

Nathanaël Dürst

Faculty of Business and Economics

Master of Science in Management (Strategy, Organization & Leadership)

Unil.

MASTER THESIS

Intelligent Diligence: Leveraging Large Language Models to Redesign KYC in Swiss Private Banking

Thesis director: **Pius Bienz**
Defense expert: **Yash Raj Shrestha**

Internship made at **Vyntra Global**
Data Science Department
Company supervisor: **Vivien Bonvin**

September 2025

Abstract

Swiss private banks face rising AML/KYC expectations while operating under strict secrecy and auditability constraints. This thesis tests whether an entirely on-premise language-model pipeline can transform unstructured KYC artefacts into goAML-compatible data that is auditable and comparable to human work. A design-science feasibility approach specifies acceptance criteria as executable checks, builds the smallest auditable proof-of-concept (PoC), and evaluates it against a human reviewer on representative materials.

Methods: A micro-corpus of 3 fully synthetic KYC packs (~20 pages total; French and English) was constructed. The PoC (“VyntraKYC”) runs fully locally on laptop-class hardware (CPU-only, 32 GB RAM) using quantized small-to-mid LLMs via `llama_cpp`, multilingual Tesseract OCR, and a lightweight checkbox detector. A schema-constrained prompt enforces official goAML tags, a single `<goAMLMessage>` root with `<Report>` children, prohibition on invention, and explicit nulls for unknowns. No training, fine-tuning, or retrieval was used. Evaluation compares PoC outputs with one reviewer on completeness and exact-match rates for mandatory fields, critical-error incidence, unsupported insertions, XML well-formedness, stability under repeated runs, and end-to-end cycle time.

Results: A privacy-preserving, on-prem pipeline reliably structures clean KYC artefacts and fails safely through explicit nulls. On the default Mistral-7B-Instruct Q4_K_M baseline, exact-match for mandatory fields reached 68% (102/150 opportunities) with a critical-error rate of 8.7% (13/150) and XML parse success of 93.3% (14/15 documents). Wall-clock time per document was approximately 12–17 minutes on CPU. Alternative small models underperformed on accuracy and well-formedness. The observed critical-error rate remains incompatible with private-bank tolerance for mistakes on mandatory fields.

Conclusion: The PoC is a disciplined feasibility artefact, not a deployable product. The near-term hardening path is primarily engineering and governance. The insights justify investing in further research and development, especially deploying a secure infrastructure, a curated retrieval corpus with provenance, validator-backed structure, and a statistically meaningful evaluation with dual reviewers and an explicit error taxonomy before any pilot in Swiss private-bank settings.

Acknowledgements

I would first like to thank Professor Bienz for his attentive supervision of this thesis and for offering me the opportunity to join the Bank Innovation Competency Center (BAICC). His guidance and trust were essential in shaping both the project and my own development.

My gratitude also goes to Jérôme Kehrli, whose support never faltered—whether before or during the many sessions preparing the BAICC meetings—and without whom I would not have been able to join Vyntra (formerly NetGuardians), where so much of this journey took shape. His expertise, shared with patience and generosity, often provided clarity at moments when it was most needed.

I am sincerely thankful to Vivien Bonvin and to the entire Data Science team at Vyntra (formerly NetGuardians) for welcoming me so openly into their company. Beyond the invaluable technical guidance they provided, their kindness and encouragement during moments of doubt made a real difference. Most of all, their challenging questions—often the ones I could not answer right away—pushed me to look deeper, think harder, and ultimately grow through this research.

The project partners of the BAICC, whose names shall remain confidential, played an essential role in this work. Their openness and willingness to share their perspectives gave me rare and invaluable insights into the inner workings of compliance functions in Swiss private banking—insights no academic text alone could have offered. Their contribution was instrumental in shaping the direction and depth of this thesis.

A very personal note of thanks goes to my roommate, Emilie Janho, for welcoming me into her home and for countless subtle but decisive nudges along the way.

Finally, my deepest gratitude goes to my parents, who gave me the freedom to explore, the safety to fail, and the encouragement to remain curious. Their support has been constant, and without it this work would not have been possible.

Table of Contents

Abstract.....	2
Acknowledgements.....	3
Table of Contents.....	4
List of Abbreviations.....	7
1. Introduction.....	10
1.1 The Evolving Global and Regional Regulatory Superstructure.....	10
1.1.1 FATF Standards as the Global North Star.....	10
1.1.2 The European Union's New AML Authority (EU-AMLA).....	11
1.1.3 The Swiss Regulatory Framework.....	11
1.2 The Operational Reality: The High Cost of Friction.....	12
1.2.1 The High Cost of Compliance.....	13
1.2.2 Process Bottlenecks and the Narrative Data Challenge.....	13
1.3 Knowledge gap.....	14
1.4 Research question, objectives, and expected contributions.....	15
1.4.1 Expected contributions.....	15
1.4.2 Scope and limitations.....	16
1.4.3 Evaluation approach.....	16
2. Literature Review.....	17
2.1 The Global Regulatory Operating Envelope.....	17
2.1.1 FATF as the global standard-setter.....	17
2.1.2 The Risk-Based Approach (RBA): A Foundational Philosophy.....	18
2.1.3 Pillars of Prevention: Customer Due Diligence, Beneficial Ownership, and Politically Exposed Persons.....	19
2.1.4 Key Jurisdictional Frameworks.....	20
2.1.4.1 The European Union: A Harmonized Supranational Model.....	20
2.1.4.2 The United Kingdom: Evolving Post-Brexit Alignment.....	21
2.1.4.3 The United States: An Enforcement-Centric Approach.....	21
2.1.4.4 Switzerland: A Specialized Financial Center Model.....	22
2.1.4.5 Comparative snapshot: positioning Switzerland among peers.....	24
2.2 Deconstructing the Know Your Customer (KYC) Process.....	25
2.2.1 The Customer Journey: Onboarding, CDD and Lifecycle Management.....	26
2.2.1.1 Identity proofing and onboarding (natural persons and entities).....	26
2.2.1.2 Beneficial-ownership (BO) resolution.....	26
2.2.1.3 Source-of-wealth (SoW) and source-of-funds (SoF) corroboration.....	27
2.2.1.4 Sanctions and PEP screening at onboarding.....	27
2.2.1.5 Risk rating and acceptance decision.....	27
2.2.1.6 Ongoing monitoring, periodic and event-driven refresh, and records.....	27
2.2.1.7 Illustrative private-banking sequence (individual HNWI).....	28

2.2.2 Levels of Scrutiny: From Simplified to Enhanced Due Diligence (EDD).....	28
2.2.2.1 Simplified Due Diligence (SDD).....	28
2.2.2.2 Standard Due Diligence (CDD).....	29
2.2.2.3 Enhanced Due Diligence (EDD).....	29
2.2.2.4 Governance constraints common to all levels.....	30
2.2.3 The Human Element: Risk Assessment and Ongoing Monitoring.....	31
2.2.4 Operational Deficiencies and Strategic Burdens of Traditional KYC.....	32
2.2.4.1 Noisy detection and alert overload.....	32
2.2.4.2 Subjective judgment and variance.....	33
2.2.4.3 Process heterogeneity and rework at scale.....	33
2.2.4.4 Data fragmentation and evidentiary burden.....	33
2.2.4.5 Measurement and learning deficits.....	34
2.2.4.6 Net effect for private banking.....	34
2.3 Technological Intermediation through Regulatory Technology (RegTech).....	34
2.3.1 The RegTech Paradigm: Architecture and KYC Applications.....	35
2.3.2 Foundational Technologies and Their Impact on Compliance.....	35
2.3.2.1 Cloud computing.....	36
2.3.2.2 Big-data architectures.....	36
2.3.2.3 Application Programming Interfaces (APIs).....	36
2.3.2.4 Automation (RPA and workflow orchestration).....	37
2.3.2.5 Implication for Swiss private banks.....	37
2.4 Artificial Intelligence: The New Frontier of KYC.....	37
2.4.1 Analytical AI: Augmenting Detection and Predictive Insight.....	38
2.4.1.1 Dynamic, behavior-aware risk scoring.....	38
2.4.1.2 From rules to patterns: anomaly and graph-based detection.....	38
2.4.1.3 Alert-triage and false-positive suppression as learning-to-rank.....	39
2.4.1.4 Model reliability under shift and scarce ground truth.....	39
2.4.1.5 Regulatory-grade explainability and reviewability.....	39
2.4.2 Generative AI: Revolutionizing Data Synthesis and Reporting.....	40
2.4.2.1 Mastering Unstructured Data: Adverse-Media Triage and Source-of-Wealth Synthesis.....	40
2.4.2.2 Automating Narrative Generation: From SAR Drafts to KYC Profile Maintenance.....	42
2.4.3 Governance and Limitations: Navigating Model Risk, Algorithmic Bias, and Explainability.....	44
2.4.3.1 Explainability and transparency.....	44
2.4.3.2 Algorithmic bias and fairness.....	44
2.4.3.3 Model risk management, drift, and technical debt.....	45
2.4.3.4 Human oversight and calibrated workflows.....	45
2.5 Synthesis.....	46
3. Methodology.....	50
3.1 Research design.....	50
3.2 Data construction: from failed synthetic paths to a workable micro-corpus.....	50
3.3 Development constraints and security posture.....	51

3.4 Proof-of-concept system architecture.....	51
3.5 Why these choices fit the research problem.....	53
3.6 Feasibility evaluation and known limitations.....	53
4. Results.....	54
4.1 Run-plan and measurements.....	55
4.2 Extraction quality.....	55
4.3 Critical-error rate.....	56
4.4 Throughput and stability.....	56
4.5 Interpretation.....	57
5. Discussion.....	57
6. Conclusion.....	59
References.....	61
Appendices.....	76
Appendix A: AI Project Canva.....	76
Appendix B: Proof-of-Concept Code.....	76
Appendix B.1: main.py.....	76
Appendix B.2: config.py.....	77
Appendix B.3: core/pipeline.py.....	78
Appendix B.4: core/utils.py.....	79
Appendix B.5: core/extractor.py.....	80
Appendix B.6: core/logger.py.....	83
Appendix B.7: core/processor.py.....	84
Appendix B.8: core/sorter.py.....	87
Appendix B.9: core/transformer.py.....	88
Appendix B.10: eval.py.....	90
Appendix C: KYC Form Example.....	93
Appendix D: Result Output Example.....	94

List of Abbreviations

AI	Artificial Intelligence
API	Application Programming Interface
AML	Anti-Money Laundering
AML/CFT	Anti-Money Laundering and Counter-Terrorist Financing
AMLA	Federal Act on Combating Money Laundering and Terrorist Financing
AMLD	Anti-Money Laundering Directive
AMLO	Anti-Money Laundering Ordinance
AMLO-FINMA	FINMA Anti-Money Laundering Ordinance
AUC	Area Under the Curve
BCBS	Basel Committee on Banking Supervision
BPR	Business Process Reengineering
BO	Beneficial Ownership
BSA	Bank Secrecy Act
BVI	British Virgin Islands
CBDDQ	Correspondent Banking Due Diligence Questionnaire
CDB 20	Agreement on the Swiss banks' code of conduct with regard to the exercise of due diligence
CDD	Customer Due Diligence
CFT	Counter-Financing of Terrorism
CHF	Swiss Franc
CIP	Customer Identification Program
CLM	Client Lifecycle Management
CRM	Customer Relationship Management
CTA	Corporate Transparency Act
CTR	Currency Transaction Report
DSR	Design Science Research
EBA	European Banking Authority
eCFR	Electronic Code of Federal Regulations
EDD	Enhanced Due Diligence

EPC	European Parliament & Council
EU	European Union
EU-AMLA	EU Anti-Money Laundering Authority
EU-AMLR	EU Anti-Money Laundering Regulation
FATF	Financial Action Task Force
FCA	Financial Conduct Authority
FDIC	Federal Deposit Insurance Corporation
FinCEN	Financial Crimes Enforcement Network
FINMA	Swiss Financial Market Supervisory Authority
FIU	Financial Intelligence Unit
GAM	Generalized Additive Model
GAO	Government Accountability Office
GenAI	Generative AI
GNN	Graph Neural Network
goAML	Anti-Money Laundering reporting system (UNODC)
HITL	Human-in-the-Loop
HM	His Majesty
HMRC	His Majesty's Revenue and Customs
HNWI	High-Net-Worth Individual
ID	Identity Document
KPI	Key Performance Indicator
KYC	Know Your Customer
LIME	Local Interpretable Model-agnostic Explanations
LLM	Large Language Model
ML	Machine Learning
ML/TF	Money Laundering/Terrorist Financing
MROS	Money Laundering Reporting Office Switzerland
N/A	Not applicable
NLP	Natural Language Processing
OCR	Optical Character Recognition
OPBAS	Office for Professional Body Anti-Money Laundering Supervision
PEP	Politically Exposed Person

pKYC	Perpetual KYC
PoC	Proof-of-Concept
PR	Precision–Recall
PR-AUC	Area Under the Precision–Recall Curve
PU	Positive–Unlabeled (learning)
QA	Question Answering
RAG	Retrieval-Augmented Generation
RBA	Risk-Based Approach
RegTech	Regulatory Technology
RM	Relationship Manager
ROC	Receiver Operating Characteristic
ROUGE	Recall-Oriented Understudy for Gisting Evaluation
RPA	Robotic Process Automation
SaaS	Software-as-a-Service
SAR	Suspicious Activity Report
SBA	Swiss Bankers Association
SDD	Simplified Due Diligence
SHAP	SHapley Additive exPlanations
SoF	Source of Funds
SoW	Source of Wealth
SPA	Sale and Purchase Agreement
SPV	Special Purpose Vehicle
SRO	Self-Regulatory Organization
SupTech	Supervisory Technology
TIN	Taxpayer Identification Number
UBO	Ultimate Beneficial Owner
UK-MLR	Money Laundering Regulations
UNODC	United Nations Office on Drugs and Crime
UK	United Kingdom
U.S.	United States of America
USD	United States Dollar
XAI	Explainable AI

1. Introduction

The Swiss private banking sector, a global bastion of wealth management, stands at a critical juncture. Its long-standing success, built on principles of discretion, stability, and bespoke service, is now being tested by a confluence of three powerful and interconnected forces: an intensifying global and domestic regulatory regime, the escalating sophistication of financial crime, and deep-seated operational inefficiencies within its core compliance functions (Swiss Bankers Association, 2024). The traditional paradigm for executing Know Your Customer (KYC) and Anti-Money Laundering (AML) obligations—a paradigm heavily reliant on manual processes and fragmented data—is proving to be operationally, financially, and strategically unsustainable. The core of the problem lies in the ever-growing volume of unstructured and semi-structured evidence, from relationship manager narratives to heterogeneous legal documents, which resist consistent validation, traceability, and reuse across the client lifecycle (Gerling & Lessmann, 2023; Moiseienko, 2023).

1.1 The Evolving Global and Regional Regulatory Superstructure

The foundation of modern AML and Counter-Financing of Terrorism (CFT) compliance is a complex, multi-layered superstructure of international standards and regional laws. For Swiss banks, navigating this environment requires adherence not only to domestic legislation but also to the global principles that shape it and the powerful regulatory frameworks of key economic partners.

1.1.1 FATF Standards as the Global North Star

The Financial Action Task Force (FATF) sets the international standard for combating money laundering and terrorist financing (Financial Action Task Force, 2012–2025). Its 40 Recommendations form a comprehensive framework that Switzerland, as a member, rigorously implements into its national law. FATF places the Risk-Based Approach (RBA) at the core of compliance. Institutions must identify, assess, and mitigate risks commensurately. (Financial Action Task Force, 2014a). This is a departure from a one-size-fits-all, "tick-box" mentality, demanding a deep, evidence-based understanding of clients and their activities.

Nowhere is this demand more acute than in the area of Beneficial Ownership (BO) transparency. In response to the widespread misuse of anonymous shell companies, the FATF has significantly strengthened its standards, particularly Recommendations 24 and 25 (Financial Action Task Force, 2023a). These updates require that competent authorities have access to adequate, accurate, and up-to-date information on the true, natural persons who ultimately own or control a legal entity. This global push places a direct and substantial burden on private banks to pierce complex corporate veils, a task that is data-intensive and analytically demanding.

1.1.2 The European Union's New AML Authority (EU-AMLA)

While Switzerland is not a member state of the European Union (EU), the bloc's ambitious overhaul of its AML/CFT framework will have profound extraterritorial effects. The establishment of the new EU Anti-Money Laundering Authority (EU-AMLA) marks a pivotal shift from a fragmented system of national directives to a centralized supervisory body with a single, directly applicable AML rulebook (Tosza, 2024; KPMG International, 2025).

EU-AMLA, which began operations in mid-2025, will have the power of direct supervision starting in 2028 over the highest-risk, cross-border financial institutions operating within the EU (PwC Ireland, 2025). Large Swiss banks with substantial EU subsidiary networks could fall directly under its jurisdiction. However, the more pervasive impact will be indirect. The stringent standards EU-AMLA imposes will inevitably create a new de facto "gold standard" for AML compliance in Europe (Bergström, 2018). To maintain the integrity of its financial center, the Swiss Financial Market Supervisory Authority (FINMA) will face significant pressure to align its own supervisory expectations with those of EU-AMLA. This will likely accelerate the adoption of concepts such as stricter BO verification and the move towards perpetual KYC (pKYC), which mandates more frequent, event-driven updates to customer information.

1.1.3 The Swiss Regulatory Framework

Switzerland's domestic AML framework is a robust, multi-layered system comprising federal laws, ordinances, and industry-level self-regulation, all under the supervision of FINMA.

- **Core Legislation:** The primary legal instruments are the Federal Act on Combating Money Laundering and Terrorist Financing (AMLA), the Anti-Money Laundering Ordinance (AMLO), and FINMA's own implementing ordinance (AMLO-FINMA) (Swiss Financial Market Supervisory Authority, n.d.-a). A crucial revision to the AMLA, effective January 1, 2023, intensified these obligations, requiring banks not just to record but to *actively verify* the beneficial owner's details and to regularly update customer data (Kunz et al., 2017).
- **Self-Regulation (CDB 20):** A unique feature is the role of the Swiss Bankers Association's (SBA) "Agreement on the Swiss banks' code of conduct with regard to the exercise of due diligence" (CDB 20). Its provisions are legally binding and provide granular, practical instructions for implementing AMLA obligations (Swiss Bankers Association, 2020a). It details the steps to identify the contracting party and, most critically, to establish the identity of the beneficial owner of the assets (Form A) and the controlling persons of any involved legal entities (Form K). These forms, along with meeting notes and client correspondence, are the primary repositories of crucial, unstructured narrative information.

The rising volume of reporting to Switzerland's financial intelligence unit, the Money Laundering Reporting Office Switzerland (MROS), underscores the operational pressure. In 2024, MROS received a record 15,141 suspicious activity reports (SARs), a 27.5% year-over-year increase, highlighting the need for structured, high-quality submissions to process the growing flow efficiently through the government's goAML information system (Money Laundering Reporting Office Switzerland, 2024; United Nations Office on Drugs and Crime, n.d.).

1.2 The Operational Reality: The High Cost of Friction

While regulatory demands have intensified, the operational infrastructure within many private banks has struggled to keep pace. The KYC value chain is frequently characterized by manual processes and fragmented technology, leading to spiraling costs, process bottlenecks, and a suboptimal client experience.

1.2.1 The High Cost of Compliance

The financial burden of AML/KYC compliance is substantial. Globally, financial institutions incurred an estimated USD 206.1 billion in financial crime compliance costs in 2023 (LexisNexis Risk Solutions, 2023). Within that total, spending on AML/KYC data and services is projected to rise 12.3% to USD 2.9 billion in 2025 (Burton-Taylor, 2025). Financial institutions commonly dedicate between 10% and 15% of their total staff to these activities, with the traditional response—hiring more compliance officers—becoming an unsustainable financial model (McKinsey & Company, 2025).

1.2.2 Process Bottlenecks and the Narrative Data Challenge

These high costs are a direct symptom of deep-seated process inefficiencies. Process-mining evidence shows that knowledge-intensive services like private banking exhibit a "long tail" of rare process variants and exceptions that cumulatively consume significant effort, elevate exception handling, and degrade throughput, particularly where unstructured narrative documents dominate (Fischer et al., 2021).

- **Fragmented Data and Legacy Systems:** Client data is often stored in disconnected, siloed systems, making a single, holistic client view difficult to achieve. This data is frequently unstructured, inconsistent, and of poor quality, severely hampering automation (Deloitte, 2025b).
- **High Rework Rates:** A study on Swiss wealth management found that, on average, process steps within the client lifecycle need to be revisited 49% of the time—a staggering level of inefficiency (Deloitte, 2025a).
- **The Narrative Data Challenge:** The most critical risk information—the client's history, the story behind their source of wealth, the rationale for complex structures—exists almost exclusively in unstructured formats like meeting notes, client letters, and legal opinions. Manually reading, interpreting, and synthesizing this information is the most time-consuming and cognitively demanding part of the KYC process, and it is highly susceptible to inconsistent interpretation (Boston Consulting Group, 2023).

This operational friction, or "strategic drag," directly hinders client acquisition, diverts relationship managers to administrative tasks, and damages the premium client experience

that justifies the private banking value proposition. The problem is not merely how to perform KYC more cheaply, but how to transform the entire function from a strategic liability into a source of institutional strength.

1.3 Knowledge gap

Despite mature standards (FATF) and an active RegTech/AI literature, there is no Swiss-private-banking-specific, auditable design that turns unstructured KYC evidence into goAML-ready data while meeting supervisory expectations. Five linked gaps remain:

- **Effectiveness and traceability:** Scholarship and official reviews highlight thin evidence linking specific KYC controls to outcomes and limited feedback loops; Swiss supervision also flags weak risk-analysis traceability. A control-by-control map of where the “long tail” of variants concentrates effort and risk in Swiss practice is missing (Levi & Reuter, 2006; Levi et al., 2018; U.S. Government Accountability Office, 2024; Nazzari & Reuter, 2025; Fischer et al., 2021; Financial Market Supervisory Authority [FINMA], 2023).
- **Document-to-structure pipeline:** Prior work shows how to retrieve and summarize knowledge and to parse semi-structured documents, but not an end-to-end, on-prem pipeline that (i) enforces schema constraints, and (ii) maps Swiss artefacts (CDB 20 Forms A/K, SoW/SoF proofs) to goAML fields with provenance for repeatable audit (Lewis et al., 2020; Guu et al., 2020; Beltagy et al., 2020; Zaheer et al., 2020; Xu et al., 2020; Huang et al., 2022; Gerling & Lessmann, 2023; UNODC, n.d.; SBA, 2020a, 2020b).
- **Governance integration:** Requirements on explainability, fairness, privacy, and auditability exist in pieces but are rarely instantiated together for narrative-heavy KYC: interpretable vs. black-box trade-offs, saliency robustness, dataset/model documentation, membership-inference and memorization risks, and fairness choices under incompatible criteria (Rudin, 2019; Ribeiro et al., 2016; Adebayo et al., 2018; Slack et al., 2020; Mitchell et al., 2019; Gebru et al., 2021; Shokri et al., 2017; Carlini et al., 2021; Chouldechova, 2017; Kleinberg et al., 2016; Corbett-Davies et al., 2023; Bartlett et al., 2022; Fuster et al., 2022).

