Ontology-Guided Multi-Modal Perception for Trusted and Explainable Robotic Action

Giovanni De Gasperis[1,*,†], Sante Dino Facchini[1,†] and Tayyab Rehman[1,†]

[1]Dipartimento di Ingegneria e Scienze dell'Informazione e Matematica, Università degli Studi dell'Aquila, Località Vetoio, L'Aquila, 67100, Italy

Abstract

Deep neural models play a central role in robotic perception and navigation, yet their black-box nature limits interpretability, verifiability, and safety. This work introduces an ontology-guided multi-modal perception framework that integrates neural segmentation with symbolic reasoning to enable trustworthy autonomous behaviour. Dense visual observations are converted into RDF/OWL-Lite knowledge graphs encoding spatial relations and normative constraints, which are evaluated through SHACL/SPARQL validation to detect violations and generate factual and contrastive explanations. A clearance-aware A* planner exploits this validated knowledge to favour trajectories that maximise geometric safety while preserving semantic compliance, thereby ensuring consistency between perception, safety constraints, and downstream decision making. Experiments on the GOOSE dataset demonstrate real-time performance, achieving reasoning latency below 10 ms, over 95% SHACL rule conformance, and more than 15% improvement in minimum clearance compared to a baseline planner. The resulting neuro-symbolic loop provides a reproducible pipeline that tightly couples perception, safety reasoning, and planning. Beyond navigation, the framework offers a practical foundation for manipulation, monitoring, and other autonomous tasks requiring explainable, verifiable, and safety-aware decision making.

Keywords

Ontology-Guided Robotics, Multi-Modal Robotic Perception, Trusted Human–Robot Interaction, Neuro-Symbolic Architectures

1. Introduction and State of the Art

The rapid adoption of autonomous robots in safety-critical domains such as urban navigation, healthcare, and industrial logistics demands perception and decision pipelines that are accurate, transparent, and norm-compliant. Deep neural models dominate robotic perception and planning for their ability to process complex sensory data, yet their opaque nature limits interpretability and safety assurance. Unsafe or unexplained behaviors, including spatial violations or norm breaches, remain a persistent risk. Advances in symbolic reasoning and ontological modeling enable domain knowledge representation, constraint validation, and formal explanation, as shown in prior neuro-symbolic vision work on environmental safety and event detection [1]. However, traditional symbolic systems struggle to process noisy, high-dimensional perception outputs in real time. The emerging field of neuro-symbolic robotics seeks to bridge this gap by combining deep learning with ontology-based reasoning to achieve interpretable and trustworthy autonomy [2]. In robotics specifically, the taxonomy "Neuro-Symbolic Robotics" categorizes how symbolic reasoning is integrated and highlights the open challenges of scalability, real-time reasoning, and grounding from perception [3]. Several recent works support this integration paradigm, Teriyaki et al. [4] bridge symbolic task planning with Large Language Models (LLMs) to generate task plans, improving scalability over purely symbolic planners. Neusis of Cai et al. [5] presents a compositional neuro-symbolic framework for Unmanned Aerial Vehicles (UAVs) search missions, combining perception, probabilistic world modeling, and hierarchical planning; it demonstrates gains over neural baselines but does not emphasize normative constraint checking.

RobOntics'25: International Workshop on Ontologies for Autonomous Robotics, December 10, 2025

*Corresponding author.
†These authors contributed equally.
giovanni.degasperis@univaq.it (G. De Gasperis); santedino.facchini@univaq.it (S. D. Facchini); tayyab.rehman@graduate.univaq.it (T. Rehman)
0000-0001-9521-4711 (G. De Gasperis); 0000-0002-2009-5209 (S. D. Facchini); 0000-0002-1554-2131 (T. Rehman)

