

SHOOSMITHS

AI for in-house lawyers

A guide to using AI in the real world—and taking your business to the next level

www.shoosmiths.com

**FOR
WHAT
MATTERS**



Contents

01	Introduction	3
02	Key findings from our survey	4
03	10 principles	12
	Principle 1: Understand what AI is and what it can and can't do	13
	Principle 2: Don't believe all the AI hype, but don't use this as an excuse to do nothing	17
	Principle 3: Having an AI strategy	19
	Principle 4: Secure leadership buy-in for AI	21
	Principle 5: Cultivate trust in AI	23
	Principle 6: Keep a human in the loop	25
	Principle 7: Don't give with one AI hand and then take with the other	27
	Principle 8: Procure wisely	29
	Principle 9: 'Good enough' in AI is often enough	31
	Principle 10: Train your people in AI	33
04	Conclusion	35



Introduction

AI is here to destroy us.

Or it's here to save us.

So which one is it?

As it turns out: it's neither. We're now at the stage where a bit of AI realism is what businesses want and need to see.

In this guide, we set out what you need to know, as an in-house lawyer, about using AI in your business—particularly generative AI (GenAI)—whether in your team or more widely. We've been using AI for many years now, so we know the pitfalls to avoid and the opportunities to be seized.

We also set out the results of our survey with in-house lawyers so you can benchmark where you are against your in-house peers. Some of the results may well surprise you.

The guide also considers the key themes of the General Counsel Strategy Summit 2024: strategy, leadership, innovation and trust. As we set out in this guide, if you get these key things right, you'll be in a strong position to take your use of AI to the next level—and possibly to places that you just can't imagine now.

Enjoy the conference. And please do come and have a chat with us if you want to discuss how we can help you or, of course, if you just want a chat generally. There's never been a more exciting time to be an in-house lawyer. We can't wait to show you why!

All the best,

David Jackson
CHIEF EXECUTIVE

Tony Randle
PARTNER, CLIENT TECHNOLOGY & SERVICE IMPROVEMENT

Key findings from our survey

60% say that AI is made available in their businesses



A small majority of businesses make AI available to their people


AI has the potential to transform productivity by automating routine tasks, allowing individuals to focus on more complex and creative matters. It can also process large amounts of data rapidly, which could help in-house teams with making better decisions and improving productivity.

One (hidden) issue, however, is that many employees may be using AI 'in the shadows'. In other words, it hasn't been officially made available by their businesses. One BBC report noted in October 2023 that "workers are finding ways to access the technology in secret, whether by simply blocking their screens from view or employing higher-tech solutions".

An AI policy is therefore essential, but, as the next survey result shows, most businesses don't have an AI policy.

Key findings from our survey

Only about **3 in 10** GCs say that they have an AI policy



A large majority of businesses don't have an AI policy


Despite many businesses applying greater focus on risk mitigation nowadays, it's a major weakness in their risk strategies that roughly seven in 10 businesses don't have an AI policy.

It's perhaps understandable given the patchy legal framework in this area, but 18 months after ChatGPT burst onto the scene, it's no longer sustainable. To seize the opportunities of AI responsibly, businesses need the guardrails from an AI policy.

No policy means that people in your business may be using AI in a risky way or in ways which are insecure, unethical or in breach of the law: think data breaches, biased decision-making, and non-compliance with regulatory standards.

Key findings from our survey

Right now, AI is having a small impact on in-house teams **(2 out of 10)**



The average in-house team is behind their business on AI

We asked: how much of a positive impact is AI having on your team now on a scale of one to 10 with one being not at all and 10 being very much? The average score is two out of 10.

Given the survey results elsewhere, this shouldn't perhaps come as a surprise. In September 2023, The Lawyer noted that only seven per cent of respondents were using AI in their teams, but 20% planned to do so within the 12-month period after that. Progress is happening, albeit slowly.

The risk is that in-house teams could well fall behind their businesses in their approach to AI. This means these teams might struggle to provide timely legal insights, which could cause their businesses to miss out on opportunities or fall behind competitors which use AI for better-informed and faster decision-making. They might also face inefficiencies, leading to increased operational costs and too much time spent on routine tasks.

Key findings from our survey

Businesses' approach to AI is at a more advanced stage **(5 out of 10)**



The average business is still getting to grips with AI

We asked: how progressive is your business' approach to AI on a scale of one to 10 with one being inactive and 10 being very progressive? The average score? Around five out of 10.

It's clear that the average business is at a more advanced stage than their in-house teams. If your business hasn't set off yet, the message is: don't worry, you're not too far behind. But now's not the time to wait either.

Key findings from our survey

In-house lawyers see a bigger future impact for AI, but perhaps not as much as you'd expect **(6 out of 10)**



The average in-house team sees a future with AI, but not one in which it dominates

We asked a future-facing question: how much of a positive impact will AI have on your team in 2030? The average score is about six out of 10.

Do in-house teams still need to be convinced by the promise of AI? Has there been far too much hype?

Key findings from our survey

1/3 have no autonomy to implement AI; **1/3** have limited autonomy; and **1/3** have full autonomy



A third of in-house lawyers have full autonomy to implement AI in their teams

This survey result is equally split. In-house teams with less or no autonomy could miss out on AI's advantages, potentially hindering their businesses' overall competitive position and the teams' positions within them.

In-house lawyers ideally need some autonomy to implement AI in their teams. That's because they know their teams' requirements and can tailor solutions to enhance their efficiency and drive innovation specific to their needs.

Key findings from our survey

Main feeling
about AI: it's OK,
but not perfect



The main sentiment on AI: realism

This is good news. Putting AI on a pedestal and expecting it to be perfect means that in-house teams are on a hiding to nothing.

