January 2026

AI Unlocked: A Practical Guide to Generative AI

AI and generative AI: A beginner’s white paper for CA’s
Introduction

The only constant is change, and this has never been more evident than with the transformative impact following the release of ChatGPT in November 2022. Although Artificial Intelligence (AI) has been studied since the 1950s, the advancements since 2022 have propelled it into a significant technology, especially in the realm of knowledge work, dramatically altering the way we work across various industries.

Recent developments, particularly in Generative AI (GenAI), have enabled machines to not only analyse but also create content such as text, images, and code, thereby unlocking new avenues for productivity and innovation.

Professionals are increasingly recognising the importance of these technologies, as they can enhance tasks ranging from drafting reports to designing products and automating tasks that previously required human judgment through agentic frameworks and multimodal capabilities. In essence, AI and GenAI are becoming indispensable tools in the modern professional's toolkit, offering unprecedented efficiency, creative support.

As AI continues to evolve, it has sparked essential discussions about its limitations and responsible use. Professionals across various sectors, including finance and healthcare, recognise that while AI can enhance productivity and offer innovative solutions, its application must be approached with care. This white paper introduces fundamental concepts of AI and Generative AI (GenAI) in a beginner-friendly way, detailing their mechanisms, applications, and best practices for ethical and effective usage. No matter your field, understanding the potential and pitfalls of AI is crucial for leveraging these technologies while fulfilling your professional responsibilities.

Understanding AI: A brief history and types

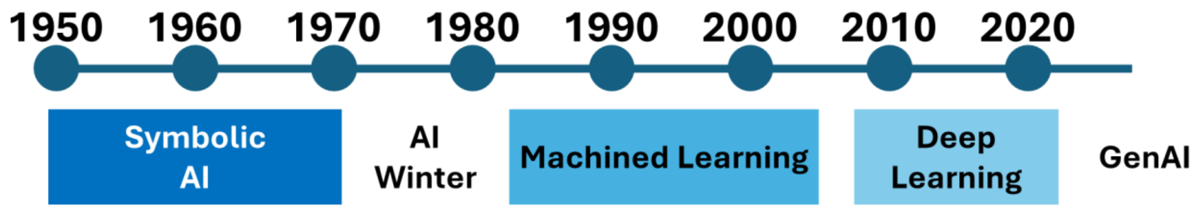

AI is not a single technology, but a broad field of computer science aimed at creating machines that can perform tasks requiring human-like intelligence. The field has evolved over decades.

Symbolic AI (1950s–1980s)

These were rule-based systems programmed logical rules and facts into an inference engine that could derive answers. For example, an expert system might contain hundreds of if-then rules to diagnose equipment failures. Symbolic AI achieved some successes but proved brittle, as it could not learn or adapt beyond its fixed rules. This led to periods known as “AI winters” when progress stalled.

Machine learning (1980s – 2000s)

Machine learning (ML) took off in the 1990s and 2000s as data availability and computing power grew. Instead of manually coding every rule, ML algorithms allow computers to learn patterns from data. Developers feed examples to algorithms (e.g. thousands of labelled emails), and the system statistically “learns” how to categorise or predict outcomes. Early ML methods included decision trees, support vector machines, and logistic regression. A key advantage of ML was that it could handle tasks that were impractical to define with rigid rules like recognising faces or spam.

Deep learning (2010s)

Within ML, a subset called **deep learning** emerged, inspired by the structure of the human brain’s neural networks. Deep learning uses multi-layered neural networks to automatically learn representations of data. Although neural network concepts date back to the 1960s, only recently have we had enough data and computing power to train very large networks effectively.

Deep learning’s rise was fuelled by the collision of big data, improved computing capacity and algorithmic innovations opening the door to the next generation of AI.

Generative AI (mid-2020s)

Generative AI (GenAI) refers to models that generate new data (text, images, audio, and video) rather than just predicting a label or outcome. It is a product of deep learning evolution. In fact, GenAI systems are often large foundation models trained on vast datasets, making them capable of producing remarkably human-like creations. GenAI models represent a “step-change” in AI: Unlike earlier narrow models, they can generalise across multiple tasks and handle unstructured data (text, images, code) all within one model.

