Rule-based Anti-Money Laundering in Financial Intelligence Units: Experience and Vision*

Luigi Bellomarini[1], Eleonora Laurenza[1,2], and Emanuel Sallinger[3]

[1] Banca d'Italia
[2] Financial Intelligence Unit for Italy
[3] TU Wien & University of Oxford

Abstract. Money laundering is a major threat to the good functioning of financial systems. Despite huge technological investments, with machine learning at the heart of the Fintech revolution, we are still lacking explainable solutions in fighting money laundering, especially for Financial Intelligence Units (FIUs). This paper is based on the joint committment of the Fintech community and academia in applying state-of-the-art rule-based reasoning to counteract money laundering. We report a visionary position about the application of logic-based Knowledge Graphs and reasoning with languages in the Datalog+/- family in the anti-money laundering (AML) domain. After motivating the impact and the importance of an explainable rule-based solution, we pin down the core AML problems in the form of high-level decision tasks. We envision that the FIU knowledge is modeled as the ground truth of a KG, so that AML tasks are formulated and carried out as reasoning tasks, addressing specific quality desiderata. We provide technical zoom and concrete exemplification of the approach with a real money laundering case. We discuss relevant research and technological challenges.

Keywords: Knowledge Graphs · Reasoning · Anti-money laundering · Fintech.

1 Introduction

This paper is based on our experience with the Financial Intelligence Unit (FIU) for Italy, one of the most relevant in the world, and the University of Oxford in applying state-of-art rule-based reasoning to counteract money laundering. Recent guidelines of the Financial Action Task Force [18] make us aware that the AML and FIU communities are permeable to and in strong need for such declarative and fully explainable approaches. Yet, it is our experience that no systematic approach, methodology or system has been developed at any FIU or shared in the literature. This paper contributes a visionary opinion based on a real-world experience, with the ambitious goal of steering processes and systems at FIUs towards a rule-based direction, for explainable, efficient and ethical AML decisions.

Money laundering is the process of making illegally-gained money proceeds, that is dirty money, appear legal and clean [21], by obfuscating the origin of the money. *Anti-money laundering* is a global challenge, which involves financial and non-financial intermediaries of the private sector (banks, trusts, money transfers, casinos, securities

*The views and opinions expressed in this paper are those of the authors and do not necessarily reflect the official policy or position of the Italian FIU or Banca d'Italia.

dealers, real estate brokers, public notaries, dealers , etc.), Financial Intelligence Units and law enforcement agencies. Our interest here is in FIUs, the government agencies that act as intermediaries between the private sector and law enforcement agencies. FIUs collect and analyze raw transactional data, namely, *Suspicious Transaction Reports* (STRs), filed by the intermediaries and conduct financial intelligence investigation to produce detailed, inspectable and logically consequential inquiry reports of the money laundering cases as well as justifiable hypotheses on the underlying crime (technically the *predicate offences*) for legal follow-up.

The current outlook on worldwide money laundering is disconcerting: The International Monetary Fund [28], estimates the amount of laundered money to be between 1.7 and 4.5 trillion USD. Meanwhile, according to Refinitiv [33], over the last year there has been an increase of over 50% in AI-based Fintech investment in improvements in contrasting financial crime. Machine learning (ML) approaches are gaining ground [11]: novel supervised [34] and unsupervised [30] techniques are being proposed, with a recent focus on graph deep learning [37], natural language processing applications [26] and social network analysis [13]. Despite these efforts, the global progress in fighting financial crime is insufficient to address the uprising emergency [3]. In fact, AI vendors [36] and ML literature [11,14,31] have been concentrating on "compliance use cases" to support financial intermediaries (e.g., to accommodate the Fifth AML Directive [16]), while the entire AML community is looking at the AI technology in its growing demand for viable solutions able to fulfill the requirements of the FIU business.

FIU use cases are far from the compliance ones and involve several **core quality desiderata**: 1. *Complex reasoning tasks* on sophisticated money laundering cases involving a *network* of many heterogeneous entities should be possible; 2. *Fully explainable* elicitation of the specific financial fraudulent schemes is desired in order to sustain legal follow-ups and avoid unethical or unjustified conclusions; 3. A careful balance of *inductive approaches* and the extensive *domain experience* must be stricken; 4. Solutions should be *scalable* to support very high volumes. None of the above requirements are currently met by any existing framework or set of techniques for AML [11].