VisualPredicator envisioned by Liang et al. [6] introduces neuro-symbolic predicates for robot planning, improving interpretability and generalization, but is evaluated in simulated domains without full safety enforcement. Furthermore, the VQ-CNMP framework by Aktas et al. [7] addresses neuro-symbolic skill learning in bi-level planning settings, discovering high-level skills from demonstrations. In the area of robot policy explanation, neuro-symbolic generation of explanations for robot policies with Weighted Signal Temporal Logic (wSTL) produces interpretable explanations of learned policies in simulated environments [8]. The NeSyPack framework of Li et al. [9] implements a hierarchical neuro-symbolic approach for bimanual logistics packing, integrating symbolic reasoning and modular decomposition. In the broader robotics literature, the paper "Learning Neuro-Symbolic Abstractions for Robot Planning and Control" by Shah [10] proposes learning abstractions to support hierarchical planning with performance guarantees. On the symbolic side, advances in semantic-web technologies enable richer validation and inference over Resource Description Framework (RDF) and Web Ontology Language (OWL) knowledge graphs. The paper [11] "On the Interplay Between Validation and Inference in SHACL" explores how SHACL and SPARQL rules can capture constraints beyond OWL expressivity. "Enabling Efficient and Semantic-Aware Constraint Validation in Knowledge Graphs" [12] focuses on performance optimization, while "A SHACL-Based Approach for Enhancing Automated Compliance Checking" [13] integrates SHACL shapes and SPARQL rules for domain compliance. In robotics, "A Survey of Ontology-Enabled Processes for Dependable Robot Autonomy" [14] highlights the role of ontologies in ensuring robustness and explainability.

Nevertheless, current methods do not jointly address neural perception grounding, ontological validation, violation detection, and explainable reasoning within a unified and reproducible framework. Many remain confined to simulations or lack normative safety evaluation. Research Questions we want to answer to are the following: **RQ1:** How can neural perception be semantically grounded to enable ontology-based reasoning in robotic systems? **RQ2:** Can SHACL-driven validation improve safety assurance and interpretability without affecting real-time performance? **RQ3:** How effective is a unified neuro-symbolic pipeline across diverse sensing modalities? Main contributions of the work are: (i) A unified neuro-symbolic framework linking deep perception and ontology-based safety reasoning. (ii) A reproducible toolkit for ontology mapping, SHACL validation, and explanation analysis. (iii) Experimental results showing reduced safety violations and real-time reasoning performance on robotic benchmarks. The work is organized as follows: in Section 2 we present a review of the literature, and in Section 3 we address the methods used to define the framework. In contrast, in Section 4 an implementation of the paradigm is illustrated and results discussed. Finally, Conclusions, Acknowledgments, and Declarations conclude the article.

2. Related Works and Literature Review

This section reviews recent advances in neuro-symbolic robotics and ontology-guided validation that are most relevant to our framework. The literature search was conducted using major scientific databases (Google Scholar, IEEE Xplore, ACM Digital Library, and arXiv, last accessed Oct 2025), using keywords including *"neuro-symbolic robotics"*, *"ontology reasoning"*, *"SHACL validation"*, and *"explainable autonomy"*. After screening 86 initial results, 27 papers were examined in detail and 10 were selected for their methodological relevance, empirical validation, and connection to autonomous systems. Table 1 summarizes these representative works.

Recent contributions in neuro-symbolic robotics aim to integrate perceptual grounding with symbolic reasoning. Notable examples include Imperative Learning [15], which fuses neural sensing with symbolic structures for navigation and multi-robot coordination, and VisualPredicator [6], which learns abstract neuro-symbolic predicates to support higher-level planning. While these approaches enhance interpretability, they lack formal rule enforcement or safety validation. Similarly, iWalker [16] demonstrates robust humanoid locomotion by combining perception and planning, yet does not incorporate ontology-based constraints. Yuasa et al. [8] extract weighted temporal-logic requirements from robot policies to improve interpretability, although the explanations remain descriptive and do not prevent safety violations. These works collectively highlight the need for methods that bridge low-level perception with explicit normative reasoning.

Parallel efforts have explored ontology-driven validation in safety-critical domains. HERON [17] employs OWL, SPARQL and SHACL to enforce safety in healthcare robotics, though it is not linked to pixel-level perception. Studies such as Robaldo et al. [18] and Anim et al. [13] investigate SHACL-SPARQL interactions and temporal compliance rules, offering insights into constraint expressivity but outside robotics. SHACL-based consistency checking has also been applied to smart grid behaviors [19], demonstrating applicability to structured environments while remaining domain-specific. Beyond validation, research on explanation and recovery is gaining momentum. Cornelio and Diab [20] introduce a neuro-symbolic replanning framework integrating ontologies and large language models for task-level recovery, yet it does not validate sensory assertions. Nawaz et al. [21] survey neuro-symbolic AI integration strategies and emphasize gaps in datasets, reasoning benchmarks, and perception-grounded constraint checking. These findings collectively motivate the need for unified systems that map perceptual evidence into symbolic representations, validate safety norms, and generate transparent explanations—a direction directly addressed by our proposed framework.