Realism in what AI can and can't do means that you understand its practical limitations and capabilities; and can set appropriate expectations, applying them in a way that complements human skills and enhances decision-making processes.

Key findings from our survey

Most popular AI tools used are ChatGPT and Copilot



OpenAI's ChatGPT and Microsoft Copilot are in-house lawyers' most used AI tools

Other AI tools from the following are used, but much less frequently:

- Anthropic's Claude
- Casetext's CoCounsel (part of Thomson Reuters)
- Counsel AI's Harvey
- Google's Gemini (ex-Bard)
- Kira
- Luminance
- ThoughtRiver

10 Principles

PRINCIPLE 1

Understand what AI is and what it can and can't do

Like a microwave, it can be used for some tasks—and is remarkable at what it can do—but it can't do it all.

Key takeaway: don't waste time trying to make AI do something that it is (currently) bad at

PRINCIPLE 2

Don't believe all the AI hype, but don't use this as an excuse to do nothing

AI is surrounded by a mix of exaggerated claims and underestimated potential. Recognise the hype for what it is: a blend of heady aspirations and sensationalised marketing. However, this shouldn't lead to inaction or complacency.

Key takeaway: embrace the practical applications where AI excels, and remain open to the continuous evolution of its abilities

PRINCIPLE 3

Have an AI strategy

Having a clear AI strategy is essential. It's your roadmap to define how you'll use AI, set goals, and ensure responsible implementation.

Key takeaway: create an adaptable AI strategy that aligns with your objectives and allows for future objectives too

PRINCIPLE 4

Secure leadership buy in for AI

To navigate AI successfully, you need to gain the support of the senior leadership in your business. Their commitment ensures alignment, resources, and sustained progress.

Key takeaway: you need to show why AI matters to the business and how it can improve the bottom line

PRINCIPLE 5

Cultivate trust in AI

If people don't trust AI, they're less likely to use it, limiting its potential benefits. Users need to believe that AI will function as expected and handle tasks effectively. Trust is built when AI consistently demonstrates reliability.

Key takeaway: trust is the bedrock upon which successful adoption rests

PRINCIPLE 6

Keep a human in the loop

You need to be in the loop at the right stage. Given the predictive nature of GenAI, it is essential to keep a keen eye on it and spot errors (particularly in a legal context!).

Key takeaway: don't just give a task to AI and then forget about it

PRINCIPLE 7

Don't give with one AI hand and then take with the other

You can use GenAI to summarise emails. Great! So don't then encourage the use of GenAI to create emails which are longer than necessary, taking you back to square one.

Key takeaway: productivity gains aren't a given. Be aware of the pitfalls and educate your people to avoid them

PRINCIPLE 8

Procure wisely

Don't enter into long-term contracts if you can avoid it. Although the scope to negotiate with large tech suppliers might be limited, try to create breaks to keep your future options open. This may cost more in the short term, but you will reap the benefits later.

Key takeaway: flexibility in procurement allows you to adapt swiftly to advancements and seize new opportunities

PRINCIPLE 9

'Good enough' in AI is often enough

Remember that 'good' can be a powerful ally. Striving for flawless AI can lead to diminishing returns. Sometimes, achieving 'good enough' is the true victory.

Key takeaway: don't focus on perfection. Instead focus on practical solutions that meet your needs effectively

PRINCIPLE 10

Train your people in AI

Empower your team with AI knowledge and skills—don't presume that they'll know about AI. Training, such as on crafting great prompts, is key to unlocking AI within your business.

Key takeaway: equip your people with AI training, such as on prompts

Understand what AI is and what it can and can't do

Understand what AI is

Just as you don't need to be a rocket scientist to understand what rockets can do, you don't need a PhD in AI to see what it can do.

This guide gives you some tips to work out how AI can help your team and your business.

Our biggest tip of all? Speak with people you know who have started their AI journey. These interactions are more than just conversations, they can give you invaluable first-hand insight and practical knowledge: things that rarely feature in guides or books (which are often out of date in any event).

Understand what AI can and can't do






You can't have it all. Similarly, AI can't do it all.

Take GenAI. A simple way to grasp the possible impact of this remarkable technology is to compare it to something many of us use in our daily lives: the microwave.

It's something that we hardly notice in our kitchens, but we often overlook how remarkable it is. It can reduce the time it takes to do some things by almost 90%, but it has its limitations. In particular, understanding how large language models (LLMs)—the models which underpin GenAI—work is the first step in understanding how this type of AI can be introduced into your business (See “Myth busting: all AI is the same” below).

What GenAI is good at:

Our experience is that GenAI is currently particularly good at the following. It still needs proper human supervision though (see Principle 6):

Task	Details
 Summarising long documents or chains of emails	Using AI can accelerate how you gather insights, which can help to take your discussions with other teams in your business to the next level. Imagine AI condensing a 50-page report into one page of key takeaways—in less time it takes to make a cup of tea
 Summarising meetings held in Teams	This is particularly useful for those who are unable to make a meeting or who join it late. The underlying transcript of the meeting can also be referred to if needs be. Time savings can be considerable
 Creating initial drafts of action points from meetings held in Teams	The result? All participants can concentrate more on what's being discussed. There's no more fretting about who's agreed to do what
 Re-writing texts in a different tone or style:	Nuanced writing helps to build relationships and manage conflict. In more complex or sensitive matters, the style and tone can significantly impact the reader's reaction to what's written. But getting the nuance right can be time-consuming. The highly-developed linguistic power of LLMs can make your writing resonate with the intended audience, at game-changing speeds
 Finetuning and polishing texts	AI can help improve and refine documents. Following your business' style guide or expectations of internal teams has never been easier

