Advanced Certification in AI in Tax Law (France)

AI Auditing and Risk Assessment

Artificial Intelligence Auditing is the systematic examination of AI-driven processes to verify that they operate as intended, comply with applicable regulations, and align with organisational risk tolerance. In the context of French tax law, auditors must assess whether an AI system that, for example, automatically classifies taxable transactions respects the principles of legality, proportionality, and non-discrimination set out in the French Tax Code and the GDPR. An audit typically begins with a review of the system's documentation, proceeds to technical testing of model outputs, and concludes with a risk-based recommendation. The auditor's role is to provide an independent assurance that the AI does not introduce hidden liabilities or expose the tax administration to sanctions.

A closely related concept is Risk Assessment, which involves identifying, analysing, and prioritising potential adverse outcomes associated with AI-enabled tax processes. Risks may stem from data quality, model bias, regulatory non-compliance, or operational failures. The assessment is usually expressed in terms of likelihood and impact, allowing the tax authority to allocate resources to the most critical control points. For instance, a risk assessment of an AI tool that predicts VAT fraud might reveal a high impact if false negatives lead to revenue loss, and a moderate likelihood given the model's historical performance.

Model Governance refers to the set of policies, procedures, and organisational structures that oversee the lifecycle of AI models. Effective governance ensures that model development, deployment, monitoring, and retirement are performed under clear accountability. In French tax law, model governance must incorporate statutory duties such as the duty of care (obligation de moyen) and the duty of result (obligation de résultat). A governance framework typically includes a model inventory, version control, change-management procedures, and designated roles such as model owner, data steward, and compliance officer.

Explainability (or interpretability) is the ability to articulate how an AI system arrives at a specific decision. Explainability is crucial for tax auditors because they must justify the outcomes of automated assessments to taxpayers and supervisory bodies. Techniques such as SHAP values, LIME explanations, or rule extraction can be employed to surface the most influential features in a model that flagged a corporate entity for potential transfer-pricing adjustments. However, explainability must be balanced against model complexity; highly accurate deep-learning models may be less transparent than simpler logistic-regression models.

Transparency extends beyond explainability by requiring that the entire AI pipeline, from data ingestion to model deployment, is open to scrutiny. Transparency obligations are reinforced by the French Data Protection Authority (CNIL) and the EU's Digital Services Act, which demand that users be informed about automated decision-making that significantly affects them. In practice, a tax administration might publish a high-level description of the AI system used for automated tax notice generation, including the categories of data processed, the logic of the decision-making, and the recourse mechanisms available to taxpayers.

Bias denotes systematic errors that cause an AI model's predictions to favour or disadvantage particular groups. In tax administration, bias can manifest as disproportionate audit rates for certain industries, regions, or company sizes. Detecting bias involves statistical tests such as disparate impact analysis or fairness metrics like equalized odds. For example, an AI system trained on historical audit data may inherit the historical over-audit of small retailers in the Île-de-France region, leading to a self-reinforcing cycle unless corrective measures are applied.

Fairness is the normative principle that seeks to mitigate bias and ensure equitable treatment across taxpayers. French tax law embodies fairness through the principle of horizontal equity, which requires that taxpayers in similar situations be taxed similarly. When evaluating AI-driven tax tools, auditors must verify that the fairness criteria embedded in the model align with legal standards. This may involve adjusting the model's loss function to penalise unfair outcomes or implementing post-processing techniques that re-balance predictions across protected attributes.

Data Quality encompasses the accuracy, completeness, timeliness, and consistency of the data used to train and operate AI models. Poor data quality can lead to erroneous tax assessments, potentially exposing the administration to legal challenges. In the French context, data sources may include the Système d'Imposition des Entreprises (SIE), customs declarations, and electronic invoicing (Factur-X) feeds. Auditors must verify that data pipelines incorporate validation rules, deduplication processes, and error-handling mechanisms to maintain high data integrity.

