

Computational Ontology of Network Operations

Alessandro Oltramari and Lorrie Faith Cranor
CyLab, Carnegie Mellon University
Pittsburgh, USA

Robert J. Walls and Patrick McDaniel
Department of Computer Science
Pennsylvania State University
University Park, USA

Abstract—In this article we outline an ontology of secure operations in cyberspace, describing its primary characteristics through some basic modeling examples. We make the case for adopting a rigorous semantic model of cyber security to overcome the current limits of the state of the art, namely lack of comprehensive knowledge representation and effective automatic reasoning functionalities.

Keywords— cyber security, ontology, situation awareness.

I. INTRODUCTION

In computer network defense, human responders play a role as central as that of the intrusion detection systems (IDS) they use. In this regard, only a holistic approach to human-machine interaction in the cyber environment can improve the situation awareness and the decisions of cyber analysts [1]. By and large, the human factors of cyber security rely on the perception of the cyber elements into play and on the explicit representation of their semantics [2]. This article, in particular, focuses on the second aspect: we overview an ontology of cyber operations, which is instrumental for cyber defenders to better comprehend, predict and prevent cyber attacks. The article is organized as follows: Section II builds the case for the adoption of ontologies in the cyber security realm; Section III outlines the structure of an ontology for the ARL Cyber Collaborative Research Alliance¹ program, focusing on domain specific examples; Section IV outlines future research agenda.

II. RELATED WORK

Every science strives to build rigorous models of specific phenomena [3]: accordingly, the object of a science of cyber security is the “cyberspace”, conceived as a dynamic series of computer network events [4]. Inasmuch as ontologies are formal models of a domain, building ontologies of network events (and related properties) is critical for the transformation of cyber security into a science. In 2010, the DoD sponsored a study to examine the theory and practice of cyber security, and evaluate whether there are underlying fundamental principles that would make it possible to adopt a more scientific approach. The study team concluded that the most important requirement would be “the construction of a common language and a set of basic concepts about which the security community can develop a shared understanding. A common language and agreed-upon experimental protocols [5] will facilitate the testing of hypotheses and validation of concepts”. The need for controlled vocabularies and ontologies to make

progress toward a science of cyber security is also recognized in [6] and [7]. From our point of view, where the human component is essential, the analysis needs to be focused on the different roles that attackers, users, defenders, computer systems and enforced policies play in the context of cyber security. Although in recent years there have been a growing number of papers on ontologies for cyber security and cyber warfare, the reported work is by and large conceptually incomplete. To the best of our knowledge, Obrst and colleagues [8] provide the most comprehensive description of the requirements for a wide-ranging ontology of cyber security, whose vision has actually inspired our paper (the scale of the project and its difficulties are also discussed in [7]). Efforts that have been made toward developing ontologies of cyber security typically do not utilize existing middle-level models such as UCORE ontology². The most important step in understanding a complex new domain involves producing accessible terminological definitions and classifications of entities [6]: discussions on cyber security often begin with the difficulties created by misused terminology (such as characterizing cyber espionage as an attack). In this regard, the Joint Chiefs of Staff created a list of cyber term definitions that has been further developed and improved in a classified version³. None of these definitions, however, are formulated as an ontology. Likewise, various agencies and corporations (NIST⁴, MITRE⁵, Verizon⁶) have formulated enumerations of types of malware, vulnerabilities, and exploitations. In particular MITRE, which has been very active in this field, maintains two dictionaries, namely CVE (Common Vulnerabilities and Exposure⁷) and CWE (Common Weakness Enumeration⁸), a classification of attack patterns (CAPEC - Common Attack Pattern Enumeration and Classification⁹), and an XML-structured language to represent cyber threat information (STIX - Structure Threat Information Expression¹⁰). A brief discussion of an ontology for DDoS attacks and a general ontology for cyber warfare are discussed in [9] and [10].

