Fonoa

The Complete Guide to AI in Indirect Tax

Why transparency is the new compliance

Contents

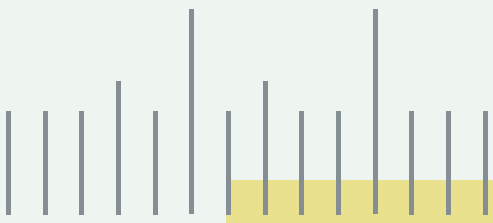

What AI means for tax

Tax leaders are standing at an inflection point. When we surveyed nearly 200 indirect tax professionals about AI in August 2025, the results revealed both opportunity and uncertainty:

How do tax leaders actually feel about AI?

Curious but cautious

Excited, this could change everything

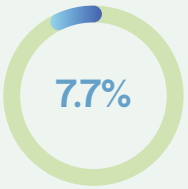

Nervous, not sure how to start

Skeptical, it's mostly hype

While a third are excited about AI's transformative potential, the majority (just over half) remain curious but cautious. It's a rational response to a fundamental shift in how tax compliance works.

Has your leadership team asked for an AI plan?

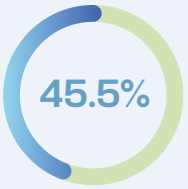

Yes, multiple times

Yes, once or twice

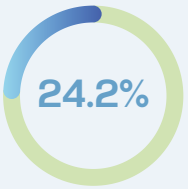

Not yet, but I'm preparing

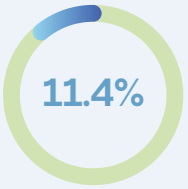

No, and not expecting it soon

Nearly two-thirds of tax leaders have already been asked about their AI strategy, with 45% fielding multiple requests. The message is unmistakable: the question isn't whether to adopt AI in tax, but how to do it responsibly.

The indirect tax function is at a crossroads.

For decades, compliance operated on sampling. Companies validated subsets of transactions. Tax authorities conducted spot-check audits.

That model is breaking down.

Real-time reporting and e-invoicing mandates now make all transaction data visible to authorities. When 100% of your data is scrutinized, manual review and rule-based systems can't keep pace.

Probabilistic AI becomes necessary to analyze millions of transactions and adapt to complexity. But it introduces uncertainty.

How can enterprises maintain accuracy, governance, and risk management in this shift?

In this guide, we explore that question, along with:

- The core differences between deterministic and probabilistic reasoning
- The implications for tax compliance accuracy
- Where AI adds the most value (and where it remains risky)
- How to manage risks and controls in an AI-driven tax function
- The policy and governance considerations
- How AI error rates compare to human error

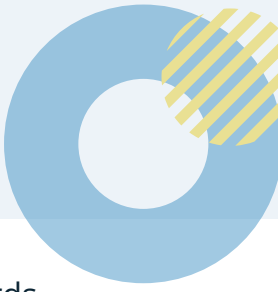

Drawing on academic research, industry case studies, and emerging standards, this guide offers a balanced global perspective to help tax leaders rethink accuracy, risk, and governance in a probabilistic world.

Whether you're excited about AI's potential or approaching it with caution, this guide will help you build a strategy that balances innovation with the rigor of tax compliance demands.

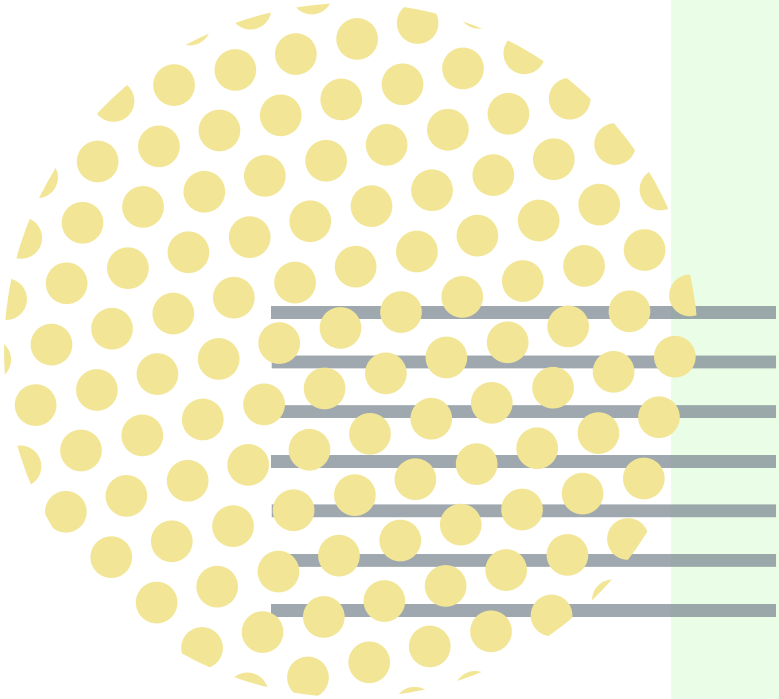

How AI introduces a new logic to tax decisions

Before you can decide where to deploy AI or how to control it, you need to understand a fundamental shift: AI makes tax decisions in a completely different way than the systems you're used to.