Contribution. In this paper, we leverage our experience with the Italian FIU and our work on state-of-art *Knowledge Representation and Reasoning (KRR)* languages such as *Warded Datalog*$^{\pm}$ [5], a language of the Datalog$^{\pm}$ family. Datalog$^{\pm}$ extends Datalog with existential rules, while introducing syntactic restrictions to guarantee decidability and tractability of the reasoning task. We have successfully applied *Warded Datalog*$^{\pm}$ as a core language for reasoning on *Knowledge Graphs* (KGs) [5], a data model for which applied reasoning use cases are growing [19]. We suggest that the favourable features of such languages and the availability of engineered systems such as VADALOG [7], enable reasoning on complex cases of general interest to FIUs and provide a visionary position motivating the application of a rule-based approach in order to fulfill the above desiderata. In particular, the main contributions of this work are:

- A novel **comprehensive formulation of AML use cases** as high-level decision tasks, with unprecedented attention to the aspects of relevance for FIUs.
- An introduction of our vision on a **rule-based reasoning approach** to AML. In particular, we envision that the enterprise knowledge available to FIUs and intermediaries is modeled as the ground truth of an AML KG. Domain experience from financial an-

alysts as well as machine learning models should be represented as *reasoning rules* over the KG, so that AML tasks are formulated and carried out as *reasoning tasks*.
- A **technical zoom** and concrete exemplification of the **benefits of a rule-based approach** to AML thanks to an anonymized and stylized Italian *money laundering case*.
- In the light of the described desiderata, a deep discussion of the most relevant **research and technological challenges** in a reasoning approach to AML: *data wrangling*, *compliance checks*, *analytics*, *detection of seen and unseen laundering patterns*, *handling uncertainty*, and taking *ethical decisions*.

Overview. Section 2 includes the use cases and our vision. Section 3 analyzes a real-world case from the Italian FIU. Section 4 focuses on challenges and opportunities for a rule-based approach to AML. Relevant related work is extensively discussed throughout the paper and summed up in Section 5. In Section 6 we draw our conclusions.

2 A Rule-based Vision on AML

We start in Section 2.1 by getting to the core of the complex set of AML-related processes a FIU carries out. We capture them in the form of *business use cases*, presented as straightforward research problems. In Section 2.2, we introduce a *visionary position* and model business use cases as reasoning tasks over a comprehensive rule-based Knowledge Graph for AML.

2.1 AML Use Cases: Getting to the Core

AML is the task of preventing the process of money laundering as a whole, in all its phases. It is frequently presented only through the lens of regulatory compliance, as the burden falls primarily on intermediaries, responsible for meeting *Know Your Customer* (KYC) standards posed by regulations or conducting *Customer Due Diligence* to avoid the related high penalties. On the other hand, FIUs, at the heart of the money-laundering contrast system, conduct many intertwined activities, culminating in *financial intelligence* analysis and disseminating the results to the judicial authority. At the core of such activities, we see the following two high-level decision tasks:

1. **Suspicion assessment**. Decide whether a financial *transaction* or a *case* meets a definition of "suspicious".
2. **Suspicious offence determination**. Decide whether a suspicious *predicate offence* underlies a transaction or case.

In both the tasks we need to distinguish between *transaction* and *case*. While the former is a money/asset transfer together with its related attributes (e.g., transaction amount, KYC profiles of the involved subjects), cases are sets of cohesive transactions along with their context, which may include: The network of involved subjects, assets, companies, shareholding structures, etc. The definition of suspicious may either derive from a national regulatory body or be independently adopted by the FIU. It can be either based on certain criteria or on probabilistic ones. A detailed example of case is in Section 3.

Task (2) complements Task (1) in that it makes the explainability requirement explicit: Besides knowing the fully detailed cause of suspiciousness, we want to decide on possible underlying predicate offences. More advanced characterizing tasks of practical utility for the FIUs can be formulated as *derived tasks* extending the previous ones:

- **Suspiciousness scoring**. Produce a heuristic measure (e.g., a score) of the level of suspiciousness of a transaction or a case (see Task (1)).
- **Suspicion classification**. Classify a case or transaction by its underlying predicate offences or produce a heuristic measure (e.g., a score) of the level of confidence that a specific predicate offence underlies a case or transaction (see Task (2)).
- **Risk-driven optimization**. Perform a risk-based scheduling or planning of the overall AML activity to optimize an enterprise-level (FIU or intermediary) objective function based on suspiciousness (see Task (1)) or predicate offence (see Task (2)). For example, process first more suspicious cases or higher-impact crimes.
- **Reconstruction**. Produce an explanation about the suspiciousness of a transaction or case (see Task (1)) or for the presence of a given predicate offence (see Task (2)).

2.2 Towards Rule-based Knowledge Graphs for AML

In order to lay out our reasoning framework for AML, let us point out the basic notions of Knowledge Graphs and Knowledge Representation and Reasoning (KRR) languages.