Breakthroughs like the Transformer architecture (introduced by Google in 2017) allowed training of extremely large language models that can compose coherent text and even solve problems by “understanding” context. Likewise, new approaches for image generation (from Generative Adversarial Networks in 2014 to Diffusion models around 2021) have enabled AI to create photorealistic images from scratch. The rapid rise of GenAI in the last few years comes from this confluence of large datasets, advanced algorithms (transformers, diffusion), and immense compute availability (tens of thousands of GPUs in data centres) – all accelerating progress on a monthly basis.

In summary, AI has evolved from rule-based systems to learning systems, to deep learning and now generative models. GenAI sits at the frontier, leveraging deep learning to not only perceive patterns but produce novel outputs, and that is why it’s capturing so much attention today.

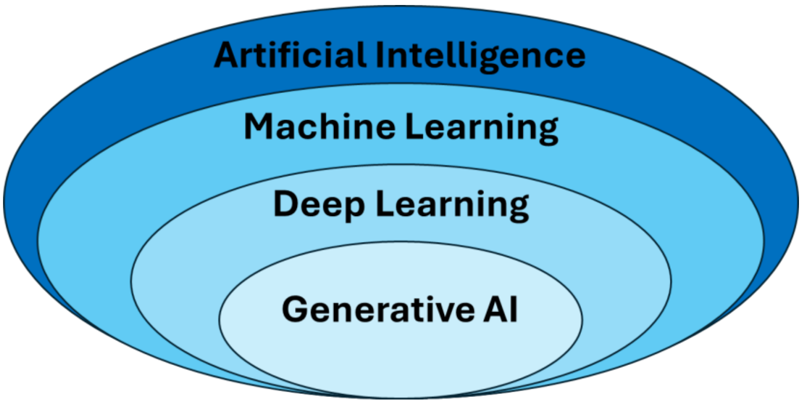

How GenAI works

Generative AI may seem magical – producing a fluent essay or a realistic painting at your request – but under the hood it relies on two key technical approaches: **Transformer models** and **Diffusion models**.

Transformer models (text generation)

Modern generative text AI relies on transformer networks, first popularised by the 2017 paper “Attention is All You Need.” These transformers use an attention mechanism to efficiently track how words relate to each other in parallel, rather than processing them one by one. This design captures context more effectively and scales up to handle enormous datasets.

A standout example is one of the very first large language models, GPT-4 Turbo Turbo, with 175 billion parameters, trained on hundreds of billions of words. In its pre-training phase, it learned grammar and facts simply by predicting the next word. It can then be fine-tuned or improved further with techniques like reinforcement learning from human feedback (RLHF), where people rate AI outputs to guide better responses. Thanks to this approach, models like GPT-4 Turbo Turbo and GPT-4 Turbo generate human-like text for a wide range of tasks — from answering questions to writing code — though they require massive data, computing power and investment to train.

Diffusion models (image generation)

Diffusion models generate images by learning to reverse the process of adding random noise to an image. During training, the model repeatedly adds noise to images, then practices removing it to restore the original. Once trained on millions of images, it can start from pure noise and denoise step by step until a coherent image forms.

It's like gradually blurring a photo until it's just static, then running that process in reverse to reconstruct the picture. Stable Diffusion XL and DALL·E 3 use this approach for text-to-image generation. You provide a text prompt (e.g., "castle on a floating island, oil painting style"), and the model refines random noise until it matches patterns for "castle," "floating island," and "oil painting." This method often produces high-quality, detailed results. Beyond art, diffusion models also support tasks like inpainting (filling in missing parts) and upscaling images by learning how pixels relate to each other.

The role of data, computing, and tuning

Crucial to both transformers and diffusion models is the scale of data and computing power. Generative models are usually trained on very large datasets. Language models ingest text from books, articles, websites. Image models train on billions of image-caption pairs scraped from the web. The diversity and size of these datasets give models a broad base of knowledge – but also introduce concerns (the models may pick up biases or errors present in the data, as we discuss later). Training is done on powerful hardware: Thousands of GPUs (Graphics Processing Units) or TPUs run for days or weeks.