² <http://www.slideshare.net/BarrySmith3/universal-core-semantic-layer-ucore/>

³ <http://publicintelligence.net/dod-joint-cyber-terms/>

⁴ <http://www.nist.gov/>

⁵ <http://www.mitre.org/>

⁶ <http://www.verizon.com/>

⁷ <https://cve.mitre.org/>

⁸ <http://cwe.mitre.org/>

⁹ <https://capec.mitre.org/>

¹⁰ <https://stix.mitre.org/language/version1.1.1/>

¹ <http://www.arl.army.mil/www/default.cfm?page=1417>

III. A THREE-LEVEL ONTOLOGY FOR THE CYBER-SECURITY RESEARCH ALLIANCE

Regardless of the important role played by the initiatives mentioned in the previous section, without a shared formal semantics the sprawling definitions they bring about are hard to maintain and port into machine-readable formats. In order to overcome this problem, in the context of the Cyber Collaborative Research Alliance we are developing CRATELO, a three-level modular ontology of cyber security. CRATELO is constituted of a suite of integrated domain ontologies (collectively indicated as OSCO), designed on the basis of DOLCE top level [11] extended with a security-related middle ontology (SECCO). These top, middle and domain level ontologies add up to 330 classes¹¹, connected by 162 relationships (132 object properties and 30 datatype properties) and encoded in OWL-DL. The logical expressivity of CRATELO is SRIQ, a decidable extension of the description logic SHIN (more details in [12]).

For reasons of space, in the remaining of the paper we limit ourselves to describe examples from OSCO (we refer the reader to [17] for a thorough account of CRATELO's architecture).

A. Ontology of Secure Cyber Operations (OSCO)

The purpose of OSCO is to model cyber operations into a framework of meaningful and reusable knowledge patterns that can improve the situation awareness of analysts. But what is a cyber operation? In a document released in 2010, the Joint Chiefs of Staff describes it as the "employment of cyber capabilities where the primary purpose is to achieve objectives in or through cyberspace". This definition includes computer network operations and activities to operate and defend the Global Information Grid" [13]. Drawing on this broad concept and relying on DOLCE and SECCO [17], in OSCO we represent a CYBER_OPERATION ψ as an OPERATION *executed* by a CYBER_OPERATOR φ , who can play either the role of DEFENDER in a DEFENSIVE_CYBER-OPERATION or the role of ATTACKER in an OFFENSIVE_CYBER-OPERATION. In the context of cyber security we can also distinguish between those OFFENSIVE_CYBER-OPERATIONS whose MISSION-PLANS satisfy the OFFENSIVE_REQUIREMENT of remaining undetected, and those that don't: we use the class CYBER_EXPLOITATION to denote the former, and CYBER-ATTACK for the latter¹². As Lin points out [14], from a technical viewpoint cyber attacks and cyber exploitations are very similar: they use the same access paths and focus on the same vulnerabilities. The difference is on the delivery and execution of the PAYLOAD that must be performed undetectably in CYBER_EXPLOITATIONS (e.g., port scanning or SQL injections). In a simplified scenario where an SQL injection attack is launched, a defensive cyber operation of INTRUSION_DETECTION can be divided into four essential sub-actions (and corresponding tasks): 1) block the IP address of the attacker; 2) escalate the level of response; 3) block all external connections and 4) redirect the incoming traffic to a honeypot for further inspection. A team of cyber analysts with different duties and privileges performs those actions: for

instance, labels like 'L1', 'L2', 'L3', etc. usually indicate the incremental levels of expertise of cyber analysts. Accordingly, 1) would only be performed by L1 analysts; 2) can only be performed by L1 analysts toward L2 analysts or by L2 toward L3; 3) can only be executed by L2 analysts and 4) only by L3. Gauging which action fits better the situation is not a one-shot decision, but rather a multi-stage evaluation process where the situational awareness of cyber analysts frequently changes. Also, each of those sub-actions has incremental costs and inversely proportional risks: for instance, if blocking all the connections to a web server eliminates the risks of a reiterated attack, suspending the network traffic has a severe impact on the system functionality (e.g., no data access for authorized third parties): escalation, in this context, is an effective means to prevent risk mismanagement. Although this simplified scenario gives only a partial account of the actions that actual analysts have at their disposal, using an ontology of cyber security like CRATELO to model intrusion detection can represent a mean to improve situational awareness and fill the *semantic gap* in our understanding of the cognitive demands in the cyber world [15]. Let's now illustrate a specific example of intrusion detection at the machine level.

1) A multi-level attack

In this section we guide the reader through an ontological analysis of an intrusion detection example, including features such as network topology, address space, communication protocols, etc., which a cyber analyst is commonly aware of.

