GENERATIVE AI: BEYOND THE THEORY

Webinar 3

Simplifying the Complex: Making Sense of GenAI

27TH FEBRUARY 2024

ADDLESHAW GODDARD

MORE IMAGINATION MORE IMPACT

CONTACTS

KERRY WESTLAND
Partner, Head of the Innovation Group

ELLIOT WHITE
Director of Innovation & Legal Technology

MIKE KENNEDY
Senior Manager, Research & Development

AGENDA

1. Key Concepts and Definitions
2. Prompting Strategies and Prompt Engineering
3. Hallucinations
4. LLM Deep Dive: Training, Models and Tokens
5. A Secure Approach

KEY CONCEPTS & DEFINITIONS

WEBINAR 1 RECAP

GENERATIVE AI is a type of machine learning that can be used to create new content. This can include audio, text, code, images, and videos.

This is made possible by **LARGE LANGUAGE MODELS** which are algorithms that have learnt how to predict the next word based on previous words in a sequence (the language pattern) and generate relevant and new conversational text in relation to a prompt or query.

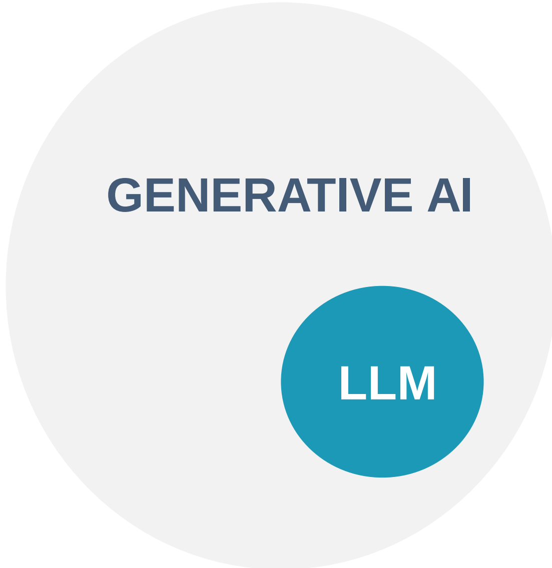

Examples

OpenAI | ChatGPT | Claude | Bing | Gemini | Lamda | Copilot

KEY CONCEPTS

TOKENISATION is a way of turning text into a sequence of tokens, that are then converted into numbers that can be understood by computers.

CONTEXT WINDOW is the maximum amount of text (in tokens) a model can take in at once.

RETRIEVAL AUGMENTED GENERATION (RAG) is a technique that combines retrieval tools with GenAI. This works by using the input to retrieve relevant information from a body of knowledge which is then used as context to inform a response back from the model.

CHUNKING refers to the process of dividing text into pieces that can be processed by a model. These chunks could be sentences, paragraphs, clauses, or any segment of text. The aim of chunking is to improve the model's ability to process text by either ensuring chunks fit in the context window, or only feeding relevant chunks that answer a specific prompt.

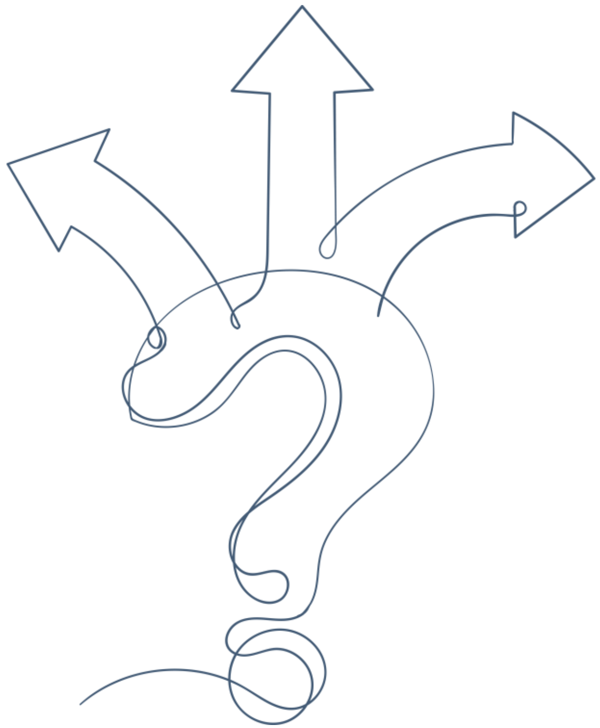

KEY CONCEPTS: PROMPTS

SYSTEM PROMPT is a set of instructions to the specific tool that will always prepend itself to any instructions of prompts entered into the tool. This is not editable to the end user and does not change.

You are a UK based lawyer who is an expert in M&A. Any responses that are given will be in UK English. You are responsible for identifying risks within Legal Agreements. You will be given some text that may or may not include a specific risk, and you will be asked a question about this text. Keep your answers concise. You must include a clause reference in your answer.

INSTRUCTION PROMPT is the specific request, query, or instruction that the user inputs into the tool. This is completely in the user's control.

Is there an exclusivity restriction within the agreement?

An exclusivity restriction is a clause or provision that limits one or more parties from engaging in similar agreements or transactions with other entities. The exclusivity clause limits licenses, distribution rights, and other rights to specific parties. It grants to a party only the rights outlined in the contract.

KEY CONCEPTS: VECTORS

VECTOR EMBEDDINGS are like unique, numeric fingerprints for data, capturing their most important traits

VECTOR DATABASES are specialised libraries that store these fingerprints and are incredibly good at quickly finding and comparing them, making them perfect for tasks that require understanding and comparing complex data at a large scale.