After initial training, organisations often apply **fine-tuning** to specialise a model. Fine-tuning means taking a pre-trained model and training it a bit more on a narrower dataset or task. For instance, a general language model might be fine-tuned on legal documents to better handle legal questions, or on conversational data to become a friendly chatbot. This is far cheaper and faster than training from scratch. Fine-tuning can also involve reinforcement learning or techniques like prompt tuning (finding the right prompt phrases to elicit better answers without changing model weights).

Challenges and limitations of GenAI

While generative AI is impressive, it is far from perfect. Understanding its limitations and risks is critical for responsible use. Key challenges include model biases, "hallucinations" (factually incorrect outputs), lack of true understanding and ethical pitfalls.

- **Bias and fairness:** GenAI models learn from vast datasets that inevitably contain human biases, ranging from cultural, gender, racial, etc. As a result,

models can reproduce or even amplify those biases in their outputs. For example, an AI autocomplete might associate certain jobs with one gender or produce stereotypical images when prompted with certain professions.
This isn't because the AI has intent; it's statistical absorption of patterns in training data. However, the impact is real: Using biased AI in hiring or legal decisions could lead to unfair outcomes. Additionally, if not carefully designed, generative models might lack diverse perspectives – for instance, predominantly English training data could make them less accurate or useful for other languages and communities. Both the makers and users of GenAI have to be conscious of these biases. Techniques like dataset curation, bias testing, and fine-tuning with more representative data are used to mitigate bias.

Content filters and guardrails can be used to prevent blatantly hateful or discriminatory outputs. But no filter is foolproof. It's crucial for users to be aware that AI outputs aren't neutral truth – they reflect the data's skew. Therefore, in applications like lending decisions or CV evaluations, one should avoid blindly trusting AI without implement safeguards and checks for biased behaviour. Responsible AI means ensuring the technology's use does not end up reinforcing societal inequalities.

- **Hallucinations and accuracy issues:** Generative AI has a well-known tendency to "hallucinate". This means when the AI produces an output which sounds plausible but is false or nonsensical.

  These errors happen for a few reasons. Sometimes the AI's training data might have had inaccuracies. Other times, the model might not have specific training knowledge but is able to produce the nearest combination of words based on the guess that fits the prompt. The transformer's architecture doesn't have a built-in fact-checker; it just generates text that looks right based on patterns. For example, if asking for a biography of a person with limited information, it might fabricate career details because it "knows" how a biography is usually written.

  Hallucinations are problematic in professional use – CA's, solicitors or doctors cannot afford fabricated citations or diagnoses. These hallucinations amplify the current need for GenAI outputs to be verified by humans, especially in high-stake use. Solutions include "human in the loop" controls and retrieval-augmented generation (providing the model with verified reference text to ground its answers). An increasingly common practice is to use a two-tiered model approach or "judge" – one generates an answer, another model (or humans) check it. The use of "models as a judge" can support larger scale processes, which are not practical for human in the loop review, this can introduce additional and unintended consequences.

  The use of GenAI assistance does not negate professional responsibilities and best practice and healthy scepticism: Treat AI answers as a draft or suggestion, unless independently confirmed.

- **Lack of explainability:** Deep learning models, including GenAI, operate as complex black boxes with millions or billions of parameters. This means that when an AI gives a certain output, it's often unclear why it produced that result. The model can't easily provide reasoning or justification in a transparent way (any explanation it gives is itself a generated output, not an actual trace of its computation).

  This lack of explainability is a challenge in domains that require reasoning or trust. For example, if an AI medical assistant suggests a treatment plan, doctors need to know the rationale – but the AI can't show a clear chain of logic like a human expert could. Similarly, if an AI denies a loan application, regulations would require evidence for the decision; a black-box model makes this difficult.

  In practical use, the best mitigation is limiting use of GenAI in scenarios where a rationale is needed or ensuring a human review and provides the reasoning. In critical decisions, think of the AI as an assistant that suggests an answer while the human expert remains responsible for the reasoning and final judgment.

- **Ethical concerns:** Generative AI raises a host of ethical issues. One concern is the potential for misuse in generating harmful content – such as very realistic fake news, deepfake images or videos, or automated spam and phishing emails. This puts an onus on AI developers to incorporate safety limits (for instance, refusing prompts to generate extremist propaganda or private personal data) and for policy makers to consider regulations.