Let's assume that the outcome of a forensic investigation conducted over a given computer network called *MyNetwork* is the following: within a series of thirty normal operations observed during a certain interval of time Δt , and collectively indicated as *sequence_detection1*, three events, namely *event-gn84*, *event-sx00f* and *event-px-5c*, were specifically classified as cyber threats by an IDS. In our example, the bulk of this information is stored in OSCO at time $\Delta t+1$ and can be thereafter retrieved by submitting an appropriate SPARQL query to the knowledge base¹³. Figure 3 indicates a query devised to assess the succession of the events ("order"), as well as the nature of the launched attacks (hacking of a workstation connected in *MyNetwork*, unauthorized access to *MyNetwork*, defacement of a website hosted in *MyNetwork*).

In the real world this kind of identification processes are operated by network intrusion prevention systems (NIPS): for the sake of the example we hypothesize that the information about our multi-stage attack was gathered using Snort, a popular open source NIPS. It's out of scope to present here all the architectural features and functionalities of Snort¹⁴: what is mostly relevant for illustrating our example is Snort's key components, i.e. "detection rules", according to which targeted actions are executed as a consequence of particular conditions being verified by the NIPS rule engine. In particular, we assume that *SNORTrule-a1* was triggered at Δt : as Figure 2 illustrates, this detection rule specifies that, if there is a communication coming from a source node with IP address '108.200.181.118' to a destination node inside

¹¹ Figure 5 contains a partial visualization of CRATELO's taxonomy.

¹² As the example exposes, one of the key design principles underlying CRATELO is to separate cyber operations from the abstract generalizations used to describe them, i.e., plans, tasks, requirements.

¹³ <http://www.w3.org/TR/rdf-sparql-query/>

¹⁴ We refer the reader to the documentation provided at: <http://www.snort.org>

MyNetwork with IP address ‘192.168.3.12’, this state of affairs flags an alert in the defense system (we postulate that the source’s IP identifies a hacker located in a known ‘rogue state’). Note that IP addresses are associated to the source and the destination nodes via a suitable OWL datatype property ‘has_IP_address’. The same modeling choice has been applied for representing the connection ports used during data transmission. Figure 1 visualizes how Snort rules are modeled in OSCO.

It’s important to highlight that the reason why *SNORTrule-a1* is an instance of type *COMMUNICATION* is that OSCO focuses on triggered detection rules, namely rules that have been fired by Snort as a consequence of changes in the cyber environment. In other words, OSCO models only those computer network events of *sequence_detection1* that are associated with specific threats or attacks. From the standpoint of usability this shows that ontology-based reasoning must not be seen as a replacement for rule-based detection, but rather as an additional tool for alert correlation. Although Snort rules are generally more complex than the one overviewed in this section, what is interesting is that, in CRATELO, low-level events occurring at the network level can be represented together with high-level decisions performed by cyber analysts in the *INTRUSION_DETECTION_CYBER_OPERATION*. Note that, by enriching CRATELO with representation of Snort rules, the OWL encoding is extended with DL-safe rules, which are SWRL rules (Semantic Web Rule Language¹⁵) modified to keep the ontology logically decidable.

The snapshot in Figure 2 shows that CRATELO can be already used in combination with a Protégé¹⁶ built-in reasoner Hermit 1.3.8¹⁷ to demonstrate some basic inference functionalities. Note that both the inference-based query mechanism underlying the example requires about 2 milliseconds in a MacBook Pro with 2.3 Ghz Intel Core i7 and 16 GB RAM.

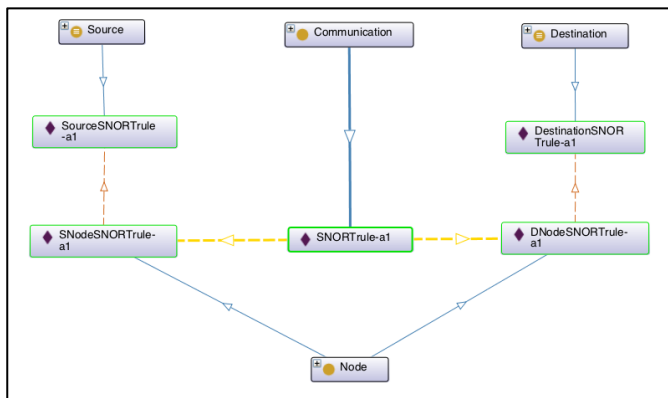


Figure 1: The components of *SNORTrule-a1*. Boxed diamonds indicate instances; boxed circles represent classes; blue arcs indicate ‘instance-of’ relation; yellow arcs ‘has participant’ relation; brown arcs ‘has role’ relation.

