Final Report on the PRIN 2020 Project RIPER

Aniello Murano
Università degli Studi di Napoli Federico II, Napoli, Italy
nello.murano@gmail.com

September 30, 2025

Abstract

This document reports the scientific activities and research outputs of the University of Naples Federico II (UNINA) research unit within the PRIN 2020 project *Resilient AI-Based Self-Programming and Strategic Reasoning* (RIPER; project code 20203FFYLK), carried out between June 1, 2022 and May 31, 2025. The UNINA unit, led by Prof. Aniello Murano, contributed to the logical and formal foundations of resilience in multi-agent artificial intelligence systems. The central goal was to study how autonomous agents can reason, coordinate, and synthesize strategies under uncertainty, partial observability, probabilistic dynamics, and quantitative or fuzzy evaluations of system properties. The resulting research program combined concurrent game structures, strategic logics, model checking, strategy synthesis, and knowledge representation. It produced theoretical results, prototype tools, and applications in domains ranging from social choice and mechanism synthesis to stochastic multi-agent systems, security, decision support, and medical AI.

Keywords: resilient artificial intelligence; strategic reasoning; multi-agent systems; concurrent game structures; ATL; Strategy Logic; model checking; strategy synthesis; uncertainty.

1 Introduction

Resilience is a central requirement for artificial intelligence systems that operate in open, uncertain, and dynamically changing environments. A resilient AI system should not merely achieve a target objective in an idealized model; rather, it should continue to operate, adapt, and reorganize its behavior when observations are partial, actions are unreliable, other agents act strategically, or the environment deviates from nominal assumptions. This article reformulates the final scientific report of the UNINA unit of the PRIN 2020 project RIPER as a scientific contribution centered on formal methods for resilient AI.

The overall objective of RIPER was to develop methodologies for the design of resilient AI systems capable of self-programming and autonomously adapting their behavior in response to environmental uncertainty and change. The project integrated planning, strategy synthesis, learning, and formal reasoning. At the national level, the project was coordinated by Prof. Giuseppe De Giacomo and involved, in addition to UNINA, units affiliated with the Universities of Brescia and Trento, coordinated respectively by Prof. Alfonso Gerevini and Prof. Marco Roveri. Within this framework, the UNINA unit focused on the logical and game-theoretic foundations of resilience, with particular emphasis on strategic reasoning, knowledge representation, and the verification of strategic abilities under uncertainty.

The scientific perspective adopted by the unit builds on a broad line of work on strategic reasoning in multi-agent systems, including the formal analysis of strategic abilities, private and public information, and human-friendly or natural strategies [22, 26, 42, 45]. It also connects to work on robust and quantitative strategic logics, fuzzy strategic behavior, and discounting in Strategy Logic [25, 28, 34, 35].

2 Scientific Objectives of the UNINA Unit

The research activity of the UNINA unit was organized along three main scientific axes.

First, the unit studied formal models of uncertain multi-agent systems based on concurrent game structures (CGSs). In these models, uncertainty arises from imperfect information, probabilistic behavior, and weighted or fuzzy interpretations of atomic propositions. Such extensions are necessary to capture agents that only observe partial views of the global system, actions whose effects are uncertain, communication channels with variable reliability, and properties that are satisfied to different degrees rather than in a purely Boolean manner. This modeling direction is connected to research on partially known environments, imperfect-information games, stochastic settings, and parametric Markov-chain families [10, 15, 21, 23, 40].

Second, the unit developed and analyzed strategic logics for resilience. The main focus was on extensions and variants of Alternating-time Temporal Logic (ATL) and Strategy Logic. These formalisms were enriched with operators and semantic mechanisms capable of expressing quantitative aspects, robustness, fault tolerance, and adaptive recovery. Relevant results include robust ATL, quantitative ATL, obstruction ATL, intuitionistic ATL, first-order coalition logic, and variants of temporal and computation-tree logics [4, 8, 13, 17, 25, 35, 39].

