

raison: Developing Reliable Decision-Making Agents

Demonstration Track

Pavlos Moraitis
Argument Theory
Paris, France
Université Paris Cité
Paris, France
pavlos.moraitis@u-paris.fr

Nikolaos Spanoudakis
Argument Theory
Paris, France
Hellenic Mediterranean University
Chania, Greece
nspan@hmu.gr

Antonis Kakas
Argument Theory
Paris, France
University of Cyprus
Nicosia, Cyprus
antonis@ucy.ac.cy

ABSTRACT

This paper presents the *raison* platform, a high-level technological environment for the development of automated, reliable and explainable decision-making agents. The research underlying the platform and its technological progress has now reached a mature stage that allows the platform to be used for the development of complex real-life applications without writing a single line of code.

KEYWORDS

Tools and platforms supporting agent development; Computational argumentation; AlaaS; Automated decision-making; xAI

ACM Reference Format:

Pavlos Moraitis, Nikolaos Spanoudakis, and Antonis Kakas. 2026. *raison*: Developing Reliable Decision-Making Agents: Demonstration Track. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 3 pages. <https://doi.org/10.65109/CYPN9399>

1 THE RAISON PLATFORM

The recent developments in AI have given rise to a renewed interest in building autonomous agents. In this effort, one of the main challenges is related to the reliability of the decisions and actions of the agents. The platform of *raison* addresses this challenge by providing an environment for supporting the development of reliable and explainable, decision making agents. It has recently been developed by the “Argument Theory” start-up company.

The wider aim of *raison* is to enable individuals to innovate in their own field by providing them with the capability of developing decision-making systems directly from their high-level understanding of their expertise without writing a single line of code. The developers and users of the *raison* platform need not be aware of the underlying technology of the platform, much in the same way as in using Large Language Models (LLMs). In comparison to LLMs, in *raison* we concentrate on building systems under the human’s expertise of the application problem offering *absolute reliability* with respect to the reasoning capability of the developed systems. This reliability is guaranteed by the use of *Gorgias* [3, 4, 13], a generic inference engine that is endowed with both *argumentative* and *abductive* reasoning capabilities. This system implements the

structured preference-based argumentation framework called **Logic Programming with Priorities** [3]. The core technology of *raison* consists of two components:

- A Front-end *authoring tool* for modeling complex decision-making problems through *natural language dialogue*, directly in the application domain. Developers can construct decision policies—i.e., decision and preference arguments expressing the relative strength of one argument over another—simply by answering the platform’s questions in natural language, without writing any code. This dialogue-driven approach automates *knowledge acquisition* and *preference elicitation*, enabling the scalable development of symbolic AI systems.
- A Back-end with explainable argumentation as a decision-making engine, using the Gorgias system, offered as a *Service over the Internet (AlaaS)*. The explanations are derived from the knowledge supplied by the developer throughout the knowledge acquisition and preference elicitation process.

These components are complemented by services for *automated code* and *API generation*, as well as *deployment*. The *raison* platform complemented with these services (tools) supports the development of automated decision-making agents by assisting developers in the specification, implementation, deployment and testing of their systems. Moreover, being offered as a service, this decision making capability can be interfaced with other “standard” agent platforms, e.g. the Java Agent Development Framework (JADE) [2], or other technologies, e.g. Internet of Things (IoT) [17], to build agentic application systems.

The authoring tool of the platform is based on the **Software Development via Argumentation (SoDA)** methodology for knowledge engineering and acquisition [4, 15, 16]. In this methodology, knowledge, e.g., a Decision Policy, is captured in terms of **hierarchies of Scenario Based Preferences (SBP)**, namely statements expressing which options are preferred under an application scenario and how such preferences change in successively more specific scenarios [4]. An *option* can be a decision, an action, a belief or a goal. Consider, for example, the following part of a Decision Policy for a Salary Negotiation Agent: “*If the offered salary is above the expected salary, accept the offer. Refuse an offer when the offered salary is low. Also refuse an offer when the offered salary is close to the expected salary, but accept it if the yearly salary increase plan, brings within two years the salary well above the expected salary.*”

Decision arguments in *raison* have the form *Scenario* \triangleright *Option*. In the above decision policy, *Options* = {*accept, refuse*} are the two possible decisions while *Scenarios*={the offered salary is above the expected salary, the offered salary is low, the offered salary is



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/CYPN9399>

close to the expected salary, the yearly salary increase plan brings within two years the salary well above the expected salary} are the possible scenarios supporting the options (i.e. the decisions) and the preferences.