Data Provenance tracks the origin, lineage, and transformation history of data elements throughout the AI lifecycle. Provenance is essential for both compliance and auditability, as it enables investigators to reconstruct the exact dataset that generated a particular model output. For tax AI systems, provenance records should capture the source system (e.G., The French tax portal), extraction date, any anonymisation steps, and the schema version. Maintaining detailed provenance logs supports the CNIL's requirement for traceability and aids in resolving disputes over automated tax decisions.

Model Validation is the process of assessing whether a model meets its intended performance criteria before deployment. Validation involves techniques such as cross-validation, hold-out testing, and out-of-sample evaluation. In tax applications, validation must also consider regulatory compliance: A model that predicts eligibility for tax credits must be validated against the statutory definition of eligibility, not merely against statistical metrics. Validation reports should document the test data, performance results, and any identified limitations.

Performance Metrics are quantitative measures used to evaluate a model's effectiveness. Common metrics include accuracy, precision, recall, F1-score, and area under the ROC curve. For tax risk models, precision (the proportion of flagged cases that are truly non-compliant) and recall (the proportion of actual non-compliant cases that are correctly flagged) are particularly relevant. Auditors should assess whether the chosen metrics reflect the tax authority's risk appetite and operational constraints, such as the capacity of audit teams.

Overfitting occurs when a model captures noise in the training data rather than the underlying pattern, resulting in poor generalisation to new data. In tax fraud detection, an overfitted model might achieve high

accuracy on historical audit records but fail to detect novel fraud schemes. Auditors must examine evidence of regularisation techniques, such as dropout or L2 penalties, and review validation results on temporally separated data to ensure the model can adapt to evolving taxpayer behaviour.

Underfitting is the converse problem where a model is too simplistic to capture relevant relationships, leading to consistently low performance. An underfitted model for VAT compliance could miss critical indicators such as sudden spikes in invoicing volume. Auditors should verify that the model complexity is appropriate for the problem domain, that feature engineering has been adequately performed, and that the model has been trained on a representative sample of tax data.

Compliance in the AI auditing context refers to adherence to legal, regulatory, and internal policy requirements. For French tax AI systems, compliance obligations include the GDPR, the French Tax Code, the CNIL's recommendations on algorithmic decision-making, and sector-specific statutes such as the anti-tax-avoidance law (Loi de lutte contre la fraude fiscale). Auditors must map each compliance requirement to concrete controls within the AI system, such as data minimisation, consent management, and documentation of the decision-making logic.

General Data Protection Regulation (GDPR) imposes stringent rules on the processing of personal data, including the use of automated decision-making. Article 22 of the GDPR grants data subjects the right to obtain human intervention, express their point of view, and contest decisions made solely by algorithms. In the tax domain, personal data may include taxpayer identifiers, income details, and transaction histories. Auditors must confirm that AI systems provide a lawful basis for processing, implement appropriate safeguards, and maintain records of the decisions made for each data subject.

French Tax Code (Code général des impôts) provides the statutory framework governing taxation in France. AI models that calculate taxable income, determine eligibility for deductions, or assess penalties must be aligned with the code's provisions. Auditors need to verify that the model's logic correctly interprets legal clauses, handles exceptions, and respects the hierarchy of norms (e.G., Constitutional principles versus administrative decrees). Any deviation may constitute a legal error, exposing the administration to judicial review.

Fiscal Data refers to the structured information collected by tax authorities, such as income statements, balance sheets, VAT declarations, and customs records. Fiscal data is often highly sensitive, requiring robust security measures and strict access controls. Auditors should evaluate whether the AI pipeline encrypts data at rest and in transit, enforces role-based access, and logs all data handling activities to satisfy both security standards and the CNIL's guidelines on data protection.

Automated Decision-Making (ADM) denotes processes where AI systems generate outcomes without human intervention. In tax law, ADM can be used for issuing tax notices, granting tax credits, or selecting audit targets. While ADM can increase efficiency, it also raises concerns about accountability and recourse. Auditors must ensure that ADM systems incorporate a "human-in-the-loop" or "human-on-the-loop" mechanism where tax officials can review and override AI decisions when necessary.

