AI-Driven Multi-Agent Systems for Automated Regulatory Analysis of Crypto Projects

Mario Trerotola
Politecnico di Torino
Turin, Italy
mario.trerotola@polito.it

Davide Calvaresi
HES-SO Valais-Wallis
Sion, Switzerland
davide.calvaresi@hevs.ch

Abstract*—The rapid expansion of the crypto-asset market is stretching supervisory authorities and institutional investors, and although MiCAR mandates comprehensive disclosures, the sheer volume and heterogeneity of project whitepapers push regulators toward a largely passive, alert-driven form of oversight that leaves little room for timely, preventive intervention. To address this challenge, this article presents an AI-driven multi-agent system (MAS) that automates Markets in Crypto-Assets Regulation (MiCAR) due diligence by synergizing large language models (LLMs) with specialized software agents. This MAS features a supervisory agent orchestrating specialized peers for evidence retrieval, document normalization, semantic extraction, and rule-based verification, culminating in an auditable compliance report and an associated knowledge graph. Key techniques include retrieval-augmented generation (RAG) for contextualizing legal texts, self-critique prompting to enhance LLM reliability, and containerized microservices for scalable deployment. Evaluated on a diverse corpus of public whitepapers, the system significantly reduces processing latency and expert workload, achieving compliance assessment accuracy comparable to human experts for critical MiCAR requirements. Furthermore, the system supports longitudinal monitoring by dynamically incorporating regulatory updates into its rule repository, ensuring ongoing alignment with evolving MiCAR standards. The MAS architecture's transparent division of labor enables fault isolation, parallel processing, and human-in-the-loop validation, providing superior robustness and interpretability over conventional monolithic AI models. By translating complex legal obligations into reproducible computational workflows, our approach advances Regulatory Technology (RegTech), offers actionable intelligence to market participants, and fosters a more transparent and accountable digital finance ecosystem. Future work will focus on multilingual capabilities and enhancing adversarial robustness against deceptive disclosures.

I. Introduction

The emergence of crypto-assets has ushered in a new era of digital finance, reshaping capital markets and democratizing access to innovative financial instruments. These assets, built on distributed ledger technologies, offer transformative potential for cross-border payments, decentralized applications, and financial inclusion. However, the rapid proliferation of diverse token offerings has simultaneously created significant challenges in ensuring transparency, investor protection, and market stability. Without clear standards and robust oversight, investors may face inadequate disclosure, and markets can be exposed to operational and systemic risks.

In response to these concerns, the European Union introduced the Markets in Crypto-Assets Regulation (MiCAR) to establish a coherent legal framework across its member states. MiCAR aims to ensure that issuers provide comprehensive information regarding token features, governance structures, reserve mechanisms, and risk factors. By codifying strict requirements for documentation and operational safeguards, the regulation seeks to foster market integrity while protecting retail and institutional participants. Nevertheless, implementing MiCAR effectively presents a formidable challenge: the sheer volume and heterogeneity of whitepapers and associated project documentation demand extensive scrutiny. Manual reviews are resource-intensive, prone to inconsistency, and struggle to keep pace with the dynamic nature of crypto offerings, creating a critical bottleneck for supervisory authorities and institutional investors.

To address this scalability and complexity gap, this article introduces an innovative hybrid-AI Multi-Agent System (MAS) designed for the automated regulatory analysis of crypto projects under MiCAR. Our MAS architecture combines the sophisticated semantic interpretation capabilities of Large Language Models (LLMs) with a distributed ensemble of specialized, task-oriented software agents. These agents collaboratively manage distinct stages of the due diligence process—from intelligent information retrieval and comprehensive document processing to rule-based compliance verification and auditable reporting. By decomposing the intricate regulatory assessment into a cooperative workflow of autonomous yet coordinated entities, our approach aims to deliver a scalable, transparent, and adaptable solution for continuous regulatory monitoring.

