Generative AI:
A (practical) legal guide

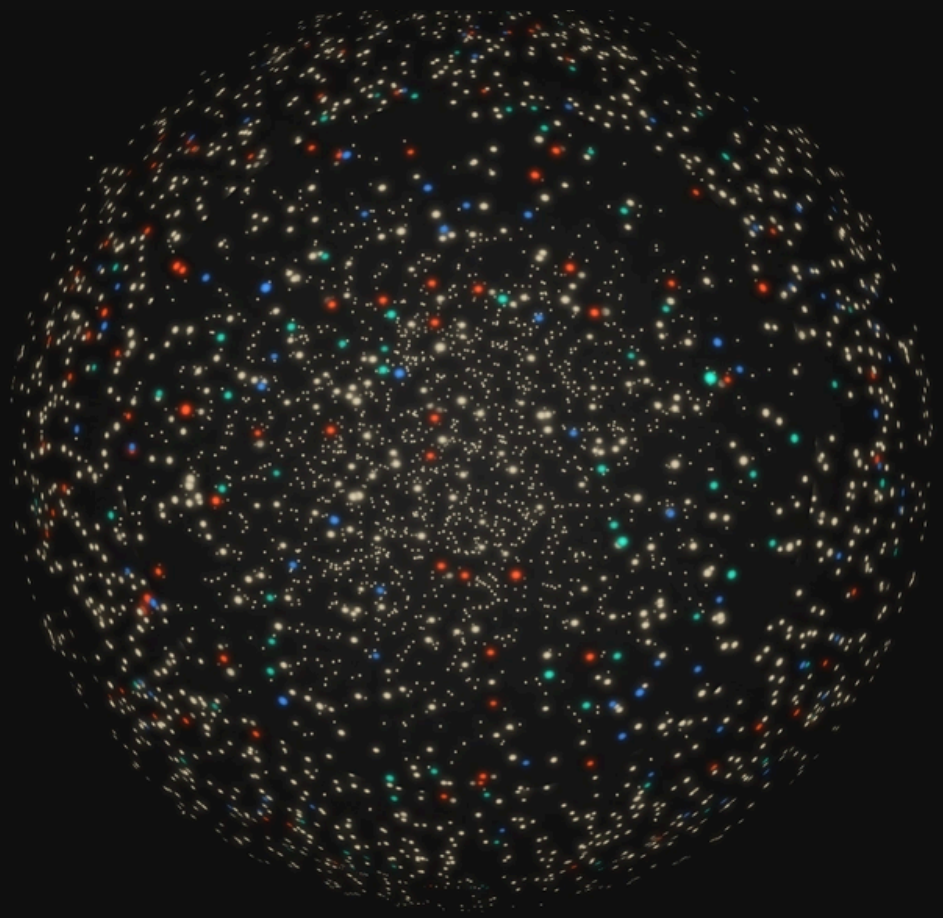

Hudson Gavin Martin
www.hgmlegal.com

Table of Contents

Introduction

The rise and proliferation of Generative AI in just the last two years has been one of the fastest and most significant technology developments of recent times. The technology itself isn't new, but the accessibility of easy to use tools like OpenAI's ChatGPT has driven a surge in popularity.

Gen AI tools are only going to become more popular as we create more data, in terms of volume and complexity, than we ever have before:

> **"The unmanageably huge volume and complexity of data (unmanageable by humans, anyway) that is now being generated has increased machine learning's potential, as well as the need for it."**
>
> **McKinsey & Company**

While there are enormous benefits that can be gained from using Generative AI in business, it does also present challenges and risks that need to be managed.

> **"AI knows no boundaries, respects no privacy, and fears no consequences"**
>
> **ChatGPT**
> **(In response to "write a pithy quote about the danger of AI re privacy")**

We know that businesses need practical advice on how to navigate the legal implications and considerations of using Generative AI in a commercial context.

How lawyers think about Generative AI

The potential benefits of using Gen AI tools are obvious. As lawyers, we need to assess the (sometimes less obvious) risks.

