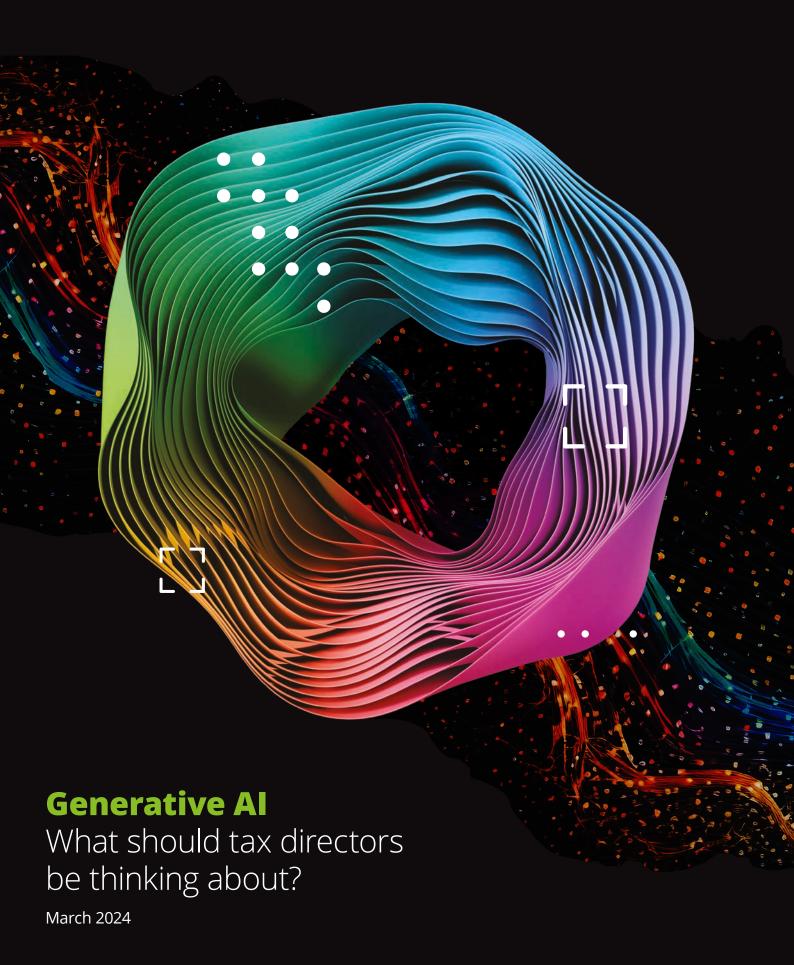
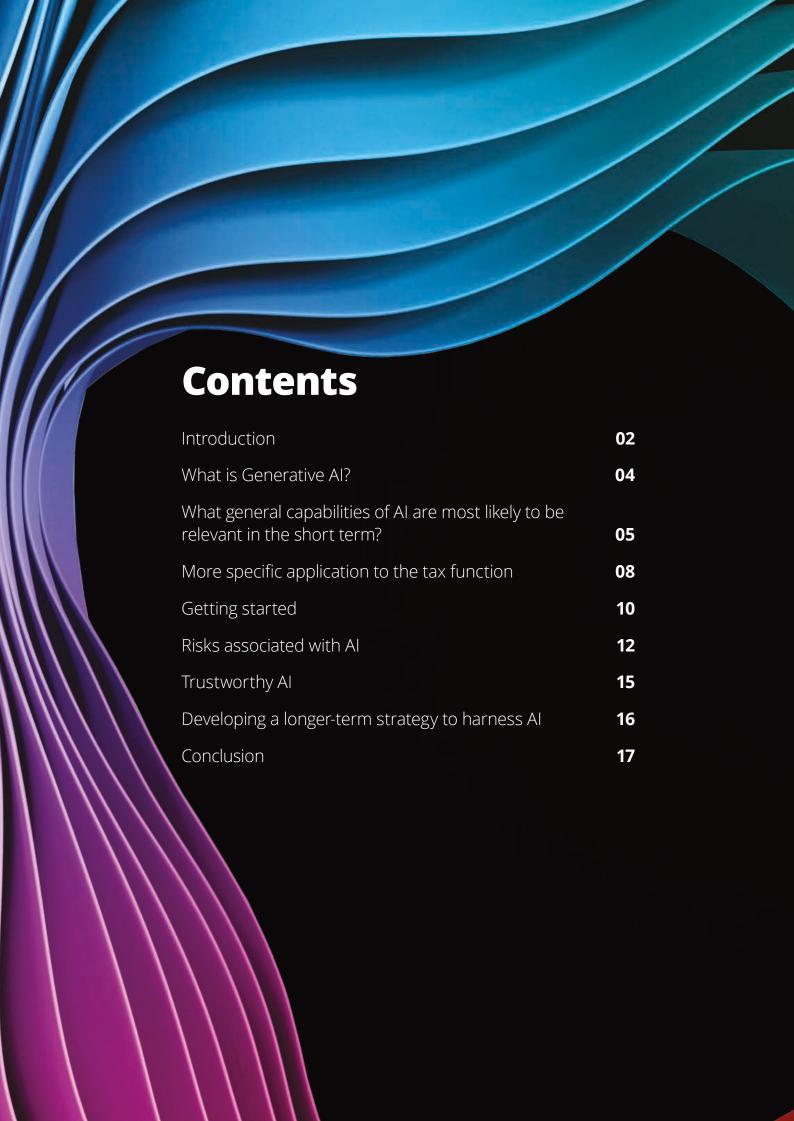
Deloitte.