By responding to the natural-language questions posed by *RAISON* through an advanced authoring tool, developers can define options and their associated scenarios to construct decision arguments. By applying a *conflict resolution process*, they can then establish hierarchies of preference arguments—expressing unconditional or conditional (contextual) preferences—to resolve conflicts between competing decision or preference arguments at different levels (i.e. meta-preferences). Once this process is completed, developers can verify the resulting decision policies and trigger (with a *run* button) their automatic translation into Gorgias language by *RAISON* (see Listings 1 and 2). The generated code is immediately automatically deployed on the platform, making the developed system ready for testing and use—either through a run screen in *RAISON* or via the API. In other words, the authoring tool of the *RAISON* platform offers a “no-code programming” environment for the software development of the decision-making modules of agents. Developers can work in two modes: *basic* and *advanced*. In basic mode, they use only natural language to represent options and argument (decision and preference) scenarios. This representation is automatically translated by *RAISON* into propositional Gorgias code (see Listing 1).

Listing 1: Basic mode Gorgias code.

```
rule (r1 , accept , []) :-
    the_salary_offered_is_above_the_expected_salary .
rule (r2 , refuse , []) :- the_offered_salary_is_low .
rule (r3 , refuse , []) :-
    the_offered_salary_is_close_to_the_expected_salary .
rule (r4 , accept , []) :-
    the_offered_salary_is_close_to_the_expected_salary .
rule (p1 , prefer (r3 , r4) , []) .
rule (p2 , prefer (r4 , r3) , []) :-
    the_yearly_salary_increase_plan_brings_within_two_years_the_salary_well_above_the_expected_salary .
rule (c1 , prefer (p2 , p1) , []) .
complement (accept , refuse) .
complement (refuse , accept) .
```

This level of representation is often sufficient for domains based on qualitative knowledge, such as legal, medical or regulatory compliance contexts. In advanced mode, developers can combine natural language with Prolog, enabling them to refine (or personalize) the automated representation using first-order logic (FOL)—introducing variables, predicates, functions, or mathematical expressions. Based on this new formal representation, *RAISON* automatically generates a new code (see Listing 2, which personalizes Listing 1). This is particularly useful in domains where formal precision is required, such as finance, trading, risk evaluation, or negotiations (i.e. our example). This dual-mode approach supports collaboration between users with different roles and technical expertise. For example, a human resources manager may define recruitment decision policies in natural language (Listing 1 for our example), while a technical expert formalizes them using first-order logic (Listing 2 for our example). The FOL representation is done through a dedicated GUI. **Both automatically generated codes are completely transparent to the developers (basic or advanced).** Interested readers can view videos showcasing the development of various use-cases and tutorials available here: www.youtube.com/@Argument-Theory.

Decision-making is then carried out via argumentation-based reasoning of building arguments for options that are acceptable, i.e., that are stronger or at least as strong, with respect to preference arguments expressing preferences for arguments supporting contrary options (see rules r1 and r2) or contrary preferences (see rules p1 (unconditional preference) and p2 (conditional/contextual preference)) at different hierarchical levels. Rule c1 expresses unconditional meta-preference.

By examining the corresponding rules (i.e., those with the same names) in Listings 1 and 2, one can observe how expressions such as “the offered salary is above the expected salary”, “the offered salary is low”, “the offered salary is close to the expected salary” or “the yearly salary increase plan brings within two years the salary well above the expected salary” are formally represented. Alternative formal encodings, however, are also possible. *Complements* represent the conflict between competing options. *Brackets* in the rules allow to declare *beliefs* (defeasible knowledge) or *abducibles* (hypotheses) when applicable.

Listing 2: Advanced mode Gorgias code.

```
rule (r1 , accept , []) :- offered_salary (O) , expected_salary (E) , O>=E .
rule (r2 , refuse , []) :- offered_salary (O) , expected_salary (E) , O=<0.7*E .
rule (r3 , refuse , []) :- offered_salary (O) , expected_salary (E) , E>O ,
    O>0.7*E .
rule (r4 , accept , []) :- offered_salary (O) , expected_salary (E) , E>O ,
    O>0.7*E .
rule (p1 , prefer (r3 , r4) , []) .
rule (p2 , prefer (r4 , r3) , []) :- offered_salary (O) , expected_salary (E) ,
    yearly_salary_increase (X) , O*((1+X)**2) >1.5*E .
rule (c1 , prefer (p2 , p1) , []) .
complement (accept , refuse) .
complement (refuse , accept) .
```

One main advantage of using an argumentation-based computation is the fact that this naturally gives informative **explainable solutions** [8, 11, 14]. Gorgias can explain its reasoning process by tracing the decision and preference arguments it employed to reach a decision, as well as the data on which the decision was based.