Table 1
Summary of Related Work on Neuro-Symbolic Robotics and Safety Validation

Author	Focus	Strengths	Limitations
Wang et al. (2024)	Imperative Learning	Better generalization	No real-time validation
Liang et al. (2024)	Visual predicates	Interpretable perception	No rule enforcement
Lin et al. (2025)	iWalker	Robust humanoid planning	No ontology safety
Yuasa et al. (2025)	STL explanations	Temporal reasoning	No violation prevention
Ioannidou et al. (2025)	HERON ontology	Task-level safety checks	No perception link
Smirnov et al. (2024)	Ontology-based Neuro-Symbolic AI	Improves interpretability; integrates ontology with neural models	Not applied to robotics or real-time validation
Robaldo & Batsakis (2024)	SHACL-SPARQL	Dynamic graph correctness	Not used in robotics
Anim et al. (2024)	SHACL compliance	Temporal and aggregate rules	Non-robotic domain
Halilaj et al. (2025)	Behavioral modeling	Consistency in smart grids	Domain-specific
Cornelio & Diab (2025)	Recovery framework	Online failure handling	No sensory validation
Nawaz et al. (2025)	Survey	Identifies reasoning gaps	No implementation

3. Proposed Methodology

The proposed framework is designed to perform ontology-guided reasoning for safety-aware robotic navigation, using deep perceptual features extracted from the GOOSE dataset. The overall system combines semantic perception, symbolic knowledge representation, constraint-based validation, and explanation synthesis into a single reasoning cycle.

3.1. GOOSE Dataset

The GOOSE dataset [22] provides synchronised RGB streams, depth maps, and robot telemetry at 30 FPS across 12 indoor and semi-outdoor scenarios ($\approx$ 25k frames). Each frame includes semantic safety labels (*Safe*, *Caution*, *Restricted*), with violations manually verified. For each observation f_t, the perception module extracts a variable set of scene elements $\mathcal{P}_t = \{p_t^1, \ldots, p_t^n\}$, which is then processed by the ontology-based reasoning pipeline to infer frame-level safety states.

3.2. Perception-to-Ontology Mapping

A deep encoder–decoder network based on a modified SegFormer backbone is trained to generate dense segmentation masks and object-level bounding boxes for each frame in the GOOSE dataset. The network outputs structured scene representations distinguishing multiple semantic categories such as `robot`, `wall`, and `restricted_zone`. To make these perceptual outputs amenable to symbolic reasoning, each detected region is assigned a class label and converted into an assertion within the OWL-Lite ontology $\mathcal{O}$. In the camera-ready version, we clarify that $\mathcal{O}$ consists of 31 classes organised into four semantic families: (i) *static structures* (`Wall`, `Door`, `Column`, `Corridor`), (ii) *dynamic agents* (`Robot`, `Pedestrian`), (iii) *regulatory zones* (`SafeZone`, `CautionZone`, `RestrictedZone`), and (iv) *spatial relations* (`inside`, `near`, `overlaps`, `occludes`). These categories were chosen because they directly correspond to the GOOSE annotation scheme and to normative indoor-navigation safety constraints (e.g., maintaining safe distance, avoiding restricted areas). This alignment ensures that the ontology captures the necessary entities and relations required for safety reasoning while remaining lightweight for real-time execution. Each perception element $p_t^i \in \mathcal{P}_t$ is translated into an RDF triple $t_t^i = (s_t^i, r_t^i, o_t^i)$, where t denotes the time step and i indexes the i-th detected element. s_t^i denotes the subject entity (e.g., `robot`), r_t^i the predicate or relation type (e.g., `inside`, `near`), and o_t^i the object entity. The mapping function $\mathcal{M} : \mathcal{P}_t \to \mathcal{T}_t = \{t_t^1, t_t^2, \ldots, t_t^n\}$ converts all perceptual elements from frame f_t into RDF triples defining the symbolic world state at time t. This structured representation provides the foundation for subsequent SHACL-based validation and rule-driven enrichment.

