

Beyond Risk Propagation: A Unified Approach

Takeaways From the Cybersecurity Domain

Alessandro Mosca^{1,*}, Mattia Fumagalli^{1,*}, Gal Engelberg^{2,4,*}, Victor D. Corvalan³, Dan Klein², Pnina Soffer⁴, Diego Calvanese¹ and Giancarlo Guizzardi⁵

¹Free University of Bozen-Bolzano, Bolzano, Italy

²Accenture, The Center of Advanced AI, EMEA

⁴University of Haifa, Haifa, Israel

³Accenture CIO, Buenos Aires, Argentina

⁵Semantics, Cybersecurity & Services (SCS), University of Twente, The Netherlands

Abstract

Risk propagation encompasses a plethora of techniques for analyzing how risk spreads in a given system. Due to the complexity and variety of the domain of application, risk propagation turns out to be a conceptually complex notion. So far several design and implementation solutions in this area have focused on how risk can be quantified, and in what sense it can be propagated in a network of correlated events. However, situations that are usually considered for the propagation of risk involve key concepts of different types, which are rarely limited to a chain of events and their probabilities. In this paper, we provide a novel account of risk propagation via an ontology-driven approach. The proposal stems from a well-founded ontological analysis and aims at modeling the phenomenon of risk propagation according to multiple epistemic dimensions, which involve objects, assets, the agents involved, and their objectives. We test our approach on an implementation and we show how the proposed solution can be used to aid in addressing multiple risk analysis tasks, including a demonstrative case from the cybersecurity domain.

Keywords

Risk propagation, Risk modeling, Risk assessment, Conceptual modeling

1. Introduction

In our daily life, we make a myriad of decisions. A key factor in any decision is the *management of risk*. Risk management requires us to understand whether and how we can reach some goals, how the plans for achieving these goals can be affected, or if unwanted events can happen that would dent these goals. Moreover, we want to assess the likelihood that these unwanted events occur, so that we can monitor and perhaps prevent their occurrence, as well as predict and mitigate their potentially negative impact. Making a decision involves thinking about what may happen in the future and weighing the consequences of our choices against each scenario we believe to be possible.

Deciding under conditions of uncertainty is not a trivial task. For this reason, assessing the risk and making decisions when we have a large number of events and multiple dependencies among them, require dedicated tools that allow for a quantitative analysis of probabilities. An application for integrating the quantification of risk and its probability into decision-making processes is provided by what is usually called *Risk Propagation* (henceforth RP) [1]. RP encompasses a plethora of techniques and approaches that aim to enable the management of complex problems, involving several variables, probabilities, events, objects, and relationships between them. Typically, RP solutions involve the implementation of a probabilistic model (like, for example, a *Bayesian network (BN)* [2] or a *Fault*

Proceedings of the Joint Ontology Workshops (JOWO) - Episode XI: The Sicilian Summer under the Etna, co-located with the 15th International Conference on Formal Ontology in Information Systems (FOIS 2025), September 8-9, 2025, Catania, Italy

*Corresponding author.

✉ almosca@unibz.it (A. Mosca); mfumagalli@unibz.it (M. Fumagalli); gal.engelberg@accenture.com (G. Engelberg); victor.d.corvalan@accenture.com (V. D. Corvalan); dan.klein@accenture.com (D. Klein); spnina@is.haifa.ac.il (P. Soffer); diego.calvanese@unibz.it (D. Calvanese); g.guizzardi@utwente.nl (G. Guizzardi)



© 2025 Copyright © 2025 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0)

Tree [3, 4]) for analyzing how risk spreads in a given system, which is then used in risk analysis and management to calculate the cascading effect of risk within a system.

RP is, however, a complex notion [5, 6], which when properly explored, unfolds a range of details that should not be underestimated. Firstly, having a clear understanding of what risk precisely is, how it is associated with certain probabilities, how it can be assigned to certain events, and how it may propagate in a network of events having certain kinds of relationships, is not a straightforward task. Moreover, situations that are usually considered for the propagation of risk involve concepts of different types, which are rarely limited to events and probabilities. Whenever we want to consider a risky situation, indeed, we cannot avoid talking about the objects involved, agents, and their objectives, which may also be the subjects of assessment themselves. Similarly, the values involved in the calculation of risk are rarely traceable to a single measure. We have values related to the probability that certain events occur, but also values related to the notion of expected loss or impact, mitigation action costs, vulnerability reduction capacities, values that relate to a certain threshold of vulnerability, and so forth. In this scenario, a general risk propagation approach that covers the multiple dimensions involved in the assessment of risk is still missing.

In this paper, we address this challenge by providing an ontology-driven risk propagation approach that stems from a well-founded ontological analysis of the notion of risk provided in previous work. We also introduce a list of prototypical queries elicited with the active support of domain experts, and a detailed analysis of the solutions that are currently offered within the realm of the so-called risk propagation technologies. As we shall see, our approach will advance the state of the art in this field by proposing a solution that can be used (i) to extend the expressiveness of networks used for risk propagation by covering concepts like assets, agents, objects, events, and their mutual relations, and accounting for the different ways by which risk can be associated with those concepts; (ii) to disambiguate and improve the understandability of risk assessment tasks by making explicit how risk values are connected to other critical values, which, in turn, can be derived from the subjective judgment of multiple assessors. We have implemented our approach in an application and have shown how the proposed solution can be used to aid in addressing multiple risk analysis tasks, including a demonstrative case from the cybersecurity domain.

The remainder of the paper is organized as follows. Section 2 describes the challenges via a running example. Section 3 details our proposal, namely the theory that supports our approach to RP. Section 4 dwells into the prototypical information needs, translated as queries, that emphasize the implications of our approach. Section 5 shows how the theory can be used in practice. In Section 6, we discuss the related work. Finally, in Section 7, we reflect on results and future work.

2. Motivations

The concept of “risk propagation” is notably intricate. Unlike a physical phenomenon that spreads like a virus, moving from one object to another, risk is inherently a *measure*. It is always associated, by a specific subject, with a certain probability and it is the output of an estimation grounded in the notion of “*expected utility*” [7]. In this regard, as discussed in [5] extensively, all approaches named under the label of “risk propagation” can be traced back to some sort of graph models: the value representing risk is here associated with a node (or an edge), and it is updated according to a certain function, which may take as inputs the values associated with neighboring nodes or edges. This is what commonly lies beneath what is meant by “propagation” and, looking at the literature, *BNs* [2], as well as other formalisms like *Markov chains* [8], represent well-established support techniques in this regard.

