Foundations of Trustworthy AI – Integrating Reasoning, Learning and Optimization
TAILOR
Grant Agreement Number 952215

Foundations, Techniques, Algorithms and Tools for Integrating Learning, Reasoning and Optimisation v.2 Report

Document type (nature)	Report
Deliverable No	4.2
Work package number(s)	4
Date	July 2024
Responsible Beneficiary	P 5 - KU Leuven
Editor(s)	Luc De Raedt
Publicity level	Public
Short description	Foundations, techniques, algorithms and tools for integrating learning, reasoning and optimisation v.2

History			
Revision	**Date**	**Modification**	**Authors**
1.0	July 2024	Major	Luc De Raedt, Vincent Derkinderen, Francesco Giannini, Andrea Borghesi, Michele Lombardi

Document Review		
Reviewer	**Partner ID / Acronym**	**Date of report approval**
Nic Wilson	UCC	05 July 2024
Umberto Straccia	CNR	05 July 2024

Table of Contents

Introduction to the deliverable

The question addressed in this WP is how to integrate learning, reasoning and optimisation, that is, how to computationally and mathematically integrate different AI paradigms. The most apparent difference between paradigms lies in the representations that are used and so an operational way to answer the question is to tightly integrate different representations as to offer both learning, reasoning and/or optimisation in common frameworks. This theme will therefore design representational systems with accompanying inference, learning and optimisation algorithms that can support trustworthy artificial intelligence.

As part of this WP, we organised several meetings and workshops, which have led to insights into the foundations and challenges connected to this WP. It has also resulted in a number of interesting tutorials and survey papers, that have partly or fully been inspired by TAILOR and that have led to novel insights and collaborations. As part of our goal to integrate the different paradigms and gaining a broader understanding on the foundational challenges, we have also studied the vast landscape of relevant systems, tasks, benchmarks and applications to which those systems have been applied. Those results are available in deliverable 4.4, which is available at: https://tailor-network.eu/wp-content/uploads/2024/05/D-4.4-Integrated-learning-reasoning-and-optimisation-in-practice-v.2.pdf

In this deliverable, we instead specifically report on the foundational challenges that are associated with integrating learning, reasoning and optimisation, and our progress in tackling those challenges through TAILOR. This deliverable is structured according to the four tasks of this WP. First, we set the scene of this WP, explaining that within TAILOR we believe that AI cannot rely on a single paradigm if it needs to be trustworthy and hence, it needs the ability to both learn and reason. Then, we provide for each task a summarising exposition on our results achieved within TAILOR, and the current state-of-the-art.

Task 4.1 - Integrating representations for learning and reasoning
Task 4.2 - Integrating approaches to learning and optimisation
Task 4.3 - Learning and reasoning with embeddings, knowledge graphs & ontologies
Task 4.4 - Learning and reasoning for perception, spatial reasoning, and vision

This is the second version of this deliverable. The first version is available at: https://tailor-network.eu/wp-content/uploads/2022/07/D4.1-Foundations-techniques-algorithms-and-tools-for-integrating-learning-reasoning-and-optimisation-v.1.pdf

Organisation

The following people are responsible for this deliverable:

Partner ID / Acronym	Name	Role	Other
KU Leuven	Luc De Raedt	Leader of Task 4.1	Vincent Derkinderen, Robin Manhaeve
University of Bologna	Michela Milano	Leader of Task 4.2	Andrea Borghesi Michele Lombardi
University of Siena/CINI	Marco Gori	Leader of Task 4.3	Francesco Giannini
RWTH	Bastian Leibe	Leader of Task 4.4	-

Scientific Tasks

Setting the Scene: Integrating paradigms and communities

There is an ongoing debate in AI about the distinction between learning and reasoning systems, between data- and knowledge-driven approaches, symbolic versus subsymbolic techniques, and solvers and learners. They are often also related to Kahnemann's System 1 vs 2 distinction. While there are (often subtle) differences between these oppositional terms, they all point in the direction of a dichotomy between two types of tasks and techniques in AI, that we will refer to as learning and reasoning. For instance, the symbolic AI or the logic paradigm has concentrated on developing sophisticated and accountable reasoning methods, the subsymbolic or neural approaches to AI have concentrated on developing powerful architectures for learning and perception, and constraint and mathematical programming have been used for combinatorial optimisation. While deep learning provides solutions to many low-level perception tasks, it is unsuitable for complex reasoning tasks that require guarantees on correctness or optimality; while for logical and symbolic methods, it is just the other way around. Symbolic AI may be more explainable, interpretable and verifiable, but it is less flexible and adaptable. There is a growing belief in the AI community that the next wave in AI will bridge the gap between learning and reasoning. Indeed, researchers such as Marcus, Darwiche, Levesque, Tenenbaum, Geffner, Bengio, Le Cun, and Kautz, have all argued that the integration of learning and reasoning is the next challenge for AI [Bengio et al. 2021, Kautz 2020, Marcus 2020]. Although they all agree on the nature of the challenge, they often disagree about how to tackle this challenge. Researchers coming from the neural network community usually argue for incorporating reasoning inside neural networks, while researchers from more traditional areas in AI often want to extend symbolic representations with learning abilities and neural networks. Learning and reasoning are different paradigms studied by different communities. Within TAILOR we believe that AI cannot rely on a single paradigm if it needs to be trustworthy and hence, it needs the ability to both learn and reason. Therefore the quest for integrated learning and reasoning abilities boils down to computationally and mathematically integrating different AI paradigms. The most apparent difference between these paradigms lies in the representations that are used and so one operational way to answer the question is to tightly integrate different representations to offer learning, reasoning and optimisation in a common framework. WP 4 therefore **designs representational systems with accompanying inference, learning and optimisation algorithms that can support trustworthy artificial intelligence.** The integrated or "unified" representations should be able to address the whole AI cycle from low-level perception to high-level reasoning, they should be able to use data as well as knowledge, and most of all, should produce trustworthy AI. With respect to trustworthiness of representations, the most critical dimension is their explainability. The quest for integrated representations and paradigms in artificial intelligence is akin to systems biology in the sense that it aims at understanding AI by putting the pieces together, rather than focussing on the individual representations and building blocks. It thus constitutes a kind of systems AI.

The work on integrated representations and abilities has been fragmented into activities in many specialised subcommunities all with their specialised workshops:

- **Neuro-Symbolic Learning Artificial Intelligence (NeSy)** wants to bridge the gap between neural networks and logical and symbolic approaches to reasoning. Although many promising NeSy models and representations have been introduced, they are still limited in their reasoning and explanation abilities because they have either pushed the logic inside the neural network, or use logic constraints to regularise the network instead of for reasoning. **Statistical Relational AI (StarAI) e**xtends probabilistic graphical models with first order logic, and is in fact pursuing similar goals as probabilistic programming. Its selling point is that it integrates probabilistic and logical reasoning with statistical learning. Within Task 4.1 we study NeSy and StarAI approaches and especially focus on their further integration.
- **Embeddings and ontological reasoning** (Task 4.3). Embeddings are among the most powerful techniques in deep learning and they are routinely applied in numerous applications concerning natural language and knowledge graphs.
- **Combinatorial optimisation.** For what concerns **learning and optimisation** (Task 4.2), the constraint programming community has contributed frameworks, such as empirical model learning (EML), and smart predict and optimise, which learn constraint satisfaction and constraint programming models from the available data and then use these models for **combinatorial optimisation.**
- **Perception, spatial reasoning and vision**. Deep learning has contributed solutions to numerous traditional computer vision and perceptual tasks. But integrating the vision with reasoning is still an open problem as witnessed by the many challenging datasets on reasoning in a computer vision context. This leads to Task 4.4 on learning and reasoning **for perception, spatial reasoning and vision.**

Therefore, the goal of WP 4 is not only to integrate the different paradigms but also the underlying communities as depicted in the figure below.

We report on the four scientific tasks in the next four sections. Each task starts with a short review of the state of the art, typically based on a survey and/or tutorial provided by one of the TAILOR partners. *Because the work under Tasks 4.1 and 4.3 is closely related and many works contribute to both tasks, we describe our contributions to Tasks 4.1. and 4.3. together, followed by our contributions to Task 4.2 and Task 4.4*.

We list the workshops organised for this work package in the Appendices.

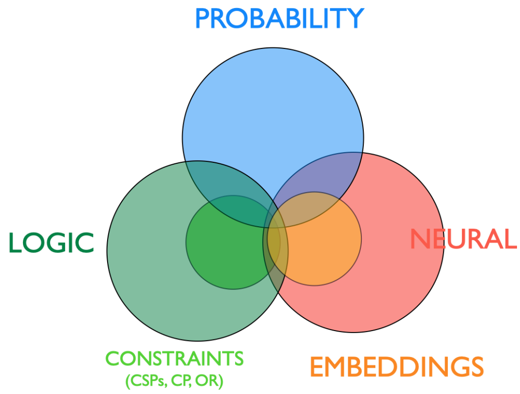

Task 4.1: Integrating Representations for Learning and Reasoning

State of the art. One view and survey on the state of the art for integrating representations for learning and reasoning is provided in detail by Marra et al. (2024) (with an earlier shorter version appearing in De Raedt et al. (2020). This survey was also the basis for tutorials at AAAI 2021, IJCAI 2021, the TAILOR Summer School in 2021, ESSAI 2023, DeepLearn 2023, NSSS 2023, and keynotes at e.g. ECAI 2020, CCAI 2021, ESWC 2021, KR 2023, etc. (see http://dtai.cs.kuleuven.be/tutorials/nesytutorial/). The survey bridges the gap within neuro-symbolic AI (NeSy), where the goal is to incorporate symbolic reasoning into neural networks with statistical relational AI (StarAI). While NeSy already has a long tradition, it has attracted a lot of attention in the last few years from various communities.

But also the field statistical relational learning and artificial intelligence (StarAI) has a long tradition in integrating learning and reasoning. Rather than focusing on how to integrate logic and neural networks, StarAI is centred around the question of how to integrate logic with probabilistic graphical models. Despite the common interest in combining logic or symbolic reasoning with a basic paradigm for learning, i.e., probabilistic graphical models or neural networks, it is surprising that there are not more interactions between these two fields. This discrepancy motivated De Raedt et al. (2020), and Marra et al. (2024) to point out the similarities between these two endeavours and in this way stimulate more cross fertilisation. They started from the literature on StarAI because, arguably, there is more consensus on what the key concepts, challenges and issues are in StarAI than in NeSy. They argue that essentially the same issues and techniques that arise in StarAI have to be addressed in NeSy as well. Their key contribution is that they identify a set of seven dimensions that these fields have in common and that can be used to categorise both StarAI and NeSy approaches. These seven dimensions are concerned with (1) directed vs undirected models, (2) grounding vs proof based inference, (3) integrating logic with probability and/or neural computation, (4) logical semantics, (5) learning parameters or structure, (6) representing entities as symbols or sub-symbols, and, (7) the type of logic used. They provide evidence for their claim by positioning a wide variety of StarAI and NeSy systems along these dimensions and pointing out analogies between them. This, in turn, allowed them to identify interesting opportunities for further research, by looking at areas across the dimensions that have not seen much work yet. Of course, they also identify important differences between StarAI and NeSy, the most important one being that the former operates more at the symbolic level, lending itself naturally to explainable AI, while the latter operates more at the sub-symbolic level, lending itself more naturally for computer vision and natural language processing. This is directly relevant to TAILOR, as StarAI is more directly applicable to trustworthy AI than NeSy.

Task 4.3: Learning and Reasoning with Embeddings, Knowledge Graphs & Ontologies

State of the art. In recent years, Knowledge Graphs (KG) have become a growing trend for knowledge representation, and several works on different methods to deal with ontologies and graph-structured data have been proposed. Historically, many techniques applied to large-scale knowledge graphs came from the Statistical Relational Learning (SRL) community, as SRL methods can be easily exploited on existing KGs to predict new facts from existing ones or even correct "noisy" known facts in the graph. For instance, this class of approaches was already discussed in the survey of [Nickel et al. 2016], that focuses on scalable SRL techniques that are applicable to graphs with millions of nodes and billions of edges. Another trend that has grown greatly in the last few years relies on exploiting low-level distributional semantics for nodes and edges, which is the fundamental idea behind Knowledge Graph Embeddings (KGE). Among others, a recent survey on the state of the art of KGE is considered by Dai et al. (2020), which extends a previous well-established survey of Wang et al. (2017), by describing various applications to which KG embeddings apply and comparing the performance of different methods on these applications. In particular, Dai et al. (2020) split the review of existing KGE models by considering at first the embedding models that only use the information coming from observed triplets in the KG, and then considering classes of advanced models that utilise additional semantic information to improve the performance of the original methods. The general class of embedding models just relying on factual triplets can be categorised in the following three groups: (i) translation-based models (ii) tensor factorization-based models and (iii) neural network-based models. For each of these groups, a set of prominent examples of classic and lately devised models are introduced, as well as a variety of applications that may benefit from these methods. Concerning KGE models exploiting additional information, Dai et al. (2020) mostly focuses on textual descriptions and relation paths, while in Wang et al. (2017) the authors also considered the cases of entity types and logical rules. Combining KGE with additional semantic information, and/or other kinds of approaches providing a different reasoning scheme, is currently one of the fundamental research directions that is widely under investigation, and some of the main advancements in this area have been collected in some surveys. For instance, [Li et al. 2020] consider approaches extending KGE models with hybrid reasoning techniques like symbolic reasoning and statistical reasoning, whereas Zhang et al. (2021) and Delong et al. (2023) focus on neuro-symbolic reasoning techniques. In particular Delong et al. categorise the state-of-the-art according to the way the logic rules are used by Neural-Symbolic systems into three main groups: (i) *Logically-Informed Embedding Approaches*, such as UniKER [Cheng et al. 2020] and RUGE [Guo et al. 2018], (ii) *KG Embeddings and Deep Learning with Logical Constraints*, such as QLogicE [Chen et al. 2022] and R2N [Marra et al. 2021], and (iii) *Learning Rules for Graph Reasoning*, such as RNNLogic [118] and NCRL [Cheng et al. 2023].

