Is the Rule-based Order in AI Dead, or is Still Kicking?

Vassil Vassilev

Abstract

This article focuses on the possibility of revival of the classical rule-based approach in AI for building heterogeneous AI systems. It is based on personal experience while working on several research and innovation projects in diverse areas - business process management, unauthorized intrusion detection, malicious interference protection, digital forensics and diagnostics, data management, etc. Despite their differences, all these areas share something in common: they require multiple operations to be executed in a single transaction and incorporate heuristic rules for different purposes, related to data, knowledge, and operation management. The position of this article is that combining the classical rule-based approach from the early days of AI with more recent developments in AI, such as data-focused machine learning and utility-based reinforcement learning, as well as utilization of the recent technological developments on the cloud and at data centers can be beneficial for widening the real-world application of AI. The challenges, which this complexity creates require joined efforts of academic researchers, industrial engineers, and business enablers. Collaboration between them across the board can be highly beneficial and the author is looking for opportunities in this direction.

Keywords

Heterogeneous AI; Knowledge-based Systems; Semantic Technologies; Heuristic Rules; Integration, Synchronization and Automation

1. Introduction

The history of AI witnessed several shifts in its dominant paradigm - from the initial amazement of the movements of digital amoebas to the decision to embed our knowledge in them to make them more active to the relief of letting them learn themselves to the complete relying on the unknown intelligence exhibited by chatting boxes... Despite these twists and turns, one baseline remained unchanged - the rule-based order still rules in AI!

This paper reflects on the experience of several AI-based projects from recent years completed under the eye of the author at GATE Institute of Sofia University and the Cyber Security Research Centre of London Metropolitan University. They all have something in common - in one way or another, they rely on rules.

The different paradigms adopted in these projects solve different tasks - in decision-making for choosing alternatives for continuation of the operations, in planning the activities for achieving the goals, in controlling the execution of the operations to stay on track, in learning from the environment to improve the outcomes, in looking back at the experience to improve the planning and in explaining the results to bring confidence in the solutions ... This easily leads us to the belief that the way forward is the hybridization and the practical way of uniting multiple paradigms to achieve this is through the use of rules. Of course, this is a complex undertaking, but it does not need to be chaotic.

The paper analyses the experience in hybridization within several projects completed or still underway in Sofia and in London, executed in collaboration between the Cyber Security Research Centre of London Metropolitan University and the GATE Institute of Sofia University. After brief description of several projects, showing the place of the heuristic knowledge in them the paper discusses three fundamental models of heuristic rules, which can meet the requirements of many similar projects. At the end of the paper, the author focuses on three main directions of interest for research and technological development, which are worth investigating further and can be considered an invitation for collaboration.

2. Recent Project Experience

Over the last ten years, the author has been involved in a number of projects which incorporate AI methods in system development. They typically involve some kind of heuristic rules at

different phases of the system lifecycle – design, implementation, deployment, and operation. Below is a short summary of the most important of them from AI perspective.

Assessing the Logical Vulnerability of Transactional Systems. This was the last of a series of projects funded by Lloyds Banking System completed at the Cyber Security Research Centre of London Metropolitan University during 2018-2019 [8]. It aimed at incorporating security threat intelligence into a logical model of financial transaction processing under duress and analyzing the logical vulnerability of the security policies at the design phase. The most important achievement of the project was the development of a method for modeling and analysis of transaction processing, based on combining general-purpose ontology of transactions under security threats with an efficient algorithm for analysis of the security policies which utilizes semantic indexing. In this project several types of rules were utilized: rules for choosing suitable methods for detection, rules for imposing security policies for counteracting security threats, and rules for analyzing the logical vulnerability of security policies.