For decades, tax technology has operated on deterministic logic: fixed rules that always produce the same answer. Probabilistic AI introduces a new approach: learning from patterns and making predictions with varying degrees of confidence. The difference shapes everything from accuracy expectations to risk management.

Deterministic vs probabilistic reasoning overview

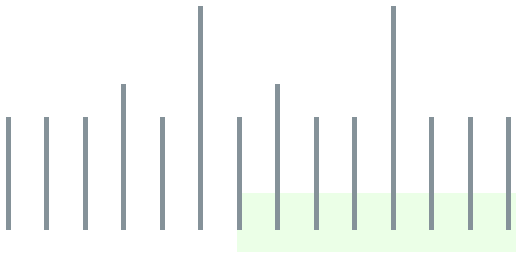

Consider how each of these logic systems compare.

	Deterministic systems	Probabilistic AI systems
What is it?	Operate on predefined rules and always produce the same output for a given input.	Learn from data and make inferences with degrees of confidence rather than hard-coded truth.
Strengths	Clear, consistent, seemingly more trustworthy, easy to audit.	Adaptability and breadth, automatically updates/improves, can consider a much wider array of inputs to infer tax outcomes, potentially catching edge cases.
Weaknesses	No streamlined way to handle exceptions or regulation changes, can be convoluted and risk omissions or conflicts, can't anticipate/adapt to new situations and rule changes.	Doesn't guarantee 100% consistency or accuracy, often "black boxes," meaning the reasoning for any single output can be hard to explain and audit.
Keep in mind	Gives the illusion of certainty: the output is certain given the rules, but if the rules are incomplete, that certainty can be misleading. The onus of adapting to new legislation is on the people monitoring the system, it will not evolve alone.	Only as good as its training data. If the data is biased or not representative, the AI's suggestions may systematically err (e.g. consistently under-reporting tax on certain transactions not prevalent in the training data).
TL;DR	Gives precise, repeatable answers but can't handle the unexpected.	Offers broad, adaptive analysis but introduces a margin of error and demands careful oversight.

How each logic applies to tax

Deterministic systems deliver precision but can't adapt easily

A deterministic system applies hard-coded logic to reach a tax decision: if goods category = X and ship-to country = Y, then apply tax rate Z. The strength is predictability. A transaction processed at 2 a.m. produces the same result as one at 2 p.m. Every decision traces to a specific rule, which auditors value for transparency.

The drawback is brittleness. Tax laws change frequently, and each new exception ("unless the customer is a charity" or "unless the goods are promotional samples") makes the rule base more complex. When faced with an unfamiliar case, a deterministic system either misclassifies or falls back to the most basic rule.

All long-standing tax engines work this way, offering reliability but requiring constant updates. The certainty they provide can be misleading if the rules are incomplete. Crucially, if the tax engine makes it hard to introduce change, then each change becomes a battle.

Probabilistic AI adapts and scales but isn't flawless

Probabilistic AI learns patterns from historical data and applies them to new inputs. Instead of a single fixed outcome, it assigns probabilities: "76% likely to be zero-rated" or "90% likely to require VAT." This allows it to capture subtle patterns in free-form invoice descriptions, contract clauses, or historical transaction behavior.

The key advantage is adaptability. As new data becomes available (new tax rulings, emerging transaction types), the model updates without human intervention. This can detect patterns and edge cases that deterministic systems miss.

The weaknesses lie in potential lack of consistency, explainability, and data quality. AI may produce different outputs for the same input. Poor training data or the wrong guidance can lead to systematic errors. In some cases, the model might produce confidently wrong outputs. In a compliance setting, these risks require careful oversight.

Measuring compliance in an AI world

The most important question about AI in tax: is it accurate? The real answer requires reframing the question itself.