The re-imagining and improvement of processes through deployment of AI is no longer just for technology specialists – it is becoming a core part of the role of a Tax Director.

Introduction

With the possible exception of Pillar Two, "Generative AI" (or "GenAI") was the buzz phrase of 2023 for tax professionals.

Al-enabled solutions have been available to tax teams for some time, but the emergence of large language models underpinning user-friendly chat interfaces has led to an explosion of interest and investment in new solutions. The re-imagining and improvement of processes through deployment of AI is no longer just for technology specialists – it is becoming a core part of the role of a Tax Director.

Rather than focus on a list of products and detailed use cases within tax, the aim of this publication is to help Tax Directors develop a framework for thinking strategically about how GenAl might be integrated within the tax function in the short, medium and long term. Even discounting some of the hyperbole about GenAl, it is sobering to think that the technology could develop faster than the time it typically takes for internal budgets to be secured for investment – simply keeping up will be a major challenge. It will also be critical for tax departments to integrate their AI strategy with other areas of the businesses, such as finance, ERM, ERP systems and so on. Tax processes must be integrated so that the AI is able to access and generate integrated data that reflects the entire value chain and risk profile/appetite of the organisation in a consistent and reliable way.

Leaders of tax functions should ideally be developing a point of view on how AI can be harnessed, whether through leveraging business-wide solutions or more bespoke tax-specific applications. As well as focusing on the opportunity, Tax Directors will need to develop a strong understanding of the risks associated with deploying GenAl in each potential use case within the business, and its "trustworthiness" especially from a quality and ethical perspective. Leaders should also be thinking about how they build a culture within their team that embraces new skills, new ways of working and new approaches to problem-solving.

Al-enabled solutions have been available to tax teams for some time, but the emergence of large language models underpinning user-friendly chat interfaces has led to an explosion of interest and investment in new solutions.







What is **Generative Al?**

GenAl is a subset of artificial intelligence that empowers machines to create novel content, encompassing text, code, voice, images, videos and processes.

The basic difference between traditional AI and GenAI is that the latter produces complex outputs based on all manner of input questions, hence the emphasis on the "creation" of content.

Though some forms of GenAl are already well established, it was the large language model (LLM) underpinning an accessible chat interface that triggered GenAl's watershed moment - notably the public launch of ChatGPT3.5 in late 2022.

In the context of business operations, GenAl can leverage data to create outputs and insights that can potentially serve the business and its stakeholders more efficiently and to a higher quality.



In the context of business operations, GenAl can leverage data to create outputs and insights that can potentially serve the business and its stakeholders more efficiently and to a higher quality.



What general capabilities of AI are most likely to be relevant in the short term?

Currently, the limiting factor is not so much the AI itself but the ability to safely and effectively harness its capabilities and embed it within a process or an organisation.

There are some obvious areas where GenAI will be impactful across a range of business functions without the need for too much tailoring.

These more generic applications could be adopted business-wide and will be fuelled by data that is readily available in high volume from reliable, permissible internal and external sources. With these applications, the focus should be on identifying how the tools can be applied to tax function activities.

