ESANN 2026 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium) and online event, 22-24 April 2026, i6doc.com publ., ISBN 9782875870964. Available from http://www.i6doc.com/en/.

Neuro Symbolic AI and Complex Data

Luca Oneto[1], Nicolò Navarin[2], Luca Pasa[2], Davide Rigoni[2], Davide Anguita[1] *

1 - University of Genoa - Via Dodecaneso 35, 16146, Genova - Italy
2 - University of Padua - Via Trieste 63, 35121, Padova - Italy

Abstract. The widespread use of Artificial Intelligence (AI) for decision-making on complex data - such as images, text, graphs, and nonlinear systems - has enabled significant progress across many application domains. However, purely data-driven approaches often struggle to provide structured reasoning, interpretability, robustness, and effective integration of domain knowledge. Neuro-Symbolic AI addresses these challenges by combining sub-symbolic learning with symbolic reasoning, allowing models to incorporate logical rules, constraints, ontologies, and expert knowledge. This tutorial offers a concise overview of Neuro-Symbolic AI for complex data, presenting core concepts, representative methods, and key application areas, including constraint-aware learning, AI for science, socially responsible AI, and knowledge-guided inference. It discusses both the opportunities and limitations of current approaches, highlighting how neuro-symbolic integration can improve transparency, trustworthiness, and human-centric alignment in AI systems. This tutorial aims to foster cross-disciplinary understanding and support the development of robust and explainable AI solutions for real-world problems.

1 Introduction

Deep learning has reshaped the practice of AI by enabling end-to-end learning from high-dimensional observations such as pixels and text tokens. Yet, modern deployments increasingly demand properties that are not naturally guaranteed by purely data-driven systems: (i) *structured reasoning* over entities and relations, (ii) *interpretability* that supports debugging and accountability, (iii) *robustness* under distribution shift and noisy labels, and (iv) *integration of domain knowledge* expressed as rules, constraints, and ontologies. These requirements are common across complex-data domains: medical imaging and clinical text, safety-critical perception for robotics, scientific discovery over graphs and simulations, and policy-constrained decision-making.

Neuro-Symbolic AI, or NeSy, aims to unify the strengths of sub-symbolic models — such as representation learning and differentiable optimization — with those of symbolic AI, including expressive knowledge representation, logic-based inference, and the enforcement of explicit constraints. This integration has

*This work is partially supported by (i) project ELSA – European Lighthouse on Secure and Safe AI funded by the European Union's Horizon Europe under the grant agreement No. 101070617, (ii) project SERICS (PE00000014) and FAIR (PE00000013) under the NRRP MUR program funded by the EU - NGEU, (iii) project FISA-2023-00128 funded by the MUR program "Fondo italiano per le scienze applicate", and (iiii) project DEEP-GRAPH: Design and Theory of Deep Graph Learning funded by the Italian Ministry of University and Research under the PRIN program - PRIN 2022 – CUP B53C24006570006

ESANN 2026 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium) and online event, 22-24 April 2026, i6doc.com publ., ISBN 9782875870964. Available from http://www.i6doc.com/en/.

been discussed as a "third wave" of AI systems that is motivated by trust and accountability concerns [1], and it is tightly connected to statistical relational learning (SRL) traditions [2]. This tutorial is centered on NeSy methods that *explicitly* incorporate symbolic structures — logical rules, ontological axioms, programs, or constraint satisfaction components — to improve learning and inference on complex data.

This tutorial is organized around: (i) core representations for knowledge and constraints; (ii) a taxonomy of integration patterns; (iii) representative methods with pointers to canonical systems; and (iv) applications in constraint-aware learning, knowledge-guided inference over graphs, AI for science, and socially responsible AI. The article concludes by discussing the limitations and open challenges in building robust and explainable NeSy systems.

2 Complex Data and the Limits of Purely Neural Systems

Complex data refers to modalities and structures where meaning emerges from *relations* and *composition*: objects and their interactions in images, long-range dependencies in text, multi-hop relations in knowledge graphs, or physical constraints in dynamical systems. While neural models can approximate complicated functions, recurring issues arise.