A KG can be defined as a semi-structured data model characterized by three components: 1. a *ground extensional component (EDB)*, with constructs, namely *facts*, to represent data in terms of a graph or a generalization thereof; 2. an *intensional component (IDB)*, with *reasoning rules* over the facts of the ground extensional component; 3. a *derived extensional component*, produced in the *reasoning process*, which applies rules on ground facts [7]. In this work, we refer to *logic-based* KGs, i.e., where the intensional component is defined with a logic-based KRR language. We envisage the adoption of VADALOG, a reasoning language revolving around *Warded Datalog*$^{\pm}$, a language of the Datalog$^{\pm}$ family that extends Datalog with existential quantification. Warded Datalog$^{\pm}$ captures full *Datalog* [10,27], so it supports *full recursion*, essential to navigate graph structures; at the same time it allows *ontological reasoning*, being able to express SPARQL queries under set semantics and the entailment regime for OWL 2 QL while guaranteeing *scalability* thanks to PTIME data complexity for the reasoning task [24]. VADALOG supports additional features of practical utility: a form of *aggregation* [35], *stratified negation*, *Boolean conditions*, *mathematical expressions*, *probabilistic reasoning*, embedded functions and *arbitrary machine learning models* [7].

The IDBs are defined in terms of existential rules, i.e. first-order sentence of the form: $\forall \bar{x} \forall \bar{y} (\varphi(\bar{x}, \bar{y}) \rightarrow \exists \bar{z}\, \psi(\bar{x}, \bar{z}))$, where φ (the *body*) and ψ (the *head*) are conjunctions of atoms with constants and variables. Heads can also equate variables. We write $\varphi(\bar{x}, \bar{y}) \rightarrow \exists \bar{z}\, \psi(\bar{x}, \bar{z})$ and replace $\wedge$ with comma to denote conjunction of atoms. The semantics of such a rule is intuitively as follows: if there is a fact $\varphi(\bar{t}, \bar{t}')$ in the derived extensional component (denoted as $\Sigma(D)$ and which initially coincides with the EDB), then there exists a tuple $\bar{t}''$ such that the facts $\psi(\bar{t}, \bar{t}'')$ are also in $\Sigma(D)$. Roughly, the *reasoning* process is the generation of $\Sigma(D)$ to answer specific queries.

A KG for AML. We are now ready to apply the above definitions to our domain. We see the set of AML tasks of interest to a FIU as reasoning over an encompassing KG, modeling all the relevant domain objects and interconnections including: transactions, enterprise data stores, compliance rules, suspicious patterns, etc. In particular, transaction data and Suspicious Transaction Reports (STRs) should be represented as EDBs as well as data from enterprise knowledge systems. For intermediaries, they include

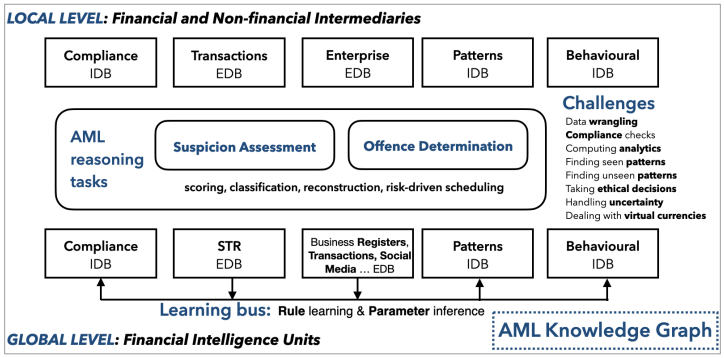

Fig. 1: An AML KG, with the main tasks and challenges.

transaction, enterprise and KYC data; EDBs derive from: STRs, business, person, asset, real estate and invoice registers, facts from social networks and media, newspapers or follow-up feedback from law enforcement authorities. IDBs should be used to represent and operationalize official regulations — in a "RegTech" approach — and encode custom criteria, including money laundering patterns (e.g., circular wire transfers, pyramidal control structures), domain rules and suspicious behaviours.

AML Intensional Knowledge and Reasoning. We believe that most of the money laundering patterns, suspicious behaviours and financial business rules can be described with a KRR language like VADALOG, supporting full recursion, ontological reasoning, probabilistic reasoning, and machine learning models. We envision that some reasoning rules are *designed* by financial analysts and domain engineers, while others are *learnt from data*, e.g., with *statistical relational learning* approaches [32].

Crafted rules embody valuable domain knowledge that cannot be induced from the data (e.g., a compliance regulation, the internals of a money laundering pattern, a complex domain rule) or is well known to the analysts (e.g., money laundering patterns, tactics for financial trail obfuscation) and inducing it from data would result in lower accuracy and explainability. The following two VADALOG rules, for instance, are part of the intensional component of an *invoice KG* and encode (a simplified version of) the domain knowledge to detect off-the-books slush funds, set up via false invoices:

$$Transaction(x, y, o) \rightarrow \exists z\ Invoice(z, y, x, o), \tag{1}$$

$$Invoice(z, y, x, o), \neg Transaction(x, y, o) \rightarrow PotentialSlushFund(y) \tag{2}$$

Normally, whenever there is a transaction from a subject x to y for a product/service o, then y issues an invoice z to x for o (Rule (1)). Yet, if we have an invoice for o, but no transaction from x to y exists, likely y is creating a slush fund (Rule (2)). This is an elementary case, representing a typical domain knowledge snippet. We shall see a full-fledged intensional component for a real AML case in Section 3.

