Reasoning with small data samples for organised crime

Wilmuth Müller, Dirk Pallmer, Dirk Mühlenberg, Ioannis Loumiotis, Konstantina Remoundou, et al.

Wilmuth Müller, Dirk Pallmer, Dirk Mühlenberg, Ioannis Loumiotis, Konstantina Remoundou, Pavlos Kosmides, Konstantinos Demestichas, "Reasoning with small data samples for organised crime," Proc. SPIE 11413, Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications II, 114130D (21 April 2020); doi: 10.1117/12.2557543

SPIE.

Event: SPIE Defense + Commercial Sensing, 2020, Online Only

Reasoning with Small Data Samples for Organised Crime

Wilmuth Müller*[a], Dirk Pallmer[a], Dirk Mühlenberg[a], Ioannis Loumiotis[b], Konstantina Remoundou[b], Pavlos Kosmides[b], Konstantinos Demestichas[b]

[a]Fraunhofer IOSB, 76131 Karlsruhe, Germany;
[b]Institute of Communication and Computer Systems (ICCS), Athens, Greece

ABSTRACT

Building upon the possibilities of technologies like big data analytics, representational models, machine learning, semantic reasoning and augmented intelligence, our work presented in this paper, which has been performed within the collaborative research project MAGNETO (Technologies for prevention, investigation, and mitigation in the context of the fight against crime and terrorism), co-funded by the European Commission within Horizon 2020 programme, is going to support Law Enforcement Agencies (LEAs) in their critical need to exploit all available resources, and handling the large amount of diversified media modalities to effectively carry out criminal investigation.

The paper at hand focuses at the application of machine learning solutions and reasoning tools, even with only small data samples.

Due to the fact that the MAGNETO tools have to operate on highly sensitive data from criminal investigations, the data samples provided to the tool developers have been small, scarce, and often not correlated. The project team had to overcome these drawbacks.

The developed reasoning tools are based on the MAGNETO ontology and knowledge base and enables LEA officers to uncover derived facts that are not expressed in the knowledge base explicitly, as well as discover new knowledge of relations between different objects and items of data.

Two reasoning tools have been implemented, a probabilistic reasoning tool based on Markov Logic Networks and a logical reasoning tool. The design of the tools and their interfaces will be presented, as well as the results provided by the tools, when applied to operational use cases.

Keywords: Reasoning, logical reasoning, Markow Logic Networks, ontology, knowledge base, law enforcement agencies

1. INTRODUCTION

The "Multimedia Analysis and correlation enGine for orgaNised crime prEvention and investigation" (MAGNETO) formalizes a structural framework to enhance the operational capability of Law Enforcement Agencies (LEAs) in their fight against organized crime and terrorist organizations.

The main ambition of the project MAGNETO and its consortium is to empower LEAs (Law Enforcement Agencies) with the capability to process, manage, analyze, correlate and reason from large datasets characterized by heterogeneity. In particular, the technical goals defined by the Consortium are: development of solutions enabling the exploration of data from various sources, their indexing, enrichment through meta-information and contextualization, development of tools supporting semantic information fusion and inference based on processed data and development of a human-machine interface (HMI) enriching the situational awareness and operational capabilities of LEAs.

The computable framework deals with knowledge in a formalized manner. In the paradigm of semantic technologies, the metadata that represent data objects are expressed in a manner in which their deeper meaning and interrelations with other concepts are made explicit, by means of an ontology. This approach provides the underlying computing systems with the capability not only to extract the values associated with the data but also to relate pieces of data one to another, based on the details of their inner relationships. Thus, using reasoning processes new information shall be extracted. The semantic information model that is based on the MAGNETO ontology, allows, therefore, navigation through the data and discovery of correlations not initially foreseen, thus broadening the spectrum of knowledge capabilities for the LEAs.

Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications II, edited by Tien Pham, Latasha Solomon, Katie Rainey, Proc. of SPIE Vol. 11413, 114130D
 · doi: 10.1117/12.2557543

2. MAGNETO`S COMMON REPRESENTATIONAL MODEL (CRM)

For MAGNETO a Common Representational Model (CRM) has been developed to represent the knowledge that has been extracted from the heterogeneous data used by LEAs. The semantic web technology Resource Description Framework [1] has been used to specify the model by the means of an ontology.

The development of the ontology has been driven by adopting successful concepts in similar fields and by the uses cases that have been defined by the law enforcement agencies. In [2], the author describes a core ontology based on NATO standards to improve military intelligence analysis. The main concepts of this core ontology have been selected as a basis of the MAGNETO ontology.