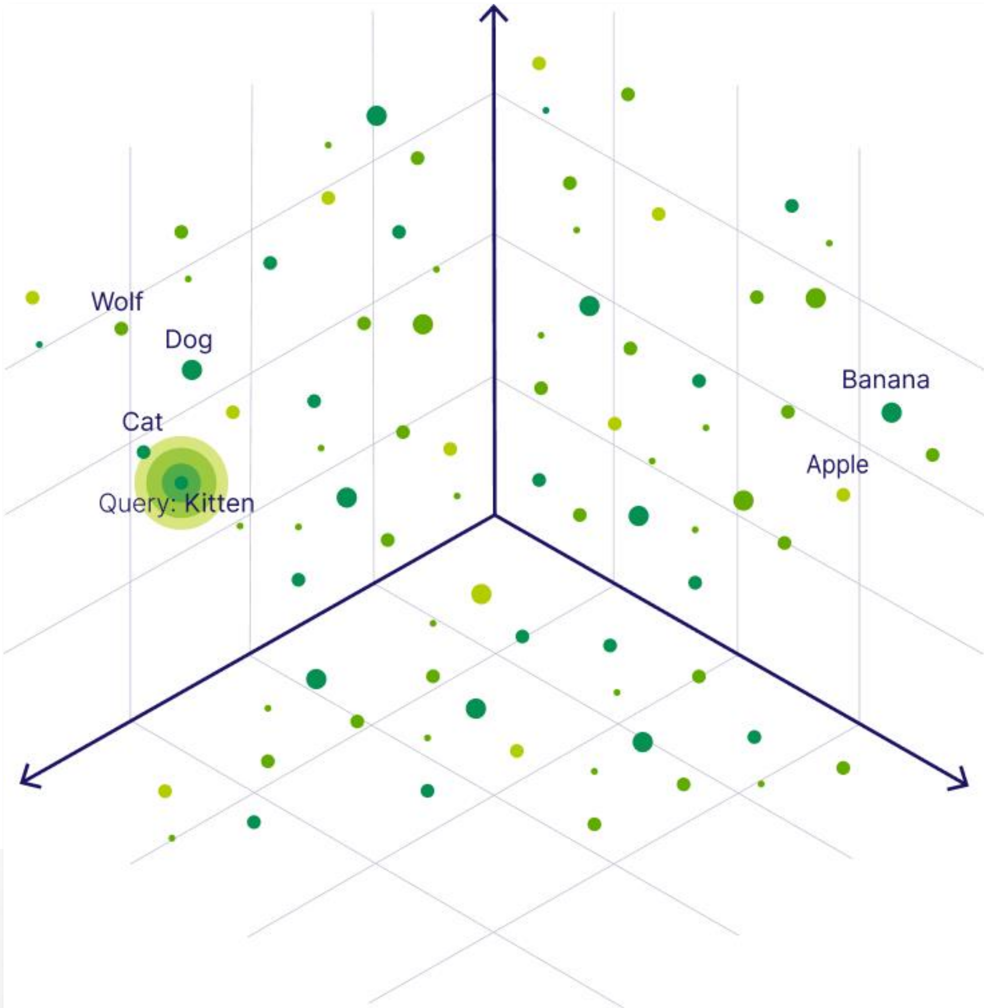

cat = [1.5, -0.4, 7.2, 19.6, 3.1, ..., 20.2]

kitten = [1.5, -0.4, 7.2, 19.5, 3.2, ..., 20.8]

PROMPTING

THE IMPORTANCE OF PROMPTING

A Prompt is the text used to communicate with an AI system. It is an instruction, usually in words, that allows us to tell the AI what we want and specifically how we want it.

There is a technique to communicating with an AI, learning how they understand language and how to get the best result – this has been coined "**Prompt Engineering**".

A well-constructed prompt is focused on the specificity of what a person is requesting of the Generative AI. The more detailed instructions provided to the AI the better the response you will get.

"CAVEAT PROMPTOR"