Understand what AI is and what it can and can't do CONTINUED

What GenAI is currently less good at:

Use cases for GenAI need to be considered carefully: the current iterations of LLMs mean that they tend to 'hallucinate'. In other words, they make things up, including legal authorities. For example:

- in June 2023, a US federal judge ordered two attorneys to each pay a fine of \$5,000 after they submitted legal briefs using case citations invented by ChatGPT such as 'Varghese v. China Southern Airlines Co., Ltd., 925 F.3d 1339 (11th Cir. 2019)'
- in the UK, in *Harber v Commissioners for His Majesty's Revenue and Customs* [2023] UKFTT 1007 (TC), a respondent cited case law which was invented by GenAI. In this case, the Tribunal said that its decision wasn't influenced by the use of the fake cases, but this won't be the only case of its kind in the UK

As more lawyers use GenAI, more guidance is being published. In December 2023, for example, **guidance for judicial office holders on AI** was published ("Judicial Guidance"). This guidance recommends not using GenAI tools (yet) for:

- legal research: "AI tools are a poor way of conducting research to find new information you cannot verify independently. They may be useful as a way to be reminded of material you would recognise as correct", or
- legal analysis: "the current public AI chatbots do not produce convincing analysis or reasoning"

Keep up to date with developments

Technology is advancing all the while. GenAI is getting better at many things—at a fair clip! As OpenAI notes, "the amount of compute used in the largest AI training runs has been increasing exponentially with a 3.4-month doubling time". So keeping an eye on the market is important.

For example, some AI tools, such as Harvey, are now used for legal research. Both Thomson Reuters and LexisNexis are also incorporating GenAI in their products.

But these AI tools can come at a significant cost: AI tools are expensive to build.

Understand the law and guidance

At the time of writing, the law on AI is 'light touch' and patchy in the UK, but this is likely to change in the medium to long term. That said, it's unclear what will happen to AI regulation after the next UK general election and whether, for example, more obligations will be put on a statutory footing.

One in-house lawyer in the survey told us that depending too much on GenAI, which is trained on various sources of information online, may create more problems later for users in terms of infringement of intellectual property rights.

In the short-term, however, the approach in the UK is to take a sector-by-sector approach.

In the legal sector:

- the Solicitors Regulation Authority (SRA) and others have already provided guidance on AI, which should be considered carefully
- in April 2024, the Legal Services Board also issued its approach to AI regulation, which means that the bodies it regulates, such as the SRA, may well issue further rules and guidelines in the near future

In the meantime, some laws do apply to AI. Under data protection law, for example, individuals have the right not to be subject to a decision based solely on automated processing, which produces legal effects concerning them or significantly affects them.

In the EU, the far-reaching Artificial Intelligence Act was adopted by the European Parliament in March 2024. The 'Brussels effect' could well mean that this becomes the de facto standard for AI regulation in all of Europe, like the General Data Protection Regulation (GDPR). Watch this space.

Understand what AI is and what it can and can't do CONTINUED

Myth busting: all AI is the same

AI models are computer programs that learn from data and can do different things with it.

There are broadly two main types of AI models: discriminative and generative.

A discriminative model learns the patterns and boundaries in data. It can categorise new comparable data by references to categories the user sets:

- an example in the context of law is that it can take, for example, many commercial contracts and compare them to pre-identified principles; eg, are there caps on liability?
- it's trained on examples of data that are pre-categorised and labelled by the user. It learns how to place new data into the same categories by reference to the common features of the training data in each category
- these forms of AI take a relatively long time to train and configure as they need many examples of data with labels to learn well, and they can sometimes get confused by new data that is different from what they learned
- the model is often used in legal document analysis and categorisation, including commercial contract review and due diligence exercises (for instance Shoosmiths' Cia[®], powered by AI platform ThoughtRiver) as it is an accurate means of identifying a comprehensive but finite list of issues and performing simple associated tasks, such as providing the relevant corrective drafting
- discriminative AI models do not create new content

Generative models create new content, such as text, images, audio and video:

- they are trained on data but, unlike discriminative AI, the data does not need to be labelled. Rather than simply learning the patterns or boundaries they use the patterns of the data to make new data that resembles the original data. They do this with the use of complex algorithms that allow it to produce new text or images with similar statistical properties to the data on which they trained. But as the Judicial Guidance notes, many LLMs 'view' of the law "is often based heavily on US law although some do purport to be able to distinguish between that and English law"
- they are used for creative tasks such as drafting correspondence or documents based on input criteria or producing a precis of lengthy texts (eg, a summary of a transcript)
- they can sometimes make up data that is wrong ('hallucinations') or of low-quality, and due to the complexity of the algorithms creating the content, it can be difficult to check their output. This means that it could, for example, give wrong or deceptive information about the law in general or how it could apply
- most law firms and legal teams are working hard to produce suitable guard rails for any use of this technology. At present, this type of AI is good at low risk administrative tasks (such as the first draft of an email or a summary of a meeting) but, as the guide shows, there are many things that it isn't good at, although this is likely to change in the future

In summary, discriminative models classify existing data, while generative models create new data based on learned patterns.