Different systems fail in different ways

Accuracy in tax isn't simply about being right or wrong on each decision. Humans, deterministic systems, and AI all make mistakes, but in different ways:

- A person who works as a tax analyst might misread a law, miss a detail due to fatigue, or mistype a code.
- A rules-based system errs by omission, defaulting to a basic answer when it encounters something outside its programmed scope.
- AI will give an answer for every case, often catching more issues overall, but with some errors that can be either subtle misinterpretations or confident but wrong outputs.

Breadth can outweigh per–item perfection

AI's strength is reach. It examines all transactions, filings, or regulatory updates, whereas humans and fixed systems work from limited samples.

An AI reviewing expense claims might flag 5% as anomalous (some false positives), but catches every genuinely problematic case. A human auditor checking only 10% of claims might be 100% accurate on those checked, but misses issues entirely.

The business decision: is 93% accuracy at scale better than 97% accuracy on a small subset? Often, yes.

Confidence thresholds make AI + human workflows manageable

An AI might classify 1,000 invoices at 90% accuracy. Tax staff review the uncertain or low-confidence 10%. The end result approaches 100% accuracy, achieved far faster than manual review alone.

Leading practices use confidence scoring: auto-process cases above 90% confidence, route the rest for human judgment. This preserves speed while controlling risk.

System–level compliance matters more than transaction–level perfection

The key is judging compliance at the system level, not transaction by transaction. If AI misclassifies a few cases but prevents significant underpayments elsewhere, total compliance improves.

Evidence from auditing shows AI enables “whole population testing,” examining 100% of data rather than samples. Research shows 24-90% error rates in complex spreadsheets used for tax calculations. Studies show AI-assisted audits lead to fewer errors and restatements, while human-only processes have hidden error rates most teams never measure.

Early tax case studies show AI catching VAT underreporting that traditional monitoring missed, saving penalties despite occasional false alarms. The goal: better decisions at scale, not perfection in isolation.

Use cases: Where AI is valuable vs. risky

Not every tax process is equally suited to AI. Large enterprises need to map where deterministic precision is critical versus where AI's strengths can be safely leveraged.

When expected outcome of using AI is:	Consider it this type of use case:
Low reward, some/high risk	Determinism-critical (e.g., Binary compliance steps)
Medium reward, some risk	Hybrid/Layered (e.g., Tax research)
High reward, low risk	AI-Advantaged (e.g., Anomaly detection)

In the next pages, we dive deeper into each of these use cases to help you consider examples most relevant to your business.

The most effective model keeps deterministic controls for bright-line compliance while deploying AI where scale, pattern recognition, or speed outweighs the cost of occasional errors.

Ultimately, the goal should be to avoid over-reliance on AI while not missing any opportunities.

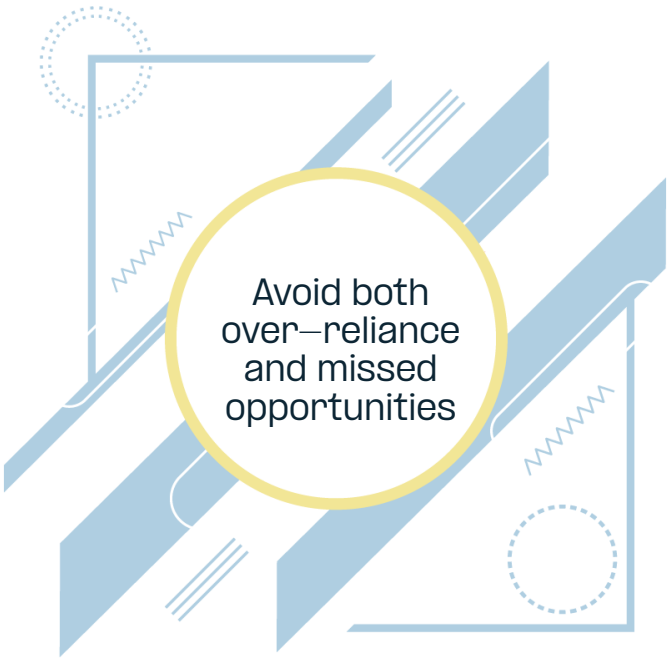

Type 1: Determinism–critical use cases

These are areas where errors carry high risk and the rules are clear.

When to use AI here: AI can assist by pre-checking rates or suggesting updates, but deterministic logic makes the final decision.

Rule of thumb: If a missed flag could trigger a major fine, rules stay in charge.

