

Normative-Aware Ontology for Assistive Robotics: Bridging Legal Accountability and Cognitive Autonomy in Elderly Care

Alberto Borboni^{1,†}, Giorgio Pedrazzi^{2,*,†} and Fabio Zanoletti^{3,†}

¹University of Brescia, Department of Industrial and Mechanical Engineering, Via Branze 38, Brescia, 25123, Italy

²University of Brescia, Department of Law, Via Battaglie 58, Brescia, 25123, Italy

³University of Brescia, Department of Industrial and Mechanical Engineering, Via Branze 38, Brescia, 25123, Italy

Abstract

As assistive robots move into home-based rehabilitation and elderly care, technical autonomy must be reconciled with legal accountability, data protection, privacy, and safety. This paper proposes an ontology-centered normative decision architecture that embeds compliance-by-design into robot planning and monitoring. The approach combines a precedence-aware constraint lattice with a multi-layer reasoning stack to operationalize legal requirements alongside clinical workflows. The ontology is modular and adopts IEEE 7007 concepts to structure explanation artifacts and responsibility attributions for automated decisions. Two execution primitives translate norms into behavior to enable, block, or mitigate actions and generate audience-appropriate explanations. Through illustrative scenarios grounded in the Italian regulatory context and a post-stroke telerehabilitation setting, the architecture shows how explicit normative precedence supports gated compliance and traceable accountability. Limitations are discussed and future directions are outlined on multi-jurisdictional profiles, interoperable evidence exchange, and governance tooling, advancing a portable, auditable substrate for safe, explainable, and norm-adaptive assistive robotics in home rehabilitation.

Keywords

Assistive Robotics, Home-based Rehabilitation, Telerehabilitation, Legal Ontology, Compliance-by-design

1. Introduction

Legal ontologies formalize legal concepts, norms, and reasoning patterns for computational systems. In eldercare, rehabilitation, and assistive-care settings, which are subject to health, privacy, safety, and consumer protection regulations, ontologies are particularly useful because these systems are characterized by complex interactions. Legal, ethical, and regulatory issues have arisen as robotics has become a part of daily life. One of the most sensitive and pressing applications is assistive robotics for elderly care. These systems, from mobile service robots to cognitive companions, promote autonomy, safety, and well-being for aging/vulnerable populations [1]. However, legal, ethical, and technical requirements complicate their deployment, so developing an ontology for elderly assistive robotics is pragmatic and strategic.

Robotic ontologies can be used in industry to improve the functionality and reliability of mechatronic systems [2], for collaborative manufacturing [3] and task/motion planning [4], in medicine for surgical workflow support, hospital service, and triage [5] and in domestic settings. In care-focused domains robotic ontologies are used to support ambient-assisted living [6], service interaction, and well-being monitoring [5, 7]. Most assistive robotics ontologies represent tasks, environments, and behavior knowledge. These ontologies model physical, functional, and contextual data but rarely include explicit representations of legal obligations, ethical principles, and accountability constraints. The IEEE 7007-2021 standard [8] aims to formalize ethical robotics ontologies, addressing data privacy, transparency,

RobOntics 2025: Workshop on Ontologies in Autonomous Robotics, December 8, 2025, Naples, Italy

*Corresponding author.

†These authors contributed equally.

✉ alberto.borboni@unibs.it (A. Borboni); giorgio.pedrazzi@unibs.it (G. Pedrazzi); fabio.zanoletti@unibs.it (F. Zanoletti)

ORCID 0000-0001-7069-1095 (A. Borboni); 0000-0002-1634-9966 (G. Pedrazzi); 0009-0006-6192-5746 (F. Zanoletti)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

and accountability. While some ontologies and conceptual frameworks integrate legal aspects into robotics and assistance, real-world applications are rare, and most work is still theoretical. In assistive robotics, sensitive data (health, behavior, location) is continuously processed, multi-actor systems (robot, caregiver, platform) must comply with privacy-by-design principles, and cross-border deployment requires interoperable legal vocabularies like DPV. Recent research emphasizes the growing need for evaluation methods and ontology development in the field to ensure consistent, transparent, and machine-actionable representations of legal requirements across diverse robotic platforms and regulatory contexts.