  Another ethical aspect is intellectual property. GenAI models are trained on vast amounts of potentially copyrighted text and images scraped without explicit permission from authors and artists. This training data results in models which can produce content in the style of certain authors or artists, raising questions about infringement and ownership.

  Best practice is to treat AI outputs with the same IP diligence as human outputs: run plagiarism checks on AI-written text, and avoid using generated images for commercial purposes without understanding the training source. Some GenAI providers are developing features to cite sources for generated text (via retrieval augmentation) or to watermark AI-generated content to distinguish it from human-made.

  Privacy is another ethical dimension: GenAI can potentially output personal information seen in training data (though rare). And if users input sensitive data into an AI service, that data might be seen by the service provider or used to further train models, posing confidentiality risks. Responsible AI isn't just a buzzword – it's an active practice of anticipating how things can go wrong and putting measures in place to prevent harm.

- **Performance constraints:** Despite their power, GenAI models have practical limits. They can struggle with very long inputs or outputs – each model has a context length limit. If you feed a document longer than the model's limit, it might lose the beginning context, affecting coherence.

They can also get things wrong when prompted in certain ways (e.g.: Complex multi-step math problems or logic puzzles) which the model wasn't explicitly designed for. Models also lack common sense knowledge in some cases, or rather, they lack the lived experience and dynamic reasoning that humans develop.

Speed and cost are factors too: Running large models, especially locally, requires significant memory and compute. Using cloud APIs costs money per request, which can add up for heavy use. There's often a latency of a few seconds for a response, which might not be ideal for real-time needs. And while newer models are being optimised, deploying a GenAI system at scale (say, a chatbot for millions of users) incurs major infrastructure costs.

Finally, model updates present a challenge: If you rely on a specific AI model, and it gets updated, its behaviour might change slightly, which could affect consistency in outputs. While GenAI is powerful, it has quirks and limits that professionals must work within.

- **Carbon cost:** GenAI models have a significant environmental impact, from the carbon footprint of training to emissions during their use. Training state-of-the-art LLMs is extremely energy-intensive, often consuming massive electricity and producing substantial CO_2 emissions – for example, training GPT-4 Turbo Turbo (175 billion parameters) used about 1,287 MWh of power and was estimated to emit roughly 500 metric tons of CO_2. Once deployed, LLMs continue to need power to run the models.

  There is a direct trade-off between computational cost and carbon impact, meaning that reducing the compute (and energy) per token of output both lowers expenses and curbs emissions. In practice, more efficient models and optimisations can reduce both costs and environmental impact in tandem.

  To mitigate LLMs' environmental impact, strategies are being pursued such as improving energy efficiency (in both hardware and data centre operations), optimising model architectures and usage (e.g. using smaller or specialised models, when possible, to avoid unnecessary computation), and transitioning to renewable energy sources for powering AI infrastructure. The carbon accounting for LLM based operations will increase in importance as the overall financial cost of running these models reduce.

In facing these challenges, the role of human oversight cannot be overstated. Users of GenAI should be trained to critically evaluate AI outputs, cross-check facts, and understand basics like prompt engineering to reduce errors. Organisations deploying GenAI should implement an AI governance process – guidelines and guardrails ensuring the AI is used in line with ethical and legal standards.

It's also wise to start with low-risk applications and gradually expand as confidence grows. By acknowledging limitations and embedding safeguards, we can reap GenAI's benefits while managing its downsides.

Common GenAI services and techniques

As GenAI usage grows, new techniques and frameworks have been developed to extend its capabilities and improve its reliability. Here we introduce a few key approaches – Retrieval-augmented generation (rag), prompt chaining, agentic AI frameworks, and other emerging methods:

- **Retrieval-augmented generation (RAG):** The most popular way to currently tackle the hallucination and knowledge cut-off issues of language models is RAG. In RAG, the AI model is augmented with an external knowledge source (like a database or document repository) that it can query to fetch relevant information, which is then provided as context for generation. Essentially, RAG marries traditional search with text generation. For example, instead of asking the model "What are the latest tax law changes in 2025?" and hoping it knows, a RAG system would first search a curated tax law database for 2025 updates, then feed those excerpts into the prompt, asking the model to base its answer only on that information.