- **Equipment**: Any item of materiel used to equip a person, organization or place to fulfil its role. An equipment may be subdivided into a number of subsidiary parts.
- **Organization**: An organizational entity or grouping which has a common purpose and which may have a recognizable hierarchical structure.
- **Place**: This category encompasses two aspects of a place. It defines a place as: The point or area on the earth's surface, or in space, occupied by a unit, person or equipment and delineated by a specific set of coordinates. A place may be a natural or a man-made feature, an area or a reference point.
- **Person**: A description of the physical characteristics and of the private and professional attributes of an individual. This will consist of, amongst other matters, details of the identification, education, family relationships, career and individual behavior patterns of the person.
- **Event**: A description of an incident or occurrence of some significance. An event may consist of a number of smaller events and is therefore capable of sub-division.

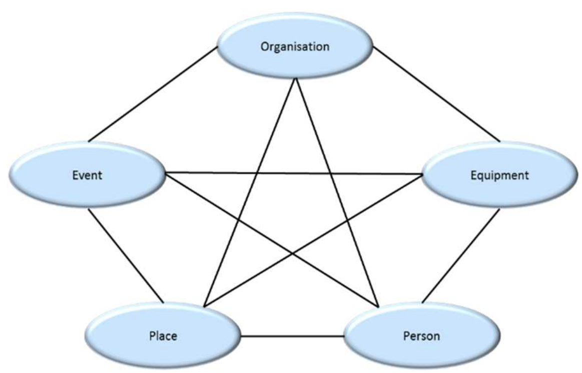

Figure 1: Main concepts of MAGNETO's Ontology.

Several existing ontologies have been identified and integrated into the MAGNETO ontology, the process and the outcome have been described in detail in [3] and [4].

The instances representing the high-level information in a given investigation are stored in a RDF triple-store provided by the Apache Jena [5] framework. The information extracted from heterogeneous data sources is communicated to the triple-store via the interface implemented by an Apache Fuseki server. Ideally, the entire triple-store is held in the memory to guarantee computationally efficient runtime. However, nowadays real-world cases have to process a vast amount of information preventing to keep everything in memory at the same time. A solution to this is the TDB2 storage concept of Apache Jena, which has theoretically no size limitations and allows 100's of millions of triples.

3. REASONING FOR INFERRING NEW KNOWLEDGE

The semantic reasoning technique aims at the enrichment of existing information, as well as the discovery of new knowledge and relations between different objects and items of data. Reasoning is a procedure that allows the addition of rich semantics to data, and helps the system to automatically gather and use deeper-level new information. Specifically, by logical reasoning MAGNETO is able to uncover derived facts that are not expressed in the knowledge base explicitly, as well as discover new knowledge of relations between different objects and items of data.

A reasoner is a piece of software that is capable of inferring logical consequences from stated facts in accordance with the ontology's axioms, and of determining whether those axioms are complete and consistent. Reasoning is part of the MAGNETO system and it is able to infer new knowledge from existing facts available in the MAGNETO knowledge base. In this way, the inputs of the reasoning systems are data that are collected from all entities in the MAGNETO environment, while the output from the reasoner will assist crime analysis and investigation capabilities.

4. RULES FOR SEMANTIC REASONING

In order for a reasoner to infer new axioms from the ontology's asserted axioms a set of rules should be provided to the reasoner. Rules are of the form of an implication between an antecedent (body) and consequent (head). The intended meaning can be read as: whenever the conditions specified in the antecedent hold, then the conditions specified in the consequent must also hold, i.e.: $antecedent \Rightarrow conoeouent$

The antecedent is the precondition that has to be fulfilled that the rule will be applied, the consequent is the result of the rule that will be true in this case. Both the antecedent and consequent consist of zero or more atoms or predicates. The antecedent is a single predicate or a conjunction of predicates, separated the character ^. An atom or predicate is of the form *C(x), P(x,y)* where *C* is an OWL class description (concept) or data range, *P* is an OWL property or relation, *x* and *y* are either variables, instances or literals, as appropriate.

An example of an antecedent is "$\text{isChildOf}(\text{s1}, \text{p}) \wedge \text{isChildOf}(\text{s2}, \text{p})$" A Conjunction of terms means that the two terms (called literals) are connected with a logical "AND", this means that the antecedent is fulfilled if both predicates are true. The logical AND is represented by the comma character or the character "^".