When considering graph-structured data, a popular research line that in these years has been the subject of many investigations concerns the study of Graph Neural Networks (GNN). For instance, a recent tutorial on "Graph Neural Networks and neural-symbolic computation" has been held by Prof. Marco Gori at the UCA Deep Learning School in 2021 https://univ-cotedazur.fr/evenement/deep-learning-school. The tutorial was split in two phases: at first, the general theory and most common applications of Graph Neural Networks have been presented, especially highlighting their connections with neural-symbolic models; secondly, it has been carried out a laboratory activity focusing on the presentation of available frameworks to define GNN models and classic learning tasks like, Node Classification, Graph Classification and Link Prediction.

One recent survey on the state of the art of GNNs is provided by Maggini et al. (2024), where an alternative view of neural network computational scheme, and in particular of GNNs, is formulated as a satisfaction problem of architectural constraints. More specifically, architectural constraints are considered as a unifying principle to define different neural architectures that also relates to the theoretical framework for BackPropagation formulated using Lagrangian optimization. In this setting, Graph Neural Networks' learning process is stated as the outcome of a joint process where the state computation on the input graph is expressed by a constraint satisfaction mechanism that does not require an explicit iterative procedure and the network unfolding. When injected into the original convergence-based Graph Neural Networks [Scarselli et al. 2009], the approach simplifies the learning procedure, avoiding the need to explicitly compute the fixed point of the state transition function during each epoch of the learning procedure. However, the approach is also extended to the case of Layered GNNs, in which multiple representations of each node are computed by a pipeline of constraints, that is related to a multi-layer computational scheme.

Achievements w.r.t Tasks 4.1 and 4.3.

The key achievements for Tasks 4.1 and 4.3 can be grouped following the topics already listed in the survey and the representations mentioned in the figure. We first discuss the LOR – learning, reasoning and optimization aspects, and then the TAI – Trustworthy AI aspects, first for pure Statistical Relational AI, then for its extensions towards Neurosymbolic AI, and finally also for knowledge graphs.

StarAI:integrating Logic and Probability for which it is useful to also further distinguish the type of representation, i.e., propositional from relational or first-order logical representations, and the type of contribution, i.e., improving the model learning process, the inference process, model analysis and trustworthiness, etc.

StarAI Representations: a) relational representations:
Within TAILOR there has been work on StarAI representations (and as a subfield, probabilistic programming representations). For instance, Azzolini et al. (2022b) investigate a new expressive StarAI formalism that is based on ASP and allows statistical statements such as "80% of the birds fly", while Azzolini and Riguzzi (2024) worked on extending ASP semantics with imprecise probabilities. Hierarchical probabilistic programs have also been considered [Fadja et al. 2021], which can be directly converted to an arithmetic circuit/neural network. GCLN [Ventola et al. 2021] introduces the first connection between relational rule

models and probabilistic circuits, the latter of which represent a sort of probability distribution through a computational graph (strongly relating to arithmetic circuits and sum product networks). When satisfying certain structural properties, such circuits enable tractable probabilistic inference, and are therefore a strong and popular (final) representation (of other modelling representations).

StarAI Representations: b) propositional representations:

Many knowledge representation languages deal with knowledge over a Boolean domain. A new recent addition to this collection of languages are so-called switch-list representations, whose succinctness and complexity (of transformations-, and of several common queries for efficient reasoning) has been investigated by Čepek (2023) and Čepek et al. (2023). In the probabilistic setting, SPLL has been proposed as a new probabilistic programming language [Pfanschilling et al. 2022]. This language performs inference through the use of sum product networks (SPNs, a kind of probabilistic circuit), for which they specifically overcome the limitation of loops, to support recursive calls. This results in an expressive language that still maintains the same tractability of inference. Pellegrini et al. (2021) introduce Learnable Aggregation Functions, a fully differentiable layer for set aggregation that can approximate several extensively used aggregators as well as more complex functions.

StarAI Inference.

As efficient inference is crucial for the adoption and general usage of StarAI models, several works within TAILOR have specifically focused on developing more efficient inference algorithms [Azzolini et al. 2022a, Azzolini and Riguzzi, 2023b]. Various of these algorithms use knowledge compilation to compile the initial model representation into another (intermediate) one that is more efficient to reason over [Fadja et al. 2021; Azzolini and Riguzzi, 2023a, Yang et al. 2022a, Ahmed et al. 2022, Ahmed et al. 2023]. For instance, Venturato et al. (2024) leverages such a representation to perform probabilistic reasoning and learning in a dynamic decision-making setting. A compilation based approach is particularly interesting for (later) integrating with neural networks, as some of the compiled target representations are differentiable and can therefore easily be combined and learned alongside a neural network. Knowledge compilation also very much relates to probabilistic circuits, which can be learned from data or compiled from knowledge. Vergari et al. (2021) report on a unified, extensive study on efficient algorithms for carrying out probabilistic and information-theoretic computations for a very general family of probabilistic and statistical-relational models. Spallitta et al. (2022) instead focus on SMT formulas, which extend propositional logic with atoms that may contain, for example, arithmetics such as $(x + y < 2)$. They introduced a novel algorithm for hybrid probabilistic inference that combines SMT-based enumeration with an effective encoding of the problem structure. This work was then extended with a simpler and more effective encoding [Spallitta et al. 2024]. As inference algorithms can also be sampling-based, the work of Sansone (2022) is also relevant. He presents LSB, a new sampling method for purely discrete domains. It is a local Markov Chain Monte Carlo (MCMC) sampling method that can autonomously adapt to the target distribution, thereby reducing the number of target evaluations required to converge.

StarAI learning.

Importantly, TAILOR also contributed to the learning of StarAI models (i.e., structure and/or parameter learning). For instance, EMPLiFi [Yang et al. 2022b] was proposed as a more efficient expectation maximisation-based parameter learning technique for probabilistic logic programs. Gentili et al. (2023) instead focuses on structure learning, proposing a method to speed up learning by first generating a large set of probabilistic rules and then pruning it by using parameter learning with regularisation. iSPN [Zecevic et al. 2021] focuses more on learning in the context of causality, which is important for creating more robust and trustworthy models. They consider the problem of learning interventional distributions (i.e., answering causal queries) with tractable probabilistic models (gated SPN). [Massidda et al. 2024b] introduced a method to learn causal abstractions from data, leveraging induced constraints to improve scalability. Massidda et al. (2024a) introduce a constraint-free continuous optimization scheme for acyclic structure learning.

Trustworthy StarAI.

Sokol and Flach (2024) have studied interpretable representations in XAI, making a range of recommendations for designing trustworthy and interpretable representations. Orthogonal to that work is the survey of Silva Filho et al. (2023), which provides a detailed overview of the principles and practice of classifier calibration, including introductory material and up-to-date technical details of the main concepts and methods.

To improve trustworthiness and explainability in the context of probabilistic logic programs (PLP), Vidal et al. (2022) studied how to generate minimal explanations as programs. A complementary approach, where the explanation is instead a proof tree, was later proposed by Vidal et al. (2024), which also introduced a new query-driven inference mechanism for PLP. While these works focus on PLPs, there is also work that instead focuses on probabilistic circuits (recall that these formed an interesting representation to compile into, for inference, or to learn from data): Bocklandt et al. (2024) considers the pruning of a learned probabilistic circuit to obtain a logical discriminative classifier that is easier to inspect and validate by a domain expert. More generally, the work of Noé et al. (2024) proposes a method to explain the probabilistic prediction of multi-class probabilistic classification models, by quantifying the contribution of each feature's value to the predictions. The method offers an extension of the Shapley values – which explain unary predictions – to a multiclass solution. To further investigate a (learned) model, and specifically to formally verify its properties, note that the task of model checking is highly relevant. Azzolini and Riguzzi (2022) performed model checking in the context of a routing network – the lightning payment network. They study the problem of reasoning about the existence of a network path that can route a payment, leveraging multiple models based on PLP for the lightning network. In case the routing in such a network can be modelled as a relational MDP, the work of Yang et al. (2022a) is applicable. Yang et al. (2022a) contributed PCTL-REBEL, a StarAI approach to model checking for relational MDPs in a lifted manner; work that also relates to planning and WP 5. Clifford et al. (2023) explore the mismatch between standard post-hoc, model-agnostic surrogate explainers and the black box models. The mismatch is due to the differences between the data distributions used to train and evaluate the surrogate explainers. The paper encourages users to take care in understanding the alignment between training and evaluation objectives to construct more faithful surrogate explainers. Besides explanations, there is also argumentation. For instance, Donadello et al. (2022) studied a realistic use case concerning healthy eating habits. They integrate machine learning and argumentation techniques for the purpose of adapting the response of an automated persuasion system to the needs of a subpopulation of users. This is accomplished by learning to identify a utility function per subpopulation. Also taking into consideration user utilities and preferences, is the work of Chesani et al. (2022a, 2022b, 2023). They present an approach founded in a logic-based framework, for the discovery of declarative process models from positive and negative examples, taking into consideration user preferences on activities. Yamagata et al. (2024) propose a preference-based reinforcement learning algorithm that jointly estimates the trainer's reliability and preferred actions. Finally, note that PLP-based modelling approaches can enhance the learning process of classifiers by modelling more explicitly their data assumptions. Verreet et al. (2023a) studied this, by exploiting PLP-based modelling and inference methods in the context of PU learning – a particular learning setting that consists of only positively-labelled (P) and unlabelled (U) data. Specifically, they studied how to formulate and integrate more realistic assumptions, through PLP-based modelling, to learn better classifiers in the context of PU data.

StarAI approaches to NeSy.

As indicated in the survey [De Raedt et al. 2020, Marra et al. 2024] already described under Task 4.1, many NeSy approaches take inspiration from StarAI or can be interpreted in these terms. Ahmed et al. (2022, 2023), for instance, introduce semantic probabilistic layers for neuro-symbolic learning. In this approach, one converts knowledge into an efficient representation to be used during learning (i.e., an arithmetic or probabilistic circuit), and incorporates the representation within the loss function of the neural network. As a result, through the loss function, the neural network will be encouraged to incorporate the knowledge in its predictions (i.e., satisfy the compiled knowledge). DeepStochLog [Winters et al. 2022] forms a different approach, being an extension of the StarAI inspired stochastic logic programs (which are based on probabilistic definite clause grammars, akin to probabilistic unification based grammar) towards neural networks. The extension is based on the notion of a neural predicate [Manhaeve et al. 2021]. This same predicate was also used in DeepSeaProbLog [De Smet et al. 2023a], where the authors contributed support for a discrete-continuous domain. However, while learning a NeSy model that is based on neural predicates, there is a risk that the learned concept associated with each neural predicate does not precisely align with their expected semantics [Manhaeve et al. 2021, Marconato et al. 2023a]. 'Reasoning shortcuts', for instance, allow the model to (unexpectedly) exploit the problem specification (i.e., the loss function through its structure or data) during the learning process [Marconato et al. 2023a], yielding a low loss but still an undesirable result. Consider for example the MNIST digit task with the following constraint, digit1 + digit2 = digit3. Without additional supervision or information, the model could learn to map every MNIST digit to `0` and would achieve a high consistency but an undesirable alignment. Achieving a correct alignment is an important challenge that affects the robustness, interpretability and trustworthiness of such models, and has consequently been more thoroughly studied. For instance, Marconato et al. (2023b) characterises the appearance of reasoning shortcuts and experimentally verifies the effectiveness of several mitigation strategies. Another relevant challenge in this domain is knowledge drift, a form of concept-drift that occurs in relational data. [Bontempelli et al. (2022) introduces an effective solution to this problem, combining automated detection with an interactive human-in-the-loop disambiguation strategy. Finally, Marconato et al. (2023a) considers NeSy in the context of continual learning, and introduces a learning strategy to prevent catastrophic forgetting of concept semantics.

StarAI inspired NeSy Systems.

In addition to the NeSy systems mentioned in the previous paragraph, we list a few systems recently developed within TAILOR. Relational Neural Machines [Diligenti et al. 2022] integrate learning and reasoning on an undirected graphical model, recovering classic supervised learning and Markov Logic Networks as special cases. NSFR [Shindo et al.

2021] is a new approach that takes advantage of differentiable forward-chaining using first-order logic. SLASH [Skryagin et al. 2022] consists of Neural-Probabilistic Predicates and logic programs which are united via answer set programming. The probability estimates resulting from these predicates act as the binding element between the logic program and raw input data, thereby allowing SLASH to answer task-dependent logical queries. RD2GCN [Dhami et al. 2022] connects graph neural networks with statistical relational learning, thereby paving a way for moving towards relational graph neural networks with rich structural information. To explain the global behaviour of graph neural networks, Azzolin et al. (2023) propose an approach that results in Boolean combinations of learned graphical concepts. Giannini et al. (2023a) demonstrate how fuzzy logic relaxation simplifies numerical properties and enhances integration between logical knowledge and learning objectives. Extending these benefits to neuro-symbolic methods, they introduce novel loss functions that achieve faster convergence rates than the existing methods. VAEL [Misino et al. 2022] bridges the gap between variational autoencoders (VAE) and probabilistic logic programs. This NeSy system offers capabilities which go beyond the traditional properties of deep generative models. The reasoning component of VAEL provides an inductive bias for the latent space of a VAE, which allows to structure the representation so that the model can generalise to previously unseen tasks. Furthermore, this hybrid generative model offers the capability to learn from smaller amounts of training data compared to purely neural-based approaches. Constrained Adversarial Networks [Di Liello et al. 2020] enrich GANs with a certifier implemented via a semantic loss layer that allows them to learn to generate structured objects satisfying known constraints in expectation. Verreet et al. (2023b) instead focused on the scalability challenge of logic-based NeSy learning by presenting a sampling based paradigm.

Efficient NeSy Inference. Several NeSy systems achieve an end-to-end differentiable approach with neural networks. Key to learning in such systems is the ability to efficiently compute gradients. Maene et al. (2024b) investigated the complexity of this task for probabilistic reasoning, and introduced a novel gradient estimator with probabilistic guarantees. Also relevant to this work is the well known Log-Derivative trick, which is a technique to compute the gradient of an expression through another expression, where computing the gradient of the latter is hopefully easier. De Smet et al. (2023b) has extended this by tailoring the estimation of gradients to those over categorical distributions. This extension results in both theoretically and empirically unbiased, low-variance estimates.