PROMPT EFFECTIVENESS

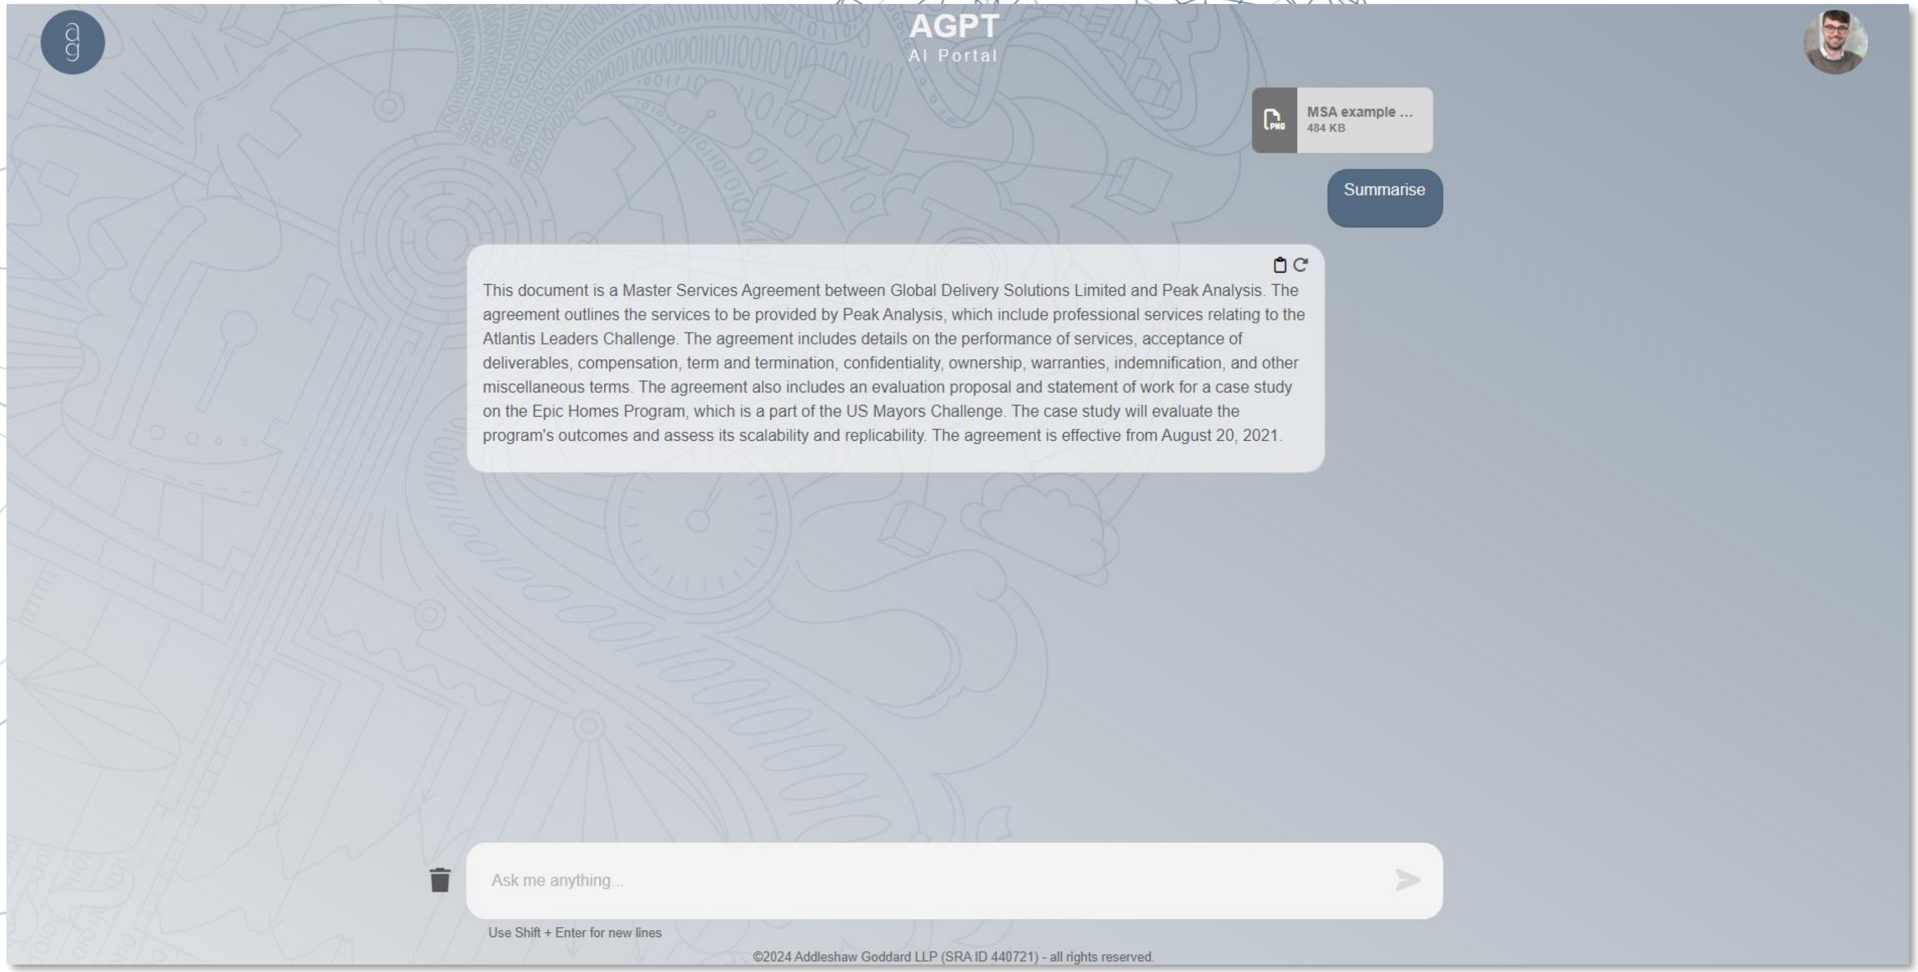

Broad, and directionless instruction to summarise a document will bring back a summary that allows the LLM to decide what it thinks it relevant, rather than what you most want to know

PROMPT EFFECTIVENESS

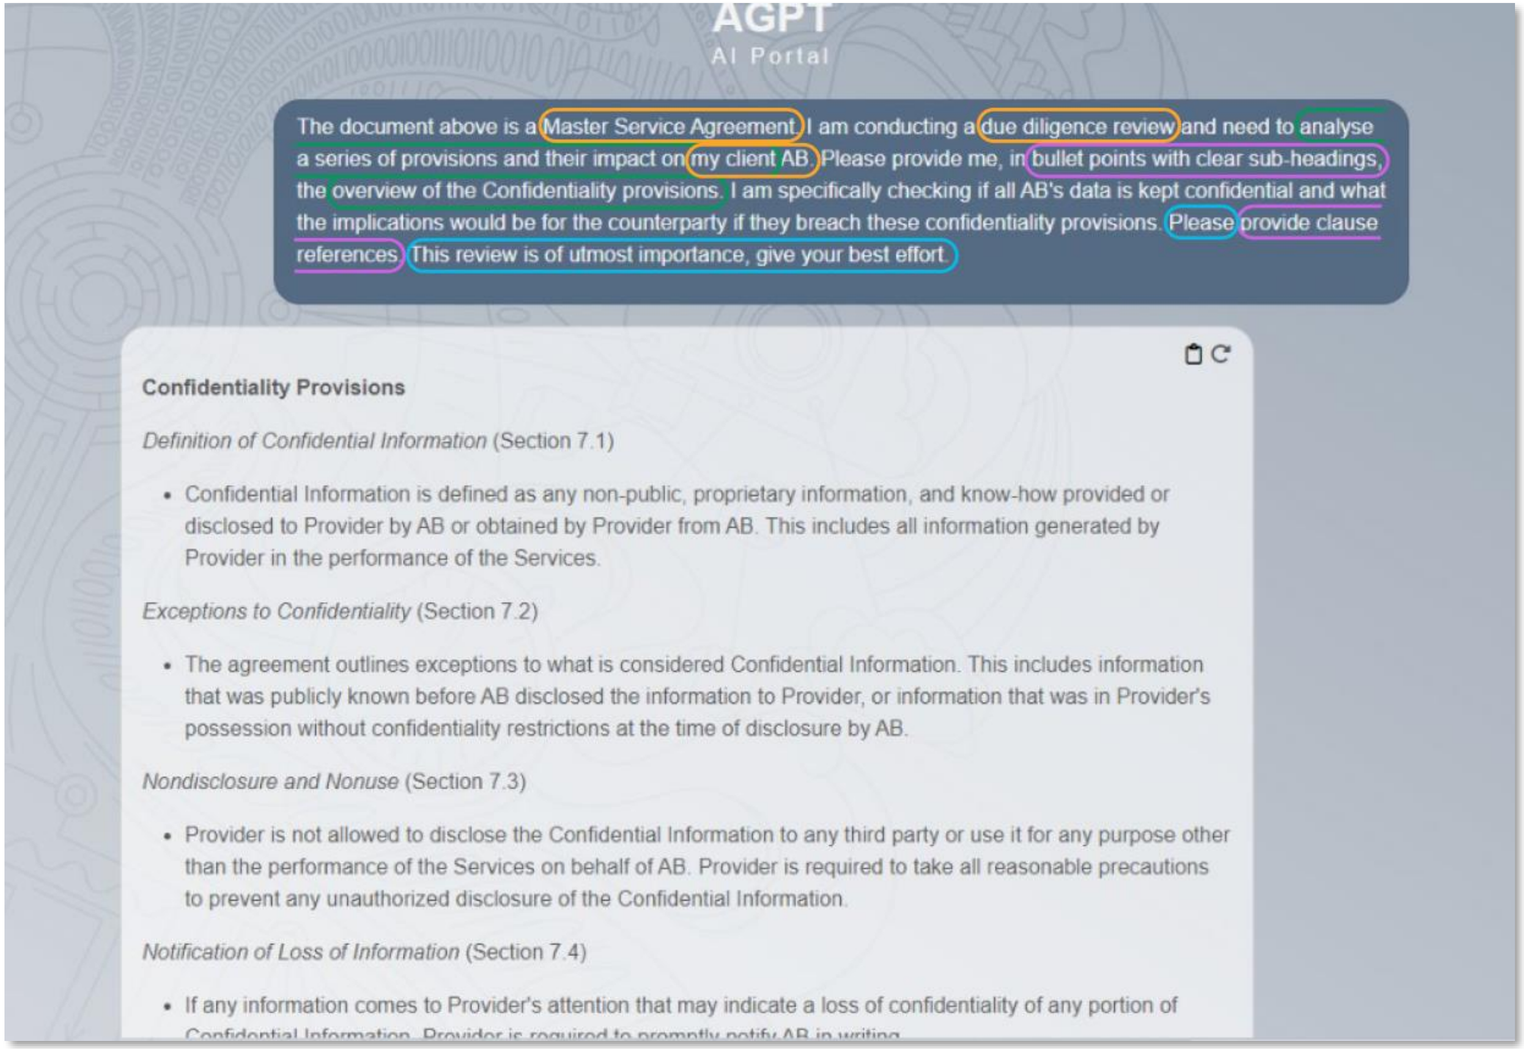

Provide Context

Outline Your Focus

Define Response Structure

Utilise Emotional Drivers

Clear and concise instruction to summarise a document focusing on key points that the user wants to know, telling the LLM what to focus on.