Modern KRR languages allow for general and robust rules that capture both seen and unseen money laundering patterns. When rules embed parametric machine learning models, such parameters can in turn be inferred from data. On the other hand, learning rules from data can sensitively expedite ordinary rule design process. The *learning bus* in Figure 1 denotes such a hybrid deductive/inductive approach. AML tasks described in Section 2.1 should be modeled as reasoning tasks on the AML KG and used

to develop value added services or APIs to support the various business processes of intermediaries and FIUs. In particular, we envisage the development of *KG-enhanced* AML applications that rely on the outcome of Tasks (1-2) for crucial decisions. For example, tools for STR workflow processing at intermediaries should query the AML KG for Task (1) in order to finalize transactions and decide, according to regulatory or custom criteria, whether to file STRs to their local FIUs. Symmetrically, FIUs, upon receiving STRs from intermediaries, should trigger Task (1) to perform the follow-up assessments. Task (2) should be used by FIUs when specific cases need to be pursued, for autonomous investigations or instances issued by the enforcement authorities.

3 Reasoning on a Real Money Laundering Case

In the daily investigation duties of a FIU, deciding on the suspiciousness of an STR is a prominent activity: The Italian FIU handles $\sim 100K$ such cases per year, for which timely and explainable decisions are mandatory [20]. This causes Task (1) to emerge in terms of practical utility, with the need to develop effective decision heuristics. Hence, *suspiciousness scoring* is of strategic relevance among derived sub-tasks in Section 2.1. Let us now see how it can be carried out with a rule-based approach by analyzing a real pattern of money laundering cases, anonymized and stylized.

Description of the Data. The extensional component of the KG is in Figure 2. Solid black edges and nodes represent the EDB. In this case 14 companies, 4 financial intermediaries, and 4 individuals are involved. Depending on the case, the reasoning process may need to explore millions of entities in a KG with $\sim$ 22M nodes and $\sim$ 25.5M links.

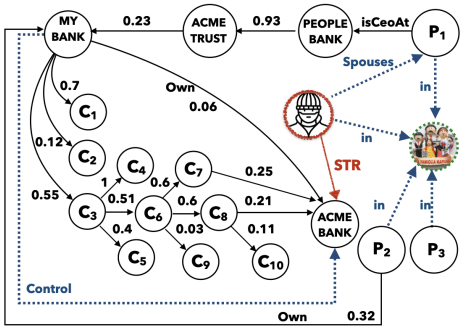

Fig. 2: An example of AML KG.

Description of the Case. Dashed edges and nodes in Figure 2 exemplify the derived extensional component. The case is triggered by an STR s (the red edge in Figure 2), reporting a loan instance from an individual x having a criminal record (the red node), to Acme Bank. Our goal is to score and explain suspiciousness of s). We start from the search for a typical laundering pattern: *A person x who is issuing a loan request to a bank b of which he/she is the ultimate beneficial owner (UBO), may intend to launder unclean money via the bank.*

A UBO is a person who ultimately, i.e., directly or via his/her undertakings, owns or controls a given entity (Acme Bank, in our case) on whose behalf a transaction (the loan) is being conducted [18]. In our case, if x were a UBO of Acme Bank (b), x might be requesting a fake loan to a bank he/she exerts control upon, with the likely intent to justify money illegally gained. What typically happens is that perpetrators also try to conceal their ultimate ownership. As we shall see, KG reasoning allows to go beyond the literal recognition of the pattern and detects the suspicious behaviour in a general sense, overcoming ownership concealment.

The following set of VADALOG rules (explained in detail in the next paragraphs) encode the intensional component representing the domain we have just described:

$$p :: Person(i_1, f_1^1, \ldots, f_n^1), Person(i_2, f_1^2, \ldots, f_n^2),$$
$$p = \#sim(f_1^1, \ldots, f_n^1, f_1^2, \ldots, f_n^2) \rightarrow Spouses(i_1, i_2). \quad (1)$$
$$Person(i, \ldots) \rightarrow \exists f\ Family(f), In(i, f). \quad (2)$$
$$Spouses(i_1, i_2), In(i_1, f_1),\ In(i_2, f_2) \rightarrow f_1 = f_2. \quad (3)$$
$$Control(x, x). \quad (4)$$
$$Own(x, y, w), w > 0.5 \rightarrow Control(x, y). \quad (5)$$
$$Control(x, z), Own(z, y, w), msum(w, \langle z \rangle) > 0.5 \rightarrow Control(x, y). \quad (6)$$
$$IsCeoAt(x, y) \rightarrow Own(x, y, 1.0). \quad (7)$$
$$Own(x, z, w_1), Own(z, y, w_2) \rightarrow Own(x, y, msum(w_1 \cdot w_2, \langle z \rangle)). \quad (8)$$
$$Own(x, y, w), In(x, f), j = msum(w, \langle x \rangle) \rightarrow Own(f, y, j). \quad (9)$$
$$w :: STR(s, b, x), Loan(s), In(x, f), Control(f, b) \rightarrow Suspicious(s). \quad (10)$$