NeSy for Trustworthy AI.
These NeSy approaches can be applied to create more safe and trustworthy AI. For example, Yang et al. (2023) proposed a method for safe reinforcement learning via probabilistic logic shields, where a probabilistic logic is used to model logical safety constraints as differentiable functions that can seamlessly be applied to any policy gradient algorithm. Delfosse et al. (2023) similarly proposed a reinforcement learning framework that uses differentiable logic reasoners to reason and learn action distributions. The resulting agent is inherently more explainable and interpretable while learning to solve environments using policy-gradient methods.

Concept-based NeSy. Focusing on explainability, Barbiero et al. (2022) make neural models explainable through an entropy-based approach that generates First-Order Logic (FOL) explanations. This method helps reverse engineer algorithms, identify vulnerabilities, and improve system design. It also aids scientific discovery by distilling formal knowledge from advanced networks. Ciravegna et al. (2023a) present Logic Explained Networks (LENs), a general approach to Explainable AI for neural architectures [Ciravegna et al. 2023b, 2023c]. They used human-understandable predicates, and provide concise explanations through simple FO logic formulas. LENs can interpret their own behaviour or that of other black-box models across different domains like tabular data, computer vision, and natural language processing using supervised or unsupervised learning. As the local explanations of LENs can be noisy and verbose, Jain et al. (2022) proposed LEN^p, which improves LEN by perturbing input words and testing it on text classification. A human survey confirmed that this approach resulted in more useful and user-friendly logic explanations than LIME's feature scoring. Concept bottleneck models (CBMs) form another approach to enhance the interpretability of neural networks. Recently, a new method was proposed to transform existing trained models into CBMs efficiently, enabling concept-based explanations and interventions [Dominici et al. 2024]. Case studies demonstrate the ability of this approach to improve interpretability while maintaining high classification performance. Concept Bottleneck models also appeared in the work Giannini et al. (2024). They introduced a mathematical framework using category theory, enabling precise definitions of key XAI concepts and processes. It incorporates interpretable Neural-Symbolic AI methods like rule-based systems and the aforementioned Concept Bottleneck Models for providing logic-based explanations of neural networks. This framework allows for modelling various learning schemes, formally defining explanations, establishing XAI taxonomies, and analysing overlooked aspects of explanatory methods.

NeSy applications for a broader understanding. To better understand the capabilities of the developed NeSy systems, many have also started to create more interesting, mature benchmarks and categorisations. Since deliverable 4.4 is dedicated to this topic, we only briefly mention the additional work of Valenti et al. (2024), who introduce ChemAlgebra, a benchmark dataset to assess the reasoning capabilities of deep learning models in chemical applications, and Vermeulen et al. (2023), who provide an overview and classification of the tasks used in state-of-the-art NeSy systems, showing that most tasks fall in one of five categories, and that few are compared on the same benchmarks.

StarAI, NeSy, Embeddings and Ontologies.

The StarAI and NeSy methods are often also integrated with embeddings, knowledge graphs and/or graph neural networks. We list a few of such integrations developed within TAILOR. For instance, DeepSoftLog [Maene et al. 2024a] provides DeepProbLog with the additional capability of reasoning over learnable embeddings through soft-unification; and Concept Embedding Models [Zarlenga et al. 2022] have been introduced to learn interpretable high-dimensional concept representations. Barbiero et al. (2023) introduced Deep Concept Reasoner (DCR), an interpretable concept based model that enhances the interpretability of high-dimensional concept embeddings by using syntactic rule structures. Unlike traditional models, DCR does not directly predict tasks but instead constructs rules from concept embeddings and evaluates them for semantically consistent predictions in a differentiable way. Relational Reasoning Networks (R2N) [Marra et al. 2021] perform learning and reasoning in latent spaces, by means of embedding representations of logic atoms and first-order logic formulas. R2N has later been used to enhance knowledge graph embeddings that struggle with complex dependencies [Diligenti et al. 2023]. For example in the biological domain, PharmKG is a dataset on which this work shows significant improvement. From the application perspective, there is the relevant work of Galassi et al. (2021), who proposed to use a neuro-symbolic approach to mine argument components and relations from textual corpora; and of Bacciu et al. (2023a) who addresses the problem of drug repurposing by leveraging graph embeddings to encode proteins' and drugs' information based on gene ontology data and structural similarities. Zeinalipour and Gori (2023) instead study electrocardiogram (ECG) signals, introducing three innovative classification techniques using deep graph neural networks. Finally, Macková and Pilát. (2024) introduced a new dataset for product mapping, i.e., the task of determining whether two product descriptions from different e-shops represent the same product. This is a task well suited to using embeddings.

Integrating (knowledge) graphs. Within TAILOR we also explored the challenges associated with (knowledge) graphs and their integration into larger paradigms. For instance, LGE [Dhami et al. 2021] is a structure learning algorithm to learn embeddings for relational data (knowledge graphs), taking into account the local information in the graph. Bacciu et al. (2023b) introduced a unifying interpretation of downsampling in regular and graph data, while McDonald et al. (2024) encoded a trained GNN into a MIP optimisation model. Keskin et al. (2023) introduces an innovative approach to establishing correlations between equational properties of algebraic structures, represented through graphs, and specific sub-portions of their topological structure. The methodology utilises graph neural architectures to validate theorems or conjectures, supported by XAI metrics to substantiate these assertions. Giannini et al. (2023b) demonstrate AI's potential impact in pure mathematical research by introducing an interpretable GNN to investigate equational and topological conjectures in universal algebra.

Reason-able embeddings are recently proposed learnable embeddings for knowledge bases (KBs), that are capable of casting multiple KBs into a single latent space using a transferable neural reasoner [Adamski and Potoniec, 2023]. To understand the limits of reason-able embeddings, Potoniec. (2023a) investigated their performance using a set of synthetic KBs, using proof length as a measure of reasoning complexity. With the aim of establishing a theoretical bound on their applicability, [Potoniec, 2023b] constructed a dataset that is experimentally shown to be hard for them.

NeSy with embeddings, analysing trustworthy AI. The rise of large language models has also increased the relevance of studying safe and trustworthy AI. Within TAILOR there has been an interest in analysing such models. For example, Poulis et al. (2024) used a description logic to construct a natural language reasoning dataset, called DELTA_D, which they then used to systematically investigate the reasoning ability of two large language models and a supervised fine-tuned DeBERTa-based model. The rise of large language models at the same time also led researchers of other fields to consider whether such models could help to improve the challenges within their field or vice versa. This then raises important questions in regard to the integration of these models into StarAI and NeSy approaches. For example, large language models could improve the quality of interacting with a human in the loop. In this way, Mechqrane et al. (2024) uses large language models to improve query-based constraint acquisition. The language model component is here used to interpret the user's answers given in natural language, to reduce the number of queries posed to the user. Hazra et al. (2024) leveraged a large language model for plan generation, guided by an heuristic search that evaluates the feasibility and optimality of each action. Di Maio et al. (2024) introduce Continual In-context Knowledge Large Language Models (CIK-LLM), which combine LLMs' inference capabilities with temporal knowledge graphs (TKG). Its local exploration of these graphs, which can continually integrate evolving knowledge, ensures scalability and enhances transparency, while preventing erroneous inferences and hallucinations.

Creative Ne(Sy) tasks (requiring embeddings). Other creative task that were studied include the automatic generation of educational crosswords (in a language different from English) using language models [Zeinalipour et al. 2023a, Zeinalipour et al. 2023, Zeinalipour et al. 2024a], the automatic generation of quizzes from educational texts [Zeinalipour et al. 2024b], spatial reasoning on puzzle diagrams by employing logic programming to derive knowledge [Buscaroli et al. 2022], and the automatic evaluation of (robotic) dance performances where various prediction method were tested to predict the audience evaluation, with the aim of helping in the creation of new dance performances [De Filippo et al. 2022].

Logic, StarAI and Ontologies. Various contributions integrate logical and ontological knowledge representations and study their properties. In Cima et al. (2021) the notion of abstraction in ontology-based data integration is studied and techniques for computing abstractions are presented. An abstraction is an abstract representation whose aim is to explain the semantics of a concrete computation expressed as a query on a set of data sources. In Cima et al. (2021b) the problems of query definability is investigated in the context of ontology-based data management. The problem of query definability is the one of deriving a query characterising a dataset given in input, and one notable application of such a problem is to explain the semantics of a black-box classifier. Console et al. (2021) presents a characterization of a broad class of ontologies based on properties of their models, i.e., the structures that satisfy their axioms. This characterization provides boundaries on the absolute expressive power of such ontologies and defines the relative expressive power of different sub-languages. The latter results provide algorithms for the rewritability problem, i.e., checking whether a given ontology can be equivalently rewritten in a less expressive and better behaved language. Other contributions use fuzzy logic as an alternative for probability. In particular, Cardillo and Straccia (2022) present a method that, given an OWL ontology and a target class T, addresses the problem of learning fuzzy concept inclusion axioms that describe sufficient conditions for being an individual instance of T (and to which degree). To do so, it presents Fuzzy OWL-Boost that relies on the Real AdaBoost boosting algorithm adapted to the (fuzzy) OWL case.

Further contributions related to machine learning include CRISPS [Teso and Vergari, 2022], which is a novel class of deep probabilistic classifiers specifically designed for supporting different forms of interactive machine learning in an efficient and reliable manner; ReliefE [Skrlj et al. 2022], which performs distance-based feature ranking in high-dimensional spaces via Riemannian manifold embedding; and the approaches for ensemble and distance-based feature ranking for unsupervised [Petković et al. 2021], and semi-supervised learning [Petković et al. 2022a].

Task 4.2: Integrating Representations for Learning and Optimization

State of the art. To the best of our knowledge, the most recent survey focused on the state of the art for integrating representations for learning and optimisation is provided in [Teso et al. 2022]. It bridges the gap between approaches to combinatorial optimization and machine learning and argues that regret minimization provides a unifying view on this newly emerging field. More specifically, they consider combinatorial optimisation problems that are only partially-specified. They survey the case where the objective function or the relations between variables are not known or are incompletely specified. The challenge is to learn them from available data, while taking into account a set of hard constraints that a solution must satisfy, and that solving the optimisation problem (esp. during learning) is computationally very demanding. Their survey overviews four seemingly unrelated approaches, that can each be viewed as learning the objective function of a hard combinatorial optimisation problem: 1) surrogate-based optimisation, 2) empirical model learning, 3) decision-focused learning ('predict + optimise'), and 4) structured-output prediction. They formalise each learning paradigm, at first in the ways commonly found in the literature, and then bring the formalisations together in a compatible way using regret. They discuss the differences and interactions between these frameworks, highlight the opportunities for cross-fertilization and survey open directions in the field.
While other focused surveys on the topic of integrated representations appear to be lacking, related fields such as Decision Focused Learning (DFL) have received considerable more attention. The most up-to-date overview of work in the area is provided by [Mandi et al., 2023]. Notably, research in DFL has incorporated some of the insights from research conducted on integrated differentiation in the context of TAILOR, in particular highlighting the viability of using black-box optimization (and specifically black-box differentiation or surrogate learning) to implement DFL approaches [Silvestri et al. 2023].

Achievements w.r.t Tasks 4.2.

Several avenues for combining learning and optimization representations were explored, in a variety of domains.

Optimizations and machine learning. There are some key contributions that combine machine learning with optimization algorithms. First, Mulamba et al. (2021) introduced a strategy based on the predict-and-optimise paradigm that leverages solution caching. Second, De Filippo et al. (2022) is another example of integrating optimization and machine learning models through Empirical Model Learning, in which the authors propose an approach to automatically perform hardware dimensioning and configuration for online algorithms in the energy system domain, under an heterogeneous set of constraints. The machine learning models are used to predict the online algorithms' performance on different hardware configurations and optimization is used to find the optimal matching of computing resources and algorithm configuration, while respecting user-defined constraints (e.g., cost, time, solution quality). Related is also the Hybrid Offline/Online Optimization for Energy Management via Reinforcement Learning by Silvestri et al. (2022). Third, Kumar et al. (2021) describe a learning approach for acquiring mixed-integer linear programming models from historical data that leverages both gradient-based and combinatorial search for learning. Finally, Boros et al. (2024) studied the integration of reasoning with optimisation in the context of game theory. They provided polynomial algorithms to recognise voting forms and voting correspondences generated by voting schemes in cases when either the number of candidates or the number of voters is equal to 2. They prove that for two voters, the unique voting correspondence has distinct rows. Pourkhajouei et al. (2023) developed a novel active preferences learning approach for interleaving optimization and elicitation of preferences for situations when the user model is noisy. In other works on this topic, a very important tool is minimax regret; the methods in Wilson (2023) enable minimax regret to become more robust.

Constraints. Task 4.2 is also concerned with the use of constraints in learning and reasoning. Here, a couple of works at TU Delft and UNIPI have investigated the use of graph neural networks to perform constraint reasoning [van Driel et al. 2021] and to generate meaningful graph-structured counterfactuals and interpretations from graph neural networks [Numeroso and Bacciu 2021], which is also relevant to NeSy. Morettin et al. (2021) contributed a survey on hybrid probabilistic inference with logical and algebraic constraints under the unifying paradigm of Weighted Model Integration, which is also highly relevant to Task 4.1, probabilistic circuits and StarAI. Finally, De Canditiis and De Feis (2021) contributed an approach to anomaly detection in multichannel data using sparse representation in RADWT frames and (Lombardi et al. 2020) analyse regularised approaches for constrained machine learning.