The consequent is usually a single predicate or a disjunction of predicates. In this example the consequent could be "isSiblingOf(s1, s2)". So the rule, which expresses that children of the same Parent are siblings, is written as:

$\text{isChildOf}(\text{s1}, \text{p}) \wedge \text{isChildOf}(\text{s2}, \text{p}) \Rightarrow \text{isSiblingOf}(\text{s1}, \text{s2})$

If the evidence in the knowledge base is

isChildOf(Benny, Jacob), isChildOf(Joseph, Jacob)

Then the result of applying the rule will be:

isSiblingOf(Benny, Joseph)

Some of the rules are predefined by the MAGNETO ontology definition. Following rules will be generated automatically:

- Taxonomy related rules on classes: If the concept "Car" is subclass of the concept "Vehicle", the rule "Car(x) => Vehicle(x)" is generated.
- Taxonomy related rules on properties: If the property "isSonOf" is a sub-property of "isChildOf", then the rule "isSonOf(s,p) => isChildOf(s,p)" is generated.
- Domain and Range related rules: Relations often have a single concept class for the domain or the range defined. The domain defines the concept that the relation arrow starts from, the range defines the concept that the arrow points to. For example, the relation "involvesPerson" has the domain "Event" and the range "Person". Therefore, it connects an event with a person that is involved in this event. From the definition of the relation "involvesPerson", two rules result:

  involvesPerson(e,p) => Event(e)

  involvesPerson(e,p) => Person(p)

5. EXAMPLE OF A RULE FOR ORGANIZED CRIME

The Police Academy in Szczytno (WSPOL) has proposed a rule for the Use Case 4 concerning the fuel crime, which is one of the main areas of activity of organized crime groups (OCG) and has the greatest impact on the depletion of tax revenues. One of the modus operandi used by perpetrators most commonly is the de-colorization of heating oil. The company buys heating oil from Germany and resells it to other companies for heating purposes. The use of heating oil is a subject to a lower rate of excise duty in Poland. Despite the declarations of going abroad, in fact, heating oil does not leave Poland, but is transported to places where the oil is discolored (the removal of colorant). Then, this oil is delivered to petrol stations and sold as a fuel, and so committing a tax fraud. Figure 2 illustrates the contents of the knowledgebase in MAGNETO for the described use case.

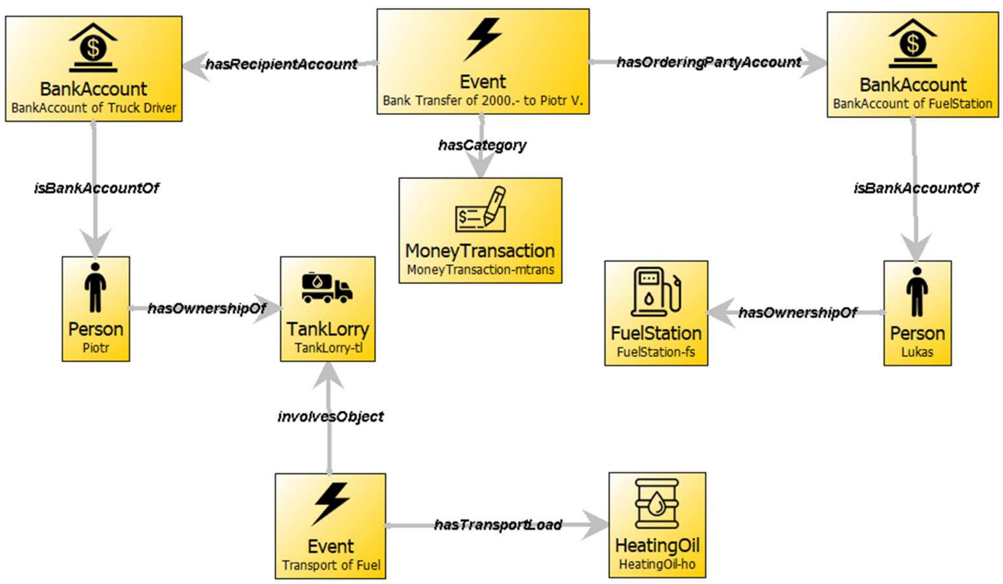

Figure 2: Example ontology population for the decolorization of heating oil rule

The following rule is proposed to detect illegal trade of heating oil has been implemented:

If the owner of a petrol station has transferred money from his account to the account belonging to the owner of the vehicle transporting the heating oil, the goods may have been used as fuel. When implementing this rule, the antecedent and the consequent must be mapped to the concepts and relations modelled in the MAGNETO ontology. As the creation of a new instance cannot be created by a rule, the outcome of the rule is the creation of four relations assigning the existing events to a crime category and assigning a suspect to the events.

