). But there is much data that remains out of reach of the general public, whether because it is more technically challenging to access, behind a paywall, unavailable for purchase, or perhaps not even stored. If data is both unique and scarce, it can be of great value to those who can access and utilize it. **If your organization has privileged access to novel data, therefore, you have a competitive advantage when it comes to utilizing AI, as this data can be used to improve the performance of the model, or even create entirely new strategies and products because of the capabilities AI offers.**

Applications of AI

McKinsey estimates that generative AI could automate work activities accounting for 60-70% of employees' time and generate $2.6-4.4 trillion dollars of value across all industries (Chui et al., 2023). This impact is astonishing, but not surprising when we consider more the types of tasks generative AI is good at, and how prevalent they are in many western economies. Generative AI, by virtue of its ability to understand natural language, is particularly good at knowledge work. Specifically, this means it excels at work tasks which involve information in some way, whether that's gathering, summarizing, synthesizing, or applying it. And this makes sense given how models work. All models are built to identify relationships and patterns in the data on which they are trained.

LLMs are trained on a massive number of text documents of various kinds, so they are very familiar with the patterns in such documents and, by consequence, language in general. Since such a significant portion of our economy made up of knowledge-based, the scale of the impact estimated by McKinsey makes a lot of sense. To better understand precisely where AI will impact your organization's work, let's look at some knowledge-based tasks that are common in the modern workplace.

70%

Generative AI could automate work activities accounting for 60-70% of employees' time and generate $2.6-4.4 trillion dollars of value across all industries

1

First, **a large number of work tasks involve communicating known information to a new or different audience.** Emails can likely be largely AI-generated from a few initial prompts before some finer details of the AI's output are manually tweaked. Generating natural language product documentation or other artifacts from a codebase is similar and can potentially save even more time. While seemingly minor, work tasks of this type are common elements of almost every role, and AI's impact here is primarily in efficiency gain - making faster what was previously necessarily but administratively burdensome and rote. Such tasks might be classified as using generative AI to synthesize information for the prompt's stated purpose.

Second, **information gathering and summarization is a ubiquitous component of many roles**. Those working in research and development can more easily stay on top of new developments within their industry by having relevant articles or papers summarized, and filter through the mass of media to identify those sources that are worth more of their time. Those in sales and marketing will similarly be interested in tracking market trends, analyzing competitors, and preparing strategies. Finally, every role across a business will be involved in some number of meetings. Automatically generating meeting summaries saves time and, even more importantly, ensures no information from the meeting is lost. In such tasks, A machine can review and summarize far more data than a human ever could, meaning more information and therefore more informed decisions than would previously be possible.

2

3

Third is **information application which is leveraging AI and the information it has access to in order to solve problems and formulate new ideas.** While perhaps less obvious, this use of AI is perhaps also the most exciting, it opens up new possibilities for how people can work and how certain tasks can be achieved. Here, **AI is not just accelerating work, but making it better**. Consider, for example, personalized interactions with customers. An AI is able to access, almost instantly, information about company and customer behavior to give them a tailored service experience which a human, realistically, could not possibly replicate. Humans are still valuable and capable in ways that AI is not, but AI brings something entirely new to these processes to help solve problems.

These are some **general uses of AI** in work tasks which are likely already being performed within your company, **but generative AI also opens the door to other non-AI technology that may previously have been inaccessible**. Something we've already discussed regarding AI is the low barrier to entry for its use. This is in contrast to its technological precursors - statistical or machine learning models - which required workers to be both knowledgeable in their industries and sufficiently technically skilled to build, validate, and maintain such a model. These types of models will still be relevant moving forward as they are often applied to different types of problems to generative AI, such as demand forecasting, churn prediction, fraud detection, and a whole host of others. This said, generative AI's ability to write code significantly lowers the barrier to entry in these areas as well, allowing software engineers and data scientists to be more efficient and knowledgeable, as well as **allowing non-technical workers to engage with these forms of technology in a way that was previously unavailable.**

A **comprehensive approach to applying AI within your organization, therefore, will cover not only your current working processes** but also require consideration of how it may present **new opportunities in areas previously inaccessible given your organization's structure and means.** Furthermore, AI will give rise to products, services, and ways of working that are entirely new. This may direct a new business venture, dictate a change in strategy, or simply offer a way to utilize AI through use of external AI-powered services relevant to your organization. Maintaining an awareness of, and readiness for, such developments is likely to be another source of value for many organizations.
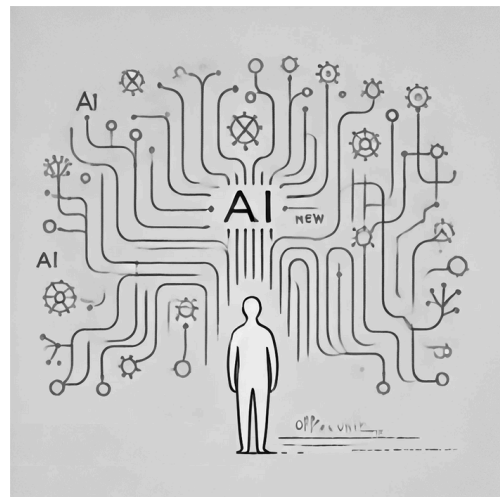

Illustration created by DALL-E

The Risks of AI

While powerful, AI is not without its risks. There are six main areas of risk relating to the use of AI, these being: a) data privacy; b) legality; c) opacity; d) quality of output; e) security; f) social. Depending on how you intend to use AI, it is quite possible that any or all of these may require consideration.