Threat Intelligence for Unauthorized Intrusion Detection. This project was initiated in the autumn of 2019 by the Cyber Security Research Centre of London Metropolitan University in collaboration with a local company that specializes in managing data centers for commercial businesses using IaaS services provided by AWS public cloud. The objective of the original project was to perform security analytics in real time using data from the company's customers' networks on the cloud. Due to the outburst of the pandemic though, the original project was substantially modified. Thanks to GATE Institute of Sofia University, which provided additional resources the project was revived and continued as a joint project until its completion in 2020 [6]. The security data was generated by simulation in London and sent for processing to Sofia in real-time, where it was analyzed – first in real-time, using statistic methods of correlation, and later on offline, using machine learning algorithms for forensic investigation of the accumulated data. The use of heuristic rules in the project was linked to the orchestration of data management operations along the pipeline for data processing from the moment the data was ingested on the cloud server up to the moment its analysis was concluded and it was accumulated in one of the databases of the cloud server. This project created the prototype of the future GATE Data Platform which in its third enterprise version currently serves the needs of data processing of GATE Institute.

Risk Assessment in Transactional Systems. This project was coined by Lloyds way back in 2019, based on the success of the previous projects, and started as an internal project at the Cyber Security Research Centre of London Metropolitan University, but due to the outbreak of the pandemic, it was completed as a collaboration project at GATE Institute of Sofia University in 2020 [5]. This collaboration leverages combining the technological experience of the Cyber Security Research Centre in ontological modeling of transactional systems with the experience of GATE Institute in decision-making and stochastic optimization. Jointly we developed a new method for modeling transactions under duress by adopting the Partially Observable Markov Decision Process model of stochastic planning (POMDP). Instead of directly solving the problem using standard iterative approximation, though, by restricting the model to the intelligence graphs developed for vulnerability analysis we were able to solve the problem using an efficient recurrent algorithm. Although the heuristic rules were embedded directly in the ontology of the transaction process itself, lots of heuristics were applied manually to configure the model and to interpret the results of the computation. This project currently continues as an innovation project under the Cyber ASAP scheme of Innovate UK with potential commercialization aimed at automation of the default configuration of the model during the modeling phase.

Air Pollution Monitoring and Environment Factors Analysis. This project started in Sofia in 2021, at that time considered to be the most polluted capital in the EU. It served as the first pilot application of the prototype of GATE Data Platform for the provision of comprehensive information in real-time about the state of air pollution in the city using information obtained directly from the monitoring stations across Sofia [4]. From a technological perspective, the most interesting thing was the combination of real-time and offline data processing of both static data, retrieved from a model of the urban infrastructure in ontological format, with external dynamic data, coming from the sensor stations across the city in real-time, and the internal event data, resulting from the interactive operations on the system interface. This concept was later on re-deployed by the same team in 2022 in London, where additional static data about the medical prescription inside the catchment area was used for analysis of the correlation between local air pollution and respiratory diseases [3]. In both cases a bunch of heuristics was used to integrate, synchronize, orchestrate and interpret the data coming from different data sources.

In addition to the above projects, in which the heuristics play an operational role and can be used directly for the automation of various analytical tasks, we have been working on some projects in which the rule-based approach can have a more methodological role. This includes the internal project for developing GATE Data Platform, the EU DiverSea project for developing an integrated architecture for analysis of the maritime biodiversity of EU coastal seas, the DIANA CoDe project for countering disinformation in media space, etc. The common theme in all these projects is the orchestration of the data management workflows, which allows automation by applying a variety of methods for data processing on the data platform, such as data unification, semantic data enrichment, dataflow synchronization, and operation orchestration.

3. Rule Design Considerations

In this section, we will consider three different approaches for modeling transactional systems which differ substantially in their complexity but can be used as a baseline in applications in which there is a need to account for domain-specific heuristics. They have been major research themes in AI and Computer Science for some time on both theoretical and technological levels but the complexity of the problems leaves plenty of space for further investigations. At the same time, they illustrate the different levels of depth in representing heuristic knowledge and its impact on the overall system architecture. Their choice reflects the need to incorporate rules in our own projects, and our own expertise and can be considered an invitation for further research and collaboration.