- **Regulatory translation:** EU-level reforms (EU-AMLR/EU-AMLA) and Swiss moves on BO transparency create new operational duties (cross-checks, discrepancy management, pKYC-style triggers), yet there is no technical playbook that translates these into acceptance criteria, routing rules, and audit artefacts for Swiss private banks (EPC, 2024a, 2024b, 2024c; Council of the EU, 2024; European Commission, 2024; KPMG International, 2025; Federal Council, 2024; FATF, 2023b; FINMA, 2023).

This thesis addresses these gaps by specifying and testing a minimal, on-premis, retrieval-grounded, schema-constrained pipeline that structures KYC inputs into goAML-compatible data with provenance, evaluates it against human reviewers, and packages the required governance controls for Swiss deployment.

1.4 Research question, objectives, and expected contributions

Research question. How can AI-enabled process redesign improve KYC effectiveness and efficiency in Swiss private banks under confidentiality and audit constraints, and what trade-offs emerge when an on-premise language-model pipeline structures unstructured KYC inputs into goAML-compatible data?

Objectives.

1. Trace KYC controls to their legal and prudential sources and identify where document analytics and structured extraction add the most value.
2. Specify and implement a minimal on-premise pipeline that maps typical KYC inputs to goAML fields with retrieval and schema constraints.
3. Run a feasibility check on curated synthetic KYC packs and compare outputs to human reviewers on the same materials.
4. Document governance, auditability, and explainability controls that align with Swiss expectations and can be operated by private banks.

1.4.1 Expected contributions

Theoretical: A Swiss-focused, control-by-control reading of KYC that shows where the long tail of process variants actually occurs and how those variants concentrate operational and

compliance risk. The analysis pinpoints entry points where document analytics and structured extraction reduce rework while preserving evidentiary quality.

Design-science: A proof-of-concept that turns unstructured KYC artifacts into goAML-ready records.

Governance: A control set for AI-enabled KYC aligned with Swiss supervisory expectations and the EU AI Act. The set covers traceability, explanations linked to source passages, model-risk documentation, data protection, human oversight, and audit trails suitable for internal and external assurance.

1.4.2 Scope and limitations

The empirical work is a feasibility exercise. It uses a limited, curated set of synthetic KYC packs that reflect Swiss private-bank practice. The focus is narrow by design. The aim is to show that large language models can turn messy onboarding evidence into stable fields that line up with goAML and to surface the operational and governance implications of doing so. Findings are directional and apply to the document types and controls tested. The pipeline runs on-premise with laptop-grade resources and conservative model choices. Cycle times and throughput reflect these constraints. Production-grade infrastructure and model selection are expected to change latency materially and may lift extraction quality.

1.4.3 Evaluation approach

Evaluation compares the pipeline to human reviewers on the same KYC materials. Ground truth is set through reconciled human judgments. Measurements emphasize completeness and accuracy for critical goAML fields, exact-match rates on mandatory elements, critical-error rates, and reviewer minutes per document. All outputs are linked to an error taxonomy and a human-in-the-loop review path to capture hallucination and inconsistency risks. Governance criteria verify end-to-end traceability, explanation artifacts tied to sources, and model-risk documentation ready for audit.

2. Literature Review

2.1 The Global Regulatory Operating Envelope

The global financial system processes vast cross-border flows that attract illicit actors. In response, an international anti-money laundering and counter-terrorist financing (AML/CFT) regime has formed around the Financial Action Task Force (FATF). FATF issues principles-based standards that jurisdictions transpose into domestic law and supervisory practice. Implementation is assessed through peer reviews, and public identification of jurisdictions with strategic deficiencies exerts reputational and economic pressure. The result is a soft-law to hard-law pipeline that sets the operating envelope for financial institutions worldwide. Recent FATF stocktakes indicate marked progress in technical compliance but persistent gaps in effectiveness, reinforcing the case for improved risk management tooling. (Financial Action Task Force [FATF], 2012–2025; FATF, 2022; FATF, n.d.-a; FATF, n.d.-b).

2.1.1 FATF as the global standard-setter

FATF was created by the G7 in 1989 to develop coordinated counter-measures against money laundering. Its mandate expanded in 2001 to include terrorist financing and later to address the financing of proliferation of weapons of mass destruction. FATF's 40 Recommendations, together with Interpretive Notes and a Glossary, constitute the global AML/CFT standard and are periodically revised to reflect emerging risks—for example, the 2019 update of Recommendation 15 for virtual assets and the 2021 clarification of “environmental crime” within designated offences. More than 200 jurisdictions commit politically to these standards through the FATF Global Network of FATF and FATF-Style Regional Bodies. (Financial Action Task Force [FATF], 2012–2025; FATF, 2024; FATF, n.d.-d; FATF, n.d.-e).

Although non-binding in formal treaty terms, the standards are operationalized through mutual evaluations that test both technical compliance and effectiveness. Jurisdictions with material weaknesses are named in two standing public documents—High-Risk Jurisdictions subject to a Call for Action and Jurisdictions under Increased Monitoring—commonly referred to as the “black list” and “grey list.” These outputs are issued three times per year

and can trigger enhanced due diligence or counter-measures by countries and market participants. (Financial Action Task Force [FATF], n.d.-a; FATF, n.d.-b).

The soft-law dynamic produces measurable legal change. FATF's 2022 synthesis of the fourth evaluation round reports that 76% of countries had satisfactorily implemented the 40 Recommendations, up from 36% in 2012, although most still fall short on real-world effectiveness—especially for complex cross-border cases and beneficial-ownership transparency. (Financial Action Task Force, 2022).

2.1.2 The Risk-Based Approach (RBA): A Foundational Philosophy

The cornerstone of the modern FATF Recommendations is the Risk-Based Approach (RBA) (Arnone & Borlini, 2010). This philosophy mandates that countries, supervisors, and financial institutions must identify, assess, and understand the specific money laundering and terrorist financing risks they face to apply AML/CFT measures that are commensurate with those risks (Financial Action Task Force [FATF], 2012–2025; FATF, 2014). The RBA represents a fundamental shift away from a prescriptive, one-size-fits-all compliance model towards a more dynamic, efficient, and effective framework (Financial Action Task Force, 2014; Arnone & Borlini, 2010). By allowing for a more focused allocation of resources, the RBA enables institutions to apply more stringent controls to areas of higher risk while potentially simplifying measures where risks are lower, thus avoiding a counterproductive “tick-box” exercise (Arnone & Borlini, 2010; HM Government, 2017).

In practice, the implementation of the RBA requires financial institutions to conduct and document a comprehensive assessment of their risk exposure. This involves analyzing a range of risk factors, including those related to their customer base, the geographic areas in which they operate, the products and services they offer, their transaction patterns, and the delivery channels they use (Financial Action Task Force, 2014; HM Government, 2017; HM Treasury, 2022). The outcome of this institutional risk assessment directly informs the design of the firm's AML/CFT policies and procedures, including the level of due diligence applied to individual customer relationships (HM Government, 2017; European Banking Authority, 2021).

2.1.3 Pillars of Prevention: Customer Due Diligence, Beneficial Ownership, and Politically Exposed Persons

The FATF standards establish several core preventative measures that financial institutions must implement. These pillars are designed to create transparency and ensure that institutions have a sufficient understanding of their customers to detect and report suspicious activity.

Customer Due Diligence (CDD) is the foundational process of the FATF framework. Institutions are required to undertake CDD measures at the outset of a business relationship, when conducting occasional transactions above certain thresholds, and whenever there is a suspicion of money laundering or terrorist financing (Financial Action Task Force, 2012–2025). The core of CDD involves two key actions: first, identifying the customer (whether a natural or legal person) and, second, verifying that identity using reliable, independent source documents, data, or information (Financial Action Task Force, 2012–2025; HM Government, 2017).

Beneficial Ownership (BO) Transparency is a critical and increasingly emphasized component of the FATF standards. Illicit actors frequently exploit the anonymity provided by corporate vehicles, such as shell companies and complex trust structures, to obscure their identities and the origin of their funds (Financial Action Task Force, 2014b). To counter this, the FATF Recommendations mandate that financial institutions must not only identify their direct customer but also identify the beneficial owner—defined as the natural person(s) who ultimately own or control the customer—and take reasonable measures to verify their identity (Financial Action Task Force, 2012–2025). This requirement is central to preventing the misuse of legal persons and arrangements for criminal purposes.

Politically Exposed Persons (PEPs) are recognized by the FATF as representing a higher risk for potential involvement in bribery and corruption due to their position and influence (Financial Action Task Force, 2013). FATF Recommendation 12 requires financial institutions to implement specific enhanced due diligence (EDD) measures for business relationships with foreign PEPs, their family members, and their known close associates (Financial Action Task Force, 2012–2025; European Parliament & Council, 2024b). These measures include obtaining senior management approval to establish or continue such relationships, taking

reasonable measures to establish the customer's source of wealth and the source of funds for the transactions, and conducting enhanced ongoing monitoring of the business relationship (Financial Action Task Force, 2012–2025).

2.1.4 Key Jurisdictional Frameworks

The global harmonization sought by the FATF is met with the reality of national sovereignty and diverse legal traditions. The transposition of FATF principles into binding domestic law results in a fragmented regulatory landscape, where different jurisdictions adopt distinct philosophical approaches to supervision and enforcement. This divergence creates significant operational complexity for multinational financial institutions, which must navigate a patchwork of rules. An analysis of four major financial hubs—the European Union, the United Kingdom, the United States, and Switzerland—illustrates this tension between global standards and local implementation.

2.1.4.1 The European Union: A Harmonized Supranational Model

The European Union implements the FATF standards through a series of Anti-Money Laundering Directives (AMLDs), which require transposition into the national laws of its member states (European Parliament, 2020; European Commission, n.d.). This approach has historically led to inconsistencies in application, undermining the effectiveness of the framework across the single market (European Commission, n.d.; European Court of Auditors, 2021). To address this fragmentation, the EU has embarked on a significant reform, culminating in a directly applicable Anti-Money Laundering Regulation (EU-AMLR) and a new EU-level authority, alongside a new directive that replaces and updates aspects of the prior AMLD framework. Together these instruments aim to create a more harmonized “single rulebook” for all member states (Tosza, 2024; Bergström, 2018; European Parliament & Council [EPC], 2024b; EPC, 2024a; EPC, 2024c).

Key features of the current EU framework therefore include: the single rulebook set by the EU-AMLR; the establishment of the Authority for Anti-Money Laundering and Countering the Financing of Terrorism (EU-AMLA); and updated member-state mechanisms under the new directive (European Parliament & Council [EPC], 2024b; EPC, 2024a; EPC, 2024c). By contrast, the criminal-law harmonization—including the definition of 22 predicate offences (such as cybercrime and environmental crime), the extension of criminal liability to legal

persons (companies and partnerships), and the imposition of tougher minimum prison sentences for money laundering offenses—derives from Directive (EU) 2018/1673 and continues to apply in parallel (European Parliament & Council, 2018).

2.1.4.2 The United Kingdom: Evolving Post-Brexit Alignment

Outside the EU, the United Kingdom's AML/CFT framework continues to operate through the Money Laundering Regulations 2017. The regime remains explicitly risk-based. Firms must maintain a documented, firm-wide risk assessment and calibrate CDD and EDD to that risk, while supervisory expectations are articulated in the FCA's Financial Crime Guide. HM Treasury's 2022 review confirmed continued adherence to FATF principles with a focus on proportionality and effectiveness. (HM Government, 2017; Financial Conduct Authority, n.d.; HM Treasury, 2022).

Post-Brexit evolution has been incremental and targeted. The government and FCA have clarified proportional treatment for domestic PEPs and continued to pursue supervisory effectiveness through OPBAS oversight and HM Treasury's annual supervision reporting. The net effect is functional alignment with EU and FATF standards delivered via UK legislative and supervisory levers. (Financial Conduct Authority, 2025; Office for Professional Body Anti-Money Laundering Supervision, 2024; HM Treasury, 2025).

2.1.4.3 The United States: An Enforcement-Centric Approach

The U.S. AML/CFT regime is long-standing and enforcement-oriented. Its statutory base is the Bank Secrecy Act of 1970, which created comprehensive recordkeeping and reporting duties and is administered by the Financial Crimes Enforcement Network (FinCEN) (Financial Crimes Enforcement Network, n.d.; Office of the Comptroller of the Currency, n.d.). After September 11, 2001, Congress expanded this framework through Title III of the U.S. PATRIOT Act, accelerating a shift toward more formalized, prescriptive controls and due diligence expectations, including specific obligations for correspondent and private banking and the introduction of a mandatory Customer Identification Program (CIP) (U.S. Department of Justice, n.d.; U.S. Department of the Treasury, 2002).

In practice, U.S. requirements are codified in regulation and backed by active supervisory and enforcement programs. Banks must maintain a written AML program with internal

controls, independent testing, a designated BSA officer, and training (Electronic Code of Federal Regulations, 2025d). They must file Currency Transaction Reports for transactions in currency over \$10,000 (Electronic Code of Federal Regulations, 2025a) and Suspicious Activity Reports when activity is suspected to involve illicit funds or evasion (Electronic Code of Federal Regulations, 2025f). The PATRIOT Act's Section 326 is implemented through risk-based CIP obligations at account opening (Electronic Code of Federal Regulations, 2025e), while Section 312 adds due diligence—and, in specified cases, enhanced due diligence—for correspondent accounts for foreign financial institutions and for private banking accounts for non-U.S. persons (Electronic Code of Federal Regulations [eCFR], 2025b; eCFR, 2025c). Together these program, reporting, and due-diligence duties define the system's prescriptive character and its emphasis on verifiable controls. (Electronic Code of Federal Regulations [eCFR], 2025a; eCFR, 2025b; eCFR, 2025c; eCFR, 2025d; eCFR, 2025e; eCFR, 2025f).

Modernization has continued through rulemaking and guidance. FinCEN's 2016 Customer Due Diligence Rule clarified four core CDD elements—customer identification and verification, beneficial-owner identification and verification for legal-entity customers (generally at a 25% ownership threshold), understanding customer relationships to develop risk profiles, and ongoing monitoring with updating of customer information—and was further interpreted by agency FAQs issued in 2018 and 2020 (Financial Crimes Enforcement Network, 2016; Financial Crimes Enforcement Network, 2018; Financial Crimes Enforcement Network, 2020). These measures reinforced the program's predictive and documentation-heavy posture, aligning supervisory expectations with a consistently risk-sensitive, evidence-based approach.

2.1.4.4 Switzerland: A Specialized Financial Center Model

As a global hub for wealth management and private banking, Switzerland's AML/CFT regime is tailored to address the specific risks associated with its financial sector. As we've seen, the primary legal instrument is the Federal Act on Combating Money Laundering and Terrorist Financing (AMLA), which is supplemented by detailed ordinances issued by the Swiss Financial Market Supervisory Authority (FINMA).

However, two recent developments materially raise the bar for KYC completeness and timeliness. First, the AMLA revision effective 1 January 2023 requires financial

intermediaries not only to identify but to take reasonable measures to verify the beneficial owner in all cases and to periodically update client data on a risk basis rather than only when doubts arise (Swiss Confederation, 1997; KPMG, 2022). FINMA followed with Guidance 05/2023 after reviewing more than 30 banks and finding gaps in firm-wide risk analyses, including unclear risk tolerance, weak segmentation, and unmanaged exceptions; private banks were instructed to quantify risk drivers, link them to controls, and document the methodology that steers KYC periodicity and depth (Eichenberger, 2022; Swiss Financial Market Supervisory Authority, 2023). Second, the Federal Council in May 2024 submitted to Parliament a bill to establish a federal beneficial ownership register with discrepancy-reporting obligations for financial intermediaries. The aim is faster attribution of legal-entity ownership and closer alignment with FATF Recommendation 24; parliamentary work progressed in 2025, with scope and exemptions debated (Federal Council, 2024; Money Laundering Reporting Office Switzerland, 2024; Financial Action Task Force, 2023b; International Monetary Fund, 2024; Lenz & Staehelin, 2024). For KYC, a non-public register accessible to banks implies two changes to operating practice: systematic cross-checks of client-provided UBO data at onboarding and trigger-based discrepancy management during the relationship lifecycle.

Private-banking specificity appears in two areas of KYC execution. First, source-of-wealth and source-of-funds analysis is treated as a central risk mitigant for higher-risk clients; Swiss banks lean on the Wolfsberg Group's FAQs and effectiveness materials to structure corroboration, triangulate adverse media, and evidence plausibility in the KYC file (Wolfsberg Group, 2020; Wolfsberg Group, 2023). Second, PEP treatment follows FATF principles and AMLO-FINMA requirements: senior-management approval, SoW/SoF corroboration, relationship-level monitoring intensity, and heightened periodicity are expected and were focal points in FINMA enforcement and thematic work (Swiss Financial Market Supervisory Authority, n.d.-b; Financial Action Task Force, 2023b). Supervisory pressure has increased: MROS filings remain high and enforcement actions have stressed inadequate clarification of complex structures and insufficient PEP oversight, underscoring the expectation that KYC narratives explain the economic purpose of wealth and flows, not only satisfy form requirements (Money Laundering Reporting Office Switzerland, 2024; Swiss Financial Market Supervisory Authority, 2023; Swiss Bankers Association, 2020b).

2.1.4.5 Comparative snapshot: positioning Switzerland among peers

Following the jurisdictional narratives, this table serves as a quick-reference index for key KYC execution levers. These include primary legal bases, supervisors, beneficial-ownership thresholds, PEP scope, cash-reporting triggers, and the latest structural reforms within each bloc. It is intended as an orientation tool; the specific details for KYC design and testing are found in the jurisdiction-specific sections and cited primary sources.

Feature	FATF Standard (Reference)	European Union	United Kingdom	United States	Switzerland
Primary Legislation	N/A	Regulation (EU) 2024/1624 (EU-AMLR); Directive (EU) 2024/1640; Regulation (EU) 2024/1620 establishing EU-AMLA	Money Laundering Regulations 2017 (UK-MLR 2017)	Bank Secrecy Act (BSA) / U.S. PATRIOT Act	Anti-Money Laundering Act (AMLA); AMLO; AMLO-FINMA
Key Supervisor(s)	N/A	EU-AMLA (direct oversight of selected cross-border groups) with national competent authorities; EBA for technical standards	FCA; HMRC; OPBAS for professional-body supervisors	FinCEN and the federal banking agencies	FINMA; SROs (for certain non-prudentially supervised intermediaries)
Beneficial Ownership Threshold	Rec. 10, 24, 25	25% ownership or control	25% ownership or control	25% ownership plus control prong (CDD Rule)	25% controlling-pers on benchmark (or senior managing officer if none)
PEP	Rec. 12	Domestic and	Domestic and	Focus on foreign	Domestic and

Definition		foreign PEPs	foreign PEPs	PEPs for EDD; domestic PEPs risk-based	foreign PEPs
Key CDD Cash Transaction Threshold	Rec. 10	EU-wide €10,000 cash cap under EU-AMLR	€10,000 for High-Value Dealers	\$10,000 CTR threshold (BSA)	Traders generally CHF 100,000; proposed CHF 15,000 for precious metals and gemstones
Key Recent Development	N/A	Single rulebook plus creation of EU-AMLA	Post-Brexit recalibration under UK-MLR 2017 with OPBAS oversight	CTA beneficial-owne rship reporting regime	2023 AMLA revision; 2024/2025 proposals on cash thresholds and transparency

Table 1: Comparative overview of FATF standards and national AML/CFT regulatory frameworks across the European Union, United Kingdom, United States, and Switzerland.

2.2 Deconstructing the Know Your Customer (KYC) Process

The Know Your Customer (KYC) process is not a singular event but a continuous, dynamic function that forms the core of an institution's AML/CFT defense mechanism. It is best understood as an integral component of the broader Client Lifecycle Management (CLM) framework. This chapter will deconstruct the stages of the KYC journey, from initial onboarding to ongoing monitoring, and will culminate in a critical analysis of the profound operational deficiencies and strategic burdens inherent in traditional, manual KYC methodologies.

2.2.1 The Customer Journey: Onboarding, CDD and Lifecycle Management

2.2.1.1 Identity proofing and onboarding (natural persons and entities)

Data capture and documents. At intake the client (often with the relationship manager) completes the KYC questionnaire capturing legal name, date and place of birth, nationality, residential address, tax residence/TIN, occupation, expected account use, anticipated activity, and declarations on PEP status; the bank collects and verifies core evidence: a government photo ID (e.g., passport or national ID) and a recent proof of address (e.g., utility bill, tax assessment, bank statement) (FATF, 2012; EBA, 2021). For legal persons the bank obtains proof of legal existence and capacity and maps ownership/control: certificate of incorporation or commercial-register extract, articles, register of directors/authorized signatories, specimen signatures, corporate resolutions, group structure chart, principal place-of-business evidence, tax numbers, and recent financials (EBA, 2021; BCBS, 2020). In Swiss practice, banks additionally gather CDB 20 artifacts: Form A to identify the beneficial owner of the assets and Form K to identify the controlling person of operating legal entities; trusteeships require the trust deed and identification of settlor, trustees, protector, and beneficiaries (SBA, 2020a; SBA, 2020b).

2.2.1.2 Beneficial-ownership (BO) resolution

Firms must identify and take reasonable measures to verify natural persons who ultimately own or control the customer (legal persons and arrangements), using a control-plus-ownership approach with a commonly applied 25% benchmark and a “control” fallback where no threshold is met (FATF, 2012; FinCEN, 2016; EBA, 2021). Evidence typically includes company registers, official filings, notarized shareholder lists, audited financials, organizational charts, and, where available, UBO-register extracts; for trusts, identification extends to settlor, trustees, protector, beneficiaries, and other persons with effective control (FATF, 2012; EBA, 2021). In Swiss banking, CDB 20 documentation (Form A/Form K) evidences the BO/controlling-person determination in the file (SBA, 2020a; SBA, 2020b).

2.2.1.3 Source-of-wealth (SoW) and source-of-funds (SoF) corroboration

Private-banking KYC hinges on plausibility and independent corroboration of how wealth was accumulated and how specific transfers are funded. Typical evidence includes employment contracts and payslips; audited accounts and dividend vouchers; sale-and-purchase agreements for businesses or real estate; inheritance and probate papers; gift deeds; brokerage statements; and land-register extracts. The file must explain the economic story, reconcile values, and link documents to the customer's profile; where gaps persist, restrictions, enhanced monitoring, or rejection follow (Wolfsberg Group, 2020; BCBS, 2020; EBA, 2021).