Model Drift describes the phenomenon where a model's performance degrades over time due to changes

in underlying data distributions. For tax AI, drift may arise from legislative reforms, economic shifts, or evolving taxpayer behaviours. Auditors should verify that the organization has instituted continuous monitoring, periodic retraining, and threshold alerts to detect drift early. A practical example is monitoring the false-positive rate of an AI system that flags large cash transactions; a sudden increase may indicate that the model no longer reflects the current regulatory environment.

Monitoring encompasses the ongoing observation of AI system behaviour in production. Effective monitoring tracks performance metrics, data quality indicators, and compliance flags. In a French tax context, monitoring dashboards might display the number of automated tax notices generated per day, the proportion of those contested by taxpayers, and the latency between data ingestion and decision output. Auditors assess whether monitoring is sufficiently granular to identify anomalies and whether escalation procedures are defined.

Audit Trail is a chronological record of all actions taken by the AI system and its supporting infrastructure. The audit trail should capture data uploads, model training runs, parameter changes, inference requests, and any manual overrides performed by tax officials. Maintaining a comprehensive audit trail satisfies the CNIL's requirement for accountability and enables forensic analysis in case of disputes. Auditors verify that the trail is immutable, time-stamped, and stored in a secure, tamper-evident repository.

Documentation is the collection of artefacts that describe the AI system's design, development, deployment, and maintenance. Documentation should include a model card, data sheet, risk register, and governance charter. For tax AI, documentation must also reference the specific articles of the French Tax Code that the model operationalises. Auditors use the documentation as a primary source to evaluate compliance, assess risk controls, and verify that the system's intended purpose matches its actual use.

Stakeholder refers to any individual or entity with an interest in the AI system's outcomes. In tax AI, stakeholders include taxpayers, tax auditors, policy makers, the CNIL, and the public at large. Engaging stakeholders early helps identify concerns such as perceived unfairness, privacy expectations, and operational constraints. Auditors should review evidence of stakeholder consultations, such as minutes of working groups or public comment periods, to ensure that the AI system reflects a balanced perspective.

Accountability is the principle that an organisation must be answerable for the decisions made by its AI systems. Accountability mechanisms include clear role definitions, documented decision pathways, and the ability to provide evidence of compliance. In the French tax administration, accountability is reinforced by the principle of responsibility (principe de responsabilité), which obliges public entities to justify their actions. Auditors assess whether accountability structures are embedded in the AI governance model and whether they are effectively enforced.

Ethical AI embodies a set of values that guide the responsible development and deployment of AI, encompassing respect for human dignity, fairness, transparency, and sustainability. While ethical AI is not a legal requirement per se, it aligns with the French Republic's constitutional commitments to liberty, equality, and fraternity. Auditors may evaluate ethical considerations by reviewing the organisation's AI ethics guidelines, the inclusion of ethical impact assessments, and the presence of mitigation strategies for identified ethical risks.

Algorithmic Accountability is a specific facet of accountability that focuses on the technical and procedural aspects of algorithmic decision-making. It requires that the logic of the algorithm be documented, that its performance be regularly evaluated, and that mechanisms exist for remediation when errors occur. In tax law, algorithmic accountability ensures that taxpayers can challenge automated assessments on substantive grounds, not merely procedural ones. Auditors should verify that the system provides actionable explanations and that remediation pathways are clearly defined.

Regulatory Impact Assessment (RIA) is a systematic analysis of the potential effects of new regulations, including those that mandate the use of AI in tax administration. An RIA may examine how AI-driven compliance checks affect taxpayer behaviour, administrative workload, and data protection obligations. Auditors can use RIA findings to calibrate risk thresholds, design appropriate controls, and anticipate legal challenges. For example, an RIA on the introduction of AI-based automatic VAT reconciliation might reveal a need for additional safeguards to protect small businesses from erroneous penalties.

