
Advanced Certification in AI in Tax Law (France)

AI and Tax Dispute Resolution

Artificial Intelligence (AI) refers to the set of computational techniques that enable machines to perform tasks that would normally require human intelligence. In the context of tax law, AI is employed to analyse massive volumes of fiscal data, identify patterns of non-compliance, and support decision-making in dispute resolution. Understanding AI terminology is essential for tax professionals who must evaluate the reliability, legality and practical impact of these technologies.

Machine Learning (ML) is a subset of AI that focuses on algorithms that improve automatically through experience. Instead of being explicitly programmed for each task, ML models learn from historical data to make predictions or classifications. In tax dispute resolution, supervised ML models can predict the likelihood that a tax audit will result in a liability, while unsupervised models can cluster similar cases to reveal systemic issues.

Deep Learning (DL) is a branch of ML that uses artificial neural networks with many layers—hence “deep.” These networks excel at processing unstructured data such as text, images or audio. For example, a deep-learning model can read the full text of a tax notice and extract relevant facts, dates and monetary amounts, dramatically reducing manual review time.

Neural Network is the computational architecture that underlies deep learning. It consists of interconnected nodes (neurons) organized in layers: an input layer, one or more hidden layers, and an output layer. Each connection carries a weight that is adjusted during training to minimise prediction error. In tax analytics, a neural network might be trained to recognise patterns of invoice manipulation that signal potential fraud.

Supervised Learning is a type of ML where the algorithm is trained on a labelled dataset—each example includes both input features and the correct output (the label). In tax contexts, a supervised model could be trained on past audit outcomes (e.g., “no adjustment,” “partial adjustment,” “full adjustment”) to predict future audit results for new cases.

Unsupervised Learning does not rely on labelled outcomes. Instead, the algorithm discovers hidden structures in the data. Clustering techniques such as K-means or hierarchical clustering are common unsupervised methods. Tax authorities might use clustering to group taxpayers with similar risk profiles, allowing targeted interventions.

Reinforcement Learning involves an agent that learns to make a sequence of decisions by receiving rewards or penalties. Though less common in tax practice, reinforcement learning could be used to optimise the scheduling of audit resources, rewarding the system for prioritising high-risk dossiers that lead to successful recoveries.

Natural Language Processing (NLP) is the field of AI that enables computers to understand, interpret and generate human language. NLP tools can parse tax legislation, extract statutory definitions, and compare

them with taxpayer submissions. For dispute resolution, NLP can be employed to summarise lengthy arguments, flag contradictory statements, or suggest precedent-based counter-arguments.

Computer Vision refers to AI techniques that interpret visual information from images or videos. In tax investigations, computer-vision algorithms can analyse scanned receipts, invoices or property photographs to verify authenticity, detect alterations, or estimate asset values.

Training Data is the set of examples used to teach a machine-learning model. High-quality training data must be accurate, representative and free from bias. In French tax law, training data may comprise historical audit files, tax returns, and court judgments, anonymised to protect personal data.

Validation Set is a subset of data used to tune model hyper-parameters and prevent over-fitting. Over-fitting occurs when a model learns the noise in the training data rather than the underlying pattern, resulting in poor performance on new cases. Regular validation ensures that AI predictions remain robust when applied to fresh dispute scenarios.

Test Set is a separate dataset reserved for final performance evaluation. It provides an unbiased estimate of how the model will behave in real-world applications, such as predicting the outcome of a tax appeal before the case reaches the Conseil d'État.

Feature Engineering is the process of selecting, transforming and creating variables (features) that improve model performance. In tax analytics, features might include the taxpayer's historical compliance rate, the proportion of deductible expenses, the presence of offshore entities, or the textual similarity between a declaration and the statutory definition of a deductible expense.

Algorithmic Bias occurs when a model systematically favours or disfavors certain groups, often due to imbalanced training data. In the French tax context, bias could manifest as higher false-positive rates for small businesses if the training data over-represents large corporate audits. Recognising and mitigating bias is a legal and ethical imperative.

Explainability (or interpretability) describes the degree to which a model's internal logic can be understood by humans. Black-box models, such as deep neural networks, often lack transparency, raising concerns about the admissibility of AI-generated evidence in tax disputes. Techniques like SHAP values or LIME provide local explanations that highlight which features most influenced a particular prediction.