Third, the unit investigated algorithmic verification and synthesis problems for these models and logics. The emphasis was on automatic synthesis of resilient strategies, strategic model checking under uncertainty, and the computational complexity trade-offs between expressiveness and decidability. This work is linked to scalable verification of Strategy Logic, verification of stochastic multi-agent systems, model checking of knowledge and time, module checking for quantitative games, and the repair or verification of general game descriptions [9, 14, 16, 18, 33].

3 Formal Models of Multi-Agent Systems under Uncertainty

A first research line extended classical concurrent game structures to scenarios in which uncertainty is intrinsic. In standard CGSs, several agents choose actions concurrently, and the joint action determines the next state. For resilient AI, this abstraction must be refined so that agents can reason under uncertainty about the actual state of the system, the behavior of other agents, and the effect of their own choices.

3.1 Imperfect Information and Knowledge

In imperfect-information settings, agents do not have access to the true global state. Their strategies can depend only on the information they observe and on the knowledge accumulated over time. This constraint is essential for modeling real AI systems, where sensors, communication, and distributed control architectures often provide only partial views of the environment. The UNINA unit studied this setting in relation to knowledge-based temporal reasoning, public and private information, and the verification of games with imperfect information using Strategy Logic [14, 23, 45].

The resilient behavior of an agent is therefore tied to what the agent knows, what it does not know, and how it can update its strategic choices as knowledge changes. This interaction between strategic reasoning and knowledge representation was a central theme of the unit's contribution.

3.2 Probabilistic Dynamics

A second source of uncertainty is probability. Probabilistic transitions model non-deterministic behavior of the environment, unreliable actions, noisy observations, and uncertain communication. Within RIPER, probabilistic models were studied as a foundation for reasoning about whether a coalition can guarantee that a goal is reached with probability at least p, even in the presence of failures or adversarial behavior.

This direction includes natural strategic ability in stochastic multi-agent systems, robust strategies for stochastic systems, forgetful strategies, entropy-based accounts of strategic uncertainty, and strategic abilities of forgetful agents [10, 12, 15, 16, 36]. These results clarify how resilience can be formalized when agents must operate with limited memory, stochastic outcomes, and probabilistic guarantees.

3.3 Weighted, Quantitative, and Fuzzy Models

A third modeling dimension concerns atomic propositions that are not simply true or false. In many AI systems, properties come with degrees: a service can have a quality level, a safety requirement can be partially satisfied, and an adverse event can have different levels of severity. Weighted and fuzzy extensions of strategic models allow such graded evaluations to be represented formally.

The UNINA unit investigated these issues in relation to quantitative ATL, fuzzy strategic behaviors, robust automata and temporal logics, and discounting in Strategy Logic [25, 28, 34, 41]. This modeling layer supports refined notions of resilience, where the objective is not only to remain functional, but to maintain an acceptable or optimal level of functionality under perturbations.

4 Strategic Logics for Resilience

The second major research line concerned strategic logics for specifying complex properties of resilient multi-agent systems. Strategic logics make it possible to express statements such as: a coalition of agents has a strategy to guarantee a goal; a system can degrade gracefully while maintaining a minimal service level; or a set of agents can reorganize cooperation after a fault.

4.1 ATL and Its Extensions

ATL and its variants provide a natural language for reasoning about the abilities of coalitions. Within the resilience setting, ATL-style modalities can be enriched with probabilistic thresholds, quantitative objectives, imperfect information, and robustness constraints. The unit contributed to the study of robust ATL, ATL for dynamic gaming environments, obstruction ATL, quantitative ATL, and intuitionistic ATL [8, 17, 20, 25, 35].

These logics support the formalization of fault tolerance and graceful degradation. For example, one can express that a system should preserve a minimum level of functionality after a disturbance, or that a coalition should have a strategy that succeeds despite the obstructive actions of an environment or adversary.