To elaborate on our point, consider the graph shown in Figure 1. It illustrates a basic framework for distributing risk across a network, with nodes symbolizing typical events in the cybersecurity domain. Notice that this model represents a straightforward and common structure of a BN, readily comparable to the familiar “*burglary and earthquake*” example found in the literature [9]. The graph in question is directed and acyclic. The nodes that do not have dependencies (e.g., *unauthorizedAccess* (U) and *malwareAttack* (K)) are said to be independent nodes. The nodes being the *target* of an edge are said to




Figure 1: An example of a Bayesian network-based risk propagation model.

be dependent on the corresponding *source* nodes. For instance, securityAlert (S) depends both on the U and K events, and the fact that the ‘admin is notified by the alarm’, i.e., iTAdminsCalls, depends on S. In this scenario, the risk propagation mechanism consists of “updating” a risk value ascribed to a certain node according to what *happens* to other nodes. Following the example, the rationale grounding the quantification is quite straightforward. On the one hand, nodes are associated with a *probability value* (e.g., the probability that an U occurs, denoted as $P(U)$, is 0.7 and the probability that a K occurs, denoted as $P(K)$, is 0.2). On the other hand, the risk quantity can be derived as “the product of the probability assigned to the given node and another value, usually quantifying a loss” [10], which is not explicitly represented in the model. In Figure 1, for example, considering malwareAttack as an event that results in a time loss, rated as “5” within a threshold ranging from 0 to 10, one could determine the risk associated with malwareAttack as $P(K) \cdot 5$. Similarly, the event securityAlert can result in another time loss, rated as “7” within a threshold from 0 to 10, and the risk associated with it can be determined as $P(S) \cdot 7$.

Now, in what sense can the risk linked to one node vary depending on the alterations in the risk associated with other neighboring nodes or relationships? For example, how might the risk of securityAlert be influenced *given that* the risk linked to malwareAttack changes? The most explicit way to depict such an influence is to embrace the concept of *conditional probability*. Consider, for example, $P(S | U, K)$ in Figure 1, i.e., the *probability that a security alert occurs given an unauthorized access and a malware attack*. Here the risk of having the securityAlert, is calculated as discussed above, but considering external factors, like, for instance, having a malware attack. Thus, an increase in the probability (and then risk) associated with having the malware attack involves an increase in the probability (and then risk) associated with having the security alert. From this perspective, again, the probability inferences enabled by the graph affect the risk values associated with connected nodes. These encode conditional probabilities and form the backbone supporting the “propagation effect”.

What further complicates matters is that the existing approaches do not even consider the simple distinctions we have conveyed through this basic example. Consider that approaches implemented via BNs, such as the one utilized in the provided example, are among the most transparent. There are alternative approaches that do not even explicitly articulate how probability values are updated across the graph [11, 12, 6]. Probability, loss, and risk are always conflated into one single measure. Moreover, many auxiliary concepts crucial for elucidating the assumptions underpinning risk quantification and update are overlooked, thereby impeding the effective governance of subsequent risk assessment and risk management activities.

3. Modeling Risk Dimensions

In this section, we focus on the concepts that are germane to the disambiguation of risk propagation semantics and the consequent covering of the expressiveness gap we discussed. Figure 2 is a diagrammatic and simplified representation of the model we propose. The model is explicitly inspired by the *Common Ontology of Value and Risk* (COVER) [13]. We took COVER as primitive because (i) it was itself subject to validation and proper comparison to the literature of risk in risk analysis and management at large (e.g., [14, 15, 16]); (ii) it is based on a widely-used foundational ontology, namely the *Unified Foundational Ontology* (UFO) [17]; (iii) it embeds a domain-independent conceptualization of risk; (iv) it is built upon widespread definitions of risk and shows how risk calculation depends on risk assessment.

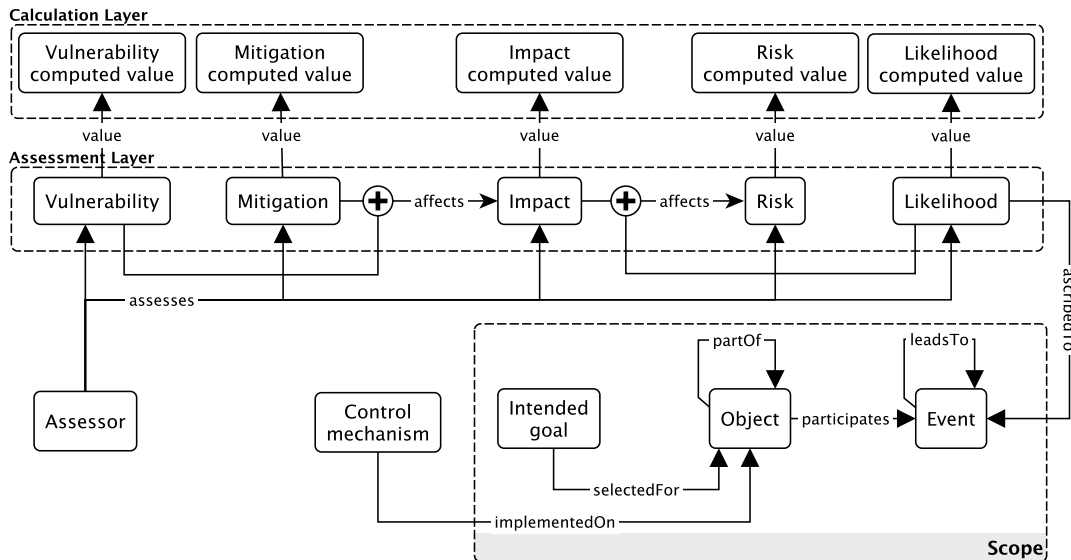


Figure 2: A diagram of the proposed conceptualization. The ‘ \oplus ’ symbols in the diagram denote the *combination* of the input assessment predicates. For example, ImpactAssessment is generated through the combination of VulnerabilityAssessment and MitigationAssessment.