Model Governance encompasses the policies, procedures and controls governing the development, deployment and monitoring of AI models. For tax authorities, model governance must address data protection, auditability, version control, and compliance with French regulations such as the Loi Informatique et Libertés and the GDPR.

Model Drift refers to the gradual degradation of a model's accuracy as the underlying data distribution changes over time. In tax dispute resolution, legislative reforms, economic shocks or shifts in taxpayer behaviour can cause drift, necessitating periodic retraining and validation of the AI system.

Data Governance is the overall management of the availability, usability, integrity and security of data used

in AI applications. It includes data lineage (tracking the origin and transformations of data), data quality assurance, and compliance with confidentiality obligations imposed by the French tax code and the EU's GDPR.

Big Data denotes data sets that are too large or complex for traditional processing tools. Tax administrations increasingly rely on big-data platforms to store and analyse millions of electronic tax returns, social-media signals, and cross-border transaction records. AI techniques such as distributed learning enable the extraction of insights from these massive repositories.

Structured Data is information organised in a predefined format, such as relational tables with columns for "taxpayer ID," "tax year," "declared revenue." Structured data is readily ingested by most ML algorithms and forms the backbone of risk-scoring engines used by the French Directorate General of Public Finances (DGFIP).

Unstructured Data includes free-text documents, scanned PDFs, email correspondence and multimedia files. NLP and computer-vision methods are required to transform unstructured data into a structured form that can be analysed alongside traditional fiscal variables.

Data Anonymisation is the process of removing or encrypting personally identifiable information to protect privacy while retaining analytical value. French tax data must be anonymised before being used for model training, in compliance with the GDPR's principle of data minimisation.

Tax Assessment is the formal determination by the tax authority of the amount of tax due. AI can assist assessors by automatically calculating taxable income, applying the appropriate tax rates, and flagging discrepancies between the declared amount and the computed liability.

Tax Audit (or fiscal inspection) is an in-depth examination of a taxpayer's records to verify compliance. AI-driven audit selection models rank taxpayers according to risk scores, allowing auditors to focus limited resources on cases with the highest probability of non-compliance.

Tax Notice is the written communication issued by the tax authority indicating a proposed adjustment, penalty or collection action. AI tools can draft tax notices by populating template fields with case-specific data, ensuring consistency and reducing manual drafting errors.

Tax Litigation encompasses the procedural steps taken when a taxpayer disputes an assessment. In France, litigation proceeds through the administrative courts (Tribunal administratif, Cour administrative d'appel) and may culminate before the Conseil d'État. AI can support litigation by predicting the likely outcome at each stage, based on historical rulings.

Burden of Proof in French tax disputes traditionally rests with the taxpayer, who must demonstrate that the declared amounts are correct. However, recent jurisprudence has shifted certain evidentiary burdens to the administration, especially in cases involving complex transfer-pricing arrangements. AI-generated analytics can provide the administration with robust evidentiary support, but must respect procedural fairness.

Standard of Proof in French administrative tax law is "probability" (probabilité sérieuse) rather than "beyond

reasonable doubt.” Predictive models that output a probability score can be calibrated to align with this standard, offering a quantitative basis for administrative decisions.

Evidence includes all documents, data extracts and expert opinions presented to support a claim. AI-derived evidence, such as an anomaly-detection report, must be admissible under French procedural rules. This often requires a clear audit trail, documentation of the model’s methodology, and assurance that the underlying data were reliable.

Tax Credit is an amount subtracted directly from the tax liability, whereas a deduction reduces the taxable base. AI can automatically calculate eligible credits (e.g., research-tax credit) by matching expenses against statutory criteria, thereby reducing the risk of misapplication.

Tax Deduction reduces the taxable income. AI systems can verify that claimed deductions meet the legal definition (e.g., “frais professionnels”) by analysing supporting invoices and cross-checking them with the French Code général des impôts (CGI).

Tax Evasion is the illegal act of deliberately misrepresenting taxable income to reduce tax liability. AI-driven fraud detection systems flag suspicious patterns—such as unusually low declared revenues relative to industry averages—prompting deeper investigations.

Tax Avoidance involves arranging affairs within the legal framework to minimise tax liability. While not illegal, aggressive avoidance may be challenged under the “abuse of law” principle (abus de droit). AI can identify structures that closely resemble prohibited arrangements, assisting tax advisors in assessing compliance risk.

Transfer Pricing refers to the pricing of transactions between related entities across borders. French tax law requires documentation that conforms to the OECD guidelines. AI can automate the benchmarking process, comparing intercompany prices against comparable uncontrolled transactions extracted from global databases.