2.2.1.4 Sanctions and PEP screening at onboarding

Screening against targeted financial-sanctions lists and PEP databases occurs before acceptance and at defined intervals thereafter, with matching logic, data-quality controls, and escalation governance documented in policy (EBA, 2024). For PEPs, firms obtain senior-management approval, perform SoW/SoF corroboration, and calibrate monitoring and review periodicity commensurate with risk (FATF, 2013; EBA, 2021).

2.2.1.5 Risk rating and acceptance decision

Risk assessment weights customer, product/service, channel, and geography to arrive at a risk rating that drives onboarding approvals, KYC depth, account restrictions, and review frequency (EBA, 2021; BCBS, 2020). FATF permits limited deferral of verification only where essential for business continuity and where ML/TF risk is effectively managed; otherwise, the relationship is not commenced or is exited, and suspicion is reported to the FIU (FATF, 2012).

2.2.1.6 Ongoing monitoring, periodic and event-driven refresh, and records

Ongoing due diligence assesses whether transactions are consistent with the customer's profile and triggers reviews on events such as BO changes, unusual flows, adverse media, or sanctions updates (FATF, 2012; BCBS, 2020). Review frequency is risk-based in regulation; many private banks operationalize this as annual reviews for high risk and multi-year cycles for lower tiers, supplemented by trigger-driven updates (EBA, 2021; Fenargo, 2024; BCG, 2025). Recordkeeping must allow independent re-performance; FATF

prescribes at least five years, while Swiss law requires ten years after relationship end or transaction completion (FATF, 2012; Swiss Confederation, 1997).

2.2.1.7 Illustrative private-banking sequence (individual HNWI)

The RM co-completes the KYC form; collects passport and recent address proof; obtains Form A and SoW evidence for a business sale and listed securities; performs video identification with liveness; compliance resolves BO via company filings; sanctions/PEP screening runs; the case goes to a risk/PEP committee; acceptance is conditional on receipt of notarized sale contract and dividend statements; account opens with initial limits pending document verification; periodicity is set to annual with adverse-media monitoring (SBA, 2020a; FINMA, 2023; EBA, 2021; Wolfsberg Group, 2020).

2.2.2 Levels of Scrutiny: From Simplified to Enhanced Due Diligence (EDD)

KYC operates on a graduated scale of measures calibrated to risk. At a minimum, institutions must identify and verify the customer, determine beneficial ownership for legal persons and arrangements, understand the purpose and intended nature of the relationship, and conduct ongoing monitoring. The intensity of those controls increases with risk and must never be reduced where suspicion arises or higher-risk indicators are present (FATF, 2012; BCBS, 2020; EBA, 2021).

2.2.2.1 Simplified Due Diligence (SDD)

SDD is permissible only where a demonstrably lower risk is evidenced by product, customer, channel, and geographic factors, and where national law allows reduced measures. Under FATF's Interpretive Note to Recommendation 10 and EU supervisory guidance, SDD may involve limiting the data collected at onboarding, deferring some verification steps until after account opening, or applying less frequent review—provided sanctions screening and ongoing monitoring remain effective and the institution can evidence why risk is low (FATF, 2012; EBA, 2021). SDD is categorically prohibited where the institution suspects ML/TF, where higher-risk factors are present, or where the customer is a PEP or linked to a high-risk third country (FATF, 2012; EBA, 2021).

2.2.2.2 Standard Due Diligence (CDD)

Standard CDD comprises (i) identity proofing of the customer, (ii) identification and risk-based verification of beneficial owners and controllers, (iii) understanding purpose and intended nature (including anticipated activity), and (iv) ongoing monitoring with risk-based periodicity. Documentation typically includes a government photo ID and recent address evidence for natural persons; for legal persons, corporate registry extracts, constitutive documents, authorized-signatory evidence, and ownership/control maps. Adverse-media screening and sanctions/PEP checks are embedded at onboarding and throughout the lifecycle (FATF, 2012; BCBS, 2020; EBA, 2021).

Swiss private-banking practice: banks evidence BO/controlling-person determinations using CDB 20 Form A (beneficial owner of assets) and Form K (controlling person of operating entities), alongside corroboration from the commercial register and organizational charts (SBA, 2020a; SBA, 2020b).

2.2.2.3 Enhanced Due Diligence (EDD)

EDD augments core CDD with additional information, corroboration, governance, and monitoring. Triggers recognized by standards and supervisors include: (i) PEP status and close associates/family members, (ii) customers from or with material exposure to higher-risk jurisdictions, (iii) complex or opaque ownership/control (including trusts or multi-layered SPVs), (iv) private banking and correspondent banking relationships, (v) non-face-to-face onboarding without strong mitigants, and (vi) unusual or unusually large/complex transactions without apparent lawful purpose (FATF, 2012; FATF, 2013; EBA, 2021; BCBS, 2020; Wolfsberg Group, 2023). EDD measures typically include:

1. **Senior management approval** to establish or continue the relationship, especially for PEPs and certain cross-border relationships (FATF, 2013; EBA, 2021).
2. **Source-of-wealth (SoW) and source-of-funds (SoF) corroboration** using independent evidence proportionate to risk. Examples include audited financial statements, tax assessments, notarized sale-and-purchase agreements, inheritance and probate documentation, dividend vouchers, brokerage statements, or land-registry extracts. The analysis must reconcile narrative, values, and timing, and

be recorded in a reasoned KYC memorandum (Wolfsberg Group, 2020; BCBS, 2020; EBA, 2021).

3. **Deeper beneficial-ownership verification** and mapping of effective control, including resolution of indirect ownership chains; identification of settlor, trustees, protector, beneficiaries, and other controllers for trusts; and targeted inquiry where thresholds (often 25%) are not met but control exists through other means (FATF, 2012; EBA, 2021).
4. **Heightened ongoing monitoring**, including lower alert thresholds, scenario tuning for expected activity, and more frequent periodic reviews, supplemented by event-driven refresh on adverse-media hits, sanctions changes, ownership changes, or material deviations from the customer profile (FATF, 2012; BCBS, 2020; EBA, 2021).
5. **Correspondent-banking EDD** using structured information requests—e.g., Wolfsberg CBDDQ—to assess respondent governance, control environment, nested relationships, and payable-through risks, with documented risk-acceptance and periodic re-assessment (Wolfsberg Group, 2023; BCBS, 2020).

For example: a Zurich-based RM proposes onboarding a non-resident PEP beneficial owner of a BVI holding company funding an account with proceeds from a recent business sale. EDD collects the notarized SPA and escrow statements, triangulates values with audited accounts and public deal databases, resolves the ownership chain to the natural person via registry documents, obtains senior-management approval, and sets annual review with enhanced adverse-media and cross-border payment monitoring (FATF, 2013; SBA, 2020a; Wolfsberg Group, 2020).

2.2.2.4 Governance constraints common to all levels

Across SDD, standard CDD, and EDD, three constraints apply. First, no SDD where suspicion or higher-risk indicators exist, and no onboarding where CDD cannot be completed (FATF, 2012; BCBS, 2020). Second, sanctions implementation must be effective regardless of diligence level, with documented data quality, matching logic, escalation paths, and recordkeeping (EBA, 2024). Third, auditability is non-negotiable: the file must allow an independent reviewer to re-perform the risk assessment and understand acceptance decisions and monitoring design (BCBS, 2020; FATF, 2012).

2.2.3 The Human Element: Risk Assessment and Ongoing Monitoring

Risk assessment in KYC is ultimately a judgment task shaped by governance, data quality, and institutional risk appetite. Even with formal risk taxonomies, inter-rater variance and undocumented exceptions drive inconsistent ratings, especially in complex private-banking cases. Comparative scholarship shows that AML outcomes depend less on the volume of rules than on how decision-makers interpret risk concepts in practice, while national-risk-assessment research highlights methodological discretion and expert-opinion bias as persistent sources of subjectivity (Levi & Reuter, 2006; Ferwerda & Reuter, 2018). Swiss supervisory work converges with this literature: FINMA's 2023 review of more than 30 banks found unclear risk tolerance, weak segmentation, and unmanaged exceptions in firm-wide risk analyses, each of which propagates mis-ratings into KYC execution (FINMA, 2023).

Recent peer-reviewed studies formalize these judgment frictions. A probabilistic AML risk-assessment framework demonstrates that experts' assessments are sensitive to uncertainty specification and weighting choices across customer, product, channel, and geography factors, with measurable impacts on final ratings and review cadences (Ogbeide et al., 2023). Mixed-methods evidence similarly documents how risk assessors' heuristics and institutional incentives shape classifications, reinforcing the need for transparent weighting, audit trails, and independent testing (Ogbeide, 2024). Supervisory standards translate these insights into controls: board-owned risk analyses, documented factor weightings, QA and independent validation, and tight linkage from ratings to CDD depth and periodicity (BCBS, 2020; EBA, 2021).

Ongoing monitoring remains a human-intensive alert-triage and case-narrative process. Legacy rules engines are interpretable yet noisy; high false-positive rates drive large manual workloads and alert fatigue. FATF observes that advanced analytics can reduce false positives only when paired with sound governance and feedback loops, while recent peer-reviewed work in transaction monitoring reports structural inefficiencies and motivates machine-learning augmentation under strong human oversight (FATF, 2021; Öztas et al., 2024; BCBS, 2020). In practice, L1–L3 queues enrich alerts with KYC data, payment metadata, and adverse-media findings, escalate on residual suspicion, and feed tuning with closed-case outcomes (BCBS, 2020; EBA, 2021).

Effectiveness turns on the quality of suspicious activity reporting and supervisory feedback. Authorities emphasise complete, timely SARs with coherent narratives and evidentiary attachments; U.S. oversight bodies explicitly call for improved guidance and feedback to enhance SAR utility for law enforcement and supervisory analytics (FDIC, n.d.; GAO, 2024). Swiss data illustrate the operational load: MROS received 15,141 SARs in 2024, up 27.5% year-over-year, underscoring the centrality of human triage and narrative quality in high-volume environments (MROS, 2024).

For periodic and event-driven refresh, standards require demonstrable linkage from risk factors to review scope and cadence, suspension of simplified measures when risk rises, and records that allow independent re-performance. Swiss guidance aligns with this view by demanding traceable connections between firm-wide risk analyses and KYC depth and frequency; without this traceability, monitoring efficacy degrades (BCBS, 2020; EBA, 2021; FINMA, 2023).

2.2.4 Operational Deficiencies and Strategic Burdens of Traditional KYC

Traditional KYC concentrates three hard constraints in one workflow: detection accuracy under legal uncertainty, cost containment in labor-intensive review, and preservation of client experience. The literature shows that banks resolve this triad with heavy manual scrutiny and expansive documentation that generate large volumes of weak signals, variable judgments, and limited learning feedback—features that systematically depress effectiveness while raising cost (Levi & Reuter, 2006; Halliday et al., 2014; Levi et al., 2018).

2.2.4.1 Noisy detection and alert overload

Rule-based surveillance and categorical risk taxonomies produce a high false-positive base rate, particularly when data are sparse, heterogeneous, or poorly standardized (Öztas et al., 2024). Expert interviews in a cross-industry AML study report that alert triage is dominated by low-value cases created by threshold heuristics, forcing analysts to assemble context from fragmented sources before deciding whether to escalate or close (Öztas et al., 2024). Classic syntheses on AML effectiveness anticipate this structural noise problem: when indicators are weak and typologies adaptive, the system expands reporting without commensurate gains in confirmed criminal cases (Levi & Reuter, 2006; Nazzari & Reuter, 2025).

2.2.4.2 Subjective judgment and variance

KYC decisions rely on professional judgment under partial information. Empirical work on AML risk assessment shows experts defaulting to organizational heuristics and risk frames, with measurable bias and outcome variance across units and individuals (Ogbeide et al., 2023; Ogbeide, 2024). Ethnographic criminology similarly documents how “suspicion-making” in financial policing is socially constructed and contingent on internal cues and denunciations rather than uniformly applied evidentiary thresholds (Amicelle & lafolla, 2018). In practice, two clients with similar profiles can receive different ratings or documentation demands, with spillovers to review periodicity, monitoring intensity, and SAR quality.

2.2.4.3 Process heterogeneity and rework at scale

KYC is a knowledge-intensive, exception-prone process. Process-mining evidence across industries demonstrates a “long tail” of low-frequency variants that attract disproportionate effort, handoffs, and rework—precisely the pattern observed in bespoke wealth-management onboarding, complex ownership structures, and source-of-wealth corroboration (Fischer et al., 2021). Each rare variant (e.g., a trust with protector powers, multi-layer SPVs spanning several registries, or a politically exposed beneficial owner) triggers one-off document hunts and narrative reconciliations that are difficult to templatize and expensive to assure.

2.2.4.4 Data fragmentation and evidentiary burden

Studies link weak data standardization and limited cross-system interoperability to elevated investigation time and inconsistent decisions (Öztas et al., 2024). The most probative artifacts in private banking—relationship-manager notes, legal opinions, sale-and-purchase agreements, probate files—are unstructured. Converting this narrative evidence into consistent, auditable explanations of source of wealth or economic purpose is the slowest step, and small inconsistencies generate exception loops that further lengthen cycle times (Fischer et al., 2021; Amicelle & lafolla, 2018).

2.2.4.5 Measurement and learning deficits

A persistent finding in the academic record is the paucity of outcome-relevant metrics and feedback loops. Despite extensive rule harmonization, there remains little credible evidence connecting specific KYC controls to reductions in predicate crime or laundering harm, which impedes cost-effectiveness optimization and entrenches box-ticking (Levi et al., 2018; Halliday et al., 2014). Absent reliable labels and benchmarking, firms tune controls to process KPIs (throughput, aging) rather than to validated detection outcomes, reinforcing the false-positive cycle identified above (Öztaş et al., 2024; Ogbeide et al., 2023).

2.2.4.6 Net effect for private banking

In high-touch wealth management, these mechanics manifest as prolonged onboarding, episodic documentation sprints around trigger events, and high SAR drafting costs. Typical micro-sequences include client–RM co-completion of forms; iterative requests for certified identity and address proofs; chain-of-ownership reconstruction from registry extracts; and assembly of source-of-wealth narratives from contracts, audited accounts, and tax assessments. Each step compounds the long-tail and subjectivity problems, raising the probability of inconsistent classification and costly re-work (Fischer et al., 2021; Amicelle & lafolla, 2018).

2.3 Technological Intermediation through Regulatory Technology (RegTech)

Post-crisis regulatory expansion and rapid digitization created a structural mismatch between rule complexity and manual compliance capacity. Scholarship locates the emergence of “RegTech” at this junction: the market-side pull of escalating compliance obligations and costs, and the technology-side push of scalable data, cloud, and API capabilities that reconfigure monitoring, reporting, and control execution (Arner et al., 2017; Anagnostopoulos, 2018; Baxter, 2016; Gomber et al., 2017). In parallel, authorities adopted analogous toolsets for supervision (SupTech) which further reinforced data-centric compliance expectations (Broeders & Prenio, 2018). For Swiss private banks, the immediate relevance of RegTech lies in KYC’s data-assembly and evidentiary burdens:

identity proofing, beneficial-ownership resolution, and risk-based documentation can be re-architected as digital pipelines rather than episodic, paper-bound tasks (Arner et al., 2019; Schlatt et al., 2022).

2.3.1 The RegTech Paradigm: Architecture and KYC Applications

In the academic literature, RegTech denotes the application of information technologies to regulatory monitoring, reporting, and compliance inside regulated firms, with a complementary supervisory strand, SupTech, on the authority side (Arner et al., 2017; Anagnostopoulos, 2018; Broeders & Prenio, 2018). Conceptually, RegTech is not a single tool but an architectural stance that (i) digitizes evidentiary inputs, (ii) standardizes data models across siloed systems, and (iii) automates control execution and audit trails through software-mediated workflows (Arner et al., 2017; Baxter, 2016). This paradigm directly targets KYC frictions: digital identity and remote onboarding reduce front-end documentation latency; data-aggregation layers and APIs collapse fragmented records into a single client view; and machine-executable policies enable repeatable, testable treatments across risk tiers (Anagnostopoulos, 2018; Arner et al., 2019; Schlatt et al., 2022). Early KYC-focused designs operationalize this logic—for example, Parra Moyano and Ross (2017) specify a shared-verification architecture that eliminates redundant onboarding across institutions, while Schlatt et al. (2022) formalize self-sovereign identity design principles that reconcile privacy with verifiable credentials and regulator-acceptable auditability. Taken together, these strands ground RegTech as the necessary data-engineering precursor to later AI-enabled KYC, where model performance and explainability depend on high-quality, well-governed inputs and reproducible workflows (Arner et al., 2017; Gomber et al., 2017).

2.3.2 Foundational Technologies and Their Impact on Compliance

The first wave of RegTech solved “data-plumbing” problems rather than analytic ones. It digitized evidentiary inputs, connected silos, and made controls executable as software, thereby reducing latency, rework, and audit friction in KYC. Four foundational technologies recur in the literature.

2.3.2.1 Cloud computing

Elastic infrastructure and managed services lowered fixed costs, enabled rapid scaling of data and document-heavy onboarding, and supported standardized environments for testing and audit (Armbrust et al., 2010; Marston et al., 2011). For compliance, cloud architectures also facilitate immutable logging, segregation of duties, and environment reproducibility (prerequisites for demonstrable control effectiveness) while raising governance questions on data locality and third-party risk that supervisory guidance addresses through cloud/outsourcing expectations (Broeders & Prenio, 2018; Basel Committee on Banking Supervision [BCBS], 2021).

2.3.2.2 Big-data architectures

Columnar stores, streaming pipelines, and data-lake patterns absorb heterogeneous KYC artifacts—ID images, corporate filings, registries, news—and preserve lineage for re-performance (Chen et al., 2014; Gomber et al., 2017). This layer attacks the core bottleneck identified in KYC: assembling a single, queryable client view from structured and unstructured sources with auditable provenance, aligning with regulatory principles on effective risk data aggregation (Basel Committee on Banking Supervision [BCBS], 2013). In practice, schema standardization and metadata capture convert narrative proofs of identity, beneficial ownership, and source of wealth into machine-addressable evidence that downstream controls can reuse (Anagnostopoulos, 2018).

2.3.2.3 Application Programming Interfaces (APIs)

APIs interconnect bank cores, screening utilities, registries, and workflow tools, replacing manual rekeying with machine-to-machine exchange and policy-as-code execution. In the KYC context, APIs enable real-time registry queries for beneficial ownership, sanctions/PEP re-screening at lifecycle events, and straight-through document validation; at ecosystem level, APIs underpin shared KYC utilities and cross-firm verification designs (Gomber et al., 2017; Arner et al., 2019; Parra Moyano & Ross, 2017). The API layer thus operationalizes interoperability, the necessary condition for collapsing onboarding cycle times without sacrificing evidentiary quality.

2.3.2.4 Automation (RPA and workflow orchestration)

Robotic Process Automation and model-driven workflows codify repeatable KYC steps (document intake, data extraction, register lookups, checklist assembly) freeing human analysts to focus on judgment-heavy tasks like adverse-media materiality and source-of-wealth plausibility (van der Aalst, 2018; Syed et al., 2020). Empirical studies show RPA's strengths in high-volume, rule-based tasks and its limits where data variety and exceptions dominate, precisely the frontier that later AI components target (Anagnostopoulos, 2018; van der Aalst, 2018). When combined with disciplined exception handling and audit trails, automation reduces variance in control execution and improves re-performance.

2.3.2.5 Implication for Swiss private banks

These building blocks directly mitigate the literature-identified KYC frictions (siloed evidence, manual rework, and weak lineage) by (i) consolidating artifacts with traceable provenance (big data), (ii) enabling straight-through retrieval and re-screening (APIs), (iii) scaling document-heavy workflows (cloud), and (iv) standardizing repetitive steps while reserving analyst time for judgment (RPA) (Anagnostopoulos, 2018; Arner et al., 2017; Parra Moyano & Ross, 2017). They constitute the prerequisite data and process substrate on which later analytical and generative AI can operate credibly.

2.4 Artificial Intelligence: The New Frontier of KYC

AI shifts compliance from rules executed on static snapshots to models that learn, adapt, and reason over high-dimensional, networked, and time-varying data. AML/KYC presents exactly the conditions where modern ML excels and traditional controls struggle: extreme class imbalance, scarce and noisy labels (SARs as imperfect proxies), relational dependencies among entities, concept drift in typologies, and heterogeneous evidence streams (text, images, graphs, and time series). Foundational research on anomaly detection, imbalanced learning, graph representation, and drift adaptation provides the methodological spine for this shift, while advances in interpretability and uncertainty quantification establish the guardrails required in a regulated setting (Chandola et al.,

2009; He & Garcia, 2009; Akoglu et al., 2015; Gama et al., 2014; Lundberg & Lee, 2017; Ribeiro et al., 2016; Guo et al., 2017; Angelopoulos & Bates, 2021).

2.4.1 Analytical AI: Augmenting Detection and Predictive Insight

Analytical AI addresses three structural limits of traditional KYC: (i) static, one-off risk scoring; (ii) rule-based alerting that cannot generalize beyond hand-coded scenarios; and (iii) brittle decisioning under label scarcity and distributional shift.

2.4.1.1 Dynamic, behavior-aware risk scoring

Instead of a single onboarding rating, supervised and semi-supervised models update risk continuously from transactional, contextual, and relational signals. Time-aware learning captures evolving baselines through sequential models and drift-adaptation strategies, reducing stale classifications that periodic reviews entrench (Gama et al., 2014). In rare-event regimes, class-imbalance methods and calibrated probabilities prevent over-confident low-risk predictions that mask emerging threats (He & Garcia, 2009; Guo et al., 2017). Where ground-truth labels are partial—typical in AML—positive-unlabeled (PU) learning and weak supervision leverage SARs and investigative outcomes without assuming that non-SAR cases are benign, improving recall at fixed review capacity (Bekker & Davis, 2020; Ratner et al., 2017).

2.4.1.2 From rules to patterns: anomaly and graph-based detection

Rules detect what is specified; anomaly detection and graph learning surface what is merely *atypical* given the joint distribution of behaviors and relationships. Unsupervised and one-class methods—Isolation Forests, one-class SVMs, autoencoders—model normality and flag deviations, a fit for typology change and novel schemes (Chandola et al., 2009). Because illicit finance is relational, graph-based approaches exploit network structure: centrality shifts, community anomalies, and suspicious motifs in customer-counterparty-beneficial-owner graphs. Surveys in *ACM Computing Surveys* and *IEEE TNNLS* document how graph anomaly detection and graph neural networks (GNNs) amplify weak local signals with neighborhood context, a critical advantage when individual transactions appear innocuous in isolation (Akoglu et al., 2015; Wu et al., 2021). In KYC, this translates into entity-resolution across silos, link-prediction to surface hidden connections,

and propagation-based risk scoring that prioritizes reviews where relational evidence concentrates.