Control. Represented in Rule (10), the suspiciousness degree of s depends on the probability of x controlling (`Control`) b. To this end, Rules (4-6) define control with a broadly accepted formulation, also present in logic programming contexts [10]:

A company (or a person/family) x controls *a company y, if: (i) x directly owns more than 50% of y (Rule (5)); or, (ii) x controls a set of companies that jointly (i.e., summing the share amounts), and possibly together with x, own more than 50% of y (Rule (6)).*

Moreover, Rule (7) extends the notion of control with the simplifying assumption that the CEO of a company has full control over it. Rule (8) accumulates direct and indirect ownerships that x exerts on y, along all possible ownership paths. In our case, individual x does not control Acme Bank: x is actually concealing control through his/her family.

Family Relationships. Let us consider Rules (1-3). Rule (2) states that every individual belongs to a family, his/her own. Rule (3) merges families f_1 and f_2 whenever they contain two spouses, i_1 and i_2. The overall effect is clustering the person space.

Family relationships are detected by Rule (1). It contains a specialized machine learning model for *link prediction* (denoted by the `#sim` embedded function). It takes as input the features of i_1 and i_2 and returns a score p measuring how likely i_1 and i_2 are spouses. Rule (1) produces `Spouses` facts with a probability depending on p. In our case, let P_1 (in Figure 2) be the suspicious individual's partner. The family also contains P_2, P_3 and potentially more people. Knowing the family members, we can determine the overall relationship of f with Acme Bank. To this aim, Rules (9) aggregates ownership amounts originating from different family members.

Overall Pattern. P_2 directly owns 0.34 of My Bank and P_1 indirectly owns $0.21 = 1 \times 0.93 \times 0.23$ of My Bank (by Rule (8)). In total, f controls My Bank owning 0.55 of the shares. My Bank, in turn, controls Acme Bank holding with 0.52 of the shares via a *pyramidal shareholding structure*, probably set up to obfuscate the connection between the two companies. In conclusion, family f controls Acme Bank.

Now we have all the ingredients to apply Rule (10). While x is not literally the UBO of Acme Bank, we have that his/her family as a whole is. Therefore we conclude there is actual suspicion x is perpetrating money laundering by justifying unclean money with

a fake loan from b, a bank he/she controls (and therefore can force to issue the loan). The overall confidence in this conclusion depends on the certainty p in the existence of the personal relationship — the output of a link prediction model — as well as on the intrinsic reliability w of the money laundering pattern.

Discussion of Real and Artificial Cases, and Performance. We tested this kind of tasks in a VADALOG system instance running on a memory-optimized virtual machine with 16 cores and 64 GB RAM (Intel Xeon architecture). The evaluation of one single case in the real Italian company KG requires ~ 420 seconds, averaged over 100 runs. Elapsed times are compatible with production use of the solution. A massive execution on the real KG, in search of all the cases respecting the same laundering pattern successfully individuated 1365 cases having a suspiciousness score of at least 0.8, in 1600 seconds. Evaluating accuracy of the specific suspiciousness scoring model adopted in this example is out of the scope of our discussion. Assessing accuracy of STR scores is controversial and usually built as a comparison against the human analyst's performance. Nevertheless, as common in unsupervised settings, automated decisions often improve human performance, finding new unseen patterns or defusing others. This makes accuracy evaluation even more challenging and matter of dedicated studies.

The core complexity element affecting reasoning times lies in the number of "ownership paths" connecting companies and individuals. In order to assess the scalability of our approach, we developed a graph generator able to build networks with the same topological characteristics as the real KG: It adopts a variant of the Barabási–Albert model [2] for directed scale-free networks, with parameters fitted from the real KG. In the generation of the graph, the attachment likelihood tuning has been used to generate instances of increasing density, while keeping the other characteristics stable. Quantitative features have been generated by sampling from appropriate distributions, also fitted from the real KG (e.g., a Beta distribution for company shares). With this tooling, we have built 8 artificial cases showing that ownership paths can be evaluated in less than 20 seconds for graphs with 1M nodes and density similar to the real one. For graphs much denser than the real-world financial networks, elapsed time grows to $\sim$ 3000 seconds for the same number of nodes. Execution time is polynomially affected by the number of nodes, e.g., with 4600 seconds for a ~ 28M nodes graph (much larger than the real one) having realistic density. In conclusion, the solution is highly scalable and ready to support even multi-national KGs, e.g., at European level, which are comparable in number of nodes yet have dramatically lower density than our synthetic cases.

4 Challenges and Opportunities for Rule-Based AML

Let us now extend our discussion and deal with what we consider the most relevant research, technological challenges and opportunities for a rule-based view of AML.