Moreover, (v) COVER has already been connected to different domain ontologies showing its utility in clarifying some connected notions (e.g., *trust*, *prevention*, *security*). Our proposal can be seen as operationalising some concepts provided by such an ontology, in the context of a risk propagation scenario. Before illustrating how the model can be used to represent and reason about risk assessment and risk management scenarios, let us delve into the description of the main concepts composing it, along with the assumptions we adopted for their representation.

The first assumption is that the model we introduce makes use of ‘computed values’ (also called ‘scores’), whose calculation comes from the application of external resources, techniques or methods. As we have seen, BNs can be used to capture both conditionally dependent and conditionally independent relationships between random variables and, therefore, as a tool to compute how likelihood values propagate in complex event graphs. Other examples include the *Common Vulnerability Scoring System (CVSS)*¹ and the *Common Weakness Scoring System (CWSS)*², which are frameworks for consistently scoring vulnerabilities of software applications. However, having a model that works on these computed values taken in isolation is not enough to account for all the information that is combined when experts reason about risk scenarios and evaluate them. Although our model does not explicitly represent the various methods that can be adopted for the calculation of these values, assessments are introduced to explicitly relate these values to the rest of the information that characterizes any risk scenario, that is: (i) the objects participating in an event type, (ii) the subject who provides the assessment, (iii) the intended goals she has in mind while evaluating the scenario, (iv) the control mechanisms (or, mitigation strategies) that may have been adopted. As a consequence, each assessment definition in the model is characterized by the presence of an assessor-dependant ‘computed value’ function, denoted as $[...]ComputedValue_a$ (where ‘[...]’ stands for Likelihood, Vulnerability, Mitigation, Impact, and Risk), returning a value v calculated for the given assessment by the assessor a . We assume that an assessor may have multiple alternative functions at her/his disposal to calculate the necessary computed values, according to the application scenario (examples of such functions include risk-level matrices, vulnerability metrics and the CVSS and CWSS tools above mentioned).

Notice that, according with [13], the present work elaborates on top of the idea that “likelihood is a quantitative concept that inheres in types, not in individuals”. In the UFO terminology, events whose occurrence in the future we assign a probability value can be understood as ‘semi-saturated type’ which

¹<https://nvd.nist.gov/vuln-metrics/cvss/v3-calculator>

²https://cwe.mitre.org/cwss/cwss_v1.0.1.html

are defined by general properties as well as individual concepts. For instance, the type “malware attack against the db_1 ” includes the general type “malware attack” and the individual “ db_1 ”, but it can be instantiated by multiple events. Therefore, when we talk about future events in the rest of the paper, we always refer to types of events.

Our second main assumption is that the evaluation of risk is *experiential*. This means that we always *ascribe risk, or better ‘riskiness’ values, to event types, according to goals that a subject is willing to achieve*. This is why we informally introduce the notion of ‘scope’ as the conjunct of: (i) an event type $Event(e)$; (ii) an object $Object(o)$ (here a generic term for multiple types of organisational assets, such as business process and business model specifications, human and artificial agents, or devices); and (iii) an intended goal $IntendedGoal(a, o, g)$ that has been set by an assessor while typically considering a certain object that participates in an event.

Assessors always focus their assessments on specific goals they have in mind when analysing the riskiness of the occurrence of a given type of event (e.g., a malware attack), for the objects that participate in it (e.g., the db_1): we think of these goals as being defined based on *capabilities* of the objects, which could be possibly harmed by an event in the future. For instance, given a malware attack having the database db_1 participating in it, an assessor can be interested in evaluating the potential impact of such an event w.r.t. the *integrity* of db_1 , while another assessor w.r.t. its *availability*. Clearly, changing the goal (or, focusing on different capabilities) may give rise to different final risk assessments, no matter the fact the event type and participant object are fixed. Coming back to the notion of scope we informally introduced before, scopes overlapping in the event type and object components but having different goals, have to be considered as distinct scopes. The introduced model explicitly accounts for the presence of assessors, as “subjects whose job is to judge or decide the amount, value, quality, or importance of something” [18]. The assessor is the one who defines the scope of the risk analysis, and consequently, the computed values that are more appropriate for a given scenario. In the *risk assessment* case, it is always the responsibility of an assessor to decide which are the *likelihood* and *impact assessments* that have to be combined for the scope at hand, among all the many different likelihood and impact assessments possibly available for that scope³. All computed values have to be understood as being under the responsibility of an assessor: a direct consequence of this choice is that we can explicitly represent scenarios in which an event type, given a participant object and an intended goal, is perceived as risky, by one subject, and as an opportunity, by another.

Given the above assumptions and concept terminology, let us observe that our model induces graphs whose nodes are *event types* and *objects*, which are connected by edges of different kinds. As shown in Figure 2, the *LeadsTo* relationship (the inverse of which is *DependsOn*) represents the connections between events. Event-to-event edges alone are those that usually constitute the backbone structure of a BN. Moreover, according to [19], we assume that events have at least one participant, and participants are Objects. The *Participates* (the inverse of which is *HasParticipant*) edges capture the relationships between objects and events. When in the presence of composite objects and composite events, a *PartOf* relation (the inverse of which is *HasPart*) is used to model how object and events parts relate to their respective wholes [20, 21].

The formal graph resulting from the mutual relationships between event types, objects, intended goals, and assessors is the key structure that allows our model to automatically reason about the update of event probabilities and, consequently, the riskiness evaluation of certain events and the fact that participating objects can be, or not, at risk. As extensively discussed in subsection 5.1, this very same structure can be also leveraged to deploy a query-answering system that allows risk managers to: (i) analyse potentially critical scenarios across the multiple epistemic dimensions connected to the notion of risk, without having to manually combine information gathered from a number of different technologies and information systems (e.g., BNs for the propagation of events’ probabilities, system logs for the discovery of event-to-event correlations, access credentials and documents’ authorship

³Considering risk values as “[...] a function of: (i) the adverse impacts that would arise if the circumstance or event occurs; and (ii) the likelihood of occurrence” (<https://csrc.nist.gov/glossary/term/risk> is pretty common nowadays, and in line with NIST suggested practices (<https://csrc.nist.gov/glossary/term/risk>).

for the identification of the involved assessors and their respective evaluations, network architecture specifications for the detection of participant assets), and (ii) define *ad hoc* strategies to intervene on potentially critical scenarios in terms of previously tested risk treatment and mitigation actions, e.g., by simulating how certain control mechanisms are expected to reduce the vulnerabilities of the assets they have been implemented on and, as a consequence, the expected impacts of given events.