Understand what AI is and what it can and can't do CONTINUED

Actions

- understand the capabilities of AI, particularly LLMs, to appreciate their potential and limitations in legal and commercial contexts
- experiment with AI tools: see what they're good at; see what they're not so good at. Every AI tool is different and has a different risk profile. The more you test these tools, the more you will understand them and the better you will understand how they can help in your work:
 - our experience is that it is often useful to use this technology after you have undertaken the work yourself to see if you have missed something
 - as the Judicial Guidance notes, AI tools can typically help you find material that you already know is right, but you don't have access to. But they aren't a good way of doing research to find new information that you can't check. In other words, they may be more useful as a way of getting tentative confirmation of something, rather than giving you accurate facts right away
 - try out prompts and see what happens (see Principle 10)
- keep abreast of developments in AI legislation and regulations:
 - for an overview of the law surrounding artificial intelligence across a variety of jurisdictions, check out the [Legal 500 Country Comparative Guide on artificial intelligence](#). Shoosmiths wrote the UK chapter
 - for developments on AI in the UK, sign up to receive Compendium, Shoosmiths' quarterly legislation tracker, designed to give you early insight into legal changes on the horizon, enabling you to be one step ahead. [Register here to receive a copy of the next quarterly edition](#)
- be aware of the risks of intellectual property infringement when using AI

Don't believe all the AI hype, but don't use this as an excuse to do nothing

Don't believe all the hype

Reporting on AI generally tends to be somewhat hyperbolic and 2023 was a powerful case study on this tendency. In November of the previous year ChatGPT burst onto the scene. In less than one week it went from zero to one million users.

However, getting practical guidance on what worked or didn't work in practice wasn't high on the agenda. And it is still often lacking.

But do something

But things are changing. 2024 is the year of the use case.

We've trialled GenAI and we now have a much better understanding of what it can and can't do (see Principle 1). Many other law firms and in-house lawyers are also trialling or using AI tools.

The key is to start somewhere.

And there's no time like the present to do so.

However, as the actions below indicate, any decision to use AI should not be made hastily or carelessly, but thoughtfully and strategically.

Myth busting: it's all hype

Roy Amara, a researcher and computer scientist at Stanford University coined a new law to describe the effect of new technologies.

Amara said "we tend to overestimate the effect of a technology in the short run" (hence the hype), but we "underestimate the effect in the long run".

In practice, in the short term, this means that businesses need to take care not to invest too much too quickly in new AI solutions. They might not yield the expected benefits.

In the long term, however, doing nothing is a risky strategy. Failing to grasp the long-term consequences of this technology risks your business not being ready for its transformative effect. This can impact strategic thinking and risk of being left behind by competitors.

For us, it's less about the hype and more about getting the timing right.

Don't believe all the AI hype, but don't use this as an excuse to do nothing CONTINUED

Actions

- think about what tasks AI (and particularly GenAI) may be able to help you with:
 - what tasks really need human input?
 - what tasks do you really enjoy doing?
 - what tasks do you consider as being 'drudge work'? Can AI help?
- work out why you're looking at AI in the first place: start with a problem. What's the end point? What's the business need?
- swap notes with:
 - the law firms you instruct. What are they doing? Can they help you with AI?
 - any in-house lawyers you know. Again, what are they up to? Can you join forces with other in-house lawyers in some way?
- stop waiting: waiting for a better time to start; waiting for IT or other colleagues to start; waiting for anyone to start the process!
- start modestly. If you're inexperienced with using AI, be careful before committing fully
- think about a pilot: starting with a smaller group of people not only gives you time to fix any problems but also lets you show the positive results that have been achieved from the pilot
- get some wins under your belt: to capture the attention of those who matter, you need to stand out: showing immediate benefits through small-scale AI projects can secure early buy-in. Long waits for complex AI tools risk diminishing interest and support. Consistently delivering small successes maintains momentum and engagement
- bear in mind that we're not alone in facing challenges to make the most of AI. Other professions also have difficulties. Look at how they overcome them and get insights from them
- remember AI may not work straight out of the box. Some AI may need to be trained, which can take a lot of time and effort. Is your internal data up to scratch?

Have an AI strategy

Having a clear AI strategy is essential. It's your roadmap to define how you'll use AI, set goals, and ensure responsible implementation.

AI strategy

Any strategy is as much about what a business does as about what it doesn't do. AI is no different.

An AI strategy might include, or refer to, the following:

- executive summary: a brief overview of the AI strategy and its alignment with the business' overall strategy and other strategies
- an evaluation of the existing technology, data, and skill levels in the business
- vision and objectives:
 - a clear articulation of the long-term vision for AI in the business
 - specific, measurable, achievable, relevant, and time-bound (SMART) objectives
- policies for data quality, security, privacy, and ethical usage
- an overview of compliance with relevant laws, including data protection laws
- a strategy for developing or acquiring AI expertise
- a plan for the necessary AI technology and infrastructure investments
- outline of initial small-scale pilot projects to test AI concepts
- approach for integrating AI into existing processes and scaling it across the business
- guidelines to ensure AI is used ethically and responsibly
- metrics and KPIs to track the performance of AI initiatives
- process for ongoing learning and adaptation of the AI strategy

AI policy

One major task for in-house lawyers is typically to mitigate risk in their businesses. Lacking an AI policy is now a major risk.

Without an AI policy, your business is exposed to various risks when using AI, such as data breaches, biased decisions, and regulatory non-compliance.

We offer a template which serves as a starting point for businesses to develop a comprehensive AI strategy and policy tailored to their specific needs and goals. **[You can sign up to receive this here.](#)**

Allowing the use of AI should typically be subject to certain conditions, although these depend on each individual AI tool and the extent to which any data which you input into it is protected by the provider of the AI tool. Conditions might include:

- not using any personal or confidential data
- not putting your organisation's name, including any other trading names (or any product names) into it
- not asking it any legal, moral or ethical questions
- not using it for any sensitive work
- not using it where there is the risk that third parties will do the same resulting in similar copy (such as on a pitch)
- not using any data which includes third party IP or any data which is in breach of any third-party IP (or appears to be)
- being mindful of any output that is or may be biased or discriminatory, and
- being sure to consider and sense-check any output.