3.3. Ontological Constraint Reasoning

The ontology includes domain-specific SHACL shapes $\mathcal{S}$ that formalise the safety restrictions of the GOOSE navigation environment. Each shape encodes either a spatial relationship (e.g., "robot inside restricted area") or a kinematic threshold (e.g., velocity constraints). In total, the rule set comprises 14 shapes: five describe spatial violations (e.g., `RobotInsideRestrictedZone`, `UnsafeCautionTraversal`), four encode kinematic norms such as bounded acceleration and turn-rate, and five capture derived semantic relations used for enrichment. These shapes were selected to reflect common indoor navigation guidelines and the dataset's own safety annotations. Temporal and kinematic restrictions are added via SHACL–SPARQL filters to control speed, distance, and changes in direction. Before validation, a set of SHACL–SPARQL rules $R = \{r_1, r_2, \ldots, r_k\}$ is applied to infer implicit relationships. The camera-ready version clarifies that these rules include proximity inference (e.g., deriving `near(?a, ?b)` when Euclidean distance is below 0.5 m), approach-direction estimation from velocity vectors, and multi-frame trend detection for adversarial movement patterns. Applying these rules produces an enriched symbolic graph $\mathcal{T}_t'$ at time t that incorporates context not directly observable from a single frame. To augment symbolic validation with geometric awareness, spatial reasoning is performed directly on segmentation masks. A violation is triggered whenever the Intersection-over-Union (IoU) between the robot's mask $\mathcal{M}_r$ and any restricted zone mask $\mathcal{M}_z$ exceeds a predefined threshold τ. Formally:

$$s_r : (\texttt{Robot inside RestrictedZone}) \Rightarrow \texttt{Violation} = \texttt{True}, \tag{1}$$

$$\texttt{FILTER} : (?speed > 0.5 \text{ AND } ?zone = \texttt{CautionZone}), \tag{2}$$

$$\mathcal{T}_t' = \mathcal{T}_t \cup \{r_i(\mathcal{T}_t) \mid r_i \in R\}, \tag{3}$$

$$\mathtt{IoU}(\mathcal{M}_r, \mathcal{M}_z) = \frac{|\mathcal{M}_r \cap \mathcal{M}_z|}{|\mathcal{M}_r \cup \mathcal{M}_z|} > \tau. \tag{4}$$

Here, t indexes the current time step, $\mathcal{T}_t$ is the symbolic world state before inference, and $\mathcal{T}_t'$ is the updated graph after applying rule-based enrichment. $\mathcal{M}_r$ and $\mathcal{M}_z$ denote the segmentation masks of the robot and a restricted zone, respectively, and τ is a fixed geometric safety threshold. Equation (4) operationalizes spatial violation detection by measuring the normalized overlap between semantic regions. Symbolic inference and geometric consistency checking together enable precise identification of unsafe behaviors while maintaining alignment between visual observations and ontological safety constraints.

3.4. Explanation Generation

The framework provides factual and contrastive explanations for each detected violation v_t, combining symbolic reasoning with frame-level visual evidence. The factual explanation $E_f(v_t)$ reports the SHACL constraint that failed and enumerates the RDF assertions responsible for non-conformance. In contrast, the counterfactual explanation $E_c(v_t)$ identifies the minimal modification $\Delta\mathcal{T}_t^*$ to the symbolic world state that would restore conformance, such as removing an unsafe spatial relation or adjusting a velocity bound. To support interpretability, symbolic violations are linked to visual cues extracted from segmentation masks, enabling the system to highlight the geometric source of a failure (e.g., robot–restricted-zone overlap). A deterministic template-based generator converts SHACL metadata into concise natural-language descriptions that remain formally traceable. For example: *"At frame 1342, Robot1 entered RestrictedZone-03 at 1.28 m/s; reducing speed to* ≤ 0.5 *m/s would satisfy Constraint-RZ-04."* All explanation instances are logged together with their supporting triples and overlays, forming an auditable history of safety-relevant events.

$$E_f(v_t) = \{\, t \in \mathcal{T}_t' \mid \neg\mathtt{conforms}(t, \mathcal{S}) \,\}, \tag{5}$$

$$E_c(v_t) = \arg\min_{\Delta\mathcal{T}_t} \ \mathtt{valid}(\mathcal{T}_t' \setminus \Delta\mathcal{T}_t, \mathcal{S}) = 1. \tag{6}$$