EVENT(TRA), HASCATGEORY(TRA, MTRANS), MONEYTRANSACTION(MTRANS), HASORDERINGPARTYACCOUNT (TRA, BAFS), HASRECIPIENTACCOUNT (TRA, BATR), BANKACCOUNT(BATR),

PERSON(TRO), ISBANKACCOUNTOF(BATR, TRO), TANKLORRY(TL), HASOWNERSHIPOF(TRO,TL), TRANSPORT(TP),

EVENT(TPE), HASCATEGORY(TPE, TP), INVOLVESOBJECT(TPE, TL), HASTRANSPORTLOAD(TPE,HO), HEATINGOIL(HO), FUELSMUGGLING(FS)

- ⇨ HASCRIMECATEGORY(TPE, FS)
- ⇨ HASCRIMECATEGORY(TRA, FS)
- ⇨ ISSUSPECT(TRO, TPE)
- ⇨ ISSUSPECT(TRA, TRA)

6. PROBABILISTIC REASONING BASED ON MARKOV LOGIC NETWORKS

The technique employed is the so-called Markov Logic Networks (MLN), which allow probabilistic reasoning by combining a probabilistic graphical model with first-order logic. Markov Logic Networks (MLNs) have been introduced in 2006 by M. Richardson and P. Domingos, see [6]. Since then they have been an active area of research and were widely applied in different scenarios, e.g. ontology matching, statistical learning and probabilistic inference, see [7], [8] and [9]. The advantage of probabilistic reasoning is the capability to deal with uncertainties in the knowledge and the rules that the reasoning is applied on. The knowledge in the Common Representational Model (CRM) is partly based on witness statements that may contain unreliable information. Concerning the rules, there might be rare situations where the antecedent is given but the consequent does not hold. The trust in the knowledge and the rules is expressed in a weight value that is assigned to a formula: the higher the weight, the higher will be the probability of the consequents-generated result.

An MLN is a first-order knowledge base, i.e. a set of formulas $\{F_i \mid i \in I\}$ stated in first-order logic, where every formula is equipped with a corresponding weight $\omega_i \in \mathbb{R}$. These formulas serve as the base of a procedure yielding a Markov network, which then can be used to assign probabilities to possible states of instances of the underlying ontology.

The set of formulas consists of the evidence retrieved from the CRM and the rules that have been developed in cooperation with law enforcement agencies.

Grounding MLNs with complex formulas and large sets of data is a very challenging task, since the number of possible groundings may explode. Many MLN implementations have a major drawback on the poor scalability resulting from an exploding number of groundings of a first-order formula that are true given a truth assignment to all of its ground atoms.

If you want to use the following simple first-order formula

$$\forall x \forall y \forall z \forall u\colon\ R(x,y) \ \vee\ \ S(y,z)\ \vee T(z,u)$$

and assume that each logical variable has 100 objects (constants) in its domain. Grounding this formula results in 100 million ground formulas and performing any kind of inference on such a large Markov network is an extremely challenging task [10].

The state-of-the-art MLN implementation Tuffy claims to achieve a better scalability compared to the current MLN implementations scale beyond relatively small data sets (and even on many of these data sets, existing implementations routinely take hours to run). One of the reasons is that these are in-memory implementations: when manipulating large intermediate data structures that overflow main memory, they either crash or thrash badly. Tuffy leverages an RDBMS to address the above scalability and performance problems (see [11]).