Risk Register is a living document that records identified risks, their severity, mitigation actions, and status. In AI auditing, the risk register should capture technical risks (e.G., Model bias), regulatory risks (e.G., GDPR violations), and operational risks (e.G., System downtime). Auditors review the risk register to ensure that all material risks have been identified, that mitigation measures are proportionate, and that risk owners are assigned. A well-maintained risk register supports the tax authority's overall risk management framework.

Control Framework defines the set of policies, procedures, and tools used to manage and mitigate identified risks. Common control categories include preventive controls (e.G., Access restrictions), detective controls (e.G., Anomaly detection), and corrective controls (e.G., Model retraining). In tax AI, a control framework might mandate that any model update undergoes an independent validation review before deployment. Auditors examine whether the control framework is documented, communicated, and consistently applied across the AI lifecycle.

Change Management governs how modifications to AI models, data pipelines, or supporting infrastructure are introduced. Effective change management reduces the likelihood of unintended side effects, such as a regression in model accuracy after a data schema update. In the French tax environment, change management processes must be aligned with the public sector's procurement and IT governance rules, including the need for formal approvals and impact analyses. Auditors assess whether change requests are properly documented, risk-assessed, and tested before implementation.

Version Control tracks successive iterations of code, models, and configuration files. Version control enables reproducibility, facilitates rollback in case of failure, and supports auditability. For tax AI systems, version control should be applied not only to software artefacts but also to trained model artefacts, including hyper-parameters and training data snapshots. Auditors verify that the versioning system records metadata such as the model's training date, data version, and responsible developer.

Data Minimisation is a GDPR principle that requires organisations to collect and process only the data necessary for a specific purpose. In tax AI, data minimisation may involve stripping personally identifiable information (PII) from datasets used for training fraud-detection models, retaining only aggregated or pseudonymised features. Auditors check whether data pipelines implement minimisation techniques, such

as hashing taxpayer identifiers, and whether any residual PII is justified by a documented legal basis.

Consent Management handles the acquisition, recording, and revocation of consent for data processing activities. While tax authorities generally have a legal mandate to process taxpayer data without explicit consent, certain AI applications (e.G., Optional predictive tax advice services) may rely on consent. Auditors should confirm that consent mechanisms are transparent, that consent records are securely stored, and that the system respects withdrawal of consent where applicable.

Data Anonymisation transforms personal data into a form that cannot be linked back to an individual, thereby reducing privacy risks. Techniques include aggregation, perturbation, and differential privacy. In tax AI, anonymisation may be applied to historical audit data used for model training, ensuring that individual taxpayer identities remain protected. Auditors evaluate the robustness of anonymisation methods and verify that re-identification risk assessments have been performed.

Differential Privacy is a mathematical framework that provides strong privacy guarantees by adding calibrated noise to data queries. When applied to tax datasets, differential privacy enables the creation of aggregate statistics for model training while limiting the exposure of any single taxpayer's information. Auditors assess whether differential privacy parameters (e.G., Epsilon values) are set appropriately and whether the resulting model performance remains acceptable for the intended tax purpose.

Algorithmic Transparency is the broader notion that stakeholders should be able to understand the functioning of AI systems, even if the underlying mathematics is complex. Transparency initiatives may include publishing model summaries, providing API documentation, or offering sandbox environments for external scrutiny. In France, the "right to explanation" under the GDPR encourages public agencies to disclose the logic of high-impact automated decisions. Auditors verify that transparency measures are proportionate, do not compromise security, and meet legal expectations.

Human-in-the-Loop (HITL) designates a workflow where human operators review and potentially modify AI-generated decisions before they are finalised. HITL is often mandated for high-risk tax decisions, such as the issuance of large penalties or the denial of tax credits. Auditors examine whether the HITL process includes clear criteria for escalation, documented review steps, and sufficient time allowances for human assessment. Evidence of effective HITL can reduce the likelihood of erroneous or unfair outcomes.

Human-on-the-Loop (HOTL) differs from HITL in that humans monitor AI decisions after they have been enacted, stepping in only if anomalies are detected. HOTL may be appropriate for lower-risk automated processes, such as routine VAT filing confirmations. Auditors assess whether monitoring mechanisms, such as exception dashboards and alert thresholds, are in place to support HOTL oversight and whether the organisation has defined response procedures for identified issues.