Legal ontologies organize legal knowledge into machine-readable frameworks and are increasingly used to support the formalization, comparison, and operationalization of legal regulations especially in areas that require automated compliance, reasoning, and explainability. Data protection and privacy, where regulatory diversity (e.g., GDPR, LGPD, LOPDP) and technical integration are major challenges, is a popular legal ontology domain. Several privacy and data protection ontologies have been proposed. OntoPriv [9] models Ecuador’s Organic Law on Personal Data Protection (LOPDP) while maintaining alignment with the GDPR and the W3C Data Privacy Vocabulary. PrOnto [10] has been adopted in projects such as Cloud4EU to support GDPR compliance in cloud services by representing legal roles, deontic relations, legal bases, and data processing workflows. Complementary approaches, including GConsent and Geko [9], focus more narrowly on modeling GDPR consent mechanisms and their relationship with information security controls.

This paper proposes a modular, ontology-based legal decision framework for assistive robotics in home-based care, focusing on elderly care as the deployment context and rehabilitation and assistance services as representative application scenarios. This approach can inform regulatory compliance-by-design and machine-interpretable reasoning in robotic systems, laying the groundwork for regulatory sandboxes, explainability modules, and autonomous decision-support mechanisms. The proposed architecture aligns with GDPR, MDR, ISO 13482:2014, and WHO frameworks on aging and assistive technologies, focusing the deployment context on elderly care and rehabilitation. This also integrates relevant technical standards like personal care robot safety, HRI protocols, and data protection schemes.

2. Materials and Methods

The design of the proposed ontology and reasoning architecture is motivated by the need to operationalize legal, regulatory, and organizational norms as first-class constraints in assistive robotics for home-based care. In this domain, compliance requirements are not static documentation artifacts but dynamic conditions that directly affect whether a robotic plan or action can be executed at runtime. To overcome this limitation, a compliance-by-design perspective is adopted, in which normative requirements actively gate planning, execution, and monitoring decisions.

At the core of the proposed framework, reported in Figure 1, is an ontology for home-based rehabilitation, covering physiotherapy and occupational therapy performed at home, remote clinical assessment, asynchronous monitoring, and caregiver-mediated interactions. The ontology includes six modules: *Actors and Roles* (patient, caregiver, tele-physiotherapist, device supplier, provider), *Activities and Plans* (assessment, exercise, adherence check, escalation), *Data and Consent* (scope, duration, purposes, retention, subject requests), *Safety and Risk* (hazards, controls, residual risk), *Security/Cybersecurity* (access, logging, audit), and *Transparency and Accountability* (explanations, provenance, responsibility). Conceptual commitments for norms, plans, and accountability follow the IEEE 7007 ERAS domain, reusing its top-level distinctions among situations, plans, actions, communications, and accountability relations (NEP, TA, DPP, EVM), which serve as alignment targets for our classes and properties.

Compliance is modeled as a precedence-aware constraint lattice:

- Binding law (e.g., data protection, medical device).
- Regulatory guidance (e.g., authority decisions, national telemedicine rules).
- Mandatory standards (by reference or conformity).
- Recognized guidelines (clinical or organizational).