4.2 Strategy Logic and Explicit Reasoning about Strategies

Strategy Logic makes strategies first-class objects of reasoning. This greater expressiveness is useful for modeling adaptation, strategy revision, and strategic dependencies between agents. Within RIPER, Strategy Logic and related formalisms were used to reason about strategy synthesis, game descriptions, discounting, and mechanism design [9, 23, 33, 34].

Explicit strategy quantification also supports the analysis of strategic constraints that are difficult to express in ATL alone, such as the comparison of alternative strategies, the synthesis of mechanisms that induce desired behaviors, and the modeling of rational agents under incentives [1, 24, 29, 43, 44].

4.3 Knowledge, Cooperation, and Robustness

The resilient response of a multi-agent system depends not only on available actions, but also on how agents coordinate and what they know about one another. The unit therefore analyzed the interaction between strategic reasoning and knowledge representation. This interaction is visible in work on model checking knowledge and time, public and private affairs in strategic reasoning, and human-friendly or natural strategies [14, 26, 42, 45].

The logical tools developed in this context make it possible to connect resilience with cooperation. A strategy can be resilient because it replans, because it redistributes responsibilities among agents, or because it remains valid under weaker knowledge assumptions.

5 Verification, Synthesis, and Prototype Tools

On the algorithmic side, the UNINA unit contributed to model checking and strategy synthesis for the models and logics described above. The goal was to determine whether a given system satisfies a strategic specification and, whenever possible, to automatically synthesize strategies that guarantee the desired property.

5.1 Strategic Model Checking

Strategic model checking under uncertainty requires algorithms that combine temporal reasoning, coalition reasoning, imperfect information, and sometimes probability. The unit analyzed conditions under which verification remains decidable and identified the computational costs associated with the increased expressiveness of the logics. Contributions in this line include scalable verification of Strategy Logic through three-valued abstraction, verification of stochastic systems with forgetful strategies, and complexity results for model checking knowledge and time [14, 16, 33].

The research also addressed temporal foundations and paradoxical temporal structures, including work on Yablo's sequences in finite-trace and potentially infinite-trace settings [19, 27]. These studies enrich the logical background needed to reason about long-term, finite-horizon, and potentially unbounded system behaviors.

5.2 Strategy Synthesis

Strategy synthesis aims to construct a strategy that satisfies a formal specification. In the context of resilience, the synthesized strategy should remain effective under uncertainty, failures, and environmental perturbations. The UNINA unit investigated synthesis for strategic logics and mechanisms, including automated synthesis of mechanisms, synthesis with Strategy Logic, and formal verification and synthesis of mechanisms for social choice [1, 43, 44].

These results provide a formal counterpart to self-programming: the system is not only verified against a property, but can be equipped with automatically generated strategic behavior that realizes the property under the assumptions of the model.

5.3 Prototype Tools

Several techniques developed within the project were implemented as prototype tools. These prototypes provide a basis for future integration into design pipelines for resilient AI systems. Examples include S4H for synthesizing human-like strategies, FindMe as a videogame AI based on CTL with an optimized synthesis algorithm, and HYASM for verifying hierarchical systems [5, 11, 38].

Prototype implementation also supported the transfer of formal methods toward executable control policies, dynamic gaming environments, and strategy-based decision procedures [20, 26].

In this way, the formal notion of strategy becomes operational and can be linked to planning and control components developed by the other RIPER units.

6 Applications and Interdisciplinary Impact

Although the UNINA unit was primarily concerned with formal foundations, several research outputs addressed applied or interdisciplinary scenarios. In security, attack graphs were studied from formal and game-theoretic perspectives, supporting the analysis of adversarial behavior and defensive strategies [2, 31]. In privacy and stochastic modeling, parameter synthesis for families of Markov chains was applied to multi-agent systems privacy [21].

In decision support and data-driven AI, the project produced work on adaptive flood prediction in digital twin environments, multi-agent sentiment analysis, and parking allocation problems [3, 30, 32, 37]. These applications show how game-theoretic and strategic perspectives can contribute to robust decision-making in dynamic environments.