Constraint acquisition. There has also been work on assisting the user in creating constraint models. First, there is the question of which constraint language to use. Bessiere et al. (2023b) propose a new constraint acquisition method that computes a suitable constraint language as part of the learning process, eliminating the need for any advance knowledge. Second, there is the formulation of constraints, wherein Bessiere et al. (2023c) assists the user: they perform constraint acquisition through user interaction via partial queries, showing how human interaction can help to learn and improve models. Similarly, Mechqrane et al. (2024) performs constraint acquisition but exploits the advancements in large language models. Prestwich and Wilson (2024) instead developed a more statistical approach to constraint acquisition based on sequential analysis, with a Bayesian analysis. Carbonnel (2022) provided a theoretical study, introducing an algebraic framework dedicated to the analysis of the non-redundancy of constraint languages. The amount of non-redundancy of a constraint language allows them to derive a lower bound on the number of examples required to learn problems formulated on the given constraint language. Related, but on the important topic of explainability, is the work of Bessiere et al. (2022, 2023a) who study the complexity of explaining the outcome of a solver on a constraint program as a sequence of human understandable reasoning steps.

Feature-ranking and ensemble methods. TAILOR has also contributed to recent advances in feature ranking and ensemble methods across multiple contexts. Osojnik et al. (2023) introduced iSOUP-SymRF, an innovative online feature ranking technique for multi-target regression, leveraging the positions of features in iSOUP-Trees within random forests to estimate importance scores. By utilising iSOUP-Trees, which can address multiple structured output prediction tasks on data streams, iSOUP-SymRF provides feature ranking across a variety of online structured output prediction tasks. Petković et al. (2022) developed methods for feature ranking in relational contexts, by adopting the relational tree ensemble approach. They consider complex relational features: by using complex aggregates, they extend the standard collection of features that correspond to existential queries to more complex features that correspond to aggregation queries. They also compute feature importance scores and rankings, which provides insights into and explain the ensemble models which would otherwise be difficult to understand. Further extending this work, they introduced CLUSplus [Petković et al. 2023], a decision tree-based machine learning framework designed for complex predictive modelling such as multi-target prediction with hierarchically organised values. CLUSplus enables state-of-the-art predictive performance via ensemble learning, exploitation of unlabeled data via semi-supervised learning, and data understanding via feature importance and building interpretable models.

Optimizations and explainability. A key focus area has been the explanation of AI models, and neurosymbolic ones in particular. Sabbatini et al. (2024) studied how to measure the quality of symbolic knowledge extracted from data, encompassing several indicators and providing a compact score reflecting readability, completeness and predictive performance associated with a symbolic knowledge representation. Sabbatini et al. (2022a) explored rule extraction techniques that associate human-interpretable knowledge with accurate predictions from opaque models. Assessing the readability of extracted knowledge quantitatively remains unresolved. Developing such a metric would enable automatic comparisons and parameter autotuning for knowledge extractors. In a similar fashion, Sabbatini et al. (2022b) presented GridREx, a pedagogical algorithm to extract knowledge from black-box regressors, along with PEDRO, an optimisation procedure to automate the GridREx hyper-parameter tuning phase with better results than manual tuning. Sabbatini et al. (2023) analyse a recurrent design adopted by symbolic knowledge extractors for black-box predictors—that is, the creation of rules associated with hypercubic input space regions [Sabbatini et al. 2023a, 2023b, 2023e]. This kind of partitioning may lead to suboptimum solutions when the data set at hand is sparse, high-dimensional, or does not satisfy symmetric constraints. Different knowledge-extraction workflows were proposed, involving clustering approaches, highlighting the possibility of outperforming existing knowledge-extraction techniques in terms of predictive performance on data sets of any kind. Along the same task of symbolic knowledge extraction, they showed that hypercube-based methods are flexible enough to support classification problems, proposed a general model for them, and discussed how they support symbolic knowledge extraction on datasets, predictors, or learning tasks of any sort [Sabbatini et al. 2023c, 2023d, 2024]. Betti et al. (2024) describes a novel approach to learning and optimization, where time plays a crucial role. A neural network learning scheme is presented, focused on problems in which data are streamed over time and where online learning constraints are present. This work creates the bases for investigating learning and optimization in a context that shares principles inherited from optimal control theory.

Fairness and AI. One area that was subject of several studies was the interplay of fairness and AI. Fairness in machine learning is crucial but often analysed with complex methods. Maggio et al. (2023) introduced GEOFFair, a geometric framework that represents distributions, ML models, and fairness constraints as vectors and sets. GEOFFair provides an intuitive understanding of fairness, visualises mitigation techniques, and studies fairness properties, including geometric distances and trade-offs. Again in the field of AI and fairness, Giuliani et al. (2023b) explored fairness over continuous protected attributes. First, they show that the Hirschfeld-Gebelein-Renyi indicator (the only one currently available for such a case) is valuable but subject to a few crucial limitations regarding semantics, interpretability, and robustness. Second, they introduce a family of indicators that are: 1) complementary to Hirschfeld-Gebelein-Renyi in terms of semantics; 2) fully interpretable and transparent; 3) robust over finite samples; 4) configurable to suit specific applications.

Automated Planning. TAILOR members have also contributed to planning in AI. One specific problem that was studied in this area is plan recognition: the problem of recognizing a goal task and an agent's plan based on the observed actions. This problem is relevant as these techniques can be employed in multiagent systems, behaviour recognition, computer security, and other fields related to artificial intelligence. Hierarchical task networks (HTN) describe the decomposition hierarchy of tasks in planning problems. In HTN plan recognition, a prefix of the plan (actions observed so far) is given as an input, and the aim is to find a task that decomposes into a sequence of actions with the given prefix. Work in this area was done by Pantůčková et al., who proposed to use techniques from formal grammars [Pantůčková et al. 2023a] and constraint programming [Pantůčková et al. 2023b] to solve this problem. Another highly related problem that was contributed to, is the task of hierarchical plan verification: determining whether an action sequence is causally consistent and can be obtained by a decomposition of a goal task. Recognizing that the task decomposition structure is very close to a parsing tree of context-free grammar (CFG), recent work has proposed to use techniques from formal grammars [Lin et al. 2023, Ondrčková et al. 2023a, 2023b, Pantůčková et al. 2024]. TAILOR members have also made fundamental contributions to the semantics for the HTN planning domain [Ondrčková et al. 2023c, 2024] and on how to express control knowledge in automated planning in a way that improves efficiency of planning without the need to modify the planner itself [Chrpa et al. 2023].

Application in Music generation. The contributions of TAILOR span diverse applications of machine learning in complex domains. One example is in music generation, where the recent advances in deep learning have led to large end-to-end neural architectures, producing convincing outputs but with drawbacks like low user control, lack of global structure, and high computational costs. Giuliani et al. (2022, 2023a) propose a novel approach combining samples under user-defined constraints. Samples are generated using deep learning models (e.g., via Transformer-like neural networks), while the combining of the samples is done via Constraint Programming. By modelling the task as a job-shop problem, the method achieves interesting results with low computational costs and is genre-independent, though it can accommodate genre-specific constraints. Such a tool offers several benefits in enhancing human creativity, as they provide the opportunity to keep human artists in the creative loop as well as to reduce computational costs and hardware requirements.

Other applications studied within TAILOR include survival analysis [Roy et al. 2022], the prediction of thermal power consumption [Stevanoski et al. 2023], and the development of surrogate models for the inversion of radiative transfer models, which is important to interpreting satellite observations [Brence et al. 2023]. Together, these studies illustrate the application of advanced machine learning techniques to enhance predictive accuracy and computational efficiency.

- Roy et al. (2022) reframe survival analysis as a multi-target regression task and employ semi-supervised predictive clustering trees, demonstrating superior performance across eleven real-world datasets compared to existing methods.

- Stevanoski et al. (2023) address the prediction of thermal power consumption for the Mars Express spacecraft through multi-target regression on telemetry data streams, evaluating the efficacy of various modelling approaches in handling time resolution and adaptability to change.
- Inversion of radiative transfer models is key to interpreting satellite observations of air quality and greenhouse gases, but is computationally expensive. In the study of Brence et al. (2023), they develop two surrogate models for this problem. Their results show that surrogate models are able to accurately emulate Sentinel 5P spectra within a millisecond or less, as compared to the minutes or hours needed to simulate the full physical model. Interestingly, they also found that models trained on the smaller (n = 1000) uniformly sampled dataset can perform as well as those trained on the larger (n = 50000), more focused dataset.

We also highlight the Dagstuhl seminar organised by TAILOR partners [Frejinger et al. 2023], which focused on strengthening the connection between the discrete operational research community and the machine learning community, to help bridge the gap between predictive and prescriptive analytics.

Task 4.4: Learning and Reasoning for Perception, Spatial Reasoning and Vision

State of the art. Interfacing traditionally statistical-learning based perception and traditionally logic-based reasoning has been a long-term goal in AI research, with important applications, e.g., in robotics. The core challenge in this integration is that perception naturally has to deal with noise in the input signal, whereas reasoning is seen as an abstract process based on facts. The question is therefore at which points and how to express uncertainty, and numerous solutions have been proposed for this over the years.

With the widespread adoption of deep learning techniques, the overall output quality of vision approaches has increased tremendously across all visual tasks. Yet, statistical classification approaches are often overconfident in the sense that they may yield high output scores even if the classification result is incorrect. In order to improve upon this situation, one research direction has been to develop approaches for quantifying the estimation uncertainty [Kendall & Gal, 2017]. Another research direction has been a move towards more easily interpretable, explainable AI decisions, e.g., through the use of counterfactual explanations [Hendricks et al. 2018]. A number of benchmark tasks has been defined on which the performance of such approaches is systematically compared, including Visual Question Answering tasks [Agrawal et al. 2015; Jabri et al. 2016].

More recently, it has been observed that large transformer-based statistical language models, such as BERT [Devlin et al. 2019] or GPT-3 [Brown et al. 2020] and their successors, are capable of providing responses to questions (or “prompts”) that often look like the effects of a reasoning process to human observers. Whether such models indeed exhibit rudimentary reasoning capabilities is a subject of intense scientific debate [Marcus and Davis, 2020; Bender et al. 2021, Piantadosi and Hill, 2022; Mitchell and Krakauer, 2023; Newport, 2023]. Whatever the answer to this question will turn out to be, their capability to derive meaning from text is undoubted, which has enticed researchers to combine such models with, e.g., visual inputs in order to imbue visual scene understanding with common-sense knowledge derived from text. Examples for such integrations include models for image captioning [Mokady et al. 2021], text-to-image synthesis [Reed et al. 2016], more recently OpenAI's DALL-E model for image generation from text captions [Ramesh et al. 2021; Ramesh et al. 2022], or image segmentation using complex textual prompts [Shindo et al. 2024].

One recent survey on advances of continual learning and optimisation in computer vision is provided in Qu et al. (2021). In contrast to batch learning where all training data is available at once, continual learning represents a family of methods that accumulate knowledge and learn continuously with data available in sequential order. Similar to the human learning process with the ability of learning, fusing, and accumulating new knowledge coming at different time steps, continual learning is considered to have high practical significance. Hence, in this survey, they present a comprehensive review of the recent progress of continual learning in computer vision tasks. In particular, the works are grouped by their representative techniques, including regularisation, knowledge distillation, memory, generative replay and parameter isolation-based techniques. For each category of these techniques, both its characteristics and applications in computer vision are presented. This survey concludes that while continual learning in image classification and segmentation is a valuable topic to be explored, successful applications of continual learning to other computer vision problems such as Visual Question Answering (VQA) are valuable as well. In VQA, a system must produce an answer to a natural language question about an image, which requires capabilities such as object detection, scene understanding, and logical reasoning. It is ultimately desirable for a practical VQA system to be adaptable to new domains and to continuously improve as more data becomes available.

Achievements w.r.t Tasks 4.4

Important achievements in regards to learning and reasoning for perception, spatial reasoning and vision are two surveys: the already mentioned survey on recent advances of continual learning in computer vision [Qu et al. 2021], where existing works are categorised by their representative techniques including regularisation, knowledge distillation, memory, generative replay and parameter isolation; and a survey on recent deep learning based human action recognition methods [Sun et al. 2022], where existing techniques are categorised based on the input data modality. For spatio-temporal and dynamic graphs, there is the important contribution of Gravina and Bacciu (2024), whose survey explores the field of representation learning for such graphs. They conduct a fair performance comparison among the most popular approaches on node and edge-level tasks, establishing a sound baseline for evaluating new architectures and approaches.

Other contributions for Task 4.4 concern human pose estimation and human activity recognition. For instance, Gong et al. (2022) learn human body/hand pose estimation from images and propose a Meta Agent Teaming Active Learning (MATAL) framework to actively select and label informative images for effective learning. It formulates the image selection procedure as a Markov Decision Process and learns an optimal sampling policy that directly maximises the performance of the pose estimator. Ma et al. (2022) perform semi-supervised human pose estimation from videos. Specifically, this paper introduces a Motion Transformer (MT) module to perform cross frame reconstruction, aiming to learn motion dynamic knowledge in videos. Besides, a reinforcement learning-based Frame Selection Agent (FSA) is designed to harness informative frame pairs on the fly to enhance the pose estimator under the cross reconstruction mechanism. Finally Li et al. (2021) report on continual human action recognition from skeleton sequences. This work proposes an Elastic Semantic Network (Else-Net) to learn new actions by decomposing human bodies into several semantic body parts. For each body part, the proposed Else-Net constructs a semantic pathway using several elastic cells learned with old actions, or explores new cells to store new knowledge.

On the topic of perception for visual reasoning, Shindo et al. (2023) contributed a framework for integrating symbolic logic with neural networks to solve visual reasoning tasks. The proposed reasoning pipeline is fully differentiable and can incorporate gradient-based learning methods. Sha et al. (2024) proposes a neurosymbolic predicate invention framework that can discover new relational concepts from visual scenes. By inventing new concepts from scenes, the resulting model can learn classification rules of complex visual patterns with less background knowledge. Shindo et al. (2024) focuses on image segmentation, proposing a new framework that can handle complex textual prompts. They combine large-scale neural networks with neuro-symbolic reasoners to conduct reasoning on the multi-modal data. The resulting model can identify target objects specified by abstract prompts, e.g. “an object that is on the table, and that is behind a mug’’ while major neural baselines fail to solve due to their limited reasoning capabilities.

