The AI conundrum: AI’s role in indirect tax

General Features | Indirect Tax



05 January 2026

AI offers powerful opportunities for indirect tax teams, but only when applied with clear controls and human judgement.

Key Points

What is the issue?

Indirect tax teams face pressure to adopt AI while ensuring accuracy, governance and accountability in a high-stakes environment where judgment and defensibility remain essential.

What does it mean to me?

AI can accelerate research, support processes and highlight anomalies, but it cannot replace human judgment. Tax professionals must design guardrails, ensure

auditability and integrate AI responsibly into existing controls.

What can I take away?

Use AI where it is actionable and auditable, automate only well-structured controls, and treat AI as a decision-support tool that enhances – rather than substitutes – professional expertise.

Every tax professional I speak to feels the same tension. On one hand, there's relentless pressure to 'do something with AI'. Boards are hearing about miraculous productivity gains and vendors are promising self-driving tax engines. We're being handed tools from IT with the expectation that we know when and how to use them.

On the other hand, indirect tax is not a playground. Paying tax correctly is part of your business's licence to operate and must be correct. We are dealing with high-volume, high-value transactions, rules that shift by jurisdiction and sector, and regimes that expect demonstrable governance. If things go wrong, it is the business and, in practice, the tax function that must stand behind the position taken.

And so, we have the AI conundrum for indirect tax: how do we take advantage of genuinely powerful new tools without outsourcing judgment or undermining accuracy and control?

What are we trying to use AI for?

When we think about the use cases for AI in tax, we often mix four very different needs, all of which represent different aspects of our roles:

- **Process and governance:** We're capturing, validating and routing data, documenting controls and generating audit trails.
- **Technical advice:** We're interpreting legislation, guidance and case law.
- **Tactical decisions:** We're defining the VAT treatment of particular flows or coding a new product.
- **Strategic decisions:** Finally, we're redesigning our supply chains, rebalancing our sourcing strategies and refining our operating model.

These needs sit on a spectrum between certainty and probability. Some tasks (like applying a clear VAT treatment where rules and facts are stable) are largely

deterministic. Others are inherently judgmental. We're weighing ambiguous facts, competing authorities, historic experience and commercial risk appetite.

The latest AI tools excel at pattern recognition and content generation. They can synthesise vast amounts of text and spot anomalies in large data sets far more quickly than a human. But they still struggle with context, risk appetite and accountability – precisely the things that dominate high-value indirect tax work.

Our use of AI needs to add real, incremental value. It needs to protect our businesses and our reputation.

The role of AI in technical advice

In the technical sphere, modern AI models are already very good at research and interpretation. They can search legislation, judgments and guidance, extract the relevant parts and summarise them in plain language. For the busy practitioner, this can be a genuine time-saver, especially when combined with a well-curated internal knowledge base.

However, that does not mean AI can replace the formation of a technical view. When we apply law to facts, we are not simply matching keywords. We're assessing the quality of evidence, interpreting grey areas and mapping positions to our individual business's cultures and appetite for risk. AI can outline possible arguments and suggest comparable precedents. But it does not understand commercial context, reputational risk or the reality of dealing with a particular tax authority. It cannot, in any meaningful sense, 'own' the judgment.

The same is true for scenario analysis. AI can model 'what if' situations (for example, different territorial footprints or supply chains) and estimate the VAT impact, provided the inputs are structured and of reasonable quality. Yet the moment a scenario involves policy trade-offs, stakeholder dynamics or regulatory uncertainty, you are back in human territory.

And when it comes to judgement and accountability, deciding what we think, documenting why and being prepared to stand behind that advice, AI is fundamentally weak. It has no skin in the game, no professional status to uphold and no mechanism for responsibly balancing risk and reward. Those belong, and will continue to belong, to us – the human professionals.

The message here is not that AI has no place in technical work. It does, as an accelerator and a challenger. We can be a naturally sceptical bunch, and I think that is right given where the technology is just now. We should challenge, question and hold ourselves to account on using AI tools responsibly.

The role of AI in process and governance

Where AI is more obviously ready for primetime is in the world of process and governance. Its ability to create quality first drafts of process and governance documentation, or even tax strategies, is impressive.

Indirect tax is awash with documents and structured data from invoices and purchase orders to customs entries and general ledger postings. Here, mature optical character recognition (OCR) and machine-learning models can capture and classify information with impressive accuracy, drastically reducing manual keying and re-work. Once the data is in the system, process automation can drive consistent treatment. Rules-based engines, enhanced with AI where appropriate, can apply VAT codes, route exceptions and trigger alerts when transactions fall outside expected patterns.

AI shines at decision support. Given a large volume of transactions, it can surface anomalies or trends that a human might never spot, such as suppliers whose coding behaviour changes over time, or jurisdictions where recovery rates systematically differ from expectations.

But again, there are caveats. Off-the-shelf AI systems struggle with explainability and governance. If a model flags an invoice as ‘high risk’ or suggests a particular treatment, can we see why? Can we show a tax authority the evidence considered, the thresholds used and the human-in-the-loop activity?