To bring the *RAISON* decision-making services into the practice of building real world agents, the platform exposes an **API** with two major services: a) **Get application metadata**: this service is of type GET and returns the ids of the elements that make up the possible scenarios and the ids of the options b) **Query an application**: this service is of type POST and allows the agent to query the developed application for allowed options within a given context composed of scenario elements. These services allow the integration of a decision policy to an agent decision-making module.

The *Gorgias* technology has been applied in a wide variety of real-life applications (see e.g. [18], [1], [12], [5], [7], [10]). Some recent applications are built using the *RAISON* platform ([6], [9]). These applications have been developed using the *RAISON* under a closed, regulated access. Since last summer, the *RAISON* platform is also open for beta testing. The *RAISON* platform currently supports decision problems involving several dozen options, several hundred decision rules, and several thousand lines of generated Gorgias code overall (including preference and meta-preference rules). Recently, we initiated research work on a *neuro-symbolic approach* aimed at enhancing the platform’s efficiency and usability, when human expertise is not available, through the introduction of novel features. *RAISON* will open to a global paying audience through Amazon Marketplace in the coming months.

REFERENCES

- [1] Sofia Almpani, Yiannis Kiouvrekis, Petros Stefaneas, and Panayiotis Frangos. 2022. Computational Argumentation for Medical Device Regulatory Classification. *International Journal on Artificial Intelligence Tools* 31, 01 (2022), 2250005. <https://doi.org/10.1142/S0218213022500051>
- [2] Federico Bergenti, Giovanni Caire, Stefania Monica, and Agostino Poggi. 2020. The first twenty years of agent-based software development with JADE. *Autonomous Agents and Multi-Agent Systems* 34, 2 (2020), 36. <https://doi.org/10.1007/s10458-020-09460-z>
- [3] Antonis C. Kakas and Pavlos Moraitis. 2003. Argumentation based decision making for autonomous agents. In *Proc. of 2nd Int. Joint Conf. on Autonomous Agents & Multiagent Systems, AAMAS, ACM*, 883–890. <https://doi.org/10.1145/860575.860717>
- [4] Antonis C. Kakas, Pavlos Moraitis, and Nikolaos I. Spanoudakis. 2019. GORGAS: Applying argumentation. *Argument & Computation* 10, 1 (2019), 55–81. <https://doi.org/10.3233/AAC-181006>
- [5] Erisa Karafilii, Linna Wang, and Emil C. Lupu. 2020. An Argumentation-Based Reasoner to Assist Digital Investigation and Attribution of Cyber-Attacks. *Forensic Science International: Digital Investigation* 32 (2020), 300925. <https://doi.org/10.1016/j.fsidi.2020.300925>
- [6] Vasileios G Mandikas and Nikolaos I Spanoudakis. 2025. An Effective Root-Finding Toolbox Using Computational Argumentation. In *2025 6th International Conference in Electronic Engineering & Information Technology (EEITE)*. IEEE. <https://doi.org/10.1109/EEITE65381.2025.11166218>
- [7] Julien Marçais, Nikolaos Spanoudakis, and Pavlos Moraitis. 2011. Using argumentation for ambient assisted living. In *International Conference on Engineering Applications of Neural Networks*. Springer, 410–419. https://doi.org/10.1007/978-3-642-23960-1_48
- [8] Leila Methmani, Virginia Dignum, and Andreas Theodorou. 2024. Clash of the Explainers: Argumentation for Context-Appropriate Explanations. In *Artificial Intelligence. ECAI 2023 International Workshops*. Springer Nature Switzerland, Cham, 7–23. https://doi.org/10.1007/978-3-031-50396-2_1