This paper details the system's modular architecture, outlining the specific roles, interactions, and LLM integration within its constituent agents, and describes the methodology employed for evaluating its performance against key MiCAR provisions. We present empirical results from applying the system to a representative corpus of crypto-asset projects, demonstrating its potential to significantly reduce processing latency and reviewer workload while achieving high accuracy in identifying compliance information. Ultimately, this work contributes to the advancement of Regulatory Technology (RegTech) by offering a robust and interpretable framework capable of delivering timely, actionable compliance insights, surfacing critical disclosure gaps, and effectively supporting human-in-the-loop validation processes, thereby fostering a safer and more accountable digital finance ecosystem.

II. Regulatory Background

The Markets in Crypto-Assets Regulation (MiCAR), formally Regulation (EU) 2023/1114 of 31 May 2023, constitutes the European Union's first comprehensive legal framework for crypto-assets. Published in the *Official Journal* on 31 May 2023 and entering into force on 29 June 2023, MiCAR establishes a phased application regime: the substantive provisions governing asset-referenced tokens (ARTs) and e-money tokens (EMTs) became applicable on 30 June 2024, while the requirements for other crypto-assets (commonly referred to as utility tokens) will apply from 30 December 2024. The primary objectives of MiCAR are to harmonise investor protection, enhance market integrity, and mitigate systemic risks arising from unregulated token offerings.

Article 3 of MiCAR delineates three principal categories of crypto-assets:

(a) **Utility tokens**, regulated under Title II, i.e. digital assets that grant the holder access to a current or future product or service on a blockchain network, without constituting a claim on the issuer;

(b) **Asset-referenced tokens (ARTs)**, regulated under Title III, i.e. tokens whose value is explicitly linked to one or more assets (such as fiat currencies, commodities, or other crypto-assets) in order to stabilise price fluctuations;

(c) **E-money tokens (EMTs)**, regulated under Title IV, i.e. digital representations of fiat currency denominated in a legally established currency, issued on a distributed ledger and intended for use as a means of payment.

MiCAR imposes baseline obligations on all issuers, including the preparation, notification, and publication of a detailed white paper; restrictions on marketing communications; and prohibitions against market abuse under Title VI. Titles II–IV then layer on specific prudential, organisational, and disclosure requirements—ranging from capital and reserve safeguards to mandatory environmental impact statements—tailored to each asset class. Table I shows selected provisions and annex requirements.

Whereas the European Banking Authority (EBA) and the European Securities and Markets Authority (ESMA) coordinate and harmonise the implementation of MiCAR at the European level, investor protection and market transparency are entrusted to the National Competent Authorities (NCAs). Accordingly, it is imperative to establish criteria that specify the provision of services within a given Member State.

III. Multi-Agent System Architecture

A *multi-agent system* (MAS)[1] is a distributed ensemble of autonomous software entities that perceive their environment, communicate, and act in pursuit of local goals while contributing to a shared global objective. Figure 1 shows the four principal agents that underpin our MiCAR compliance pipeline: *Search*, *Project Research*, *Web Crawler*, and *Compliance*. Each agent encapsulates a narrowly scoped responsibility and exchanges JSON-serialised messages via a lightweight message bus, enabling parallel execution and isolation.

A central design question was **where to embed large-language-model (LLM) capabilities** as opposed to relying on deterministic micro-services. LLMs excel at semantic interpretation and reasoning, but their latency and computational cost can bottleneck an otherwise scalable workflow. Conversely, deterministic components—such as headless browsers or rule-based filters—offer predictability and speed but lack deep contextual understanding. Our guiding principle was therefore to *reserve LLMs for tasks that demand discourse-level reasoning or legal nuance*, and to employ conventional software techniques wherever simpler pattern-matching or data transformation suffices.

Table II operationalises this principle with enriched explanations of the trade-offs that informed each choice.

This hybrid allocation reduces end-to-end latency by more than 40% compared with an all-LLM baseline (see Section III)[2] while preserving the explanatory power required by regulators (see Section VII) .

A. Initial Prototype Technology Stack

To evaluate the feasibility of our multi-agent framework, we developed an initial prototype leveraging a carefully selected set of tools and libraries. This technology stack was chosen to ensure modularity, scalability, and semantic accuracy throughout the regulatory analysis pipeline.