Without training these models for the needs of our roles, we may end up with faster processes but weaker control, a situation no Head of Tax wants to defend.

Two pillars: actionability and auditability

To turn AI from a clever toy into a defensible component of our control framework, I find it useful to think in terms of two pillars: actionability and auditability.

Actionability: Actionability is about whether AI outputs drive real, measurable decisions.

- Do they integrate with our existing workflows and systems?
- Do they trigger clear next steps, or do they just sit in dashboards?
- Have we defined how exceptions are handled and who has the authority to override?

Auditability: Auditability is about whether we can explain and defend those AI-assisted decisions.

- Do we log data inputs, model versions and rationale?
- Can we map AI-enabled steps to our existing control frameworks (for example, VAT governance, SOX compliance or internal tax policies)?
- Can we reconstruct the decision path for a specific transaction months or years later?

High-quality AI in tax must score well on both. Actionable but unauditable AI is a regulatory time-bomb. Auditable but unactionable AI is a glossy report that no one uses.

Turning auditability into practice

So, what does auditability look like in practice? To get there, tax and technology teams need to log every decision – not only the final VAT code, but also the key data points and logic applied.

They must design explainability checkpoints that require the system to produce a ‘because…’ statement before an outcome is accepted, and they should embed human review loops so that, at defined risk points such as new flows, high-value transactions or new jurisdictions, a human reviews and either accepts or declines the AI’s suggestion.

It is also important to align AI-enabled steps within existing control frameworks, mapping them to the control objectives that internal and external auditors already understand.

Finally, teams must store the reasoning alongside outcomes so that, when the tax authority asks, ‘Why was this treated this way?’, the organisation can offer a

substantiated explanation rather than simply saying ‘the system said so’.

Done well, this approach does not just make AI safer; it can actually enhance your ability to evidence good governance.

Turning actionability into practice

On the actionability side, practical steps include integrating AI outputs into core systems – such as enterprise resource planning (ERP) platforms, workflow tools or robotic process automation – so that suggested VAT codes, risk scores or anomaly flags appear where people already work.

Organisations should define clear exception paths that set out when humans must intervene, what options they have and how their decisions are recorded. They also need to set measurable thresholds that link AI insights to KPIs such as error-rate reduction, time saved or audit findings, enabling them to demonstrate incremental value.

Closing the feedback loop is equally important, feeding user actions – whether they accept, amend or reject a suggestion – back into the model so that it learns where its suggestions are helpful and where they are off the mark.

Finally, every AI-supported decision must still have a named human owner. The ‘system’ is never an appropriate control owner. If AI is not changing decisions, reducing effort or improving control, then it is, at best, an experiment – and at worst, it is an added layer of complexity without meaningful benefit.

A practical test…

A simple question I encourage tax teams to ask is: ‘Which control would we be comfortable automating tomorrow?’ First, is the control automation-ready?

- Are the inputs structured and consistent enough for a system to handle reliably?
- Are the decision criteria explicit and relatively stable over time?
- Would automation improve, or at least preserve, the audit trail?

If the answer is ‘no’ to these, you do not yet have a technology problem, you have a process and data problem. Only then should you ask whether the control is AI-ready.

- Does it require nuanced human judgment or contextual understanding?
- Do you have enough historic data (including examples of ‘good’ and ‘bad’ decisions) to train or calibrate a model?
- Crucially, could you still explain the ‘why’ behind outcomes to an auditor or authority?

Many high-volume, low-complexity VAT controls will pass both tests. These are your early candidates for AI-enabled automation and improvements. Higher-risk, judgment-heavy decisions will sit further out on the horizon, where they’re appropriate for decision support, but not yet decision replacement.

What does this mean for the tax profession?

For tax professionals, all of this has two important implications.

First, our core skills remain central. The market may be excited by AI’s ability to draft memos and summarise legislation, but the real value in indirect tax lies in structuring transactions, articulating defensible positions, making judgments and building governance that can withstand scrutiny. AI does not reduce the need for that work; if anything, it makes it more visible.

Second, we have a new design responsibility. As tax professionals, we can no longer be passive recipients of technology. We need to be involved in designing the decision flows, guardrails and audit trails that determine how AI is actually used. If we are not, others will make those choices for us, and we will still be the ones signing off the returns.

That means asking awkward questions of vendors, collaborating closely with finance and IT, and being precise about which parts of our work we are willing to automate and on what terms.

In conclusion

AI will not, in the foreseeable future, deliver a push-button, fully autonomous indirect tax function, but nor should we dismiss it as hype. Used thoughtfully, it can help us

process more data, spot more issues and explain our decisions more clearly than ever before. It can even take on the aspects of our roles we enjoy the least – now wouldn’t that be nice!

The conundrum is resolved when we stop asking whether AI can ‘do tax’ and start to ask some more grounded questions:

- What decisions are we trying to support?
- How will AI outputs become actionable in real workflows?
- Can we audit and explain those decisions months or years later?

If we can answer those questions with confidence, AI becomes not a threat to professional judgement but a tool that amplifies it, enabling indirect tax teams to be not just faster, but more controlled, more transparent and ultimately more valuable to the organisations we serve.

© Getty images