PROMPTING STRATEGIES: PROVIDING EXAMPLES

Provide a few examples of the task you want the model to perform. These examples serve as a guide for the model to understand and execute similar tasks. For instance, when asking for a Termination for Convenience provision, provide an example of what would be covered in such provision, including example rights given to parties.

What are the termination for convenience rights in this contract? As an example, In a contract, a Termination for Convenience provision may allow one party to terminate the contract without cause, meaning without any fault on either side. This provision is typically used in business agreements where circumstances may change over time, making the contract no longer beneficial or feasible for one party. The party terminating the contract is often required to provide notice to the other party and may also be required to compensate the other party for work performed up to the point of termination. This provision protects the terminating party from breach of contract claims, as it allows for termination outside of any specific breaches or failures by the other party. However, it's important to note that the use of a Termination for Convenience clause should be carried out in good faith and not as a means to avoid contractual obligations.

In this Master Services Agreement, the termination for convenience rights are outlined in Section 6.2. According to this clause, AB has the right to terminate the agreement, in whole or in part, for any reason or for no reason by giving a 30-day prior written notice to the Provider. This allows AB to end the contract without having to prove that the Provider has breached any terms of the agreement.

Since this is a Firm Fixed Price Contract, if AB terminates the contract without cause, the Provider will invoice for the remaining Budget and Billing Schedule as outlined in the Statement of Work (SOW). This ensures that the Provider is compensated for the work performed up to the point of termination.

This clause provides AB with flexibility in case circumstances change over time and the contract is no longer beneficial or feasible. However, it also implies a responsibility to use this provision in good faith and not as a means to avoid contractual obligations.

PROMPTING STRATEGIES: BREAKING QUESTIONS UP

Asking one question at a time ensures a more focused and detailed response from the model. Avoid asking a list of questions.

This technique is crucial for long documents. Current technology requires chunking the documents and using a retrieval method to find correct chunks. This means the LLM will not be fed all of the correct information for all questions.

Prompt: What are the specific payment terms outlined in the MSA?

1. Payments are to be made based on the Budget and Billing Schedule set forth in the Statement of Work (SOW) upon the satisfactory completion of work.
2. There is no specific frequency of payments mentioned (like monthly, quarterly, etc.), it's based on the completion of work and submission of invoices by the Provider.
3. Payments for any undisputed amount in the invoices are to be made no later than 60 days after receipt of Provider's Monthly Invoice or 30 days after AB's receipt of the respective fees paid by BFF under the BFF Agreement, whichever is later.
4. The MSA does not mention any late payment penalties or interest charges.
5. There are no provisions mentioned for early payment discounts.
6. There are no provisions mentioned for payment in instalments or milestones.
7. The MSA does not specify a particular payment method or currency.
8. There is no provision for retaining a certain percentage of payment as a form of security or guarantee.
9. Any changes to the SOW and/or subsequent price changes must be mutually agreed upon by the Provider and AB in writing in the form of a formal Contract Amendment (CA) (Section 1.2).
10. The MSA does not specify provisions for dispute resolution related to payment issues.
11. AB may terminate this Agreement, in whole or in part, for any reason or for no reason by giving 30-day prior written notice to Provider (Section 6.2).
12. Provider will submit invoices to AB at the designated time of payment based on the Budget and Billing Schedule (Section 5.2).
13. Provider will provide AB with reasonable access to Provider's premises, records, and personnel so that AB may confirm that Provider has properly invoiced AB under this Agreement (Section 5.5).
14. The MSA does not mention any provisions for indemnification or liability related to payment.
15. The MSA does not mention any provisions for confidentiality or non-disclosure related to payment information.

The specific payment terms outlined in the Master Services Agreement (MSA) contract are as follows:

1. Services Fees: AB will compensate the Provider for the Services according to the Budget and Billing Schedule set forth in the Statement of Work (SOW). This is a Firm Fixed Price Contract, and AB will not reimburse the Provider for any costs or expenses unless specified in the SOW and approved by AB before incurring a specific cost or expense (Section 5.1).
2. Payment Schedule: The Provider will submit invoices to AB at the designated time of payment based on the Budget and Billing Schedule upon the satisfactory completion of work as set forth in the SOW and Section 3. AB will pay any undisputed amount set forth in such invoices no later than 60 days after receipt of Provider's Monthly Invoice or 30 days after AB's receipt of the respective fees paid by BFF under the BFF Agreement, whichever is later (Section 5.2).
3. Reimbursements for Expenses: AB will not reimburse the Provider for any costs or expenses unless specified in the SOW and with AB's approval before incurring a specific cost or expense (Section 5.3).
4. Travel Expenses: AB will not reimburse the Provider for any costs or expenses related to travel unless specified in the SOW and with AB's approval before incurring a specific cost or expense (Section 5.4).
5. Inspection by AB: The Provider will provide AB with reasonable access to Provider's premises, records, and personnel so that AB may confirm that the Provider has properly invoiced AB under this Agreement (Section 5.5).

PROMPTING STRATEGIES: FOLLOW UP

Drafting a prompt that can give you a high-level initial overview and then following up for more detail can deliver accurate answers.

Follow up allows for a more conversational review and speeds up human validation. It can be used to highlight sources, clause references or provide a detailed reasoning.

Initial 8 question review response

⚠ These suggestions are provided for assistance purposes only, you retain ultimate responsibility for your review and ensuring it is accurate.