  By grounding the model in up-to-date, specific data, RAG significantly improves factual accuracy and reduces hallucinations. It also allows models to cite sources – since the retrieved documents can be referenced, the AI's answer can include footnotes or URLs, building trust with users.

  RAG often uses a vector database to store embeddings of documents, enabling semantic search to retrieve passages related to the query. What this means practically is the knowledge base is converted firstly into "chunks" which are smaller sections of the document and then converted into number representations which hold semantic weighting. As an example, if you are trying to find the answer to the question "how much holiday can I take" in an HR policy document, using traditional word searches, you will only find an answer if the word "holiday" appears. With a vector based semantic search all semantically similar words or "chunks" will also come up as responses. In this example it would include "vacation", "time off" and "annual leave".

  Although this technique improves the responses and accuracy there are important consequences to be aware of. RAG systems are only as good as the underlying information the retrieve (if a policy is poorly written or vague, the responses will reflect this). RAG is also better suited to searching parts of knowledge stores and as a result may answer on out of context chunks based on the retrieval process. Filtering techniques can be used to reduce this, but as stated previously human review and validation is still required.

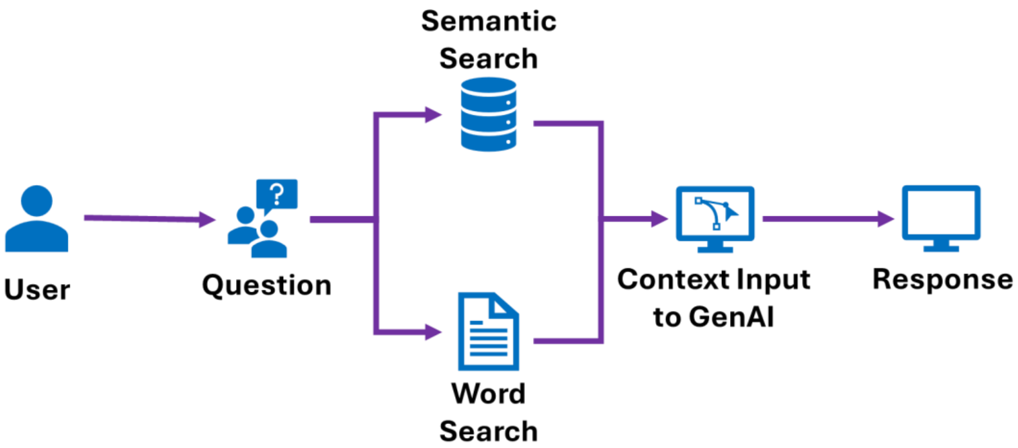

- **GraphRAG** uses a graph-based knowledge structure in addition to standard text search, allowing the AI to see not just isolated snippets but also how those snippets connect, which can greatly improve the accuracy and depth of answers. Unlike typical RAG, which retrieves paragraphs or chunks from a database, GraphRAG organizes information into a network of entities and relationships. By following these connections, the system can piece together multiple facts or concepts that might otherwise be overlooked.

  This approach is particularly useful for complex questions that rely on combining disparate pieces of knowledge—for instance, linking items from different sections of a policy. The result is a more comprehensive retrieval, with fewer gaps or hallucinations. GraphRAG still incorporates semantic search techniques and benefits from the same up-to-date grounding as regular RAG, but it adds an extra layer of structure that yields more reliable, “big-picture” insights. However, just like other RAG methods, GraphRAG depends on the quality of the underlying data and requires human oversight to verify context and correctness.

- **Prompt chaining:** Prompt chaining is a simple yet powerful concept: instead of asking a model to perform a complex task in one go, you break the task into a sequence of smaller prompts, using the output of one step as the input to the next. This chain-of-thought approach can guide the model through multi-step reasoning or formatting tasks that it might otherwise get wrong if asked directly. For example, imagine you need an AI to analyse a lengthy article and produce a concise report with key points and recommendations.

  Rather than one giant prompt (“Read this 10-page article and give me a report…” which risks the model losing focus or context), you could do something like: First prompt the model to summarise each section of the article, then feed those summaries into another prompt asking for overall key insights, then a final prompt to draft recommendations based on those insights. Each step simplifies the model’s job. Prompt chaining is useful when an LLM struggles with depth of reasoning or multi-part instructions – by structuring the interaction, you essentially hand-hold the model through the task.