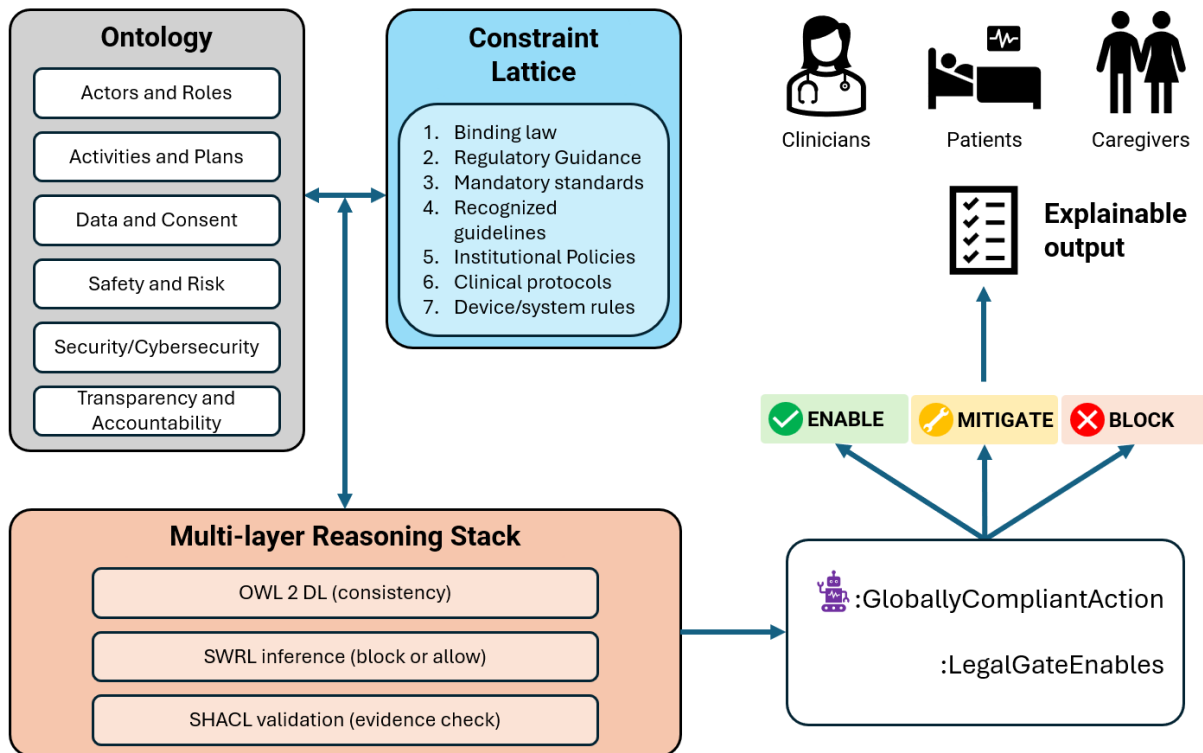


Figure 1: Schematic workflow for the normative-aware ontology for post-stroke telerehabilitation. The ontology integrates six modular domains with a precedence-aware constraint lattice representing binding laws, standards, and organizational rules. Compliance reasoning is performed through a multi-layer reasoning stack combining OWL 2 DL (consistency), SWRL inference (permission and prohibition), and SHACL validation (evidence checks). The two execution primitives translate normative constraints into runtime behavior, enabling, blocking, or mitigating robotic actions and producing traceable, audience-appropriate explanations for clinicians, patients, and caregivers.

- Institutional policies.
- Clinical protocols (pathway/condition specific).
- Device/system rules (e.g., runtime safety, UI constraints).

Each constraint instance is typed (e.g., *LegalConstraint*, *StandardConstraint*) with a priority level and scope (*session*, *dataset*, *device*). The lattice uses property overrides and preconditions with an irreflexive, acyclic ordering in OWL 2 DL. Higher-order constraints override or gate lower ones. A plan is enacted only if (i) higher-priority preconditions hold and (ii) no overriding constraint blocks it. A precedence-aware constraint lattice was selected to explicitly represent the hierarchical nature of normative sources in healthcare, where binding law must override standards, guidelines, and internal protocols in case of conflict. This design choice enables transparent conflict resolution and prevents lower-priority clinical or technical rules from defeating higher-order legal obligations.

Conflicts are handled through a multi-layer reasoning stack to balance conceptual consistency, explainable normative inference, and closed-world validation of operational evidence:

- *OWL 2 DL* is used to ensure logical consistency and ontological coherence across the domain and normative models. This layer captures the structural relations among actors, actions, plans, and constraints, and prevents unsatisfiable configurations, such as actions being simultaneously classified as mandatory and prohibited. OWL reasoning can provide a stable, explainable foundation for higher-level compliance decisions, ensuring that normative knowledge remains logically well-formed.
- *SWRL rules* provide lightweight inferences (e.g., prohibition if consent is missing; risk control for active home devices). These rules are used to derive permissions, prohibitions, and conditional

requirements based on runtime context, such as the presence or absence of valid consent, the operational mode of the session, or the risk class of a device. SWRL was selected to maintain transparency and traceability of inferred outcomes, enabling explicit links between triggering conditions and the resulting compliance status without resorting to procedural or black-box decision logic.

- *SHACL shapes* validate closed-world aspects (e.g., consent duration, identity proofing), returning severity levels (violation/warning) used as decision evidence. This layer checks the presence, completeness, and validity of required evidence, such as consent duration, identity proofing artifacts, or documented risk control measures. SHACL validation results are associated with severity levels (e.g., violation or warning) and are used as decision evidence to gate or condition action execution. By separating evidentiary validation from conceptual and inferential reasoning, the stack could support runtime compliance while preserving explainability and auditability.