Data efficiency and compositional generalization. When tasks demand systematic recombination of learned concepts, e.g., relational reasoning and multi-step deduction, purely neural approaches often require large datasets and may generalize poorly out of distribution. NeSy methods can inject compositional priors via rules or programs, supporting more systematic behavior like differentiable rule learning in ILP-style systems [3].

Constraints, safety, and accountability. In many domains, correctness is not only empirical but also normative: predictions should satisfy hard or soft constraints, including logical consistency, adherence to physical laws, and compliance with fairness criteria. Expressing such constraints in symbolic form and enforcing them during learning is a direct route to more controllable systems [4, 5].

Reasoning with explicit knowledge. Knowledge graphs, ontologies, and rule bases capture structured facts and relations; leveraging them for inference can improve transparency and enable auditing. Differentiable reasoning over symbolic structures, like Neural-LP [6], TensorLog [7], and differentiable proving [8], provides one family of bridges between learning and logic.

3 Symbolic Knowledge: What Is Being Integrated?

NeSy methods differ substantially in the symbolic artifacts they take as input and generate as output. Depending on the task and modeling assumptions, they may consume hand-crafted logical theories, ontologies, or constraint sets,

ESANN 2026 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium) and online event, 22-24 April 2026, i6doc.com publ., ISBN 9782875870964. Available from http://www.i6doc.com/en/.

or instead rely on more lightweight templates that are refined during learning. Symmetrically, their outputs can range from enriched knowledge bases and induced rule sets to probabilistic programs or structured constraints that can be fed into downstream reasoning or optimization pipelines. This variation in consumed and produced artifacts shapes how prior knowledge is injected, how interpretability is realized, and how the resulting systems can be verified and integrated into larger workflows. Follow a more in-depth description of them.

First-order logic and rule sets. Rules can encode type constraints, relational implications, and domain heuristics. In practice, rules are often treated as *soft* constraints, weighted or fuzzy, to accommodate noise and exceptions [4, 9].

Ontologies and description logics. Ontologies formalize taxonomies and relations, e.g., `is-a` and `part-of`, supporting consistency checking and inheritance-based reasoning. NeSy systems may map neural predictions to ontology-aligned labels or enforce ontological constraints in training.

Probabilistic logics and SRL formalisms. Probabilistic extensions of logic support uncertain facts and weighted rules, such as Markov Logic Networks [10] or ProbLog [11]. These frameworks are natural homes for integrating learned perceptual modules.

Constraints and optimization problems. Many applications can be expressed as constraint satisfaction or structured prediction. Here, symbolic knowledge appears as constraint sets, sometimes solved by external solvers or relaxed into differentiable objectives (e.g., semantic loss [5]).

4 A Taxonomy of Neuro-Symbolic Integration Patterns

A practical way to characterize NeSy systems is to analyze *where* and *how* symbolic reasoning is embedded within the overall computational pipeline. Different approaches inject symbolic structure at distinct stages: as constraints during training, as a post-hoc guide, as differentiable reasoning modules tightly coupled with neural components, or as logic programs orchestrating neural predicates. This perspective highlights integration patterns rather than specific models, allowing the collection of methods that share similar points of contact between sub-symbolic learning and symbolic inference. In the following, several recurring patterns that exemplify these design choices are outlined.

Constraints as Regularization A common pattern is to train a neural model while penalizing violations of logic constraints:

$$\min_{\theta} \mathcal{L}_{\text{task}}(\theta) + \lambda \sum_{k} \phi_k(\theta), \qquad (1)$$

ESANN 2026 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium) and online event, 22-24 April 2026, i6doc.com publ., ISBN 9782875870964. Available from http://www.i6doc.com/en/.

where ϕ_k quantifies the degree to which predictions violate constraint k. Semantic-Based Regularization (SBR) operationalizes this idea with fuzzy logic constraints and multi-objective learning [4]. Semantic loss derives a principled penalty from a Boolean formula by measuring (log-)probability mass assigned to satisfying assignments [5]. Logic Tensor Networks (LTNs) use differentiable many-valued logic to integrate learning and reasoning, with applications such as semantic image interpretation [9, 12].