3.1 State-space Model of the Transactions

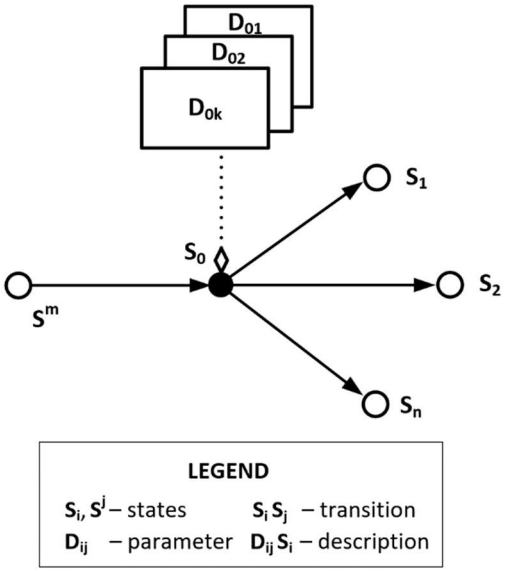

Figure 1: Inventory for Formulating Structural and Parametric Deterministic Rules

The simplest template for formulating policy rules can be constructed in the classical state-space model of dynamic systems, which was introduced in AI in the situation calculus [11]. It adopts a functional view of the actions as mappers of the global *states*, infamously and misleadingly called *situations*. As a sequence of actions, the formal model of the transactions in this approach is a directed graph, leading from the initial state to a commit state - in the case of success, or, to a rollback state instead - in the case of failure. In **Fig. 1** we show a fragment of such a transaction graph, focused on a single node representing a particular state, together with the related nodes. On top of such a model, we can formulate multiple heuristic rules (*policies*) to support the use of the transaction model for various purposes – design, monitoring, synchronization, orchestration, automation, etc.

The rules can be formulated as functional dependencies of the results of the actions on the values of the state parameters. potentially also accounting on the history:

$$\begin{aligned}
&\text{Rule}_{01}\text{: } S_1=f_1\,(D_{01}S_0=V_{11},D_{02}S_0=V_{12},\ldots,D_{0k}S_0=V_{1k},\, S^m)\\
&\text{Rule}_{02}\text{: } S_2=f_2\,(D_{01}S_0=V_{21},D_{02}S_0=V_{22},\ldots,D_{0k}S_0=V_{2k},\, S^m)\\
&\qquad\qquad\vdots\\
&\text{Rule}_{0n}\text{: } S_n=f_n\,(D_{01}S_0=V_{n1},D_{02}S_0=V_{n2},\ldots,D_{0k}S_0=V_{nk},\, S^m)
\end{aligned}$$

where f_i are the transition functions corresponding to the potential actions that change the current state. The rules are usually formulated in a logical format, which allows to combine explicitly AND and OR conditions on the states. The security policies in such a case can be interpreted conveniently as an AND/OR graph.

This template is basic – it does not represent the types of state descriptors, it does not account for potential classifications, and it cannot distinguish synchronous and asynchronous activities in the same model. This limits its use to static tasks during the design phase, which considers only synchronous or asynchronous activities, but it is not suitable for operational control and automation in real-time applications. From a purely theoretical perspective, it also faces several principal problems, such as qualification, ramification, and frame problems [11], which may lead to practical complications.

At the same time, this template is universal, and can be helpful at the initial design, linked to the structural configuration of the workflows and their parametrization. We have investigated a series of structural patterns, providing sufficiently informative ground for specifying structural heuristics concerning the composing and controlling of the workflows. To increase the efficiency of the rule-based systems that use deterministic rules in this format, we also extensively indexed the rules against the states and their parameters. Since this can be done entirely incrementally, it can be very useful for business process management [10]. This approach is also suitable for automation of the scheduling of data processing pipelines on general-purpose data platforms, since it does not account for fine-grained domain-specific knowledge which requires more complex heuristics and can be easily mapped to the tools used to control the data processing pipelines, such as AirFlow we used [1].

3.2 Situation-based Model of the Transactions

The situation-based model of transaction processing is based on several unrelated early research efforts in philosophical logic and linguistic semantics. This approach changes the semantic reference to the world as a point in the global state space and considers it as a set of references relative to the context of discourse instead (called *model sets* in the original Hintikka's approach [12] and *situations* in Barwise and Perry's approach [13]).