- **Agentic AI frameworks:** While prompt chaining involves the user orchestrating multiple steps, what if the AI could orchestrate its own steps to achieve a goal? That's the idea behind agentic AI frameworks – giving AI a level of autonomy to decide which actions to take (which could include calling tools or issuing new prompts to itself) in order to complete a task. An "AI agent" is a system that can interact with its environment (through tools like search engines, calculators, databases) and perform multi-step reasoning without constant human intervention.

  For instance, instead of just answering a question, an agent could decide: "I should do a web search, then use a calculator, then compose an answer using those results." Such an agent might handle a query like "What was the highest temperature in London this week and convert it to Fahrenheit?" by actually searching for London weather data, finding the max temperature, then converting it. Conventional GenAI interactions would not reliably get that right without the search. With more autonomy comes more unpredictability. One must carefully constrain agents to not go off-track or misuse tools.

Optimising GenAI for work

To get the most out of generative AI in your daily work, it's important to learn how to effectively interact with these models and integrate them thoughtfully into your workflow. This involves mastering prompt techniques, setting realistic expectations, and avoiding pitfalls like over-reliance or misuse. Here are some best practices and tips for professionals:

1. Craft clear and specific prompts: The output of a GenAI is only as good as the prompt you give it. A well-crafted prompt can dramatically improve relevance and accuracy.

When asking the AI to do something, be explicit about what you want. Include necessary context, specify the format of the answer, and define the role or style if needed. For example, instead of asking "Tell me about climate change", you might prompt: "You are an environmental policy analyst. Summarise the three biggest impacts of climate change on coastal cities, in 2-3 paragraphs, in an academic tone, citing factual data." This gives the model context (analyst perspective), scope (coastal cities, three impacts), length, tone and an instruction to focus on data. Such specificity guides the AI and reduces ambiguity.

2. Use examples and iteration: One powerful technique in prompting is to show the AI examples of the output you want. This is called few-shot prompting. For instance, if you want it to generate responses in a certain format, you can provide a demonstration: "Q: [example question]\nA: [well-structured answer]\nQ: [your question]\nA:". The model will infer that it should follow the shown format for the new answer. You can also give a sample of the output.

Iterate with the model. You might start with an initial query, get an answer, then refine your prompt or ask follow-up requests to improve the output. Treat it as a collaboration: Get something, give feedback (via another prompt), and so on. Rather than one huge prompt for a complex output, chain a couple of prompts: First brainstorm ideas, then refine a chosen idea, then polish wording. Each step can be a new prompt where you instruct the model based on the last output. This interactive, stepwise approach tends to yield better final results and gives you more control.

3. Set realistic expectations and maintain oversight: As emphasised earlier, AI is a powerful assistant, not an infallible oracle. Always approach outputs critically. Never use the raw output blindly. For professional use, you (or someone on your team) should review and edit AI-generated content. This is important not only to catch factual errors or typos but also to ensure the output meets your needs and tone. Use AI for the heavy lifting of a draft, then apply your expertise to refine it. Also, be realistic about what tasks you delegate to AI. It's great for first drafts, summaries, boilerplate generation, idea generation, and even making sense of data. But for final decisions, nuanced judgments, or highly sensitive communications, a human touch is essential.

By setting this expectation with yourself and your team, you avoid over-reliance. Viewing AI as a copilot – one that sometimes goes off course – will help maintain necessary oversight.

4. Protect confidentiality and use AI responsibly: When integrating AI into your workflow, especially public or third-party models, be mindful of what data you input. Assume that anything you paste into a cloud AI service could be seen by humans or used to train models (unless you have guarantees from the provider otherwise). So, don't input sensitive client information or company secrets into a public chatbot.

Many AI providers now offer enterprise plans where data isn't retained or used for training – if you plan heavy use with proprietary data, look into those options. It is recommended organisations to have AI usage policies, so all colleagues have a clear understanding of the expectations and permitted use.