This is easier said than done and will require a deep practical understanding of the function's activities (what it spends its time doing), the inherent limitations in the GenAl technology, as well as a strong overlay of creativity, vision and self-challenge. External perspective, whether from elsewhere in the business or an outside adviser, will really enrich this process through shared experience and constructive challenge of the status quo.



The focus should be on identifying how the tools can be applied to tax function activities.





Here are some of the applications of AI that are likely to fall into this category:



Textual content generation

Creating content based on input text examples, documents, data, or a specific theme. Examples include drafting memos, policies, advice and communications.



Classification

Assigning a category or label to a given input. This could include the classification of contracts and documents for due diligence purposes, or "tagging" data to create a smart, searchable knowledge management database.



Summarisation

Producing a concise summary of a long text or a collection of texts. Applications include summarising due diligence reports, structure reports, domestic and international tax developments, case law. This can be extended to summaries of content generated in other formats, such as the production of meeting notes by converting complex oral data into text form.

- •
- •
 - •
 - •
 - •



 \neg \vdash



Transformation

Converting content into a new type, format, or style. For example, using source data and narrative reporting from various parts of the business to create filing submissions that have to be prepared in a certain format (such as Country by Country Reporting).



Q&A

Providing a natural language answer to a natural language question based on a text or knowledge base (subject to some of the risks and limitations that we highlight later). This could be employed for tax self-service requests (e.g., in areas such as global mobility) or to answer repetitive questions that require a level of technical insight (e.g. the appropriate VAT or customs duty treatment of a supply).



Extraction

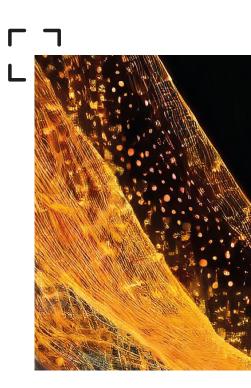
Deriving specific information or data points from a given input. A key use case would be tax due diligence.



Reasoning

Producing logical technical analysis given a context or knowledge base. Research on tax technical areas could be significantly enhanced.

More specific application to the tax function

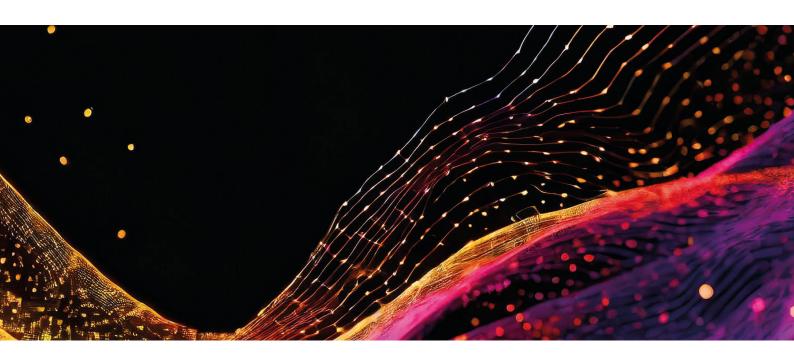


As well as identifying use cases for business-wide AI tools, tax departments undertake a variety of specific activities that lend themselves to AI enablement.

This may involve the creation or procurement of tax-specific Al, or the adaptation of tools that are deployed in other similar areas of the business (e.g., there may be compliance tools that can assist the legal, risk management and the tax departments, or which can be easily adapted between these areas).

Tax leaders should consider involving specialist process mappers to determine each task that the department undertakes, and its interdependencies with other systems and teams and, from that, potential AI use cases. For instance, Deloitte has designed a Digital Artifact Generation/Validation method to help innovation leaders determine whether an idea can be turned into a beneficial use case leveraging GenAl.

Deloitte has designed a Digital Artifact Generation/Validation method to help innovation leaders determine whether an idea can be turned into a beneficial use case leveraging GenAl.



At the core of this method are two of the most critical elements to consider: the human effort required to complete a task without GenAl, and the necessary effort to validate or fact check the output from the GenAl. This leads to a two-dimensional classification, categorising use cases based on the required human effort and the ability of the user to validate the results.