- **Python** & **SPADE**[1]: The core agent infrastructure was implemented in Python using the SPADE library for XMPP-based messaging. SPADE's lightweight asynchronous channels enabled robust inter-agent communication and dynamic discovery, facilitating parallel execution and fault tolerance.
- **Playwright**[2] & **ChromeDriver**[3]: A custom web crawler based on Playwright with headless ChromeDriver was developed to retrieve fully rendered HTML from target websites. By simulating realistic network conditions and filtering irrelevant paths (e.g., '/blog', '/media'), we ensured high-fidelity document capture for downstream processing.
- **GPT-4.1o via OpenAI API**: Large Language Model functionalities were accessed through GPT-4.1o. In the Search Agent, GPT-4.1o performs semantic filtering and deduplication of raw URLs; in the Compliance Agent, it executes schema mapping to the MiCAR regulatory framework, supporting explainable extractive reasoning.
- **LangChain**[4]: It has been used to define and orchestrate prompt chains and agent workflows. The framework's

[1] SPADE MAS framework documentation, covering installation, XMPP-based agent API, and examples: https://SPADE-mas.readthedocs.io/en/latest/

[2] Playwright official guide, including API reference for browser automation and network emulation: https://playwright.dev/docs/intro

[3] ChromeDriver download, setup instructions, and compatibility matrix with Chromium versions: https://chromedriver.chromium.org/

[4] LangChain documentation for prompt templates, chain composition, and memory management, with code examples: https://python.langchain.com/en/latest/

TABLE I
COMPARATIVE OVERVIEW OF SELECTED MICAR REQUIREMENTS

Category	Key regulatory requirements (MiCAR refs.)	Mandatory white-paper elements (Annex refs.)
Other crypto-assets (utility tokens)	• Issuer must be a legal person (Art. 5). • No prudential licence required (Consideration 26). • White paper notification and publication (Arts. 8–9). • 14-day withdrawal right; safeguarding of client funds (Arts. 11–14). • Marketing must be fair, clear, non-misleading (Art. 7).	• Annex I: Issuer details; tokenomics; offer terms; holder rights; technology; environmental impact; risk factors.
Asset-referenced tokens (ART)	• Prior authorisation (Arts. 15–17). • Own funds $\geq$2% (3% if significant) (Art. 35). • 1:1 reserve in high-quality assets (Art. 36). • Continuous redemption right (Art. 39). • Governance, risk and continuity frameworks (Art. 34).	• Art. 19 & Annex II: Governance; token features; offer/trading details; redemption terms; technology; reserve composition; environmental disclosure.
E-money tokens (EMT)	• Issuer must be a credit or e-money institution (Art. 48). • White paper notice $\geq$20 days before issuance (Art. 51). • Issuance and redemption at nominal value (Art. 49). • $\geq$30% of funds in segregated deposits; remainder in low-risk assets (Arts. 46–47). • No interest on tokens (Art. 50); additional oversight for significant EMTs (Arts. 54–56).	• Art. 51 & Annex III: Issuer identity; token description; offer mechanics; holder rights; insolvency procedures; technology; risk assessment; environmental impact.

TABLE II
LLM USAGE ACROSS AGENTS

Agent	LLM?	Design Rationale
Search	Yes	Emulates human investor queries and semantically filters search results, boosting precision/recall beyond keyword heuristics.
Project Research	Yes	Disambiguates project mentions in heterogeneous web pages and links them to official domains using retrieval-augmented generation.
Web Crawler	No	Pure browser automation delivers deterministic, high-throughput HTML snapshots; LLM reasoning would add latency without benefit.
Compliance	Yes	Maps free-text disclosures to MiCAR's structured schema and produces explainable verdicts demanded by auditors and regulators.

abstractions for prompt templates and chain composition facilitated the seamless integration of LLM API calls across the various agents. The prompting strategy adopted leverages template-based constructs and chaining logic, ensuring that context-specific tasks, such as the refinement of search results and compliance verification, are handled with precision. To ensure structured outputs from the LLM, we utilized Pydantic models for schema validation, integrating LangChain's key-description schema. This configuration allowed for the effective use of the function-calling capabilities of the ChatGPT 4.1 model, thereby ensuring that the responses were both contextually relevant and structured according to predefined data models.

This combination of deterministic components for high-throughput tasks and LLM-driven modules for context-sensitive inference underpins the reproducibility and extensibility of our initial testbed. Further implementation details are discussed throughout this article.