Risk Scoring is the assignment of a numerical value that reflects the likelihood of non-compliance or dispute. The scoring model integrates multiple features—financial ratios, historical behaviour, textual analysis of filings—and can be visualised in dashboards for auditors and legal counsel.

Predictive Analytics encompasses statistical techniques that forecast future events based on historical data. In tax dispute resolution, predictive analytics can estimate the probability of a successful appeal, the expected monetary recovery, and the time needed to reach settlement.

Anomaly Detection identifies observations that deviate markedly from the norm. Techniques such as isolation forests or autoencoders are applied to transaction streams to uncover hidden irregularities—e.g., sudden spikes in deductible expenses that lack supporting documentation.

Decision Support System (DSS) is a software platform that aggregates data, analytical models and visualisations to aid human decision-makers. A tax-DSS may combine risk scores, legal precedents, and financial impact analyses to guide auditors in selecting the most appropriate resolution pathway.

Expert System is an AI program that emulates the reasoning of a human expert by applying a set of rules to a knowledge base. Early tax-compliance tools were rule-based expert systems that encoded the CGI articles. Modern expert systems integrate ML components while preserving a traceable rule set for legal compliance.

Ontology in the AI context is a formal representation of concepts and relationships within a domain. A tax ontology defines entities such as "taxpayer," "tax base," "deduction," and links them to legislative articles. Ontologies enable semantic search and reasoning, facilitating the retrieval of relevant statutes during dispute preparation.

Semantic Analysis examines the meaning of text beyond keyword matching. By mapping the language of a tax notice to the underlying legal concepts, semantic analysis helps ensure that AI-generated summaries preserve the nuance required for accurate legal argumentation.

Case Law Database stores judicial decisions, including the reasoning of administrative courts. AI-enhanced legal research tools can perform similarity searches, clustering cases by factual pattern, and surfacing the most persuasive precedents for a given dispute.

Legal Reasoning involves applying statutory provisions to specific facts, interpreting ambiguous language, and balancing competing principles. While AI cannot replicate the full depth of judicial reasoning, it can assist by suggesting applicable articles, highlighting relevant factual parallels, and estimating the weight of contradictory arguments.

Settlement is the resolution of a tax dispute without proceeding to full litigation, often through negotiation of a reduced liability or payment plan. AI can propose settlement offers by modelling the expected outcome of a court decision and the financial impact on the taxpayer, thereby facilitating mutually acceptable agreements.

Mediation and Arbitration are alternative dispute-resolution mechanisms recognised under French law. AI-driven platforms can schedule sessions, generate briefing documents, and even provide neutral, algorithm-based recommendations for dispute resolution, subject to parties' consent.

Procedural Timeline outlines the sequential steps in a tax dispute, from initial notice to final appeal. AI tools can map a case's progress against the statutory deadlines, issuing alerts when filing periods are approaching, thus preventing procedural defaults.

Evidence Chain of Custody documents the handling of evidence from collection to presentation in court. For AI-generated evidence, the chain must include details of data extraction, preprocessing, model selection, and parameter settings to satisfy evidentiary standards.

Data Privacy is a cornerstone of EU law, especially under the GDPR. Any AI system processing personal tax data must implement lawful bases (e.g., legal obligation, public interest), minimise data exposure, and provide mechanisms for data subjects to exercise their rights (access, rectification, erasure).

Confidentiality Obligation under French tax law requires tax authorities to protect taxpayer information. AI platforms must enforce strict access controls, encryption at rest and in transit, and audit logs that record

who accessed which data and when.

Regulatory Compliance for AI in tax disputes involves adhering not only to fiscal statutes but also to sector-specific AI regulations, such as the EU's proposed AI Act. The Act categorises AI systems based on risk; a tax-risk-assessment tool would likely be deemed "high-risk," subject to conformity assessments, transparency obligations and post-market monitoring.

Transparency Requirement mandates that AI-driven decisions be explainable to affected parties. In practice, this means providing the taxpayer with a clear statement of the factors that influenced the AI's recommendation, the confidence level of the prediction, and the possibility to contest the result.

Accountability obliges the entity deploying the AI system to accept responsibility for its outputs. For a French tax administration, accountability entails maintaining documentation, performing regular audits of model performance, and establishing a governance board that reviews ethical and legal implications.