References

[Adamski and Potoniec, 2023] D. M. Adamski, and J. Potoniec (2023). Reason-able Embeddings: Learning Concept Embeddings with a Transferable Neural Reasoner. In: Semantic Web 2023 [link] **- PUT**

[Agrawal et al. 2015] A. Agrawal, J. Lu, S. Antol, M. Mitchell, C. L. Zitnick, D. Batra, and D. Parikh (2015). VQA: Visual Question Answering. In ICCV 2015 [link]

[Ahmed et al. 2022] K. Ahmed, S. Teso, K.-W. Chang, G. Van den Broeck, and A. Vergari (2022). Semantic Probabilistic Layers for Neuro-Symbolic Learning. In NeurIPS [Link] **- UNITN**

[Ahmed et al. 2023] K. Ahmed, S. Teso, P. Morettin, L. Di Liello, P. Ardino, J. Gobbi, Y. Liang, E. Wang, K.-W. Chang, A. Passerini, and G. Van den Broeck (2023). Book chapter in Compendium of Neurosymbolic Artificial Intelligence. [Link, Overview] **- UNITN**

[Amayuelas et al. 2021] A. Amayuelas, S. Zhang, X. S. Rao, and C. Zhang (2021, October). Neural methods for logical reasoning over knowledge graphs. In ICLR 2021

[Azzolin et al. 2023] S. Azzolin, A. Longa, P. Barbiero, P. Liò, and A. Passerini (2023). Global Explainability of GNNs via Logic Combination of Learned Concepts. In ICLR 2023 [link] **- UNITN**

[Azzolini and Riguzzi, 2022] D. Azzolini, and F. Riguzzi (2022). Probabilistic Logic Models for the Lightning Network. In Cryptography [link] **- CINI**

[Azzolini et al. 2022a] D. Azzolini, E. Bellodi, and F. Riguzzi (2022). MAP inference in probabilistic answer set programs. In AIxIA 2022 - Advances in Artificial Intelligence [link] **- CINI**

[Azzolini et al. 2022b] D. Azzolini, E. Bellodi, and F. Riguzzi (2022). Statistical statements in probabilistic logic programming. In Logic Programming and Nonmonotonic Reasoning [link] **- CINI**

[Azzolini and Riguzzi, 2023a] D. Azzolini and F. Riguzzi (2023). Inference in probabilistic answer set programming under the credal semantics. In AIxIA 2023 - Advances in Artificial Intelligence [link] **- CINI**

[Azzolini and Riguzzi, 2023b] D. Azzolini and F. Riguzzi (2023). Lifted inference for statistical statements in probabilistic answer set programming. In International Journal of Approximate Reasoning, 2023 [link] **- CINI**

[Azzolini and Riguzzi, 2024] D. Azzolini and F. Riguzzi (2024). Inference in Probabilistic Answer Set Programs with Imprecise Probabilities via Optimization, In UAI 2024 [link] **- CINI**

[Bacciu et al. 2023a] D. Bacciu, F. Errica, A. Gravina, L. Madeddu, M. Podda, and G. Stilo (2023). Deep Graph Networks for Drug Repurposing with Multi-Protein Targets. In IEEE Transactions on Emerging Topics in Computing [link] **- UNIPI**

[Bacciu et al. 2023b] D. Bacciu, A. Conte, and F. Landolfi (2023). Generalizing Downsampling from Regular Data to Graphs. In AAAI 2023 [link] **- UNIPI**

[Barbiero et al. 2022] P. Barbiero, G. Ciravegna, F. Giannini, P. Lió, M. Gori, and S. Melacci. (2022). Entropy-based logic explanations of neural networks. In AAAI [link] **- CINI**

[Barbiero et al. 2023] P. Barbiero, G. Ciravegna, F. Giannini, M. E. Zarlenga, L. C. Magister, A. Tonda, P. Lió, F. Precioso, M. Jamnik, and G. Marra (2023). Interpretable neural-symbolic concept reasoning". In ICML 2023. [link] **- KUL, CINI**

[Bender et al. 2021] E. M. Bender, T. Gebru, A. McMillan-Major, and S. Shmitchell (2021). On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? 🦜. In FAccT'21 [link]

[Bengio et al. 2021] Y. Bengio, Y. LeCun, and G. Hinton (2021). Deep Learning for AI. In Communications of the ACM, Volume 64, Issue 7. DOI: 10.1145/3448250 [link]

[Bessiere et al. 2022] C. Bessiere, C. Carbonnel, M. C. Cooper, E. Hebrard (2022). Complexity of Minimum-Size Arc-Inconsistency Explanations. In CP 2022 [link] **- CNRS**

[Bessiere et al. 2023a] C. Bessiere, C.Carbonnel, E. Hebrard, M. Cooper (2023). Complexity of minimum-size arc-inconsistency explanations, In Constraints [link] **- CNRS**

[Bessiere et al. 2023b] C. Bessiere, C. Carbonnel, and A. Himeur (2023). Learning Constraint Networks over Unknown Constraint Languages. In IJCAI 2023 [link] **- CNRS**

[Bessiere et al. 2023c] C. Bessiere, C.Carbonnel, A. Dries, E. Hebrard, G. Katsirelos, N. Narodytska, C.-G. Quimper, K. Stergiou, D.C. Tsouros, and T. Walsh (2023). Learning Constraints through Partial Queries, In Artificial Intelligence [link] **- CNRS**

[Betti et al. 2024] A. Betti, M. Casoni, M. Gori, S. Marullo, S. Melacci, and M. Tiezzi (2024). Neural Time-Reversed Generalized Riccati Equation. In AAAI 2024: 7935-7942. [link] **- CINI**

[Bocklandt et al. 2024] S. Bocklandt, V. Derkinderen, K. Vanderstraeten, W. Pijpops, K. Jaspers, and W. Meert (2024). Pruning-based extraction of descriptions from probabilistic circuits. In arXiv:2311.13379v2 [link] **- KUL**

[Bontempelli et al. 2022] A. Bontempelli, F. Giunchiglia, A. Passerini, and S. Teso (2022). Human-in-the-loop Handling of Knowledge Drift. In Data Mining and Knowledge Discovery [link] **- UNITN**

[Boros et al. 2024] E. Boros, O. Čepek, V. Gurvich, K. Makino (2024). Recognizing distributed approval voting forms and correspondences. In Annals of Operations Research. Volume 336, pp. 2091-2110. Springer, (DOI: 10.1007/s10479-023-05430-2) [link] **- CUNI**

[Brence et al. 2023] J. Brence, J. Tanevski, J. Adams, E. Malina, S. Džeroski (2023). Surrogate models of radiative transfer codes for atmospheric trace gas retrievals from satellite observations. In Machine learning. 2023, 112, str. 1337–1363. ISSN 1573-0565, DOI: 10.1007/s10994-022-06155-2 [link] **- JSI**

[Brown et al. 2020] T. Brown, et al. (2020). Language Models are Few-Shot Learners. In NeurIPS 2020 [link]

[Buscaroli et al. 2022] R. Buscaroli, F. Chesani, G. Giuliani, D. Loreti, and P. Mello (2022). A Prolog application for reasoning on maths puzzles with diagrams. In Journal of Experimental & Theoretical Artificial Intelligence [link] **- UNIBO**

[Carbonnel 2022] C. Carbonnel (2022). On redundancy in constraint satisfaction problems. In CP 2022 [link] **- CNRS**

[Cardillo and Straccia 2022] F. A. Cardillo, U. Straccia (2022). Fuzzy OWL-Boost: Learning fuzzy concept inclusions via real-valued boosting. In Fuzzy Sets and Systems 2022 [link] **- CNR**

[Čepek, 2023] O. Čepek (2023). Non-standard knowledge representation languages. In Proceedings of the 23rd Czech-Japan Seminar on Data Analysis and Decision Making under Uncertainty (CJS 2023), Matsuyama, Japan [link] **- CUNI**

[Čepek et al. 2023] O. Čepek, J. Weigle (2023). *A faster implementation of EQ and SE queries for switch-list representations.* In Annals of Mathematics and Artificial Intelligence 2023. DOI: 10.1007/s10472-023-09915-5 [link] **- CUNI**

[Chen et al. 2022] P. Chen, Y. Wang, X. Yu, and R. Feng (2022). Qlogice: Quantum logic empowered embedding for knowledge graph completion. In Knowledge-Based Systems, 239, 107963.

[Cheng et al. 2023] K. Cheng, N. K. Amed, and Y. Sun (2023, April). Neural Compositional Rule Learning for Knowledge Graph Reasoning. In ICLR

[Cheng et al. 2020] K. Cheng, Z. Yang, M. Zhang, and Y. Sun (2020). UniKER: A unified framework for combining embedding and Horn rules for knowledge graph inference. In ICML Workshop on Graph Representation Learning and Beyond.

[Chesani et al. 2022a] F. Chesani, C. Di Francescomarino, C. Ghidini, D. Loreti, F. M. Maggi, P. Mello, M. Montali, E. Palmieri, and S. Tessaris (2022). Discovering Business Processes models expressed as DNF or CNF formulae of Declare constraints. In CILC 2022: 37th Italian Conference on Computational Logic [link] **- UNIBO**

[Chesani et al. 2022b] F. Chesani, C. Di Francescomarino, C. Ghidini, G. Grundler, D. Loreti, F. M. Maggi, P. Mello, M. Montali, and S. Tessaris (2022). Shape Your Process: Discovering Declarative Business Processes from Positive and Negative Traces Taking into Account User Preferences. In EDOC 2022 [link] **- UNIBO**

[Chesani et al. 2023] F. Chesani, C. Di Francescomarino, C. Ghidini, D. Loreti, F. M. Maggi, P. Mello, M. Montali, and S. Tessaris (2023). Binary Discovery of Declarative Business Processes with ASP Preferences. In Thirty-Seventh AAAI Conference on Artificial Intelligence. AAAI-23 Bridge Program. B5: Artificial Intelligence and Business Process Management [link] **- UNIBO**

[Chrpa et al. 2023] L. Chrpa, R. Barták, J. Vodrážka, M. Vomlelová (2023). Attributed Transition-Based Domain Control Knowledge for Domain-Independent Planning (Extended Abstract). ICDE 2023: 3879-3880 [link] **- CUNI**

[Cima et al. 2021] G. Cima, M. Console, M. Lenzerini, and A. Poggi (2021). Abstraction in Data Integration. In LICS 2021 [link] **- UOR**

[Cima et al. 2021b] G. Cima, F. Croce, and M. Lenzerini (2021). Query Definability and Its Approximations in Ontology-based Data Management. In CIKM 2021 [link] **- UOR**

[Ciravegna et al. 2021] G. Ciravegna , P. Barbiero , F. Giannini, M. Gori, P. Lió, M. Maggini, and S. Melacci (2021). Logic explained networks. In arXiv 2021 [link] **- CINI**

[Ciravegna et al. 2023a] G. Ciravegna, F. Giannini, P. Barbiero, M. Gori, P. Lio, M. Maggini, S. Melacci (2023). Learning Logic Explanations by Neural Networks. In Compendium of Neurosymbolic Artificial Intelligence. IOS Press [link] **- CINI**

[Ciravegna et al. 2023b] G. Ciravegna, P. Barbiero, F. Giannini, M. Gori, P. Lió, M. Maggini, and S. Melacci (2023). Logic explained networks. In Artificial Intelligence Journal [link] - **CINI**

[Ciravegna et al. 2023c] G. Ciravegna, P. Barbiero, F. Giannini, M. Gori, P. Lió, M. Maggini, and S. Melacci (2023). Logic explained networks. In NeSy 2023, 17th International Workshop on Neural-Symbolic Learning and Reasoning, 2023 (432-433). [link] **- CINI**

[Clifford et al. 2023] M. Clifford, J. Erskine, A. Hepburn, P. Flach, and R. Santos-Rodríguez (2023). Reconciling Training and Evaluation Objectives in Location Agnostic Surrogate Explainers. In Proceedings of the 32nd ACM International Conference on Information and Knowledge Management (CIKM '23). Association for Computing Machinery, New York, NY, USA, 3833–3837 [link] **- UNIBRIS**

[Console et al. 2021] M. Console, P. G. Kolaitis, and A. Pieris (2021). Model-theoretic Characterizations of Rule-based Ontologies. In PODS 2021 [link] **- UOR**

[Dai et al. 2021] Y. Dai, S. Wang, N. N. Xiong, and W. Guo (2021). A survey on knowledge graph embedding: Approaches, applications and benchmarks. In Electronics 2020 [link]

[De Canditiis and De Feis 2021] D. De Canditiis, and I. De Feis (2021). Anomaly Detection in Multichannel Data Using Sparse Representation in RADWT Frames. In Mathematics 2021 [link] **- CNR**

[De Filippo et al. 2022] A. De Filippo, P. Mello, and M. Milano (2022). Do You Like Dancing Robots? AI Can Tell You Why. In PAIS 2022 Volume 351 [link] **- UNIBO**

[De Raedt et al. 2020] L. De Raedt, S. Dumančić, R. Manhaeve, G. Marra (2020). From Statistical Relational to Neural-Symbolic Artificial Intelligence. In IJCAI-20 Survey track [link] **- KUL, TUD**

[De Smet et al. 2023a] L. De Smet, P. Z. Dos Martires, R. Manhaeve, G. Marra, A. Kimmig and L. De Readt (2023). Neural Probabilistic Logic Programming in Discrete-Continuous Domains. In UAI [link] - **KUL**

[De Smet et al. 2023b] L. De Smet, E. Sansone, and P. Z. Dos Martires (2023). Differentiable Sampling of Categorical Distributions using the CatLog-Derivative Trick. In NeurIPS 2023 [link] **- KUL**

[Delfosse et al. 2023] Q. Delfosse*, H. Shindo*, D. Singh Dhami, K. Kersting (2023). Interpretable and Explainable Logical Policies via Neurally Guided Symbolic Abstraction. In NeurIPS [link] **- TUDA**

[Delong et al. 2023] L. N. Delong, R. F. Mir, M. Whyte, Z, Ji, and J. D. Fleuriot (2023). Neurosymbolic AI for reasoning on graph structures: A survey. arXiv preprint arXiv, 2302.