This shift from a global absolute reference to a single point in the state space towards a local relative reference to a set of points in the same space, which share common characteristics, is a true revolution. It allows changing the previous functional interpretation of the actions as mapping the states to a more general relational interpretation as mapping only situations. In such a setting the change of situations does not necessarily change the states, which allows the actions and the states to be considered independently. At the same time, we can continue using the useful metaphor for the transactions as directed graphs. **Fig. 2** shows a fragment of such a graph, centered around a single node that represents one particular situation along the transaction path, together with the associated nodes of the graph, related to it. Like the state-space graph, the intelligence graph represents the actions as edges of the graph, while the nodes represent the situations together with their descriptors.

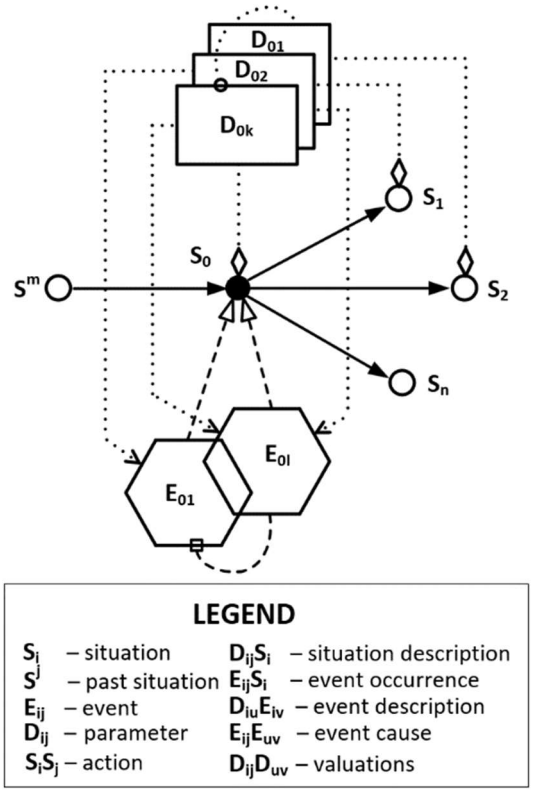

Figure 2: Inventory for Formulating Static and Dynamic Deterministic Rules

We have adopted the situational view at the actions a long time ago, but the practical value of this approach came out after we managed to formulate the transactional model in terms of Description Logic (DL) [14] and to represent it in serialized form as an ontology using standard languages of the semantic technologies multi-layered cake - RDF/RDFS and OWL [15]. For this purpose we introduced three separate vocabularies into the description logic with appropriate axiomatization of the theory:

- modeling the situations as DL concepts (interpreted semantically as a set and represented in the ontology as OWL classes)
- modeling the actions to be DL properties (interpreted semantically as relations between the situation sets and represented in the ontology as OWL properties)
- introducing events in the DL theory as another type of DL concept to model the asynchronous activities

This approach has substantial advantages over the state space-based approach. Firstly, by modeling the actions as relations between the situations we avoid some of the hurdles of the state-space approach since the description of the situations does not need to be exhaustive, which neutralizes the qualification problem. Secondly, utilizing the possibility of using conceptual types in DL it is possible to distinguish static and dynamic concepts explicitly and this way, to model both synchronous and asynchronous activities as independent. This conveniently supports event-driven control and real-time operation in transactional systems. Thirdly, by removing the explicit syntactic parametrization of the actions, which are no longer functions of their own parameters but conceptual relations between situations we can adopt the implicit semantic binding of their parameters to the parameters of the situations on the meta-level, which resolves the frame problem in an unexpected and elegant way [9]. Finally, the possibility to use taxonomies in DL theories and the corresponding serialized ontologies for both concepts and properties allows us to employ object orientation in modeling and implementation as well.