1. Inventive Names Inc - Explanation: *The counterparty to the agreement is Inventive Names Inc as referred to in the Parties section (page 1).*

2. Yes - Explanation: *There is a cap on liability. The Supplier will not be liable for any direct or indirect damage arising out of or in connection with this agreement as referred to in clause 5 (page 3).*

3. No - Explanation: *There is no exclusivity provision mentioned in this agreement.*

**4. Either party may terminate this agreement by giving one month written notice to the other party. Additionally, either party may terminate this agreement with immediate effect by giving written notice to the other party if the other party commits a material breach of any term of this agreement. The Supplier may also terminate this agreement with immediate effect by giving written notice to the Customer if the Customer fails to pay any amount due under this agreement on the due date for payment as referred to in clause 2.2, 2.3, 2.4 (page 2)*.

5. No - Explanation: *There is no Change of Control provision or restriction mentioned in this agreement.*

6. No - Explanation: *There is no restriction on transfer of intellectual property. The Customer will retain ownership in all Intellectual Property Rights in the Customer Materials as referred to in clause 6.1 (page 3).*

7. Yes - Explanation: *The Customer shall indemnify the Supplier in full against all liabilities, costs, expenses, damages and losses suffered or incurred by the Supplier arising out of or in connection with any claim brought against the Supplier for actual or alleged infringement of a third party's Intellectual Property Rights arising out of the receipt or use of the Customer Materials in the performance of this agreement as referred to in clause 6.3 (page 3).*

8. Yes - Explanation: *There are several onerous terms including the limitation of liability, the indemnification requirement, and the strict payment terms as referred to in clauses 5, 6.3, and 3.6 (pages 2-3).*

💡 If I have made a vague or inaccurate attempt at an answer, put in the question reference (e.g. **1.1**) and I will have a more thorough attempt.

Drilling into Q8 for more detail

Under the terms of this agreement, there are several provisions that could be considered onerous or burdensome for the parties involved.

Firstly, the limitation of liability clause in clause 5 (page 3) could be considered onerous as it absolves the Supplier from any liability for direct or indirect damages arising from this agreement. This could potentially leave the Customer without recourse in the event of a significant loss or damage.

Secondly, the indemnification requirement in clause 6.3 (page 3) could also be seen as onerous. This clause requires the Customer to indemnify the Supplier against all liabilities, costs, expenses, damages, and losses arising from any claim of infringement of a third party's Intellectual Property Rights due to the use of the Customer Materials. This could potentially expose the Customer to significant financial risk.

Lastly, the payment terms in clause 3.6 (page 2) could be seen as onerous. This clause stipulates that if the Customer fails to pay any sum due under this agreement on the due date, the Customer shall pay interest on the overdue sum at a rate of 4% per year above the Bank of England's base rate, and a £40 administration fee will be added to any overdue sum. This could potentially result in significant additional costs for the Customer in the event of a late payment.

These clauses could potentially place a significant burden on the parties, particularly the Customer, and should be carefully considered and negotiated where possible.

HALLUCINATIONS

HALLUCINATIONS

A ‘Hallucination’ is content generated by an LLM that is presented as accurate but is made-up, or inconsistent with the input data

The large language model creates text based on its understanding of the prompt and the broader context within it

Not a search tool, it relies on the content provided or the source it has been specifically connected to

The creativity is a consequence of the model and how it creates sentences

Creativity vs **Accuracy**

OR

Hallucinations vs **Copyright Infringement**

HALLUCINATIONS: COMMON EXAMPLES

Using a tool such as ChatGPT to carry out Legal Research (such as finding case law) will have a high likelihood of resulting in hallucinations. Without proper connections and access to legal databases, the generated content will not be reliable.

Querying points of law also requires direct source material. Asking whether exempt credit agreements are covered by a specific section of the CCA can be reasoned to be false (due to the use of ‘exempt’), but is not so simple in reality.

Even with source material, occasionally these tools can hallucinate or bring back inaccuracies in document review. Either by omission or by including additional restrictions or rights that may be commonplace but are not actually in the source document. Often caused by search, chunking or RAG errors.

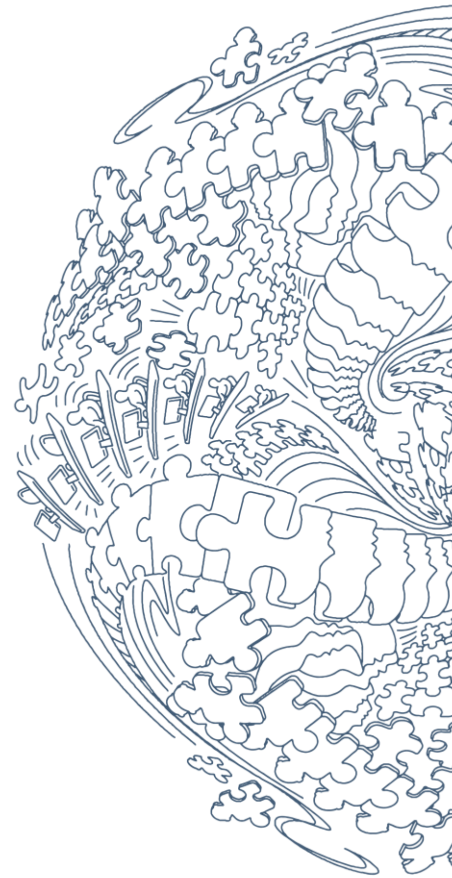

HALLUCINATIONS: MITIGATION STRATEGIES

PROMPT ENGINEERING. Following good practice with prompts and being clear in what is expected back.

CONNECTION TO SOURCE DATA. Utilising legal databases before using GenAI tools, and/or using the GenAI tools that have direct integrations with sources of data such as case law or regulation.

USING FOR THE RIGHT TASKS. GenAI is not a silver bullet for every task. There are specific things it is good at and many things it is not suited for.

HUMAN VALIDATION. Check outputs and spend time verifying information, saving 75% of time and being right is better than saving 100% of time and being wrong.

USING THE LATEST MODELS. Whilst hallucinations will never fully disappear, newer models are reducing the prevalence, staying up to date with your tech stack will help take advantage of progress.

HALLUCINATIONS: RECAP

Currently not possible to guarantee zero replication of wording from sources.

Also not possible to guarantee 100% correct retrieval of wording from sources.

There are ways to mitigate hallucinations, but taking responsibility for your use of GenAI is the most important thing.

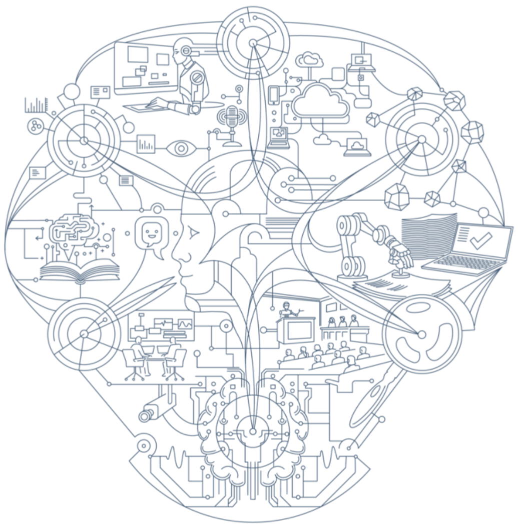

"…your affiant greatly regrets having utilized generative artificial intelligence to supplement the legal research performed herein and will never do so in the future without absolute verification of its authenticity…"

Steven Schwartz, New York Attorney

DEEP DIVE INTO LLMS

HOW DOES AN LLM WORK: TRAINING

Large Language Models (LLMs) have been trained on extremely large amounts of content to give a foundational understanding of language.