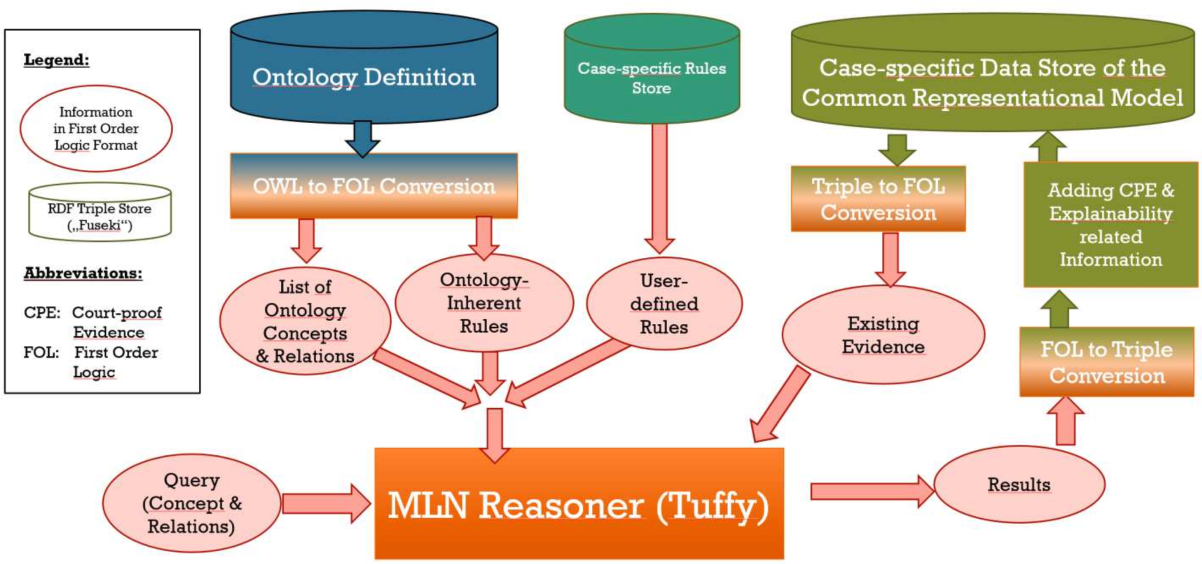

Figure 3: Workflow of MAGNETO's MLN reasoning module

In [12] the MLN implementation Tuffy has been integrated into an information fusion component for fusing information acquired by a distributed surveillance system with prior information contained in intelligence databases. Information given in the form of an OWL ontology, such as a taxonomy of defined concepts as well as relations, have proven to be easily convertible into FOL formulas and integrable into an MLN model. In a first step, OWL modeling constructs such as concepts and object properties (relations) have to be converted into unary and binary First-Order-Logic predicates, respectively, and taxonomy relationships are reformulated in terms of rules. This has been performed with the help of an existing transformation tool called Incerto [13]. In the final step, the transformed model has to be converted into an input form suitable for the MLN reasoner Tuffy, e.g., by creating separate input files for probabilistic evidences and weighted formulas as well as representing the formulas in conjunctive normal form.

The reasoning module developed for MAGNETO follows a similar approach, it is also based on the open-source implementation of MLN reasoned Tuffy, developed by Stanford University. It requires the input of text files in first-order logic. An adapter has been developed to integrate the MLN reasoner into the MAGNETO framework. The adapter connects to the knowledge stored in the Fuseki RDF Triple Store for input and output (see Figure 3).

The reasoning module needs following parameters as input - all URLs exposed by the FUSEKI data store containing the ontology and the knowledge base:

- OntologyURL: the Ontology File and the URL of the FUSEKI datastore containing the ontology
- InstanceURL: the URL of the FUSEKI datastore containing the instances that reflect the evidence collected so far. This is also the URL that the result will be written to.
- RulesURL: the rules file or an URL to the FUSEKI datastore containing an instance of the concept RulesFile
- QueryURL: the query file or an URL to the FUSEKI datastore containing an instance of the concept QueryFile

7. LOGICAL REASONING

For the logical reasoning process the Pellet reasoner **Fehler! Verweisquelle konnte nicht gefunden werden.** has been employed, which can support the SWRL rules. It has been published under the AGPL v3 license and it is a free open-source Java-based reasoner for OWL 2 and Semantic Web Rule Language. Pellet uses a tableau-based decision procedure to provide many reasoning services (e.g. subsumption, satisfiability, classification, instance retrieval, conjunctive query answering) along with the capability to generate explanations for the inferences it computes. It has bindings for both OWL API and Jena libraries, and supports the Protégé tool, command line and OWL API too. It supports reasoning services such as realization, classification, satisfiability, conjunctive query answering, entailment, consistency and explanation.

In order to allow the reasoner to infer the new axioms the following APIs have been used:

- *Ontology Web Language API [14]:* It is a standard Java library and API for working with OWL ontologies (OWL API main repository, 2019). It provides a standard interface to OWL reasoners, so different reasoners can be used without any update in the implementation.
- *Semantic Web Rule Language (SWRL) API[16]*: The SWRLAPI is a Java API for working with the OWL-based SWRL rule and SQWRL query languages [17]. It provides a collection of software components and Java-based APIs that allows among other to create, edit and manage the SWRL rules.

The above APIs provide the appropriate tools to create a Pellet Reasoner and apply the SWRL rules in order to infer the new knowledge in MAGNETO.

A simple flow of the reasoning procedure as described above is presented in Figure 4.