[Devlin et al. 2019] J. Devlin, M.W. Chang, K. Lee, and K. Toutanova (2019). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In NAACL-HLT 2019 [link]

[Dhami et al. 2021] D. S. Dhami, S. Yan, G. Kunapuli, and S. Natarajan (2021). Non-Parametric Learning of Gaifman Models. In ILP 2021 [link] **- TUDA**

[Dhami et al. 2022] D. S. Dhami, S. Yan, S. Natarajan (2022). A Statistical Relational Approach to Learning Distance-based GCNs. In AAAI workshop 2022 [link] **- TUDA**

[Di Liello et al. 2020] L. Di Liello, P. Ardino, J. Gobbi, P. Morettin, S. Teso, and A. Passerini (2020). Efficient Generation of Structured Objects with Constrained Adversarial Networks. In NeurIPS 2020 [link] **- UNITN, KUL**

[Di Maio et al. 2024] C. Di Maio, A. Zugarini, F. Giannini, M. Maggini, S. Melacci (2024). Tomorrow Brings Greater Knowledge: Large Language Models Join Dynamic Temporal Knowledge Graphs. Accepted at Collas 2024 **- CINI**

[Diligenti et al. 2022] G. Ciravegna, F. Giannini, S. Melacci, M. Maggini, and M. Gori (2022). A Constraint-Based Approach to Learning and Reasoning. In AAAI 2020 [link] **- CINI**
[Diligenti et al. 2023] M. Diligenti, F. Giannini, S. Fioravanti, C. Graziani, M. Falaschi, and G. Marra (2023). Enhancing Embedding Representations of Biomedical Data using Logic Knowledge. In IJCNN 2023 [link] **- KUL, CINI**

[Dominici et al. 2024] G. Dominici, P. Barbiero, F. Giannini, M. Gjoreski, and M. Langheinrich (2024). AnyCBM: How to Turn Any Black Box into a Concept Bottleneck Model. Accepted at XAI2024. [link] **- CINI**

[Donadello et al. 2022] I. Donadello, A. Hunter, S. Teso, and M. Dragoni (2022). Machine Learning for Utility Prediction in Argument-Based Computational Persuasion. In AAAI [link] - **UNITN**

[Fadja et al. 2021] A. N. Fadja, F. Riguzzi, and E. Lamma (2021). Learning hierarchical probabilistic logic programs. In Machine Learning 2021 [link] **- CINI**

[Frejinger et al. 2023] E. Frejinger, A. Lodi, M. Lombardi, and N. Yorke-Smith (2023). Report from Dagstuhl Seminar 22431: Data-Driven Combinatorial Optimisation. In Dagstuhl reports, May 2023 [link] **- KUL, TUD**

[Galassi et al. 2021] A. Galassi, M. Lippi, and P. Torroni (2021). Investigating Logic Tensor Networks for Neural-Symbolic Argument Mining. In StarAI workshop 2021 [link] **- CINI**

[Gentili et al. 2023] E. Gentili, A. Bizzarri, D. Azzolini, R. Zese, and F. Riguzzi (2023). Regularization in probabilistic inductive logic programming. In ILP 2023 [link] **- CINI**

[Giannini et al. 2023a] F. Giannini, M. Diligenti, M. Maggini, M. Gori, G. Marra (2023). T-norms driven loss functions for machine learning". In Applied Intelligence Journal [link] **- KUL, CINI**

[Giannini et al. 2023b] F. Giannini, S. Fioravanti, O. Keskin, A. M. Lupidi, L. C. Magister, P. Lió, and P. Barbiero. "Interpretable Graph Networks Formulate Universal Algebra Conjectures." In NeurIPS 2023. [link] **- CINI**

[Giannini et al. 2024] F. Giannini, S. Fioravanti, P. Barbiero, A. Tonda, P. Liò, and E. Di Lavore (2024). Categorical Foundation of Explainable AI: A Unifying Theory. Accepted at XAI 2024 [link] **- CINI**

[Giuliani et al. 2022a] L. Giuliani, A. De Filippo, A. Borghesi (2022). Towards Intelligent Music Production: A Sample-based Approach. In CREAI 2023, Workshop on Artificial Intelligence and Creativity [link] **- UNIBO**

[Giuliani et al. 2022b] Luca Giuliani, Allegra De Filippo, Andrea Borghesi, Paola Mello, Michela Milano, "A Multi-modal Perspective for the Artistic Evaluation of Robotic Dance Performances", in CREAI 2022, Workshop on Artificial Intelligence and Creativity, 2022, pp. 84 - 93 (atti di: Workshop on Artificial Intelligence and Creativity co-located with 21st International Conference of the Italian Association for Artificial Intelligence, Udine, 28 November 2022 - 3 December 2022) - **UNIBO**

[Giuliani et al. 2023a] Giuliani L, Ballerini F, De Filippo A, Borghesi A. MusiComb: a Sample-based Approach to Music Generation Through Constraints. In 2023 IEEE 35th International Conference on Tools with Artificial Intelligence (ICTAI) 2023 Nov 6 (pp. 194-198). IEEE. **- UNIBO**

[Giuliani et al. 2023b] Giuliani L, Misino E, Lombardi M. Generalized disparate impact for configurable fairness solutions in ML. In International Conference on Machine Learning 2023 Jul 3 (pp. 11443-11458). PMLR. **- UNIBO**

[Gong et al, 2022] J. Gong, Z. Fan, Q. Ke, H. Rahmani, and J. Liu (2022). Meta agent teaming active learning for pose estimation. In CVPR 2022 [link] **- LU**

[Gravina and Bacciu, 2024] A. Gravina and D. Bacciu (2024). Deep learning for dynamic graphs: models and benchmarks. In IEEE Transactions on Neural Networks and Learning Systems 2024 [link] **- UNIPI**

[Guo et al. 2018] S. Guo, Q. Wang, L. Wang, B. Wang, and L. Guo (2018, April). Knowledge graph embedding with iterative guidance from soft rules. In AAAI (Vol. 32, No. 1).

[Hazra et al. 2024] R. Hazra, P. Z. Dos Martires, and L. De Raedt (2024). SayCanPay: Heuristic Planning with Large Language Models Using Learnable Domain Knowledge. In AAAI 2024. [link] **- KUL**

[Hendricks et al. 2018] Generating Counterfactual Explanations with Natural Language. arXiv 2018

[Jabri et al. 2016] Revisiting Visual Question Answering Baselines. ECCV 2016

[Jain et al. 2022] R. Jain, G. Ciravegna, P. Barbiero, F. Giannini, D. Buffelli, P. Lio (2022). Extending Logic Explained Networks to Text Classification. In EMNLP 2022 [link] **- CINI**

[Kautz, 2020] H. Kautz (2020). The Third AI Summer. A talk presented during the Robert S. Engelmore Memorial Lecture Award at AAAI-20. [link]

[Kendall & Gal, 2017] What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision? In NeurIPS 2017

[Keskin et al. 2023] O. Keskin, A. M. Lupidi, S. Fioravanti, L. C. Magister, P. Barbiero, P. Lio, and F. Giannini (2023). Bridging Equational Properties and Patterns on Graphs: an AI-Based Approach. In ICML 2023 Workshop: Topological, Algebraic and Geometric Learning Workshop [link] **- CINI**

[Kumar et al. 2021] Learning Mixed-Integer Linear Programs from Contextual Examples. Data Science Meets Optimization Workshop. 2021 [link] **- KUL, UNITN**

[Li et al. 2021] Else-net: Elastic semantic network for continual action recognition from skeleton data. ICCV 2021 [link] **- LU**

[Li et al. 2020] Li, W., Qi, G., & Ji, Q. (2020). Hybrid reasoning in knowledge graphs: Combing symbolic reasoning and statistical reasoning. Semantic Web, 11(1), 53-62.

[Lin et al. 2023] S. Lin, G. Behnke, S. Ondrčková, R. Barták, and P. Bercher (2023). On Total-Order HTN Plan Verification with Method Preconditions - An Extension of the CYK Parsing Algorithm. In AAAI 2023: 12041-12048 [link] **- CUNI**

[Lombardi et al. 2020] An Analysis of Regularized Approaches for Constrained Machine Learning. arXiv 2020 [link] **- UNIBO**

[Ma et al. 2022] Remote: Reinforced motion transformation network for semi-supervised 2d pose estimation in videos. AAAI 2022 [link] **- LU**

[Macková and Pilát. 2024] K. Macková, and M. Pilát (2024). ProMap: Product Mapping Datasets. In ECIR (2) 2024: 159-172 [link] **- CUNI**

[Maene et al. 2024a] J. Maene, and L. De Raedt (2024). Soft-Soft-Unification in Deep Probabilistic Logic. In NeurIPS 2023 [link] **- KUL**

[Maene et al. 2024b] J. Maene, V. Derkinderen, and L. De Raedt (2024). On the Hardness of Probabilistic Neurosymbolic Learning. In ICML 2024 [link] **- KUL**

[Maggini et al. 2024] M. Maggini, M. Tiezzi, and M. Gori (2024). A Lagrangian Framework for Learning in Graph Neural Networks. In AI in the Age of Neural Networks and Brain Computing 2024 [link] **- CINI**

[Maggio et al. 2023] Maggio A, Giuliani L, Calegari R, Lombardi M, Milano M. A geometric framework for fairness. - **UNIBO**

[Mandi et al. 2023] Jayanta Mandi, James Kotary, Senne Berden, Maxime Mulamba, Victor Bucarey, Tias Guns, Ferdinando Fioretto: Decision-Focused Learning: Foundations, State of the Art, Benchmark and Future Opportunities. CoRR abs/2307.13565 (2023) [link]

[Manhaeve et al. 2021] R. Manhaeve, S. Dumančić, A. Kimmig, T. Demeester, and L. De Raedt (2021). Neural probabilistic logic programming in DeepProblog. In Artificial Intelligence [link]

[Marconato et al. 2023a] E. Marconato, G. Bontempo, E. Ficarra, S. Calderara, A. Passerini, and S. Teso (2023). Neuro-Symbolic Continual Learning: Knowledge, Reasoning Shortcuts and Concept Rehearsal. In ICML 2023 [link] - UNITN

[Marconato et al. 2023b] E. Marconato, S. Teso, A. Vergari, and A. Passerini (2023). Not All Neuro-Symbolic Concepts Are Created Equal: Analysis and Mitigation of Reasoning Shortcuts. In NeurIPS 2023 [link] **- UNITN**

[Marcus, 2020] G. Marcus (2020). The Next Decade in AI: Four Steps Towards Robust Artificial Intelligence. In Arxiv:2002.06177 [link]

[Marcus and Davis, 2020] G. Marcus and E. Davis (2020). GPT-3, Bloviator: OpenAI’s language generator has no idea what it’s talking about. In Technology Review August 22, 2020 [link]

[Marra et al. 2021] G. Marra, M. Diligenti, F. Giannini (2021). Learning Representations for Sub-Symbolic Reasoning. In arXiv:2106.00393v2 [link] **- KUL, CINI**

[Marra et al. 2024] G. Marra, S. Dumančić, R. Manhaeve, and L. De Raedt (2024). From Statistical Relational to Neurosymbolic Artificial Intelligence: A Survey. In Artificial Intelligence [link] **- KUL, TUD**

[Massidda et al. 2024a] R. Massidda, F. Landolfi, M. Cinquini, and D. Bacciu (2024). Constraint-Free Structure Learning with Smooth Acyclic Orientations. In ICLR 2024 [link] **- UNIPI**

[Massidda et al. 2024b] R. Massidda, S. Magliacane, and D. Bacciu (2024). Learning Causal Abstractions of Linear Structural Causal Models. In UAI 2024 [link] **- UNIPI**

[McDonald et al. 2024] T. McDonald, C. Tsay, A. M. Schweidtmann, and N. Yorke-Smith (2024). Mixed-integer Optimisation of Graph Neural Networks for Computer-Aided Molecular Design. In Computers & Chemical Engineering [link] **- TUD**

[Mechqrane et al. 2024] Y. Mechqrane, C. Bessiere, I. Elabbassi (2024). Using Large Language Models to Improve Query-based Constraint Acquisition. In IJCAI 2024 **- CNRS**

[Misino et al. 2022] E. Misino, G. Marra, and E. Sansone (2022). VAEL: Bridging Variational Autoencoders and Probabilistic Logic Programming. In NeurIPS 2022 [link] **- KUL & UNIBO**

[Mitchell and Krakauer. 2023] M. Mitchell, and D. C. Krakauer (2023). The debate over understanding in AI's large language models. In Proceedings of the National Academy of Sciences 120(13) [link]

[Mokady et al. 2021] R. Mokady, A. Hertz, and A. H. Bermano (2021). Clipcap: Clip prefix for image captioning. In arXiv:2111.09734. [link]

[Morettin et al. 2021] P. Morettin, P. Z. Dos Martires, S. Kolb, and A. Passerini (2021). Hybrid probabilistic inference with logical and algebraic constraints: a survey. In IJCAI Survey Track 2021 [link] **- KUL, UNITN**

[Mulamba et al. 2021] M. Mulamba, J. Mandi, M. Diligenti, M. Lombardi, V. Bucarey, and T. Guns (2021). Contrastive Losses and Solution Caching for Predict-and-Optimize. In IJCAI 2021 [link] **- KUL**

[Newport, 2023] C. Newport (April 2023). What Kind of Mind Does ChatGPT Have?. In The New Yorker. Retrieved 12 June 2023 [link]

[Nickel et al. 2016] M. Nickel, K. Murphy, V. Tresp, and E. Gabrilovich (2016). A Review of Relational Machine Learning for Knowledge Graphs. In IEEE 2016 [link]

[Noé et al. 2024] P.-G. Noé, M. Perello-Nieto, J.-F. Bonastre and P. Flach (2024). Explaining a probabilistic prediction on the simplex with Shapley compositions. In ECAI 2024 **- UNIBRIS**

[Ondrčková et al. 2023a] S. Ondrčková, R. Barták, P. Bercher, and G. Behnke (2023). Lessons Learned from the CYK Algorithm for Parsing-based Verification of Hierarchical Plans. In FLAIRS 2023 [link] **- CUNI**