These models are trained at a point in time and are focused on understanding language and relationships between words. **This understanding does not have an expiry date.**

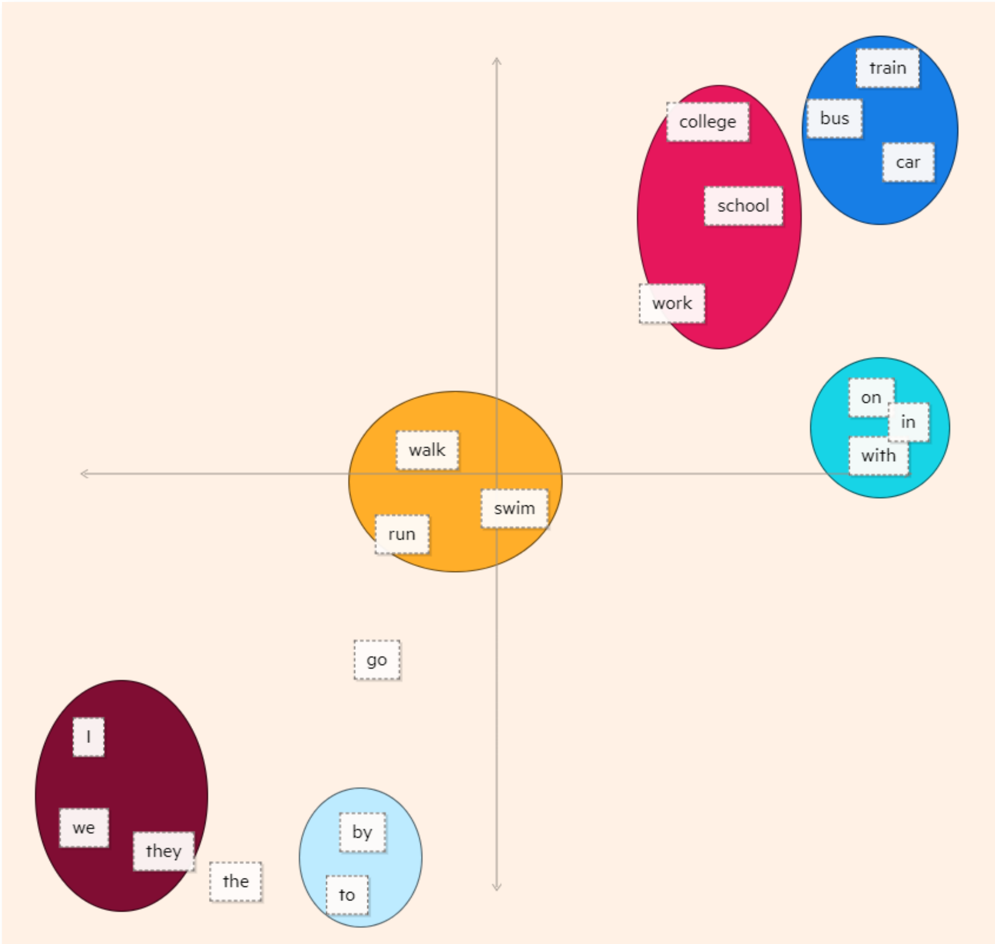

Data entered into these models for training purposes is a drop in the ocean, and you are very unlikely to be able to reproduce data entered into something like ChatGPT. Individual data points get absorbed into the 1.76 trillion parameters that are used to train a model.

Source: https://ig.ft.com/generative-ai/

HOW DOES AN LLM WORK: TRANSFORMER

Initially developed by Google in 2017. The Transformer has been a real driving force behind Generative AI.

Key feature of LLMs is **Self-Attention**. This means that a model can look at all words in a sentence at the same time and apply the context of the entire sentence when understanding or generating language.

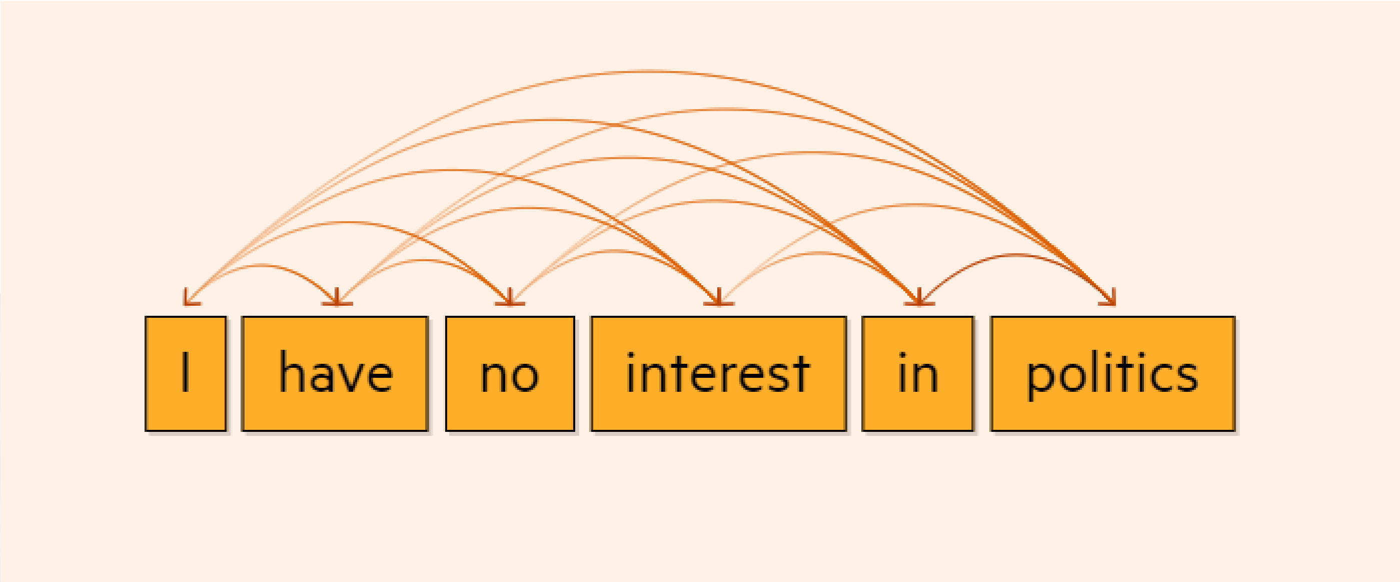

Source: https://ig.ft.com/generative-ai/

HOW DOES AN LLM WORK: UNDERSTANDING TEXT

This means that the **intended meaning** of a word or a sentence can be understood much better than previous Natural Language Processing tools.