If a task requires effort to execute but is easy to validate, it has the hallmarks of a good use case. Here are some examples of tax processes that could, in future, be enhanced by GenAl:

- Statutory tax compliance: There is already significant scope for deployment of traditional AI on the preparation process and associated claims, and this will be an area of fast-paced development across all territories. GenAI will enable more targeted risk management, greater efficiency and it should be able to take an integrated approach to the global data to identify opportunities for improvement. Interface with tax authority systems has been a key theme in this area for many years and the addition of GenAI will bring some interesting opportunities and challenges for governments and taxpayers.
- HMRC enquiries (and equivalent processes in other territories): Analysis of case data (e.g., underlying documents, internal correspondence) for rapid early case assessment, thereby saving money and time associated with protracted correspondence, and potentially being used as a tool to mitigate penalties. Even better, a similar approach can be taken during the compliance process to inform the positions taken in submitted tax returns.

- HMRC Risk Reviews: Generation of data and narrative to support the risk review process. In due course we can expect the way that HMRC operates these reviews to include a high Al dependency, both in terms of how HMRC develops its own focus areas and in terms of the tools that are deployed on the taxpayer's data.
- M&A: Using LLMs to conduct more detailed due diligence across significantly larger data sets, increasing the accuracy of target valuations and required contractual protections.
- International structures: Gaining a real-time understanding of the evolving tax regimes across all relevant jurisdictions and the generation of preliminary risk assessment prompts (reducing the risks, costs and time periods associated with manual assessment of complex data across multiple jurisdictions).
- **Knowledge management**: Generation of updates and briefings for the business, generation of training materials, interrogation of knowledge assets (reduces risk and upskills teams, improves the employee proposition).
- Integrated systems: Tax will become more fundamentally embedded in the broader Enterprise Risk Management and Enterprise Resource Planning frameworks, as those frameworks also transform with GenAl.
- **Procurement**: Automated RFP (request for proposal) response scoring and assessment (eases the burden of a painful process and encourages objectivity), comparison of work outputs with contractual scope and fees.

• • •

Getting started

Tax Directors are questioning where to begin, especially in the context of tax-specific GenAl.

Being part of a community with other tax leaders and advisers addressing similar challenges will be a huge advantage in keeping up to date with trends and opportunities, but each business's needs will be different and so here are some key criteria to consider when evaluating GenAl use cases.



Demand

- Is the process understood in sufficient detail to identify activities where AI may play a part? Ideally each process should be atomised to its most granular level.
- Is the volume of work sufficient to justify a solution? As with any
 automation, there may be incremental gains each time an activity
 is performed but real economies of scale are achieved only when
 the activity is performed at a reasonably high frequency and/or
 over an extended period.
- Does the current work require enough human effort to warrant an Al solution? As Al is essentially replacing tasks that humans may otherwise perform, a greater return on investment will be achieved if the work that is being replaced currently requires a reasonable amount of human effort.
- Is there scope for AI to deliver higher quality (lower risk) outcomes, to reduce "key person" risk and bring other sustainability improvements?
- There are many other digital and tech-based solutions that are not GenAl enabled, and which may be a viable and cost-effective alternative to both human hours and GenAl. For example, taking advantage of off the shelf coding platforms to build digital and robotic solutions. The current focus on GenAl should catalyse a broader look at the available options. In practice, a process will be tackled with a combination of technologies of which GenAl will be one aspect.
- What cultural barriers exist within the organisation and/or the team and how can these be navigated?

• • •

Feasibility

- Does the work involve the creation of a digital artifact? GenAl is about generating content. For GenAl it is therefore key to focus on activities that involve such content generation (e.g., a tax compliance claim or a report summary). Without this feature, traditional Al and other digital solutions could deliver the desired outcomes.
- How much human effort is required to verify the accuracy of the Al result? GenAl is based on statistics and is not 100% accurate. The outputs from GenAl therefore require validation (see Risks below). If a disproportionate amount of effort is required to validate the outputs of the Al, it may make more sense to continue with the more traditional approach. Good GenAl use cases will be those where the output can be validated with lower amounts of human effort.
- Who possesses the subject matter expertise to feed into the
 model and validate its outputs? When considering use cases,
 thought needs to be given to who has the underlying subject
 matter expertise to create the most appropriate solution is it
 publicly available, can it be acquired/licensed, or does it need to
 be created? Is the organisation willing to invest in building the
 appropriate skills within the team in core areas such as "prompt
 engineering" (so that teams can extract the most effective
 outputs from GenAl) as well as in more complex areas?