[Ondrčková et al. 2023b] S. Ondrčková, R. Barták, P. Bercher, and G. Behnke (2023). On the Impact of Grounding on HTN Plan Verification via Parsing. In ICAART (3) 2023: 92-99 [link] **- CUNI**

[Ondrčková et al. 2023c] S. Ondrčková, and R. Barták (2023). On Semantics of Hierarchical Planning Domain Models with Decomposition Constraints and Empty Methods. In ICTAI 2023: 349-353 [link] **- CUNI**

[Ondrčková et al. 2024] S. Ondrčková, K. Pantůčková, and R. Barták (2024). Handling Empty Decomposition Methods in Hierarchical Planning, In FLAIRS 2024, Frontiers in Artificial Intelligence and Applications, 2023, Vol. 37 [link] **- CUNI**

[Osojnik et al. 2023] A. Osojnik, P. Panov, and S. Džeroski (2023). iSOUP-SymRF : symbolic feature ranking with random forests in online multi-target regression. In Discovery science: 26th international conference, DS 2023, Porto, Portugal, October 9-11, 2023 : proceedings. Cham: Springer, cop. 2023. Str. 48-63, ilustr. Lecture notes in computer science (Internet), Lecture notes in artificial intelligence, 14276. ISBN 978-3-031-45275-8. ISSN 1611-3349, DOI: 10.1007/978-3-031-45275-8_4 [link] **- JSI**

[Pantůčková et al. 2023a] K. Pantůčková, and R. Barták: Using Earley Parser for Recognizing Totally Ordered Hierarchical Plans. ECAI 2023: 1819-1826 [link] **- CUNI**

[Pantůčková et al. 2023b] K. Pantůčková, S. Ondrčková, and R. Barták (2023). Parsing-Based Recognition of Hierarchical Plans Using the Grammar Constraint. In FLAIRS 2023 [link] **- CUNI**

[Pantůčková et al. 2024] K. Pantůčková, S. Ondrčková, and R. Barták (2024). Using Earley Parser for Verification of Totally Ordered Hierarchical Plans. In FLAIRS 2024, Frontiers in Artificial Intelligence and Applications, 2024, Vol. 37 [link] **- CUNI**

[Pellegrini et al. 2021] G. Pellegrini, A. Tibo, P. Frasconi, A. Passerini, and M. Jaeger (2021). Learning Aggregation Functions. In IJCAI 2021 [link] **- UNITN**

[Peng et al. 2021] H. Peng, H. Huang, L. Xu, T. Li, J. Liu, H. Rahmani, Q. Ke, Z. Guo, C. Wu, R. Li, M. Ye, J. Wang, J. Zhang, Y. Liu, T. He, F. Zhang, X. Liu, and T. Lin (2021). The Multi-Modal Video Reasoning and Analyzing Competition. In ICCV Workshop 2021 [link] **- LU**

[Petković et al. 2021] M. Petković, D. Kocev, B. Škrlj, and S. Džeroski (2021). Ensemble- and Distance- Based Feature Ranking for Unsupervised Learning. In Int. J. Intell. Syst. 2021 [link] **- JSI**

[Petković et al. 2022a] M. Petković, D. Kocev, B. Škrlj, and S. Džeroski (2022). Feature ranking for semi-supervised learning. In MLJ 2022 [link] **- JSI**

[Petković et al, 2022b] M. Petković, M. Ceci, G. Pio, B. Škrlj, K. Kersting, S. Džeroski (2022). Relational tree ensembles and feature rankings. In Knowledge-based systems. vol. 251, str. 109254-1-109254-19. ISSN 0950-7051. DOI: 10.1016/j.knosys.2022.109254 [link] **- JSI**

[Petković et al. 2023] M. Petković, J. Levatić, K. Dragi, M. Breskvar, S. Džeroski (2023). CLUSplus: a decision tree-based framework for predicting structured outputs. In SoftwareX.

Dec. 2023, 101526, vol. 24, str. 1-7, ilustr. ISSN 2352-7110. DOI: 10.1016/j.softx.2023.101526. [link] **- JSI**

[Pfanschilling et al. 2022] V. Pfanschilling, H. Shindo, D. Singh Dhami, K. Kersting (2022). Sum-Product Loop Programming: From Probabilistic Circuits to Loop Programming. In KR [link] **- TUDA**

[Piantadosi and Hill. 2022] S. T. Piantadosi, and F. Hill (2022). Meaning without reference in large language models. In Arxiv:2208.02957 [link]

[Potoniec, 2023a] J. Potoniec (2023). Is the Proof Length a Good Indicator of Hardness for Reason-able Embeddings? In NeSy 2023, Proceedings of the 17th International Workshop on Neural-Symbolic Learning and Reasoning [link] **- PUT**

[Potoniec, 2023b] J. Potoniec (2023). Are Quantified Boolean Formulas Hard for Reason-Able Embeddings? In Progress in Polish Artificial Intelligence Research 4 pp. 343-348 2023 [link] **- PUT**

[Poulis et al. 2024] A. Poulis, E. Tsalapati, M. Koubarakis (2024). Transformers in the Service of Description Logic-based Contexts. Under review at an international AI conference. [link] **- UOA**

[Pourkhajouei et al. 2023] S. Pourkhajouei, F. Toffano, P. Viappiani and N. Wilson (2023). An Efficient Non-Bayesian Approach for Interactive Preference Elicitation Under Noisy Preference Models. In Proceedings of the 17th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU 2023), Arras (France), September 2023. [link] **- UCC**

[Prestwich and Wilson, 2024] S. Prestwich, and N. Wilson (2024). A Statistical Approach to Learning Constraints. In International Journal of Approximate Reasoning 2024 [link] **- UCC**

[Qu et al, 2021] H. Qu, H. Rahmani, L. Xu, B. Williams, and Jun Liu (2021). Recent Advances of Continual Learning in Computer Vision: An Overview. In arXiv 2021 [link] **- LU**

[Ramesh et al. 2021] A. Ramesh, M. Pavlov, G. Goh, S. Gray, C. Voss, A. Radford, M. Chen, I. Sutskever (2021). Zero-Shot Text-to-Image Generation. In ICML [**link**]

[Ramesh et al. 2022] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen (2022). Hierarchical Text-Conditional Image Generation with CLIP Latents. In arXiv 2022 [link]

[Reed et al. 2016] Generative adversarial text to image synthesis ICML 2016

[Roy et al. 2022] B. Roy, Tomaž Stepišnik, C. Vens, S. Džeroski (2022). Survival analysis with semi-supervised predictive clustering trees. In Computers in Biology and Medicine. 2022, vol. 141, str. 105001-1-105001-19. ISSN 0010-4825. DOI: 10.1016/j.compbiomed.2021.105001 [link] **- JSI**

[Sabbatini et al. 2022a] Sabbatini F, and Calegari R. Evaluation metrics for symbolic knowledge extracted from machine learning black boxes: A discussion paper. arXiv preprint arXiv:2211.00238. 2022 Nov 1. [link] - **UNIBO**

[Sabbatini et al. 2022b] Sabbatini F, and Calegari R. Symbolic knowledge extraction from opaque machine learning predictors: GridREx & PEDRO. In Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning 2022 Jul 30 (Vol. 19, No. 1, pp. 554-563) [link] - **UNIBO**

[Sabbatini et al. 2023a] Sabbatini F, and Calegari R. ExACT explainable clustering: unravelling the intricacies of cluster formation. In Proceedings of the 2nd International Workshop on Knowledge Diversity, KoDis 2023 (pp. 2-8) [link] - **UNIBO**

[Sabbatini et al. 2023b] Sabbatini F, and Calegari R. Bottom-up and top-down workflows for hypercube-and clustering-based knowledge extractors. In International Workshop on Explainable, Transparent Autonomous Agents and Multi-Agent Systems 2023 May 29 (pp. 116-129). Cham: Springer Nature Switzerland [link] - **UNIBO**

[Sabbatini et al. 2023c] Sabbatini F, Ciatto G, Calegari R, and Omicini A. Towards a unified model for symbolic knowledge extraction with hypercube-based methods, Intelligenza Artificiale 17 (2023) 63–75 [link] - **UNIBO**

[Sabbatini et al. 2023d] Sabbatini F, and Calegari R. Symbolic Knowledge-Extraction Evaluation Metrics: The FiRe Score. In ECAI 2023 2023 (pp. 2033-2040). IOS Press [link] - **UNIBO**

[Sabbatini et al. 2023e] Sabbatini F, and Calegari R. Unlocking Insights and Trust: The Value of Explainable Clustering Algorithms for Cognitive Agents. In CEUR WORKSHOP PROCEEDINGS 2023 (Vol. 3579, pp. 232-245). CEUR-WS [link] - **UNIBO**

[Sabbatini et al. 2024] Sabbatini F, and Calegari R. On the evaluation of the symbolic knowledge extracted from black boxes. AI and Ethics. 2024 Jan 9:1-0 [link] - **UNIBO**

[Sansone, 2022] E. Sansone (2022). LSB: Local Self-Balancing MCMC in Discrete Spaces. In ICML 2022 [link] **- KUL**

[Scarselli et al. 2009] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini (2009). The Graph Neural Network Model. In IEEE Transactions on Neural Networks [link]

[Scavuzzo et al. 2022] L. Scavuzzo, F. Y. Chen, D. Chételat, M. Gasse, A. Lodi, N. Yorke-Smith, and K. Aardel (2022). Learning to Branch with Tree MDPs. In NeurIPS 2022 [link] **- TUD**

[Sha et al. 2024] J. Sha, H. Shindo, K. Kersting, and D. Singh Dhami (2024). Neural-Symbolic Predicate Invention: Learning Relational Concepts from Visual Scenes. In Neurosymbolic Artificial Intelligence Journal (NAIJ), 2024 [link] **- TUDA**

[Shindo et al. 2022] H. Shindo, D. Singh Dhami, and K. Kersting (2022). Neuro-Symbolic Forward Reasoning. In arXiv 2022 [link] **- TUDA**

[Shindo et al. 2023] H. Shindo, V. Pfanschilling, D. Singh Dhami, and K. Kersting (2023). alphaILP: Thinking Visual Scenes as Differentiable Logic Programs. Machine Learning Journal (MLJ) vol. 112, pp. 1465–1497 [link] **- TUDA**

[Shindo et al. 2024] H. Shindo, M. Brack, G. Sudhakaran, D. Singh Dhami, P. Schramowski, K. Kersting (2024). DeiSAM: Segment Anything with Deictic Prompting. In AAAI Workshop on Neuro-Symbolic Learning and Reasoning in the Era of Large Language Models (NuCLeaR@AAAI) [link] **- TUDA**

[Silva Filho et al. 2023] T. Silva Filho, H. Song, M. Perello-Nieto, R. Santos-Rodriguez, M. Kull, and P. Flach (2023). Classifier calibration: a survey on how to assess and improve predicted class probabilities. In Machine Learning 112, 3211–3260 [link] **- UNIBRIS**

[Silvestri et al. 2022] Hybrid Offline/Online Optimization for Energy Management via Reinforcement Learning. CPAIOR 2022 [link] **- UNIBO**

[Silvestri et al. 2023] Mattia Silvestri, Senne Berden, Jayanta Mandi, Ali Irfan Mahmutogullari, Maxime Mulamba, Allegra De Filippo, Tias Guns, Michele Lombardi: Score Function Gradient Estimation to Widen the Applicability of Decision-Focused Learning. CoRR abs/2307.05213 (2023) [link] **- UNIBO**

[Skrlj et al. 2022] ReliefE: feature ranking in high-dimensional spaces via manifold embeddings. Machine Learning 2022 [link] **- JSI**

[Skryagin et al. 2022] SLASH: Embracing Probabilistic Circuits into Neural Answer Set Programming. KR 2022 [link] **- TUDA**

[Sokol and Flach, 2024] K. Sokol, and P. Flach (2024). Interpretable representations in explainable AI: from theory to practice. In Data Mining and Knowledge Discovery [link] **- UNIBRIS**

[Spallitta et al. 2022] G. Spallitta, G. Masina, P. Morettin, A. Passerini, and R. Sebastiani (2022). SMT-based Weighted Model Integration with Structure Awareness. In UAI 2022 [link] **- UNITN, KUL**

[Spallitta et al. 2024] G. Spallitta, G. Masina, P. Morettin, A. Passerini, and R. Sebastiani (2024). Enhancing SMT-based Weighted Model Integration by Structure Awareness. In Artificial Intelligence 2024 [link] **- UNITN**

[Stevanoski et al. 2023] B. Stevanoski, D. Kocev, A. Osojnik, I. Dimitrovski, and S. Džeroski (2023). Data stream mining for predicting the thermal power consumption of the Mars Express spacecraft. In Acta astronautica. Sep. 2023, vol. 210, str. 411-427, ilustr. ISSN 1879-2030, DOI: 10.1016/j.actaastro.2023.04.037 [link] **- JSI**

[Sun et al. 2022] Human Action Recognition from Various Data Modalities: A Review. IEEE Transactions on Pattern Analysis and Machine Intelligence 2022 [link] **- LU**

[Teso and Vergari, 2022] S. Teso, and A. Vergari (2022). Efficient and Reliable Probabilistic Interactive Learning with Structured Outputs. AAAI Workshop 2022 [link] **- UNITN**

[Teso et al. 2022] S. Teso, L. Bliek, A. Borghesi, M. Lombardi, N. Yorke-Smith, T. Guns, and A. Passerini (2022). Machine Learning for Combinatorial Optimisation of Partially-Specified Problems: Regret Minimisation as a Unifying Lens. In CoRR abs/2205.10157 [link] **- UNITN, KUL, TU Delft, UNIBO**

[Valenti et al. 2024] A. Valenti, D. Bacciu, A. Vergari (2024). ChemAlgebra: Algebraic Reasoning on Chemical Reactions. In IEEE WCCI 2024 [link] **- UNIPI**

[van Driel, 2021] R. van Driel, E. Demirović, and N. Yorke-Smith (2021). Learning Variable Activity Initialisation for Lazy Clause Generation Solvers. In CPAIOR 2021 [link] **- TU Delft**

[Ventola et al. ILP 2021] F. Ventola, D. S. Dhami, and K. Kersting (2021). Generative Clausal Networks: Relational Decision Trees as Probabilistic Circuits. In ILP 2021 [link] **- TUDA**