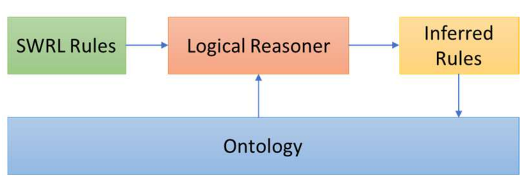

Figure 4: Logical Reasoning

To illustrate the application in MAGNETO, a simple example of the reasoning us shown here. Specifically, assume a simple rule that a person is suspect for a murder event if the person is an enemy of the victim. This rule can be described using SWRL as follows:

$$Event(?z) \wedge isEnemyOf(?x,?y) \wedge isDead(?y, true) \wedge involvesPerson(?z,?y) \Rightarrow isSuspect(?x,?z)$$

In this simple example, assume that person Dieter is enemy of Karl who has been murdered (Figure 5).

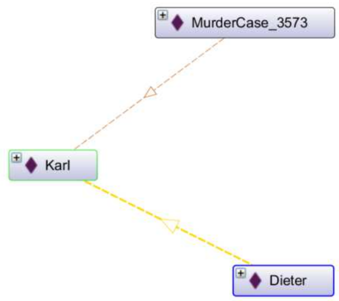

Figure 5: Example of a murder case in MAGNETO Common Representational Model.

Using the reasoner with the above simple rule, we can infer the suspect of the case, as it can be seen in the following result.

```
person : Dieter
inferred class for Dieter: magnetoModelObject
asserted class for Dieter: Person
inferred object property for Dieter: isSuspect -> MurderCase_3573
inferred object property for Dieter: involvesEntity -> MurderCase_3573
inferred object property for Dieter: socialRelation -> Karl
asserted object property for Dieter: isEnemyOf -> Karl
inferred object property for Dieter: involvesPerson -> MurderCase_3573
inferred object property for Dieter: hasEventObjectProperty -> MurderCase_3573
inferred object property for Dieter: hasPersonObjectProperty -> Karl
Is Dieter suspect for MurderCase_3573 ? : true
```

Then by adding the inferred axioms in the knowledge base the new relationships can be used (Figure 6) in order to query the ontology and find all the necessary information.

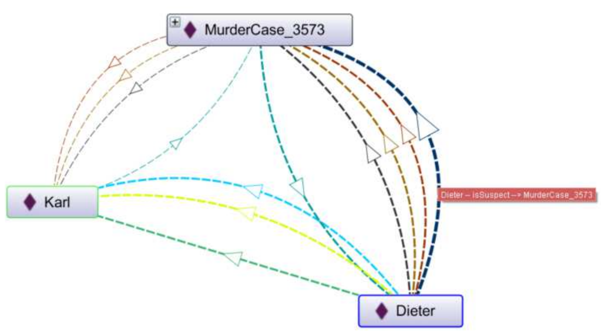

Figure 6: Example of a murder case in MAGNETO ontology after the application of the reasoner.

Thus, we can query the results using the DL query in Protégé in order to find the suspect or by using SPARQL. It is expected that for the purposes of MAGNETO the SPARQL will be used by the tools that LEAs will interact with.

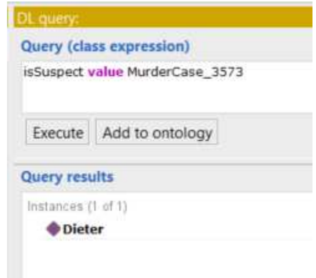

Figure 7: Querying in Protégé for the murder example.

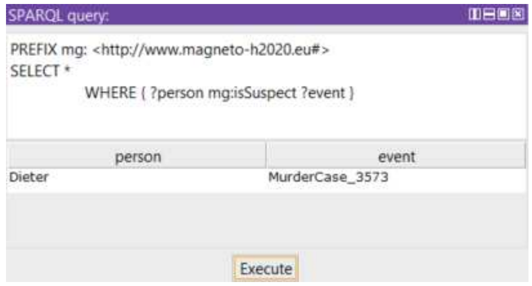

Figure 8: Querying using SPARQL for the murder example.

8. ANNOTATIONS FOR COURT-PROOF EVIDENCE AND EXPLAINABILITY

For research projects receiving funding from EU, it is mandatory that the developed solutions can meet the legal and ethical requirements. Two requirements shall be highlighted here and the implementation approach in the MAGNETO project shall be briefly described:

- **Explainability** is defined as "explaining of the workings of the system at both the global level as well as in relation to particular cases and circumstances". This requirement is a very problematic for applications of AI, since the calculations that lead to the result are very complicated and not comprehensible to the end-user without expert knowledge in AI.
- **Court-proof evidence** means that the law enforcement agencies are supported to present the facts and evidence that the accusation is based on.