Depending on the AI tool in question, you may need to limit or restrict the use of certain AI tools in the following circumstances, in particular:

- providing legal advice
- working on internal work with legal implications, or
- undertaking any legal research.

As with any policy, there may be exceptions and exclusions for certain AI tools.

Have an AI strategy CONTINUED

In-house insights: staying in the loop

An in-house lawyer in our survey was concerned that non-lawyers in their business are using AI, but they aren't asking for the guidance they need from legal. This could result in bad outcomes for their business—in numerous ways, expected and unexpected.

This is why an AI policy is an essential tool in managing how AI is used, together with guidance on when in-house legal support needs to be requested. Taking the actions set out below can help businesses to minimise risk and maximise the opportunities AI presents.

Actions

- download our AI policy and related documents (see above). Our carefully crafted draft GenAI policy is designed to serve as a foundational reference for in-house lawyers. It will empower you to develop a policy tailored to your business's specific needs and challenges
- consider whether any other strategies and policies need to be amended when introducing an AI strategy and AI policy
- regularly review the AI strategy and AI policy and make any necessary amendments in light of developments

Secure leadership buy-in for AI

To succeed with AI, you need to get the backing from senior leaders in your business.

Their visible endorsement ensures that AI initiatives tie in with the business' long-term goals:

- in most sectors and industries, AI will disrupt existing processes and systems. Senior leadership buy-in is vital to manage this change, foster a culture of innovation, and guide the business through a smooth transition
- AI projects require significant investment too. Senior leadership buy-in ensures that the necessary financial and human resources are allocated, ensuring the project's viability and sustainability
- AI implementation also comes with its share of risks and legal implications. Senior management's involvement is essential to navigate these challenges, ensuring that AI solutions comply with regulatory requirements and ethical standards. Inevitably the majority of the day-to-day implementation of AI will be undertaken by your IT and projects teams

On our Copilot trial, weekly emails from the CEO (see Principle 10) have become a great way to show our business that AI matters to it—and the people who use it too.

In-house insights: don't mention 'AI'

What happens when you refer to 'AI'?

Nowadays, it's such a catch-all term, it's become as meaningless as statements like 'we must use more technology'. (What technology? How? Where?)

As a result, one in-house lawyer in the survey tells us that the term 'AI' is losing its cachet with senior leaders and therefore impact.

When talking about AI, it makes sense to be specific and avoid being generic. In other words, be clear about what solution you're referring to so that senior leaders can assess its relevance.

Secure leadership buy-in for AI CONTINUED

Actions

- talk to senior leaders to find out what they want and hope from AI (although, as referenced above, try to be specific, and not generic)
- explain the benefits of specific AI solutions to senior leaders in clear language—try to avoid jargon—to gain their backing, dealing with any downsides and offering clear guidance on, for example, the legal consequences of using AI
- consider writing a business case for senior leaders which highlights the benefits, costs, and risks associated with AI. Include potential return on investment, productivity gains, and improvements in legal compliance. (You will find that good suppliers might be able to help you construct a business case for a particular product once you have independently ascertained its value to the business)
- define clear KPIs for AI projects in consultation with senior leaders. In this respect, consider using any materials developed by AI providers. Microsoft, for example, has suggested Copilot survey questions which mean businesses can compare them with Microsoft's benchmark results. It also has a Copilot dashboard, powered by Viva Insights, which provides metrics to help determine, eg, time saved
- obtain agreement for the required budget and resources from senior leaders, including training on AI
- work with senior leaders to implement change management strategies to help the legal team and the rest of the business adapt to AI
- create systems and processes to monitor progress and evaluate the effect of AI
- regularly review AI implementation with senior leaders and be prepared to adjust strategies based on performance data and feedback. Share key insights and trends with senior leaders to inform future AI use. After all, senior leadership buy-in isn't a one-off event!
- document the difficulties, achievements, and insights that occur during the AI implementation. This will ensure that you have the information to hand when asked tricky questions
- work with senior leaders to develop risk management strategies for AI implementation

Cultivate trust in AI

Trust is something that is too often overlooked.

If people don't trust AI and its output, they're less likely to use it, limiting its potential benefits.

Businesses also need trust from their staff that AI will make many of the tasks that they do easier, but that it won't ultimately replace them.

As AI deployment goes from small to large-scale, the issue of trust is likely to grow and will need to be managed carefully.

But the trust that we have in AI mustn't be misplaced. The SRA in its [Risk Outlook series](#) (published on 20 November 2023) refers to, for example, the problem that "people often put more trust in computers than in humans".

To build trust in AI, for example, at Shoosmiths we've committed to:

- adhere to any relevant ethical guidance governing the usage of AI technology
- comply with all relevant laws, including data protection laws
- promote the responsible use of AI technology
- regularly evaluate our AI usage and identify any emerging risks
- undertake periodic reviews of our AI related policies and procedures, as necessary, to reflect any regulatory or legislative changes

We also place a lot of importance on the human side of the AI equation. So we're looking to ensure, for example, that our lawyers realise the importance of both technical and emotional intelligence (EQ) skills to use AI effectively. For us, using AI and EQ together—our trump card in an era of accelerated change—will help the users of AI have confidence in its output.

Myth busting: GenAI is intelligent

GenAI isn't intelligent as such—it just seems it. In essence, as a large statistical model, it doesn't 'know' what it is doing.

So we asked ChatGPT whether it is intelligent. The chatbot 'admitted' that "it doesn't possess true intelligence or consciousness" and that it "lacks understanding and self-awareness".