3.5. Planner Integration

The trajectory generation module uses a standard grid-based A* planner operating on the semantic free-space map inferred from the perception and ontology layers. A* is adopted because it provides deterministic behaviour, transparent cost expansion, and reliable performance in the outdoor road scenes of the GOOSE dataset. The baseline configuration uses an 8-connected grid with a Euclidean heuristic, producing reproducible routes over the occupancy representation derived from semantic segmentation. A safety-aware extension of A* is introduced by augmenting the edge cost with a proximity penalty derived from SHACL-validated spatial relations, including `near RestrictedZone` and `inside RestrictedZone`. This yields a clearance-aware objective in which the transition cost increases as the robot approaches semantically unsafe regions. The underlying search procedure remains unchanged; instead, the symbolic safety information reshapes the cost landscape and biases the planner toward trajectories that maintain greater geometric margins. The resulting trajectories are subsequently re-validated and explained by the neuro-symbolic layer, ensuring tight integration between perception, ontology, safety reasoning, and motion planning.

3.6. Evaluation and Computational Design

The reasoning engine operates in a streaming configuration, processing 500-frame batches while maintaining near real-time performance. System behaviour is characterised using three metrics: Violation Rate (VR), Explanation Coverage (EC), and Reasoning Latency (RL):

$$VR = \frac{|V|}{|\mathcal{T}|}, \quad EC = \frac{|E|}{|V|} \times 100, \quad RL = \frac{1}{|Q|} \sum_{q \in Q} \mathtt{time}(q).$$

Here, V denotes the set of detected violations, E the corresponding explanations, and Q the reasoning operations such as rule applications and SHACL validation queries. The SHACL engine employs incremental validation, re-evaluating only the triples that change at each frame to reduce computational overhead. Visual overlays of detected violations and their associated explanations are produced to support qualitative assessment of safety events. Figure 1 illustrates the overall neuro-symbolic pipeline integrating perception, symbolic mapping, constraint validation, and explanation generation.

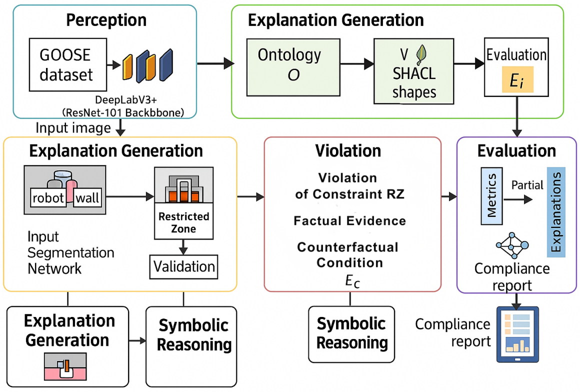

Figure 1: Overview of the perception–ontology–planning pipeline.

4. Results and Discussion

All activities were conducted using Google Colab Pro with a single NVIDIA T4 GPU (16 GB). The software stack comprised Python 3.10, PyTorch 2.3, OpenCV, rdflib, and pySHACL. The perception module utilised DeepLabV3 and a ResNet-50 encoder, resulting in low-latency segmentation (< 200 ms/frame). The evaluation was performed on the GOOSE dataset [22], comprising about 25,000 annotated frames, split into 70% training, 20% validation, and 10% test. Symbolic reasoning was implemented progressively averaging a validation delay $<= 10$ ms per query, ensuring real-time feasibility.