As a result of the reasoning, new relations (object properties) linking existing instances, are created in the CRM. The reasoning supplies a confidence value for these inferred relations that must be reflected in the CRM, as well as the fact that the relation is created by a reasoning process. However, the ontology's data model based on triples cannot attach any additional information linked to the object property itself, neither as data nor as object property, so the concept/class "RelationDescription" has been introduced to hold the information about the new relation. The "RelationDescription" instance is then linked to the domain and the range instance of the inferred object property. Figure 9 illustrates this for the inferred relation "hasCrimeCategory". The concept class "RelationDescription" has the data property "hasRelationName" that holds the name of the inferred object property, as well as the confidence value. The data property "hasReasoningRules" will contain a list of rules that can create a new relation of this type (in the example type "hasCrimeCategory").

Unfortunately, the reasoners used in MAGNETO do not provide any information, which of the rules in the list produced the result. It might be a single "crucial" rule; it might be several rules that add up to the confidence value that has been calculated by the MLN. The fulfilment of the antecedent of the "crucial" rule might be based on the evidence stored in the CRM, or the fulfilment might result from the application of a second rule that has "previously" generated the evidence, and the second rule might make use of knowledge created by third rule. In the end, the result might be based on the application of a chain of rules, but this chain cannot be delivered from the MLN reasoning tool.

The "RelationDescription" is linked to an instance of the concept "ReasoningProcess", that holds data properties describing some details, i.e. the date of the processing (reasoning). The "ReasoningProcess" owns the object property "hasReasoningRules" that holds the complete set of rules that has been supplied to the reasoning process.

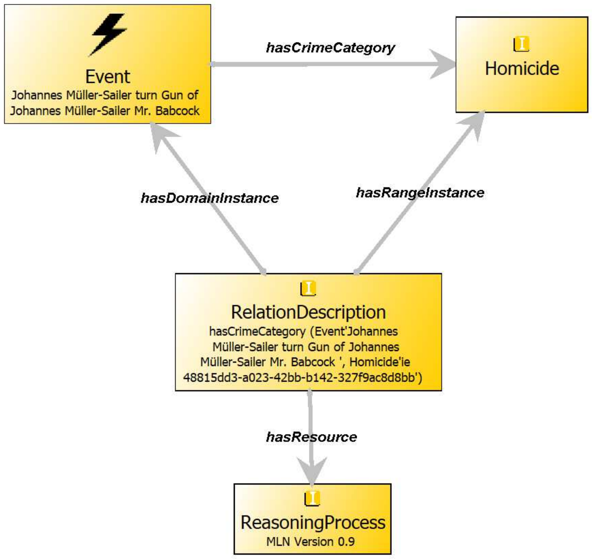

Figure 9: Annotations for inferred relations for achieving explainablity

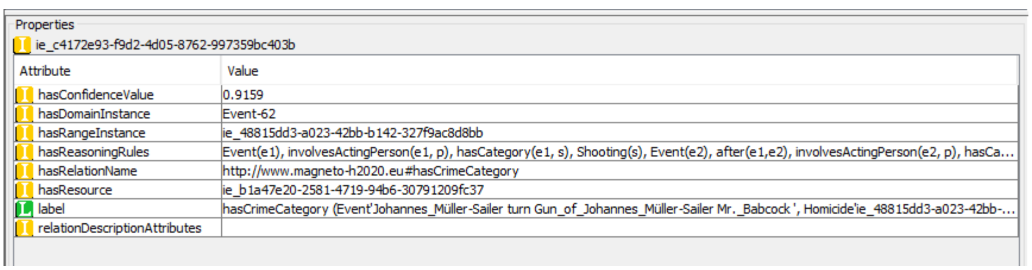

Properties
ie_c4172e93-f9d2-4d05-8762-997359bc403b

Attribute	Value
hasConfidenceValue	0.9159
hasDomainInstance	Event-62
hasRangeInstance	ie_48815dd3-a023-42bb-b142-327f9ac8d8bb
hasReasoningRules	Event(e1), involvesActingPerson(e1, p), hasCategory(e1, s), Shooting(s), Event(e2), after(e1,e2), involvesActingPerson(e2, p), hasCa...
hasRelationName	http://www.magneto-h2020.eu#hasCrimeCategory
hasResource	ie_b1a47e20-2581-4719-94b6-30791209fc37
label	hasCrimeCategory (Event'Johannes_Müller-Sailer turn Gun_of_Johannes_Müller-Sailer Mr._Babcock ', Homicide'ie_48815dd3-a023-42bb-...
relationDescriptionAttributes	

Figure 10: Data properties of the RelationDescription instance of Figure 9

9. CONCLUSION

Two reasoning tools have been developed for the MAGNETO platform to generate new knowledge by applying rules on the evidence stored in the Common Representational Model (CRM):

- The logical reasoning tool is based on the binary model of the evidence and its conclusions being true or false.
- The probabilistic reasoning tool is based on Markov Logic Network. A numerical confidence value can be assigned both for the evidence and the rules, and the conclusions are also rated with a confidence level.