Since the semantic cake maps the interpretation of different serialized representations onto the same semantic domain, the interpretation of rules becomes semantically consistent with the interpretation of the ontology itself and can be easily integrated with it. The standard way of doing this is by using SWRL as a modeling language [15]. This allows to turn the models of transactional systems into *intelligence graphs*, which embed data, facts, conceptual and heuristic knowledge in a single repository, similarly to the knowledge graphs, which combine data, facts and conceptual knowledge.. This approach allows to incorporate much richer heuristics knowledge through the use of a separate domain ontology. We successfully used this approach to analyze the logical vulnerability of security policies in banking, formulated as SWRL rules on top of the ontology [8] and for combining domain-specific and problem-specific knowledge to generate a more informative presentation of geolocation information [3-5]. It has the potential to address the problem of explanation in AI through the use of a separate ontology of causality, where a simple black-box approach can bring to the surface a deeper explanation based on the causal relation between events, actions, and situations participating in the workflow of data processing operations [7]. This methodology is currently under development at GATE Institute within the framework of the EU Horizon DiverSea project.

3.3 Stochastic Model of the Transactions

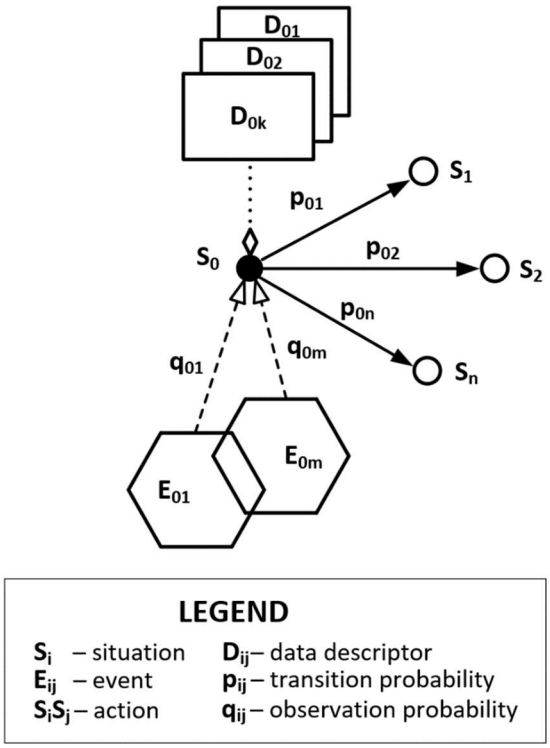

Figure 3: Inventory for Formulating Non-deterministic Rules

Up to the moment we have considered only deterministic templates, which allow us to rely on purely logical methods for modeling the heuristics as rules. However, in many practical applications, it is impossible to formulate the rules precisely due to the impossibility of assessing the environmental conditions or the subjective expertise with certainty and precision. Such is the case in many typically analytical tasks which require assessing the risks, associated with the normal functioning of transactional systems due to a variety of factors, such as unauthorized intrusion, malicious intervention, unexpected faults, imprecise measurements, or insufficient trustworthiness. In several of our projects in both the UK and Bulgaria which were focused on controlling financial transactions, we tried to overcome the limitation of the deterministic

models by associating the degree of probability directly in the heuristic rules, considering the transactions as Markov processes [5]. **Fig. 3** shows a fragment of a graph that models transactions under stress as a Partially Observable Markov Decision Process (POMDP) [16].

This model introduces non-determinism in two dimensions of the transactions – the transition between situations (labeled using p_{ij} as a probability of transitioning from situation S_i to situation S_j) and the observations of the events in each situation (labeled using q_{ij} as a probability of observing event E_j in situation S_i). The non-deterministic transition in this case is interpreted as caused by the conditions within the state, which can be both anticipated, but unknown (i.e., results of tests), or unexpected, but recoverable (like security threats, malicious interferences, device malfunctioning, etc.). The non-determinism of the observation is interpreted as either probable appearance of the event (in the case of unexpected), or probability for detection (in the case of anticipated).

Unfortunately, introducing some degree of expectation of the possible transitions, such as probability for choosing an alternative route for continuing the transactions, and assigning a degree of imprecision of the observations, such as probability for detection of the events along the transactions, introduces additional difficulties for using this transactional model. Apart from the computational complexity of the algorithms, used for evaluating global characteristics of the transactions which typically require approximate solutions of Bellman-type of equations, the additional difficulty comes from the dependence of these probabilities. While the first problem in some cases can be solved by reduction of the original POMDP problem to an MDP problem, which has precise solution using an efficient recurrent algorithm as we have shown in [5], the second problem remains completely outside the mathematical brilliance. At the same time, this creates a whole new world of opportunities for engineering pragmatism.