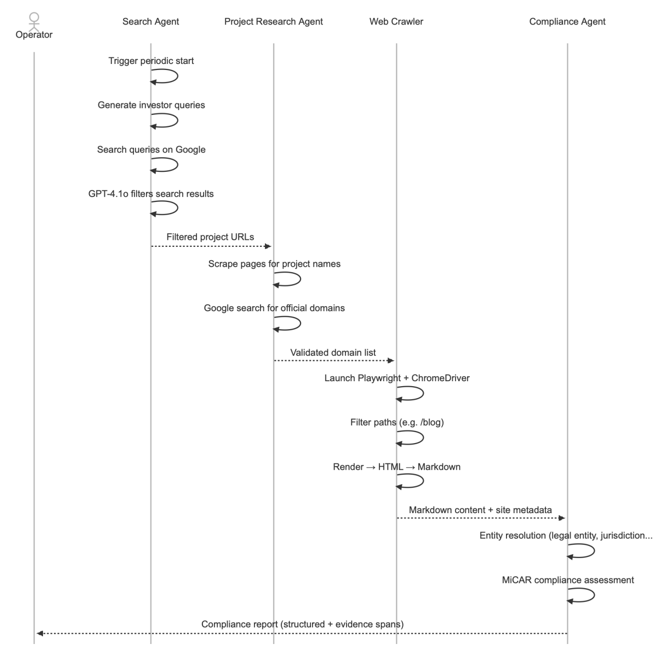

Fig. 1. Sequence diagram of the multi-agent architecture for automated regulatory analysis.

IV. SEARCH AGENT

The supervision pipeline begins with an *investor-behaviour simulation* designed to emulate real-world information-seeking patterns through the generation of natural language queries. Leveraging GPT-4.1 [3], the prompt employed to synthesize these queries incorporates the following key elements:

- A **semantic focus** on the activity of interest.
- A **temporal anchor**, which biases the retrieval towards the most recent content.

The activities analyzed in this study are related to newly issued crypto projects. The main focus includes crypto pre-sales, ICOs, crypto launchpads, as well as official project websites and informational content that presents such tokens. For each of these topics, two distinct queries are synthesized.

Query execution is performed using the Google Search engine to retrieve organic search results. Each result contains a page title, a URL, and a brief content snippet.

To reduce noise and minimize the number of irrelevant results passed downstream, a large language model (GPT-4.1) is used to batch-process and filter the retrieved entries. Temporal filtering is performed using the default search filter, while result filtering is carried out by formatting the search results into Markdown format. This format includes the URL, result title, and result snippet. The competence filtering is then applied using the Large Language Model (LLM), which includes only websites related to crypto pre-sales, ICOs, crypto launchpads, as well as official project websites and informational content presenting such tokens. The filtered URLs are returned in a JSON format. This filtering step ensures that only links explicitly referring to actual crypto projects are retained.

The refined list of search results is then passed to the *Project Research Agent* for deeper semantic analysis and structured extraction.

V. Project Research Agent

Following the initial filtering of search results, the next phase involves identifying and extracting structured information about one or more crypto projects potentially mentioned within each source. The objective is to determine, for each project referenced, both its name and official website—information which is essential to enable further crawling and detailed analysis (as described in Section VI).

To automate this step, we developed an AI-powered agent, capable of interpreting web content with minimal supervision. The agent operates over each filtered search result and is equipped with two specialized tools:

- **Web Scraper Tool** — Scans a webpage for mentions of crypto projects, extracting project names from headers, titles, and introductory content. It associates each project with a hyperlink, assuming external links point to the official website.
- **Official Website Finder —** Searches for the official URL when not explicitly provided, using semantic similarity, domain consistency, and structural cues like product pages or roadmaps to validate and return the most likely website.

Unlike single-entity extraction pipelines, the agent is designed to handle multi-entity extraction scenarios. A single search result or webpage may refer to multiple distinct crypto projects—for example, in comparative reviews, listicles, or rankings. The agent identifies each project individually and performs a targeted lookup to retrieve the corresponding official link, when necessary.