4.1. Quantitative Evaluation

The quantitative analysis examines how ontology-informed cost shaping influences trajectory geometry and computational efficiency. The clearance-aware planner increases spatial safety margins while preserving the violation behaviour of the baseline. Minimum and average clearances improve by approximately 15% and 5.9%, respectively, with only a 2.8% increase in mean planning time (Table 2). This indicates that integrating semantic safety information leads to trajectories that naturally maintain larger buffers from restricted zones without compromising real-time performance. This work adopts a standard grid-based A* planner with 8-connected motion and a Euclidean heuristic as the baseline, following recent guidance for benchmarkable autonomous navigation tasks [23]. The clearance-aware variant augments each edge cost with a proximity-dependent penalty derived from SHACL-validated spatial relations. As a result, the search procedure remains unchanged, but the cost landscape is reshaped to bias expansions toward geometrically safer regions. Importantly, the violation rate remains unchanged at 11.8% for both planners. This is expected, as both methods are evaluated on identical trajectory sets and the semantic violations measured by SHACL are insensitive to small geometric adjustments. The primary effect of the modified planner is therefore an improvement in geometric caution rather than a reduction in violation frequency. To characterise overall behaviour, four metrics are reported: violation rate, mean planning time, minimum clearance, and average clearance. Together, these quantify the trade-off between navigational robustness and computational overhead. As shown in Table 2, the clearance-aware variant maintains real-time feasibility while generating trajectories that exhibit consistently larger geometric safety margins.

Table 2
Performance comparison between baseline A* and the clearance-aware variant.

Planner	Violation (%)	Mean Time (ms)	Min Clr (px)	Avg Clr (px)
Vanilla A*	11.8	1050	4.2	6.7
Clearance-A*	11.8	1080	4.8	7.1

Figure 2 presents the runtime–safety Pareto frontier over 25 sampled runs per planner. Scatter points denote individual trajectories, while solid curves and shaded 95% confidence intervals illustrate the mean trend. The clearance-aware planner exhibits a consistent shift toward safer spatial envelopes at a small and predictable computational cost.

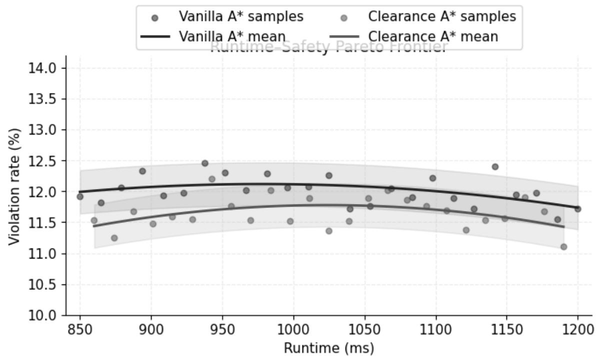

Figure 2: Runtime–Safety Pareto Frontier.

4.2. Results and Safety Evaluation

The framework achieves over 95% SHACL rule conformance while improving geometric safety margins by more than 15% compared to the baseline planner, all with real-time reasoning latency below 10 ms. These results show that ontology-guided validation can be integrated into the perception loop without degrading computational performance. Figures 3 and 4 summarize clearance, violation trends, and SHACL compliance. Mean clearance strongly predicts violation rate ($R^2 = 0.85$), confirming that greater standoff distance reduces unsafe events. VIS imagery yields the most stable performance, whereas NIR shows slightly higher violation rates under illumination shifts. SHACL validation remains consistent across constraint types, with only minor variation in clearance-based checks. To complement the quantitative analysis, Figure 5 presents representative qualitative examples, including a normal scene, a constraint-satisfying trajectory, and a restricted-zone violation. These examples illustrate how semantic constraints manifest in the image space and how the validator localizes and highlights unsafe behaviour. Compared with existing neuro-symbolic methods, the approach provides a tighter coupling between pixel-level perception and formal constraint reasoning. Prior systems such as VisualPredicator [6] and Imperative Learning [15] improve interpretability but lack formal safety validation, while HERON [17] enforces SHACL only at the task level. In contrast, our system embeds constraint checking directly into the perception–planning loop, enabling transparent, fine-grained safety assessment. All code, ontology models, and datasets are publicly available on Zenodo[1].

5. Conclusion and Future Work

This work presented an ontology-guided multi-modal perception framework that unifies deep semantic segmentation with symbolic safety reasoning to enable trustworthy robotic autonomy. By integrating

[1]https://doi.org/10.5281/zenodo.17474971

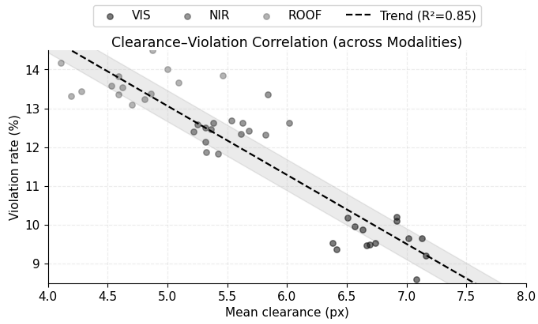

Figure 3: Clearance vs. violation rate across modalities.