The nature and extent of these risks will depend on the context – the nature of the use case, the nature of the datasets and the nature of the tool.

From a legal risk assessment perspective, it's important as part of your Generative AI governance to ask the following questions at the outset of your project (and as part of your ongoing due diligence):

- What is the use case?
- What data is going in?
- What content is coming out?
- Who is involved?
- What are their terms?
- What governance is in place?

To help you understand this analysis, we have worked through these questions in the context of a fictional case study.

Case study: HealMe.AI

Imagine Hudson Gavin Martin has received instructions about the following new Gen AI project.

Overview

HealMe.AI is a telehealth business that has been operational for five years, primarily focused on delivering virtual clinical healthcare services. Initially established to address the unique needs of rural communities where access to healthcare providers can be challenging, the demand for HealMe.AI's services has surged post-COVID as more individuals seek convenient virtual healthcare options.

Objectives

The primary objective of this project is to optimise HealMe.AI's operations by integrating the OpenAI GPT model as a chatbot. The model will be trained on internal confidential data to provide initial medical diagnoses.

Implementation

- **Chatbot integration:** The OpenAI GPT model will serve as a chatbot for patients, allowing users to upload their personal health information and receive automated responses.
- **Model training:** The GPT model will be trained on HealMe.AI's internal, confidential information, including patient notes.
- **Code development:** Developers will leverage the GPT model to assist in the development of the underlying code base.
- **Media campaigns:** HealMe.AI's media team will utilise the tool to craft marketing campaigns. In certain cases, they will collaborate with third party creatives to generate content using Generative AI.

Understand the use case

In other words, what is the problem you are trying to solve by using an AI tool? As lawyers, this is the first question we'll ask. Like most applications of technology, it's not the tech itself that is the issue, it's how you use it.

The answer to this question drives:

- The AI tool you use (and whether it is public/private/enterprise).
- The way you deploy it.
- The legal risk profile (considering intellectual property, data protection, and third party risks).
- How much the lawyers care!

So, in our HealMe.AI case study, we would immediately note some key features about the commercial context we've been given:

There will be more than one use case at once.	This adds complexity and governance risk.
The business will be using Open AI / GPT.	This is a reasonably well understood tool, but there are potential risks around data and intellectual property.
At least one use case is using and ingesting very sensitive information (health information), and providing very sensitive content back (diagnoses).	So, we not only need to be concerned about general data privacy risks but also any specific obligations under health legislation. Correct outputs will also be critical.
The other marketing use case seems less risky from a data perspective but involves third party creatives and publication of AI-generated content.	So, there may potentially be some intellectual property and third party risk there. We also need to consider reputational risk from AI-generated content – how much human oversight is going on?
Developers will also be using the model to develop the code base.	This is often the sort of nugget that gets casually dropped into conversation! There is significant litigation about this issue currently, so our ears prick up.

Inputs and outputs

In terms of inputs, there are two types of input data we need to understand:

- What data has the tool been trained on?
- What data are your users inputting into the tool? And in this second bucket, we also want to know if the tool is training itself on your user data.

From a data perspective, as lawyers we're concerned about understanding what information is going into Gen AI tools because we need to know what obligations attach to that information.

Where you are inputting personal information you will need to consider the privacy obligations that attach to it – particularly in the case of sensitive personal information, like the health information in the HealMe.AI case.

The Privacy Commissioner in New Zealand recommends completing a privacy impact assessment before using personal information in any AI tool, including an understanding of what data sources the AI tool was trained on and how relevant and reliable those sources are for your purposes.

Be mindful that data inputted into an AI tool may constitute disclosure outside New Zealand for the purposes of privacy legislation.

In addition to this, any third party data you input may have obligations under contracts with those third parties.

So as part of using the AI tool it is important to understand what obligations you have in using the data and how you are going to comply with them.

The other consequence of inputting information is that it may be confidential information that is of value to you – that data being compromised could have significant financial implications.

Inputs and outputs

In relation to outputs, we want to understand what content will be produced so we can assess the risks. Issues with outputs will often stem from issues with input data.