The project also produced medical AI contributions in which the logical-strategic component played a decision-making role within hybrid or neuro-symbolic methods. In particular, adaptive neuro-fuzzy clustering and convolutional neural-network approaches were developed for leukemia prediction [6, 7]. These works illustrate the potential of combining formal reasoning and learning-based components in settings where resilience, interpretability, and reliability are important.

7 Training, Recruitment, and Community Building

The RIPER project had a significant impact on advanced training at UNINA. PRIN funds supported research fellowships and contracts for young scholars working on topics directly connected to *Resilient AI-Based Self-Programming and Strategic Reasoning*. These activities focused especially on formal aspects of automated strategic reasoning and on the development of models and logics for resilience in multi-agent systems.

The project contributed to the creation of a cohesive and interdisciplinary research group on logic, multi-agent systems, and resilient artificial intelligence. It supported the scientific growth of PhD students, postdoctoral researchers, and young researchers, and promoted their integration into national and international research communities in artificial intelligence, logic in computer science, and multi-agent systems.

The results were disseminated through journal articles, international conference papers, workshops, and seminars. This dissemination increased the visibility of RIPER and of the UNINA unit within the scientific community.

8 Conclusions and Future Directions

The UNINA unit of PRIN RIPER provided relevant contributions to the formal understanding of resilience in artificial intelligence systems. It proposed and studied uncertain multi-agent models based on concurrent game structures with imperfect information, probability, and weighted or fuzzy valuations. It developed and extended strategic logics, including ATL, Strategy Logic, and related variants, to capture sophisticated properties of robustness, adaptivity, and fault tolerance. It also analyzed verification and synthesis problems for resilient strategies and laid the groundwork for automatic tools supporting the design of reliable AI systems.

These results are a key component of the integrated RIPER vision: combining formal methods, planning, learning, and robotics in order to build AI systems that can achieve complex goals while continuing to function and adapt under uncertainty, partial information, and environmental perturbations.

The research lines initiated in the project will continue beyond its formal conclusion. Future work will address strategic logics that are increasingly expressive while remaining algorithmically tractable, tighter integration between formal reasoning and learning methods, and applications to real-world cooperative multi-agent systems and critical infrastructures, where resilience is not an optional property but a fundamental requirement.

Acknowledgements

This docuument reports activities carried out within the PRIN 2020 project *Resilient AI-Based Self-Programming and Strategic Reasoning* (RIPER; project code 20203FFYLK) at the University of Naples Federico II research unit.

References

[1] Munyque Mittelmann, Bastien Maubert, Aniello Murano, Laurent Perrussel: Formal verification and synthesis of mechanisms for social choice. Artif. Intell. 339: 104272 (2025)

[2] Davide Catta, Jean Leneutre, Vadim Malvone, Aniello Murano: A formal approach to attack graphs. Ann. Math. Artif. Intell. 93(4): 589-610 (2025)

[3] Youcef Djenouri, Michal Tomasz Godziszewski, Fabio A. A. Andrade, Gautam Srivastava, Ahmed Nabil Belbachir, Aniello Murano: Game-Theoretic Consensus Deep Learning for Adaptive Flood Prediction in Digital Twin Environments. IEEE J. Sel. Top. Appl. Earth Obs. Remote. Sens. 18: 21355-21366 (2025)

[4] Andrea Capone, Laura Bozzelli, Davide Catta, Vadim Malvone, Aniello Murano: An Intuitionistic Version of Computation Tree Logic, EUMAS 2025

[5] Marco Aruta, Vadim Malvone, Aniello Murano: S4H: A Tool for Synthesizing Human-Like Strategies, EUMAS 2025

[6] Marco Aruta, Ciro Listone, Giuseppe Murano, Aniello Murano: ADNF-Clustering: An Adaptive and Dynamic Neuro-Fuzzy Clustering for Leukemia Prediction, Healthcom 2025