We have used this model successfully for quantitative assessment of the risks in online transactions under security threats and for analyzing the impact of the precision of intrusion detection for the purpose of designing security infrastructure with guaranteed low risks [2,6]. In our most recent project funded by UK Innovate UK project (CyDRA), we have adopted this model for the design of the system architecture of a commercial software product we are currently developing specifically for vulnerability analysis and security risk assessment of cyber systems.

4. Opportunities for Further Development

From the above considerations, it is clear that the problem for modeling, controlling, and analyzing the data processing in transactional systems is far from simple. But it also creates a lot of opportunities for further development of the technologies and creating the methodological basis for creating heterogeneous and distributed AI systems. While there is no chance to create a generic solution that is universally applicable to all problems requiring intelligence, it is also obvious that hybridization is the way forward. Based on such an understanding in this section we will discuss the opportunities for achieving this from both the conceptual and technical side.

4.1 Extending the Rule-based Systems with External Ontological Models

Knowledge is power but knowing the problem is only one of the conditions for finding solution of complex problems. The ontological models embody domain expertise which can be brough by the external experts to the software and data engineers to contextualize the solutions, make the AI models more informative and the algorithms for problem solving more efficient. This direction of research is known as ontological logic programming (OLP). It promises to go beyond the classical logic programming paradigm of problem solving, but to avoid some of the difficulties for adopting it in practice it should balance better between theoretical generality and practical applicability. Our own direction of research to achieve such a balance is the use of templates in the modeling language and the use of indexing mechanisms in the heuristic rule-based inference.

4.2 Adding Heuristics for setting the Default Probabilities in Stochastic Models

One of the practical difficulties to adopt more powerful stochastic models for solving real-life problems involving incomplete, imprecise and fuzzy knowledge is the setting of default parameters of the stochastic models. The mathematical models behind POMDP and MDP consider the prior probabilities as independent, while in reality they are subject of logical, temporal and causal dependencies between the situations, events and actions. By accounting these dependencies more realistic distribution of the probabilities can be set up. This is a subject of a separate analysis which can lead to adoption of logically consistent heuristics for parametrization of the model at the design stage. We are currently working on a set of templates for contextualizing the menu-driven interface of the transaction modeler, which makes use of such dependencies. This would allow implicit accounting of domain-independent and even domain-specific heuristics directly during the modeling process, which would increase the quality of the model without the need of complex modeling experience.

4.3 Utilization of Rule-based System Architectures for Explanation

The adoption of multi-layer architecture for data processing in which the data management, the data analysis and the explanation of the results appear on different levels requires complex explanation, which combines the logics of the separate levels. However, the explanation has its own logic going back to the philosophical studies of scientific explanation by Carnap, Quine, Hempel, and others, and although it is not realistic to expect full depth coverage of this phenomena, the ontology of causal dependencies between conditions, events, actions, and effects can be developed even without philosophical depth, purely from technical common sense. Such an ontology can be used for a black-box type of explanation generation which can address many concerns in contemporary literature, linked to some hard legal issues of adoption of AI and the need for developing Explainable AI.

5. Conclusion

So, the conclusion is simple: the rule-based approach from the early days of AI is not dead, and the heuristic rules are still kicking in multiple places; even more – the rule-based systems are increasingly more important to handle the real-life complexity of digital reality.

The projects considered in this paper have a wide scope: from purely security issues (fraud detection, unauthorized intrusion detection, vulnerability analysis, and risk assessment), to environment issues (air pollution monitoring, biodiversity analysis), and to their wider social impact (in healthcare, business management and legal practice). However, all of the projects have something in common - they involve complex transaction management which combines multiple paradigms for data processing. Such a system can be orchestrated by a centralized system, which requires system policies based on rules. The common denominators here are two: knowledge modeling, which can be a basis for formulating and applying both domain-specific and problem-specific policy rules, and the format of the rules themselves, which affects the algorithms for data processing.