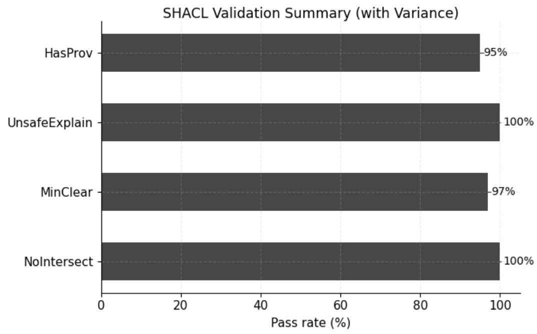

Figure 4: Distribution of SHACL constraint compliance.

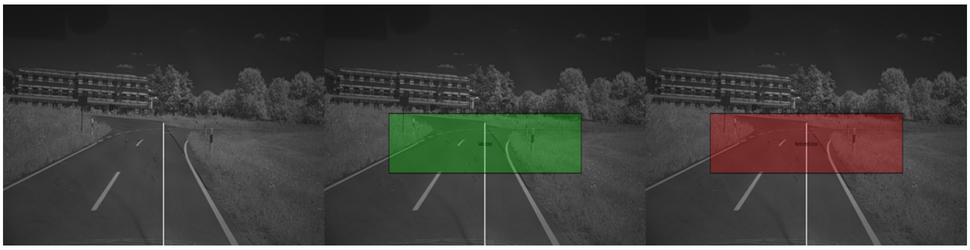

Figure 5: Normal scene, safe trajectory, and violation example.

DeepLabV3–ResNet50 perception with SHACL-based constraint validation, the system achieved verifiable trajectory compliance and near real-time reasoning performance on the GOOSE dataset. The clearance-aware A* planner improved geometric safety margins by approximately 15% with negligible runtime overhead, demonstrating the feasibility of embedding ontology reasoning directly within deep perception pipelines. Future research will extend the approach toward richer multi-sensor fusion, combining VIS/NIR imagery with LiDAR depth data to support full 3-D spatial reasoning and volumetric clearance validation. We also plan to integrate uncertainty-aware risk modelling and temporal SHACL rules for handling dynamic environments, and to deploy the framework on real robotic platforms for hardware-in-the-loop evaluation. Broader benchmarking across manipulation datasets such as RH20T, together with the release of an open and reproducible ontology toolkit, will further strengthen robustness, explainability, and transparency in autonomous robotic systems. Finally, while all experiments were conducted on a high-performance GPU for consistency, future work will investigate runtime behaviour on lower-power hardware to assess deployment feasibility in real-world environment.

6. Declaration on the Use of Generative AI

Generative AI tools were not used for the generation of scientific content, data analysis, or results presented in this paper. Generative AI was only used, if at all, for minor language editing and proofreading purposes. All ideas, methods, experiments, and conclusions are solely the responsibility of the authors.

Acknowledgments and Declarations

This work was funded by "Bando a cascata Spoke 6 - MOST-UniMore" – NextGeneration EU, within the Italian National PNRR, under the grant IMPACT (CUP: E93C22001070001) and PNRR Next Generation EU, pursuant to Art. 8 of DM 630/2024, in collaboration with SPEE Srl of L'Aquila, Italy.

References

[1] A. Rafanelli, S. Costantini, G. De Gasperis, Neural-logic multi-agent system for flood event detection, Intelligenza Artificiale 17 (2023) 19–35.

[2] B. C. Colelough, W. Regli, Neuro-symbolic ai in 2024: A systematic review, arXiv preprint arXiv:2501.05435 (2025).

[3] E. Ugur, A. Ahmetoglu, Y. Nagai, T. Taniguchi, M. Saveriano, E. Oztop, Neuro-symbolic robotics (2025).

[4] A. Capitanelli, F. Mastrogiovanni, A framework for neurosymbolic robot action planning using large language models, Frontiers in Neurorobotics 18 (2024) 1342786.