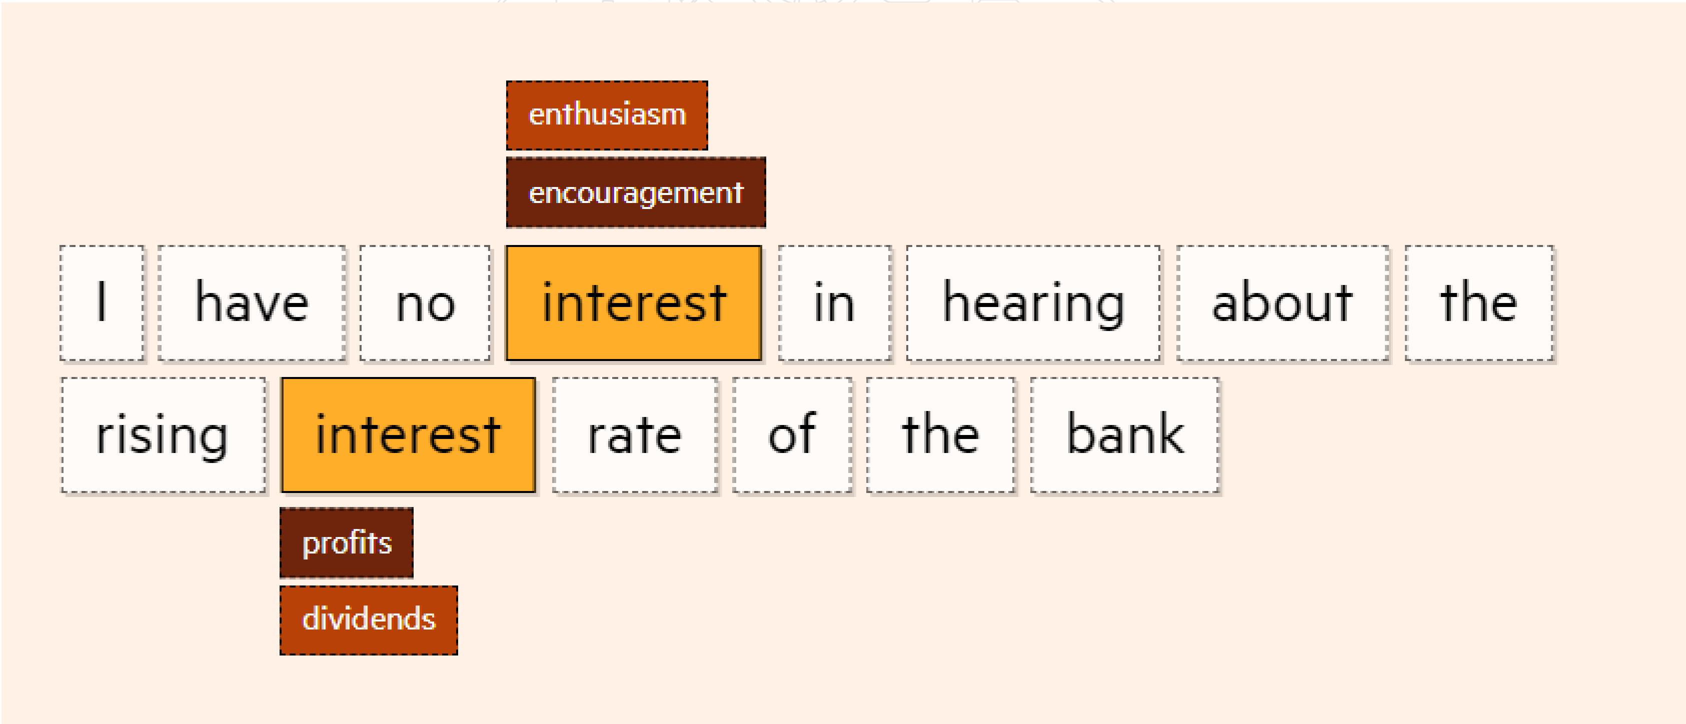

Trained on enough data, this can **mimic language understanding and reasoning capability.**

Source: https://ig.ft.com/generative-ai/

LLMS: DOMAIN SPECIFIC OR NOT?

FINE TUNING is a process where a pre-trained model, which has already learned some patterns and features on a large dataset, is further trained on a smaller, domain-specific dataset.

DOMAIN SPECIFIC LLMS are models that have been trained on a domain specific dataset, such as solely financial data.

BloombergGPT: A Large Language Model for Finance

Shijie Wu[1,*], Ozan İrsoy[1,*], Steven Lu[1,*], Vadim Dabravolski[1], Mark Dredze[1,3], Sebastian Gehrmann[1], Prabhanjan Kambadur[1], David Rosenberg[2], Gideon Mann[1]
[1] Bloomberg, New York, NY USA
[2] Bloomberg, Toronto, ON Canada
[3] Computer Science, Johns Hopkins University, Baltimore, MD USA

Abstract

The use of NLP in the realm of financial technology is broad and complex, with applications ranging from sentiment analysis and named entity recognition to question answering. Large Language Models (LLMs) have been shown to be effective on a variety of tasks; however, no LLM specialized for the financial domain has been reported in literature. In this work, we present BloombergGPT, a 50 billion parameter language model that is trained on a wide range of financial data. We construct a 363 billion token dataset based on Bloomberg's extensive data sources, perhaps the largest domain-specific dataset yet, augmented with 345 billion tokens from general purpose datasets. We validate BloombergGPT on standard LLM benchmarks, open financial benchmarks, and a suite of internal benchmarks that most accurately reflect our intended usage. Our mixed dataset training leads to a model that outperforms existing models on financial tasks by significant margins without sacrificing performance on general LLM benchmarks. Additionally, we explain our modeling choices, training process, and evaluation methodology. We release Training Chronicles (Appendix C) detailing our experience in training BloombergGPT.

February 20, 2024 · 9 min read

Meet KL3M: the first Legal Large Language Model.

KL3M is the first model family trained from scratch on clean, legally-permissible data for enterprise use.

LLMS: TOKENS

You are a UK based lawyer who is an expert in M&A. Any responses that are given will be in UK English. You are responsible for identifying risks within Legal Agreements. You will be given some text that may or may not include a specific risk, and you will be asked a question about this text. Keep your answers concise. You must include a clause reference in your answer.