Data privacy is likely to be the primary concern for many users of AI. All generative AI models will initially be trained on some dataset in order to first build it. Data privacy here is not of concern to users since the model is not being trained on anything that is not either publicly available or approved by the model's creators[1]. However, once built the model can learn from and incorporate data it has seen. This can happen in a few ways. Short-term memory is a common one, where the model remembers for a brief period of time (such as the duration of a chat session) your inputs and considers these when providing answers to subsequent prompts. Long-term memory can also be implemented with some models where interactions are stored and referenced more or less permanently. Provided this data is only referenced by the model during your personal chat sessions, data privacy is of little concern (although data leaks should still be considered a risk as with all online activity). **However, if this data is somehow made available to other users through the model, this is a far greater issue, and can lead to your private data - whether PII, IP, or other sensitive information - being visible to unintended parties. This is a known risk if the model provider is using user submitted data to retrain their model at a future date[2].** While model providers are somewhat addressing these issues (OpenAI has stated that their models will not be trained on API data, for example (OpenAI, n.d.)), it is still in**cumbent on the user to ensure that their data remains secure.**

A similarly significant concern is legal issues relating to AI. While you don't want to leak sensitive data to the model, you also don't want to accidentally (or unknowingly) receive it from the model. **A lack of information regarding the sources the model used to generate its output can lead to unwitting use of another company's IP.** This applies not only to copyrighted material, but also to open-source code or software if the open-source licenses are breached (Ovington, 2023). Many models do not provide any information regarding the sources used to generate their response. Some model implementations make use of RAG (retrieval-augmented generation), a technique for supplementing the model's knowledge base with additional data sources, thereby providing the model with the ability to cite sources. Again, though, it is incumbent on the user to verify that the way the model's outputs are used is legal with regards to intellectual property and licensing regulations.

---

[1] Notwithstanding concerns you may have over your own digital footprint
[2] For example, the Samsung IP Leak. See: (Ray, 2023)

Issues of source attribution relate closely to the third risk in the above list: opacity. **Opacity refers to the fact that users (and often, in fact, builders) have little or no insight into how the model is generating its output. Naturally, this makes it far harder to identify when copyrighted or open-source material is being used,** but it also poses a challenge regarding the quality of the model's output. If you are unable to see from where and how the model generated its output, you either must choose to trust it as sufficient and appropriate for your intended use case, or else validate it yourself, reducing the efficiency benefit gained from using generative AI in the first place. Some model providers address this by listing some of the sources they use to train the model[1], and in cases of fine-tuned models this is part of the value proposition - that the data on which the model is trained or tuned is well suited to your use case. But, trust in the technology is still required.

Quality of output is the risk that users will perhaps be most familiar with. In many cases, if you know something about the topic you are interacting with the model about, you will notice when certain information is incorrect. This may manifest as incorrect details in a competitor summary or as malfunctioning code in a function which, i**f you are a domain expert, are often easy to identify and resolve. However, non-experts (those who are more likely to use generative AI for certain work tasks, such as writing code) will have a harder time spotting these issues**. In these cases, over-reliance on the technology can introduce bugs or incorrect information into your work that may prove to have insidious consequences. Furthermore, **inaccuracy is not the only quality related issue; depending on how models are built and on which data they are trained, they can also be biased.** This could be bias in the form of discrimination against a particular group, as a lack of representative data in a particular domain, or even in model behavior by favoring certain types of responses.

[1] For example, offering to sell a car for $1 (Notopoulos, 2023)

Do you trust your AI?

If you need to validate the output yourself it reduces the efficiency benefit gained from using generative AI in the first place.

Finally, like other great technological innovations before it, AI will have a huge impact on the workforce as a whole and the individuals within it. For many, the prospect of job displacement or reskilling is a real concern. **Challenger has tracked 16,989 job cuts due to AI since May of 2023 through Sept. 2024,** for example (Challenger, Gray & Christmas, Inc., 2024), and 49% of users of AI believe their jobs will not exist in 10 years (BCG, 2024). This concern is validated on the business side, with some companies anticipating significant employee layoffs as the technology continues to improve (Muirhead, 2023).