[7] Marco Aruta, Ciro Listone, Giuseppe Murano, Aniello Murano: L-CNN: A Convolutional Neural Network for Leukemia Prediction, KES IN MED 25

[8] Andrea Capone, Laura Bozzelli, Davide Catta and Aniello Murano: An Intuitionistic Version of Alternating-Time Temporal Logic, KR 2025

[9] Yifan He, Munyque Mittelmann, Aniello Murano, Abdallah Saffidine and Michael Thielscher: Repairing General Game Descriptions, KR 2025

[10] Raphaël Berthon, Joost-Pieter Katoen, Munyque Mittelmann, Aniello Murano: Robust Strategies for Stochastic Multi-Agent Systems. AAMAS 2025: 2437-2439

[11] Marco Aruta, Vadim Malvone, Aniello Murano, Vincenzo Pio Palma, Salvatore Romano: FindMe: A Prototype Videogame AI based on CTL with an Optimized Synthesis Algorithm. AAMAS 2025: 2997-2999

[12] Wojciech Jamroga, Michal Tomasz Godziszewski, Aniello Murano: Strategies, Credences, and Shannon Entropy: Reasoning about Strategic Uncertainty in Stochastic Environments. IJCAI 2025: 126-134

[13] Davide Catta, Rustam Galimullin, Aniello Murano: First-Order Coalition Logic. IJCAI 2025: 4410-4418

[14] Laura Bozzelli, Bastien Maubert, Aniello Murano: On the Complexity of Model Checking Knowledge and Time. ACM Trans. Comput. Log. 25(1): 8:1-8:42 (2024)

[15] Raphaël Berthon, Joost-Pieter Katoen, Munyque Mittelmann, Aniello Murano: Natural Strategic Ability in Stochastic Multi-Agent Systems. AAAI 2024: 17308-17316

[16] Francesco Belardinelli, Wojtek Jamroga, Munyque Mittelmann, Aniello Murano: Verification of Stochastic Multi-Agent Systems with Forgetful Strategies. AAMAS 2024: 160-169

[17] Davide Catta, Jean Leneutre, Vadim Malvone, Aniello Murano: Obstruction Alternating-time Temporal Logic: A Strategic Logic to Reason about Dynamic Models. AAMAS 2024: 271-280

[18] Wojciech Jamroga, Munyque Mittelmann, Aniello Murano, Giuseppe Perelli: Playing Quantitative Games Against an Authority: On the Module Checking Problem. AAMAS 2024: 926-934

[19] Michal Tomasz Godziszewski, Davide Catta, Aniello Murano: Temporal Truth in the Limit: Yablo's Paradox in LTLf and over Potentially Infinite Traces. EUMAS 2024: 110-128

[20] Marco Aruta, Aniello Murano, Salvatore Romano: ATL for Dynamic Gaming Environments. EUMAS 2024: 129-137

[21] Francesco Spegni, Luca Spalazzi, Roberto Rosetti, Aniello Murano: Parameter Synthesis for Families of Markov Chains with an Application to Multi-agent Systems Privacy. EUMAS 2024: 159-178

[22] Aniello Murano: Formal Aspects of Strategic Reasoning. KoDis+CAKR+SYNERGY@KR 2024

[23] Yifan He, Munyque Mittelmann, Aniello Murano, Abdallah Saffidine, Michael Thielscher: Verification of General Games with Imperfect Information Using Strategy Logic. KR 2024

[24] David Hyland, Munyque Mittelmann, Aniello Murano, Giuseppe Perelli, Michael J. Wooldridge: Incentive Design for Rational Agents. KR 2024

[25] Angelo Ferrando, Giulia Luongo, Vadim Malvone, Aniello Murano: Theory and Practice of Quantitative ATL. PRIMA 2024: 231-247

[26] Marco Aruta, Vadim Malvone, Aniello Murano: Development of Natural Strategies in Strategic Logics (short paper). AI4CC-IPS-RCRA-SPIRIT@AI*IA 2024