From the perspective of contemporary technological advancements, the adoption of above principles requires adoption of a suitable platform for data processing, which utilizes virtualization, containerization, and orchestration of software services in a cloud environment. The recent shift of attention to AI also opens a wide horizon for automatic configuration of the software services, control of the operational pipelines, layered visualization, and causal explanation of the results on such a platform. This leads to truly heterogeneous and distributed systems.

We have successfully implemented elements of such a complex solution at the two research centers in London and in Sofia, proving its viability. Although we are committed to continue working after the above principles within our own centers, we are also very keen at collaborating with other research groups sharing similar views. Particular potential for this exists in two current projects (DiverSea and CyDRA), The large scope of DiverSea project, which covers most of the shores of Europe, together with the wide representation of partners creates an excellent opportunity for regional follow-up projects to apply and deepen the methodology for investigating the biodiversity on a regional scale. On the other hand, the potential application of CyDRA software currently under development creates an opportunity for direct use of the models, methods and algorithms in other application domains where the focus is on transactional information processing under various factors of risks. Particularly interesting is the possibility for a follow-up project in the domain of healthcare, where the assessment of the risks of developing certain diseases can substitute the second opinion of the medical professionals. Another interesting option here is the possibility to use the vulnerability analysis and risk assessment for designing production lines with guaranteed safety of operation, which is critical not only in manufacturing but also in food production.

Acknowledgments

The above considerations have been in the making for nearly a decade at two different institutions – the Cyber Security Research Centre of London Metropolitan University and GATE Institute of Sofia University. The projects referred to here have been funded from multiple sources in the UK (Lloyds Bank, Higher Education Investment Fund, and Innovate UK, specifically the CyDRA project under CyberASAP programme[1]), the EU (Horizon 2020 WIDESPREAD Programme, Horizon Europe 2023 Programme, specifically the DiverSea project under BIODIV Programme[2]), and Bulgaria (BG Government Operational Programme Science and Education for Smart Growth). The author is grateful to all these organizations for their continuing support and trust. However, all views expressed in the paper are of the author and in no way reflect the official positions of these organizations on the problems discussed here.

During these years I was lucky to have excellent students. Many of them later became my PhD students and project collaborators, like Pawel Gasiorowski, Viktor Sowinski-Mydlarz, Kanana Ezekiel, Karolina Bataytite, Artur Nascicionis, Khalid Mohamed, Sorin Radu, Hristo Hristov, Martin Hristev, Dion Mariyanayagam, Sabin Nakarmi and Reza Baghaeishiva, to mention some of them. In addition to the understanding and the support from colleagues and friends at Sofia University and London Metropolitan University, working with my students was a decisive factor in completion of the projects. Special thanks to my old friend Doncho Donchev from GATE Institute, who created the algorithm for reducing the original POMDP problem to a more tractable MDP problem.

A final word about the role of industrial partnership. The academic environment provides the working environment for conducting research and technological advancement, but it is too sterile and often prevents from facing the reality of real-world conditions. The author is grateful to many people from the industry, who helped to overcome this constraint, but first and foremost, I would like to mention Tony Phipps and Matt Lane, with whom I had many hours of interesting, inspiring, and productive discussions.

---

[1] Innovate UK, CyDRA: Risk and Vulnerability Assessment in Transactions under Security Threats, Cyber Security Academic Startup Accelerator Program, Project 10139273 (2024).

[2] EU Executive Agency, DiverSea: Integrated Observation, Monitoring and Prediction Architecture for Functional Biodiversity of Coastal Seas, HORIZON-CL6-2022-BIODIV-01 Programme, Project 101082004 (2022).