3.1. Assessment Types

The primitive notions of computed values we introduced and related assumptions allow us to represent multiple assessment options.

The first assessment is the one concerning the *likelihood* of an event, namely the LikelihoodAssessment, and it is formally defined as follows:

$$\begin{aligned} & \forall e \text{LikelihoodComputedValue}(\text{lih}_a(e)) \rightarrow \text{Event}(e). \\ & \forall e \text{LikelihoodComputedValue}(\text{lih}_a(e)) \rightarrow \exists a \text{Assessor}(a) \wedge \text{LikelihoodAssessment}(e, a, \text{lih}_a(e)). \end{aligned} \quad (1)$$

The value $\text{lih}_a(e)$ represents the expected probability value for an event type, as elaborated by the assessor a . As discussed before, the LikelihoodComputedValue predicate represents the fact that the value is a likelihood value is computed according to the specific function lih_a chosen by the assessor a (e.g., the value may be derived by the conditional probabilities encoded in a BN or by the application of machine learning techniques to time series data). As will be seen below, each type of assessment depends on different ways of calculating the associated values, which may be different depending on the scenarios and the assessors. For the sake of readability, we assume all the variables appearing in the antecedents of the following definitions to be universally quantified.

The second type of assessment we introduce is the VulnerabilityAssessment, which is modeled as follows:

$$\begin{aligned} & \text{VulnerabilityComputedValue}(\text{vul}_a(e, o, g)) \rightarrow \text{Event}(e) \wedge \text{Object}(o) \wedge \text{Goal}(g). \\ & \text{VulnerabilityComputedValue}(\text{vul}_a(e, o, g)) \rightarrow \exists a \text{Assessor}(a) \wedge \text{Participates}(o, e) \wedge \text{IntendedGoal}(a, o, g) \wedge \\ & \quad \text{VulnerabilityAssessment}(e, o, g, a, \text{vul}_a(e, o, g)). \end{aligned} \quad (2)$$

The intended semantics of vulnerability we assume here is consistent with the definitions provided by different available standards, such as [22, 23, 24], for instance, and can be generally expressed as “*the propensity or predisposition of something to be adversely affected*”. In our setting, an object can show a certain degree of vulnerability with respect to the occurrence of a given event type and a capability that has been identified as a subject goal, only. We can say, for instance, that a database system is highly vulnerable to a certain kind of cyber attack if we have in mind the goal of preserving its *integrity* and, at the same time, that is fully safe for the same kind of attack if we focus our attention on its *availability*. Therefore, an object vulnerability assessment always requires the specification of: (i) the event type the assessor is taking into consideration for the vulnerability assessment, $\text{Event}(e)$; (ii) the object for which one has to evaluate how vulnerable it is, $\text{Object}(o)$; (iii) the subject responsible for the assessment, $\text{Assessor}(a)$; (iv) the fact that the considered object and event must be connected by an occurrence of the participation relation, $\text{Participates}(o, e)$; (v) the goal the assessor is interested in for that object, $\text{IntendedGoal}(a, o, g)$; (vi) a vulnerability computed value, $\text{VulnerabilityComputedValue}(\text{vul}_a(e, o, g))$.

The MitigationAssessment represents our third type of assessment and is modelled as follows:

$$\begin{aligned} & \text{MitigationComputedValue}(\text{mit}_{a_1}(e, o, g, cm)) \rightarrow \text{Event}(e) \wedge \text{Object}(o) \wedge \text{Goal}(g) \wedge \text{ControlMechanism}(cm). \\ & \text{MitigationComputedValue}(\text{mit}_{a_1}(e, o, g, cm)) \rightarrow \exists a_1, a_2 \text{Assessor}(a_1) \wedge \text{Assessor}(a_2) \wedge \\ & \quad \text{ImplementedOn}(cm, o) \wedge \text{IntendedGoal}(a_2, o, g) \wedge \text{MitigationAssessment}(e, o, g, a_1, \text{mit}_{a_1}(e, o, g, cm)). \end{aligned} \quad (3)$$

This type of assessment, even if formally independent from other assessments, usually assumes that an analysis of the possible vulnerabilities of an object in a potentially threatening scenario has already been carried out and supports the explicit representation of the counter-measures that have been put in place to decrease the identified weaknesses. Notice that we assume here the presence of two, possibly distinct, assessors: a_1 , who takes care of the mitigation assessment, and a_2 , who set the goal component of the scope to be analyzed. Control mechanisms cm are ImplementedOn objects (as shown in Figure 2) at a cost, with the main objective of mitigating specific object vulnerabilities. A control

mechanism can be a specific software component when digital devices or applications are concerned, for instance, or a domain-related mitigation strategy in scenarios where natural hazards are the sources of harm in a given ecosystem. Similarly to what happens in the likelihood and vulnerability cases, the `MitigationComputedValue` reports the value representing how much a control mechanism protects the object, given a specific event type and intended goal. In general, we say that control mechanisms mitigate the impact a certain type of event may have on the objects they are implemented on, and for specific goals. For example, if a system is equipped with anti-malware software, this will reduce its vulnerability to malware attacks and, therefore, the expected negative impact such attacks may have in terms of its integrity. On the other hand, the very same control mechanism may not have any positive effect against other types of attack such as unauthorized accesses, if we consider the expected impact on the system integrity.