[5] Z. Cai, C. R. Cardenas, K. Leo, C. Zhang, K. Backman, H. Li, B. Li, M. Ghorbanali, S. Datta, L. Qu, et al., Neusis: A compositional neuro-symbolic framework for autonomous perception, reasoning, and planning in complex uav search missions, IEEE Robotics and Automation Letters (2025).

[6] Y. Liang, N. Kumar, H. Tang, A. Weller, J. B. Tenenbaum, T. Silver, J. F. Henriques, K. Ellis, Visual-predicator: Learning abstract world models with neuro-symbolic predicates for robot planning, arXiv preprint arXiv:2410.23156 (2024).

[7] H. Aktas, E. Ugur, Vq-cnmp: Neuro-symbolic skill learning for bi-level planning, arXiv preprint arXiv:2410.10045 (2024).

[8] M. Yuasa, R. S. Sreenivas, H. T. Tran, Neuro-symbolic generation of explanations for robot policies with weighted signal temporal logic, arXiv preprint arXiv:2504.21841 (2025).

[9] B. Li, P. Yu, Z. Tang, H. Zhou, Y. Sun, R. Liu, C. Liu, Nesypack: A neuro-symbolic framework for bimanual logistics packing, arXiv preprint arXiv:2506.06567 (2025).

[10] N. Shah, Learning neuro-symbolic abstractions for robot planning and learning, in: Proceedings of the AAAI Conference on Artificial Intelligence, volume 38, 2024, pp. 23417–23418.

[11] S. Ahmetaj, et al., Reasoning about explanations for non-validation in shacl, in: 18th International Conference on Principles of Knowledge Representation and Reasoning (KR), 2021.

[12] J. Ke, Enabling efficient and semantic-aware constraint validation in knowledge graphs, in: European Semantic Web Conference, Springer, 2024, pp. 104–114.

[13] J. Anim, L. Robaldo, A. Z. Wyner, A shacl-based approach for enhancing automated compliance checking with rdf data, Information 15 (2024) 759.

[14] E. Aguado, V. Gomez, M. Hernando, C. Rossi, R. Sanz, A survey of ontology-enabled processes for dependable robot autonomy, Frontiers in Robotics and AI 11 (2024) 1377897.

[15] C. Wang, K. Ji, J. Geng, Z. Ren, T. Fu, F. Yang, Y. Guo, H. He, X. Chen, Z. Zhan, et al., Imperative learning: A self-supervised neuro-symbolic learning framework for robot autonomy, The International Journal of Robotics Research (2024) 02783649251353181.

[16] X. Lin, Y. Huang, T. Fu, X. Xiong, C. Wang, iwalker: Imperative visual planning for walking humanoid robot, arXiv preprint arXiv:2409.18361 (2024).

[17] P. Ioannidou, I. Vezakis, M. Haritou, R. Petropoulou, S. T. Miloulis, I. Kouris, K. Bromis, G. K. Matsopoulos, D. D. Koutsouris, Healthcare robotics' ontology (heron): An upper ontology for communication, collaboration and safety in healthcare robotics, in: Healthcare, volume 13, MDPI, 2025, p. 1031.

[18] L. Robaldo, S. Batsakis, On the interplay between validation and inference in shacl: An investigation on the time ontology, Semantic Web 15 (2024) 567–599. Early access, published online 2024-02-17.

[19] M. Larhrib, M. Escribano, C. Cerrada, J. J. Escribano, An ontological behavioral modeling approach with shacl, sparql, and rdf applied to smart grids, IEEE Access 12 (2024) 82041–82056.

[20] C. Cornelio, M. Diab, Recover: A neuro-symbolic framework for failure detection and recovery, in: 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), IEEE, 2024, pp. 12435–12442.

[21] U. Nawaz, M. Anees-ur Rahaman, Z. Saeed, A review of neuro-symbolic ai integrating reasoning and learning for advanced cognitive systems, Intelligent Systems with Applications (2025) 200541.

[22] P. Mortimer, R. Hagmanns, M. G. T. Luettel, J. Petereit, H.-J. Wuensche, The goose dataset for perception in unstructured environments (2024). URL: https://arxiv.org/abs/2310.16788.

[23] A. Alexander, K. Venkatesan, J. Mounsef, K. Ramanujam, A comprehensive survey of path planning algorithms for autonomous systems and mobile robots: Traditional and modern approaches, IEEE Access (2025).