References

[1] V. Vassilev, V. Sowinski-Mydlarz, P. Gasiorowski, et al. (2023), "Building A Big Data Platform using Software without License Costs". In: L. Castro (ed.), *Open Source Horizons*, London: IntechOpen, 2024, pp. 29-52; DOI:10.5772/intechopen.1003268

[2] D. Dunsin, M. Ghanem, K. Ouazzane, and V. Vassilev (2023), "A Comprehensive Analysis of the Role of AI and Machine Learning in Modern, Digital Forensics and Incident Response", Forensic Science International, Volume 48, March 2024, Elsevier, pp. 1-22; DOI: 10.1016/j.fsidi.2023.301675

[3] V. Vassilev, B. Virdee, K. Ouazzane, et al. (2022), "Data Platform and Urban Data Services on Private Cloud". In: Senjyu, T., So-In, C., Joshi, A. (eds.), *Smart Trends in Computing and Communications*, Springer LNNS, Vol. 650, 2023, pp. 263-275; DOI: 10.1007/978-981-99-0838-7_23

[4] V. Vassilev, V. Sowinski-Mydlarz, D. Petrova-Antonova, et al. (2022), "Towards first urban data space in Bulgaria". In: Proc. 8th Int. Smart Cities Conference (ISC2), Paphos, Cyprus, IEEE, 2022, pp. 1-7. DOI: 10.1109/ISC255366.2022.9922237

[5] V. Vassilev, D. Donchev and D. Tonchev (2022), "Risk Assessment in Transactions under Threat as a Partially Observable Markov Decision Process". In: L. Amorosi et. Al. (eds.), *Optimization in Artificial Intelligence and Decision Sciences*, Springer, 2022, pp. 50-66; DOI: 10.1007/978-3-030-95380-5_18

[6] V. Vassilev, K. Ouazzane, V. Sowinski-Mydlarz, et al. (2021), "Network Security Analytics on the Cloud". In: Proc. 13th Int. Conf. on Cloud Computing, Data Science & Engineering (Confluence), IEEE, 2023, pp. 151-156; DOI:10.1109/Confluence56041.2023.10048889.

[7] V. Vassilev, S. Ilieva, V. Sowinski-Mydlarz, et al. (2021), "AI-based Hybrid Data Platforms", In: Curry, E., Scerri, S. and Tuikka, T. (eds.), *Data Spaces: Design, Deployment, and Future Directions*, Springer, 2022, pp. 147-170; DOI: 10.1007/978-3-030-98636-0_8

[8] V. Vassilev, V. Sowinski-Mydlarz, P. Gasiorowski, et al. (2019), "Intelligence Graphs for Threat Intelligence and Security Policy Validation of Cyber Systems". In: P. Bansal, et al. (eds.), *Advances in Intelligent Systems and Computing*, Vol. 1164, Springer, 2020, pp. 125-139; DOI: 10.1007/978-981-15-4992-2_13

[9] K. Bataityte, V. Vassilev, and O. Gill (2019), "Ontological Foundations of Modelling Security Policies for Logical Analysis". In: I. Maglogiannis, et al. (eds.), *IFIP Advances in Information and Communication Technology*, Vol. 583, Springer, 2020, pp. 368-380; DOI: 10.1007/978-3-030-49161-1_31

[10] K. Ezekiel, V. Vassilev, K. Ouazzane, et al. (2017), "Adaptive business rules framework for workflow management". Business Process Management Journal, 2018: 25 (5), pp. 948-971; DOI: 10.1108/BPMJ-08-2017-0219

[11] J. McCarthy and P. Hayes, "Some philosophical problems from the standpoint of artificial intelligence", In: B. Meltzer and D. Michie (eds.), *Machine Intelligence 4*, Edinburgh University Press (1969), pp. 463–502; ISBN: 9780852240625

[12] J. Hintikka (1962), *Knowledge and Belief: An Introduction to the Logic of the Two Notions.* King's College Publications (2005); ISBN: 9781904987086

[13] J. Barwise and J. Perry (1983), *Situations and Attitudes.* MIT Press (1983); ISBN: 9780262021890

[14] F. Baader, I. Horrocks, C. Lutz, and U. Sattler (2017), *An Introduction to Description Logic.* Cambridge University Press (2017); ISBN: 9781139025355

[15] P. Hitzler, M. Krötzsch, and S. Rudolph (2009), *Foundations of Semantic Web Technologies.* Chapman & Hall/CRC (2009); ISBN: 9781420090505

[16] Mausam and A. Kolobov (2012), *Planning with Markov Decision Processes: An AI Perspective.* Morgan Claypool Publishers (2012); ISBN: 9781608458868