Each automated decision (*enable/mitigate/block*) includes a justification graph linking the plan to enabling or defeating constraints, and provenance of data and policies. Following IEEE 7007, explanations are artifacts “provided for” an audience, addressing transparency concerns, with agents “accountable for” disclosed content to ensure traceable, audience-appropriate explanations for home rehabilitation.

The implementation will produce: (i) an OWL 2 DL ontology that defines domain concepts, relations and compliance constraints; (ii) SHACL shapes to validate data and configurations against constraints; and (iii) SWRL rules that derive policy-based inferences and permissions. The package supports three modes (in-person, telemedicine, caregiver-only) and a legal gate *:LegalGateEnabled* computed by SHACL when strong constraints hold and the class *:GloballyCompliantAction* identifies actions with valid legal and technical outcomes.

3. Results

We initially instantiated the ontology for Italy regulation framework: EU GDPR (Reg. 2016/679), Italian Data Protection Code, National Data Protection Authority’s guidance on health data and telemedicine consent, EU MDR (Reg. 2017/745) for home-use medical devices, and national telemedicine frameworks (State–Regions agreements, AGENAS guidance). These are modeled as *LegalConstraint* and *AuthorityGuidance* with priority above *StandardConstraint*. We considered the following technical standards applied in Italy: ISO 13485 - QMS for MD, ISO 14971 - Risk Management for MD, IEC 62304 - MDSW Lifecycle, IEC 62366 - Usability for MD, IEC 60601 - Medical Electrical Equipment safety, ISO/IEC 27001 - ISMS, and HL7 FHIR profiles for Healthcare Data exchange standards. They are represented as *StandardConstraint* applied to certain device categories, software risk classes, or interface types.

A post-stroke telerehabilitation pathway with caregiver assistance was encoded:

- (1) Referral and enrollment.
- (2) Remote baseline assessment (range of motion, functional tests) with identity proofing.
- (3) Weekly tele-supervised exercises using a CE-marked device and mobile app.
- (4) Asynchronous homework with adherence checks.
- (5) Continuous monitoring (diary, telemetry).
- (6) Progression and escalation (e.g., pain triage → clinician call).

Three operational modes, *in-person/telemedicine/caregiver-only*, are modeled as mutually exclusive, pre-gated by *:LegalGateEnabled* and ranked by risk, with home use requiring specific controls. Compliance reasoning is implemented as a structured pipeline in which different reasoning paradigms are applied in a controlled and defined order. For each session or planned action, the ontology is instantiated with the relevant activities, actors, devices, constraints, and available evidence. An OWL 2 DL reasoner is first employed to guarantee logical consistency and to perform monotonic classification under the open-world assumption, ensuring that incompatible normative states cannot co-exist while avoiding violation inferences due to missing information. Subsequently, SHACL validation is applied to enforce

closed-world requirements that depend on the explicit availability and validity of evidence, such as consent scope and duration, risk-control documentation, usability evidence, or identity proofing. SHACL outcomes are materialized as decision evidence and determine whether the legal gate holds. Finally, SWRL rules are used to materialize policy effects by deriving prohibitions, permissions, or mitigations without procedural code. An action is classified as *:GloballyCompliantAction* only when no higher-priority constraint defeats it and both legal and technical conditions are satisfied. The pipeline is re-evaluated whenever relevant data or evidence changes, supporting precedence-aware, traceable, and runtime-aligned compliance decisions.

The following illustrative examples are used to report how the proposed reasoning pipeline operates in representative scenarios:

- *Consent-based legal gating* (GDPR). A caregiver-only session is classified in OWL as a *Video-RequiredSession*. SHACL validation detects that no *VideoConsent* evidence is available for the required processing purpose. This validation outcome is materialized as decision evidence and consumed by SWRL rules, which derive a violation of the *GDPRVideoRequirement LegalConstraint*, ranked above clinical workflow rules in the constraint lattice. As a result, the session is inferred as a *ProhibitedAction*. An IEEE 7007-aligned explanation documents that the caregiver-session protocol was defeated by a binding legal constraint due to an insufficient consent scope.
- *Risk-based action gating* (ISO 14971). A balance and weight-shifting exercise is classified in OWL as involving an Class IIa home-use medical device and is therefore associated with applicable MDR risk-management constraints. SHACL validation checks for mandatory risk-control evidence (e.g., caregiver presence, environment setup, fall-prevention briefing) and initially reports a violation due to missing documentation. SWRL rules translate this outcome into a blocked action state. Once the required evidence is provided and SHACL validation succeeds, SWRL rules materialize a positive technical outcome, and the exercise is classified as a *GloballyCompliantAction*, enabling session execution.
- *Usability-driven mitigation* (IEC 62366). A caregiver-operated task is classified in OWL as subject to usability engineering requirements. SHACL validation verifies the presence of usability evidence specific to the Italian context (instructions, readability assessment, competency records). When such evidence is incomplete, SHACL produces a warning rather than a violation. SWRL rules consume this outcome and infer the action as *PermittedWithMitigation*, triggering compensatory measures (e.g., guided tutorials) while preserving traceability of the justification.
- *Transparency-driven explanation and accountability* (IEEE 7007/7001). A tele-supervised session is interrupted due to insufficient network quality. OWL reasoning classifies the event under safety-relevant conditions. SHACL validation confirms the availability of network logs and threshold parameters used in the decision. SWRL rules apply precedence relations (safety over continuity) to justify automatic rescheduling. An IEEE 7007-aligned explanation artifact is generated, explicitly linking the rescheduling decision to input data, applied thresholds, and normative precedence, and is addressed to patients and caregivers.

Beyond action-level gating, the instantiated ontology captures core data protection and security requirements relevant to Italian home-based rehabilitation, illustrating the scope of normative aspects that can be represented and reasoned over. Each dataset is linked to a clearly defined processing purpose, such as rehabilitation or quality improvement, ensuring that data are used only for their intended goals. Data retention is limited in time, with automated SHACL checks verifying that active policies on storage and deletion are respected. Requests from data subjects, such as access, export, or anonymization, are treated as actionable events and are documented with supporting evidence like export logs or anonymization receipts. Security measures enforce role-based access, apply the principle of data minimization, and maintain detailed audit logs. These logs are treated as key elements within the explanation framework, allowing every access or automated decision to be transparently justified in accordance with the local Data Protection Authority expectations.

The instantiation also enables validation and inspection of compliance states through standard reasoning and query mechanisms. Validation occurs at two levels. Logical: a DL reasoner ensures satis-

fiability and computes *:GloballyCompliantAction* as the conjunction of legal and technical compliance. Closed-world: SHACL asserts required fields and gates (e.g., consent period, ID proofing, risk evidence).

In the Italian home rehabilitation context, the ontology supports traceable compliance from plans to justifications, reduces normative contradictions through explicit precedence relations, and clarifies accountability via explanation artifacts. Limitations include the effort to maintain updated guidance, jurisdictional variations, and balancing open-world reasoning with closed-world validation. Still, combining an IEEE 7007-aligned conceptual core with OWL/SHACL/SWRL implementation and explicit legal gating (*:LegalGateEnabled*) yields a robust, auditable foundation for safe, lawful home rehabilitation in Italy.

4. Discussion

This work proposed an ontology-centered normative decision architecture that bridges legal accountability and technical autonomy in assistive robotics for home-based rehabilitation and elderly care. By combining a precedence-aware constraint lattice with a multi-layer reasoning stack, the proposed approach could operationalizes legal and organizational requirements as first-class constraints on robotic plans and actions. The Italian instantiation and the post-stroke telerehabilitation pathway show that concepts such as consent scope, safety risk controls, and role-based access could be enforced at runtime while preserving traceability and explainability. In practice, the gate *:LegalGateEnabled* and the classification *:GloballyCompliantAction* translates abstract compliance into executable conditions that either enable or defeat actions and generate audience-appropriate explanations for clinicians, patients, and caregivers (see the schematic in Figure 1).

Relative to existing ontologies for service or assistive robotics, which largely emphasize environment, task, and object knowledge, our contribution centers on integrating normative structures and evidence requirements into the same operational loop as planning and monitoring [11]. Moreover, whereas privacy and legal ontologies typically address documentation or design-time checks, here the legal layer is intertwined with scheduling, tele-supervision, and escalation decisions in a concrete clinical workflow. The modularization into *Actors and Roles*, *Activities and Plans*, *Data and Consent*, *Safety and Risk*, *Security/Cybersecurity*, and *Transparency and Accountability* facilitates alignment to recognized standards and guidance, and supports the reuse of submodules across conditions, devices, and service models [12]. The adoption of IEEE 7007 concepts helps standardize explanations and accountability attributions so that every allow/deny decision can be justified with a justification graph and provenance records, strengthening trust and auditability.