- [9] Ioannis Michalakis. 2025. *Argumentation-based decision support system for systems deployment: case study in the Ministry of Digital Governance*. Diploma Thesis. Technical University of Crete. <https://doi.org/10.26233/healink.tuc.104026>
- [10] Pavlos Moraitis and Nikolaos Spanoudakis. 2007. Argumentation-based agent interaction in an ambient-intelligence context. *IEEE Intelligent Systems* 22, 6 (2007), 84–93. <https://doi.org/10.1109/MIS.2007.101>
- [11] Nikolaos Spanoudakis, Antonis C. Kakas, and Adamos Koumi. 2022. Application Level Explanations for Argumentation-based Decision Making. In *1st International Workshop on Argumentation for eXplainable AI co-located with 9th International Conference on Computational Models of Argument (COMMA 2022), Cardiff, Wales, September 12, 2022 (CEUR Workshop Proceedings, Vol. 3209)*. Kristijonas Cyras, Timotheus Kampik, Oana Cocarascu, and Antonio Rago (Eds.). CEUR-WS.org.
- [12] Nikolaos I. Spanoudakis, Elena Constantinou, Adamos Koumi, and Antonis C. Kakas. 2017. Modeling Data Access Legislation with Gorgias. In *Advances in Artificial Intelligence: From Theory to Practice - 30th International Conference on Industrial Engineering and Other Applications of Applied Intelligent Systems, IEA/AIE 2017, Arras, France, June 27-30, 2017, Proceedings, Part II (Lecture Notes in Computer Science, Vol. 10351)*, Salem Benferhat, Karim Tabia, and Moonis Ali (Eds.). Springer, 317–327. https://doi.org/10.1007/978-3-319-60045-1_34
- [13] Nikolaos I. Spanoudakis, Georgios Gligoris, Antonis C. Kakas, and Adamos Koumi. 2022. Gorgias Cloud: On-line Explainable Argumentation. In *Computational Models of Argument - Proceedings of COMMA 2022, Cardiff, Wales, UK, 14-16 September 2022 (Frontiers in Artificial Intelligence and Applications, Vol. 353)*, Francesca Toni, Sylvia Polberg, Richard Booth, Martin Caminada, and Hiroyuki Kido (Eds.). IOS Press, 371–372. <https://doi.org/10.3233/FAIA220178>
- [14] Nikolaos I. Spanoudakis, Georgios Gligoris, Adamos Koumi, and Antonis C. Kakas. 2023. Explainable argumentation as a service. *Journal of Web Semantics* 76 (2023), 100772. <https://doi.org/10.1016/j.websem.2023.100772>
- [15] Nikolaos I. Spanoudakis, Antonis C. Kakas, and Pavlos Moraitis. 2016. Applications of Argumentation: The SoDA Methodology. In *ECAI 2016 - 22nd European Conference on Artificial Intelligence, 29 August-2 September 2016, The Hague, The Netherlands (Frontiers in Artificial Intelligence and Applications, Vol. 285)*, Gal A. Kaminka, Maria Fox, Paolo Bouquet, Eyke Hüllermeier, Virginia Dignum, Frank Dignum, and Frank van Harmelen (Eds.). IOS Press, 1722–1723. <https://doi.org/10.3233/978-1-61499-672-9-1722>
- [16] Nikolaos I. Spanoudakis, Konstantinos Kostis, and Katerina Mania. 2021. Web-Gorgias-B: Argumentation for All. In *Proceedings of the 13th International Conference on Agents and Artificial Intelligence, ICAART 2021, Volume 2, Online Streaming, February 4-6, 2021*, Ana Paula Rocha, Luc Steels, and H. Jaap van den Herik (Eds.). SCITEPRESS, 286–297. <https://doi.org/10.5220/0010269402860297>
- [17] Nikolaos I. Spanoudakis, Panteleimon Krinakis, and Dionysia Kolokotsa. 2024. A Methodology for Applying Decision Policies into Smart Buildings with the use of Computational Argumentation and IoT Technologies. In *5th International Conference in Electronic Engineering, Information Technology & Education (EEITE 2024)*. IEEE. <https://doi.org/10.1109/EEITE61750.2024.10654451>
- [18] Panayiotis Tanos, Ioannis Yiangou, Giorgos Prokopiou, Antonis Kakas, and Vasilios Tanos. 2024. Gynaecological Artificial Intelligence Diagnostics (GAID) GAID and Its Performance as a Tool for the Specialist Doctor. *Healthcare* 12, 2 (2024). <https://doi.org/10.3390/healthcare12020223>