We have all heard about the risk of hallucinations, where the Gen AI tool "makes things up" – often in a very convincing way. We would be especially concerned about hallucinations in medical diagnoses in the case of HealMe.AI.

The risk of bias in using Gen AI relates to several factors, including training data bias (the models might learn from data containing historical biases or stereotypes leading to outputs that reinforce those biases) and a lack of diversity in the training data (so the AI produces outputs that don't adequately represent or address a full range of perspectives and experiences). In relation to HealMe.AI, we would be asking whether appropriate training data is being used and whether the medical information has been screened for potential bias. For example, if most of our patient information comes from males, does this affect the quality of medical advice provided to females?

The greater the risks attaching to outputs, the more we would be looking for a "human in the loop" in the use case. This would be essential in the HealMe.AI use case of providing medical diagnoses to patients.

The IP analysis of inputs and outputs can be complex.

We would be thinking about all the following risks in relation to HealMe.AI and ensuring they were adequately covered off by contract (to the greatest extent possible):

- **Infringement of copyright** in creating the outputs.
- **Ownership of generated content** particularly if the output is being used commercially or integrated into other products.
- **Use of proprietary data** in training the tool.
- **Licensing issues** when the outputs will be used for commercial purposes (sometimes prohibited by the provider's terms of service).

Who's involved and what are their terms?

Like any technology, each solution has different advantages and disadvantages, with different vendors who prioritise different things.

In the context of our case study, we would be asking why HealMe.AI is using GPT to provide health information. Is the business comfortable it's trained on the right data? Is there a better tool / LLM for the use case?

For example, GPT is a 'generalist' model trained on a wide dataset. This makes it great for ChatGPT, where it will give you a semi-decent answer to any prompt in the world. But does it work for a telehealth solution that needs very accurate answers based on a much smaller, relevant dataset?

Also, generalist tools – which are fed on wide datasets and often publicly available – are likely to offer less contractual protection from a data security, IP infringement and output accuracy perspective than private/enterprise options. (It's worth noting that many public tools are evolving to provide quite strong protections, and standing behind their training data and outputs more because they have greater control over the training datasets.)

We would also be asking questions about, not just the provider of the AI tool, but other vendors in the supply chain using AI. For example, in the HealMe.AI media team use case, we know they are working with third party creatives who are using AI in the background. HealMe.AI may only have a contract with the creative agency – how do we know how they are using AI and what terms they have agreed to? And what protections can they give HealMe.AI as the end customer?

Governance:
The balance between risk and reward

In a commercial context, governance is the first thing to consider in relation to the use of AI in your business.

We are asked a lot about this. In practice, each organisation is different – the right question is how do you balance the risk and opportunity in the best way for **your** business?

We have seen a spectrum from outright bans of Gen AI tools, through to full-scale, commercial adoption as a business model. Fundamentally, there is no “right” and “wrong” approach. Something that looks like best practice for one business might be heavy handed for your use cases, and you risk losing the commercial benefits and opportunities that can flow from AI. There can also be significant risk in not embracing AI, or just moving too slowly.

Importantly, one of the first steps your business should take is a stocktake of your current governance framework to determine whether you need new policies or procedures to manage the responsible and effective use of AI in your workplace, or whether you can rely on what you already have.

Acting for HealMe.AI, we would first be reviewing their existing, relevant policies, which might include:

- Information security
- Technology procurement
- Data privacy
- IP and trade secrets
- Customer care
- Organisational values.

If need be, we would draft HealMe.AI a new framework for the responsible, compliant, and effective use of AI tech. The factors we would consider incorporating into policy and procedures would include ethical use, legal compliance, quality control, transparency and ‘explainability’, risk management (including reputational risk), consistency and accountability.

HGM Gen AI Policy

At Hudson Gavin Martin, we decided that AI has a big enough impact on how we work that a new policy was warranted, as we didn't have anything in our existing governance framework that covered it adequately.

Given the nature of our practice, we could keep that policy high-level, and principles-based.