The agent's operation follows the multi-step process: (1) identifiable project names extract from the page; (2) association with a hyperlink if available; (3) if the link is missing or uncertain, invoke the official website discovery mechanism; (4) structured list of project names and validated URLs.

This modular and extensible design ensures that the downstream crawler receives a comprehensive and accurate set of targets, minimizing redundancy and maximizing semantic coverage. It also lays the foundation for future enhancements such as integrating project categorization, token symbol recognition, or smart contract address resolution.

VI. Web Crawler Design

The web crawler is designed to generate a structured and contextual snapshot of a given website domain, capturing the fully rendered content of each page for downstream analysis. This approach is particularly useful in regulatory and compliance scenarios, where semantic fidelity and metadata consistency are essential.

The implementation relies on Playwright [4] enabling robust scraping of JavaScript-based web pages. To optimize the data collected, certain URL paths—such as those containing `/blog` or `/media`—are excluded, as they typically contain redundant or non-informative content with respect to the objectives. It is possible that some websites may not allow large-scale scraping of multiple web pages or may present CAPTCHA challenges when unusual activity is detected. To address this, specific request headers, such as `Accept-Language` and `Referrer`, have been incorporated to simulate a browser session that appears to be controlled by a human user. Additionally, the pages are scraped within the same browser session. Furthermore, the use of rotating proxies or specialized services for bypassing CAPTCHA mechanisms and rate limits has been implemented.

The agent is composed of three main components:

- **Controller** — Manages the crawling process, controlling fetch frequency, identifying internal links, and appending metadata (title, meta description, canonical URL, crawl timestamp) to the extracted content in Markdown format.
- **Headless Browser Worker** — Utilizes Playwright with ChromeDriver to render client-side JavaScript, ensuring consistent and reliable crawling.
- **Parser** — Converts the rendered HTML to Markdown, preserving document structure and reducing token usage for language model processing.

The crawler outputs Markdown-formatted content along with structured metadata for each page. Additionally, it returns the list of internal links discovered for potential inclusion in the traversal queue. This output is optimized for both human interpretability and integration into downstream pipelines for classification, summarization, or content indexing.

VII. Compliance Agent

The Markdown corpus extracted from a project's official website is ingested by the *Compliance Agent*, which performs two successive tasks:

a) **Entity resolution**. The agent:
   - determines the legal entity responsible for the project;

- identifies the website's jurisdiction;
- enumerates the languages in which content is offered;
- detects whether a direct purchase mechanism for the crypto-asset is provided.

b) **Regulatory assessment**. Using a predefined output schema aligned with MiCAR, the agent evaluates the presence, absence, or partial fulfilment of every mandatory disclosure item. Each verdict is returned in a human-readable format and is accompanied by the precise textual span that substantiates the decision.

The structured data populate a templated compliance report for supervisory analysts. Each analysis is returned in a human-readable format, accompanied by the exact textual span supporting the decision. Additionally, the agent generates a concise *explanation string* summarizing *why* a field was populated in that manner. This rationale, produced by the underlying LLM using attention-based feature attribution techniques, helps detect hallucinations and is passed to the Operator Dashboard, enabling analysts to trace the decision-making path and identify potential errors [5, 6, 7].

VIII. Operator Dashboard

A dashboard exposes the findings to human analysts. Key features include:

- **Portfolio View**: sortable table enumerating all monitored projects with categorization.
- **Evidence Explorer**: side-by-side display of extracted field values and their supporting text spans, with hyperlinks to the original web pages.
- **Alert Engine**: rule-based notifications triggered when any mandatory MiCAR requirement evaluates to *false*.

By integrating automated extraction, formal rule evaluation and interactive visual analytics, the dashboard operationalises a risk-based supervisory workflow that aligns with MiCAR's objectives of investor protection and market integrity.

IX. Results

To assess the effectiveness of the proposed multi-agent pipeline, we issued ten natural–language queries—listed in Table III—designed to mimic realistic information-seeking behaviour by retail and institutional investors.