Teacher-Student and Rule Distillation Another approach is to construct a "teacher" that enforces rules and distill its behavior into a "student" neural network, transferring rule knowledge through training targets. A canonical example is the iterative distillation framework of [13], which injects first-order logic rules into neural models for NLP tasks.

Differentiable Reasoning Modules Here, symbolic inference is implemented (or approximated) in a differentiable way, enabling end-to-end training. Neural theorem proving replaces symbolic unification with differentiable similarity in embedding spaces and trains proof procedures by gradient descent [8]. Neural logic programs such as Neural-LP learn to compose differentiable operators corresponding to rule application for knowledge-base reasoning [6], building on differentiable deductive databases like TensorLog [7].

Neural-Probabilistic Logic Programming Probabilistic logic programming provides a clean interface: deep networks act as *neural predicates* whose outputs become probabilistic facts, while a logic program handles compositional reasoning. DeepProbLog is a widely used representative, blending deep learning with ProbLog-style inference [14].

Program Induction and Neuro-Symbolic Perception-Reasoning In multimodal domains, a neural perception module builds an intermediate symbolic representation (objects, attributes), and a symbolic program executes reasoning. The Neuro-Symbolic Concept Learner (NS-CL) learns object-centric representations and maps language to executable programs for visual question answering [15].

Interpretable Logic-Parameterized Networks Logical Neural Networks (LNNs) assign semantic meaning to neurons via weighted real-valued logic, enabling omni-directional inference and constraint handling while remaining trainable [16]. Related ideas include differentiable ILP and neural logic architectures that learn rules and generalize systematically [3, 17].

5 Representative Methods for Complex Data

This section connects the above patterns to concrete complex-data settings.

ESANN 2026 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium) and online event, 22-24 April 2026, i6doc.com publ., ISBN 9782875870964. Available from http://www.i6doc.com/en/.

Images and Structured Vision Computer vision models excel at recognition but typically lack explicit relational reasoning. NeSy methods address this by introducing structured representations (objects, relations) and constraints (e.g., `part-of`, mutual exclusivity, taxonomy consistency). LTNs have been applied to semantic image interpretation where background knowledge improves robustness under noisy labels [9, 12]. NS-CL shows another path: neural perception yields object-centric "slots"; language is parsed into a symbolic program executed over the scene [15]. A recurring theme is that intermediate symbolic structure supports both generalization and explanation: answers can be traced back to objects, relations, and executed steps.

Text, Rules, and Weak Supervision NLP often benefits from expert heuristics (patterns, gazetteers, grammar constraints) that are hard to learn reliably from limited labels. Rule injection can improve label efficiency and calibrate behavior. Hu et al [13] demonstrate logic-rule distillation for tasks such as sentiment analysis and named entity recognition, using rules to constrain predictions without requiring hard-coded post-processing. Constraint-based learning can also enforce structural properties (e.g., BIO tagging consistency, entailment constraints) by adding penalties such as semantic loss [5] or related constraint objectives.

Graphs and Knowledge-Guided Inference Graphs are a natural substrate for NeSy: they expose entities and relations explicitly, and symbolic knowledge can specify multi-hop rules and constraints. SRL formalisms like MLNs [10], ProbLog [11], and PSL/HL-MRFs [18, 19] provide weighted logical templates for collective inference. On the neural side, differentiable rule-learning methods such as Neural-LP [6] and differentiable proving [8] learn to reason over incomplete knowledge bases, often yielding human-readable rule patterns. DeepProbLog couples such symbolic reasoning with perception modules (e.g., a CNN recognizing digits) to solve tasks requiring both recognition and compositional reasoning [14].