System Prompt
92 tokens

“The Purchaser shall have exclusive rights to supply each Product to the distributor channel in the UK & Ireland, and non-exclusive rights to supply to the distributor channel in other European States, at the price of the Purchaser's choosing; provided that (without prejudice to the Supplier's rights under paragraph 3.2) if the Purchaser appears, in the Supplier's reasonable discretion, likely to fail to meet the MPO for that Product in the period set out in the relevant Annex, its rights to supply that Product to the distributor channel in the agreed Territory shall become non-exclusive. The Purchaser may not market the Product to any prospective customers other than customers within this Territory, or to any prospective customer outside this Territory, without the prior agreement of the Supplier.”

Input Text
202 tokens

Is there an exclusivity restriction within the agreement?

An exclusivity restriction is a clause or provision that limits one or more parties from engaging in similar agreements or transactions with other entities. The exclusivity clause limits licenses, distribution rights, and other rights to specific parties. It grants to a party only the rights outlined in the contract.

Instruction Prompt
94 tokens

Yes, there is an exclusivity restriction within the agreement. The Purchaser has exclusive rights to supply each Product to the distributor channel in the UK & Ireland, and non-exclusive rights to supply to the distributor channel in other European States, subject to meeting the MPO. This is outlined in the clause relating to "Exclusive Rights" and "Non-Exclusive Rights" in the agreement. (Clause reference: Paragraph 3.1)

Output
106 tokens

LLMS: TOKENS - PRICING

You are a UK based lawyer who is an expert in M&A. Any responses that are given will be in UK English. You are responsible for identifying risks within Legal Agreements. You will be given some text that may or may not include a specific risk, and you will be asked a question about this text. Keep your answers concise. You must include a clause reference in your answer.

System Prompt
92 tokens
$0.006

Services Agreement – 4 pages – 1494 words

Input Document
1893 tokens
$0.05619

What are the Supplier's termination rights? Set each one out with a header in bold text and a specific clause reference.

Instruction Prompt
24 tokens
$0.003

Termination for Material Breach (Clause 2.3)

The Supplier may terminate this agreement with immediate effect by giving written notice to the Customer if the Customer commits a material breach of any term of this agreement.

Termination for Non-Payment (Clause 2.4)

The Supplier may terminate this agreement with immediate effect by giving written notice to the Customer if the Customer fails to pay any amount due under this agreement on the due date for payment.

Output
115 tokens
$0.014

LLMS: PRICING - WORKED EXAMPLE

Contract Review

- 1000 Documents
- 25 Legal Queries
- 20 Page Average Document Size

≈ $9,629
Token cost using GPT-4

Additional costs

+ *Tech cost - Azure infrastructure implementation or potential licence for tools*
+ *People cost for human validation of output and the actual advice (i.e. lawyers)*
+ *People cost for prompt engineering and time spent scoping*

LARGE LANGUAGE MODELS (LLMS)

LLM	Developer	Token Limit	~Words	Access
GPT-4-Turbo	OpenAI	128,000	100,000	Direct API Microsoft Azure
GPT-4	OpenAI	32,000	25,000	Direct API Microsoft Azure Microsoft Copilot ChatGPT App
GPT 3.5-Turbo	OpenAI	8,000	6,000	Direct API Microsoft Azure ChatGPT App
Claude 2	Anthropic	200,000	150,000	Direct API Amazon Bedrock Claude App
Gemini Pro	Google	32,000	25,000	Direct API Gemini App
Gemini 1.5 Pro	Google	1,000,000	700,000	Direct API
Llama 2	Meta	4,096	3,000	Direct API

A SECURE APPROACH

OUR APPROACH: AGPT

AGPT builds on some understandings on LLMs and how we apply this technology within a law firm.

Focus on security and data protection. Hosted in Microsoft Azure in the UK. AGPT has no data storage and is processed internally within AG.

Businesses can build their own tools with dedicated resource and commitment from the business. Conversations with clients are an important part of this, and driving understanding has been key.

OUR APPROACH: THIRD PARTIES

Robust information security policies for suppliers alongside an IT team that will assess against standards ensures safety when contracting with third parties.

Many of the suppliers offering GenAI tools are US based, or use US hosted models and data centres.

It is necessary to have a clear understanding of any risks, data processing procedures and data storage arrangements when contracting.

OUR APPROACH: PUBLIC TOOLS

Significant time has been spent on drafting policies and guidance for the firm to have clear rules when engaging with public tools.

There are use cases that involve public data or non-confidential information where tools like ChatGPT are very useful.

Public tools often have an opt-out option to ensure no data entered is being fed back to the model. Some tools will also offer self-hosting and enterprise licencing.

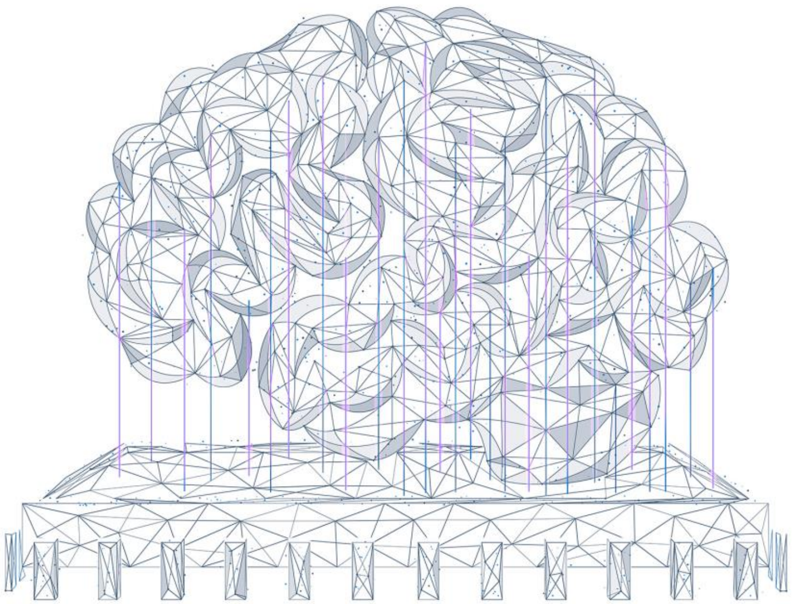

THIS SERIES

WEBINAR 1:

GENAI ESSENTIALS

WEBINAR 2:

EXPLORING GENAI IN LEGAL

WEBINAR 3:

SIMPLIFYING THE COMPLEX

CLIENT WORKSHOPS:

AI Essentials

Use Case Deep Dive

Build Your Business Case

Managing Risk / Maximising Opportunity

Find out more and register interest

QUESTIONS

?