Ethical AI principles include fairness, non-discrimination, respect for human autonomy, and avoidance of harmful outcomes. Tax professionals must assess whether AI models inadvertently penalise vulnerable groups, such as micro-enterprises, and implement safeguards—e.g., threshold adjustments or human-in-the-loop reviews.

Human-in-the-Loop (HITL) is a design paradigm where AI suggestions are reviewed and approved by a qualified professional before final action. In tax dispute resolution, a HITL approach ensures that AI-generated settlement offers are vetted by senior tax counsel, preserving legal oversight.

Model Audit is a systematic review of an AI model's design, data, performance metrics and compliance with regulatory standards. A model audit report typically includes a risk-assessment matrix, bias analysis, validation results, and recommendations for remediation.

Audit Trail records all actions performed on a dataset or model, including data imports, feature transformations, model training runs, and prediction outputs. Maintaining a comprehensive audit trail is essential for both internal governance and external scrutiny by supervisory authorities.

Continuous Learning describes AI systems that update their parameters automatically as new data become available. While this can improve relevance, it also raises concerns about model drift and the need for ongoing validation—especially when legislative changes alter the legal landscape.

Model Versioning tracks successive iterations of an AI model, enabling rollback to a prior state if a new version exhibits undesirable behaviour. Version control tools (e.g., Git) are employed to manage code, hyper-parameters and training datasets, ensuring reproducibility.

Data Quality encompasses accuracy, completeness, consistency, timeliness and relevance. Low-quality data—such as outdated address records or missing invoice numbers—can degrade model performance and lead to erroneous dispute outcomes. Data-quality dashboards help monitor key indicators and trigger remediation workflows.

Data Lineage maps the flow of data from source systems (e.g., SAP, Cegid) through transformation pipelines to the final analytical dataset. Understanding lineage is crucial when a taxpayer challenges the provenance of AI-derived evidence, as it demonstrates that the data were derived from official tax filings.

Cross-Border Taxation involves rules governing transactions between French entities and foreign jurisdictions. AI can assist in identifying potential double-taxation issues by analysing treaty databases, matching transaction dates with treaty applicability, and suggesting appropriate relief mechanisms.

OECD BEPS Action Plan (Base Erosion and Profit Shifting) sets out measures to prevent tax avoidance through profit shifting. AI tools can automate the detection of BEPS-related risks, such as intra-group financing arrangements that lack economic substance, by scanning corporate structures against the OECD guidelines.

Digital Services Tax (DST) is a French levy on revenues generated by digital platforms. AI can monitor online revenue streams, calculate the DST liability, and flag inconsistencies between reported figures and third-party data (e.g., Google Analytics).

Fiscal Year is the twelve-month period used for tax reporting. In France, the fiscal year generally coincides with the calendar year, but companies may adopt a different accounting period subject to approval. AI systems must align data aggregation with the correct fiscal year to avoid miscalculations.

Tax Base is the amount on which the tax rate is applied. Accurate determination of the tax base requires applying numerous deductions, exemptions and adjustments. AI can automate the sequential application of these rules, reducing manual errors and ensuring compliance with the CGI.

Taxable Income is the portion of the tax base that remains after all allowable deductions and exemptions. Machine-learning models can predict taxable income based on historical filing patterns, assisting auditors in spotting abnormal deviations that merit further examination.

Tax Liability is the final amount due after applying rates, credits and penalties. AI can generate liability forecasts under different scenarios (e.g., with or without a contested deduction), facilitating strategic decision-making for both the taxpayer and the administration.

Penalty is a monetary sanction imposed for non-compliance, such as late filing or underpayment. AI can estimate penalty amounts by referencing statutory formulas, and can also predict the likelihood of penalty mitigation based on past leniency trends.

Interest accrues on unpaid tax amounts. AI-driven calculators incorporate the French tax-interest rate schedule, automatically updating for rate changes announced by the Ministry of Economy and Finance.

Tax Appeal (recours) is the procedural right to challenge an administrative decision before a higher authority. AI can assist by generating a structured appeal dossier, aligning arguments with relevant statutes, and suggesting supporting jurisprudence.

Administrative Court (Tribunal administratif) adjudicates disputes between taxpayers and public authorities.

AI tools can monitor case filings, deadlines, and procedural requirements, ensuring that the taxpayer's rights are exercised within the statutory timeframes.