Nonlinear Systems, Physics, and Scientific Constraints Many scientific and engineering problems are governed by known equations, conservation laws, and invariants. Even when full symbolic models are unavailable, partial constraints can regularize learning. Physics-Informed Neural Networks (PINNs) encode differential equation residuals as training penalties, improving data efficiency and physical plausibility [20]. While PINNs are not always framed as "symbolic reasoning", they exemplify a broader NeSy principle: *learning under structured constraints derived from domain knowledge.* A promising direction is to combine symbolic rule/constraint layers with neural surrogates for simulators, enabling reasoning about mechanisms, interventions, and uncertainty.

6 Key Application Areas

This section connects the above patterns to concrete application areas.

ESANN 2026 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium) and online event, 22-24 April 2026, i6doc.com publ., ISBN 9782875870964. Available from http://www.i6doc.com/en/.

Constraint-Aware Learning and Robustness Constraint-aware training aims to prevent implausible outputs, such as impossible label combinations or violations of ontological axioms, and to reduce sensitivity to noisy supervision. SBR and semantic loss provide generic toolkits for incorporating constraints into learning objectives [4, 5]. A practical benefit is *test-time consistency*: some approaches enforce constraints not only during training but also during inference, improving global coherence of predictions.

AI for Science NeSy provides a framework for coupling data-driven discovery with explicit scientific knowledge. Recent advances in deep learning highlight the impact of strong inductive biases: for instance, AlphaFold's protein-structure predictions demonstrate how structured modeling assumptions and large-scale learning can reshape biology [21]. At the same time, emerging generative models for *de novo* drug design, which aim to generate novel molecular structures with targeted properties within a large chemical space, make the explicit integration of chemical rules, reaction knowledge, and mechanistic constraints a natural setting for NeSy approaches. Indeed, de novo molecular design requires exploring a vast chemical space while satisfying multiple, often competing, objectives such as potency, selectivity, and synthetic accessibility. Generative models such as variational autoencoders [22–25], normalizing flows [26], and diffusion models [27, 28], applied to molecular graphs or SMILES strings [29] — a line notation for describing molecular structures — can navigate this space and optimize differentiable property predictors. However, without explicit enforcement of chemical and pharmacological rules, they frequently generate invalid, non-synthesizable, or pharmacologically implausible compounds. NeSy approaches offer a principled way to integrate medicinal-chemistry rules, reaction templates, and chemical knowledge into these generative pipelines. Symbolic components can encode constraints on valence and stereochemistry, forbidden or preferred substructures, and retrosynthetic feasibility via reaction rules or planning systems [30], while neural components propose and refine candidate structures and predict properties. Such hybrids promise more data-efficient search, improved scientific plausibility, and greater transparency: generated hypotheses or molecular candidates can be accompanied by rule-based rationales or candidate synthetic routes, facilitating expert review and iterative design. NeSy methods more broadly can enable explicit reasoning over hypotheses, integrate curated knowledge bases, and produce explanations aligned with scientific concepts, such as rules over interaction graphs or mechanistic constraints.

Socially Responsible AI: Fairness and Accountability Many fairness and accountability criteria can be expressed as constraints: treating similar individuals similarly [31], or enforcing equality of opportunity across groups [32]. NeSy approaches support such requirements by (i) encoding normative constraints in a human-auditable form and (ii) enforcing them through training objectives or constrained inference. However, symbolic constraints do not automatically solve governance problems: selecting appropriate constraints, verifying them under distribution shift, and handling conflicts between constraints remain open challenges.

7 Opportunities, Limitations, and Open Challenges

Despite recent progress, NeSy remains a challenging and active research area. Current methods present a range of conceptual and practical limitations that hinder their reliability, scalability, and adoption in real-world pipelines. These limitations cut across algorithmic aspects (e.g., inference and optimization), knowledge engineering (e.g., acquisition and maintenance of rules and ontologies), and evaluation (e.g., how to measure benefits beyond accuracy). This section highlights some challenges and failure modes that need to be addressed for NeSy to deliver robust, trustworthy, and widely deployable systems.