Once the vulnerability and the mitigation assessments have been provided, the `ImpactAssessment` can be introduced as follows:

$$\begin{aligned} \text{ImpactComputedValue}(\text{imp}_a(v_1, v_2)) &\rightarrow \text{VulnerabilityAssessment}(e, o, g, a_1, v_1) \wedge \\ &\quad \text{MitigationAssessment}(e, o, g, cm, a_2, v_2). \\ \text{ImpactComputedValue}(\text{imp}_a(v_1, v_2)) &\rightarrow \text{ImpactAssessment}(e, o, g, a, \text{imp}_a(v_1, v_2)). \end{aligned} \quad (4)$$

The notion of *impact* in 4 is consistent with existing standards such as [22] where it is defined as “*the magnitude of harm that can be expected to result from the consequences of unauthorized disclosure of information, unauthorized modification of information, unauthorized destruction of information, or loss of information or information system availability*”. Accordingly, the `ImpactAssessment` is modeled as resulting from the quantification of the potential loss as expected by an assessor given a scope (i.e., an event type, an object and an intended goal) and also considering control mechanisms that might be set. As 4 shows, the impact value is derived in our model from the combination of pre-computed information about the vulnerability of the object concerning the given event and goal, and the active presence of a control mechanism eventually implemented. The definition, in order to be meaningful, requires the involved assessments (i.e., the impact, the vulnerability and the mitigation ones) to insist on the same scope. One key observation, which relates to the effective applicability of our model, is that the vulnerability and mitigation values could have been previously assessed by different assessors (i.e., a_1 and a_2). The assessor a is meant to be the one who decides which vulnerability and mitigation assessments have to be considered, and how their respective values must be combined.

The final assessment is the one about risk. We model `RiskAssessment` as follows:

$$\begin{aligned} \text{RiskComputedValue}(\text{ris}_a(v_1, v_2)) &\rightarrow \text{ImpactAssessment}(e, o, g, a_1, v_1) \wedge \text{LikelihoodAssessment}(e, a_2, v_2). \\ \text{RiskComputedValue}(\text{ris}_a(v_1, v_2)) &\rightarrow \text{RiskAssessment}(e, o, g, a, \text{ris}_a(v_1, v_2)). \end{aligned} \quad (5)$$

As mentioned at the beginning of this section, the risk assessment is provided based on the combination of the likelihood value associated with an event type of interest, the impact value computed over the same event type and related to a given object (here again, notice that the likelihood of the event type and the impact it is expected to cause could have been suggested by different assessors, namely a_1 and a_2).

As a final remark, the proposed model implies that assessing the risk of an event type requires the assessor to consider and appropriately integrate various pieces of information. This process should be tailored to the domain of interest and the characteristics of the entities involved. In addition to the likelihood of a particular event type involving an object, the assessor should also take into account pertinent information, such as the expected impact the event type would have on the selected goal. Accordingly, we claim that our model does not take “risk” as a self-standing concept. We rather prefer to say that an event type has a certain level of risk, once a scope is taken into account, and that the assessment of such level is provided by an assessor. In doing that, the assessor who is in charge of the risk scenario evaluation is also the one who combines the likelihood, impact, vulnerability, and mitigation computed values associated with these predicates. Moreover, nothing prevents the risk assessor from re-using values that have been established/assessed by others (artificial or human agents) for these predicates. Therefore, what we understood as risk propagation here is exactly the result of the combination of values and concepts represented by the introduced model.

4. Risk Propagation Assessment Queries

The gap highlighted in Section 2 is supported by an analysis of the literature concerning risk propagation. This analysis was conducted also by exploiting the one provided in previous work [13, 5], which proposes well-founded definitions of the phenomena of risk assessment and propagation. Through our examination, we identified several issues that existing approaches cannot address or can address individually, which would benefit from a holistic approach like the one we are proposing. These issues mainly concern the desirability of accompanying a risk propagation task with operations that may be central to the risk assessment and subsequent decision-making process.

In this context, to exemplify the role of our solution, we explicitly extracted a series of queries from the articles about related work. To select the papers of interest, we first retrieved a list of articles (from 2000 to 2023) having the term “risk propagation” in the abstract or title. We then manually refined the selection by searching for articles addressing related topics, such as “risk spreading” or “risk cascading effect”, and dwelt on the most cited ones. Finally, we applied a process for extracting possible queries of interest, while also taking advantage of automatic support for text extraction. The steps we adopted in this extraction process were as follows: we manually extracted and curated queries from the selected papers; (i) we employed the research paper titled *On the Semantics of Risk Propagation* [5] and the concepts provided by the COVER ontology to provide chat-GPT 3.0 [25] with a uniform background terminology (this step aimed to decrease the potential natural language heterogeneity of the output); (ii) we used the generative model to extract other representative queries addressed in the selected papers, starting from the manually identified ones as examples (which were supplied as part of the prompt); (iii) we eliminated implausible or non-sense queries; (iv) the queries generated through automated means and those extracted from research papers were harmonized into a cohesive formalization to eliminate redundancy; (v) the final compilation of queries underwent thorough a review process involving domain experts in the realm of security and risk assessment. In this last step, we also selected a small subset of queries that we considered prototypical and ideal for demonstrating the potential of our proposal through a proof-of-concept. We have made available the information collected through the process described above⁴. Below we report the list of prototypical queries with a brief discussion of the reference work from which they were extracted.

Q1. *Retrieval of event’s likelihood, and impact values.* In [26], the authors propose an ontology-based BN model to capture the causal relationships between supply chain risk events. The model allows for the retrieval of the *likelihood* and *impact values* of risk events, which are key inputs for effective supply chain risk management. For example, what is the likelihood value associated with securityAlert event? What is the expected *loss* associated with malwareAttack event?

Q2. *Retrieval of objects participating in risky events.* In [27], the authors discuss how to trace back to the objects involved in a chain of events. However, the issue remains unresolved there because the adopted model has the form of a BN and presents nodes only as events. For example, what are the *objects participating* in the malwareAttack event (e.g., software and devices)?

Q3. *Retrieval of risk event pathways given various observations, as an object participating in a root event.* The papers [28, 29] describe techniques for automatically generating attack graphs that illustrate potential multi-stage, multi-host attack paths in enterprise networks. Similarly, the research on risk propagation mechanisms in supply chains [27] demonstrates how risks can cascade through interconnected components. By modelling these risk propagation pathways, it may be possible to retrieve likely risk event sequences given specific observations, which is a valuable capability for risk management. For example, what is the cascading effect of the cryptoLocker participation in a malwareAttack event?

Q4. *Assessing how the likelihood of an event impacts the riskiness of events that are dependent on that event in the given graph.* The elaboration of this query was inspired by the same work inspiring the elicitation of Q3. For example, checking how the likelihood associated with unauthorizedAccess event affects the riskiness of the securityTeamCalls.