5. Continuously learn and adapt: GenAI capabilities and best practices are evolving rapidly. Keep yourself updated on new features or techniques that can improve outcomes. For example, prompt engineering forums and guides share tips on how to elicit certain styles or how to circumvent common issues. Many AI providers publish best practice guides which are good resources to use.

If a model isn't giving good results, be open to exploring different techniques, styles or models. Approach it with an open mind willing to explore but accepting that the desired outcome may either only be partially possible or not at all

Ethical and regulatory considerations

The rapid adoption of AI, especially GenAI models, has spurred action in terms of governance, policy and regulation. The EU AI Act came into force in 2024, and similar frameworks are under discussion in the UK and globally. Professionals must be mindful of the ethical and legal frameworks surrounding AI use, which continue to evolve rapidly. This includes compliance with data privacy laws, understanding emerging AI-specific regulations, and adhering to best practices in AI governance within organisations. Below, we outline key considerations:

AI governance and policies: Many organisations are now developing internal AI governance programs to ensure responsible use of AI. This typically involves cross-functional committees or working groups that create guidelines on where and how AI can be used, review high-risk AI applications, and monitor outcomes for any ethical or compliance issues.

Part of governance is also training employees – making sure users of AI understand things like not to expose confidential data, to verify outputs, and to avoid bias. On a technical level, organisations might maintain a registry of AI models in use and perform periodic audits. These audits could check for things like disparate impact (to ensure an AI model isn't unintentionally discriminating against a group) or security of the AI systems.

Privacy and data: Privacy laws such as GDPR and the EU AI Act (enacted in 2024) apply fully to AI usage. If you are using personal data to feed an AI model or service, you must ensure you have the right legal basis to do so, and that individuals' rights are respected.

When using GenAI for client work, do not input personally identifiable information (PII) unless the tool is approved for that purpose and compliant. If you were to, say, summarise a client's HR records with an AI, you may be transmitting personal data to a third-party AI provider – which could violate privacy laws or contractual confidentiality.

Some GenAI models can "memorise" bits of training data. There's an interesting risk: If a model was trained on sensitive data, it might inadvertently regurgitate some of it when prompted a certain way. This is why many companies do not allow feeding proprietary data into public model training.

Using public AI models for client work: One of the most common questions professionals have is: "Is it safe and appropriate to use tools like ChatGPT or other public AI services for my client work?" The answer is nuanced. While these models can be incredibly helpful, using them in a client context introduces risks related to confidentiality, compliance, and quality control that you must manage.

Risks of using public AI models:

Confidentiality and data privacy: When you use a public AI and input client information or documents, that data is sent to a third-party server. While reputable providers have policies and security measures, you typically do not have full control over that data. It may be stored on the server, and even if not used for training, it could reside in logs or backups.

This raises serious issues: If you copy-paste a client's strategy memo into a model to summarise, you may have just exposed sensitive client info to an external system without authorisation. Many client contracts (and professional ethics rules) have strict confidentiality clauses. Violating those can lead to loss of trust, legal liability, or regulatory penalties.

Even if the AI provider is trustworthy, any transmission of non-public data could be considered a breach of confidentiality if not permitted by the client or law. In short, you must treat a public AI service as you would any external service: Don't share data that isn't approved for external sharing. If you wouldn't email that text to a random person, don't feed it to a random AI either.

Transparency and disclosures: An emerging best practice is being transparent when AI is involved in automated decisions or content creation. In some jurisdictions, this is becoming a legal requirement.

For professional services, consider disclosing AI assistance where relevant. For example, if a financial planning report was drafted with AI help, the firm might note internally (or even to the client in some cases) that "This document was prepared with the assistance of an AI tool and has been reviewed by [Person/Role]." Such transparency can build trust, as clients appreciate knowing that while you leveraged advanced tools, you also applied your expertise to ensure quality.

When GenAI contributes to a decision – say, an AI pre-screens job applicants or flags transactions for fraud review – informing the subjects (applicants or customers) might be required by law (as part of algorithmic transparency initiatives) or at least recommended to maintain fairness and allow recourse if there's an error.

In summary, the ethical and regulatory landscape for AI is adapting with the technology. It is important for professionals don't view AI as a wild west where "anything goes" – existing laws on privacy, discrimination, liability and others do apply to AI uses.