Distributed Reasoning. End-to-end AML processes involve the cooperation of different FIUs and intermediaries (e.g., within the European FIU.net[4]). Processes should be designed as distributed reasoning tasks: At FIU level, an overall scheduling of AML activities based on STR and predicate offence scores could be performed, so as to concentrate computational resources only (or first) on the most relevant cases. FIUs and intermediaries (and law enforcement agencies in the longer term) would cooperate in the

[4] https://ec.europa.eu/home-affairs/e-library/glossary/fiunet_en

execution of Tasks (1-2) to converge to AML decisions with a global/local perspective. In particular, reasoning tasks should span multiple KGs, deployed at and maintained by the different actors. Reasoning should be performed as coordinated sub-tasks, each local to its respective data shard. For example: an intermediary could trigger *suspicion assessment* before issuing a transaction; then, this process, would in turn activate the respective *suspicion assessment* and *suspicious offence determination* on the FIU side.

Trust, Information Sharing and Data Integration. Distributed reasoning would sustain trust thanks to a form of *reasoning locality* with shared processes and intermediary-/FIU-level EDBs. This would overcome, for example, the need for the FIUs to retrieve and store transaction data from the intermediaries or to share privileged intelligence information with one another. On the other hand, the approach would encourage actors to share IDBs, e.g., in the form of AML criteria or patterns, so fostering *standardization* and *reproducibility of checks*. Moreover KGs should contribute to the construction of an *integrated information asset* for intermediaries and FIUs, to *standardize compliance checks* and contrast the proliferation of diverse implementations of the various tasks.

Data Wrangling: Building the EDBs. Data wrangling is the construction of the EDBs from the various sources involved in AML such as transaction data, enterprise stores, business registers, KYC profiles, and requires the solution to non-trivial data management problems such as *entity resolution* [12], *data fusion* [15], *natural language processing* algorithms for unstructured sources [26], etc. These tasks can be specified in the form of reasoning rules in VADALOG, a very effective *lingua franca* for data wrangling [22,29], with a a solid history in the data management community [17,23].

Computing Analytics. For tactical and strategic planning of AML activities, the availability of aggregate analytics is fundamental. VADALOG benefits from recent extensions to recursive logic formalisms, namely *monotonic aggregations* in Datalog that support queries and reasoning about the number of distinct variable assignments satisfying specific goals and conjunction of goals [35]. It is our experience that typical AML aggregate indicators can be expressed by such formalisms, which support scalable implementations with limited memory footprint.

Finding seen and unseen Money Laundering Patterns. *Graph database technology* lacks sufficient expressive power to capture many known money laundering patterns as, e.g., *RPQ-based languages* [9] do not handle full recursion; also, as they do not support ontological reasoning, patterns can be matched only on a case-by-case basis, with an unaffordable proliferation of queries. KGs represent a very good opportunity to industrialize pattern detection, which should be modeled in the form of reasoning rules of the *intensional component*. VADALOG is a good choice to represent any pattern, thanks to high expressive power, scalability and full explainability.

In order to keep up with the evolving pace of financial crime, we need to rely on *self-adapting* reasoning rules, robust and able to generalize unseen patterns as well. VADALOG copes with these requirements. First, the support for intensional predicates, existential quantification and, more generally, full ontological reasoning enables automatic interaction of seemingly unrelated domain areas to detect unforeseen illicit situations. Second, the possibility to embed machine learning models allows to detect fuzzy patterns and limit the proliferation of special cases. Reasoning can also be used to implement *scalable and explainable clustering* to group the entities (e.g., the suspicious transactions) according to their features or topological role in their network. Vice versa,

reasoning can be operated inside clusters calculated with standard techniques (e.g., *k-means*) to perform fine-grained comparisons. Another promising technique for unseen patterns is *link analysis* [11], consisting in establishing connections between entities (individuals, transactions, etc.) and using them to assess suspiciousness.

Usability and Effectiveness. VADALOG embodies many usability characteristics [24]: *plainess*, as it is based on the basic Datalog syntax; *simplicity*, as facts can be seen as database tuples; *modularity*, as rules do not rely on any form of compilation dependency or procedural ordering. Practical adoption in AML shows VADALOG to be a very *compact* formalism, able to encode in tens of rules, thousands of lines of procedural code. The semantics, highly based on first-order logic, is intuitive also for non-IT domain experts and the use of recursion is easily grasped as a form of *inductive definitions*. We observed high appreciation for the presence of inductive definitions involving existentials and aggregations, traditionally absent in *domain-specific languages* of financial and statistical realms. Although the development of complex cases may require the support of "VADALOG engineers", the language sparks users' interest in understanding core business details, with reduced IT costs, especially w.r.t. non-declarative technology.

Scalability. VADALOG offers a transparent and safe approach to scalability. When only core features are used (we are in the Warded Datalog$^{\pm}$ core), the language guarantees polynomial complexity, suitable to the majority of AML use cases. For ultra-scale scenarios, VADALOG trades minor syntactic restrictions (rules are in the *Piecewise Linear Warded Datalog*$^{\pm}$ [8] fragment) for high parallelizability. The adoption of the full range of advanced features, may come at the cost of higher complexity.