Core examples:

- **E-invoicing and transaction reporting:** Real-time reporting as required by mandates.
- **Statutory filing formats:** Electronic VAT returns must match the exact government schema. No "almost correct" exists.
- **Binary compliance steps:** Regulatory deadlines and minimum legal rates have no room for probabilistic answers.

Type 2: Hybrid use cases – layered AI plus deterministic checks

Some processes combine clear rules with open-ended analysis.

The strategy: "Layered AI" - let AI expand coverage and handle repetitive work, with deterministic checks or human review for final calls.

Core examples:

- **Transaction-level tax determination:** Leveraging AI to build and develop taxation rules, speeding up build of tax logic. Human-in-the-loop testing before going into production.
- **Regulatory monitoring:** AI scans global sources for tax law changes and summarizes them. Humans determine applicability.
- **Audit preparation:** AI flags anomalies in transaction data. Tax managers verify against rules before auditors review.
- **Tax controversy support:** AI sifts large document volumes. Humans apply legal judgment for responses.

Type 3: AI–advantaged use cases

High-volume, complex, or under-served areas where some uncertainty is acceptable versus the status quo.

When to use AI here: AI can assist by pre-checking rates or suggesting updates, but deterministic logic makes the final decision.

Core examples:

- **Anomaly detection:** Scan millions of invoices for unusual patterns. One solution processed 5-6 million invoices yearly (300x human capacity).
- **Fraud detection:** Flag suspicious vendor networks or transaction chains that rules can't capture.
- **Document processing:** Extract invoice data at 90% accuracy. Staff correct the remaining 10% (far faster than 100% manual entry).
- **Generative drafting:** Create first versions of technical memos that humans refine.

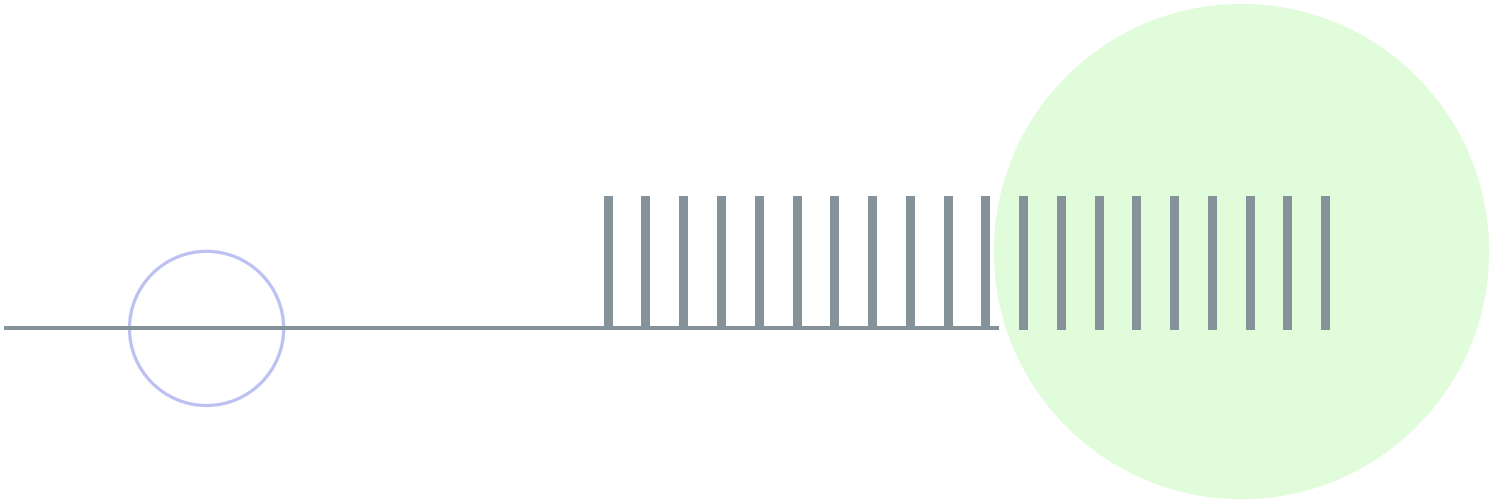

Controls for AI–driven tax operations

Adopting AI in indirect tax doesn't reduce the need for strong controls. It increases it.

In deterministic systems, governance focused on managing rule changes, testing, and access permissions. In the AI era, governance must also cover model performance, bias, explainability, and accountability. The goal: build enough confidence in outputs so stakeholders trust the system without expecting perfection.