- Should a generic solution be used, or a tax-specific one? If so, should an "off the shelf" solution be used, or should something bespoke be created?
- Have Al developments brought about new/better opportunities to outsource or to partner with an external provider?
- Data is GenAl's basic fuel. Is there sufficient volume of quality data accessible to the organisation and at what cost? What are the limitations and legal restrictions associated with leveraging and handling the outputs of the GenAl? Does the data need to be adapted before it can be leveraged in GenAl? Does the organisation have sufficient protection from the GenAl provider to allow it to process sensitive data through the tool? Where data is limited, are there opportunities to safely synthesise it or purchase synthesised data from third parties? If a data set is indeed supplemented by third-party information, has the organisation ensured that any GenAl processing fully complies with the limitations and terms of the license agreement? Organisations should be evolving their data strategies and data governance approaches to cater for the additional complexities of Al and tax teams should tap into that activity.
- How does the proposed GenAl deployment fit within the organisation's own governance and ethics framework?
 In particular, the maturity of the IT governance framework and infrastructure may enable or limit the scope to effectively deploy GenAl.

At this early stage of use of GenAI, we believe there will be a mix of approaches, depending on the specifics of each organisation and use case. Over time, as the use of the technology evolves, clear trends and common practices will emerge.

We should also recognise that GenAI is developing at such a pace that today's infeasible use case may easily become tomorrow's "big win" (e.g., through improvements in the technology, lowering of costs and changes in the organisation's own approach). Decisions therefore need to be kept under periodic review.



Over time, as the use of the technology evolves, clear trends and common practices will emerge.

Risks associated with Al



Many of these risks are common to GenAl in general but are particularly resonant in a tax context due to the fundamental requirement for absolute auditable accuracy. Here are a few of the main areas of consideration:



Data Strategy and Governance

Data Strategy and Governance is fundamental to effective and safe use of Al. What data is being used? What are the legal constraints (domestically and cross border)? What permissions are required? What assurances are there as to accuracy and origin? What data is being shared (e.g., with technology providers and other external parties, perhaps even with competitors, government and other stakeholders) and is this permissible/desirable/ understood? What data security measures are required? What local regulatory rules may apply? Is IP being protected, or might a service provider be using an organisation's IP to train its AI model for others' use? This is clearly a significant area of focus.



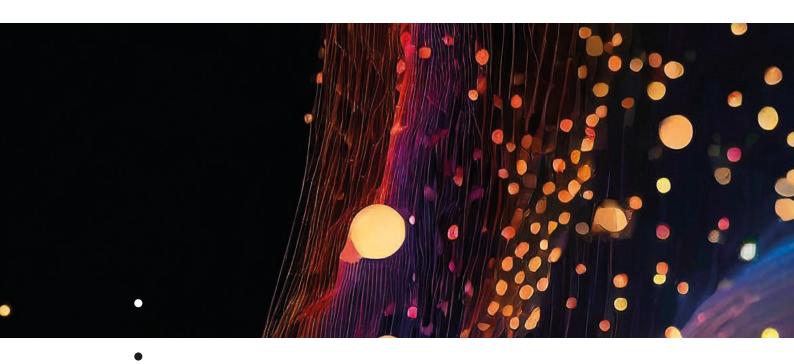
Ria

The data from which GenAl sources its conclusions may contain inherent bias, and bias may also be introduced in the course of developing GenAl solutions. Of course, whether GenAl leads to more or less bias than human-orientated systems is the real question, as the latter is also prone to bias.



Other ethical considerations

These are wide-ranging and depend to some extent on the organisation and the way that AI is to be used. Ethics may be concerned with local issues, such as whether AI creates a less even playing field in dealing with customers or staff. It is also concerned with societal issues, for example the role that social media companies may inadvertently play in the dissemination of AI generated fake news and its impact on democracy, the impact on jobs in certain sectors (which some organisations may conclude imposes a moral obligation to invest in re-training initiatives among other things). The International Monetary Fund, among other organisations, has raised concerns about whether Al could widen the gap between rich and poor nations, especially when it comes to the potential displacement of human labour. These concerns may be prompting some multinationals to reflect on the adoption and promotion of the technologies in these areas from a broader ESG perspective. Many organisations are already developing a "techno-ethics" framework as part of their governance activities and this can then be used to help guide decision-making.