Recourse Mechanism provides taxpayers with a formal path to challenge automated decisions. In French tax law, recourse may take the form of an administrative appeal (recours administratif) or a judicial review (recours contentieux). AI-driven tax systems must integrate recourse mechanisms that allow taxpayers to request a manual review, receive a clear explanation, and obtain a timely resolution. Auditors verify that the system logs recourse requests, tracks their handling, and ensures that outcomes are communicated

transparently.

Impact Assessment evaluates the potential consequences of deploying an AI system, covering legal, ethical, operational, and societal dimensions. For tax AI, impact assessments might examine how automated risk scoring influences audit allocation, whether it exacerbates existing inequalities, and how it complies with data protection obligations. Auditors review impact assessment reports to confirm that mitigation strategies have been implemented and that residual risks are within the organisation's risk appetite.

Model Explainability Techniques include methods such as feature importance ranking, partial dependence plots, and surrogate modelling. These techniques help translate complex model behaviour into human-readable insights. In the tax context, a feature importance ranking might reveal that unusually high inter-company payments are a strong predictor of transfer-pricing adjustments. Auditors evaluate whether the chosen explainability techniques are appropriate for the model type and whether they are communicated effectively to both technical and non-technical stakeholders.

Bias Mitigation Strategies encompass pre-processing (e.G., Re-sampling, re-weighting), in-processing (e.G., Fairness-aware loss functions), and post-processing (e.G., Threshold adjustments) approaches. Implementing bias mitigation in tax AI requires a careful balance: Overly aggressive mitigation could reduce the model's ability to detect genuine non-compliance, while insufficient mitigation may perpetuate discriminatory outcomes. Auditors assess the documentation of bias mitigation steps, the validation of fairness metrics after mitigation, and the ongoing monitoring for re-emergence of bias.

Explainable AI (XAI) is an emerging field that seeks to build models that are inherently interpretable while retaining high predictive performance. Techniques such as Generalised Additive Models (GAMs) or rule-based classifiers can be employed for tax risk scoring, offering both accuracy and transparency. Auditors may recommend XAI approaches when the regulatory environment places a premium on interpretability, such as when AI decisions are subject to judicial review.

Model Lifecycle Management covers the end-to-end processes of model conception, development, validation, deployment, monitoring, and retirement. Effective lifecycle management ensures that models remain aligned with evolving legal requirements and business needs. In French tax administration, lifecycle management must be coordinated with the annual budget cycle, legislative updates, and the periodic audit calendar. Auditors verify that each lifecycle stage is documented, that hand-offs are clearly defined, and that decommissioning plans are in place for obsolete models.

Decommissioning is the orderly retirement of an AI model that is no longer fit for purpose. Decommissioning may be triggered by regulatory changes (e.G., New anti-avoidance rules), performance degradation, or the identification of insurmountable bias. A decommissioning plan should include data archiving, impact analysis on downstream processes, and communication to affected stakeholders. Auditors assess whether decommissioning procedures are followed, that residual data is handled according to GDPR standards, and that any dependencies are properly migrated.

Incident Response outlines the steps to be taken when an AI system experiences a failure, security breach, or compliance violation. In the tax setting, an incident could involve the accidental exposure of taxpayer

data during model training or the generation of erroneous tax notices due to a software bug. An incident response plan must define roles (e.G., Incident manager, legal counsel), escalation paths, and reporting obligations to the CNIL and the relevant supervisory authority. Auditors review incident logs, root-cause analyses, and remediation actions to ensure that the organization learns from each event.

Security Controls protect AI systems from unauthorised access, tampering, and cyber-attacks. Controls may include network segmentation, multi-factor authentication, intrusion detection, and regular penetration testing. Tax AI systems, which process highly sensitive fiscal data, must adhere to the French National Cybersecurity Agency (ANSSI) guidelines. Auditors evaluate whether security controls are proportionate to the risk, whether they are regularly tested, and whether security incidents are promptly reported.