The illustrative examples highlight how the proposed architecture would affect operational decisions in representative home-based rehabilitation scenarios. Legal precedence on consent prevents permissive clinical protocols from overriding binding law, while ISO 14971-driven risk controls are operationalized as evidence-based gates rather than static documentation. Usability requirements condition action enablement and support automated mitigations (such as guided tutorials [13]) for caregiver-operated tasks, and transparency artifacts show how safety-first decisions can be communicated in a way that preserves accountability and trust. Together, these patterns show how compliance-by-design can be rendered operational without collapsing into procedural code or ad hoc rules.

Several limitations deserve attention. Maintaining up-to-date authority guidance, institutional policies, and jurisdiction-specific nuances introduces curatorial overhead and potential configuration drift. Balancing open-world ontological reasoning, which avoids inferring violations from missing information, with closed-world operational validation, which requires explicit evidence to enable actions, remains challenging, particularly when evidence is incomplete or delayed. Performance considerations arise as ontologies grow and SHACL/SWRL checks scale across many concurrent sessions, devices, and data streams. Additionally, while the Italian profile can provide a concrete test bed, broader generalization will require layered profiles for other jurisdictions and care settings, as well as empirical evaluation of user burden and workflow fit in clinical practice. A further limitation is the absence of empirical deployment or user-level evaluation: the proposed framework is demonstrated through

conceptual instantiation and illustrative scenarios, while real-world testing in clinical settings remains future work.

Future work could also include multi-jurisdictional profiles with automated change tracking for regulatory updates, integration with interoperable healthcare data vocabularies for robust evidence exchange, and prospective evaluations in clinical pilots measuring safety events, plan defeaters, and explanation utility. Methodologically, extending defeasible reasoning and preference handling, enriching provenance models for decision-critical data, and tooling for model governance (versioning, validation pipelines, regression tests for SHACL/SWRL) will be key. Ultimately, the goal is a portable, auditable substrate that allows assistive robots to adapt plans to human state and legal context while offering explanations that regulators, clinicians, patients, and caregivers can understand and trust.

5. Conclusions

This work introduced a normative-aware ontology that embeds compliance-by-design into assistive robotics for home rehabilitation, coupling a precedence-aware constraint lattice with a hybrid reasoning stack (OWL 2 DL for consistency, SWRL for lightweight derivations, and SHACL for closed-world validation). Instantiated to the Italian regulatory context and exercised on a post-stroke telerehabilitation pathway, the approach operationalizes consent scope, risk controls, security, and transparency as first-class constraints on plans and actions. Two execution primitives, *:LegalGateEnabled* and *:GloballyCompliantAction*, bridge legal and technical requirements to runtime behavior, producing enable/defeat decisions with IEEE 7007-aligned, audience-appropriate explanations. Preliminary results indicate traceable compliance, fewer contradictions through explicit precedence, and clearer accountability via justification graphs and provenance.

Limitations include curation effort for evolving guidance, jurisdictional tailoring, and scalability of SHACL/SWRL at larger operational scales. Future work will extend the ontology with real-world testing and multi-jurisdictional profiles, will strengthen interoperable evidence exchange with clinical systems, and develop governance tooling (versioning, validation pipelines, and regression tests) to sustain dependable deployments. Overall, the proposed framework could offer a portable, auditable substrate for safe, explainable, and norm-adaptive assistive robotics in home-based care.

Acknowledgments

Giorgio Pedrazzi is a member of the University of Insubria's PRIN 2022 Research Project of National Relevance titled "Demographic and legal changes. Towards an elder law" CUP J53D23005860006, funded by the European Union - Next Generation EU.

Alberto Borboni and Fabio Zanoletti acknowledge the support of the Italian Ministry of Enterprises and Made in Italy (Ministero delle Imprese e del Made in Italy) under the funding program No. F/310250/02/X56 - CUP: B89J23002550005 - COR: 16076042.

Declaration on Generative AI

The author(s) have not employed any Generative AI tools.

References

- [1] L. Ragno, A. Borboni, F. Vannetti, C. Amici, N. Cusano, Application of social robots in healthcare: Review on characteristics, requirements, technical solutions, *Sensors* 23 (2023). URL: <https://www.mdpi.com/1424-8220/23/15/6820>. doi:10.3390/s23156820.