Keeping your AI policy principles-based is useful in an area that is so fast moving. One of the tests we applied was "how regularly will we have to update this policy as the technology and regulations move?". Future proofing your policy and procedures as much as possible is key to ensure your AI governance is flexible and fit for purpose.

We have open sourced our AI policy and you can access it, in full, on our website. The key features of our policy are:

- We won't incorporate **confidential, sensitive or personal** information into a publicly available Generative AI tool.
- We will **check the outputs** of Generative AI before using them.
- We will be **transparent** about our use of Generative AI – with each other and with clients and stakeholders.
- We will communicate in our **own voice**.

Our policy won't be a template for everyone (or even every law firm), but we hope it's a useful reference. You can access it using the QR code below.

If you have any questions about your business's use of AI, please contact one of the specialists in our team.

A bit about us

Hudson Gavin Martin is a corporate, commercial and dispute resolution law firm boasting the largest and most experienced legal team in the country focused on technology, media and IP. Ranked Tier 1 in both TMT and IP by global legal directories Chambers & Partners and Legal 500, we support a diverse clientele including startups, scale-ups, SMEs, large national corporations, and global enterprises.

As New Zealand's largest legal team with expertise in innovative and disruptive technologies, we are industry-agnostic and committed to assisting clients across all sectors in taking advantage of the opportunities that these technologies provide.