Conseil d'État is the highest administrative court in France. AI-assisted legal research can retrieve previous Conseil d'État rulings that influence the interpretation of tax statutes, providing strategic insight for high-value appeals.

Legal Precedent (jurisprudence) guides the application of statutes. AI can rank precedents by relevance, using semantic similarity metrics, and present the most persuasive cases to counsel preparing a dispute brief.

Statutory Interpretation involves determining the meaning of legislative language. AI-enhanced tools can parse the wording of the CGI articles, identify definitions, and highlight cross-references, supporting a rigorous interpretative analysis.

Fiscal Policy shapes the overall tax structure, influencing the environment in which disputes arise. AI can model the impact of policy changes—such as a new tax credit—on compliance rates and dispute volumes, assisting policymakers in assessing reform outcomes.

Compliance Monitoring is the continuous oversight of taxpayer behaviour to ensure adherence to obligations. AI-driven dashboards provide real-time risk indicators, enabling proactive outreach before formal audits are initiated.

Risk Management in tax departments involves identifying, assessing and mitigating potential financial exposures. AI contributes by quantifying exposure probabilities, simulating worst-case scenarios, and recommending mitigation actions (e.g., targeted education campaigns).

Data Protection Impact Assessment (DPIA) is a GDPR-mandated analysis of processing activities that pose high risks to individuals' rights. Implementing an AI system that analyses personal tax data requires a DPIA that outlines the purpose, data categories, safeguards, and retention periods.

Consent Management governs the collection of explicit permission from data subjects. While tax authorities generally rely on legal obligations rather than consent, AI platforms that incorporate third-party data (e.g., credit-bureau information) must manage consent appropriately.

Encryption protects data confidentiality both at rest and in transit. End-to-end encryption mechanisms safeguard sensitive tax records when they are processed by AI services, especially when cloud-based solutions are employed.

Access Control restricts data visibility to authorised personnel. Role-based access ensures that only auditors, legal analysts or senior managers can view specific AI outputs, preserving the principle of least privilege.

Transparency Report documents the operation of AI systems, outlining data sources, model architecture, performance metrics and governance measures. Publishing a transparency report enhances public trust and demonstrates compliance with emerging AI regulations.

Algorithmic Auditing involves independent evaluation of AI systems for fairness, accuracy and legal conformity. External auditors may assess whether the tax-risk-scoring algorithm respects the principle of proportionality, a key requirement under French administrative law.

Proportionality requires that administrative measures be suitable, necessary and balanced against the interest pursued. AI-generated penalties must be calibrated so that the severity of the sanction is commensurate with the identified breach.

Legal Reasonableness is a standard that courts apply when reviewing administrative decisions. AI-based recommendations must be anchored in a rational legal analysis; otherwise, they risk being overturned for lacking reasonable justification.

Procedural Fairness guarantees that taxpayers receive an impartial hearing, the opportunity to present evidence, and a reasoned decision. AI tools that automate decision-making must incorporate safeguards—such as human review and the right to a written explanation—to uphold fairness.

Case Management System (CMS) tracks the lifecycle of disputes, from receipt of a tax notice to final resolution. Integrating AI modules into a CMS enables automated case assignment, priority setting, and predictive outcome tracking.

Workflow Automation streamlines repetitive tasks, such as generating standard letters, updating case status, or notifying parties of upcoming deadlines. AI can trigger workflow actions based on rule-based conditions (e.g., “if risk score > 0.8, assign to senior auditor”).

Knowledge Base stores domain expertise, including tax statutes, procedural rules, and internal guidelines. AI-driven chatbots can query the knowledge base to answer routine taxpayer inquiries, freeing human staff for more complex dispute work.

Chatbot leverages NLP to conduct conversational interactions. In the French tax context, a chatbot can guide a taxpayer through the steps required to file a corrective declaration, clarify the meaning of a particular article, or schedule an appointment with a tax officer.

Virtual Assistant extends chatbot capabilities by integrating with calendar systems, document repositories, and case-management platforms. A virtual assistant can retrieve a taxpayer’s filing history, suggest relevant documents for an upcoming audit, and draft a preliminary response letter.

Sentiment Analysis evaluates the tone of textual communication, detecting frustration, confusion or compliance intent. Analyzing correspondence between taxpayers and the administration can help identify cases where a more conciliatory approach may prevent escalation.

Document Classification automatically assigns categories (e.g., “financial statement,” “contract,” “receipt”) to uploaded files. Accurate classification expedites the extraction of key data fields, which are then fed into downstream risk-assessment models.