But the societal impact is not all negative. We've already seen that AI is expected to cause massive economic growth, driven in part by automation of previously human-conducted tasks and increased efficiency of workers. Coupled with this is the opportunity for workers to spend more time on tasks that are both likely more interesting and of higher value potential. Needless to say, **how AI is integrated into businesses will require considerable thought, particularly in relation to how it affects the often overlooked value of human capital within the organization.**
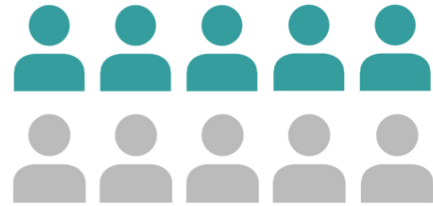
Nearly 50% of users of AI believe their jobs will not exist in 10 years

Preparing for & Adopting AI

With so much to consider with regards to how AI can positively impact your business, along with its potential risks, preparing for and adopting AI can be daunting, particularly for a company with little technical expertise. Here, we will lay out some concrete steps any organization can take to prepare for and ultimately adopt AI.

First, identify use cases within your existing organizational structure and work processes. Application of AI here is the low hanging fruit and may offer immediate efficiency gains with little investment, such as in writing emails or product documentation. Other use cases will require more investment and analysis of risk, but mapping out and estimating the value of AI in each case will be important for what you choose to prioritize and in how you shape your internal AI policies.

Second, consider how AI may encourage change in your existing structure, processes, and strategy. Particular use cases may require more investment into technology and skilled employees than others, perhaps even warranting this in areas in which your organization has not previously operated. Similarly, AI may give rise to new products or services that render certain roadmap items unnecessary, or which accelerate others either directly or indirectly by providing you access to previously unavailable technologies. This applies not only to current services/products but also to what's likely to exist in 3, 6, or 12 months time.

Third, consider the risks associated with each of your use cases and what internal policies need to be defined before each use case can be enacted. Any use of AI will require some input of data, so data privacy will be a concern for every use case, but simple policies regarding the type of data that can be shared should address this in many cases, i.e. for writing internal emails. In others, such as in using AI to generate code, opacity, use rights, and quality of output will also be concerns. Other applications may necessitate sharing of private data in some capacity, whether with an AI directly or through an intermediary such as a platform, which is relevant to cases like personalized customer interactions or scoring of sales leads. The effort of establishing and enforcing policies, the potential liability, and the effort of establishing the necessary technological infrastructure must, in each case, all be weighed against the short and long term value AI can provide.

The day to day of business and the nature of work is changing rapidly. Hopefully, this paper will have given you a greater understanding of how and why AI is the cause of this change, and ultimately how you and your organization can utilize it to accelerate your business goals and growth in the midst of it, rather than shrinking away. To make this more concrete, a worksheet is provided which you can use to map some of the aforementioned steps to your organization in order to begin to process. If the barrier to using AI still feels too great, but you understand its potential, then you need not do this on your own - we're here to help.

References

Klimek, M. (2022, June 11) Google Engineer Blake Lemoine thinks its LAMDA AI has come to life. The Washington Post. https://www.washingtonpost.com/technology/2022/06/11/google-ai-lamda-blake-lemoine/

OpenAI, How CHATGPT and our language models are developed, OpenAI. Retrieved February 16, 2024, from https://help.openai.com/en/articles/7842364-how-chatgpt-and-our-language-models-are-developed

Chui, M., Roberts, R., Yee, L., Hazan, E., Singla, A., Smaje, K., Sukharevsky, A., Zemmel, R. (2023, June 14). The Economic Potential of Generative AI: The Next Productivity Frontier. McKinsey. https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier#introduction[1]

Ray, S. (2023, May 2). Samsung Bans ChatGPT Among Employees After Sensitive Code Leak. Forbes. https://www.forbes.com/sites/siladityaray/2023/05/02/samsung-bans-chatgpt-and-other-chatbots-for-employees-after-sensitive-code-leak/?sh=6d76b8f36078

Ovington, Tristan (2023, November 22). 7 AI legal issues and how to deal with them. Walkme. https://www.walkme.com/blog/ai-legal-issues/

The Allen Institute for Artificial Intelligence. Open Language Model: OLMo. Retrieved February 16, 2024, from https://allenai.org/olmo

Notopoulos, K. (2023, December 18). A car dealership added an AI chatbot to its side. Then all hell broke loose. https://www.businessinsider.com/car-dealership-chevrolet-chatbot-chatgpt-pranks-chevy-2023-12

Challenger, Gray & Christmas, Inc. (2024, February 1). The Challenger Report January 2024 and October 2024 https://www.challengergray.com/blog/job-cuts-announced-by-us-based-companies-surge-136-to-82307-to-begin-2024-financial-tech-lead/

PwC (2022, May 2024). What 52,000 People Think About Work Today. PwC. https://www.pwc.com/gx/en/issues/workforce/hopes-and-fears-2022.html[2]

Muirhead, C. (2023, May 18). Is This the Start of the AI Jobs Bloodbath? As BT Revelas it Wants to Replace 10,000 Staff with Artificial Intelligence, Fears Grow Over the Impact on Workers and Customers. Daily Mail. https://www.dailymail.co.uk/news/article-12100605/As-BT-reveals-wants-replace-10-000-staff-AI-fears-grow-impact-workers.html

ARE YOU MAXIMIZING AI FOR YOUR SALES & MARKETING TEAMS?

Reach out to discuss your AI strategy

info@revenueinnovations.com

www.revenueinnovations.com