[27] Michal Tomasz Godziszewski, Davide Catta, Aniello Murano: Temporal (Non-)Paradox: Yablo's Sequences in LTL over Finite Traces (short paper). AI4CC-IPS-RCRA-SPIRIT@AI*IA 2024

[28] Patricia Bouyer, Orna Kupferman, Nicolas Markey, Bastien Maubert, Aniello Murano, Giuseppe Perelli: Reasoning about Quality and Fuzziness of Strategic Behaviors. ACM Trans. Comput. Log. 24(3): 21:1-21:38 (2023)

[29] Munyque Mittelmann, Bastien Maubert, Aniello Murano, Laurent Perrussel: Formal Verification of Bayesian Mechanisms. AAAI 2023: 11621-11629

[30] Aniello Murano, Mimmo Parente, Silvia Stranieri: A Multi-Agent Game for Sentiment Analysis. IPS-RCRA-SPIRIT@AI*IA 2023

[31] Davide Catta, Antonio Di Stasio, Jean Leneutre, Vadim Malvone, Aniello Murano: A Game Theoretic Approach to Attack Graphs. ICAART (1) 2023: 347-354

[32] Aniello Murano, Silvia Stranieri, Munyque Mittelmann: Multi-Agent Parking Problem with Sequential Allocation. ICAART (3) 2023: 484-492

[33] Francesco Belardinelli, Angelo Ferrando, Wojciech Jamroga, Vadim Malvone, Aniello Murano: Scalable Verification of Strategy Logic through Three-Valued Abstraction. IJCAI 2023: 46-54

[34] Munyque Mittelmann, Aniello Murano, Laurent Perrussel: Discounting in Strategy Logic. IJCAI 2023: 225-233

[35] Aniello Murano, Daniel Neider, Martin Zimmermann: Robust Alternating-Time Temporal Logic. JELIA 2023: 796-813

[36] Francesco Belardinelli, Wojciech Jamroga, Munyque Mittelmann, Aniello Murano: Strategic Abilities of Forgetful Agents in Stochastic Environments. KR 2023: 726-731

[37] Francesco Noviello, Munyque Mittelmann, Aniello Murano, Silvia Stranieri: Parking Problem with Multiple Gates. PAAMS 2023: 213-224

[38] Angelo Ferrando, Vadim Malvone, Aniello Murano, Silvia Stranieri: HYASM: A Tool to Verify Hierarchical Systems. WETICE 2023: 1-6

[39] Davide Catta, Vadim Malvone, Aniello Murano: Reasoning about Intuitionistic Computation Tree Logic. AREA@ECAI 2023: 42-48

[40] Benjamin Aminof, Aniello Murano, Sasha Rubin, Florian Zuleger: Verification of agent navigation in partially-known environments. Artif. Intell. 308: 103724 (2022)

[41] Laura Bozzelli, Aniello Murano, Adriano Peron: Context-free timed formalisms: Robust automata and linear temporal logics. Inf. Comput. 283: 104673 (2022)

[42] Francesco Belardinelli, Wojtek Jamroga, Vadim Malvone, Munyque Mittelmann, Aniello Murano, Laurent Perrussel: Reasoning about Human-Friendly Strategies in Repeated Keyword Auctions. AAMAS 2022: 62-71

[43] Munyque Mittelmann, Bastien Maubert, Aniello Murano, Laurent Perrussel: Synthesis of Mechanisms with Strategy Logic. ICTCS 2022: 47-52

[44] Munyque Mittelmann, Bastien Maubert, Aniello Murano, Laurent Perrussel: Automated Synthesis of Mechanisms. IJCAI 2022: 426-432

[45] Nathanaël Fijalkow, Bastien Maubert, Aniello Murano, Sasha Rubin, Moshe Y. Vardi: Public and Private Affairs in Strategic Reasoning. KR 2022