

OMEGA^{*}: An Ontology-Driven Tool for Explaining Multi-Agent Path Finding

Bharath Muppasani, Ritirupa Dey, Biplav Srivastava, Vignesh Narayanan

University of South Carolina, USA

bharath@email.sc.edu, deyr@email.sc.edu, biplav.s@sc.edu, vignar@sc.edu

Abstract

Multi-Agent Path Finding (MAPF) algorithms provide highly optimized solutions for coordinating multiple agents in shared environments, yet their outputs lack explainability to human stakeholders. Existing explanation approaches, such as visual trace segmentation or logic-based reasoning, remain fragmented. In this demo, we present OMEGA, an *interactive explanation platform* that generates Natural Language (NL) explanations using the novel *Multi-Agent Planning Ontology (maPO)*. Our framework transforms raw MAPF planner execution logs into a semantic knowledge graph, enabling SPARQL-based explanations of collision events, replanning strategies, and efficiency trade-offs. A lightweight web interface allows users to query, visualize, and interpret planner decisions, thereby making MAPF solutions transparent and auditable. We conducted a user study that confirms the ontology-driven explanations are significantly clearer and more preferred than raw logs, underscoring the potential of semantic technologies for explainable multi-agent systems.

Code — github.com/ai4society/ma-planning-ontology

Website — ai4society.github.io/ma-planning-ontology/

Demo Video — shorturl.at/298j5

1 Introduction

Multi-Agent Path Finding (MAPF) addresses how multiple agents move to their goals in a shared space while avoiding collisions. While both search-based and learning-based planning approaches have demonstrated strong performance in terms of solution quality, optimality, and scalability, the decision-making process remains opaque to human users. Practitioners often ask why an agent had to wait, why a specific agent was chosen to replan, and how these choices changed overall cost or makespan. Consistent with user studies (Brandao et al. 2022), insufficient explanations make MAPF plans harder to understand and use in practice.

Prior explanation approaches in MAPF provide only partial coverage, such as visual segmentation methods that verify safety but do not explain plan choices (Almagor and Lahijanian 2020), or logic-based frameworks that answer specific “why” or “why-not” queries but do not capture broader plan evolution (Bogatarkan 2021). User studies

have identified categories of explanation needs ranging from causal reasoning to performance summaries (Brandao et al. 2022). Some methods restrict planner outputs to forms that are easier to explain, but they remain tied to specific algorithms (Kottinger, Almagor, and Lahijanian 2022).

We propose OMEGA, an ontology-driven tool that explains MAPF plans and agent behaviors. It uses a novel Multi-Agent Planning Ontology (*maPO*) and an online explanation framework that together deliver a *planner-agnostic, queryable, and provenance-complete solution*. *maPO* extends the Planning Ontology (Muppasani et al. 2024) with constructs for agents, paths, collisions, alerts, and replanning strategies, and reuses W3C and IEEE standards such as OWL-Time (Pan and Hobbs 2006), PROV-O (Lebo et al. 2013), SOSA (Janowicz et al. 2019), and CORA (Schlenoff et al. 2012). Our system, OMEGA, transforms a standardized planner log into a knowledge graph, answers competency questions through SPARQL queries, and displays synchronized textual and visual explanations in a web interface. We conducted a user study and found that our ontology-driven explanations were preferred over raw logs in 94% of cases, with an average clarity rating of 4.44/5 validated by binomial and Wilcoxon tests (Wagner-Menghin 2005; Woolson 2007).

2 Building maPO

Our *maPO* was designed to capture the unique requirements of multi-agent path finding while remaining lightweight, interoperable, and planner-agnostic. It reuses established vocabularies such as OWL-Time (Pan and Hobbs 2006), PROV-O (Lebo et al. 2013), SOSA (Janowicz et al. 2019), and CORA (Schlenoff et al. 2012), while extending the Planning Ontology (Muppasani et al. 2024) with concepts tailored for MAPF.

ma:Agent represents each actor, aligned with *sosa:Platform*, and links to *ma:AgentState* capturing location and time. This supports reasoning about movement, waiting, and synchronization.