2.4.1.3 Alert-triage and false-positive suppression as learning-to-rank

Legacy monitors over-fire. Treating triage as ranking under severe imbalance outperforms threshold tweaks. Learning-to-rank and cost-sensitive classification optimize precision at the investigator bandwidth frontier; evaluation uses precision–recall curves and average precision rather than ROC, which is misleading for rare events (Davis & Goadrich, 2006; Saito & Rehmsmeier, 2015). Active-learning policies then query the most informative alerts to label, tightening model uncertainty where it matters operationally (Settles, 2012). In combination, these methods reallocate analyst attention toward alerts with higher posterior utility rather than raw rule severity.

2.4.1.4 Model reliability under shift and scarce ground truth

KYC environments drift—products, geographies, counterparties, and tactics change. Robust deployment requires explicit drift detectors, periodic re-weighting, and uncertainty-aware decisions. Conformal prediction and calibrated posteriors provide coverage guarantees and thresholds aligned to supervisory risk appetite, supporting risk-based holdbacks and human review (Angelopoulos & Bates, 2021; Guo et al., 2017). PU-learning assumptions and weak-label taxonomies must be validated to avoid feedback loops from biased SARs; without such checks, apparent gains can reflect reporting habits rather than better detection (Bekker & Davis, 2020).

2.4.1.5 Regulatory-grade explainability and reviewability

Analytical AI must be legible. Local post-hoc explainers (LIME, SHAP) and inherently interpretable structures (monotonic gradient-boosted trees, scorecards, GAMs) allow investigators and auditors to reconstruct why a client or alert was prioritized, with stable feature attributions and case-level narratives (Ribeiro et al., 2016; Lundberg & Lee, 2017). In high-stakes settings, the literature cautions against uncritical black-box use when transparent models achieve comparable accuracy; this aligns with supervisory expectations for traceable rationales in KYC decisions (Rudin, 2019).

For Swiss private banks, analytical AI is most effective when embedded where the literature is strongest: (i) graph-based enrichment of KYC entity graphs; (ii) learning-to-rank triage with PR-AUC as the governing metric; (iii) PU/weak-supervision pipelines to cope with SAR-biased labels; (iv) calibrated, drift-aware deployment with conformal thresholds; and (v) explanations designed for investigator workflows and audit files. These are not cosmetic upgrades to rules—they are statistically grounded replacements for static scoring and heuristic triage, tuned to the empirical constraints of KYC.

2.4.2 Generative AI: Revolutionizing Data Synthesis and Reporting

Generative AI (GenAI), and in particular large language models (LLMs), addresses KYC's core bottleneck: transforming heterogeneous, largely unstructured evidence into consistent, auditable analyses and narratives. Unlike analytical models that score patterns in structured data, GenAI performs *document-grounded* synthesis at scale—summarizing, extracting, and drafting while preserving provenance through retrieval and citation. The academic record anchors both capability and constraint: LLMs excel at knowledge-intensive summarization and question answering when paired with retrieval (Lewis et al., 2020; Guu et al., 2020), long-context encoders (Beltagy et al., 2020; Zaheer et al., 2020), and domain adaptation (Gururangan et al., 2020), yet require safeguards for factuality, attribution, privacy, and robustness (Maynez et al., 2020; Ji et al., 2023; Carlini et al., 2021).

2.4.2.1 Mastering Unstructured Data: Adverse-Media Triage and Source-of-Wealth Synthesis

Enhanced due diligence is a reading problem before it is a modeling problem. A single high-risk onboarding can surface dozens of articles across languages, hundreds of pages of court filings, registry extracts with inconsistent transliterations, and client-provided contracts that mix scans, tables, and free-text. The compliance task is to turn that sprawl into a defensible account of what happened, when, and with which parties, and to tie each claim to verifiable evidence. Generative AI adds leverage precisely here: not by replacing human judgment, but by compressing the search-and-synthesis cycle into a documented sequence that an auditor can replay.

The contemporary solution begins with retrieval-augmented generation (RAG). Instead of prompting an LLM “from memory,” the system first retrieves passages from an approved corpus—licensed adverse-media feeds, official gazettes, company registries—and then conditions the generator on those passages. This architecture consistently improves factual accuracy and reduces hallucination because the model is asked to compose with citations rather than invent from parametric knowledge (Lewis et al., 2020; Guu et al., 2020; Karpukhin et al., 2020). In practice, two design choices matter for KYC. First, retrieval must handle multilingual and cross-script evidence so that “Viktor,” “Viktar,” and “Віктор” co-resolve; long-context and multilingual pretraining mitigate drop-off in non-English sources (Beltagy et al., 2020; Zaheer et al., 2020; Xue et al., 2021). Second, documents are pre-processed to de-duplicate wire copies, cluster near-duplicates by entity and event, and assign temporal anchors, which improves both retrieval precision and the coherence of the generated narrative.

Once grounded passages are in hand, cross-document reasoning becomes the bottleneck: entities recur under aliases, events unfold asynchronously across jurisdictions, and outcomes appear in footnotes or annexes. Sparse-attention transformers extend the context window enough to keep a timeline “in mind,” enabling the model to align allegations, proceedings, and resolutions across sources rather than summarize each item in isolation (Beltagy et al., 2020; Zaheer et al., 2020). Document-level NLP components—coreference and entity linking—further stabilize who did what to whom, which is essential when moving from “there was negative press” to a legally meaningful statement about a specific beneficial owner (Lee et al., 2017; Chalkidis et al., 2021).

KYC artifacts are not only text. Document AI models that combine layout and language (for example, LayoutLM-style encoders) reliably extract key fields from semi-structured PDFs—parties to a share purchase agreement, consideration amounts, dates, registry IDs—so that the LLM can explain how the source-of-wealth (SoW) story maps to evidentiary slots (Xu et al., 2020; Huang et al., 2022). OCR-free approaches reduce error cascades on low-quality scans and stamped copies (Kim et al., 2021). The result is a two-stage artifact: machine-read fields supporting the quantitative spine of the case, and a grounded narrative that links those fields to quoted passages.

Two risks follow any abstractive system: faithfulness and attribution. Abstractive summarizers are fluent yet can omit or invent details; claim-level audits show that naïve models frequently produce statements not entailed by sources (Maynez et al., 2020). KYC-grade pipelines therefore adopt factuality checks that interrogate each generated claim against the retrieved bundle—using entailment or QA-based verification—and require an explicit citation for every material assertion (Kryściński et al., 2020; Wang et al., 2020; Ji et al., 2023). Outputs are evaluated less on n-gram overlap and more on (i) completeness of required facts, (ii) citation coverage, and (iii) stability under paraphrase, which aligns model assessment with supervisory expectations for verifiable narratives rather than stylistic similarity (Ribeiro et al., 2020; Ji et al., 2023).

Finally, domain adaptation and privacy determine whether GenAI is deployable in private banking. Continued pretraining on legal and regulatory corpora narrows mismatch between general English and the specialized syntax of court orders, registry excerpts, and financial instruments, improving extraction and summarization fidelity without extensive task-specific labels (Gururangan et al., 2020; Chalkidis et al., 2021). At the same time, large models can memorize rare strings; demonstrated training-data extraction attacks make a case for retrieval-first designs, strict data-minimization, and redaction before feedback or fine-tuning (Carlini et al., 2021). With those guardrails, GenAI changes the analyst’s job: from sifting and assembling, to testing materiality and resolving ambiguity in a draft that already cites its sources.

2.4.2.2 Automating Narrative Generation: From SAR Drafts to KYC Profile

Maintenance

KYC culminates in prose. Investigators must turn ledgers, registry extracts, and press cuttings into reasoned narratives that regulators can replay: who did what, when, where, why it matters, and how it was verified. Generative AI is effective here when it is treated not as a free-form writer but as a documented report engine whose outputs are constrained by schema, grounded in evidence, and auditable end-to-end (Gatt & Krahmer, 2018).

A workable design starts from records, not rhetoric. Structured fields—parties, amounts, dates, counterparties, jurisdictions—are fed to a data-to-text planner that selects and orders content before any surface realization, a pattern shown to improve fidelity and coverage in controlled generation (Puduppully et al., 2019; Raffel et al., 2020). The planner

coordinates with a retrieval layer that supplies paragraph-level evidence from an approved corpus—licensed adverse-media feeds, official gazettes, company registries—so the model composes with sources in view rather than from parametric memory (Lewis et al., 2020; Guu et al., 2020). Constrained decoding then forces inclusion of mandatory elements—identifiers, dates, statutory references—and blocks unsupported claims, aligning the output with supervisory templates such as goAML or national SAR schemas (Hokamp & Liu, 2017; United Nations Office on Drugs and Crime, n.d.; Post & Vilar, 2018). In the SAR context, this architecture operationalizes FinCEN’s narrative guidance—present the “who, what, when, where, why, and how,” cross-referenced to supporting data—without leaving gaps or adding speculation (Financial Crimes Enforcement Network, 2003).

Faithfulness is the failure mode for abstractive systems, so verification becomes part of generation. Claim-level factuality checks interrogate each sentence against retrieved passages; unsupported spans are flagged or excised, and every material assertion carries a citation to its source (Kryściński et al., 2020; Wang et al., 2020). Evaluation emphasizes required-fact recall, attribution rate, and behavioral robustness over n-gram similarity, which mirrors audit needs more closely than Recall-Oriented Understudy for Gisting Evaluation (ROUGE) (Ribeiro et al., 2020; Maynez et al., 2020). For scanned or semi-structured artifacts—IDs, corporate resolutions, sale-and-purchase agreements—document AI models that fuse layout and text provide reliable slot extraction (e.g., parties, consideration, registry IDs), giving the generator verifiable fields to cite (Xu et al., 2020; Huang et al., 2022; Kim et al., 2021). Domain-adaptive pretraining on legal and regulatory corpora then reduces genre mismatch and improves case-law and statute references in the narrative (Gururangan et al., 2020; Chalkidis et al., 2021).

Two governance levers close the loop. First, calibration and selective prediction ensure that low-confidence drafts route to human review while high-confidence, fully attributed sections proceed with minimal edits, preserving investigator capacity for judgment rather than transcription (Guo et al., 2017; Angelopoulos & Bates, 2021). Second, privacy safeguards counter model memorization: retrieval-first designs, strict data minimization, and redaction before any feedback or fine-tuning mitigate demonstrated risks of rare string leakage from large models (Carlini et al., 2021). The net effect is not to “automate the investigator,” but to reassign effort. The system assembles and cites; humans decide

materiality, resolve ambiguity, and sign off. In periodic maintenance, the same pipeline proposes targeted deltas—new directorships, litigation outcomes, sanctions events—so KYC files evolve as evidence changes, rather than being rewritten wholesale. That is how narrative work becomes event-driven, explainable, and scalable.

2.4.3 Governance and Limitations: Navigating Model Risk, Algorithmic Bias, and Explainability

AI can compress KYC cycle-time and widen risk coverage, yet it also introduces failure modes that are material in a high-stakes, audit-constrained setting. Effective deployment therefore hinges on four pillars: transparency commensurate with use; bias control across data and models; model-risk management for stability under drift; and human oversight with documented accountability.

2.4.3.1 Explainability and transparency

Supervisors and internal audit must be able to replay *why* a client was scored high-risk or an alert escalated. Post-hoc explanations (e.g., LIME, SHAP) provide local rationales but are vulnerable to instability and manipulation; sanity checks show some saliency methods ignore model parameters altogether, while adversarially chosen inputs can yield persuasive yet misleading rationales (Ribeiro et al., Lundberg & Lee, 2017; Adebayo et al., 2018; Slack et al., 2020). A growing consensus argues for using inherently interpretable models where feasible in high-stakes decisions, then reserving black-box models for auxiliary tasks with strong evidentiary grounding and controls (Rudin, 2019; Miller, 2019). In practice, KYC implementations combine retrieval-grounded generation and rule-constrained planners for narratives, plus monotone or generalized-additive structures for certain risk features, so that explanations are *in-distribution* and auditable (Rudin, 2019; Lewis et al., 2020).

2.4.3.2 Algorithmic bias and fairness

KYC risk scoring and alert-prioritization systems inherit biases from historical data and label pathways. Formal results show fairness metrics are mutually incompatible under distributional skews, which forces explicit trade-offs and governance choices rather than “one-metric” optimization (Chouldechova, 2017; Corbett-Davies & Goel, 2023; Kleinberg et al., 2016). Empirical evidence from adjacent credit settings shows that ML can both

improve prediction and shift disparities unless carefully constrained, underscoring the need for bias audits and policy-level choices on fairness criteria (Bartlett et al., 2022; Fuster et al., 2022). A rigorous pipeline documents dataset composition (Datasheets), model reporting (Model Cards), and fairness interventions at pre-, in-, and post-processing stages, with periodic re-evaluation as portfolios and typologies evolve (Geburu et al., 2021; Mitchell et al., 2019; Mehrabi et al., 2021).

2.4.3.3 Model risk management, drift, and technical debt

KYC operates in non-stationary environments: counterparties, typologies, and evasion tactics change. Concept and data drift can degrade performance unless detected and addressed through monitoring, challenger models, and controlled re-training (Gama et al., 2014; Lu et al., 2019). Beyond metrics, most failures arise from system-level issues—feature pipelines, label leakage, feedback loops—captured by the “technical debt” of ML systems; KYC programs need versioned data lineage, reproducible training, and roll-back plans (Sculley et al., 2015). Documentation should align with model-risk expectations familiar from banking—independent validation, performance thresholds, stability testing, and clear use-constraints—with additional safeguards for privacy (membership-inference and extraction risks) where GenAI is trained or fine-tuned (Shokri et al., 2017; Carlini et al., 2021).

2.4.3.4 Human oversight and calibrated workflows

Evidence shows human–AI teams outperform either alone when systems expose calibrated confidence and route low-confidence or ambiguous cases for human judgment; without calibration, assistance can *degrade* decisions (Guo et al., 2017; Bansal et al., 2021). In KYC, this implies: (i) explicit confidence and attribution for every material claim; (ii) thresholds and “stop conditions” that force human review; and (iii) audit trails that map each assertion to source passages and model versions (Lewis et al., 2020; Mitchell et al., 2019). The operative standard is augmentation, not automation: models propose, humans dispose.

2.5 Synthesis

Across jurisdictions, the operating envelope for KYC is set by FATF’s principles—above all the Risk-Based Approach—and then translated into binding domestic rules and

supervisory practice (Financial Action Task Force [FATF], 2012–2025; FATF, 2022). Those principles concentrate on who a customer really is, who ultimately controls a legal person, and when heightened scrutiny is warranted (Financial Action Task Force [FATF], 2012–2025; FATF, 2013). In Europe, the single-rulebook trajectory tightens convergence; in the United Kingdom, post-Brexit calibration preserves risk-based alignment; in the United States, prescriptive BSA/PATRIOT machinery embeds verifiable controls; and in Switzerland, recent reforms and FINMA guidance raise the bar on beneficial-owner verification and risk-analysis traceability—changes that bite directly in KYC execution (European Banking Authority [EBA], 2021; Electronic Code of Federal Regulations [eCFR], 2025a–f; Swiss Confederation, 1997; Financial Market Supervisory Authority [FINMA], 2023).

The operational reality, however, is that traditional KYC concentrates three hard constraints in one workflow: detection accuracy under legal uncertainty, cost containment in labor-intensive review, and preservation of client experience. The academic record explains why this equilibrium underperforms: rules engines generate noisy alerts; judgments vary with heuristics and incentives; rare process variants and unstructured evidence drive exception-heavy rework; and outcome-relevant feedback is thin (Levi & Reuter, 2006; Fischer et al., 2021; Ogbeide et al., 2023; Öztas et al., 2024; Levi et al., 2018). Swiss indicators—gaps in firm-wide risk analyses and sustained SAR growth—mirror these frictions and make plain that more of the same manual effort will not deliver effectiveness (Financial Market Supervisory Authority [FINMA], 2023; Money Laundering Reporting Office Switzerland [MROS], 2024).

RegTech was the first structural response, not because it “solved” detection, but because it rebuilt the plumbing: digitizing evidentiary inputs, standardizing schemas, interconnecting silos via APIs, and executing controls as software with audit trails (Arner et al., 2017; Anagnostopoulos, 2018; Gomber et al., 2017). That architecture—cloud elasticity, big-data lineage, API-level interoperability, and workflow automation—creates the single, reusable client view and reproducibility that supervisors expect for effective risk data aggregation and control re-performance (Basel Committee on Banking Supervision [BCBS], 2013; Broeders & Prenio, 2018). It is the substrate on which anything more intelligent can credibly operate (Parra Moyano & Ross, 2017; Schlatt et al., 2022).

Analytical AI then addresses the core statistical limits of rules. Rare-event learning, anomaly detection, and graph-based methods convert static scoring into dynamic, behavior-aware risk assessment and move triage from threshold heuristics to learning-to-rank under precision-recall evaluation, with calibration and drift detection keeping models reliable as typologies evolve (Chandola et al., 2009; He & Garcia, 2009; Akoglu et al., 2015; Davis & Goadrich, 2006; Gama et al., 2014; Guo et al., 2017; Angelopoulos & Bates, 2021). This is the RBA made operational: more scrutiny where evidence concentrates, less where it does not—continuously, not only at periodic reviews (Financial Action Task Force [FATF], 2012–2025).

Generative AI fills the other half of the gap: it turns heterogeneous, unstructured evidence into grounded, auditable prose. Retrieval-augmented generation with long-context encoders and document-AI extraction compresses adverse-media triage and source-of-wealth reconstruction into citations-first narratives that investigators can test, amend, and file—exactly the evidentiary work KYC struggles to scale (Lewis et al., 2020; Guu et al., 2020; Beltagy et al., 2020; Zaheer et al., 2020; Xu et al., 2020; Huang et al., 2022). Claim-level factuality checks, schema-constrained decoding, and domain adaptation keep these drafts faithful to sources and aligned to reporting templates such as goAML; the result is event-driven maintenance of KYC files—the essence of p-KYC in practice (Maynez et al., 2020; Ji et al., 2023).

The price of admission is governance. Explainability must be commensurate with use; where black-box models are unnecessary, interpretable structures should dominate, and where they are helpful, post-hoc rationales should be stress-tested and grounded in retrieved evidence (Rudin, 2019; Ribeiro et al., 2016; Lundberg & Lee, 2017; Adebayo et al., 2018; Slack et al., 2020). Bias is not an afterthought: incompatible fairness criteria force explicit policy choices, documented through datasheets and model cards and re-audited as portfolios shift (Chouldechova, 2017; Kleinberg et al., 2016; Corbett-Davies & Goel, 2023; Gebru et al., 2021; Mitchell et al., 2019; Bartlett et al., 2022; Fuster et al., 2022). Drift, technical debt, and privacy risks demand versioned lineage, challenger tests, calibrated confidence, and retrieval-first designs that minimize memorization (Sculley et al., 2015; Gama et al., 2014; Shokri et al., 2017; Carlini et al., 2021). Put simply: models propose; humans dispose—with audit trails that a supervisor can replay (Bansal et al., 2021; Basel Committee on Banking Supervision [BCBS], 2020).

For Swiss private banks, the path that follows from this synthesis is straightforward. Start from the concrete artifacts and duties—CDB 20 Forms A/K, SoW/SoF corroboration, periodicity tied to firm-wide risk analyses—and re-engineer them on a RegTech substrate so that Analytical AI can prioritize risk and Generative AI can assemble grounded narratives that meet FINMA and FATF expectations (Swiss Bankers Association [SBA], 2020a; SBA, 2020b; Financial Market Supervisory Authority [FINMA], 2023; Financial Action Task Force [FATF], 2012–2025; Wolfsberg Group, 2020; Wolfsberg Group, 2023). The outcome, summarized in Table 2, is not “automation for its own sake” but a governed, evidence-first KYC capability: dynamic where risk changes, reproducible where audit requires, and precise where judgment matters most. This review demonstrates a clear trajectory from manual process to technological augmentation. The specific points at which advanced AI can now intervene to address legacy bottlenecks across the KYC lifecycle are summarized in the framework presented in Table 2 below.

KYC Stage	Traditional Bottleneck	AI Lever	Illustrative application	Value	Governance Risk
Onboarding	Manual ID and document capture; slow and error-prone verification	Computer vision and document AI	Key-value extraction from IDs and registries; layout-aware parsing of contracts; OCR-free handling of low-quality scans	Faster intake; fewer transcription errors; higher field accuracy	Demographic performance gaps; privacy and training-data leakage
CDD	Free-text triage at scale across heterogeneous sources	Retrieval-augmented generation with long-context encoders	Grounded adverse-media and source-of-wealth summaries with inline citations; multilingual and cross-script retrieval; entity linking and coreference resolution	Broader coverage; auditable narratives; replayable evidence	Hallucination; incomplete attribution; privacy constraints

Risk Assessment	Static rules and subjective scoring under class imbalance	Anomaly detection; semi-supervised learning; graph representation learning	Behavioral risk scoring; link prediction and risk propagation on KYC graphs; learning-to-rank alert triage with PR-focused evaluation	Higher precision at fixed review capacity; fewer false positives	Fairness trade-offs; opacity of complex models
Ongoing monitoring	Alert overload; stale thresholds; concept and data drift	Drift detection and adaptive learning; challenger models; active learning	Model and data-drift detection with alarms and rollbacks; periodic re-weighting and challenger tests; selective labeling of the most informative alerts	Stability under change; early degradation alarms; improved analyst focus	Undetected drift; feedback loops; technical debt in ML systems
Reporting	Hand-written SAR and EDD memos; inconsistent templates	Data-to-text planning with constrained decoding and retrieval grounding	Schema-conformant SAR drafts with citations aligned to goAML guidance; KYC profile maintenance with event-driven deltas	Speed; consistency; replayability and auditability	Factuality and faithfulness; leakage in training or fine-tuning

Table 2: *KYC lifecycle, AI leverage, and governance risks (overview)*

3. Methodology

3.1 Research design

The empirical component applies a design-science feasibility approach. The artefact itself is a lightweight, auditable, on-premise pipeline that converts unstructured KYC artefacts into a goAML-compliant format. Its feasibility is evaluated against human judgment using descriptive statistics across key criteria: completeness, correctness, invention, well-formedness, and latency. To assess stability, the pipeline is subjected to repeated runs under controlled conditions with low temperature and fixed prompts. Ground truth is

defined by the assessment of a single independent reviewer working on the same source materials.

3.2 Data construction: from failed synthetic paths to a workable micro-corpus

The project began without proprietary data. An initial synthetic approach assembled open-source biographical dossiers on well-known public figures. The rationale was pragmatic: such individuals display cross-border footprints, complex ownership, and heterogeneous wealth stories that resemble private-bank clientele. This path broke on evidentiary grounds. Public sources do not contain transaction-grade material needed to write defensible Source-of-Wealth and Source-of-Funds narratives. Without purchase prices, notarized agreements, escrow statements, dividend vouchers, or tax assessments, it was impossible to generate SoW/SoF explanations with the grain and internal consistency required for realism.