Optical Character Recognition (OCR) converts scanned images of paper documents into machine-readable

text. OCR combined with NLP enables the extraction of numeric values, dates and entity names from legacy tax filings, facilitating their inclusion in AI analyses.

Entity Resolution reconciles multiple records that refer to the same taxpayer or related party, despite variations in naming or identifier formats. Robust entity resolution prevents double-counting of risk indicators and ensures that AI models operate on a unified view of the tax ecosystem.

Graph Analytics examines relationships among entities (e.g., shareholders, subsidiaries, related-party contracts) using graph structures. By visualising these networks, AI can uncover hidden pathways of profit shifting or undisclosed ownership, key concerns in complex tax disputes.

RegTech (Regulatory Technology) encompasses solutions that help organisations comply with regulations. AI-enabled RegTech platforms for tax dispute resolution automate compliance checks, generate audit evidence, and monitor regulatory updates in real time.

FinTech innovations, such as blockchain-based invoicing or digital payment platforms, introduce new data sources for tax authorities. AI can ingest transaction data from FinTech services, applying anti-money-laundering (AML) filters and tax-compliance checks simultaneously.

Blockchain provides an immutable ledger of transactions. When integrated with tax reporting, blockchain can assure the authenticity of invoice data, reducing opportunities for fraud. AI can verify blockchain-recorded transactions against declared amounts, flagging mismatches for audit.

Smart Contract is a self-executing contract with terms encoded on a blockchain. Tax authorities may require that certain obligations (e.g., VAT remittance) be embedded in smart contracts, and AI can monitor compliance by analysing contract execution logs.

Data Lake stores raw, uncurated data from diverse sources, enabling flexible analysis. A tax authority's data lake may contain XML tax returns, PDF audit reports, email correspondence, and sensor data from customs checkpoints. AI pipelines extract, transform and load (ETL) relevant subsets for model training.

Data Warehouse contains structured, cleaned data optimised for reporting and analytics. After ETL, tax-related data are loaded into a warehouse where AI models can efficiently query historical audit outcomes, compute aggregate risk metrics, and generate management dashboards.

ETL Process (Extract, Transform, Load) is the backbone of data preparation. Extraction pulls data from source systems; transformation applies cleansing, standardisation, and enrichment; loading deposits the final dataset into the analytics environment. Each stage must be documented to satisfy audit requirements.

Feature Importance quantifies the contribution of each input variable to a model's predictions. In a tax-risk model, feature importance charts reveal that "percentage of foreign-source income" and "frequency of amendment filings" are the strongest drivers of audit risk, informing policy adjustments.

Model Calibration aligns predicted probabilities with observed outcomes. A well-calibrated tax-risk model ensures that a 70% risk score truly corresponds to a 70% chance of audit adjustment, enhancing trust

among auditors and litigants.

Cross-Validation splits data into multiple folds to assess model robustness. By rotating training and validation sets, cross-validation mitigates the risk of over-fitting, providing a more reliable estimate of how the AI will perform on unseen tax cases.

Confusion Matrix summarises classification performance, displaying true positives, false positives, true negatives and false negatives. In tax dispute prediction, a high false-positive rate (incorrectly flagging compliant taxpayers) could lead to unnecessary audits and reputational damage.

Precision measures the proportion of positive predictions that are correct. High precision in a tax-risk model indicates that most flagged cases truly warrant investigation, reducing wasted resources.

Recall (sensitivity) gauges the ability to capture all relevant cases. High recall ensures that most non-compliant taxpayers are identified, but may increase false positives if not balanced with precision.

F1 Score harmonises precision and recall into a single metric, useful for evaluating trade-offs in tax-risk models where both false positives and false negatives carry significant costs.

ROC Curve (Receiver Operating Characteristic) plots true-positive rate against false-positive rate across thresholds. The area under the ROC curve (AUC) provides a summary measure of discriminative ability; an AUC close to 1.0 signals excellent separation between compliant and non-compliant cases.

Threshold Optimization selects the decision boundary that best balances business objectives (e.g., audit cost versus revenue recovery). In tax dispute resolution, the threshold may be set higher for high-value cases to prioritise resource allocation.

Data Enrichment supplements core tax data with external information—such as industry benchmarks, credit-rating scores, or market-price indices. Enriched datasets improve model accuracy by providing contextual cues that pure fiscal data may lack.