2) A question of trust

As an example of how trust can be studied in cyber operations, let’s consider a measurement of the network delay in *MyNetwork* during *sequence_detection1*. As we know, the delay within a network specifies how long it takes for a bit of data to travel from one node to another. If we assume that an acceptable delay in a network communication has a value included within 0.1 and 0.215 ms, this implies that any out-of-range measurement should make data communication unreliable, “therefore” untrustworthy. Note that this implication represents a simplified notion of trust adopted for the sake of this example, and that doesn’t take into account dynamic dimensions of operations like network topology, latency, data rates etc.

In this setting, OSCO defines *RELIABILITY* as a quality associated to two dimensions in the *COMMUNICATION_TRUST_SPACE*, i.e., ‘*BIT_ERROR_VALUE*’ and ‘*NETWORK_DELAY_VALUE*’. This characterization of communication reliability as an attribute of trust conceptualized into a bi-dimensional space is not OSCO-specific, but is grounded on CRATELO top-level ontology, *DOLCE*. *DOLCE* is designed to capture the conceptual primitives underlying natural language, commonsense, and naïve psychology. Accordingly, qualities are conceived as ‘inherent in’ other entities and ‘associated with’ specific values. For example, ‘shape’, ‘size’, ‘color’, ‘weight’, ‘sound’, ‘smell’ are quality types, while ‘triangular’, ‘small’, ‘red’, ‘50 pounds’, ‘70 Hz’, ‘bitter’ are value types. The relation of inheritance in *DOLCE* indicates that the color exhibited by an object (a specific quality) is treated as different from its individual color (a specific value). Further examples can be made by considering physical magnitudes, such as the diameter of the Moon and the measure of 2159.2 miles or the frequency range of the human voice and the related interval 500-2000 Hz.

If in *DOLCE* quality values denote the position of an individual quality within a conceptual space, by applying these structural distinctions to OSCO, we can model trust as a quality of the class *COMMUNICATION*. In principle trust can be represented by means of a wide spectrum of conceptual dimensions, but in our example we only focus on ‘*BIT_ERROR_VALUE*’ and ‘*NETWORK_DELAY_VALUE*’. This scenario is partially visualized in the bottom part of figure 4. We attributed a property called ‘*ReliabilityMyNetwork*’ to *MyNetwork*, assigning the value 0.3 ms to it, which is greater than the maximum delay as previously stated. After initializing the automatic reasoner Hermit in the Protégé, OSCO consistently classifies ‘*NetworkDelayMyNetwork*’ as ‘untrustworthy’, explicitly representing that the specific delay is ‘associatedWith’ the ‘*ReliabilityMyNetwork*’ individual quality. This inference, highlighted in Fig. 4 with a pale yellow mark, is derived by the dichotomy embedded in the trust space. More specifically, from a technical standpoint, this result is obtained using a closure axiom on ‘*NetworkDelayValue*’, a formal tool that constrains the kinds of sub-types a class can have (in our case either trustworthy or untrustworthy ‘*NetworkDelayValue*’). Similar arguments and

¹⁵ <http://www.w3.org/Submission/SWRL/>

¹⁶ <http://protege.stanford.edu/>

¹⁷ <http://hermit-reasoner.com/>

examples apply to other attributes of trust (i.e., availability, confidentiality, integrity, certainty) across domains. For instance, ‘Privacy’ is a component of the quality ‘Confidentiality’ in a social network, and can be represented by different dimensions in a trust space, from values of password strength to the conformity of biometric parameters.

We are currently working on an extension of OSCO that includes a model of risk parameters and system vulnerabilities similar to the one just presented for trust attributes, making extensive use of the ontological pattern quality-quality space-quality value.




Figure 2: The inferences showing how *alert1* is triggered by *SNORRule-a1*.