As GenAl models operate on the basis of mathematical probabilities, there is a risk of generating incorrect outputs. While this risk diminishes with more domain-specific models, it is crucial to validate the model's output. Human intervention, coupled with other forms of technology, must be overlaid.



"Hallucinations" Return on Investment

Determining a cost benefit analysis could be complicated in terms of factoring in the tangible and intangible aspects, and the dynamics could change very quickly as the cost of emerging technologies is driven down.



Obsolescence

Will bespoke solutions become obsolete when large technology and service providers release industry standard tools? What are the maintenance costs?



Expertise

Do you have the appropriate experts to partner with you through this journey? What skills need to be developed or hired within teams to use GenAl as safely and effectively as possible?



Inactivity

Moving too slowly carries its own risks through lost opportunity, perception that the team is resistant to change and the potential demotivation of high performing team members who may relish the professional and developmental opportunity to embrace GenAl.

For deployments of AI in tax functions there are a number of specific risk areas to consider:

- If AI tools are used to research technical topics, what are the limitations of the source data and what areas of implicit bias may be present? For example, in the UK HMRC tends to litigate only when its confidence level is high. This means that the material that is in the public domain in relation to tax disputes (i.e. tax case law) will not necessarily be representative of the broad spectrum of tax disputes and outcomes that exist between taxpayer and HMRC on any given topic, because the majority of disputes whose facts and analysis led to HMRC or the taxpayer conceding earlier in the process are not in the public domain and so will not be accessible by open Al.
- Some areas of tax are relatively niche, and the quantity and quality of data and precedent in the public domain from which AI can draw may be too limited to generate reliable conclusions – it is important that this is evident to the user (e.g., if the AI can also provide a confidence measure). It may be possible to develop synthetic data with which to "teach" the GenAl within proprietary systems to help address this problem.
- Users will need an audit trail to support decisions taken in materials such as tax returns and HMRC correspondence. If GenAl is used to generate some of the underlying data, the ultimate sources will need to be apparent to the user, ideally alongside the weighting that the AI may have placed on different sources (noting, however, that GenAl can "hallucinate" its source references in the same way that it can create fictitious outputs). For instance, sources that are more frequent and which are based on legislation should carry higher weight than less frequent sources and those based on opinions expressed by commentators.

- It will need to be transparent to all stakeholders where AI has been deployed. If junior team members have used it to research topics, then the senior reviewer needs to be aware. Best practice for research documents is that they include sources and the implications for disclosure in tax authority correspondence.
- In the tax domain, tax authorities inevitably have the upper hand when it comes to access to a wider pool of underlying tax reporting data, not just domestically but cross-border. This is nothing new. However, if AI becomes highly influential in determining tax outcomes and if governments take the position that the data they hold is "better" than which any external organisation has at their disposal, this would fundamentally change the way that tax outcomes are determined, the way that uncertain tax positions are handled and the future role that the courts will play in determining the outcomes stemming from uncertain tax positions.
- We have yet to see how governments, individually and collectively, will seek to regulate AI.



Some areas of tax are relatively niche, and the quantity and quality of data and precedent in the public domain from which AI can draw may be too limited to generate reliable conclusions - it is important that this is evident to the user.



Trustworthy Al

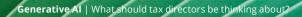
Trustworthy AI encompasses areas such as human agency and oversight, technical robustness and accuracy, privacy and data governance, transparency, fairness, non-discrimination and other ethical principles.

Applying these principles in the tax environment involves reflection in a number of areas before embracing a given GenAl use case, such as:

- Can it deliver a high impact and value add outcome without compromising on quality or introducing new risks?
- Do we fully understand the role that it is playing in a process? All tax decisions need to be supportable and auditable.
- There are significant unknowns about how GenAl really operates, behaves and will evolve. Rigorous testing, experimentation and monitoring are essential to mitigate the risks of unintended consequences.
- Do we understand every aspect of the "bargain" being made with a GenAl provider beyond the superficial contractual terms? For instance, by using the product the user might be training it and allowing proprietary data (whether specific or contextual fact patterns) to be used by the provider. In the area of tax knowledge management and systems, intellectual property ownership can be particularly difficult as solutions are normally a blend of client circumstance and the adviser's own experience and knowledge. If an adviser wishes to use its precedent work to train GenAl, untangling the adviser's own IP from the client's confidential facts and circumstances is not straight forward and will affect the quality of the outputs. Transparency is key, and while experimentation is one of the ways that GenAl improves, users need to understand to what extent they are participating in that experimentation process.