ma:AgentSubPlan stores an agent’s plan with *ma:hasPlanCost* and *ma:belongsToAgent*. It is specialized as *ma:OriginalSubPlan* (before conflict) or *ma:ResolvedSubPlan* (after replanning). A *ma:JointPlan* aggregates subplans and records *ma:hasGlobalMakespan*.

ma:CollisionEvent encodes conflict type, location, time,

^{*}Ontology driven Multi-agent Explanation Generator
Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

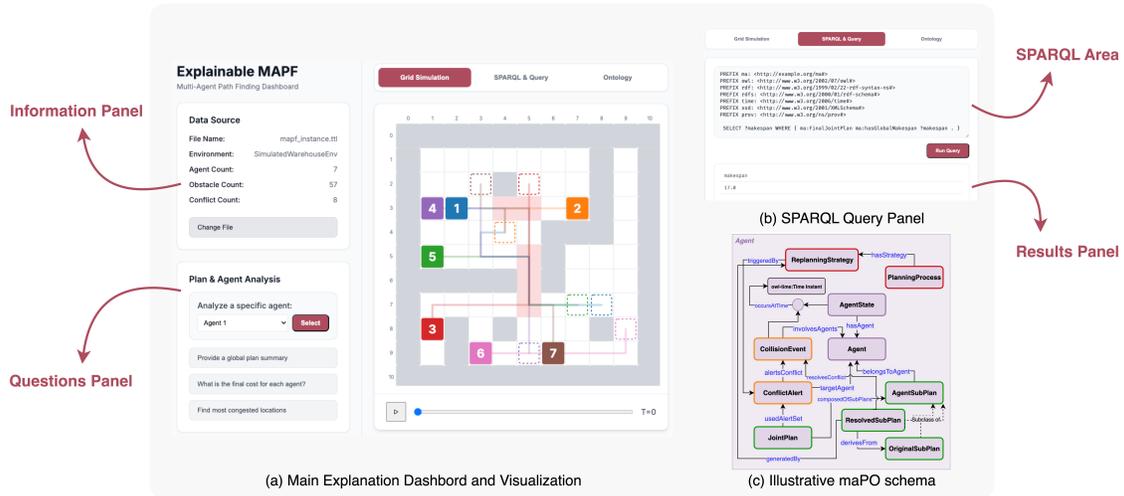


Figure 1: Demo interface for Explainable MAPF. (a) Dashboard for plan- and agent-level analysis. (b) SPARQL query panel for ontology-backed queries. (c) Illustrative *maPO* schema showing conflict–alert lifecycle (orange), replanning processes (red), and plan evolution (green). The full, up-to-date ontology is available on our explanation platform via its PURL URI.

and participants (*ma:involvesAgentsEvent*), enabling conflict-centric queries like listing all conflicts for an agent. *ma:ConflictAlert* links a conflict to the agent selected for replanning (*ma:targetAgent*) with a *ma:selectionRationale*. *ma:ReplanningStrategy* is modeled as a *prov:Activity* that turns an *ma:OriginalSubPlan* into a *ma:ResolvedSubPlan*, with provenance *prov:wasDerivedFrom* and *prov:wasGeneratedBy*, capturing causal chains from conflict to resolution.

The ontology was designed to support a suite of competency questions, including: (a) Listing all conflicts involving a specific agent and their participants. (b) Identifying all conflicts at a given location and time. (c) Comparing *ma:OriginalSubPlan* and *ma:ResolvedSubPlan* for an agent to explain plan differences. (d) Explaining why an agent waited at a location by linking idle segments to related conflicts. (e) Summarizing global plan metrics, including makespan, number of replans, and per-agent costs. By representing plans, conflicts, and strategies with explicit semantics, *maPO* enables queryable explanations that extend across planners and domains, reducing the cognitive load of interpreting raw execution logs.

3 Details of Online Explanation Framework

Our demo integrates *maPO* into a web-based framework (Figure 1). The pipeline has three stages: (i) a planner such as search or a learning-based method outputs a JSON log of environment, trajectories, conflicts, and metrics; (ii) a Python converter maps this to RDF triples consistent with *maPO*, producing a Turtle knowledge graph; and (iii) the web application ingests the graph, executes SPARQL queries, and presents explanations in natural language alongside synchronized grid visualizations.