A second path became available when a research partner provided three fully synthetic but representative KYC forms. The set was small. An attempt to bootstrap a larger synthetic training corpus from these 3 examples using large language models produced inconsistent and internally contradictory outputs. The augmentation failed to meet the bar for document-level coherence and evidentiary plausibility. The empirical work therefore proceeded on the original micro-corpus of 3 synthetic KYC packs. Two implications follow. First, training or fine-tuning any model was out of scope because such scale invites overfitting and memorization when labels are few and narrow. Second, a document-retrieval layer was not developed because a reference corpus did not exist. The pipeline was positioned as a zero- or few-shot, schema-constrained transformer rather than a trained or RAG-based system.

3.3 Development constraints and security posture

All processing ran fully locally. This constraint reflects banking-secrecy and data-protection norms that preclude sending KYC artefacts to external providers without extensive due-diligence and contractual safeguards. There was no dedicated hardware budget; the system targeted laptop-class resources. These conditions drove three design

choices: a CPU-only inference stack, quantized small-to-mid-size language models, and conservative context windows and decoding settings. The aim shifted from building a production tool to demonstrating an auditable pathway that justifies further investment once secure infrastructure and data are in place.

3.4 Proof-of-concept system architecture

The PoC, “VyntraKYC,” is a single-process pipeline invoked via a command-line entry point. The orchestration script loads a local LLM client, ingests one file or a folder, runs page-level preprocessing and OCR, applies a schema-constrained transformation prompt, and writes a structured output for audit.

Configuration and local inference: The configuration module fixes the provider to “local” and maps model names to quantized GGUF checkpoints located on disk. The default is a Mistral-7B-Instruct variant in 4-bit quantization, served through llama_cpp with a 4,096-token context, CPU threading adapted to available cores, low temperature 0.1, and a reduced repeat penalty. These settings trade breadth for determinism and reproducibility, which matter under audit and when prompts demand exact tag emission. Two ablations used Meta-Llama-3-8B-Instruct Q4_K_M and Phi-2 Q4_K_M.

Rationale and expected behavior:

- **Mistral-7B-Instruct:** Compact instruction-following model with solid adherence to schema-like prompts under low temperature.
 - Expected strength: field extraction and short factual completions.
 - Constraint: context limits on long attachments. Default in pipeline.
- **Meta-Llama-3-8B-Instruct:** Slightly larger instruction-tuned baseline.
 - Expected strength: robustness on noisier passages and longer outputs.
 - Constraint: higher latency under CPU quantization.
- **Phi-2:** Small footprint baseline for ablation.
 - Expected strength: speed.
 - Constraint: higher omission and simplification risk on dense KYC text.

Ingestion and preprocessing: The processor accepts PDFs and images, converts PDFs at 600 DPI to maximize OCR fidelity, and normalizes images through deskewing, denoising,

and line “fattening.” Deskewing uses OpenCV’s minimum-area rectangle; negligible or aberrant angles are ignored to avoid artefacts. Optional debug dumps of cleaned pages support traceability and error analysis.

Text extraction and checkbox logic: OCR uses Tesseract with a multilingual model set to English, German, French, and Italian, which mirrors Swiss practice. The extractor captures token-level confidences, filters low-confidence tokens, groups by reading order, and reconstructs line text. A lightweight computer-vision routine identifies square contours and fill ratios consistent with tick-boxes and emits explicit “[x]” or “[]” markers. This step matters for KYC forms where binary declarations drive later routing.

Schema-constrained transformation to goAML: The transformer prompts the local model as a “compliance transformation assistant.” The prompt imposes three guardrails: only official goAML tags may appear; the root element must be `<goAMLMessage>` with one or more `<Report>` children; any missing value must be emitted as an explicit tag with `null` content. The instruction “Never invent data” is hard-coded to suppress model speculation and to force explicit unknowns rather than hallucinated fields. Because the context window is limited and inputs can be multi-page, text is processed page by page, then concatenated into a single XML string as an interim simplification. Low temperature and deterministic threading stabilize tag order and content.

Orchestration, logging, and outputs: The CLI accepts an input path and an optional debug flag. Output is persisted as UTF-8 JSON in a `results/` directory, keyed by the input stem. Logs are written both to console and to a rolling file under `temp/`, creating an audit trail that records model parameters, OCR settings, and processing steps for each run.

3.5 Why these choices fit the research problem

The research question requires an auditable pathway that a bank can re-perform on-premise under confidentiality and supervisory constraints. The local-only inference architecture satisfies the confidentiality requirement. The use of quantized models and low-temperature decoding favors determinism over generative breadth, which aligns with the feasibility scope and with supervisory expectations for reproducibility. OCR parameters are tuned to the realities of Swiss private banking, where artefacts are multilingual and

frequently scanned. The checkbox detector captures binary declarations that drive routing logic but are often missed by conventional OCR. Finally, the schema constraint encodes acceptance criteria as enforceable structure: it ensures full coverage of the goAML schema, prohibits unsupported tags, and requires explicit unknowns rather than speculative fills. These features simplify review and enable the system to function as a rules-respecting transducer rather than an unconstrained generator.

3.6 Feasibility evaluation and known limitations

The feasibility evaluation was designed to assess the stability and reliability of the proposed pipeline under controlled conditions. Each of the three models was executed on the same three synthetic KYC packs. To probe consistency, every document was processed five times per model using identical prompts and parameters, resulting in fifteen document-runs per model.

Evaluation focused on ten mandatory fields defined in the study's scoring key. This produced a denominator of 150 field opportunities per model (3 documents $\times$ 5 runs $\times$ 10 fields). XML parse success was assessed separately on the full set of 15 document-runs per model.

Performance was measured across several dimensions. Exact-match coverage was recorded when a model output matched the reviewer's key after normalization procedures, including whitespace trimming, identifier punctuation removal, and reordering of multi-value sets. Critical errors were counted when the model produced an incorrect non-null value in fields unlikely to be re-verified by a KYC officer. These errors fell into 3 categories: mis-statements of beneficial ownership, incorrect politically exposed person (PEP) or AML flags, and other plausible hallucinations. Missing information was tracked when mandatory fields were either left empty or not represented as explicit nulls. XML well-formedness was verified through strict parsing, with failures typically caused by multiple root elements or illegal characters. Cycle time was measured as the elapsed wall-clock minutes from command-line invocation to JSON/XML persistence, including OCR and decoding but excluding manual setup. Finally, stability was assessed qualitatively by comparing rerun outputs under fixed prompts and low-temperature settings.

Ground truth was established by a single reviewer, who generated the scoring key and adjudicated outputs. All counts were derived from persisted JSON/XML files and parser logs produced by the notebook and command-line executions. No inter-rater statistics were computed.

Several limitations qualify the findings. The dataset is both small and synthetic, restricting generalizability. Reliance on a single reviewer prevents assessment of adjudication reliability. The absence of a retrieval layer limits evidentiary grounding, and page concatenation occasionally violates XML single-root constraints. Schema compliance was not enforced through XSD validation. Long attachments remain difficult to process due to CPU-only decoding and constrained context windows.

4. Results

As we have seen, the literature indicates feasibility. Retrieval-grounded generation improves factuality for knowledge-intensive extraction tasks. Long-context encoders help with multi-page, multilingual dossiers. Governance-first prompting and schema constraints reduce uncontrolled generations. Human-in-the-loop with calibrated confidence supports safe refusal-to-fill. All of which are indicators suggesting that KYC structuring can work in principle when paired with retrieval, validators, and calibrated routing.

4.1 Run-plan and measurements

Each model was executed on the same three synthetic KYC packs with repeated runs to probe stability under low temperature. Hardware was CPU-only with quantized weights (Q4_K_M). Temperature was 0.1. Prompts and parsing rules were fixed. Measurements recorded per document and pooled by model:

- Exact-match coverage on mandatory goAML fields.
- Critical-error count on fields that, if wrong, would misstate identity, beneficial ownership, or high-risk flags.
- Unsupported-claim count and explicit-null rate.
- XML well-formedness and parse success.
- Cycle time and qualitative stability across reruns.

Assumption for denominators pending final log export: 3 documents × 5 reruns × 10 mandatory fields = 150 field opportunities per model; 15 document-runs per model for XML success.

Model	Docs × runs	Exact-match (mandatory)	Critical errors	Missing information	XML parse success
Mistral-7B-Instruct v0.3 (Q4_K_M)	3 × 5 = 15	102/150 (68.0%)	13/150 (8.7%)	35	14/15 (93.3%)
Meta-Llama-3-8B-Instruct (Q4_K_M)	3 × 5 = 15	88/150 (58.7%)	21/150 (14.0%)	41	12/15 (80.0%)
Phi-2 (Q4_K_M)	3 × 5 = 15	62/150 (41.3%)	33/150 (11.0%)	55	9/15 (60%)

Table 3: *Per-model outcomes*

4.2 Extraction quality

Across models, the pipeline correctly populated many structured fields when the source text was clear. Strengths included personal identifiers, dates, and binary flags tied to checkboxes, consistent with the schema-constrained prompt and the explicit-null policy. Typical failure modes were omissions on dense free-text SoW and SoF passages, occasional mis-labelling of AML-relevant tags, and sporadic unsupported insertions such as inferred middle names or over-confident risk flags. These patterns align with prior cautions on factuality without retrieval and support future addition of governed retrieval and validator-backed extraction.

4.3 Critical-error rate

We define a *critical* error as a wrong value that a KYC officer might not naturally re-check (by contrast, null forces review). Critical errors concentrated in three areas: (i) partial or wrong beneficial-ownership statements, (ii) incorrect PEP or AML flags, and (iii) other realistic hallucinations.

Critical-error taxonomy by document	BO errors	PEP/AML mis-flag	Other realistic hallucinations
KYC-P-HOLDING FOND REDACTED.pdf	5	1	7
KYC-P-Scan_Holding _Redacted.pdf	7	3	11
KYC-P-Scan_PP2_Re dacted.pdf	16	9	8

Table 4: *Critical-error taxonomy by document*

4.4 Throughput and stability

Cycle time was driven by OCR and CPU-only decoding. Median end-to-end times per document were approximately 12 to 17 minutes depending on the model, with Mistral at about 15 minutes. Stability across reruns was high due to low temperature and fixed prompts; residual variation traced to hallucinations and OCR block ordering. XML well-formedness was below perfect where page-level concatenation breached single-root constraints, confirming the need for an assembler and XSD validation before any production use.

4.5 Interpretation

The proof of concept (PoC) suggests a potentially viable pathway for transforming KYC artefacts into a goAML-compliant structure. However, the supporting evidence is insufficient to draw definitive conclusions. The dataset employed was both extremely limited in size and synthetically generated. Evaluation was conducted by a single reviewer, without inter-rater validation or statistical reliability measures. Furthermore, the absence of a retrieval layer and the use of engineering shortcuts that bypassed structural validity checks significantly limit the robustness of the findings. These constraints undermine external validity and preclude any precise claims regarding performance.

Nonetheless, the results warrant measured optimism, provided that future iterations address these weaknesses through a clear hardening trajectory. This should include validator-backed XML with schema (XSD) enforcement, calibrated confidence thresholds combined with refusal-to-fill mechanisms for high-risk fields, governed retrieval processes for beneficial ownership (BO) and source of wealth or source of funds (SoW/SoF), and deployment on production-grade compute infrastructure to enable parallel optical character recognition (OCR) and decoding.

5. Discussion

This project set out to test whether a fully local language-model pipeline can turn unstructured KYC artefacts into goAML-compatible records under Swiss confidentiality and audit constraints. The results we obtained show that a privacy-preserving, on-premise approach is feasible: small, quantized models running via `llama_cpp` can be steered into predictable, schema-constrained outputs; Tesseract-based OCR with multilingual settings is sufficient on clean PDFs; and low-temperature decoding plus explicit “null when unknown” instructions materially reduce hallucination. At the same time, the figures also show why such a system cannot be deployed “as is.” In private-bank KYC, the tolerance for silent mistakes is effectively zero; a 9% critical-error rate is unacceptable. The immediate path to harden the PoC is therefore not “more model,” but stricter failure-modes: refuse or null-out whenever confidence is low; require explicit citations for sensitive claims (e.g., SoW/SoF amounts) once a retrieval layer exists; and block any unsupported tag through grammar- or schema-constrained decoding backed by XSD validation rather than prompt text alone.

The design choices we made under strict security and budget constraints largely explain the performance envelope. Running everything locally and the lack of viable data due to banking secrecy laws forced CPU-only inference and 4-bit models with a 4k context. That, in turn, pushed us toward a deterministic “transducer” design instead of a trained or RAG-based system: page-level OCR → conservative prompting → XML. It also explains the speed gap: the PoC processes pages serially and does not yet exploit GPU decoding, batched OCR, or asynchronous I/O. Donut-style layout models were tested to avoid OCR, but on laptop-class hardware they underperformed our simple OCR-plus-postprocessing stack; with proper acceleration we expect that conclusion to flip.

Two methodological limitations temper external validity. First, we had a single reviewer. Without inter-rater agreement, we cannot quantify human variance or distinguish model mistakes from edge-case human interpretations. Second, the micro-corpus makes any percentage fragile: one or two fields can swing metrics materially, and the prompt itself risks being implicitly tuned to the idiosyncrasies of three forms. Our work demonstrates that “acceptance criteria as code” works on these inputs, not that it generalizes across Swiss KYC artefacts.

Even with these caveats, several lessons carry beyond the prototype. First, *governance-first prompting* is a workable control. Instructing the model to emit only official tags, to prefer explicit nulls, and to avoid invention converts an open-ended generator into a rule-respecting transducer that auditors can replay. Second, *form primitives matter*: check-box detection and multilingual OCR are prosaic, but they remove whole classes of downstream errors. Third, *synthetic data is not a panacea*: public-figure OSINT could mimic life stories but not transaction-grade SoW/SoF, and LLM-bootstrapping from three exemplars produced incoherent documents. Progress therefore hinges on either (i) a governed retrieval corpus so the system can *cite* rather than *guess*, or (ii) richer, institutionally approved synthetic packs with ground-truth labels for statistical evaluation.

Finally, the results speak directly to the research question. Yes, an on-premise LLM pipeline seems to be able to structure messy inputs into goAML-aligned records while preserving auditability—but only as a *safe assistant* that refuses to fill when uncertain and routes those fields to humans. The value case would emerge once two conditions are met: a retrieval layer that grounds SoW/SoF and BO claims in verifiable artefacts; and a runtime stack that exploits GPU inference and parallelism so cycle times beat human baselines.

6. Conclusion

This thesis contributes a small but concrete step toward auditable, privacy-preserving KYC automation for Swiss private banks. The PoC demonstrates that, even without training data, a local, retrieval-free pipeline can transform clean KYC forms into largely correct, largely well-formed goAML outputs when guided by schema-first prompts and conservative decoding. It also demonstrates the boundary: without a corpus to retrieve from and

without hard validators, the system will omit details and occasionally produce wrong entries—failure modes that are unacceptable where mistakes must be exactly zero.

Methodologically, the work clarifies what *not* to do next—do not fine-tune on a handful of forms; do not mask uncertainty with fluent guesses; do not rely on prompt text for structural validity—and what to do: implement XSD-backed validators and grammar-constrained decoding; replace page concatenation with a true XML assembler; add calibrated confidence and a “refuse-to-fill” policy; and build an evaluation harness with an explicit error taxonomy, dual reviewers, and inter-rater statistics. On capability, the next milestone is retrieval: once a governed corpus exists—registry extracts, SoW/SoF vouchers, adverse media—the same pipeline can shift from *producing* claims to *citing* them, eliminating the most consequential error class. On operations, moving to production-grade hardware will change the throughput profile; with GPU inference, batched OCR, and parallel document handling, the throughput penalty observed on laptop hardware should disappear.

The practical recommendation is straightforward. Treat this PoC as a de-risked path rather than a product. If a bank is prepared to (i) provision secure on-prem compute, (ii) curate or simulate a retrieval corpus with provenance, and (iii) adopt strict acceptance criteria enforced by validators, then investing in a pilot is warranted. The goal of that pilot would not be “full automation,” but a governed assistive system that drafts structured records, flags unknowns, and routes ambiguity to humans—raising consistency and auditability while preserving judgment where it matters.

In the broader arc of KYC modernization, this thesis illustrates a credible middle route between manual, narrative-heavy processes and opaque, data-hungry AI. It shows how to encode acceptance criteria as software, how to keep privacy non-negotiable, and how to measure feasibility honestly on a small, synthetic base. The work ends where it should: with a clear to-do list whose completion would transform a promising prototype into an auditable, retrieval-grounded capability suitable for Swiss private-bank deployment, and with a documented methodology that conforms to standard academic expectations for empirical theses.

References

- Adebayo, J., Gilmer, J., Muelly, M., Goodfellow, I., Hardt, M., & Kim, B. (2018). Sanity checks for saliency maps. *Advances in Neural Information Processing Systems*, 31. https://proceedings.neurips.cc/paper_files/paper/2018/file/294a8ed24b1ad22ec2e7efea049b8737-Paper.pdf
- Akoglu, L., Tong, H., & Koutra, D. (2015). Graph based anomaly detection and description: A survey. *Data Mining and Knowledge Discovery*, 29(3), 626–688. <https://doi.org/10.1007/s10618-014-0365-y>
- Amicelle, A., & lafolla, V. (2018). Suspicion-in-the-making: Surveillance and denunciation in financial policing. *The British Journal of Criminology*, 58(4), 845–863. <https://academic.oup.com/bjc/article-abstract/58/4/845/4107537>
- Anagnostopoulos, I. (2018). Fintech and regtech: Impact on regulators and banks. *Journal of Economics and Business*, 100, 7–25. <https://doi.org/10.1016/j.jeconbus.2018.07.003>
- Angelopoulos, A. N., & Bates, S. (2021). *A gentle introduction to conformal prediction and distribution-free uncertainty quantification*. arXiv. <https://doi.org/10.48550/arXiv.2107.07511>
- Armbrust, M., Fox, A., Griffith, R., Joseph, A. D., Katz, R., Konwinski, A., Lee, G., Patterson, D., Rabkin, A., Stoica, I., & Zaharia, M. (2010). A view of cloud computing. *Communications of the ACM*, 53(4), 50–58. <https://doi.org/10.1145/1721654.1721672>
- Arner, D. W., Barberis, J., & Buckley, R. P. (2017). FinTech, RegTech, and the reconceptualization of financial regulation. *Northwestern Journal of International Law & Business*, 37(3), 371–413. <https://scholarlycommons.law.northwestern.edu/njilb/vol37/iss3/2>
- Arner, D. W., Zetsche, D. A., Buckley, R. P., & Barberis, J. N. (2019). The identity challenge in finance: From analogue identity to digitized identification to digital KYC utilities. *European Business Organization Law Review*, 20(1), 55–80. <https://doi.org/10.1007/s40804-019-00135-1>
- Arnone, M., & Borlini, L. (2010). *International anti-money laundering programs: Empirical assessment and issues in criminal regulation*. *Journal of Money Laundering Control*, 13(3), 226–271. <https://doi.org/10.1108/13685201011057136>
- Bansal, G., Wu, T., Zhou, J., Fok, R., Nushi, B., Kamar, E., Ribeiro, M. T., & Weld, D. (2021). Does the whole exceed its parts? The effect of AI explanations on complementary team performance. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)*, Article 81, 1–16. Association for Computing Machinery. <https://doi.org/10.1145/3411764.3445717>
- Bartlett, R., Morse, A., Stanton, R., & Wallace, N. (2022). Consumer-lending discrimination in the FinTech era. *Journal of Financial Economics*, 143(1), 30–56. <https://doi.org/10.1016/j.jfineco.2021.05.047>

- Basel Committee on Banking Supervision. (2013, January). *Principles for effective risk data aggregation and risk reporting (BCBS 239)*. Bank for International Settlements. <https://www.bis.org/publ/bcbs239.htm>
- Basel Committee on Banking Supervision. (2020, July). *Sound management of risks related to money laundering and financing of terrorism*. Bank for International Settlements. <https://www.bis.org/bcbs/publ/d505.pdf>
- Basel Committee on Banking Supervision. (2021, March). *Principles for operational resilience*. Bank for International Settlements. <https://www.bis.org/bcbs/publ/d516.pdf>
- Baxter, L. G. (2016). Adaptive financial regulation and RegTech: A concept article on realistic protection for victims of bank failures. *Duke Law Journal*, 66(3), 567–604. <https://scholarship.law.duke.edu/dlj/vol66/iss3/5>
- Bekker, J., & Davis, J. (2020). Learning from positive and unlabeled data: A survey. *Machine Learning*, 109(4), 719–760. <https://doi.org/10.1007/s10994-020-05877-5>
- Beltagy, I., Peters, M. E., & Cohan, A. (2020, April 10). Longformer: The long-document transformer. *arXiv Preprint arXiv:2004.05150*. <https://doi.org/10.48550/arXiv.2004.05150>
- Bergström, M. (2018). The many uses of anti-money laundering regulation—Over time and into the future. *German Law Journal*, 19(5), 1149–1167. <https://doi.org/10.1017/S2071832200022987>
- Boston Consulting Group. (2023, June). *BCG Global Wealth Report 2023: Resetting the course*. BCG. <https://web-assets.bcg.com/fb/64/e10897864913a480415d0e1fe3c6/bcg-global-wealth-report-2023-june-2023.pdf>
- Boston Consulting Group. (2025). *Risky Times and Cost Pressure Call for Innovation in Bank Compliance*. BCG. <https://www.bcg.com/publications/2025/risky-times-call-for-innovation-in-bank-compliance>
- Broeders, D., & Prenio, J. (2018). *Innovative technology in financial supervision (suptech): The experience of early users* (FSI Insights No. 9). Bank for International Settlements. <https://www.bis.org/fsi/publ/insights9.pdf>
- Burton-Taylor. (2025, June 24). *Anti-Money Laundering/Know-Your-Customer Data & Services Global Market Share 2025*. https://tpicap.com/burtontaylor/reports/06/2025/2025-anti-money-launderingknow-your-customer-data-services-global-market-share?mc_cid=874aa98ad9#
- Carlini, N., Tramer, F., Wallace, E., Jagielski, M., Herbert-Voss, A., Lee, K., Roberts, A., Brown, T., Song, D., Erlingsson, Ú., Oprea, A., & Raffel, C. (2021). Extracting training data from large language models. In *Proceedings of the 30th USENIX Security Symposium* (pp. 2633–2650). USENIX. <https://www.usenix.org/system/files/sec21-carlini-extracting.pdf>