In summary, before embracing any GenAl use case, teams need to have fluency with the technology itself and the broader risks. Policies need to be developed so that team members know how to use it safely within the organisation's established risk framework.

In the area of tax knowledge management and systems, intellectual property ownership can be particularly difficult as solutions are normally a blend of client circumstance and the adviser's own experience and knowledge.



Developing a longer-term strategy to harness Al

As GenAl technology becomes more pervasive, the traditional boundaries between organisational activities will blur and new ones will form.

Currently many teams exist to provide a service to the wider business (often these are teams that are referred to as "in-house", as in in-house legal functions, in-house tax functions etc.).

Al will challenge this notion of service-provision as many of the activities will instead be founded on huge amounts of centrally curated internal and external data, with that data not merely informing but also driving the operations and decision making in new ways.

Developing the example of an organisation's internal tax team, traditional tax teams often exist as a defined reporting unit because organisations often organise teams around skills-based competencies.

With the greater automation and integration that AI will bring, combining human and machine capabilities in new ways, this model may no longer be appropriate and there may be a case for the tax function to become integrated across other organisational domains such as compliance, risk, governance, supply chain.

The emphasis on governance here may seem strange, because we normally view tax governance as a part of compliance and risk. However, as we have seen, AI will bring with it a need for significant focus on Data and Information Governance to manage the risks associated with AI, including in relation to tax.



Al will bring with it a need for significant focus on Data and Information Governance.

Conclusion

Some of the AI hype of the past 12 months may subside, but the potential of GenAI is huge and it will continue to develop at pace. The potential application to the tax profession is clear, but there is much work to be done to shape the vision and the possibilities into practical and sustainable solutions that can be embedded, and which deliver real value.

There is an old consulting saying: "How do you eat an elephant? One bite at a time!" and this is absolutely true of this topic.

For those with a more plant-based disposition, this quote from St Francis of Assisi may resonate more: "Start by doing what is necessary, then do what is possible; and suddenly you are doing the impossible".

At Deloitte, with our Deloitte Institute of AI, we have a huge focus in helping our clients seize the opportunities created by GenAI in tax, whether in specific use cases or broader strategic areas, always underpinned by our trustworthy AI principles. Our goal is collaborating with our clients, whether on our own or other technologies, to ensure value led – and values led – adoption.

Authors



Karen McNicholls
Partner
Business Tax
+44 (0)20 7007 2713
kmcnicholls@deloitte.co.uk



Frankie Jell
Partner
Tax Technology Consulting
+44 (0)20 7007 2230
fjell@deloitte.co.uk



Priya Vijayasarathy
Director
Digital Innovation
+44 (0)20 7303 5078
priyavijayasarathy@deloitte.co.uk

Further contacts



Hayley McKelvey Partner Digital Innovation+44 (0)20 7303 3940
hmckelvey@deloitte.co.uk



Howard Cooke
Partner
Digital Innovation
+44 (0)118 322 2315
hcooke@deloitte.co.uk



Deloitte.

This publication has been written in general terms and we recommend that you obtain professional advice before acting or refraining from action on any of the contents of this publication. Deloitte LLP accepts no liability for any loss occasioned to any person acting or refraining from action as a result of any material in this publication.

Deloitte LLP is a limited liability partnership registered in England and Wales with registered number OC303675 and its registered office at 1 New Street Square, London EC4A 3HQ, United Kingdom.

Deloitte LLP is the United Kingdom affiliate of Deloitte NSE LLP, a member firm of Deloitte Touche Tohmatsu Limited, a UK private company limited by guarantee ("DTTL"). DTTL and each of its member firms are legally separate and independent entities. DTTL and Deloitte NSE LLP do not provide services to clients. Please see www.deloitte.com/about to learn more about our global network of member firms.

© 2024 Deloitte LLP. All rights reserved.

Designed and produced by 368 at Deloitte. J31839