There's no doubt, however, that people have developed a strong emotional response to GenAI as the way the information is presented—as though someone is typing—makes it seem as though they are talking to a 'living' machine. (Most of us say 'please' when giving it an instruction!)

Cultivate trust in AI CONTINUED

Actions

- teach your people on how trust is essential for effective AI use (see Principle 10)
- keep up with the law, ethical standards, and professional conduct rules on AI (see Principle 1). Many of these have cultivating trust at the heart of them, such as data protection rules (ie, where AI uses personal data)
- support responsible AI deployment, with support from senior leaders (see Principle 4)
- communicate openly with stakeholders about AI use
- perform ongoing evaluations of AI applications to identify and address emerging risks, which may have an impact on trust
- involve your people who use AI in discussions about its advantages and disadvantages
- regularly review and update AI policies and procedures for regulatory and legislative compliance
- stay responsible for all legal services with AI involvement, ie, keep a human in the loop (see Principle 6)
- create protocols to supervise AI decision-making processes and outcomes
- implement risk management strategies for AI, including planning for potential AI errors

Keep a human in the loop

Most AI needs a human to monitor its output. And, of course, it also needs a human to ask the right questions of it in the first place.

In most cases, we can't just let AI do what it does, without their being a 'human in the loop'. Or, to put another way, we're the thought leaders and AI is the 'thought partner'.

For example, research suggests that GenAI works better where users keep a keen eye on it and thus can spot errors. In effect, the more wary we are about the output which is created, the better. For work undertaken by lawyers, this is an understatement! Our experience is that AI can include some or a lot of inaccuracy. Checking is imperative.

We also need to give the correct instructions when using this technology (see Principle 10 on training). Just as we wouldn't get much from a trainee solicitor if we only said 'do it' without much context, the same applies to GenAI. It requires us to be as precise as possible.

In effect, having a human in the loop:

- provides context, nuance, and critical thinking that AI may lack
- can ensure a higher level of precision for tasks which require subjective judgment or complex decision-making
- means that novel situations or context are dealt with better. Equally, it means that decisions based on nuanced information are likely to be better, given that AI can often struggle with ambiguity and unexpected scenarios
- can identify and rectify biased outcomes, given that AI might inadvertently perpetuate biases present in their training data

There are also obligations to comply with and guidance to follow:

- the SRA notes, for example, that "solicitors need to remember that they are still accountable to clients for the services provided, whether or not external AI is used"
- the Judicial Guidance also makes it clear that work needs to be checked. On summaries, for example, "care needs to be taken to ensure the summary is accurate"

Myth busting: GenAI is out to get us all

The non-profit collaboration Better images of AI notes that pictures of "humanoid robots, glowing brains, outstretched robot hands, blue backgrounds, and the Terminator" aren't helping us understand or engage with AI.

Put simply, hackneyed sci-fi images of blank-faced humanoids are making many people wary of this remarkable new technology.

But the AI in use now isn't dystopian. At Shoosmiths, for example, we've developed AI tools to review contracts at much quicker speeds, freeing up lawyers' time to concentrate on more pressing matters.

Keep a human in the loop CONTINUED

Actions

- understand the AI technology and keep up to date on AI law and guidance (see Principle 1)
- implement an AI strategy and AI policy (see Principle 3)
- assess the legal, ethical and other risks associated with the AI's decision-making capabilities, especially regarding data privacy, intellectual property, and confidentiality
- establish clear procedures for human oversight to review and, if necessary, override AI-generated outcomes
- maintain transparency in the AI tool to ensure accountability and the ability to audit processes and decisions
- regularly monitor the AI's performance to ensure it aligns with your standards and legal requirements
- train your legal team and other stakeholders on the implications of AI technology and the importance of human oversight (see Principle 10)
- be ready to adapt policies and practices as AI technology and the legal landscape evolve

Don't give with one AI hand and then take with the other

It's not a given that AI will result in productivity gains.

We need to have agency with AI and not treat it as something that happens to us.

Many legal documents, for example, ballooned in size when word processing came into our offices. What effect will AI have on those legal documents now? Will lawyers' desire to deal with worse-case scenarios, turbocharged by AI, triple or quadruple the size of the documents we deal with day to day?

Take the following example: you've just received an unsolicited email from a seller, but you're not interested in what they are selling. Or you just don't have the time to look at it. You'd like to respond nicely with a 'thanks, but no thanks' email back. If requested, GenAI could create a long-form text as follows:

"I am writing to you in response to your email that you sent to me regarding your product. First of all, I would like to thank you for reaching out to us and for sharing the details of your product with us. I appreciate the time and effort that you have put into your presentation and I am impressed by the features and benefits that your product offers. I am glad that you considered us as a potential partner and that you value our expertise and experience in the field.

However, I regret to inform you that at this moment in time, we do not have the capacity to consider new products. Our team is currently focused on other projects and priorities that are aligned with our strategic goals and vision. We are unable to allocate the necessary resources to evaluate your offer and to conduct a thorough and fair assessment of your product. We have a limited budget and a tight schedule and we have to be very selective and careful about the decisions that we make.

I understand that this might not be the news that you were hoping for and I apologise for any inconvenience this may cause. Please know that this decision is not a reflection of the quality or value of your product, but rather a result of our current circumstances. I have no doubt that your product has a lot of potential and that it could benefit many customers and clients. I admire your passion and innovation and I encourage you to continue pursuing your goals and dreams.

Thank you again for your interest and for considering us. I wish you all the best in your future endeavours and hope that you will find success in bringing your product to market. Please keep in touch and let us know if there is anything we can do to support you in the future."