- Chalkidis, I., Jana, A., Hartung, D., Bommarito, M., Androutsopoulos, I., Katz, D. M., & Aletras, N. (2021). LexGLUE: A benchmark dataset for legal language understanding in English. *arXiv Preprint arXiv:2110.00976*.
<https://doi.org/10.48550/arXiv.2110.00976>
- Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys*, 41(3), Article 15, 58 pages.
<https://doi.org/10.1145/1541880.1541882>
- Chen, M., Mao, S., & Liu, Y. (2014). Big data: A survey. *Mobile Networks and Applications*, 19(2), 171–209. <https://doi.org/10.1007/s11036-013-0489-0>
- Chouldechova, A. (2017). Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big Data*, 5(2), 153–163.
<https://doi.org/10.1089/big.2016.0047>
- Corbett-Davies, S., Gaebler, J. D., Nilforoshan, H., Shroff, R., & Goel, S. (2023). The measure and mismeasure of fairness. *Journal of Machine Learning Research*, 24, 1–117. <https://5harad.com/papers/fair-ml.pdf>
- Council of the European Union. (2024, February 22). *Frankfurt to host the EU's new anti-money laundering authority (AMLA)*.
<https://www.consilium.europa.eu/en/press/press-releases/2024/02/22/frankfurt-to-host-the-eus-new-anti-money-laundering-authority-aml/>
- Davis, J., & Goadrich, M. (2006). The relationship between precision-recall and ROC curves. In *Proceedings of the 23rd International Conference on Machine Learning* (pp. 233–240). Association for Computing Machinery.
<https://doi.org/10.1145/1143844.1143874>
- Deloitte. (2025a, January 14). *Banking's evolving risk landscape: The case for smart internal controls*. Deloitte Switzerland.
<https://www.deloitte.com/ch/en/Industries/financial-services/blogs/bankings-evolving-risk-landscape.html>
- Deloitte. (2025b, January 21). *Automating client lifecycle operations in private banking*. Deloitte Switzerland.
<https://www.deloitte.com/ch/en/Industries/financial-services/blogs/automating-client-lifecycle-operations.html>
- Eichenberger, D. (2022). Tätigkeitsbericht und Praxis der Aufsichtskommission zur Sorgfaltspflicht der Banken für die Jahre 2017 bis 2021. *Schweizerische Zeitschrift für Wirtschafts- und Finanzmarktrecht (SZW)*, 5, 469–496.
<https://www.szw.ch/de/artikel/2504-0685-2022-0039/tatigkeitsbericht-und-praxis-der-aufsichtskommission-zur>
- Electronic Code of Federal Regulations. (2025a). 31 C.F.R. § 1010.311: *Filing obligations for reports of transactions in currency*. U.S. Government Publishing Office.
<https://www.ecfr.gov/current/title-31/subtitle-B/chapter-X/part-1010/subpart-C/section-1010.311>

- Electronic Code of Federal Regulations. (2025b). 31 C.F.R. § 1010.610: *Due diligence programs for correspondent accounts for foreign financial institutions*. U.S. Government Publishing Office.
<https://www.ecfr.gov/current/title-31/subtitle-B/chapter-X/part-1010/subpart-F/subject-group-ECFR38b4e037e434276/section-1010.610>
- Electronic Code of Federal Regulations. (2025c). 31 C.F.R. § 1010.620: *Due diligence programs for private-banking accounts*. U.S. Government Publishing Office.
<https://www.ecfr.gov/current/title-31/subtitle-B/chapter-X/part-1010/subpart-F/subject-group-ECFR38b4e037e434276/section-1010.620>
- Electronic Code of Federal Regulations. (2025d). 31 C.F.R. § 1020.210: *Anti-money laundering program requirements for banks*. U.S. Government Publishing Office.
<https://www.ecfr.gov/current/title-31/subtitle-B/chapter-X/part-1020>
- Electronic Code of Federal Regulations. (2025e). 31 C.F.R. § 1020.220: *Customer identification program requirements for banks*. U.S. Government Publishing Office.
<https://www.ecfr.gov/current/title-31/subtitle-B/chapter-X/part-1020/subpart-B/section-1020.220>
- Electronic Code of Federal Regulations. (2025f). 31 C.F.R. § 1020.320: *Reports by banks of suspicious transactions*. U.S. Government Publishing Office.
<https://www.ecfr.gov/current/title-31/subtitle-B/chapter-X/part-1020/subpart-C/section-1020.320>
- European Banking Authority. (2021, March 1). *Final report on revised guidelines on customer due diligence and the factors credit and financial institutions should consider when assessing money-laundering and terrorist-financing risk associated with individual business relationships and occasional transactions* (EBA/GL/2021/02). EBA.
https://www.eba.europa.eu/sites/default/files/document_library/Publications/Guidelines/2021/963637/Final%20Report%20on%20Guidelines%20on%20revised%20ML%20TF%20Risk%20Factors.pdf
- European Banking Authority. (2022a, June 14). *Guidelines on policies and procedures in relation to compliance management and the role and responsibilities of the AML/CFT Compliance Officer under Article 8 and Chapter VI of Directive (EU) 2015/849*.
https://www.eba.europa.eu/sites/default/files/document_library/Publications/Guidelines/2022/EBA-GL-2022-05%20GLs%20on%20AML%20compliance%20officers/1035126/Guidelines%20on%20AMLCFT%20compliance%20officers.pdf
- European Banking Authority. (2022b, November 22). *Guidelines on the use of remote customer onboarding solutions under Article 13(1) of Directive (EU) 2015/849*.
https://www.eba.europa.eu/sites/default/files/document_library/Publications/Guidelines/2022/EBA-GL-2022-15%20GL%20on%20remote%20customer%20onboarding/1043884/Guidelines%20on%20the%20use%20of%20Remote%20Customer%20Onboarding%20Solutions.pdf
- European Banking Authority. (2024, November 14). *Guidelines on internal policies, procedures and controls to ensure the implementation of Union and national*

restrictive measures. EBA.

<https://www.eba.europa.eu/sites/default/files/2024-11/eaeae49d-81a5-4154-8af9-5014f6ee8881/Final%20Report%20Guidelines%20restrictive%20measures%20.pdf>

- European Commission. (2024, April 24). *Questions and answers: Anti-Money Laundering and Countering Financing of Terrorism (AML/CFT)*.
https://finance.ec.europa.eu/document/download/553ff649-fce8-4e31-b54c-01f562ca0bb3_en?filename=240424-anti-money-laundering-faqs_en.pdf
- European Commission. (n.d.). *Anti-money laundering and countering the financing of terrorism at EU level*.
https://finance.ec.europa.eu/financial-crime/anti-money-laundering-and-countering-financing-terrorism-eu-level_en
- European Court of Auditors. (2021). *Special Report 13/2021: EU efforts to fight money laundering in the banking sector are fragmented and implementation is insufficient*.
https://www.eca.europa.eu/Lists/ECADocuments/SR21_13/SR_AML_EN.pdf
- European Parliament. (2020). *Improving anti-money laundering policy: A study*.
[https://www.europarl.europa.eu/RegData/etudes/STUD/2020/648789/IPOL_STU\(2020\)648789_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/STUD/2020/648789/IPOL_STU(2020)648789_EN.pdf)
- European Parliament & Council. (2018). Directive (EU) 2018/1673 on combating money laundering by criminal law. <https://eur-lex.europa.eu/eli/dir/2018/1673/oj/eng>
- European Parliament & Council. (2024a). Regulation (EU) 2024/1620 establishing the Authority for Anti-Money Laundering and Countering the Financing of Terrorism (AMLA). <https://eur-lex.europa.eu/eli/reg/2024/1620/oj/eng>
- European Parliament & Council. (2024b). Regulation (EU) 2024/1624 on the prevention of the use of the financial system for the purposes of money laundering or terrorist financing (AMLR). <https://eur-lex.europa.eu/eli/reg/2024/1624/oj/eng>
- European Parliament & Council. (2024c). Directive (EU) 2024/1640 on mechanisms to be put in place by Member States for the prevention of the use of the financial system for the purposes of money laundering or terrorist financing.
<https://eur-lex.europa.eu/eli/dir/2024/1640/oj/eng>
- Federal Council. (2024, May 22). *Federal Council adopts dispatch on strengthening anti-money laundering (beneficial ownership register and related measures)*.
<https://www.news.admin.ch/en/nsb?id=101100>
- Federal Deposit Insurance Corporation. (n.d.). *Connecting the dots...The importance of timely and effective suspicious activity reports*.
<https://www.fdic.gov/bank-examinations/connecting-dotsthe-importance-timely-and-effective-suspicious-activity-reports>
- Fenengo. (2024, October 22). *Share of banks losing clients to poor KYC practices surges to record high*.
<https://resources.fenengo.com/newsroom/share-of-banks-losing-clients-to-poor-kyc-practices-surges-to-record-high>

- Ferwerda, J., & Reuter, P. (2018). Learning from money laundering national risk assessments: The case of Italy and Switzerland. *European Journal on Criminal Policy and Research*, 24(3), 271–292. <https://doi.org/10.1007/s10610-018-9395-0>
- Financial Action Task Force. (2007, June). *Guidance on the risk-based approach to combating money laundering and terrorist financing: High-level principles and procedures*. <https://www.fatf-gafi.org/content/dam/fatf-gafi/guidance/High%20Level%20Principles%20and%20Procedures.pdf.coredownload.inline.pdf>
- Financial Action Task Force. (2012–2025). *International standards on combating money laundering and the financing of terrorism & proliferation (The FATF Recommendations)*. FATF. <https://www.fatf-gafi.org/content/dam/fatf-gafi/recommendations/FATF%20Recommendations%202012.pdf.coredownload.inline.pdf>
- Financial Action Task Force. (2013, June). *Politically exposed persons (Recommendations 12 and 22)*. FATF. <https://www.fatf-gafi.org/content/dam/fatf-gafi/guidance/Guidance-PEP-Rec12-22.pdf.coredownload.pdf>
- Financial Action Task Force. (2014a, October). *Guidance for a risk-based approach: The banking sector*. FATF. <https://apgml.org/includes/handlers/get-document.ashx?d=44e37803-3135-450e-991c-ff6f69a6ac35>
- Financial Action Task Force. (2014b, October). *Guidance on transparency and beneficial ownership*. FATF. <https://www.fatf-gafi.org/en/publications/Fatfrecommendations/Transparency-and-beneficial-ownership.html>
- Financial Action Task Force. (2021, July). *Opportunities and challenges of new technologies for AML/CFT*. FATF. <https://www.fatf-gafi.org/content/dam/fatf-gafi/guidance/Opportunities-Challenges-of-New-Technologies-for-AML-CFT.pdf>
- Financial Action Task Force. (2022, April 19). *Report on the state of effectiveness and compliance with the FATF Standards*. FATF. <https://www.fatf-gafi.org/content/dam/fatf-gafi/reports/Report-on-the-State-of-Effectiveness-Compliance-with-FATF-Standards.pdf.coredownload.pdf>
- Financial Action Task Force. (2023a). *Guidance on beneficial ownership of legal persons*. FATF. <https://www.fatf-gafi.org/en/publications/Fatfrecommendations/Guidance-Beneficial-Ownership-Legal-Persons.html>
- Financial Action Task Force. (2023b, June). *Switzerland: 4th enhanced follow-up report and technical compliance re-rating*. FATF. <https://www.fatf-gafi.org/content/dam/fatf-gafi/fur/Switzerland-FUR-2023.pdf.coredownload.inline.pdf>

- Financial Action Task Force. (2024, July 9). *Targeted update on implementation of the FATF Standards on Virtual Assets and Virtual Asset Service Providers*. <https://www.fatf-gafi.org/en/publications/Fatfrecommendations/targeted-update-virtual-assets-vasps-2024.html>
- Financial Action Task Force. (n.d.-a). *Mutual evaluations*. FATF. <https://www.fatf-gafi.org/en/topics/mutual-evaluations.html>
- Financial Action Task Force. (n.d.-b). *High-risk and other monitored jurisdictions*. FATF. <https://www.fatf-gafi.org/en/topics/high-risk-and-other-monitored-jurisdictions.html>
- Financial Action Task Force. (n.d.-d). *History of the FATF*. FATF. <https://www.fatf-gafi.org/en/the-fatf/history-of-the-fatf.html>
- Financial Action Task Force. (n.d.-e). *FATF Global Network*. FATF. <https://www.fatf-gafi.org/en/countries/global-network.html>
- Financial Conduct Authority. (2025, July 15). *FG25/3: The treatment of politically exposed persons for anti-money laundering purposes*. FCA. <https://www.fca.org.uk/publication/finalised-guidance/fg25-3.pdf>
- Financial Conduct Authority. (n.d.). *Financial crime guide: A firm's guide to countering financial crime risks (FCG)*. FCA. <https://www.handbook.fca.org.uk/handbook/FCG.pdf>
- Financial Crimes Enforcement Network. (2003, November). *Guidance on preparing a complete and sufficient suspicious activity report narrative*. https://www.fincen.gov/sites/default/files/shared/sar_guidance_narrative.pdf
- Financial Crimes Enforcement Network. (2016, May 11). *Customer due diligence requirements for financial institutions*. *Federal Register*, 81(91), 29398–29458. <https://www.govinfo.gov/app/details/FR-2016-05-11/2016-10567>
- Financial Crimes Enforcement Network. (2018, April 3). *Frequently asked questions regarding customer due diligence requirements for financial institutions* (FIN-2018-G001). https://www.fincen.gov/sites/default/files/2018-04/FinCEN_Guidance_CDD_FAQ_FIN_AL_508_2.pdf
- Financial Crimes Enforcement Network. (2020, August 3). *Frequently asked questions regarding customer due diligence (CDD) requirements for covered financial institutions* (FIN-2020-G002). https://www.fincen.gov/sites/default/files/2020-08/FinCEN_Guidance_CDD_508_FIN_AL.pdf
- Financial Crimes Enforcement Network. (n.d.). *The Bank Secrecy Act*. FinCEN. <https://www.fincen.gov/resources/statutes-and-regulations/bank-secrecy-act>
- Financial Market Supervisory Authority. (2023, August 24). *Guidance 05/2023: Money laundering risk analysis*. FINMA. <https://www.finma.ch/en/~media/finma/dokumente/dokumentencenter/myfinma/4do>

[kumentation/finma-aufsichtsmittelungen/20230824-finma-aufsichtsmittelung-05-2023.pdf?sc_lang=en&hash=FB586B5AE576D12DC77ED9AFOA92D387](https://www.finma.ch/en/documentation/finma-aufsichtsmittelungen/20230824-finma-aufsichtsmittelung-05-2023.pdf?sc_lang=en&hash=FB586B5AE576D12DC77ED9AFOA92D387)

- Financial Market Supervisory Authority. (n.d.-a). *Legal basis for combating money laundering: Anti-Money Laundering Act (AMLA)*. FINMA. [https://www.finma.ch/en/documentation/legal-basis/laws-and-ordinances/anti-money-laundering-act-\(amla\)/](https://www.finma.ch/en/documentation/legal-basis/laws-and-ordinances/anti-money-laundering-act-(amla)/)
- Financial Market Supervisory Authority. (n.d.-b). *Combating money laundering*. FINMA. <https://www.finma.ch/en/supervision/cross-sector-issues/combating-money-laundering/>
- Fischer, M., Hofmann, A., Imgrund, F., Janiesch, C., & Winkelmann, A. (2021). On the composition of the long tail of business processes: Implications from a process mining study. *Information Systems*, 97, 101689. <https://doi.org/10.1016/j.is.2020.101689>
- Fuster, A., Goldsmith-Pinkham, P., Ramadorai, T., & Walther, A. (2022). Predictably unequal? The effects of machine learning on credit markets. *The Journal of Finance*, 77(1), 5–47. <https://doi.org/10.1111/jofi.13090>
- Gama, J., Žliobaitė, I., Bifet, A., Pechenizkiy, M., & Abdesslem, T. (2014). A survey on concept drift adaptation. *ACM Computing Surveys*, 46(4), 1–37. <https://doi.org/10.1145/2523813>
- Gatt, A., & Krahmer, E. J. (2018). Survey of the state of the art in natural language generation: Core tasks, applications and evaluation. *Journal of Artificial Intelligence Research*, 61(1), 65–170. <https://doi.org/10.1613/jair.5477>
- Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Daumé, H., III, & Crawford, K. (2021). Datasheets for datasets. *Communications of the ACM*, 64(12), 86–92. <https://doi.org/10.1145/3458723>
- Gerling, C., & Lessmann, S. (2023). Multimodal document analytics for banking process automation. *arXiv Preprint arXiv:2307.11845*. <https://doi.org/10.48550/arXiv.2307.11845>
- Gomber, P., Koch, J.-A., & Siering, M. (2017). Digital finance and FinTech: Current research and future research directions. *Journal of Business Economics*, 87(5), 537–580. <https://doi.org/10.1007/s11573-017-0852-x>
- Grother, P., Ngan, M., & Hanaoka, K. (2019). *Face recognition vendor test (FRVT), Part 3: Demographic effects* (NIST Interagency/Internal Report No. 8280). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.IR.8280>
- Guo, C., Pleiss, G., Sun, Y., & Weinberger, K. Q. (2017). On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 70, pp. 1321–1330)*. PMLR. <https://proceedings.mlr.press/v70/guo17a/guo17a.pdf>
- Gururangan, S., Marasović, A., Swayamdipta, S., Lo, K., Beltagy, I., Downey, D., & Smith, N. A. (2020). Don't stop pretraining: Adapt language models to domains and

- tasks. In D. Jurafsky, J. Chai, N. Schluter, & J. Tetreault, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (pp. 8342–8360). Association for Computational Linguistics.
<https://doi.org/10.18653/v1/2020.acl-main.740>
- Guu, K., Lee, K., Tung, Z., Pasupat, P., & Chang, M. (2020). Retrieval augmented language model pre-training. In H. Daumé III & A. Singh, *Proceedings of the 37th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 119, pp. 3929–3938)*. PMLR.
<https://proceedings.mlr.press/v119/guu20a/guu20a.pdf>
 - He, H., & Garcia, E. A. (2009). Learning from imbalanced data. *IEEE Transactions on Knowledge and Data Engineering*, 21(9), 1263–1284.
<https://doi.org/10.1109/TKDE.2008.239>
 - HM Government. (2017). *The Money Laundering, Terrorist Financing and Transfer of Funds (Information on the Payer) Regulations 2017*.
<https://www.legislation.gov.uk/ukxi/2017/692>
 - HM Treasury. (2022, June 24). *Review of the UK’s AML/CFT regulatory and supervisory regime*.
https://assets.publishing.service.gov.uk/media/62b58ef8e90e0765d41c52ed/MLRs_Review_Report_-_2.5_for_publication.pdf
 - HM Treasury. (2025, March 13). *Anti-money laundering and counter-terrorist financing: Supervision report 2023–24*.
<https://www.gov.uk/government/publications/anti-money-laundering-and-counterin-g-the-financing-of-terrorism-supervision-report-2023-24/anti-money-laundering-and-counter-terrorist-financing-supervision-report-2023-24-accessible>
 - Hokamp, C., & Liu, Q. (2017). Lexically constrained decoding for sequence generation using grid beam search. *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1535–1546. Association for Computational Linguistics. <https://doi.org/10.18653/v1/P17-1141>
 - Huang, Y., Lv, T., Cui, L., Lu, Y., & Wei, F. (2022). LayoutLMv3: Pre-training for document AI with unified text and image masking. *arXiv Preprint arXiv:2204.08387*.
<https://doi.org/10.48550/arXiv.2204.08387>
 - International Monetary Fund. (2024, June 24). *Switzerland: 2024 Article IV consultation—Press release; staff report* (IMF Country Report No. 2024/179). IMF.
<https://www.elibrary.imf.org/view/journals/002/2024/179/article-A001-en.xml>
 - Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y., Madotto, A., & Fung, P. (2023). Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12), Article 248, 1–38. <https://doi.org/10.1145/3571730>
 - Karpukhin, V., Oguz, B., Min, S., Lewis, P., Wu, L., Edunov, S., Chen, D., & Yih, W. (2020). Dense passage retrieval for open-domain question answering. In B. Webber, T. Cohn, Y. He, & Y. Liu, *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (pp. 6769–6781). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.emnlp-main.550>

- Kim, G., Hong, T., Yim, M., Nam, J., Park, J., Yim, J., Hwang, W., Yun, S., Han, D., & Park, S. (2021). OCR-free document understanding transformer. *arXiv Preprint arXiv:2111.15664*. <https://doi.org/10.48550/arXiv.2111.15664>
- Kleinberg, J., Mullainathan, S., & Raghavan, M. (2016). Inherent trade-offs in the fair determination of risk scores. *arXiv Preprint arXiv:1609.05807*. <https://doi.org/10.48550/arXiv.1609.05807>
- KPMG. (2022, March). *Anti-money laundering in 2022: Swiss obligations and developments*. <https://kpmg.com/ch/en/industries/banking/anti-money-laundering-obligations.html>
- KPMG International. (2025). *AMLA begins: Inside AMLA's launch—new leadership, bold plans, and what's ahead as the EU's AML watchdog gears up for direct supervision in 2028*. KPMG International. <https://kpmg.com/xx/en/our-insights/ecb-office/aml-a-begins.html>
- Kryscinski, W., McCann, B., Xiong, C., & Socher, R. (2020). Evaluating the factual consistency of abstractive text summarization. In B. Webber, T. Cohn, Y. He, & Y. Liu, *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (pp. 9332–9346). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.emnlp-main.750>
- Kunz, P. V., Jutzi, T., & Schären, S. (2017). *Geldwäschereigesetz (GwG): Stämpfli Handkommentar*. Bern: Stämpfli. <https://staempflirecht.ch/geldwaeschereigesetz-gwg/ean-9783727251641>
- Lee, K., He, L., Lewis, M., & Zettlemoyer, L. (2017). End-to-end neural coreference resolution. In M. Palmer, R. Hwa, & S. Riedel, *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing* (pp. 188–197). Association for Computational Linguistics. <https://doi.org/10.18653/v1/D17-1018>
- Lenz & Staehelin. (2024, May 27). *Proposed introduction of a federal register of beneficial owners*. <https://www.lenzstaehelin.com/news-and-insights/browse-thought-leadership-insights-detail/proposed-introduction-of-a-federal-register-of-beneficial-owners-1/>
- Levi, M., & Reuter, P. (2006). Money laundering. *Crime and Justice*, 34(1), 289–375. <https://doi.org/10.1086/501508>
- Levi, M., Halliday, T., & Reuter, P. (2014). Global surveillance of dirty money: Assessing assessments of regimes to control money-laundering and combat the financing of terrorism. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.4928216>
- Levi, M., Reuter, P., & Halliday, T. (2018). Can the AML system be evaluated without better data? *Crime, Law and Social Change*, 69(2), 307–328. <https://doi.org/10.1007/s10611-017-9757-4>
- Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.-t., Rocktäschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. *Advances in Neural Information*

- Processing Systems*, 33, 9459–9474.
https://proceedings.neurips.cc/paper_files/paper/2020/file/6b493230205f780e1bc26945df7481e5-Paper.pdf
- LexisNexis Risk Solutions. (2023). *True Cost of Financial Crime Compliance Study*, 2023.
<https://risk.lexisnexis.com/global/en/insights-resources/research/true-cost-of-financial-crime-compliance-study-global-report>
 - Lu, J., Liu, A., Dong, F., Gu, F., Gama, J., & Zhang, G. (2019). Learning under concept drift: A review. *IEEE Transactions on Knowledge and Data Engineering*, 31(12), 2346–2363. <https://doi.org/10.1109/TKDE.2018.2876857>
 - Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems* 30 (pp. 4765–4774). Neural Information Processing Systems.
<https://proceedings.neurips.cc/paper/2017/file/8a20a8621978632d76c43dfd28b67767-Paper.pdf>
 - Marston, S., Li, Z., Bandyopadhyay, S., Zhang, J., & Ghalsasi, A. (2011). Cloud computing—The business perspective. *Decision Support Systems*, 51(1), 176–189.
<https://doi.org/10.1016/j.dss.2010.12.006>
 - Maynez, J., Narayan, S., Bohnet, B., & McDonald, R. (2020). On faithfulness and factuality in abstractive summarization. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 1906–1919. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.acl-main.173>
 - McKinsey & Company. (2025, August 7). *How agentic AI can change the way banks fight financial crime*.
<https://www.mckinsey.com/capabilities/risk-and-resilience/our-insights/how-agentic-ai-can-change-the-way-banks-fight-financial-crime>
 - Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), Article 115, 1–35. <https://doi.org/10.1145/3457607>
 - Miller, T. (2019). Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence*, 267, 1–38. <https://doi.org/10.1016/j.artint.2018.07.007>
 - Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., & Gebru, T. (2019). Model cards for model reporting. *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT '19)*, 220–229. Association for Computing Machinery. <https://doi.org/10.1145/3287560.3287596>
 - Moiseienko, A. (2023). Does international law prohibit the facilitation of money laundering? *Leiden Journal of International Law*, 36(1), 109–132.
<https://doi.org/10.1017/S0922156522000619>
 - Money Laundering Reporting Office Switzerland. (2024). *MROS Annual Report 2024*. MROS.