Privacy-by-Design is a proactive approach that embeds privacy considerations into the architecture of AI systems from the outset. In practice, this may involve designing data pipelines that default to anonymised data, limiting data retention periods, and implementing robust access controls. Auditors check that privacy-by-design principles are reflected in system specifications, that privacy impact assessments have been conducted, and that the system can demonstrate compliance with GDPR throughout its lifecycle.

Data Governance encompasses the policies, standards, and processes that ensure data is accurate, available, and secure. For tax AI, data governance must address the stewardship of fiscal data, the classification of sensitive information, and the definition of data quality metrics. Auditors assess whether a data governance council exists, whether data owners are identified, and whether data quality issues are systematically addressed through remediation workflows.

Data Stewardship assigns responsibility for specific data domains to individuals who ensure data integrity and compliance. In a tax authority, a data steward for corporate income tax filings might oversee the ingestion of annual return data, verify schema conformity, and coordinate with the model development team. Auditors verify that data stewards have the authority, training, and resources necessary to fulfil their duties, and that their responsibilities are documented in the governance charter.

Model Risk Management (MRM) is a discipline that focuses on identifying, measuring, and controlling risks associated with AI models. MRM frameworks often adopt the three-line-of-defence model: First line (model developers), second line (risk and compliance functions), and third line (internal audit). In French tax administration, MRM must align with the Autorité des Normes Comptables (ANC) guidelines for risk management. Auditors evaluate whether the MRM framework includes risk appetite statements, stress testing procedures, and independent validation checkpoints.

Stress Testing subjects AI models to extreme but plausible scenarios to assess their robustness. For tax fraud detection, stress testing might involve simulating a sudden surge in cross-border e-commerce transactions or the introduction of a new tax incentive scheme. Auditors review stress-test designs, the selection of scenario parameters, and the interpretation of results, ensuring that the organization can respond effectively to adverse conditions.

Explainability Reporting provides stakeholders with documented insights into how AI decisions are derived. Such reports may include model architecture diagrams, feature importance tables, and case studies

illustrating typical decision pathways. In the French tax context, explainability reporting supports the principle of legal certainty (principe de sécurité juridique) by allowing taxpayers to understand the basis of automated assessments. Auditors verify that explainability reports are regularly updated, that they reflect the current model version, and that they are accessible to both internal reviewers and external auditors.

Legal Certainty is a cornerstone of French administrative law, requiring that public decisions be predictable and based on clear legal rules. AI systems that automate tax determinations must be designed to uphold legal certainty, meaning that the same set of inputs should consistently produce the same outcome, and that the underlying legal rules are explicitly encoded. Auditors assess whether the model's decision logic can be traced back to specific articles of the Tax Code and whether any discretionary elements are documented and justified.

Algorithmic Auditing is the practice of independently reviewing AI systems to assess compliance, performance, and ethical considerations. Algorithmic audits may be internal (conducted by the tax authority's audit department) or external (performed by certified third-party auditors). An algorithmic audit typically includes a review of data pipelines, model documentation, performance metrics, bias analysis, and governance processes. Auditors conducting the audit must follow a structured methodology, maintain independence, and provide evidence-based findings that can be acted upon by senior management.

Regulatory Sandbox offers a controlled environment where innovative AI solutions can be tested under regulatory supervision before full deployment. The French Financial Prosecutor's Office (Parquet National Financier) has experimented with sandboxes for AI-driven anti-fraud tools. In a sandbox, the tax authority can evaluate the efficacy of a new AI model for detecting illicit tax schemes while ensuring that data protection and procedural safeguards are observed. Auditors monitor sandbox activities to ensure that the limited-scope testing does not compromise taxpayer rights.

Data Ethics Board is an interdisciplinary committee tasked with overseeing the ethical use of data and AI within an organisation. For a tax administration, the board may include legal experts, data scientists, civil-society representatives, and privacy officers. The board reviews AI project proposals, assesses potential societal impacts, and issues recommendations on fairness, transparency, and accountability. Auditors review the board's minutes, decisions, and follow-up actions to ensure that ethical considerations are integrated into the AI development lifecycle.