API Integration enables AI services to communicate with existing tax-administration systems. Secure APIs allow the risk-scoring engine to receive real-time filing data, return predictions to the case-management interface, and log interactions for audit purposes.

Micro-service Architecture decomposes the AI platform into independent, loosely coupled components (e.g., data ingestion, model inference, reporting). This design facilitates scalability, allowing the tax authority to handle spikes in filing volume during peak periods.

Scalability refers to the capacity of an AI system to maintain performance as data volume or user load grows. Cloud-based deployments, containerisation, and auto-scaling groups ensure that tax-dispute-resolution tools remain responsive during high-traffic filing seasons.

Latency measures the time between a request (e.g., “predict audit risk for taxpayer X”) and the system’s response. Low latency is critical for interactive decision-support tools used by auditors in the field, where rapid feedback influences investigative direction.

Batch Processing handles large volumes of data in scheduled jobs (e.g., nightly risk-score calculations). While not real-time, batch processing is efficient for periodic tasks such as updating the risk-assessment database before the annual audit cycle.

Real-Time Processing provides immediate analysis as data arrive. For high-value transactions (e.g., large cross-border payments), real-time AI can flag suspicious activity instantly, enabling prompt intervention by tax officers.

Service Level Agreement (SLA) defines performance expectations for AI services, including uptime, response time and support. Clear SLAs ensure that tax-dispute-resolution teams can rely on AI tools without unexpected downtime that could jeopardise case timelines.

Incident Management outlines procedures for handling system failures, data breaches or model anomalies. A well-documented incident-response plan mitigates operational risk and demonstrates compliance with regulatory expectations for continuity.

Change Management governs the introduction of new AI capabilities, ensuring that stakeholders are trained, processes are updated, and risks are assessed. In the French tax administration, change-management protocols must align with the internal audit framework and the Ministry's digital-transformation roadmap.

Stakeholder Engagement involves communicating AI initiatives to internal users (auditors, legal counsel) and external parties (taxpayers, industry groups). Transparent outreach builds confidence, clarifies the role of AI in dispute resolution, and gathers feedback for iterative improvements.

Legal Hold preserves relevant electronic evidence when litigation is anticipated. AI-generated documents, model outputs and underlying data must be placed under legal hold to prevent alteration or deletion, safeguarding their admissibility in court.

Data Retention Policy dictates how long tax-related data, including AI-derived artifacts, are stored. French law imposes specific retention periods (e.g., ten years for fiscal documents). AI systems must enforce automatic archiving or deletion in accordance with these mandates.

Data Minimisation limits the collection of personal data to what is strictly necessary for the intended purpose. When designing an AI model for tax dispute resolution, practitioners must justify each data field, discarding extraneous attributes that could increase privacy risk.

Ethical Review Board evaluates AI projects for compliance with ethical standards, societal impact, and alignment with public interest. In the context of tax disputes, the board may assess whether automated risk scoring could disproportionately affect certain economic sectors.

Bias Mitigation Techniques include re-sampling, re-weighting, and algorithmic adjustments to correct imbalanced representations. For example, if the training set contains an over-representation of large corporations, techniques such as SMOTE (Synthetic Minority Over-Sampling Technique) can balance the dataset.

Fairness Metrics (e.g., demographic parity, equalised odds) quantify how evenly a model treats different groups. Applying these metrics to a tax-risk model helps ensure that small-enterprise owners are not unfairly targeted compared to multinational corporations.

Explainable AI (XAI) frameworks provide user-friendly explanations of model decisions. Visual tools such as partial dependence plots illustrate how changes in a specific feature (e.g., “percentage of deductible expenses”) affect the predicted audit probability, aiding auditors in interpreting results.

Regulatory Sandbox offers a controlled environment where innovative AI solutions can be trialled under regulatory oversight. The French tax authority may host a sandbox to test novel dispute-resolution algorithms before full deployment, allowing iterative refinement based on supervisory feedback.

Data Sovereignty concerns the location of data storage and processing. French law and EU regulations often require that personal tax data remain within the European Economic Area. AI providers must guarantee that cloud infrastructure complies with data-localisation requirements.

Third-Party Vendor Management involves vetting external AI service providers for security, compliance and performance. Contracts should include clauses on data protection, audit rights, liability, and termination, ensuring that the tax authority retains control over sensitive fiscal information.

Incident Reporting obliges organisations to notify supervisory authorities of data breaches