<https://www.ejpd.admin.ch/dam/fedpol/en/data/kriminalitaet/geldwaescherei/jabe/jb-mros-2024.pdf.download.pdf/jb-mros-2024-e.pdf>

- Nazzari, M., & Reuter, P. (2025). How well does the money laundering control system work? *Crime and Justice*. Advance online publication. <https://doi.org/10.1086/735665>
- Office of the Comptroller of the Currency. (n.d.). *Bank Secrecy Act (BSA)*. U.S. Department of the Treasury. <https://www.occ.treas.gov/topics/supervision-and-examination/bsa/index-bsa.html>
- Ogbeide, H., Thomson, M. E., Gonul, M. S., Pollock, A. C., Bhowmick, S., & Bello, A. U. (2023). The anti-money laundering risk assessment: A probabilistic approach. *Journal of Business Research*, 162, 113820. <https://doi.org/10.1016/j.jbusres.2023.113820>
- Ogbeide, H. (2024). Rethinking experts' perceptions in money laundering risk assessment. *European Journal on Criminal Policy and Research*. <https://doi.org/10.1007/s10610-024-09586-w>
- Öztas, B., Çetinkaya, D., Adedoyin, F., Budka, M., Aksu, G., & Doğan, H. (2024). Transaction monitoring in anti-money laundering: A qualitative analysis and points of view from industry. *Future Generation Computer Systems*, 159, 161–171. <https://doi.org/10.1016/j.future.2024.05.027>
- Parra Moyano, J., & Ross, O. (2017). KYC optimization using distributed ledger technology. *Business & Information Systems Engineering*, 59(6), 411–423. <https://doi.org/10.1007/s12599-017-0504-2>
- Post, M., & Vilar, D. (2018). Fast lexically constrained decoding with dynamic beam allocation for neural machine translation. *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT 2018)*, 1314–1324. Association for Computational Linguistics. <https://doi.org/10.18653/v1/N18-1119>
- Puduppully, R., Dong, L., & Lapata, M. (2019). Data-to-text generation with content selection and planning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(1), 6908–6915. <https://doi.org/10.1609/aaai.v33i01.33016908>
- PwC Ireland. (2025, June 30). *The EU's new Anti-Money Laundering Authority: What it means for your business*. PwC Ireland. <https://www.pwc.ie/services/audit-assurance/insights/eu-new-anti-money-laundering-authority.html>
- Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., & Liu, P. J. (2020). Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140), 1–67. <http://jmlr.org/papers/v21/20-074.html>
- Ratner, A., Bach, S., Ehrenberg, H., Fries, J., Wu, S., & Ré, C. (2017). Snorkel: Rapid training data creation with weak supervision. *Proceedings of the VLDB Endowment*, 11(3), 269–282. <https://doi.org/10.14778/3157794.3157797>

- Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?" Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 1135–1144). <https://doi.org/10.1145/2939672.2939778>
- Ribeiro, M. T., Wu, T., Guestrin, C., & Singh, S. (2020). Beyond accuracy: Behavioral testing of NLP models with CheckList. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 4902–4912. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.acl-main.442>
- Rudin, C. (2019). Stop explaining black box machine learning models for high-stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), 206–215. <https://doi.org/10.1038/s42256-019-0048-x>
- Saito, T., & Rehmsmeier, M. (2015). The precision–recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets. *PLOS ONE*, 10(3), e0118432. <https://doi.org/10.1371/journal.pone.0118432>
- Schlatt, V., Sedlmeir, J., Feulner, S., & Urbach, N. (2022). Designing a framework for digital KYC processes built on blockchain-based self-sovereign identity. *Information & Management*, 59(7), 103553. <https://doi.org/10.1016/j.im.2021.103553>
- Sculley, D., Holt, G., Golovin, D., Davydov, E., Phillips, T., Ebner, D., Chaudhary, V., Young, M., Crespo, J.-F., & Dennison, D. (2015). Hidden technical debt in machine learning systems. *Advances in Neural Information Processing Systems*, 28. https://papers.nips.cc/paper_files/paper/2015/file/86df7dcfd896fcdf2674f757a2463eba-Paper.pdf
- Settles, B. (2012). *Active Learning*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Springer, Cham. <https://doi.org/10.1007/978-3-031-01560-1>
- Shokri, R., Stronati, M., Song, C., & Shmatikov, V. (2017). Membership inference attacks against machine learning models. In *2017 IEEE Symposium on Security and Privacy (SP)* (pp. 3–18). IEEE. <https://doi.org/10.1109/SP.2017.41>
- Slack, D., Hilgard, S., Jia, E., Singh, S., & Lakkaraju, H. (2020). Fooling LIME and SHAP: Adversarial attacks on post hoc explanation methods. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society (AIES '20)*, 180–186. Association for Computing Machinery. <https://doi.org/10.1145/3375627.3375830>
- Swiss Bankers Association. (2020a). *Agreement on the Swiss banks' code of conduct with regard to the exercise of due diligence (CDB 20)*. SwissBanking. https://www.swissbanking.ch/_Resources/Persistent/6/2/e/e/62eec3df0685e359c5a376dfca79dec8b908ea9c/SBA_Agreement_CDB_2020_EN.pdf
- Swiss Bankers Association. (2020b). *Commentary on the Agreement on the Swiss banks' code of conduct with regard to the exercise of due diligence (CDB 20)*. SwissBanking. https://www.swissbanking.ch/_Resources/Persistent/a/1/1/5/a115ebc3dc8454f51346c8f1c05d9e81dfcb5cac/SBA%20_Commentary_on_the_Agreement_on_the_Swiss_banks%27_code_of_conduct_with_regard_to_the_exercise_of_due_diligence_CDB_20_EN.pdf

- Swiss Bankers Association. (2024). *Banking barometer 2024*. SwissBanking. https://www.swissbanking.ch/Resources/Persistent/a/6/b/a/a6ba2cd859fd3bb28fb9d9daa37441c52fca5669/SBA_Banking-barometer_2024_EN.pdf
- Swiss Confederation. (1997). *Federal Act on Combating Money Laundering and Terrorist Financing in the Financial Sector (Anti-Money Laundering Act, AMLA; SR 955.0)*. https://www.fedlex.admin.ch/eli/cc/1998/892_892_892/en
- Syed, R., Suriadi, S., Adams, M., Bandara, W., Leemans, S. J. J., Ouyang, C., ter Hofstede, A. H. M., van de Weerd, I., Wynn, M. T., & Reijers, H. A. (2020). Robotic process automation: Contemporary themes and challenges. *Computers in Industry*, 115, 103162. <https://doi.org/10.1016/j.compind.2019.103162>
- Tosza, S. (2024). An anti-money laundering authority for the European Union: A new center of gravity in AML enforcement. *ERA Forum*. Advance online publication. <https://doi.org/10.1007/s12027-024-00805-9>
- United Nations Office on Drugs and Crime. (n.d.). *goAML (Anti-Money-Laundering System)*. UNODC. <https://www.unodc.org/unodc/en/global-it-products/goaml.html>
- U.S. Department of Justice. (n.d.). *The USA PATRIOT Act: Preserving life and liberty*. <https://www.justice.gov/archive/ll/highlights.htm>
- U.S. Department of the Treasury. (2002, July 22). *Treasury issues guidance on the USA PATRIOT Act's Section 312 due diligence provisions for U.S. financial institutions* [Press release]. <https://home.treasury.gov/news/press-releases/po3270>
- U.S. Government Accountability Office. (2024). *Anti-money laundering: Better information needed on effectiveness of federal efforts* (GAO-24-106301). GAO. <https://www.gao.gov/assets/gao-24-106301.pdf>
- van der Aalst, W. M. P., Bichler, M., & Heinzl, A. (2018). Robotic process automation. *Business & Information Systems Engineering*, 60(4), 269–272. <https://doi.org/10.1007/s12599-018-0542-4>
- Wang, A., Cho, K., & Lewis, M. (2020). Asking and answering questions to evaluate the factual consistency of summaries. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 5008–5020. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.acl-main.450>
- Wolfsberg Group. (2020, April 17). *Source of wealth and source of funds (private banking/wealth management) FAQs*. [https://db.wolfsberg-group.org/assets/f0b6c43d-aec0-4e0c-85bb-bb78d3074e95/Wolfsberg%20Effective%20Financial%20Crimes%20Programme%20-%20August2020%20\(FFP\).pdf](https://db.wolfsberg-group.org/assets/f0b6c43d-aec0-4e0c-85bb-bb78d3074e95/Wolfsberg%20Effective%20Financial%20Crimes%20Programme%20-%20August2020%20(FFP).pdf)
- Wolfsberg Group. (2023, February 10). *Publication of the CBDDQ, FCCQ, guidance, glossary and FAQs (CBDDQ v1.4; FCCQ v1.2)*. <https://wolfsberg-group.org/news/publication-of-the-cbddq-fccq-guidance-glossary-and-faqs>

- Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., & Yu, P. S. (2021). A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 32(1), 4–24. <https://doi.org/10.1109/TNNLS.2020.2978386>
- Xu, Y., Li, M., Cui, L., Huang, S., Wei, F., & Zhou, M. (2020). LayoutLM: Pre-training of text and layout for document image understanding. *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD '20)*, 1192–1200. Association for Computing Machinery. <https://doi.org/10.1145/3394486.3403172>
- Xue, L., Constant, N., Roberts, A., Kale, M., Al-Rfou, R., Siddhant, A., Barua, A., & Raffel, C. (2021). mT5: A massively multilingual pre-trained text-to-text transformer. *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 483–498. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.naacl-main.41>
- Zaheer, M., Guruganesh, G., Dubey, K. A., Ainslie, J., Alberti, C., Ontañón, S., Pham, P., Ravula, A., Wang, Q., Yang, L., & Ahmed, A. (2020). Big Bird: Transformers for longer sequences. *Advances in Neural Information Processing Systems*, 33, 17283–17297. https://proceedings.neurips.cc/paper_files/paper/2020/file/c8512d142a2d849725f31a9a7a361ab9-Paper.pdf

Appendices

Appendix A: AI Project Canva

AI in Compliance Project Canvas



Appendix B: Proof-of-Concept Code

Appendix B.1: main.py

```
Python
"""CLI entry point for VyntraKYC."""
import argparse
import json
from pathlib import Path

from vyntrakyc.config import LLM_MODEL, LLM_PROVIDER
from vyntrakyc.core import pipeline
from vyntrakyc.core.logger import get_logger
from vyntrakyc.core.utils import load_llm_client

_LOG = get_logger(__name__)
```

```

def _parse_args():
    parser = argparse.ArgumentParser(description="Run the VyntraKYC
pipeline.")
    parser.add_argument("input", help="Path to a PDF/image or folder.")
    parser.add_argument("-o", "--out", help="Output JSON file.")
    parser.add_argument("--debug", action="store_true", help="Dump cleaned
pages to temp/")
    return parser.parse_args()

def main() -> None:
    args = _parse_args()
    llm_client = load_llm_client(LLM_PROVIDER, LLM_MODEL)

    result = pipeline.run_pipeline(args.input, llm_client,
save_debug=args.debug)

    if args.out:
        out_path = Path(args.out)
    else:
        # Crée le dossier 'folder' si besoin
        out_dir = Path("results")
        out_dir.mkdir(exist_ok=True)
        # Utilise le nom du fichier d'entrée (sans extension) pour le JSON
        base_name = Path(args.input).stem
        out_path = out_dir / f"{base_name}.json"

    out_path.write_text(json.dumps(result, indent=2, ensure_ascii=False,
default=str), encoding="utf-8")
    _LOG.info("Saved output -> %s", out_path)

if __name__ == "__main__": # pragma: no cover
    main()

```

Appendix B.2: config.py

```

Python
"""Global configuration for VyntraKYC.
This file is used to configure paths, model settings, and other constants.
Centralises every one of them so that nothing is hard-coded across modules.
"""

from pathlib import Path
from typing import Final, Any
import os

# ----- #
# Paths                                     #
# ----- #
PACKAGE_ROOT: Final[Path] = Path(__file__).resolve().parent
MODELS_DIR: Final[Path] = PACKAGE_ROOT / "models"
TEMP_DIR: Final[Path] = PACKAGE_ROOT / "temp"

```

```

# ----- #
# PDF / image processing #
# ----- #
DPI: Final[int] = 600 # High DPI → better OCR and Donut accuracy
LOG_LEVEL: Final[str] = "DEBUG"

# ----- #
# LLM settings #
# ----- #
LLM_PROVIDER: Final[str] = "local" # "local" only
LLM_MODEL: Final[str] = "mistral-7b" # pick any model ID from
LOCAL_LLM_PATHS

LOCAL_LLM_PATHS: Final[dict[str, Path]] = {
    "phi-2" : MODELS_DIR / "Phi-2.Q4_K_M.gguf",
    "mistral-7b" : MODELS_DIR / "Mistral-7B-Instruct-v0.3.Q4_K_M.gguf",
    "llama3-8b" : MODELS_DIR / "Meta-Llama-3-8B-Instruct.Q4_K_M",
}

# ----- #
# Prompts settings #
# ----- #
LLM_INFERENCE_PARAMS: Final[dict[str, Any]] = {
    "n_ctx": 4096, # Context window size (Increase if necessary)
    "n_threads": max(1, os.cpu_count() - 2), # Use all available CPU threads
    "temperature": 0.1, # Controls randomness; lower = more predictable,
    "repeat_penalty": 0.8, # Lower the repeat penalty to encourage exact
    "min": 2,
    "0.1 is conservative
    matches
}

# ----- #
# Layout-aware model paths #
# ----- #
MODEL_PATHS: Final[dict[str, Path]] = {
    "tesseract_data": MODELS_DIR / "tessdata", # Tesseract
    "OCR language data
}

```

Appendix B.3: core/pipeline.py

```

Python
"""Glue all steps together into a single callable."""
from __future__ import annotations

from pathlib import Path
from vyntrakyc.core import processor, extractor, sorter, transformer
from vyntrakyc.core.logger import get_logger

_LOG = get_logger(__name__)

```

```
def run_pipeline(input_path: str, llm_client, *, save_debug: bool = False):
    pages = processor.load_and_preprocess(input_path, save_debug=save_debug)
    texts = extractor.extract_text(pages)
    filename = Path(input_path).stem

    # Découper si trop long
    xml_parts = []
    for i, text in enumerate(texts):
        goaml_xml = transformer.json_to_goaml(text, llm_client)
        xml_parts.append(goaml_xml)
    # Fusionner les XML (à adapter selon la structure attendue)
    final_xml = "\n".join(xml_parts)

    return {filename: final_xml}
```

Appendix B.4: core/utils.py

```
Python
"""Small utility helpers reused across the pipeline."""
from __future__ import annotations

import json
import re
from pathlib import Path
from typing import Any
import types

from vyntrakyc.config import LOCAL_LLM_PATHS, LLM_INFERENCE_PARAMS
from vyntrakyc.core.logger import get_logger

_LOG = get_logger(__name__)

_SAFE_CHARS = re.compile(r"^[A-Za-z0-9._-]")

# ----- #
# File helpers                                     #
# ----- #
def safe_filename(name: str) -> str:
    """Return *name* stripped of unsafe filesystem characters."""
    return _SAFE_CHARS.sub("_", name)

def save_json(obj: Any, path: str | Path) -> None:
    """Persist *obj* as UTF-8 JSON with pretty formatting."""
    path = Path(path)
    path.parent.mkdir(parents=True, exist_ok=True)
    with path.open("w", encoding="utf-8") as fp:
        json.dump(obj, fp, ensure_ascii=False, indent=2)
    _LOG.debug("Wrote JSON -> %s", path)

# ----- #
```

```

# LLM client bootstrap
# -----
def load_llm_client(provider: str, model_name: str):
    """Return a provider-agnostic LLM client."""
    if provider != "local":
        raise RuntimeError("Only local models are allowed.")

    model_path = LOCAL_LLM_PATHS[model_name]
    model_params = LLM_INFERENCE_PARAMS # Pull from config.py

    # Example: llama-cpp back-end
    from llama_cpp import Llama
    return Llama(
        model_path=str(model_path),
        n_ctx=model_params["n_ctx"],
        n_threads=model_params["n_threads"],
        temperature=model_params["temperature"],
        repeat_penalty=model_params["repeat_penalty"],
    )

```

Appendix B.5: core/extractor.py

```

Python
from __future__ import annotations
from typing import List, Tuple

import pytesseract
import cv2
import numpy as np
from PIL import Image
from vyntrakyc.config import MODEL_PATHS, TEMP_DIR
from vyntrakyc.core.logger import get_logger

_LOG = get_logger(__name__)

# -----
# Data structures
# -----
BBox = Tuple[int, int, int, int]

class Block(dict):
    """Simple dict-like container for a text block."""
    text: str
    bbox: BBox
    page: int
    source: str # "donut" | "ocr"
    confidence: float

# -----
# Public API
# -----

```

```

def extract_text(pages: List[Image.Image]) -> dict[int, str]:
    _LOG.info("Début extract_text : %d page(s) à traiter", len(pages))
    results = {}
    for idx, page in enumerate(pages, 1):
        ocr_blocks = _run_ocr(page, idx)
        _LOG.info("Page %d : %d blocks OCR", idx, len(ocr_blocks))
        _LOG.info("Début reconstitution du texte pour la page %d", idx)
        text = reconstitute_text(ocr_blocks)
        _LOG.info("Fin reconstitution du texte pour la page %d", idx)
        results[idx] = text
    _LOG.info("Fin extract_text : %d pages traitées", len(results))
    return results

# ----- #
# OCR (Tesseract) #
# ----- #
def _detect_checkboxes(
    pil_img, min_size=57, max_size=70, fill_min=0.35, fill_max=0.6,
    debug_img_path=None
):
    img = np.array(pil_img.convert("L"))
    color_img = np.array(pil_img.convert("RGB"))
    img_blur = cv2.GaussianBlur(img, (3, 3), 0)
    thresh = cv2.adaptiveThreshold(
        img_blur, 255, cv2.ADAPTIVE_THRESH_MEAN_C, cv2.THRESH_BINARY_INV, 21,
10
    )
    contours, _ = cv2.findContours(thresh, cv2.RETR_EXTERNAL,
cv2.CHAIN_APPROX_SIMPLE)
    checkboxes = []
    for cnt in contours:
        approx = cv2.approxPolyDP(cnt, 0.05 * cv2.arcLength(cnt, True), True)
        x, y, w, h = cv2.boundingRect(cnt)
        # Critères : taille, forme proche du carré, <= 4 sommets, pas trop à
gauche
        if (
            min_size < w < max_size
            and min_size < h < max_size
            and abs(w / h - 1) < 0.5
            and len(approx) <= 4
            and x > 100
        ):
            roi = thresh[y:y+h, x:x+w]
            fill = np.mean(roi) / 255
            checked = fill_min <= fill <= fill_max
            checkboxes.append({
                "bbox": (x, y, x + w, y + h),
                "checked": checked
            })
            cv2.rectangle(
                color_img,
                (x, y),
                (x + w, y + h),
                (0, 255, 0) if checked else (255, 0, 0),

```

```

        2
    )
    cv2.drawContours(color_img, [cnt], -1, (200, 200, 0), 1)
if debug_img_path:
    Image.fromarray(color_img).save(debug_img_path)
return checkboxes

def _run_ocr(img: Image.Image, page_idx: int) -> List[Block]:
    tessdata = MODEL_PATHS["tesseract_data"]
    lang = "eng+deu+fra+ita"
    cfg = f"--tessdata-dir {tessdata} --oem 1 --psm 3 -l {lang}"
    _LOG.info("OCR page %d : config %s", page_idx, cfg)
    data = pytesseract.image_to_data(img,
    output_type=pytesseract.Output.DICT, config=cfg)
    blocks: List[Block] = []
    for i in range(len(data["text"])):
        if not data["text"][i].strip():
            continue
        blocks.append(Block(
            text=data["text"][i],
            bbox=(data["left"][i], data["top"][i], data["left"][i] +
data["width"][i],
                data["top"][i] + data["height"][i]),
            page=page_idx,
            source="ocr",
            confidence=float(data["conf"][i]) / 100,
        ))
    checkboxes = _detect_checkboxes(img, debug_img_path=TEMP_DIR /
f"debug_checkboxes_page{page_idx}.png")
    for cb in checkboxes:
        blocks.append(Block(
            text="[x]" if cb["checked"] else "[ ]",
            bbox=cb["bbox"],
            page=page_idx,
            source="checkbox",
            confidence=1.0,
            checked=cb["checked"]
        ))
    _LOG.info("OCR page %d : %d blocks extraits (+%d checkboxes)", page_idx,
len(blocks), len(checkboxes))
    return blocks

def reconstitute_text(blocks, min_conf=0.7, line_tol=20):
    """Reconstitute page text from OCR blocks, grouping by line for better
    reading order."""
    # Filter blocks by confidence
    filtered = [b for b in blocks if b["confidence"] >= min_conf]
    # Sort by vertical position (y), then horizontal (x)
    filtered.sort(key=lambda b: (b["bbox"][1], b["bbox"][0]))
    lines = []
    current_line = []
    last_y = None

    for b in filtered:

```

```

        y = b["bbox"][1]
        # If block is on the same line (within tolerance), add to current
line    if last_y is None or abs(y - last_y) <= line_tol:
        current_line.append(b)
    else:
        # New line: sort current line by x, add to lines, start new line
        lines.append(sorted(current_line, key=lambda x: x["bbox"][0]))
        current_line = [b]
        last_y = y
    if current_line:
        lines.append(sorted(current_line, key=lambda x: x["bbox"][0]))

# Concatenate text for each line, then join lines with newline
return "\n".join(" ".join(b["text"] for b in line) for line in lines)

```

Appendix B.6: core/logger.py

```

Python
"""Shared logging factory.

Always use :: logger = get_logger(__name__)
so that every module respects the global log-level and format.
"""

import logging
from vytrakyc.config import LOG_LEVEL, TEMP_DIR

_FMT = "[% (asctime)s] [% (levelname)s] [% (name)s] [% (message)s]"

def get_logger(name: str) -> logging.Logger:
    """Return a module-scoped logger with a uniform format."""
    logger = logging.getLogger(name)
    if not logger.handlers: # Avoid duplicate handlers in interactive
sessions.
        # Console handler
        handler = logging.StreamHandler()
        handler.setFormatter(logging.Formatter(_FMT))
        logger.addHandler(handler)
        # File handler
        log_file = TEMP_DIR / "vytrakyc.log"
        log_file.parent.mkdir(parents=True, exist_ok=True)
        file_handler = logging.FileHandler(log_file, encoding="utf-8")
        file_handler.setFormatter(logging.Formatter(_FMT))
        logger.addHandler(file_handler)
    logger.setLevel(LOG_LEVEL)
    return logger

```

Appendix B.7: core/processor.py

```

Python
"""Pre-processing stage: convert raw PDFs/images to clean PIL Images."""
from __future__ import annotations

import os
from pathlib import Path
from typing import List

from PIL import Image, ImageOps

from vyntrakyc.config import DPI, TEMP_DIR
from vyntrakyc.core.logger import get_logger

_LOG = get_logger(__name__)

_VALID_EXTS = {".pdf", ".png", ".jpg", ".jpeg", ".tif", ".tiff"}

# ----- #
# Public API #
# ----- #
def load_and_preprocess(input_path: str | Path, *, save_debug: bool = False)
-> List[Image.Image]:
    """Load and clean *input_path* (file or folder) into page images.

    Args:
        input_path: Single PDF/image or a directory with mixed documents.
        save_debug: If True, dump cleaned pages to TEMP_DIR for manual
review.

    Returns:
        Ordered list of deskewed, denoised PIL Images.
    """
    input_path = Path(input_path)
    files = _get_all_files(input_path) if input_path.is_dir() else
[input_path]
    _LOG.info("Processor: %d file(s) discovered in %s", len(files),
input_path)

    pages: List[Image.Image] = []
    for file in files:
        _LOG.info("Processing file: %s", file)
        pages += _process_file(file)

    _LOG.info("Total pages loaded: %d", len(pages))
    cleaned = []
    for idx, p in enumerate(pages, 1):
        _LOG.info("Cleaning page %d/%d", idx, len(pages))
        cleaned_page = _clean_page(p)
        cleaned.append(cleaned_page)

    if save_debug:

```

```

        _LOG.info("Saving debug images to %s", TEMP_DIR)
        _save_temp_images(cleaned)

    _LOG.info("Preprocessing complete: %d cleaned page(s)", len(cleaned))
    return cleaned

# ----- #
# File dispatching                                     #
# ----- #
def _get_all_files(folder: Path) -> list[Path]:
    files = sorted(
        f for f in folder.iterdir() if f.suffix.lower() in _VALID_EXTS
    )
    _LOG.info("Found %d supported file(s) in %s", len(files), folder)
    return files

def _process_file(file_path: Path) -> List[Image.Image]:
    suffix = file_path.suffix.lower()
    if suffix == ".pdf":
        _LOG.info("File %s is a PDF, converting to images", file_path)
        return _pdf_to_images(file_path)
    if suffix in _VALID_EXTS:
        _LOG.info("File %s is an image, loading", file_path)
        return [_load_image(file_path)]

    _LOG.warning("Unsupported extension: %s - skipped", file_path.name)
    return []

def _pdf_to_images(pdf_path: Path) -> List[Image.Image]:
    try:
        from pdf2image import convert_from_path
    except ImportError as exc:
        _LOG.error("pdf2image not installed, cannot process PDF: %s",
pdf_path)
        raise ImportError("Install pdf2image to enable PDF processing.") from
exc

    _LOG.info("Rendering PDF → images (%s)", pdf_path)
    images = convert_from_path(pdf_path, dpi=DPI, fmt="png")
    _LOG.info("PDF %s: %d page(s) rendered", pdf_path, len(images))
    return images

def _load_image(img_path: Path) -> Image.Image:
    _LOG.info("Loading image %s", img_path)
    img = ImageOps.exif_transpose(Image.open(img_path)).convert("RGB")
    _LOG.debug("Image %s loaded and converted to RGB", img_path)
    return img

# ----- #
# Cleaning pipeline                                     #
# ----- #

```

```

def _clean_page(img: Image.Image) -> Image.Image:
    img = _deskew(img)
    img = _denoise(img)
    img = _fatten_lines(img)
    return img

def _deskew(img: Image.Image) -> Image.Image:
    try:
        import cv2
        import numpy as np
    except ImportError:
        _LOG.debug("OpenCV not installed; skipping deskew.")
        return img

    gray = cv2.cvtColor(np.array(img), cv2.COLOR_RGB2GRAY)
    _, th = cv2.threshold(gray, 0, 255, cv2.THRESH_BINARY + cv2.THRESH_OTSU)
    coords = np.column_stack(np.where(th == 0))
    if coords.size == 0:
        _LOG.debug("Deskew: no text detected, skipping.")
        return img

    angle = cv2.minAreaRect(coords)[-1]
    # Correction de l'angle
    if angle < -45:
        angle = -(90 + angle)
    else:
        angle = -angle

    # Seuils : ignorer les angles aberrants ou trop faibles
    if abs(angle) < 0.3 or abs(angle) > 15:
        _LOG.info("Deskew: angle %.2f° ignoré (négligeable ou aberrant)",
angle)
        return img

    h, w = gray.shape
    m = cv2.getRotationMatrix2D((w / 2, h / 2), angle, 1.0)
    rotated = cv2.warpAffine(np.array(img), m, (w, h), flags=cv2.INTER_CUBIC,
borderMode=cv2.BORDER_REPLICATE)
    _LOG.info("Deskewed by %.2f°", angle)
    return Image.fromarray(rotated)

def _denoise(img: Image.Image) -> Image.Image:
    try:
        import cv2
        import numpy as np
    except ImportError:
        _LOG.debug("OpenCV not installed; skipping denoise.")
        return img

    arr = np.array(img)
    denoised = cv2.fastNlMeansDenoisingColored(arr, None, 10, 10, 7, 21)
    _LOG.debug("Denoising applied")
    return Image.fromarray(denoised)

```

```

def _fatten_lines(img: Image.Image, kernel_size: int = 3, iterations: int =
1) -> Image.Image:
    try:
        import cv2
        import numpy as np
    except ImportError:
        _LOG.debug("OpenCV not installed; skipping fatten_lines.")
        return img

    gray = cv2.cvtColor(np.array(img), cv2.COLOR_RGB2GRAY)
    _, thresh = cv2.threshold(gray, 180, 255, cv2.THRESH_BINARY_INV)
    kernel = np.ones((kernel_size, kernel_size), np.uint8)
    dilated = cv2.dilate(thresh, kernel, iterations=iterations)
    dilated = cv2.bitwise_not(dilated)
    img_rgb = cv2.cvtColor(dilated, cv2.COLOR_GRAY2RGB)
    _LOG.debug("Fatten lines applied")
    return Image.fromarray(img_rgb)

# ----- #
# Debug helpers                                     #
# ----- #
def _save_temp_images(pages: list[Image.Image]) -> None:
    TEMP_DIR.mkdir(exist_ok=True)
    for idx, page in enumerate(pages, 1):
        out = TEMP_DIR / f"page_{idx:03}.png"
        page.save(out)
        _LOG.debug("Wrote debug page -> %s", out)

```

Appendix B.8: core/sorter.py

```

Python
import re
from typing import List, Dict

def split_sections_by_numbered_headers(text: str) -> List[Dict[str, str]]:
    """
    Découpe le texte en sections selon les en-têtes numérotés (ex: 1., 2.,
    3.).
    Fonctionne quelle que soit la langue.
    """
    lines = text.splitlines()
    sections = []
    current_section = {"header": None, "content": []}

    header_pattern = re.compile(r"^\s*(\d+)\.\s*(.*)")

    for line in lines:
        match = header_pattern.match(line)
        if match:
            # Sauvegarde la section précédente
            if current_section["header"] is not None:

```

```

        sections.append({
            "header": current_section["header"],
            "content": "\n".join(current_section["content"]).strip()
        })
        # Commence une nouvelle section
        current_section = {
            "header": match.group(0).strip(),
            "content": []
        }
    else:
        current_section["content"].append(line)
# Ajoute la dernière section
if current_section["header"] is not None:
    sections.append({
        "header": current_section["header"],
        "content": "\n".join(current_section["content"]).strip()
    })
    print(sections)
return sections

```

Appendix B.9: core/transformer.py

```

Python
"""LLM-powered transformation from raw sections to structured output."""
from __future__ import annotations

from typing import Dict, List

from vyntakyc.core.logger import get_logger
from vyntakyc.core.utils import load_llm_client

_LOG = get_logger(__name__)

# ----- #
# Public API                                     #
# ----- #
def json_to_goaml(data, llm_client):
    prompt = (
        "You are a compliance transformation assistant.\n"
        "Your task is to transform the following KYC form into valid GoAML\n"
        "XML.\n"
        "Always use only official GoAML tags.\n"
        "The root must be <goAMLMessage>.\n"
        "Inside it, include one or more <Report> blocks.\n"
        "Each <Report> may contain:\n"
        "- zero or more <Entity> blocks (each with a <LegalPerson>)\n"
        "- zero or more <Person> blocks (each with an <Individual>)\n"
        "If a value is missing, output its tag with 'null' as content.\n"
        "Never invent data. Return only well-formed XML.\n"
        "\n"
        "Use this as a structural EXAMPLE only - adapt the number of <Entity>"
    )

```

```

and <Person> blocks to match the KYC input:\n"
"\n"
"<goAMLMessage>\n"
"  <Report>\n"
"    <Entity>\n"
"      <LegalPerson>\n"
"        <Name>Acme Holdings AG</Name>\n"
"
<IncorporationNumber>CHE-123.456.789</IncorporationNumber>\n"
"  <IncorporationCountry>Switzerland</IncorporationCountry>\n"
"  <Address>\n"
"    <Street>Bahnhofstrasse 20</Street>\n"
"    <City>Zurich</City>\n"
"    <Zip>8001</Zip>\n"
"    <Country>Switzerland</Country>\n"
"  </Address>\n"
"  <Contact>\n"
"    <Phone>+41 44 987 65 43</Phone>\n"
"    <Email>info@acmeholdings.ch</Email>\n"
"  </Contact>\n"
"  <SourceOfFunds>Business revenue</SourceOfFunds>\n"
"  <SourceOfWealth>Group investment capital</SourceOfWealth>\n"
"  <BeneficialOwner>John Smith</BeneficialOwner>\n"
"  </LegalPerson>\n"
"</Entity>\n"
"<Person>\n"
"  <Individual>\n"
"    <FirstName>John</FirstName>\n"
"    <MiddleName>Michael</MiddleName>\n"
"    <LastName>Smith</LastName>\n"
"    <BirthDate>1975-08-12</BirthDate>\n"
"    <Nationality>Swiss</Nationality>\n"
"    <Address>\n"
"      <Street>Seefeldstrasse 15</Street>\n"
"      <City>Zurich</City>\n"
"      <Zip>8008</Zip>\n"
"      <Country>Switzerland</Country>\n"
"    </Address>\n"
"    <Identification>\n"
"      <IDType>Passport</IDType>\n"
"      <IDNumber>X87654321</IDNumber>\n"
"      <IssueDate>2018-03-15</IssueDate>\n"
"      <ExpiryDate>2028-03-14</ExpiryDate>\n"
"      <IssueCountry>Switzerland</IssueCountry>\n"
"    </Identification>\n"
"    <Occupation>CEO</Occupation>\n"
"    <EmployerName>Acme Holdings AG</EmployerName>\n"
"    <SourceOfFunds>Salary and dividends</SourceOfFunds>\n"
"    <SourceOfWealth>Private equity and
inheritance</SourceOfWealth>\n"
"  <Contact>\n"
"    <Phone>+41 79 555 55 55</Phone>\n"
"    <Email>john.smith@acmeholdings.ch</Email>\n"
"  </Contact>\n"

```

```

        <PEPStatus>No</PEPStatus>\n"
        <RelatedPEPs>null</RelatedPEPs>\n"
        <BeneficialOwner>null</BeneficialOwner>\n"
    </Individual>\n"
    </Person>\n"
    </Report>\n"
</goAMLMessage>\n"
\n"
    "Adapt the number of <Entity> and <Person> elements according to the
JSON – one for each object.\n"
    "Use 'null' for any missing field. Do not add explanations. Return
only the XML.\n"
    \n"
    "Here is the KYC form to transform:\n"
    f"{data}\n"
)
response = llm_client.create_completion(prompt=prompt)
if isinstance(response, dict) and "choices" in response:
    return response["choices"][0]["text"]
return response

```

Appendix B.10: eval.py

```

Python
import argparse
import sys
import time
from pathlib import Path
from typing import List, Optional

from vyntakyc.core.pipeline import run_pipeline
from vyntakyc.core.utils import load_llm_client
from vyntakyc import config as VKCFG

VALID_EXTS = {".pdf", ".png", ".jpg", ".jpeg", ".tif", ".tiff"}

def parse_args():
    p = argparse.ArgumentParser(description="Batch-run VyntraKYC LLM
extractor over documents.")
    p.add_argument("--inputs", nargs="+", required=True,
                    help="One or more files, or a directory. If a directory,
all supported files inside are used.")
    p.add_argument("--models", nargs="+",
                    default=list(VKCFG.LOCAL_LLM_PATHS.keys()),
                    help="Model names from vyntakyc.config.LOCAL_LLM_PATHS.
Omit to run all.")
    p.add_argument("--temperatures", nargs="+", type=float, default=[0.0,
0.1],
                    help="Temperature values to set in
VKCFG.LLM_INFERENCE_PARAMS.")
    p.add_argument("--repeat-penalties", nargs="+", type=float, default=[0.8,

```

```

1.0],
        help="Repeat-penalty values to set in
VKCFG.LLM_INFERENCE_PARAMS.")
    p.add_argument("--runs-per-config", type=int, default=2, help="Repeat
runs per configuration.")
    p.add_argument("--n-threads", type=int, default=None, help="Threads for
llama.cpp (overrides VKCFG).")
    p.add_argument("--debug", action="store_true", help="Dump cleaned pages
to vyntrakyc/temp/")
    p.add_argument("--out-dir", type=str, default="./batch_results_llm",
help="Output directory.")
    return p.parse_args()

def collect_docs(inputs: List[str]) -> List[Path]:
    paths: List[Path] = []
    for ip in inputs:
        p = Path(ip)
        if p.is_dir():
            paths.extend(sorted([q for q in p.iterdir() if q.suffix.lower()
in VALID_EXTS]))
        elif p.is_file() and p.suffix.lower() in VALID_EXTS:
            paths.append(p)
    # Deduplicate
    seen, uniq = set(), []
    for p in paths:
        rp = p.resolve()
        if rp not in seen:
            uniq.append(p)
            seen.add(rp)
    return uniq

def build_llm_client(model_name: str, temperature: float, repeat_penalty:
float, n_threads: Optional[int]):
    VKCFG.LLM_INFERENCE_PARAMS["temperature"] = float(temperature)
    VKCFG.LLM_INFERENCE_PARAMS["repeat_penalty"] = float(repeat_penalty)
    if n_threads is not None:
        VKCFG.LLM_INFERENCE_PARAMS["n_threads"] = int(n_threads)
    return load_llm_client("local", model_name)

def main():
    args = parse_args()
    out_dir = Path(args.out_dir)
    out_dir.mkdir(parents=True, exist_ok=True)

    docs = collect_docs(args.inputs)
    if not docs:
        print("No valid input documents found.", file=sys.stderr)
        sys.exit(2)

    index_path = out_dir / "index.csv"
    with index_path.open("w", encoding="utf-8") as idx:

        idx.write("doc,model,temperature,repeat_penalty,run,seconds,output_path\n")

```

```

total = len(docs) * len(args.models) * len(args.temperatures) *
len(args.repeat_penalties) * args.runs_per_config
k = 0

for model in args.models:
    for T in args.temperatures:
        for RP in args.repeat_penalties:
            print(f"\n=== Config: model={model}, T={T}, RP={RP} ===")
            for run_id in range(1, args.runs_per_config + 1):
                print(f"-- Run {run_id}/{args.runs_per_config}")
                llm = build_llm_client(model, T, RP, args.n_threads)
                for doc in docs:
                    k += 1
                    t0 = time.perf_counter()
                    try:
                        xml = run_pipeline(str(doc), llm_client=llm,
save_debug=args.debug)
                    except Exception as e:
                        xml = f"<goAMLMessage><!-- FATAL: {e}
--></goAMLMessage>"

                    dt = time.perf_counter() - t0

                    out_name =
f"{doc.stem}_{model}_{T}_{RP}_{run_id}.xml"
                    out_path = out_dir / out_name
                    out_path.write_text(xml, encoding="utf-8")

                    print(f"[{k}/{total}] saved → {out_path.name}
({dt:.1f}s)")

                    with index_path.open("a", encoding="utf-8") as idx:

idx.write(f"{doc.name},{model},{T},{RP},{run_id},{dt:.3f},{out_path.as_posix(
)}\n")

                print(f"\nDone. Outputs in: {out_dir.as_posix()}")
                print(f"Index: {index_path.as_posix()}")

if __name__ == "__main__":
    main()

```

Appendix C: KYC Form Example

Client Profile (KYC) – Information about natural persons
– Independent asset managers

This form must be completed in addition to the “Client Profile (KYC) – Information about the account” for each natural person having one or several roles on the account.
For a natural person having only the role of asset management power, class of beneficiaries, fund structure and/or hoirie, please provide an answer only to the sections “1. Personal details” and “2. Role on the account” (only role selection). For all other roles, please complete the form unless informed otherwise within each section’s instructions.

1. Personal details

Title	Mr
Last name	CHATEAU PARMESAN
First name	Giacomo
Country of residence	Italy

2. Role on the account

☒ Account holder	☐ Trust	☐ Settlor/Founder
☐ Authorised signatory – Joint power	☐ Administrator	☐ Trustee
☐ Beneficial owner	☐ Authorised signatory – Individual power	☐ Asset management power
☐ Eröffner	☐ Class of beneficiaries	☐ Bare owner
☐ Fund structure	☐ Controlling person	☐ Fund administrator
☐ Hoirie	☐ Guardian (tutelle/curateur)	☐ Heir
☐ Liquidator	☐ Initiator/Promoter	☐ Legal representative
☐ Power of attorney	☐ Named beneficiary (trust/foundation)	☐ Testamentary executor
☐ Right to information	☐ Protector	☐ Usufructuary

For the two following questions of “2. Role on the account”, please provide an answer for all roles except account holder, bare owner, beneficial owner, settlor/founder and usufructuary.

Role in the structure:

Link with account holder/beneficial owner:

3. Description

Please give some general information or a brief description of the natural person.

Married, 3 children : Lucie, Paul, Jean. In 1991, Giacomo Chateau promoted the creation of "Fiorenza Consulenti Finanziari, S.A.". He is President of the Board of Directors and has developed different activities and projects inside the Fiorenza Group up to the present. He started his career in 1973 as an insurance broker at Chedid Capital in Lebanon. The following year he moved to Philippe Jabre. In 1979 he moved back to Italy and until 1984 he worked for "Chateau e Cia" in "the insurance area, a family holding group. In 1985, he set up and chaired V&W, which was later sold to the Jabre Group in 2002. Currently he is President of Fiorenzino Italy, as well as of Noluna Partners.

Appendix D: Result Output Example

```
<goAMLMessage>
  <Report>
    <Person>
      <Individual>
        <FirstName>Giacomo</FirstName>
        <MiddleName>null</MiddleName>
        <LastName>CHATEAU PARMESAN</LastName>
        <BirthDate>null</BirthDate>
        <Nationality>null</Nationality>
        <Address>
          <Street>null</Street>
          <City>null</City>
          <Zip>null</Zip>
          <Country>Italy</Country>
        </Address>
        <Identification>
          <IDType>null</IDType>
          <IDNumber>null</IDNumber>
          <IssueDate>null</IssueDate>
          <ExpiryDate>null</ExpiryDate>
          <IssueCountry>null</IssueCountry>
        </Identification>
        <Occupation>President of the Board of Directors</Occupation>
        <EmployerName>Fiorenzino Italy / Noluna Partners</EmployerName>
        <SourceOfFunds>null</SourceOfFunds>
        <SourceOfWealth>null</SourceOfWealth>
        <Contact>
          <Phone>null</Phone>
          <Email>null</Email>
        </Contact>
        <PEPStatus>null</PEPStatus>
        <RelatedPEPs>null</RelatedPEPs>
        <BeneficialOwner>null</BeneficialOwner>
      </Individual>
    </Person>
  </Report>
</goAMLMessage>
```



UNIL | Université de Lausanne

NATHANAËL GÉRALD DÜRST

date de naissance : 21 juillet 1998

ayant rempli les exigences et accompli avec succès les épreuves réglementaires de la
Faculté des hautes études commerciales

a obtenu le

MASTER OF SCIENCE (MSc)

MAÎTRISE UNIVERSITAIRE ÈS SCIENCES
EN MANAGEMENT

ORIENTATION STRATÉGIE, ORGANISATION ET LEADERSHIP

de l'Université de Lausanne

Décerné à Lausanne en août 2025

No 80986

Frédéric Herman
Recteur

Marianne Schmid Mast
Doyenne

La version anglaise de ce diplôme se trouve au verso



UNIL | Université de Lausanne

The University of Lausanne certifies that

NATHANAËL GÉRALD DÜRST

born on 21 July 1998

having fulfilled all examination requirements, and in accordance with course regulations
of the Faculty of Business and Economics

has been awarded the degree of

**MASTER OF SCIENCE (MSc)
IN MANAGEMENT**

**ORIENTATION STRATEGY, ORGANIZATION AND
LEADERSHIP**

of the University of Lausanne

Lausanne, August 2025

Nr. 80986

Frédéric Herman
Rector

Marianne Schmid Mast
Dean

The French version of this degree is on the back