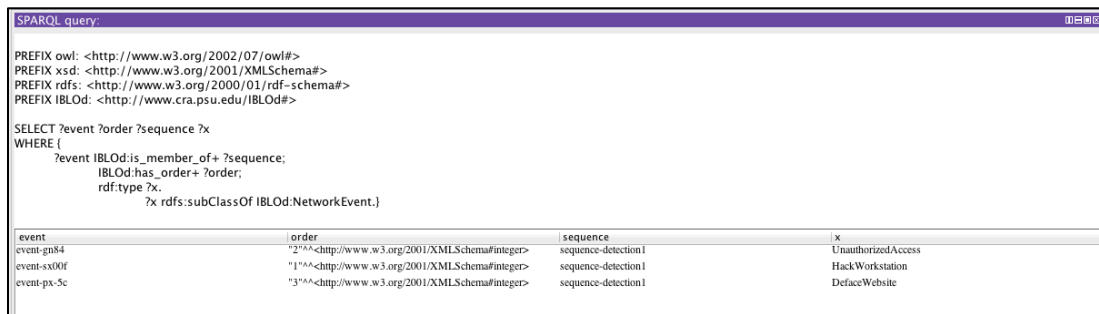


Figure 3: A SPARQL query (top-part) and the results returned (ordered components of *sequence-detection1*).

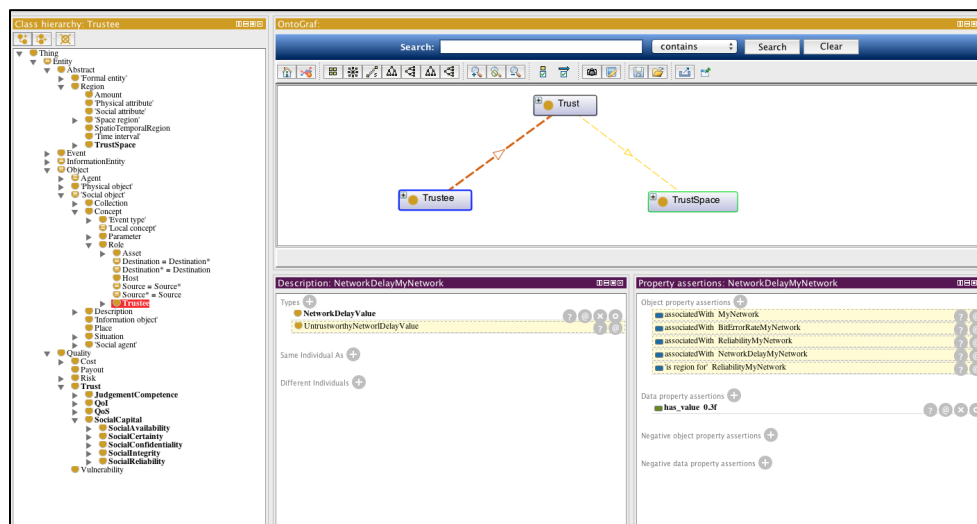


Figure 4: From the left to the bottom (clockwise): 1. the backbone taxonomy of trust dimensions; 2. the core relational schema formed by ‘has dimension’ relation (yellow dotted arc) and ‘associatedWith’ (brown line); and 3. the logical inference underlying the ‘MyNetwork’ example.¹⁸

¹⁸ Figure 1 and 4 were generated and exported using Ontograf (<http://protegewiki.stanford.edu/wiki/OntoGraf>), a visualization plug-in for Protégé. Even within the same ontology, Ontograf automatically assigns different colors to arcs when a new figure is created: this explains mismatch of colors between the two figures.

IV. CONCLUSIONS AND FUTURE WORK

Notwithstanding the proliferation of taxonomies, dictionaries, glossaries, and terminologies of the cyber landscape, building a comprehensive model of this domain remains a major objective for government agencies, private organizations, researchers and intelligence professionals. There are multiple reasons behind the discrepancy between demand and supply of semantic models of cyber security. Although we cannot thoroughly address this topic here, we are firmly convinced that a great part of the problem is the lack of balance between the ‘vertical’ and the ‘horizontal’ directions of the effort. From one side, state of the art consists of several classifications of the domain, as argued in Section II: these efforts typically yield rich catalogs of cyber attacks, exploits and vulnerabilities. On the other side, a rigorous conceptual analysis of the entities and relationships that are encompassed by different cyber scenarios would also be needed, but little work has been done on this horizontal dimension (if we exclude the ongoing MITRE initiative described by Leo Obrst and colleagues in [8]). In this paper we placed ourselves on the second perspective: instead of presenting “yet another” catalog of cyber notions, an endeavor that remains however of undisputable relevance for the cyber security community, we decided to explore the area of network operations.

Our investigation addresses cyber operations as complex entities where the human factor is as important as the technological spectrum: our ontological analysis is grounded on a bedrock of foundational concepts bond to the domain of cyber operations through an intermediate layer of core security notions.