Scalability and latency. Symbolic inference can be expensive, especially with large knowledge bases or complex constraints. Differentiable relaxations improve optimization but can trade off soundness and may be sensitive to hyperparameters.

Soundness versus flexibility. Strict logical soundness often conflicts with noisy real-world data. Soft logic and probabilistic rules help, but interpreting the semantics of "partially true" statements requires care, particularly when explanations are used for high-stakes decisions.

Knowledge acquisition and maintenance. NeSy needs knowledge: rules, ontologies, or constraints. Curating them is labor-intensive, and automatically learned rules can be brittle or reflect spurious correlations. Hybrid workflows (expert + induction) are promising but underdeveloped.

Evaluation and benchmarking. Comparing NeSy methods is difficult because benefits vary by domain: some methods improve accuracy, others improve interpretability, and others primarily improve constraint satisfaction. Better benchmarks that jointly measure performance, consistency, and explanation quality are needed.

Uncertainty and distribution shift. Many NeSy systems assume that constraints hold globally, but real environments change. Handling uncertain or context-dependent rules (and detecting when rules should be revised) is central to robust deployment.

8 Conclusion

Neuro-Symbolic AI (NeSy) provides a toolkit for building AI systems that can learn from complex data while leveraging explicit knowledge, rules, and constraints. Across images, text, graphs, and nonlinear systems, NeSy methods help bridge representation learning with structured reasoning, improving interpretability and enabling more controllable behavior. At the same time, fundamental challenges, including scalability, knowledge acquisition, and reliable evaluation, remain. A practical takeaway is to treat NeSy as a *design space*: choose the symbolic artifacts — such as rules, ontologies, and probabilistic programs — select an integration pattern like constraints, differentiable reasoning, and program induction, and evaluate against domain requirements for robustness, transparency, and human-centric alignment.

References

[1] A. S. D. Garcez and L. C. Lamb. Neurosymbolic AI: The 3rd wave. *arXiv preprint arXiv:2012.05876*, 2020.

[2] L. De Raedt, R. Manhaeve, S. Dumancic, T. Demeester, and A. Kimmig. From statistical relational to neuro-symbolic artificial intelligence. In *International Joint Conference on Artificial Intelligence*, 2020.

[3] R. Evans and E. Grefenstette. Learning explanatory rules from noisy data. *Journal of Artificial Intelligence Research*, 61:1–64, 2018.

[4] m. Diligenti, M. Gori, and C. Saccà. Semantic-based regularization for learning and inference. *Artificial Intelligence*, 244:143–165, 2017.

[5] J. Xu, Z. Zhang, T. Friedman, Y. Liang, and G. Van den Broeck. A semantic loss function for deep learning with symbolic knowledge. In *International Conference on Machine Learning*, 2018.

[6] F. Yang and W. W. Yang, Z.and Cohen. Differentiable learning of logical rules for knowledge base reasoning. In *Neural Information Processing Systems (NeurIPS)*, 2017.

[7] W. W. Cohen. TensorLog: A differentiable deductive database. *arXiv preprint arXiv:1605.06523*, 2016.

[8] T. Rocktäschel and S. Riedel. End-to-end differentiable proving. In *Neural Information Processing Systems (NeurIPS)*, 2017.

[9] L. Serafini, I. Donadello, and A. S. D. Garcez. Learning and reasoning in logic tensor networks: Theory and application to semantic image interpretation. In *Symposium on Applied Computing*, 2017.

[10] M. Richardson and P. Domingos. Markov logic networks. *Machine Learning*, 62(1-2):107–136, 2006.

[11] L. De Raedt, A. Kimmig, and H. Toivonen. ProbLog: A probabilistic Prolog and its application in link discovery. In *International Joint Conference on Artificial Intelligence*, 2007.

[12] I. Donadello, L. Serafini, and A. S. D. Garcez. Logic tensor networks for semantic image interpretation. In *International Joint Conference on Artificial Intelligence*, 2017.