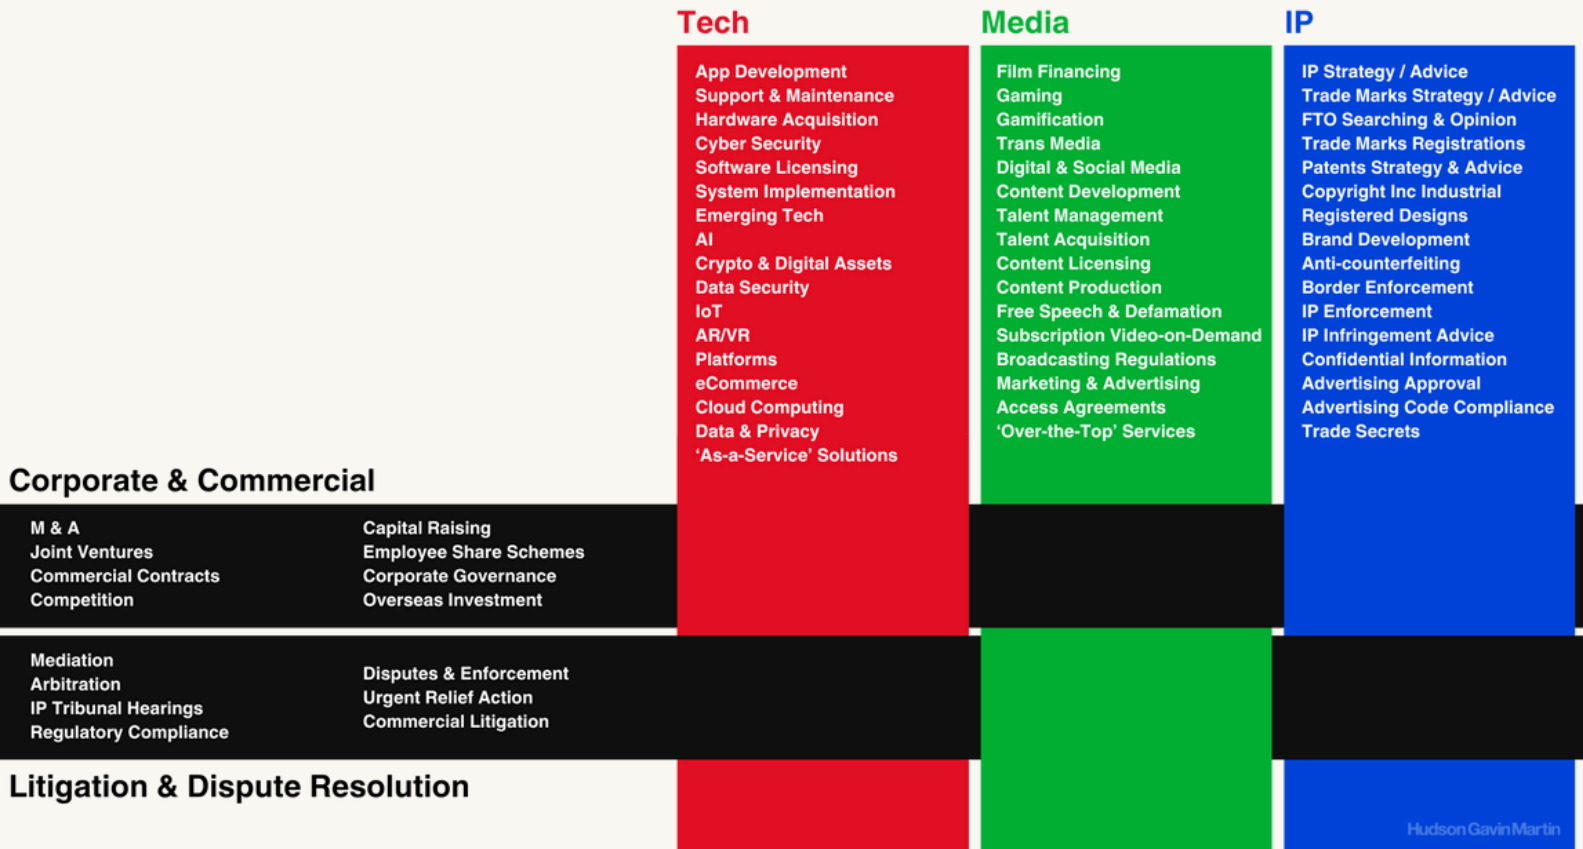

Navigate with an expert

Simon Martin
Partner
simon.martin@hgmlegal.com
A founding partner, Simon understands the challenges thrown at businesses and enjoys working with them to help maximise the value that can be gained from technology.

Ed Lim
Partner
edwin.lim@hgmlegal.com
Ed is known for his unique blend of tech and legal expertise, and offers comprehensive advice on business technology issues.

Anchali Anandanayagam
Partner
anchali.anandanayagam@hgmlegal.com
Anchali specialises in data protection and transactions where technology (including AI), media and telecommunications assets are at the core.

Andrew Dentice
Partner
andrew.dentice@hgmlegal.com
Andrew is an expert in fintech, technology transactions, data protection and commercialisation, and regulatory issues applying to technology.

Checklist:
Managing Gen AI risks

Identify your needs

- ☐ Clearly define the problem you are trying to solve with a Gen AI tool – what are your use cases?
- ☐ Identify the Gen AI tools available to meet your needs. Do thorough due diligence on the AI tools available in the market, and the vendors.
- ☐ Assess whether the AI tool is appropriate for the size of your business. How will it be deployed and what are the (total) costs? How will you train staff to use it effectively?
- ☐ Evaluate the technology – testing different AI tools is an essential part of the procurement process. Only anonymised and de-identified or fictional data should be used for testing.

Assess data risks

- ☐ Determine what data the Gen AI tool has been trained on and check if the tool is training itself on user data. This may have adverse effects on how your information will be protected.
- ☐ Evaluate what data your users will be inputting into the tool. Ensure compliance with privacy obligations for personal information. Conduct a privacy impact assessment if appropriate.
- ☐ Understand your third party data obligations under relevant contracts.
- ☐ Protect confidential information to avoid unintended financial impacts.
- ☐ Don't forget to assess risks that arise from other parties in your supply chain using AI tools.

Ensure output accuracy

- ☐ Evaluate the risks to your business based on what content will be produced and how it will be used – how will you mitigate these risks?
- ☐ Be aware of the risk of hallucinations and potential bias in the training data. Ensure accuracy in your content and that diverse perspectives are represented.
- ☐ Implement a "human in the loop" for high-risk outputs. Be mindful of reputational risk, not just legal risk.