Future work will focus on the following research steps:

- populating OSCO with a large set of cyber operations documented in the literature and learned from real-world case studies, extending the modeling primitives to risk parameters and vulnerabilities;
- designing and customizing a methodology for ontology validation based on ‘competency questions’ submitted to domain experts (along to what has been proposed in [16]);
- running cyber warfare simulations within military exercises, collecting data to be modeled with CRATELO;
- studying ontology mappings between CRATELO and other semantic models (e.g., MITRE’s Cyber Ontology Architecture and UCORE ontology), ensuring interoperability and reusability of the resource.

We are aware of the challenges ahead of us in pursuing this research agenda, which would usually be very difficult to implement. Nevertheless, we’re also persuaded that, in the broad vision framed by the ARL Cyber Collaborative Research Alliance, what we have described illustrates a realistic work plan and a necessary step toward the foundation of a science of cyber security.

ACKNOWLEDGMENTS

The authors want to thank Noam Ben-Asher and Jin Hee-Cho for their valuable help in concocting the example discussed in section III.

This Research was sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-13-2-0045 (ARL Cyber Security CRA). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

REFERENCES

- [1] P., Rivera, B., Swami, A. McDaniel, "Toward a Science of Secure Environments," *Security and Privacy*, vol. 12, no. 4, pp. 68-70, July/August 2014.
- [2] M.R. Endsley, "Toward a Theory of Situation Awareness in Dynamic Systems," *Human Factors*, vol. 37, no. 1, pp. 32-64, 1995.
- [3] M. Bunge, *Causality and Modern Science*. New York: Dover Publications, 1979.
- [4] A. Kott, "Towards Fundamental Science of Cyber Security," in *Network Science and Cybersecurity*, R. E. Pino, Ed. New York, 2014, vol. 55.
- [5] The MITRE Corporation, "Science of Cyber-Security," The MITRE Corporation, McLean, VA, Technical 2010.
- [6] D. A. Mundie and D. M. McIntire, "The MAL: A Malware Analysis Lexicon," CERT® Program - Carnegie Mellon University , Technical 2013.
- [7] Randall Dipert, "The Essential Features of an Ontology for Cyberwarfare," in *Conflict and Cooperation in Cyberspace - The Challenge to National Security*, Panayotis A Yannakogeorgos and A. B. Lowther, Eds.: Taylor & Francis, 2013, pp. 35-48.
- [8] L., Chase, P., & Markeloff, R. Obrst, "Developing an ontology of the cyber security domain," in *Seventh International Conference on Semantic Technologies for Intelligence, Defense, and Security*, 2012, pp. 49-56).
- [9] I. Kotenko, "Agent-Based modeling and simulation of cyber-warfare between malefactors and security agents in internet ," in *19th European Conference on Modeling and Simulation*, 2005.
- [10] A., Buchanan, L., Goodall, J. & Walczak, P. D’Amico. (2009) Mission impact of cyber events: Scenarios and ontology to express the relationship between cyber assets. [Online]. <http://www.dtic.mil/cgi-bin/GetTRDoc?AD=ADA517410>
- [11] C., Borgo, S., Gangemi, A., Guarino, N., Oltramari, Schneider, L. A. Masolo, "The WonderWeb Library of Foundational Ontologies and the DOLCE ontology," Laboratory For Applied Ontology, ISTC-CNR, Technical

Report 2002.

- [12] I., Kutz, O., Sattler, U. Horrocks, "The Irresistible SRIQ ," in *OWLED '05 - "OWL: Experiences and Directions"*, vol. 188, Galway, 2005.
- [13] Joint Staff Department of Defense. Joint Terminology for Cyber Operations. [Online]. http://afri.au.af.mil/cyber/Docs/panel1/Cyber_Lexicon.pdf
- [14] H. Lin, "Escalation Dynamics and Conflict Termination in Cyberspace," *Strategic Studies Quarterly*, vol. 6, no. 3, pp. 46-70, Fall 2012.
- [15] C., Ben-Asher, N., Oltramari, A., Lebiere, C. Gonzalez, "Cognitive Models of Cyber Situation Awareness and Decision Making," in *Cyber Defense and Situational Awareness*, A., Wang, C., Erbacher, R. Kott, Ed.: Springer, 2014, vol. 62.
- [16] S., Ekelhart, A. Fenz, "Formalizing Information Security

Knowledge," in *th International Symposium on Information, Computer, and Communications Security (ASIACCS '09)*, New York, pp. 183-194.