You send it to the vendor, who then uses GenAI to summarise it thus:

"[x] responded to an email sent to them and their manager regarding a product.

They appreciate the presentation and are impressed by the product, but regret to inform that they do not have the capacity to consider new products at the moment.

Their team is currently focused on other projects and priorities and are unable to allocate resources to evaluate the offer."

An exaggeration? Perhaps. But it shows how one side to a conversation can use AI to significantly increase the length of an email, with the other side then just cutting it down to size.

Aside from the seller feeling a bit better about themselves, having perhaps glanced at the niceties in the long-form email, we need to ask ourselves: how much has been gained by using AI here?

Don't give with one AI hand and then take with the other CONTINUED

Myth busting: GenAI always results in productivity gains

In 1955 C. Northcote Parkinson wrote in *The Economist* that “it is a commonplace observation that work expands so as to fill the time available for its completion”.

In his article, he observed that an individual could spend a whole day writing and sending a postcard to a family member in Bognor Regis, with time spent looking for a suitable postcard; more time composing it; plenty of time spent faffing; and, finally, 20 minutes spent working out “whether or not to take an umbrella when going to the pillar-box in the next street”. On the other hand, a busy individual could take about three minutes to do the same job.

Ultimately it's humans who need to work with GenAI. And it's humans who will decide whether GenAI will free up more time to do the work that really matters to them.

Actions

- be aware that new technology often raises the bar on what is expected from work, potentially increasing the workload needed to meet these expectations. Are you doing work that matters? Or doing work to simply fill the time allotted to it?
- set clear objectives for AI applications
- use AI to manage mundane and time-intensive legal tasks, saving time for more challenging issues that demand human expertise
- use AI to prioritise legal tasks based on urgency and importance, ensuring that time is allocated efficiently
- establish metrics and KPIs to monitor the results of using AI tools and make any changes necessary to optimise use and performance
- create a system for feedback on the AI's performance, allowing for continuous improvement and adaptation to your needs.
- stay informed about the latest AI tools and ensure adequate training is available to enhance work efficiency (see Principle 10)

Procure wisely

Procurement teams have sometimes faced the criticism of not being strategic enough. The result is that these teams are ignored or left alone without enough engagement from key stakeholders.

Put simply: procurement is much more complex than purchasing and investing in the things that a business needs; or bargaining for lower prices from suppliers. Procuring AI is no different.

What's more, it may be that not everyone in the business needs the same AI-powered solutions.

Principle 3 highlights the need to not jump in at the deep end when procuring new AI tools:

- start modestly. Don't commit fully
- always consider using pilots and trials before committing
- get some small wins under your belt
- remember that some AI may need to be trained

Shoosmiths case study: know what you are getting

We use Copilot for Microsoft 365. But this isn't the only AI-powered product that Microsoft offers. The company has a number of products which include AI, but where the functionality differs materially, for example:

- Copilot for Microsoft 365
- Teams Premium, and
- Copilot (formerly Bing Chat Enterprise).

At the time of writing, Teams Premium, for example, records meetings and produces an automatic summary of them, but you can't use additional prompts to interrogate meeting transcripts produced by AI, such as "what did Jane say?". This product is therefore sufficient for smaller, less critical meetings, but if we need to derive content from their discussions, we have found Copilot for Microsoft 365 to be more useful.

The names of many Microsoft products have changed recently, with 'Bing Chat' and 'Bing Chat Enterprise' being renamed 'Copilot' and 'Microsoft 365 Copilot' being renamed 'Copilot for Microsoft 365'. This had the potential of causing confusion, so we ensured that we communicated these changes to the relevant teams and updated our policies accordingly.

Procure wisely CONTINUED

Actions

- use pilots and trials to see whether the technology meets the needs of the business
- consider any non-tech alternatives. An (expensive) AI-powered solution may not always be the best solution. From our experience, this isn't something that is always considered in the face of shiny new tech
- consider whether everyone in the business needs the same AI-powered solutions
- consider any regulatory and procurement frameworks that may apply
- determine the needs or wants of the business (ie, the must-haves and the nice-to-haves). What are its goals? What resources are there? As many IT teams have asked over the years: 'do we buy it, do we build it or do we rent it?'
 - if you build it, will you partner with a third party?
 - if you rent it, what are the licensing terms?
- do you need a project methodology (useful if it is a large procurement)?
- work out a list of suppliers. Ask your in-house peers or law firms you instruct who they use. How happy are they with them?
- do due diligence on any suppliers, including potentially meeting with them
- determine what testing might be needed and collaborate with both technical team members and potential users of the system to test, experiment and collect feedback
- determine the terms and conditions and know how much room there is for bargaining. What KPIs are needed?

‘Good enough’ in AI is often enough

‘Good’ can be a useful ally.

Pursuing perfect AI can often result in diminishing returns, ie, the cost of edging towards 100% perfection becomes more and more expensive. As a result, reaching the sweet spot of ‘good enough’ can be the real success.

The notion that good AI is often enough reflects a pragmatic approach to the deployment and use of AI tools:

- AI that is ‘good enough’ can often be deployed faster, allowing businesses to seize opportunities and refine their models in real-time
- ‘good enough’ AI provides a balanced cost-benefit ratio, delivering value without breaking the bank
- ‘perfect’ AI may overwhelm users with its complexity. ‘Good enough’ AI, on the other hand, can help to facilitate a smoother adoption curve
- AI is never truly finished. ‘Good enough’ AI allows for continuous improvement, evolving with user feedback and changing requirements

This guide has already shown that AI is far from perfect, but it’s nonetheless a powerful tool, provided that, for example, there is a human in the loop (see Principle 6). The risks of ‘good enough’ AI also need to be considered and understood. At times, and for certain tasks, ‘good enough’ AI won’t suffice.