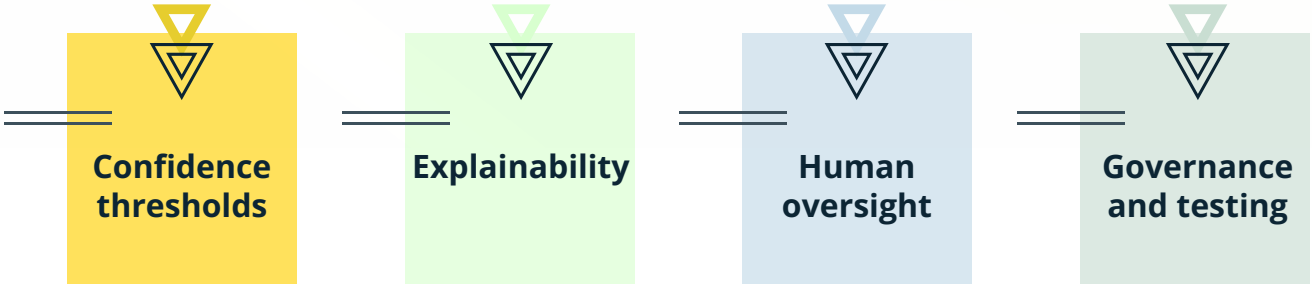

Fonoa

Confidence scores translate uncertainty into actionable workflows

Well-designed AI gives a confidence level for each output: “95% confident this transaction is exempt.”

Set thresholds for automatic posting versus human review, then fine-tune them by comparing AI predictions to actual outcomes. For instance, only autopost tax decisions above 90% confidence. Route everything else for review.

This turns AI’s probabilistic nature into practical workflows and reminds users that not all outputs carry equal weight.

Explainability and audit trails satisfy external scrutiny

Auditors, regulators, and compliance teams will ask: why did the AI make this decision?

Explainable AI techniques can highlight contributing factors, such as an unusual combination of customer location and product code. Every recommendation should be logged with its input data and model version.

Shadow mode deployment: Many companies start by logging AI recommendations while still following the old process. This builds trust by comparing and validating outputs before full deployment.

Strong AI governance can actually create more transparent audit trails than undocumented human decisions. Each AI decision gets timestamped, linked to source data, and backed by clear rationale.

Human oversight prevents automation from becoming autopilot

AI works best as decision support, not as an unchecked decision-maker.

Build in review points where human experts approve outputs. That might mean a tax manager signing off on AI-prepared returns or auditors investigating flagged transactions. Humans are essential for borderline cases and for interpreting AI findings in context.

Where to put humans in the loop: Base checkpoints on transaction materiality, decision type, volume requirements, and real-time constraints. Feedback from these reviews also helps retrain models and improve accuracy.

Bias testing and governance frameworks prevent systematic errors.

Even in tax, AI can reflect biases from training data. It might under-scrutinize certain regions because they were under-scrutinized historically. Regular testing catches systematic skew.

Security and privacy controls are critical since tax data is sensitive. Frameworks from NIST and OECD are emerging to guide trustworthy AI. Finance has managed algorithmic systems (trading, credit scoring) for years using similar controls.

AI governance essentials

- Model performance monitoring
- Bias and fairness testing
- Data security and privacy controls
- Version control and rollback capability
- Incident response procedures
- Stakeholder communication protocols

Implementation requires trust but verify at every stage

Risk management in AI is about maintaining oversight over something less predictable than fixed code. It's doable.

The mantra: trust but verify. Use AI, but check outputs through strong controls. Be ready to show any stakeholder (CFO, external auditors, tax authorities) that your AI is under control, just like any other critical process.

Learn more

Learn more about governance and AI controls in our recent blog series.

Read now

Building AI into your tax framework

The paradigm shift from deterministic to probabilistic systems in indirect tax is underway. Leadership teams are asking for AI strategies. The question has moved from “whether” to “how.”

What we’ve covered:

Tax systems can’t deliver both perfect certainty and infinite adaptability. The future belongs to hybrid approaches that leverage both rule-based precision and AI’s pattern recognition.

Accuracy must be measured differently. Analyzing 100% of transactions at 93% accuracy catches more total errors than analyzing 10% at 97%. System-level compliance matters more than transaction-level perfection.

Not all processes suit AI equally. Transaction-level tax determination demands deterministic precision. Anomaly detection across millions of invoices needs AI. Most processes benefit from layered approaches where AI expands coverage and humans make final calls.

Controls become more critical, not less. Confidence thresholds, explainability, human oversight, and governance frameworks turn AI from a risk into an advantage.