- [17] A. Oltramari, L.F. Cranor, R. Walls, P. McDaniel, "Building an Ontology of Cyber Security" in *9th International Conference on Semantic Technologies for Intelligence, Defense and Security (STIDS 2015)*, Fairfax, pp. 54-71.

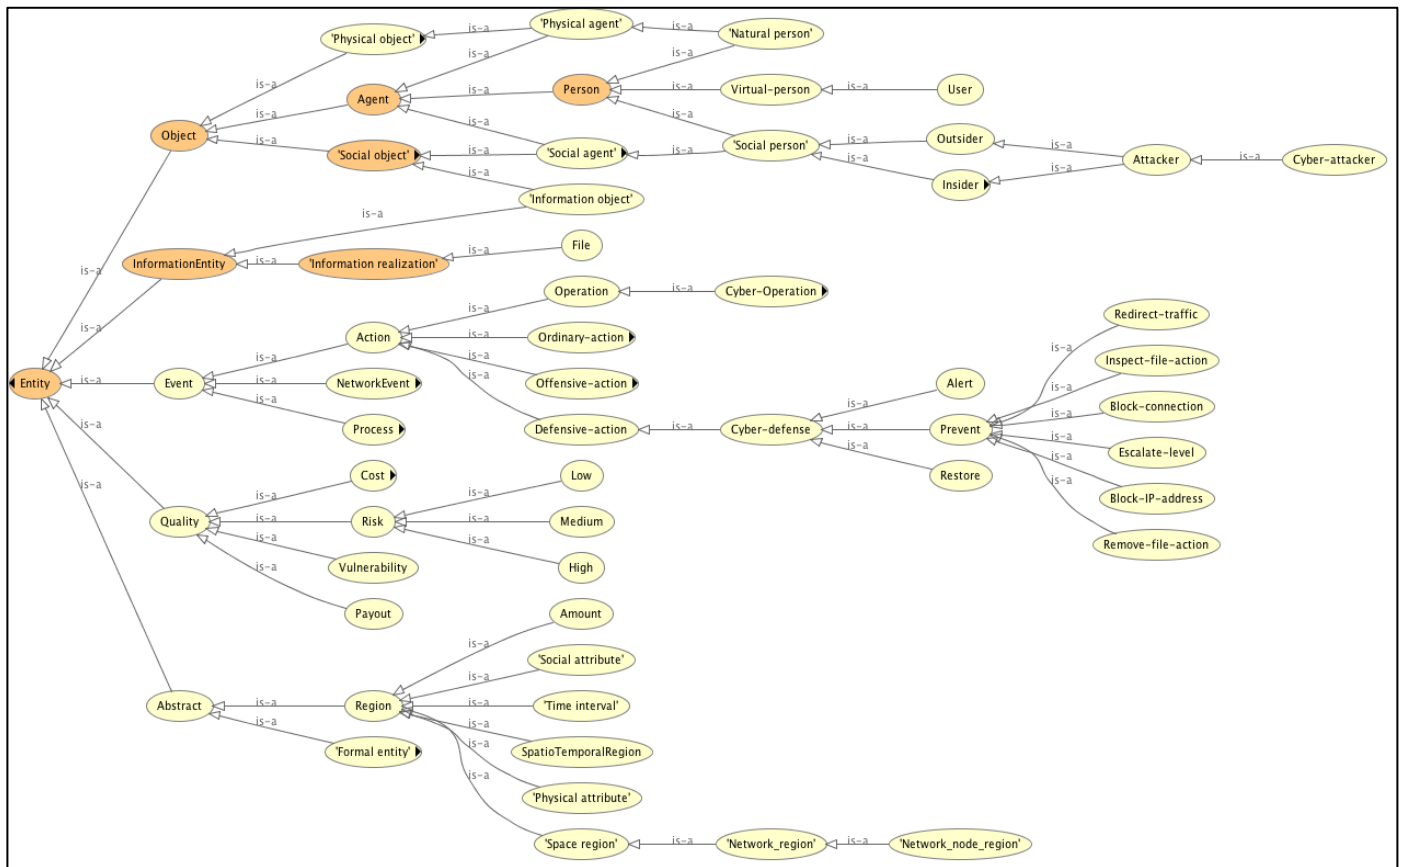


Figure 5: a representative view on CRATELO taxonomy, including classes from DOLCE top-level (far left side), SECCO (central part) and OSCO (far right side).