Shoosmiths case study: getting to an acceptable level of accuracy

Shoosmiths created an AI-powered tool, Cia, a few years ago to review certain commercial contracts faster than a human. After two years of feeding data from contracts and correcting Cia when it made mistakes (“marking Cia’s homework” as we called it), it’s now around 94% accurate.

This means, AI usually performs better than the average legal associate. In a controlled study on AI software in the United States, lawyers’ accuracy without using AI on checking non-disclosure agreements was, on average, nine percentage points lower at 85%.

'Good enough' in AI is often enough CONTINUED

Actions

- use metrics and KPIs, if you can, to understand the accuracy of AI. Compare this with the accuracy of humans
- for each AI tool risk assess how good it needs to be. In certain situations 'good enough' may not suffice, particularly where the stakes are higher, such as in high-risk areas
- identify areas where AI can have the most impact and prioritise development accordingly
- define what 'good enough' means for your project, focusing on core functionalities that deliver tangible benefits
- collect user feedback and performance data to inform iterative improvements
- encourage innovation but remain grounded in practicality, ensuring that AI solutions are viable and sustainable

Train your people in AI

Last, but by no means least, is training.

This is key to ensuring that AI works in your business. It's the bridge connecting the technology and the humans who will use it.

Training in other areas is just as important as training on AI too. One in-house lawyer in the survey expressed concern that junior team members might see their training in the job reduced as AI does some of the work typically given to such junior lawyers.

However, with EQ rising in importance in the work that we do—lawyers have countless skills which can't and won't be replaced by technology—in-house legal teams can now do more training on the 'soft' skills required. These will be needed more than ever!

Shoosmiths case study: training in AI

We use Copilot for Microsoft 365 and we were one of the first law firms globally to be on the trial for this AI tool. This is how we've trained our people and ensured continuous engagement:

- users complete mandatory eLearning modules to become familiar with the AI's current features and functionalities, and how to use them effectively and responsibly:
 - what Copilot is and what it can help with
 - how to access Copilot in different Microsoft apps, such as Word, Teams, Outlook, and PowerPoint
 - how to use the draft, summarise, and polish functionalities to generate, condense, and refine text
 - how to use prompting to guide Copilot and get the best results
 - why prompting is important and how to write effective prompts
- we encourage users to attend weekly drop-in sessions for technical updates and real-life examples of how colleagues are using AI. If users can't attend, they can watch recordings of the sessions

We set up a Teams site which is a go-to hub for collaboration, support and staying connected with other members. Here, users can:

- share successes: if they are impressed with something AI has helped them with, they can share their example in the 'Success Stories' channel. They can also visit this channel to learn top tips and best practice
- give feedback: we ask for feedback at regular intervals. However, users can submit their feedback at any point as their knowledge and experience of the AI grows
- ask for support: users use this channel to post any requests for support and access to a Q&A document
- access training resources, including videos of all training sessions, previous drop-in sessions and demonstrations.

Continues on next page...

Train your people in AI CONTINUED

Weekly emails are also sent out by the CEO which include:


- general and technical updates
- top tips
- details of the previous drop-in session

We encourage our teams to try it out. Here's some examples how:

- Outlook. Ask Copilot to: draft a communication on a topic of your choice; make your drafting more concise; improve your grammar and punctuation; summarise the content in a document; or write an executive summary
- Word. Ask Copilot to: draft a communication on a topic of your choice; make your drafting more concise; improve your grammar and punctuation; summarise the content in a document; or write an executive summary
- Teams. Use Copilot to: prepare a meeting agenda with your colleagues; summarise what was concluded in your meeting; list the action points; or generate an email to your colleagues to list the key points from your discussion

Actions

- don't assume that your team already has AI knowledge or expertise. Ensure that the training covers the basics of AI
- training on how to write effective prompts is essential to get the most out of AI in your business
- provide case studies and practical examples of how AI can be applied
- develop training on the legal and ethical use of AI, highlighting the importance of understanding the sources of training data for AI tools and checking for potential biases. Such training can also set out what your team needs to know about your AI strategy and AI policy (see Principle 3)
- train the team on how to identify and mitigate the risks associated with AI
- encourage ongoing learning on AI



Conclusion

Is there still a need for in-house lawyers in our AI-powered future?

In a word: yes!

As we look to 2030, we see in-house lawyers giving their businesses turnaround speed which was previously impossible, and so deal flow will increase.

But also, with much business-as-usual being handled by AI, in-house lawyers will be able to focus far more on strategic aspects, such as concentrating on the biggest, most complex matters, where AI will do much of the heavy lifting, but only humans can achieve the outcomes that the business needs.

Better still, in-house lawyers will be able to focus more on the processes within the business, to safeguard against problems occurring in the first place.

With AI enabling lawyers to actualise their full potential to add value to their business, you could describe this as the dawn of a golden age of in-house lawyering.

Much less firefighting.

Much more strategising.

We're excited. And we hope you are too.



**David
Jackson**

CHIEF EXECUTIVE

T +44 (0)3700 864 184

E david.jackson@shoosmiths.com



**Tony
Randle**

PARTNER, CLIENT TECHNOLOGY
& SERVICE IMPROVEMENT

T +44 (0)3700 864 243

E tony.randle@shoosmiths.com

DISCLAIMER

This information is for educational purposes only and does not constitute legal advice. It is recommended that specific professional advice is sought before acting on any of the information given.

©Shoosmiths LLP 2024

SHOOSMITHS

www.shoosmiths.com

**FOR
WHAT
MATTERS**