Q5. *Assessing the impact value associated with an event, given an observation of a vulnerability associated with some objects participating in that event.* The work of [30] demonstrates the need to assess the effect

⁴<https://purl.archive.org/brp/queries>

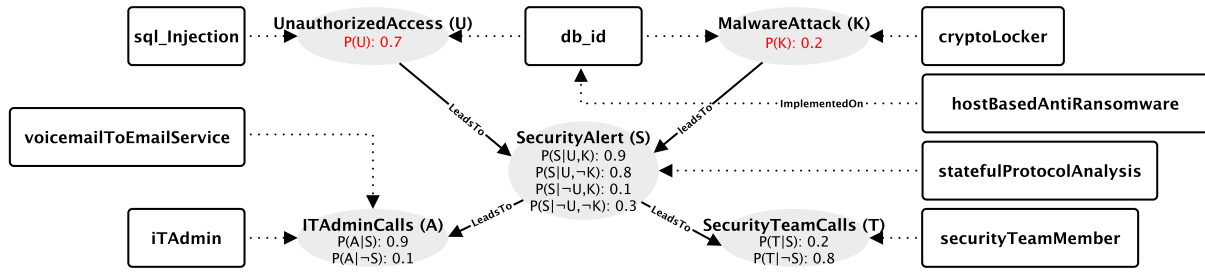


Figure 3: Extended version of the BN from Figure 1, where objects, control mechanisms (white rectangles) and participation relations (dotted edges) have been introduced.

of vulnerabilities on the overall system performance. Additionally, [31] provides a framework for systematically evaluating risks based on the observed vulnerabilities in information systems. For example, checking how a *vulnerability* associated with a database affects the impact value of unauthorizedAccess event, given the fact that the database is a participant in the unauthorizedAccess event.

Q6. Assessing the risk value associated with some given events considering multiple perspectives of a *reference assessor* given her *goals*. In the existing literature, various methodologies have been proposed to quantify risk by considering multiple impact perspectives on the same event. For instance, a single (potentially risky) event may lead to compromises in the integrity, confidentiality, and availability of an object participating in it [6, 32]. For instance, the risk of malwareAttack given *goals* like *database integrity* and/or *database confidentiality*.

Q7. Assessing how *vulnerability values* associated with an object, changed according to different goals, affect the classification of events as *risky*. The elaboration of this query was inspired by the same work inspiring the elicitation of Q6. For instance, considering a scenario in which vulnerability assessment values associated with a database, for the *integrity* and *confidentiality* are different, it turns out to be useful to identify which are the risky events in which that database is involved.

5. Theory in Action

Figure 3 provides a snapshot of a graph encoding the information introduced by Figure 1, once extended with the new possible information that our model allows. In the picture, we now have key elements such as objects (white rectangles) and relations between objects and events (e.g., participates). The figure can be complemented by further information (which currently has been omitted from the figure due to lack of space) regarding assessors, their goals and other elements, such as control mechanisms. Notice that the likelihood and conditional probability values appearing are meant to be provided by a given assessor and that different assessors may fulfil the network with different values. To demonstrate how this data can be combined and exploited, we implemented⁵ the set of rules composing the proposed model along with the assertions representing our running example in ProbLog [33], a well known *probabilistic logic programming language* in which facts and rules can be annotated with probabilities. This also enabled us to show how the resulting theory can be coupled with definite components (e.g., BNs), which are responsible for calculating probability values, as required by risk propagation scenarios.

5.1. Queries Resolution

Let us get back to the queries we introduced in Section 4 and see how the proposed approach can help to address them from a unified perspective.

Q1 concerns the possibility of calculating the probability and/or impact value of a given event. Suppose we want to calculate such a value for the event securityTeamCalls. This can be done by leveraging the predicates LikelihoodAssessment and ImpactAssessment by fixing the target event. Note that these predicates may return multiple data for the same event. On the one hand, LikelihoodAssessment will

⁵The complete implementation is accessible at: <https://purl.archive.org/brp>

return as many values as assessors are involved. For example, if we have two assessors (e.g., Tom and Dick), the query will return the likelihood value assigned by Tom and Dick respectively, where these values are generated via the BN set by the two assessors (e.g., Figure 3).⁶ On the other hand, ImpactAssessment will return as many values as there are available combinations to generate the impact value, which, as from the rule (5) in section 3, always depends on an assessor, a goal and the objects participating in the target event. Q2 concerns the possibility of retrieving objects participating in risky events. As we have seen, the graph encoded by the model can be directly used to navigate event-object connections. The notion of risky event can be naturally captured by defining a threshold from which the event is considered. Suppose that an event is defined as risky when the risk value of the scope it takes part in passes a given threshold (e.g., 0.5). The corresponding rule will be as follows:

$$\text{riskyEvent}(E, O, G) :- \text{RiskAssessment}(E, O, G, RA, RV), RV > 0.5. \quad (6)$$

Accordingly, by querying the riskyEvent predicate we will be able to check what objects participate in risky events. For instance, if malwareAttack has been classified as a risky event, given a certain scope, the query will return the object participating in that scope (see Figure 3).

Q3, in turn, can be addressed by leveraging the event-object graph built upon the participates and the dependsOn relations. If we want, for instance, to retrieve all the events depending on the event in which ransomware, like cryptoLocker, participates, a predicate like the following needs to be defined, with the corresponding query:

$$\begin{aligned} \text{participantBasedDependency}(O, E1, E2) :- \text{dependsOn}(E1, E2), \text{participates}(O, E2). \\ \text{query}(\text{participantBasedDependency}(\text{cryptoLocker}, E1, E2)). \end{aligned} \quad (7)$$

Q4 requires making explicit how the likelihood of an event impacts the *riskiness* of dependent events. This scenario can be tested by leveraging predicates like the one expressed by (7), namely riskyEvent (or atRiskObject, which can be similarly defined) and varying the probabilities encoded by the reference BN(s) connected to the theory. Following our running example we can show that the likelihood of the unauthorizedAccess event (e.g., by adopting 0.4 or 0.1 as the associated likelihood value) affects the riskiness of the securityTeamCalls event. This also impacts the riskiness of objects participating in the involved events. In fact, the riskiness of securityTeamMember object participating in the securityTeamCalls event also depends on the likelihood of the unauthorizedAccess event.