Handling Uncertainty. Mutually connected levels of uncertainty are involved in AML reasoning to define the similarity of patterns, the suspiciousness level of a transaction or the likelihood of a predicate offence, etc. VADALOG can handle uncertainty primarily with *probabilistic reasoning* [6]: Weights can be used to specify importance of rules; specifically, they are parameters of log-linear models defining the marginal probability of the resulting facts. *Embedded machine learning models*, e.g., for *graph embeddings* [25], are also supported and can provide input facts to or receive training data from the rules (see *learning bus* in Figure 1) in hybrid deductive/inductive reasoning.

Taking Ethical Decisions. Dealing with AML in a global perspective poses non-negligible ethical and legal issues in that one individual's subnetwork should be inspected only as a consequence of explicit suspicion. By contrast, some techniques (e.g., social network analysis [13,34]) move from emerging statistical evidence and then focus on specific subjects as a consequence. This behaviour can be considered intrusive misconduct for a FIU and rarely produces arguable evidence in judicial follow-ups. AML approaches based on reasoning on KGs overcome these difficulties by operating in a *query driven* fashion [7]: The traversal of the KG is driven by specific goal of the AML task and the portion of the graph that is actually analyzed is the minimal expansion of the suspicious individual/bank/transaction vicinity.

5 Related Work

Logic-based KGs and VADALOG are fully covered in [5,7]; recent works already delve into practical applications [1] and KG architecture [4]. An interesting recent collection of KG use cases, under diverse perspectives, is provided by Toma et al. [19].

The most comprehensive survey on *AI approaches* to AML can be found in Chen et et al. [11], where they focus on *machine learning techniques*, including supervised [34] and unsupervised [30] algorithms, graph deep learning [37] and NLP [26]. Most of machine learning approaches proposed in the AML literature suffer from some limitations [11]: classification exhibits affordable performance only for medium datasets and suffer from slow training; approaches based on fuzzy logic capture specific statistical models and require intensive human work for rule writing [11]; neural approaches suffer from low explainability, which makes them hardly applicable for a FIU. Such patterns need to be carefully balanced with inductively learned patterns.

Finally, the related field is *social network analysis* on money laundering data has the primary objective of computing emerging statistical properties that are specific measures referring to the network as a whole [13,34].

6 Conclusion

Money laundering is a significant financial risk for intermediaries and a severe threat for economies, governments and ultimately for personal wealth and freedom. The potpourri of rising AI techniques for AML too often neglect domain experience and fail to address requirements of relevance to FIUs. With this paper, we wish to share our experience in making tangible actions to combat money laundering with AI: We aim to systematize and stimulate the research debate by offering a unifying framework for explainable AML. Logic-based KGs are the right means towards a holistic approach, balancing the power of the inductive techniques, with the brilliancy of top-level financial analysis.

Acknowledgements. This work was supported by EPSRC programme grant EP/M025268/1, the EU H2020 grant 809965, and the Vienna Science and Technology (WWTF) grant VRG18-013.

References

1. Atzeni, P., Bellomarini, L., Iezzi, M., Sallinger, E., Vlad, A.: Weaving enterprise knowledge graphs: The case of company ownership graphs. In: EDBT (2020)
2. Barabasi, A.L., Albert, R.: Emergence of scaling in random networks. Science (New York, N.Y.) 286, 509–12 (11 1999)
3. Basel Institute on Governance: AML Index 2018. https://bit.ly/2Yd3ray (2019), [Online; accessed 17-Jan-2020]
4. Bellomarini, L., Fakhoury, D., Gottlob, G., Sallinger, E.: Knowledge graphs and enterprise AI: the promise of an enabling technology. In: ICDE. pp. 26–37. IEEE (2019)
5. Bellomarini, L., Gottlob, G., Pieris, A., Sallinger, E.: Swift logic for big data and knowledge graphs. In: IJCAI (2017)
6. Bellomarini, L., Laurenza, E., Sallinger, E., Sherkhonov, E.: Reasoning under Uncertainty in Knowledge Graphs (to appear). In: RuleML+RR (2020)
7. Bellomarini, L., Sallinger, E., Gottlob, G.: The vadalog system: Datalog-based reasoning for knowledge graphs. PVLDB 11(9), 975–987 (2018)
8. Berger, G., Gottlob, G., Pieris, A., Sallinger, E.: The space-efficient core of vadalog. In: PODS. pp. 270–284. ACM (2019)
9. Bonifati, A., Fletcher, G.H.L., Voigt, H., Yakovets, N.: Querying Graphs. Synthesis Lectures on Data Management, Morgan & Claypool (2018)
10. Ceri, S., Gottlob, G., Tanca, L.: Logic programming and databases. Springer (2012)
11. Chen, Z., Khoa, L.D.V., Teoh, E.N., Nazir, A., Karuppiah, E.K., Lam, K.S.: Machine learning techniques for AML solutions in STR detection: a review. K. Inf. Syst. 57(2), 245–285 (2018)