Checklist:
Managing Gen AI risks

Intellectual property considerations

- ☐ Address any risks related to copyright infringement in creating outputs.
- ☐ Clarify ownership of generated content, especially for commercial use.
- ☐ Ensure proper use of any proprietary data in training the AI tool.
- ☐ Review licensing terms before the commercial use of outputs.

Contractual terms

- ☐ There should be a written contract in place between you and the supplier of the AI tool.
- ☐ Get legal advice – your contractual terms should mitigate your risk to the greatest extent possible.
- ☐ At a minimum, your contract should adequately address data protection/information security, privacy, confidentiality and intellectual property rights.

Governance and policies

- ☐ Conduct a stocktake of your current governance framework to determine whether it adequately addresses responsible AI use. Review any existing policies related to information security, technology procurement, data privacy, IP and trade secrets, customer care, and organisational values. Ensure your staff have clear direction about how they may use AI at work.
- ☐ If necessary, draft new policies and procedures, incorporating ethical use, legal compliance, quality control, transparency/explainability, risk management, consistency, and accountability principles.
- ☐ Regularly update your policies and procedures to keep pace with technological and regulatory changes.

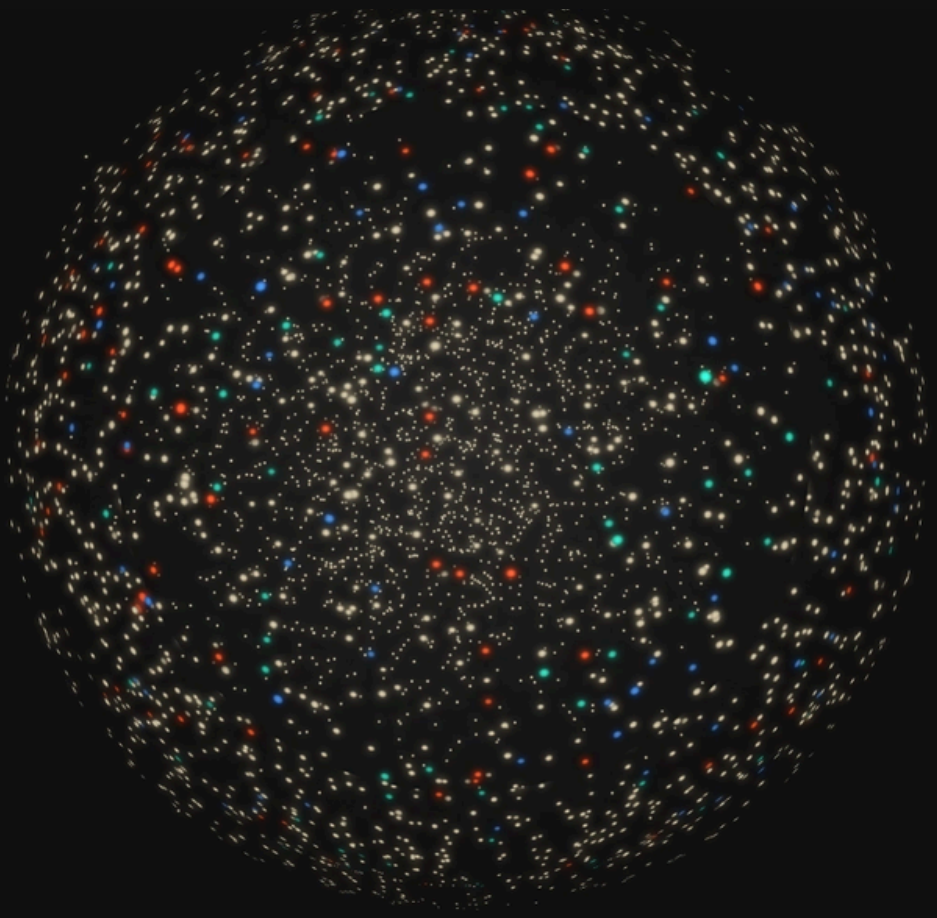

HGM Insights

Scan the QR code to see our latest Insights.
www.hgmlegal.com