Q5 serves to check the impact assessment associated with an event concerning the vulnerability of the object(s) participating in it. This information can be returned by leveraging the ImpactAssessment predicate. For instance, if we want to check how the vulnerability associated with db₁ affects the impact value of unauthorizedAccess, according to the assessor Tom and given the goal of maintaining its dbIntegrity, we can run query(ImpactAssessment(unauthorizedAccess, db₁, dbIntegrity, tom, V)). Note that the value V depends on the mitigation effects of any control mechanism implemented on db₁ at the moment of the assessment.

Q6 requires that the risk value of an event is assessed by making explicit the different assessors and related goals ascribed to the participant objects. This can be addressed by exploiting the RiskAssessment predicate, where variables RA, RV and G (see rule (6)) return the desired values. Note that for a given event (e.g., unauthorizedAccess), a given assessor (e.g., tom), and a given goal (e.g., dbIntegrity) we may have multiple risk values, since, potentially, the impact and the likelihood values may be provided by other assessors. For instance, the likelihood associated with the same event can be calculated via a BN designed by tom or via a BN designed by another assessor assigning different probability values. Through the different predicates of the theory (e.g., RiskAssessment and RiskComputedValue), and as further explained in Section 6.2, this mechanism aims to make the assumptions adopted to derive the risk values as explicit as possible, thus safeguarding the transparency and control of subsequent decision-making processes.

Q7 requires assessing how the vulnerability value changes, for a given goal, can affect the classification of events as risky. For instance, consider the scenario in which vulnerability assessment values are

⁶Note that this can be handled via ProbLog with the subquery construct.

associated with the object `db_id`, for the *integrity* and *confidentiality* goals. By leveraging the `riskyEvent` predicate it is possible to check what events are classified as risky, both focusing on the integrity and confidentiality of the database. Note that, by playing with different values for integrity and confidentiality, the set of events is classified as risky changes. The reason for this is that a higher vulnerability of the database concerning its confidentiality capability gives rise to higher impact assessment values and, therefore, to a higher risk assessment value, which may or may not exceed the risky set threshold.

6. Related Work

While there is considerable interest in risk propagation techniques for risk assessment, research on using ontologies for *risk propagation* is limited compared to *risk management*. Notable works such as [34, 31] focus more on risk management than on enhancing risk propagation with ontologies. Therefore, this section concentrates on approaches aligned with our specific goal.

Focusing on exploiting an ontology to support the risk propagation task, the work in [26] is highly related to our work. A key contribution of this work was the methodology for constructing BNs to calculate the propagation of event probabilities from an ontology created by domain experts. This ontology guides the model design for risk propagation and assessment, highlighting the importance of ontological knowledge in representing the risk domain. However, it differs significantly from our proposal. Firstly, the ontology is domain-specific, which limits the generalization of the approach to broader scenarios, as acknowledged by the authors. Secondly, the ontology comprises events and their relations only. Other concepts related to risk propagation and assessment are not considered.

Another noteworthy work is [27], which explores risk propagation mechanisms and proposes strategies for sustainable perishable products supply chains. This employs the *Tropos Goal-Risk framework* to model the supply chain and identify risk propagation. Although not an ontology, this conceptual framework supports risk analysis in the requirements phase of software development, embedding ontological analysis [35] related to risk mechanisms and countermeasures. The paper includes a case study of the yoghurt supply chain to validate the proposed approach. Results indicate that the implemented risk management strategies effectively reduce risks and enhance the sustainability of the supply chain of perishable products. Our proposal differs to [27] primarily because it focuses on using an ontology to enable reasoning about the different values involved in risk assessment and propagation tasks. Still, *Tropos Goal-Risk framework* might be integrated into our, to improve the available reasoning capabilities.

The authors of [28] propose a solution for assessing the impact of cyber attacks on military operations using *Cyber-ARGUS*, a simulation framework, and semantic technologies. *Cyber-ARGUS* integrates data from sensors in both the physical and cyber domains for impact assessment. The paper explains how data from sensors is aggregated, how node states are calculated, and how the impact is propagated throughout the network. Here, ontologies are used differently than in our approach, primarily for data representation and integration, representing concepts and relationships within the domains, and fusing sensor data. Semantic reasoning is also used to infer missing information and resolve data inconsistencies.

Although with a different focus, in the cybersecurity domain, [29] introduces an approach to represent and generate attack graphs. This approach uses an attack graph generation tool built on *MulVAL*, a logical programming-based network security analyzer. The generated attack graph encodes inference rules from cybersecurity experts, analyzing the potential pathways an adversary might take to reach a specific target. This work shares our emphasis on the importance of technology for assessing and quantifying risk. However, unlike our method, it does not integrate *ad hoc* technologies to implement risk propagation and query answering. Moreover, it does not account for the multiple values involved in risk quantification and propagation. In response to the same problem, the authors of [32] introduced an approach to assess cyber-attacks' impact on business processes. They generated an interconnected graph depicting dependencies between vulnerabilities on hosts, relations between services to hosts, and

tasks to services. The dependencies were encoded using a Datalog⁷ model extracted via MulVAL. In this work, the method proposed for modeling the impact value relates to ours. However, our approach distinguishes between the set of methods used for generating the values involved in risk propagation and assessment, and the (ontological) theory to be used for enabling reasoning over them. Our approach also extends beyond the cybersecurity domain and addresses risk from a holistic and cross-domain perspective.

7. Conclusion and Perspectives

In this article, we have discussed the limitations in the expressivity of available risk propagation solutions the expressivity. We addressed this challenge by providing a model, grounded on a well-founded ontological analysis of risk and the literature, which can be combined with already existing quantification and propagation mechanisms (e.g., BNs). Our model enables automated reasoning that leverages several typical concepts of risk assessment scenarios (e.g., the assessors, the event participating objects, the intended goals and object capabilities). While it demonstrates reliability in addressing questions such as those listed in section 4, it also has central value in making explicit relevant dimensions that contribute to the analysis of risk, which often are traced back to 'black-box' propagation mechanisms. In the future, we plan to create an application where the proposed model will be interfaced with a series of tools designed to quantify the considered computed values, extend the current conceptualization with new concepts to address a broader set of useful queries, and, subsequently, test the application on real data from different scenarios, with a particular focus on the cyber security domain.