12. Christen, P.: Data Matching Entity Resolution, and Duplicate Detection. Springer (2012)
13. Colladon, A.F., Remondi, E.: Using social network analysis to prevent money laundering. Expert Syst. Appl. 67, 49–58 (2017)
14. Dias, L.F.C., Parreiras, F.S.: Comparing data mining techniques for anti-money laundering. In: SBSI. pp. 73:1–73:8. ACM (2019)
15. Dong, X.L., Naumann, F.: Data fusion - resolving data conflicts for integration. PVLDB 2(2), 1654–1655 (2009)
16. European Parliament: Directive (eu) 2018/843 of the european parliament and of the council. https://bit.ly/3gWZYFX (2018), [Online; accessed 17-Jan-2020]
17. Fagin, R., Kolaitis, P., Miller, R., Popa, L.: Data exchange: Semantics and query answering. In: ICDT (2003)
18. FATF: Transparency and Beneficial Ownership. https://bit.ly/2UklJWj (2016), [Online; accessed 17-Jan-2020]
19. Fensel, D., Simsek, U., Angele, K., Huaman, E., Kärle, E., Panasiuk, O., Toma, I., Wahler, J.U..A.: Knowledge Graphs - Methodology, Tools and Selected Use Cases. Springer (2020)
20. Financial Intelligence Unit for Italy: Rapporto annuale 2018. https://bit.ly/3fwYH6W (2019), [Online; accessed 19-Jun-2020]
21. Fincen: History of AML Laws. https://bit.ly/2AMlNau (2015), [Online; accessed 17-Jan-2020]
22. Furche, T., Gottlob, G., Neumayr, B., Sallinger, E.: Data wrangling for big data: Towards a lingua franca for data wrangling. In: AMW (2016)
23. Golshan, B., Halevy, A.Y., Mihaila, G.A., Tan, W.: Data integration: After the teenage years. In: PODS. pp. 101–106. ACM (2017)
24. Gottlob, G., Pieris, A.: Beyond SPARQL under OWL 2 QL entailment regime: Rules to the rescue. In: IJCAI. pp. 2999–3007 (2015)
25. Grover, A., Leskovec, J.: node2vec: Scalable feature learning for networks. In: KDD (2016)
26. Han, J., Barman, U., Hayes, J., Du, J., Burgin, E., Wan, D.: Nextgen AML: distributed deep learning based language technologies to augment aml investigation. In: ACL (2018)
27. Huang, S.S., Green, T.J., Loo, B.T.: Datalog and emerging applications. In: SIGMOD (2011)
28. International Monetary Fund: World economic outlook, april 2019. https://bit.ly/3cKyuzL (2019), [Online; accessed 17-Jan-2020]
29. Konstantinou, N., Abel, E., Bellomarini, L., Bogatu, A., Civili, C., Irfanie, E., Koehler, M., Mazilu, L., Sallinger, E., Fernandes, A., Gottlob, G., Keane, J.A., Paton, N.W.: VADA: an architecture for end user informed data preparation. J. Big Data 6, 74 (2019)
30. Larik, A.S., Haider, S.: Clustering based anomalous transaction reporting. In: WCIT. vol. 3, pp. 606–610. Elsevier (2011)
31. Oeben, M., Goudsmit, J., Marchiori, E.: Prerequisites and ai challenges for model-based anti-money laundering. In: IJCAI 2019 Workshop (2019)
32. Raedt, L.D., Kersting, K., Natarajan, S., Poole, D.: Statistical Relational Artificial Intelligence: Logic, Probability, and Computation. Synthesis Lectures on AI and machine learning, Morgan & Claypool (2016)
33. Refinitiv: Innovation and fight against financial crime. https://www.refinitiv.com/en (2019), [Online; accessed 17-Jan-2020]
34. Savage, D., Wang, Q., Zhang, X., Chou, P., Yu, X.: Detection of money laundering groups: Supervised learning on small networks. AAAI Workshops (2017)
35. Shkapsky, A., Yang, M., Zaniolo, C.: Optimizing recursive queries with monotonic aggregates in deals. In: ICDE. pp. 867–878 (2015)
36. Weber, M., Chen, J., Suzumura, T., Pareja, A., Ma, T., Kanezashi, H., Kaler, T., Schardl, C.E.L..T.B.: Scalable graph learning for aml: A first look. CoRR abs/1812.00076 (2018)
37. Weber, M., Domeniconi, G., Chen, J., Weidele, D.K.I., Bellei, C., Robinson, T., Leiserson, C.E.: Anti-money laundering in bitcoin: Experimenting with graph convolutional networks for financial forensics. CoRR abs/1908.02591 (2019)