Fairness Metrics are quantitative indicators used to assess whether an AI model treats different groups equitably. Common metrics include demographic parity, equal opportunity, and calibration across groups. In the tax domain, fairness metrics might compare the false-positive rates of audit selection for small versus large enterprises, or for taxpayers residing in different départements. Auditors evaluate the selection of fairness metrics, the thresholds set for acceptable disparity, and the remediation measures applied when metrics exceed tolerable limits.

Model Documentation Standards provide a structured format for capturing essential information about AI models. The Model Cards for Model Reporting (MCMR) initiative, endorsed by the European Commission, recommends sections such as model description, intended use, performance, ethical considerations, and maintenance. Tax authorities adopting these standards can streamline internal reviews and facilitate external

audits. Auditors check that each model's documentation conforms to the prescribed template, that it is kept up-to-date, and that any deviations are justified.

Data Lineage visualises the flow of data from source to consumption, illustrating how raw fiscal records are transformed into features for model training. Data lineage diagrams help auditors trace the origin of a specific prediction, identify potential data quality issues, and verify compliance with data minimisation principles. In the French tax ecosystem, data lineage may involve multiple sources such as the impôt sur le revenu database, the TVA collection platform, and external customs data feeds. Auditors rely on accurate lineage mappings to assess the integrity of the AI pipeline.

Model Interpretability is the degree to which a human can comprehend the internal mechanics of a model. High interpretability is often achieved with simple models like decision trees or linear regressions, which align well with the need for legal justification in tax decisions. However, more complex models (e.G., Gradient-boosted ensembles) may offer better predictive power at the cost of reduced interpretability. Auditors must weigh the trade-off between performance and interpretability, documenting the rationale for selecting a particular model type.

Auditability denotes the capacity of an AI system to be examined and verified by an independent party. Auditability is achieved through comprehensive logging, version control, documentation, and transparent governance structures. In tax law, auditability supports the principle of public oversight, ensuring that taxpayers and regulators can scrutinise automated processes. Auditors assess whether the system's design facilitates auditability, that logs are retained for the legally required period, and that access to audit artefacts is controlled but not obstructed.

Regulatory Compliance Checklist is a tool that enumerates all legal and regulatory requirements applicable to an AI system. For French tax AI, the checklist may include items such as GDPR articles, CNIL recommendations, specific provisions of the Tax Code, and sector-specific directives (e.G., The anti-abuse rule for digital services). Auditors use the checklist to verify that each requirement has been addressed, that supporting evidence (e.G., Policy documents, test results) is attached, and that any gaps are mitigated with corrective actions.

Performance Degradation Monitoring continuously tracks key performance indicators to detect declines in model effectiveness. Alerts can be configured to trigger when metrics such as precision fall below a predefined threshold, prompting a review of data drift, feature relevance, or model staleness. In tax AI, performance degradation may signal emerging tax evasion tactics that the model has not been trained on. Auditors review the monitoring configuration, the responsiveness of the mitigation process, and the documentation of any remedial actions taken.

Ethical Review Process ensures that AI projects undergo systematic evaluation of potential moral implications before deployment. The process typically includes a risk-benefit analysis, stakeholder impact assessment, and alignment with organisational values. For tax authorities, ethical review may focus on the proportionality of automated sanctions, the preservation of taxpayer dignity, and the avoidance of discriminatory outcomes. Auditors verify that the ethical review has been completed, that its findings are recorded, and that any identified concerns have been addressed.

Algorithmic Impact Statement (AIS) is a concise report that summarises the expected effects of an AI system on individuals, organisations, and society. The AIS may cover aspects such as privacy, fairness, transparency, and accountability. In the French tax context, an AIS for an AI-based tax credit eligibility engine would outline how the system determines eligibility, the data sources used, the safeguards against bias, and the mechanisms for taxpayers to contest decisions. Auditors evaluate the completeness and accuracy of the AIS, ensuring it aligns with the underlying technical documentation.