TABLE III
Example natural–language queries for the Search Agent

N.	Query text
1	Best emerging crypto projects **April 2025** for long-term investment
2	Top DeFi tokens launched in **Q1 2025** with active user base
3	Most promising blockchain gaming coins to watch in **2025**
4	New asset-backed stablecoins approved in the EU in **2025**
5	Crypto projects raising funds via ICO **Spring 2025** Europe
6	Sustainable ESG-focused crypto-asset announced in **2025**
7	Low-cap utility tokens gaining traction **April 2025**
8	Regulated e-money tokens compliant with MiCAR in **2025**
9	Best layer-2 scaling solutions launched in **2025**
10	High-yield staking opportunities for new tokens **2025**

TABLE IV
Summary of CryptoIndex Token (INDX)

Category	Key Details
Project	CryptoIndex https://cryptoindex.com
Issuer	Onemore LLC, SVG; Reg.# 2345 LLC 2022; Suite 305 Griffith Corporate Centre; `contact@cryptoindex.com`
Token	CryptoIndex Token (INDX): Utility for governance, fee discounts, platform participation
Blockchain	BNB Chain, Ethereum (Multi-chain)
Whitepaper	https://cryptoindex-1.gitbook.io/cryptoindex-whitepaper
Audits	Smart contracts audited by CertiK (completed Jan 2025)
Supply	Total: 1.5 B INDX; Circulating (TGE): 339 M; Decimals: 18
Allocation	• Pre-sale: 28.9% • Team: 5% • ...
Utility	$\leq$ 35% fee discounts; staking rewards; DAO governance
Roadmap	• Q4'24: Pre-sale • Q1'25: TGE & Staking • ...
Risks	Market volatility; smart-contract vulnerabilities; regulatory uncertainty

Execution of these queries through the *Search Agent* returned 1 000 organic search results, which were subsequently filtered and deduplicated by the *Project Research Agent*. From this corpus, we identified 137 unique crypto-asset projects. The regulatory assessment performed by the *Compliance Agent* indicates that 83%[5] of the projects adherence (at least partially) to the MiCAR requirements. A finer-grained breakdown is:

- **Fully compliant (45 %)** — projects providing complete and unambiguous coverage of every mandatory white-paper item, including governance, reserve mechanisms, risk disclosures, and environmental impact statements.
- **Partially compliant (38 %)** — projects meeting core prudential and governance obligations but lacking one or more ancillary elements (e.g., continuous redemption terms, detailed environmental metrics).
- **Non-compliant (17 %)** — projects omitting critical disclosures such as legal-entity identification, reserve composition, or holder redemption rights, thereby failing to satisfy the minimum MiCAR threshold.

These figures demonstrate the system's capacity to triage large sets of heterogeneous documents and to surface substantive compliance gaps with minimal expert intervention; for example, Table IV provides a compliance summary for the CryptoIndex Token, a PoS-based ERC-20 utility token.

Performance metrics for each pipeline stage and the overall end-to-end latency are summarised in Table V.

X. Limitations and Future Work

Our initial evaluation demonstrates the promise of an AI-driven MAS for MiCAR compliance screening, but several factors limit its generalisability and operational maturity.

[5] A project is deemed *compliant* when it satisfies, fully or partially, all mandatory disclosure items enumerated by MiCAR.

TABLE V
RUNTIME STATISTICS OF THE MAS PIPELINE (N = 100 PROJECTS)

Pipeline stage	**Mean (s)**	**Median (s)**	**Std. dev. (s)**
Search Agent	8.1	7.8	1.9
Project Research Agent	21.0	20.3	3.2
Web Crawler	15.7	14.9	3.8
Compliance Agent	12.3	12.0	2.5
End-to-end (one project)	57.1	54.0	7.8

Limitations

1) **Sampling Bias of Discovery Queries.** We rely on search results generated from ten simulated investor queries. This set may not cover all real-world discovery channels—social media, specialized forums or non-English resources—leading to potential false negatives (missed projects) and false positives (irrelevant or misleading entries).
2) **Rule-Mapping Granularity.** Compliance checks rest on manually crafted mappings from text spans to MiCAR clauses. Such mappings cannot capture all interpretive nuances—delegated acts, Q&As, country-specific guidance—or the subtleties of legal language, especially in local translations.
3) **Jurisdictional Coverage.** While the system flags the presence or absence of MiCAR disclosures, it does not determine whether the issuer actually offers services within specific EU Member States—a critical criterion for applicability and NCA notification under Art. 21–22 of MiCAR.
4) **Regulatory Drift.** MiCAR is supplemented by delegated regulations, technical standards and supervisory Q&As that evolve over time. Our rule-base is updated manually, introducing latency and risking misalignment with the latest requirements.