In cooperation with LEAs, rules have been developed and supplied into the platform.

The population of the CRM’s ontology with the inferred knowledge is illustrated and the implementation of the ethical and legal requirements concerning explainability and court-proof evidence has been achieved by annotating inferred relations with information about the rules used to create them.

The next step will be the evaluation of the MAGNETO platform by the LEA end users.

ACKNOWLEDGMENTS

This work has been performed under the H2020 786629 project MAGNETO, which has received funding from the European Union's Horizon 2020 Programme. This paper reflects only the authors' view, and the European Commission is not liable to any use that may be made of the information contained therein.

REFERENCES

[1] Resource Description Framework (RDF) - Model and Syntax, retrieved from https://www.w3.org/TR/WD-rdf-syntax-971002/

[2] Valentina Dragos (2014). Developing a core ontology to improve military intelligence analysis. International Journal of Knowledge-Based and Intelligent Engineering Systems, IOS Press, 2013, 17 (1), pp.29-36.

[3] Kozik, R. et al. "The Identification and Creation of Ontologies for the Use in Law Enforcement AI Solutions – MAGNETO Platform Use Case", In: Nguyen N., Chbeir R., Exposito E., Aniorté P., Trawiński B. (eds) Computational Collective Intelligence. ICCCI 2019. Lecture Notes in Computer Science, vol 11684. Springer, Cham Proc., 335-345 (2019).

[4] Kozik, R. et al. "Common Representational Model and Ontologies for Effective Law Enforcement Solutions", Vietnam Journal of Computer Science, Vol. 07, No. 01, 1-18 (2020).

[5] Apache Jena Homepage, retrieved from https://jena.apache.org/, last accessed 2020/03/25.

[6] M. Richardson, P. Domingos, "Markov logic networks", Machine learning 62.1-2 (2006): 107-136.

[7] M. Niepert, C. Meilicke, H. Stuckenschmidt. "A Probabilistic-Logical Framework for Ontology Matching." AAAI. 2010.

[8] H. Chen et al. "Scaling up Markov logic probabilistic inference for social graphs." IEEE Transactions on Knowledge and Data Engineering 29.2 (2017): 433-445.

[9] P. Wittek, C. Gogolin. "Quantum enhanced inference in markov logic networks." Scientific reports 7 (2017): 45672.

[10] Deepak Venugopal et al.: "Just Count the Satisfied Groundings: Scalable Local-Search and Sampling Based Inference in MLNs", Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence 2015

[11] Niu et al. "Tuffy: Scaling up statistical inference in markov logic networks using an rdbms." Proceedings of the VLDB Endowment 4.6 (2011): 373-384.

[12] Kuwertz A., Mühlenberg D., Sander J., Müller W. (2018) Applying Knowledge-Based Reasoning for Information Fusion in Intelligence, Surveillance, and Reconnaissance. In: Lee S., Ko H., Oh S. (eds) Multisensor Fusion and Integration in the Wake of Big Data, Deep Learning and Cyber Physical System. MFI 2017. Lecture Notes in Electrical Engineering, vol 501. Springer, Cham

[13] Pedro Oliveira, "Probabilistic Reasoning in the Semantic Web using Markov Logic", MSc Thesis, Department of Informatics Engineering, University of Coimbra, Portugal, July 2009

[14] Pellet: An Open Source OWL DL reasoner for Java, available online: https://github.com/stardog-union/pellet, retrieved on 17.12.2018

[15] The OWL API, retrieved from http://owlapi.sourceforge.net//, last accessed 2020/03/25.

[16] Java API for working with the SWRL rule and SQWRL query languages. Retrieved from https://github.com/protegeproject/swrlapi

[17] M.J. O'Connor, R. S. (2008). The SWRLAPI: A Development Environment for Working with SWRL Rules. OWL: Experiences and Directions (OWLED), 4th International Workshop. Washington, D.C. , U.S.A.