Model Retraining Schedule defines the frequency and conditions under which an AI model is updated with new data. A regular retraining cadence (e.G., Quarterly) helps maintain relevance in the face of legislative changes or shifting economic patterns. The schedule should also specify trigger conditions, such as a detected performance drop or a significant policy amendment. Auditors check that the retraining schedule is adhered to, that retraining processes are documented, and that any new model version undergoes the same validation rigor as the original.

Explainability Dashboard provides visualisations of model explanations for end-users, such as tax auditors or compliance officers. The dashboard may display feature contributions for individual predictions, aggregate importance trends, and confidence intervals. By offering interactive exploration, the dashboard facilitates trust and enables users to identify anomalous outputs. Auditors assess the usability of the dashboard, the accuracy of the displayed explanations, and the security controls that protect sensitive data displayed on the interface.

Data Retention Policy dictates how long fiscal data and model artefacts are stored before deletion or archival. Retention periods must balance operational needs, legal obligations (e.G., The five-year prescription period for tax records), and privacy considerations. For AI models, retaining training data beyond its statutory purpose may constitute a GDPR violation. Auditors verify that the retention policy is documented, that automated deletion mechanisms are in place, and that any exceptions are justified and approved.

Algorithmic Accountability Register is a public-facing ledger that records key information about AI systems used by the tax administration, including purpose, legal basis, performance metrics, and oversight mechanisms. Publishing such a register promotes transparency and builds public trust. Auditors review the register for completeness, accuracy, and timeliness, ensuring that it reflects the current state of AI deployments and that any changes are promptly reflected.

Risk-Based Prioritisation allocates audit resources according to the assessed risk level of AI-generated outputs. High-risk predictions, such as those indicating large potential revenue loss, receive immediate human review, while low-risk outputs may be processed automatically. This approach optimises the use of limited audit personnel while maintaining regulatory compliance. Auditors evaluate the criteria used for risk categorisation, the effectiveness of the prioritisation logic, and the monitoring of outcomes to ensure that risk thresholds remain appropriate.

Model Explainability Gap identifies the difference between the level of explanation required by law or policy and the actual explanatory capability of the AI system. A large gap may indicate that the model is too opaque for the intended use case. Auditors quantify the gap by comparing legal requirements (e.G., The

need for a concise rationale for each tax notice) with the model's ability to generate such rationales. If the gap is significant, auditors recommend either enhancing explainability techniques or selecting a more interpretable model architecture.

Governance Maturity Model assesses the development stage of an organisation's AI governance practices, ranging from ad-hoc processes to fully institutionalised frameworks. The maturity model helps the tax authority identify gaps, set improvement targets, and track progress over time. Auditors may use the maturity model to benchmark the current state, recommend actions to advance governance, and verify that improvement plans are being executed.

Data Subject Rights Management ensures that individuals can exercise their GDPR-mandated rights, such as access, rectification, and erasure, in relation to AI-processed data. In tax administration, data subjects may request details about how their data contributed to an automated decision. The rights management system must be capable of locating relevant data, generating understandable explanations, and responding within statutory timeframes. Auditors test the rights-management workflow to confirm that requests are handled correctly and that no undue delays occur.

Algorithmic Transparency Portal is an online platform where the tax authority publishes technical specifications, performance reports, and audit findings for its AI systems. The portal may include downloadable model cards, data dictionaries, and compliance certificates. By providing open access, the portal supports democratic oversight and fosters stakeholder confidence. Auditors verify that the portal is kept up-to-date, that sensitive information is appropriately redacted, and that the published content accurately reflects the operational reality of the AI systems.

Model Performance Benchmarking involves comparing the AI model's results against established baselines or alternative approaches. Benchmarks may include simple rule-based systems, legacy statistical models, or industry-standard algorithms. In tax fraud detection, benchmarking helps determine whether the AI solution delivers a meaningful improvement over manual screening.