Directions for Future Work

To strengthen our framework and address the above gaps, we plan to pursue:

- **Automated Regulatory Ingestion.** Build a *regulatory-watch* microservice that periodically scrapes official EU and NCA sources (ESMA, EBA, ECB, national registers) and ingests updates to MiCAR texts, delegated acts and guidance into the rule repository via RAG-powered parsing.
- **Comparative LLM Evaluation.** Perform a parallel evaluation of multiple large language models (e.g., GPT-4, LLaMA). For each feature extracted, aggregate the outputs across all models and select the prediction endorsed by the majority as the final result.
- **Context-Aware Agent Dynamics.** Incorporate context-awareness frameworks into the MAS workflow—enabling agents to dynamically adapt to environmental signals and evolving operational contexts—by leveraging recent advances in context-aware multi-agent architectures [8]
- **Jurisdiction Inference Module.** Integrate geolocation and corporate registry lookup agents to infer the Member States in which a project actually offers services—by combining domain-based heuristics, website language metadata and automated NACE-code verification against national supervisory registers.
- **Cross-Regulatory Generalisation.** Extend the MAS architecture to accommodate both EU and non-EU regulatory frameworks (e.g. DLT Act; Singapore Payment Services Act; Hong Kong SFC's Virtual Asset regime), by parametrizing compliance schemas, data models and rule-sets for each jurisdiction, thereby creating a truly domain-agnostic RegTech platform.
- **Absence of Expert Benchmark Comparison.** No comparative analysis with a manual expert control group was conducted in this study. The focus of this contribution was on the automation of the regulatory analysis pipeline. Consequently, the system's performance was evaluated solely based on the automated pipeline. However, future work will include a comparative assessment against expert evaluations to provide a more comprehensive validation of the system's accuracy and to further refine its capabilities.

By systematically tackling these enhancements, future iterations will deliver broader language support, real-time regulatory alignment, precise jurisdictional applicability assessments, and greater resilience to adversarial inputs—culminating in a more dependable tool for continuous, pan-European crypto-asset oversight.

REFERENCES

[1] W. Van der Hoek and M. Wooldridge, "Multi-agent systems," *Foundations of Artificial Intelligence*, vol. 3, 2008.

[2] C. Qian, Z. Xie, Y. Wang, W. Liu, Y. Dang, Z. Du, W. Chen, C. Yang, Z. Liu, and M. Sun, "Scaling large-language-model-based multi-agent collaboration," *arXiv preprint arXiv:2406.07155*, 2024.

[3] R. Ye, S. Tang, R. Ge, Y. Du, Z. Yin, S. Chen, and J. Shao, "Mas-gpt: Training llms to build llm-based multi-agent systems," *arXiv preprint arXiv:2503.03686*, 2025.

[4] B. García, J. M. del Alamo, M. Leotta, and F. Ricca, "Exploring browser automation: A comparative study of selenium, cypress, puppeteer, and playwright," in *International Conference on the Quality of Information and Communications Technology*. Springer, 2024.

[5] X. Wu, H. Zhao, Y. Zhu, Y. Shi, F. Yang, T. Liu, X. Zhai, W. Yao, J. Li, M. Du *et al.*, "Usable xai: 10 strategies towards exploiting explainability in the llm era," *arXiv preprint arXiv:2403.08946*, 2024.

[6] A. Bilal, D. Ebert, and B. Lin, "Llms for explainable ai: A comprehensive survey," *preprint arXiv:2504.00125*, 2025.

[7] K. G. Barman, N. Wood, and P. Pawlowski, "Beyond transparency and explainability: on the need for adequate and contextualized user guidelines for llm use," *Ethics and Information Technology*, vol. 26, no. 3, p. 47, 2024.

[8] H. Du, S. Thudumu, R. Vasa, and K. Mouzakis, "A survey on context-aware multi-agent systems: techniques, challenges and future directions," *arXiv preprint arXiv:2402.01968*, 2024.