- [2] F. Aggogeri, A. Borboni, R. Faglia, Reliability roadmap for mechatronic systems, *Applied Mechanics and Materials* 373-375 (2013) 130 – 133. URL: <https://www.scientific.net/AMM.373-375.130>. doi:10.4028/www.scientific.net/AMM.373-375.130.
- [3] R. Pagani, C. Nuzzi, M. Ghidelli, A. Borboni, M. Lancini, G. Legnani, Cobot user frame calibration: Evaluation and comparison between positioning repeatability performances achieved by traditional and vision-based methods, *Robotics* 10 (2021). URL: <https://www.mdpi.com/2218-6581/10/1/45>. doi:10.3390/robotics10010045.
- [4] A. Borboni, F. Aggogeri, N. Pellegrini, R. Faglia, Precision point design of a cam indexing mechanism, *Advanced Materials Research* 590 (2012) 399 – 404. URL: <https://www.scientific.net/AMR.590.399>. doi:10.4028/www.scientific.net/AMR.590.399.
- [5] S. Manzoor, Y. G. Rocha, S.-H. Joo, S.-H. Bae, E.-J. Kim, K.-J. Joo, T.-Y. Kuc, Ontology-Based Knowledge Representation in Robotic Systems: A Survey Oriented toward Applications, *Applied Sciences* 11 (2021) 4324. URL: <https://www.mdpi.com/2076-3417/11/10/4324>. doi:10.3390/app11104324, publisher: Multidisciplinary Digital Publishing Institute.
- [6] F. Zanoletti, A. Borboni, Descriptive ontology for home-based smart rehabilitation environments for motor and cognitive therapy in patients with neurological disorders, *Mechanisms and Machine Science* 190 (2025) 260 – 267. URL: https://link.springer.com/chapter/10.1007/978-3-032-02106-9_29. doi:10.1007/978-3-032-02106-9_29.
- [7] D. Paulius, Y. Sun, A Survey of Knowledge Representation in Service Robotics, *Robotics and Autonomous Systems* 118 (2019) 13–30. URL: <https://www.sciencedirect.com/science/article/pii/S0921889018303506>. doi:10.1016/j.robot.2019.03.005.
- [8] IEEE Ontological Standard for Ethically Driven Robotics and Automation Systems, *IEEE Std 7007-2021* (2021) 1–119. URL: <https://ieeexplore.ieee.org/document/9611206>. doi:10.1109/IEEESTD.2021.9611206.
- [9] G. Suntaxi, K. Ojeda, F. Rodríguez, P. Hierro, J. Miño, D. Flores, OntoPriv: Towards Designing a DPV-based Legal Ontology for Knowledge Retrieval and Compliance in Privacy Legislation, *CLEI Electronic Journal* 28 (2025). doi:10.19153/cleiej.28.5.4.
- [10] M. Palmirani, M. Martoni, A. Rossi, C. Bartolini, L. Robaldo, Legal Ontology for Modelling GDPR Concepts and Norms, in: *Legal Knowledge and Information Systems*, IOS Press, 2018, pp. 91–100. URL: <https://ebooks.iospress.nl/doi/10.3233/978-1-61499-935-5-91>. doi:10.3233/978-1-61499-935-5-91.
- [11] S. Y. Tan, A. Taeihagh, A. Tripathi, Tensions and antagonistic interactions of risks and ethics of using robotics and autonomous systems in long-term care, *Technological Forecasting and Social Change* 167 (2021). URL: <https://www.sciencedirect.com/science/article/pii/S0040162521001189?via%3Dihub>. doi:10.1016/j.techfore.2021.120686.
- [12] A. U. R. Butt, T. Mahmood, T. Saba, S. A. O. Bahaj, F. S. Alamri, M. W. Iqbal, A. U. Rehman, An optimized role-based access control using trust mechanism in e-health cloud environment, *IEEE Access* 11 (2023) 138813 – 138826. URL: <https://ieeexplore.ieee.org/document/10327710>. doi:10.1109/ACCESS.2023.3335984.
- [13] M. Germanotta, M. C. Mauro, F. Falchini, F. Scotto di Luzio, L. Vollero, L. Zollo, I. G. Aprile, A robotic rehabilitation intervention in a home setting during the covid-19 outbreak: a feasibility pilot study in patients with stroke, *Journal of NeuroEngineering and Rehabilitation* 22 (2025). URL: <https://jneuroengrehab.biomedcentral.com/articles/10.1186/s12984-025-01633-9>. doi:10.1186/s12984-025-01633-9.