[13] Z. Hu, X. Ma, Z. Liu, E. Hovy, and E. Xing. Harnessing deep neural networks with logic rules. In *Annual Meeting of the Association for Computational Linguistics*, 2016.

[14] R. Manhaeve, S. Dumancic, A. Kimmig, T. Demeester, and L. De Raedt. Neural probabilistic logic programming in DeepProbLog. *Artificial Intelligence*, 298:103504, 2021.

[15] J. Mao, C. Gan, P. Kohli, J. B. Tenenbaum, and J. Wu. The neuro-symbolic concept learner: Interpreting scenes, words, and sentences from natural supervision. In *International Conference on Learning Representations*, 2019.

[16] R. Riegel, A. Gray, F. Luus, N. Khan, and Others. Logical neural networks. *arXiv preprint arXiv:2006.13155*, 2020.

[17] H. Dong, J. Mao, T. Lin, C. Wang, L. Li, and D. Zhou. Neural logic machines. In *International Conference on Learning Representations*, 2019.

[18] S. H. Bach, M. Broecheler, B. Huang, and L. Getoor. Hinge-loss markov random fields and probabilistic soft logic. *Journal of Machine Learning Research*, 18(109):1–67, 2017.

[19] A. Kimmig, S. H. Bach, M. Broecheler, B. Huang, and L. Getoor. A short introduction to probabilistic soft logic. In *Neural Information Processing Systems Workshop on Probabilistic Programming: Foundations and Applications*, 2012.

[20] M. Raissi, P. Perdikaris, and G. E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.

[21] J. Jumper, R. Evans, A. Pritzel, T. Green, and Others. Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873):583–589, 2021.

[22] M. J. Kusner, B. Paige, and J. M. Hernández-Lobato. Grammar variational autoencoder. In *International conference on machine learning*, 2017.

[23] H. Dai, Y. Tian, B. Dai, S. Skiena, and L. Song. Syntax-directed variational autoencoder for structured data. In *International Conference on Learning Representations*, 2018.

[24] D. Rigoni, N. Navarin, and A. Sperduti. Conditional constrained graph variational autoencoders for molecule design. In *IEEE Symposium Series on Computational Intelligence*, 2020.

[25] D. Rigoni, N. Navarin, and A. Sperduti. Rgcvae: relational graph conditioned variational autoencoder for molecule design. *Machine Learning*, 114(2):47, 2025.

[26] D. Rigoni, S. Yaddehige, N. Bianchi, A. Sperduti, S. Moro, and C. Taccioli. Tumflow: An ai model for predicting new anticancer molecules. *International Journal of Molecular Sciences*, 25(11):6186, 2024.

[27] M. Madeira, C. Vignac, D. Thanou, and P. Frossard. Generative modelling of structurally constrained graphs. *Neural Information Processing Systems*, 2024.

[28] S. Cognolato, D. Rigoni, M. Ballarini, L. Serafini, S. Moro, A. Sperduti, and Others. D4: Distance diffusion for a truly equivariant molecular design. In *European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning*, 2025.

[29] D. Weininger. Smiles, a chemical language and information system. 1. introduction to methodology and encoding rules. *Journal of chemical information and computer sciences*, 28(1):31–36, 1988.

[30] Binghong Chen, Chengtao Li, Hanjun Dai, and Le Song. Retro*: Learning retrosynthetic planning with neural guided a* search. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 1608–1616. PMLR, 13–18 Jul 2020.

[31] C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. Zemel. Fairness through awareness. In *Innovations in Theoretical Computer Science Conference*, 2012.

[32] M. Hardt, E. Price, and N. Srebro. Equality of opportunity in supervised learning. In *Neural Information Processing Systems (NeurIPS)*, 2016.

ESANN 2026 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium) and online event, 22-24 April 2026, i6doc.com publ., ISBN 9782875870964. Available from http://www.i6doc.com/en/.