Acknowledgement: Research leading to this work was (partially) supported by the Italian PNRR MUR project PE0000013-FAIR, *Future Artificial Intelligence Research*. Moreover, this research was partially supported by the HEU project CyclOps (GA No. 101135513); by the Province of Bolzano and the FWF through the project OnTeGra (DOI: 10.55776/PIN8884924); by the Province of Bolzano and EU through projects ERDF-FESR 1078 CRIMA, and ERDF-FESR 1047 AI-Lab; by MUR through PRIN project 2022XERWK9 S-PIC4CHU; by the EU and MUR through PNRR project PE0000013-FAIR; by Accenture, The Center of Advanced AI, EMEA.

Declaration on Generative AI

During the preparation of this work, the author(s) used GPT-3 as part of the query elicitation as specified in Section 4, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the publication's content.

References

- [1] J. Jiang, S. Wen, et al., Identifying propagation sources in networks: State-of-the-art and comparative studies, *IEEE Communications Surveys & Tutorials* 19 (2016) 465–481.
- [2] J. Pearl, S. Russell, *Bayesian networks*, 2000.
- [3] S. M. Nicoletti, E. M. Hahn, M. Fumagalli, G. Guizzardi, M. Stoelinga, Watchdog: an ontology-aware risk assessment approach via object-oriented disruption graphs, in: *International Conference on Advanced Information Systems Engineering*, Springer, 2025, pp. 314–331.
- [4] S. M. Nicoletti, E. M. Hahn, et al., Bfl: a logic to reason about fault trees, in: *2022 52nd Annual IEEE/IFIP DSN*, IEEE, 2022, pp. 441–452.
- [5] M. Fumagalli, G. Engelberg, et al., On the semantics of risk propagation, in: *International Conference on Research Challenges in Information Science*, Springer, 2023, pp. 69–86.
- [6] G. Engelberg, M. Fumagalli, et al., An ontology-driven approach for process-aware risk propagation, in: *Proceedings of the 38th ACM/SIGAPP Symposium on Applied Computing*, 2023, pp. 1742–1745.

⁷<https://docs.racket-lang.org/datalog/>

- [7] S. C. Stearns, Daniel bernoulli (1738): Evolution and economics under risk, *Journal of biosciences* 25 (2000) 221–228.
- [8] W.-K. Ching, M. K. Ng, *Markov chains, Models, algorithms and applications* (2006).
- [9] D. Niedermayer, *An introduction to bayesian networks and their contemporary applications*, in: *Innovations in Bayesian networks: Theory and app.*, Springer, 2008, pp. 117–130.
- [10] J. Von Neumann, O. Morgenstern, *Theory of games and economic behavior*, in: *Theory of games and economic behavior*, Princeton university press, 2007.
- [11] K. Shin, Y. Shin, et al., Risk propagation based dynamic transportation route finding mechanism, *Ind. Manag. Data Syst.* (2012).
- [12] G. Kavallieratos, G. Spathoulas, et al., Cyber risk propagation and optimal selection of cybersecurity controls for complex cyberphysical systems, *Sensors* 21 (2021) 1691.
- [13] T. P. Sales, F. Baião, et al., The common ontology of value and risk, in: *ER 2018*, Springer, 2018, pp. 121–135.
- [14] I. Coso, *Enterprise risk management-integrated framework*, Committee of Sponsoring Organizations of the Treadway Commission 2 (2004).
- [15] ISO, *Risk Management - Vocabulary*, ISO Guide 73:2009, 2009.
- [16] ISO, *ISO 31000:2018 - Risk management – Guidelines*, 2018.
- [17] G. Guizzardi, *Ontological foundations for structural conceptual models*, CTIT, Centre for Telematics and Information Technology, 2005.
- [18] Cambridge dictionary, 2024. URL: <https://dictionary.cambridge.org/>.
- [19] N. Guarino, R. Baratella, et al., Events, their names, and their synchronic structure, *Applied Ontology* 17 (2022) 249–283.
- [20] A. J. Cotnoir, A. C. Varzi, *Mereology*, Oxford University Press, 2021.
- [21] R. Allen, *The mereology of events*, *Sorites* (2005) 23.
- [22] NIST: Vulnerability, 2024. URL: <https://csrc.nist.gov/glossary/term/vulnerability>.
- [23] R. Andy, G. Mathias, et al., *The concept of risk in the IPCC 6th assessment report*, 2020.
- [24] UNDRR: Vulnerability, 2024. URL: <https://www.undrr.org/terminology/vulnerability>.
- [25] OpenAI, GPT-3.0, 2023. URL: <https://chat.openai.com>.
- [26] S. Cao, K. Bryceson, et al., An ontology-based bayesian network modelling for supply chain risk propagation, *Industrial Management & Data Systems* (2019).
- [27] X. Deng, X. Yang, et al., Risk propagation mechanisms and risk management strategies for a sustainable perishable products supply chain, *CAIE* 135 (2019) 1175–1187.
- [28] A. De Barros Barreto, P. C. G. da Costa, et al., Using a semantic approach to cyber impact assessment., in: *STIDS*, 2013, pp. 101–108.
- [29] X. Ou, W. F. Boyer, et al., A scalable approach to attack graph generation, in: *Proceedings of the 13th ACM CCS*, 2006, pp. 336–345.
- [30] M. A. Haque, S. Shetty, C. A. Kamhoua, et al., Modeling mission impact of cyber attacks on energy delivery systems, in: *International Conference on Security and Privacy in Communication Systems*, Springer, 2020, pp. 41–61.
- [31] O. T. Arogundade, A. Abayomi-Alli, et al., An ontology-based security risk management model for information systems, *Arabian Journal for Science and Eng.* 45 (2020) 6183–6198.
- [32] C. Cao, L.-P. Yuan, et al., Assessing attack impact on business processes by interconnecting attack graphs and entity dependency graphs, in: *IFIP Annual Conference on Data and Applications Security and Privacy*, Springer, 2018, pp. 330–348.
- [33] L. De Raedt, A. Kimmig, et al., ProbLog: A probabilistic Prolog and its application in link discovery, in: M. M. Veloso (Ed.), *Proceedings of IJCAI 2007*, 2007, pp. 2462–2467.
- [34] G. Nota, R. Aiello, et al., Ontology based risk management, in: *Decision theory and choices: A complexity approach*, Springer, 2010, pp. 235–251.
- [35] Y. Asnar, P. Giorgini, et al., Goal-driven risk assessment in requirements engineering, *Requirements Engineering* 16 (2011) 101–116.