The opportunity

Shift from reactive to proactive compliance. Focus tax teams on strategy while AI handles volume and pattern recognition. Build more resilient, efficient, compliant operations.

The risk

Deploying AI without understanding where it adds value versus where it introduces unacceptable risk. Or rejecting AI entirely and falling behind.

Your AI implementation checklist

- [ ] Educate your team on deterministic vs. probabilistic thinking Use analogies like paper map vs. GPS. Combat expectations of perfection; focus on improving overall compliance.
- [ ] Map your tax processes on the AI suitability spectrum Classify each as determinism-critical, hybrid, or AI-advantaged. Identify where bright-line rules must govern versus where AI can expand coverage.
- [ ] Start with low-risk, high-value use cases Anomaly detection, document processing, or regulatory monitoring. Early successes build momentum.
- [ ] Establish confidence thresholds and human review workflows Define what confidence level (e.g., 95%) triggers automatic acceptance versus human review. Document who reviews what and when.
- [ ] Build your AI governance framework now Assign clear accountability for AI outputs. Define policies for data training, bias testing, validation, and incident response. Map to frameworks like NIST AI RMF.
- [ ] Maintain deterministic guardrails for critical processes For transaction tax and statutory filings, keep rule-based controls around AI. Use AI in suggestion mode initially with hard limits on outputs.
- [ ] Invest in explainability and audit trails Log every recommendation with input data and model version. Require AI outputs include source references. Consider shadow mode deployment initially.
- [ ] Track system-level metrics Monitor: error rates pre/post-AI, issue detection rates, false positives, post-filing adjustments, audit findings. Stop obsessing over per-transaction perfection.
- [ ] Brief external auditors on your AI use and controls Stay current on regulatory developments. Join industry forums on tax and AI.
- [ ] Measure what matters: Is your tax function making fewer mistakes and catching more issues with AI in the mix?

For AI to thrive in tax operations, it needs structured, mature tax technology to build on. Without accurate data, reliable integrations, and robust controls, AI becomes a risk rather than an advantage.

Fonoa provides the tax technology foundation that makes AI adoption safe and effective. Our platform delivers the structured data, deterministic controls, and global coverage that turn AI from an uncertainty into a strategic capability.

Find out how Fonoa can set you up to thrive in the AI era.

Talk to the team

The platform to power your tax automation

The first global tax automation platform that lets you validate tax IDs, calculate taxes, generate compliant e-invoices, report transactions in real time, and file tax returns—all in a single solution.

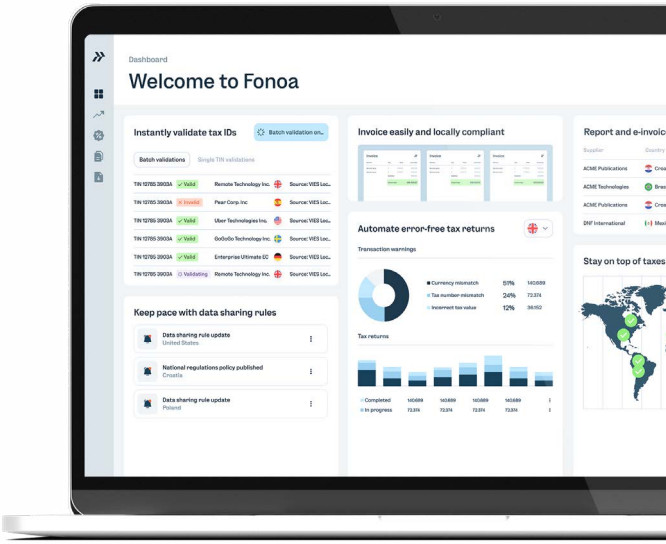

Lookup

Instantly validate tax IDs in 116+ countries to ensure you're compliant with local tax regulations.

Tax

Determine the right sales tax, VAT and GST treatment of your transactions globally with a single solution.

Returns

Automatically prepare, review and submit error-free, VAT/GST and sales tax returns everywhere from one single platform.

E-Invoicing

Report transactions and generate tax compliant e-invoices and credit-notes in real-time across countries with one standardized solution.

Data Sharing

Simplify global reporting compliance—from DAC7 in the EU to requirements in the UK, Canada, Australia, and more. We manage the entire process securely and seamlessly, from data preparation to submission.

Trusted by the world's leading digital companies

Uber teach:able zoom SQUARESPACE

Bolt Lime BIRD remote INTUIT

Learn how Fonoa automates your tax and compliance workflows for greater control and scale

Book a demo