The interface provides three views: *Grid Simulation*, *SPARQL Query*, and *Ontology*. The dashboard (Fig. 1a) shows agents moving along their paths with start/goal mark-

ers and conflict highlights. Selecting a global or agent-specific question runs a SPARQL query and returns both a textual summary and visual overlay (e.g., explaining an idle action as a response to a nearby conflict). The query panel (Fig. 1b) exposes templates for inspection and modification, supporting auditability. The ontology view (Fig. 1c) visualizes *maPO*’s classes and properties.

Because the framework only depends on the standardized JSON log schema, it applies uniformly across centralized, suboptimal, anytime, and learning-based planners. The entire system runs client-side using D3 for visualization, N3 for RDF parsing, and Comunica for SPARQL querying, ensuring lightweight deployment without server dependencies.

4 Conclusion

This demo presents *maPO* and an online explanation framework for MAPF that together provide causal, contrastive, and performance-centered explanations in a planner-agnostic way. By grounding explanations in a semantic ontology aligned with community standards, the system enables formal reasoning and interoperability. The web platform shows how explanations can be delivered interactively, combining natural language with synchronized visualizations. Empirical evaluation indicates that ontology-driven explanations improve clarity and trust compared to raw logs. Future directions include extending the framework to real robot data, integrating execution monitoring with SOSA, enhancing natural language generation, and applying the system to large-scale logistics and swarm robotics scenarios.

Acknowledgments

This work was supported in part by a J.P. Morgan Chase Faculty Award, by the Air Force Office of Scientific Research under Award No. FA9550-24-1-0228, and by the National Science Foundation under Award No. 2337998.

References

- Almagor, S.; and Lahijanian, M. 2020. Explainable multi agent path finding. In *AAMAS*.
- Bogatarkan, A. 2021. Flexible and explainable solutions for multi-agent path finding problems. *arXiv preprint arXiv:2109.08299*.
- Brandao, M.; Mansouri, M.; Mohammed, A.; Luff, P.; and Coles, A. 2022. Explainability in multi-agent path/motion planning: User-study-driven taxonomy and requirements. In *International Conference on Autonomous Agents and Multi-agent Systems (AAMAS)*.
- Janowicz, K.; Haller, A.; Cox, S. J.; Le Phuoc, D.; and Lefrançois, M. 2019. SOSA: A lightweight ontology for sensors, observations, samples, and actuators. *Journal of Web Semantics*, 56: 1–10.
- Kottinger, J.; Almagor, S.; and Lahijanian, M. 2022. Conflict-based search for explainable multi-agent path finding. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 32, 692–700.
- Lebo, T.; Sahoo, S.; McGuinness, D.; Belhajjame, K.; Cheney, J.; Corsar, D.; Garijo, D.; Soiland-Reyes, S.; Zednik, S.; and Zhao, J. 2013. Prov-o: The prov ontology.
- Muppasani, B.; Gupta, N.; Pallagani, V.; Srivastava, B.; Mutharaju, R.; Huhns, M. N.; and Narayanan, V. 2024. Building a Plan Ontology to Represent and Exploit Planning Knowledge and Its Applications. In *Eighth International Conference on Data Science and Management of Data (CODS-COMAD'24), India*.
- Pan, F.; and Hobbs, J. R. 2006. Time ontology in owl. *W3C working draft, W3C*, 1(1): 1.
- Schlenoff, C.; Prestes, E.; Madhavan, R.; Goncalves, P.; Li, H.; Balakirsky, S.; Kramer, T.; and Miguelanez, E. 2012. An IEEE standard ontology for robotics and automation. In *2012 IEEE/RSJ international conference on intelligent robots and systems*, 1337–1342. IEEE.
- Wagner-Menghin, M. M. 2005. Binomial test. *Encyclopedia of statistics in behavioral science*.
- Woolson, R. F. 2007. Wilcoxon signed-rank test. *Wiley encyclopedia of clinical trials*, 1–3.