[Venturato et al. 2024] G. Venturato, V. Derkinderen, P. Z. Dos Martires, and L. De Raedt (2024). Inference and Learning in Dynamic Decision Networks Using Knowledge Compilation. In AAAI 2024 [link] **- KUL**

[Vergari et al. 2021] A. Vergari, Y. Choi, A. Liu, S. Teso, and G. Van den Broeck (2021). A Compositional Atlas of Tractable Circuit Operations: From Simple Transformations to Complex Information-theoretic Queries. In NeurIPS 2021 [link] **- UNITN**

[Vermeulen et al. 2023] A. Vermeulen, R. Manhaeve, and G. Marra (2023). An Experimental Overview of Neural-Symbolic Systems. In ILP 2023 [link] **- KUL**

[Verreet et al. 2023a] V. Verreet, L. De Raedt, and J. Bekker (2023). Modeling PU learning using probabilistic logic programming. In Machine Learning 2023 [link] **- KUL**

[Verreet et al. 2023b] V. Verreet, L. De Smet, and E. Sansone. EXPLAIN, AGREE and LEARN (2023). A Recipe for Scalable Neural-Symbolic Learning. In ICML Workshop: Knowledge and Logical Reasoning in the Era of Data-driven Learning 2023 [link] **- KUL**

[Vidal 2022] G. Vidal (2022). Explanations as Programs in Probabilistic Logic Programming. In FLOPS 2022 [link] **- UPV**

[Vidal 2024] G. Vidal (2024). Explaining Explanations in Probabilistic Logic Programming. In Submitted for publication [link] **- UPV**

[Wan et al. 2024] Z. Wan Z, CK. Liu, H. Yang, C. Li, H. You, Y. Fu, C. Wan, T. Krishna, Y. Lin, A. Raychowdhury (2024). Towards Cognitive AI Systems: a Survey and Prospective on Neuro-Symbolic AI. In arXiv preprint arXiv:2401.01040. 2024 Jan 2. [link]

[Wang et al. 2017] Q. Wang, Z. Mao, B. Wang, and L. Guo (2017). Knowledge graph embedding: A survey of approaches and applications. In IEEE Transactions on Knowledge and Data Engineering [link]

[Wilson, 2023] N. Wilson (2023). On the Variation of Max Regret with respect to the Scaling of the Objectives, In ECAI 2023, October 2023 [link] **- UCC**

[Winters et al. 2022] T. Winters, G. Marra, R. Manhaeve, and L. De Raedt (2022). Deepstochlog: Neural stochastic logic programming. In AAAI 2022 [link] **- KUL**

[Yamagata et al. 2024] T. Yamagata, T. Oberkofler, T. Kaufmann, V. Bengs, E. Hüllermeier, and R. Santos-Rodríguez (2024). Relatively Rational: Learning Utilities and Rationalities Jointly from Pairwise Preferences. In ICML 2024, workshop on Models of Human Feedback for AI Alignment **- UNIBRIS**

[Yang et al. 2022a] W. C. Yang, J. F. Raskin, and L. De Raedt (2022). Lifted Model Checking for Relational MDPs. In Machine Learning, 111(10), 3797-3838 [link] **- KUL**

[Yang et al. 2022b] W. C. Yang, A. Jain, L. De Raedt, and W. Meert (2022). Parameter Learning in ProbLog With Annotated Disjunctions. In IDA 2022 [link] **- KUL**

[Yang et al. 2023] W. C. Yang, G. Marra, G. Rens, and L. De Raedt (2023). Safe Reinforcement Learning via Probabilistic Logic Shields. In IJCAI 2023 [link] - **KUL**

[Zarlenga et al. 2022] M. E. Zarlenga, P. Barbiero, G. Ciravegna, G. Marra, F. Giannini, M. Diligenti, Z. Shams, F. Precioso, S. Melacci, A. Weller, P. Lio, and M. Jamnik (2022). Concept embedding models: Beyond the accuracy-explainability trade-off. In NeurIPS 2022 [link] **- CINI, KUL**

[Zecevic et al. 2021] M. Zečević, D. Dhami, A. Karanam, S. Natarajan, and K. Kersting (2021). Interventional sum-product networks: Causal inference with tractable probabilistic models. In NeurIPS 2021 [link] **- TUDA**

[Zeinalipour and Gori, 2023] K. Zeinalipour, and M. Gori (2023). Graph Neural Networks for Topological Feature Extraction in ECG Classification. In Applications of Artificial Intelligence and Neural Systems to Data Science. Singapore: Springer Nature Singapore, 2023. 17-27. [link] **- CINI**

[Zeinalipour et al. 2023a] K. Zeinalipour, T Laquinta; G. Angelini, L. Rigutini, M. Maggini, and M. Gori (2023). Building bridges of knowledge: Innovating education with automated crossword generation. In ICMLA [link] **- CINI**

[Zeinalipour et al. 2023b] K. Zeinalipour, M. Z. Saad, M. Maggini, and M. Gori (2023). ArabIcros: AI-Powered Arabic Crossword Puzzle Generation for Educational Applications. In Proceedings of ArabicNLP 2023. [link] **- CINI**

[Zeinalipour et al. 2024a] K. Zeinalipour, Y. G. Keptiğ, M. Maggini, L. Rigutini, and M. Gori (2024). A Turkish Educational Crossword Puzzle. In AIED 2024 [link] **- CINI**

[Zeinalipour et al. 2024b] K. Zeinalipour, Y. G. Keptiğ, M. Maggini, and M. Gori (2024). Automating Turkish Educational Quiz Generation Using Large Language Models. In ISPR 2024 [link] **- CINI**

[Zhang et al. 2021] J. Zhang, B. Chen, L. Zhang, X. Ke, and H. Ding (2021). Neural, symbolic and neural-symbolic reasoning on knowledge graphs. AI Open, 2, 14-35.

Appendix 1: Workshops and Activities Organised by WP 4

We provide a list of the workshop and activities organised by WP 4. More detailed information about each event is provided in the subsequent appendices.

1. WP 4 Kickoff Jan 4, 21 - go to appendix
   - What is trustworthy AI ? (with Fosca Giannotti WP3)
   - Introduction to different tasks & Discussions on collaborations
2. WP 4 Task 4.2. on Learning and Optimisation March 22, 21 - go to appendix
   - Empirical Model Learning talk (Lombardi & Milano)
   - Short presentations other groups & discussion
3. Taskforce on WP 4 Challenges and Benchmarks April 30, 21
   - Needed to bridge the gap between different approaches
   - Set of specific benchmarks for different tasks are being assembled
4. WP 4 Task 4.3 on Embeddings, knowledge graphs & ontologies June 6, 21 - go to appendix
   - Six talks on these topics & Discussions
   - Link to AI4EU platform (intro by Alessandro Saffiotti)
5. WP 4 Task 4.1. on Integrated representations July 7, 21 - go to appendix
   - Keynote by Richard Evans (DeepMind) - The apperception engine
   - Connection to Industry and Roadmapping Activity (Marc Schoenauer)
   - 13 Posters
6. WP 4 Poster Session in TAILOR Conference Sept 21, 21
   - about 15 posters
7. Connecting to Research Camp "Automating Data Science" WP4-7, ERC SYNTH project of Luc De Raedt, 2-4 Feb, 2022 - go to appendix
8. WP 4 What Are the Next Measurable Challenges in AI? March 3, 22 - go to appendix
   - Focus on Datasets and Benchmarks & Deliverable 4.3 Integrated learning, reasoning and optimisation in practice v.1 [M22] —including Panel & Keynote Joaquin Vanschoren
   - Several intermediate task force meetings
9. WP 4 Task 4.4 Learning and reasoning for perception, spatial reasoning, and vision May 30, 22 + Deliverables - go to appendix
10. Summer School on Artificial Intelligence July 24-28, 2023 - go to appendix
    - Tutorial by Roman Barták (CUNI) on "Foundations of Automated Planning"
11. Workshop Boolean Seminar Liblice September 24-28, 23 - go to appendix
    - Organised by Ondřej Čepek (CUNI) with support from TAILOR
12. WP 4 Workshop on "WebCrow Project" Feb 24, 2023 - go to appendix
    - Presentation of the WebCrow Benchmark for TAILOR.
13. WP 4 Workshop on "Benchmarks for Neural-Symbolic AI" July 3, 2023 - go to appendix
    - Seven talks on Neural-Symbolic benchmarks within different domains.
    - The workshop has been collocated at the NeSy 2023 Conference.
14. WP4 Workshop on "Neuro-Symbolic Benchmarks" March 11, 2024 - go to appendix
15. WP4 Workshop on "Neuro-Symbolic AI" March 20, 2024 - go to appendix
    - Invited talk: Floris Geerts
16. Joint WP 4-5 Workshop on the Integration of Large Language Models and Reasoning", April 19, 2024 - go to appendix

- Exploration of synergies between large language models and reasoning.
- Invited speakers: Guy Van den Broeck (UCLA) and Scott Sanner (University of Toronto).

17. Workshop on Human-Interpretable AI at KDD conference, to be held on August 26, 2024 - go to appendix

Appendix 2: WP 4 Kickoff 4/1/2021

Program
13:45-14:00 - Doors open
14:00-14:10 - Introduction and overview
14:10-14:20 - Introductory talk: What is trustworthy AI? - Fosca Giannotti
14:20-14:35 - Introduction: One minute per partner
14:35-15:05 - Tasks

- 4.1: Integrated representations for learning and reasoning
- 4.2: Integrated approaches to learning and optimisation
- 4.3: Learning and reasoning with embeddings, knowledge graphs, and ontologies
- 4.4: Learning and reasoning for perception, spatial reasoning, and vision
- 4.5: Synergies industry, challenges, roadmap
- 4.6: Fostering the AI scientific community

15:05 - 15:10 - Explanation about focus groups
15:10 - 15:25 - Break
15:25 - 16:15 - Focus groups on tasks 4.1 - 4.4 (what is the problem and how to measure progress?)
16:15 - 17:00 - Presentation and discussion of each focus group on tasks 4.1 - 4.4
17:00 - 17:15 - Break
17:15 - 18:00 - Open discussion on organisation of WP activities (workshops, challenges, site, discussion groups, Tasks 4.5 and 4.6)

Appendix 3: WP 4 Task 4.2 on Learning and Optimisation 22/3/2021

Abstract. Empirical Model Learning (EML) is a technique to enable Combinatorial Optimization and decision making over complex real-world systems. The approach is based on a two-fold mechanism: 1) using a Machine Learning (ML) model to approximate the input/output behaviour of a system, and 2) embedding such an Empirical Model into a Combinatorial Optimization model. The EML approach has been employed with a measure of success to the application of Combinatorial Optimization to systems that are too complicated for an expert-designed, hand-crafted model, and to the generation of adversarial examples and certification of ML models. Specific use cases include: thermal-aware workload dispatching, transprecision computing, hardware dimensioning and algorithm configuration, epidemiological model, and NN verification. However, the method has potentially much broader applicability, such as providing an alternative approach to deal with uncertainty in optimization, enabling the definition of hierarchies of optimization systems (each one approximated via ML), black-box optimization, and parameter tuning. Research in these directions has been so far constrained by limited resources and by some notable, open, scientific problems. The goal of the workshop will be to present the expertise accumulated at UniBo on EML topics, highlight outstanding issue, promising research directions, and defining concrete steps for cooperation and advancement

Program

10:00 - 10:10: Welcome + Introduction (Luc de Raedt)
10:10 - 10:55: Talk (Michela Milano, Michele Lombardi, Andrea Borghesi)

- Group presentation
- Empirical Model Learning (the problem, application/success stories)
- Open Issues

10:55-11:00 - Break
11:00-11:45 - Proposals for concrete ideas around the task
11:45-12:30 - Follow-up discussion & Collaboration definition

Appendix 4: WP 4 Task 4.3 on Embeddings, Knowledge graphs & Ontologies 6/6/2021

Abstract. Integrating Learning and Reasoning is a fundamental problem in AI, especially in application domains dealing with relational data, such as knowledge graphs and ontologies. In particular, the notion of embedding may play a crucial role to encode relational knowledge in a latent space and to provide a more flexible representation to perform learning and reasoning. During this workshop, some relevant models and methods employing different kinds of reasoning mechanisms will be presented and discussed. In particular, the goal will be to present some of the main research activities of the CINI group in order to outline possible collaboration and research directions around Task 4.3. Moreover, thanks to the participation of Task 7.4 of the AI4EU project, it will be discussed how newly developed assets for the integration of learning and reasoning might be published on the AI4EU platform.

Program
09:30 - 09:40 – Welcome and Introduction (Marco Gori, UniSi)
09:40 - 10:25 – First Talk Session (Chair Marco Lippi)

- KENN: Knowledge Enhanced Neural Networks (Alessandro Daniele, FBK)
- Learning Representation for Sub-Symbolic Reasoning (Francesco Giannini, UniSi)
- Empirical Model Learning: embedding ML models in declarative optimization model (Michele Lombardi, UniBo)

10:25 - 10:50 – How you can publish your work on the AI4EU platform (Alessandro Saffiotti, ORU & Peter Schuller, TUW)
10:50 - 11:05 – Coffee break
11:05 - 11:50 – Second Talk Session (Chair Francesco Giannini)

- Structure Learning of Probabilistic Logic Programs (Fabrizio Riguzzi, UniFe)
- Towards Explainable Autonomous Development (Marco Lippi, UniMoRe)
- Online Learning of Planning Domain Representations from Sensor Data (Alfonso Gerevini & Leonardo Lamanna, UniBs)

11:50 - 12:25 – Open issues & Proposals around the task
12:25 - 13:00 – Follow-up discussion & Collaboration definition

Appendix 5: WP 4 Task 4.1 on Integrated Representations 21/7/2021

Invited Talk (Richard Evans). This talk attempts to answer a central question in unsupervised learning: what does it mean to “make sense” of a sensory sequence? In our formalisation, making sense involves constructing a symbolic causal theory that both explains the sensory sequence and also satisfies a set of unity conditions. The unity conditions insist that the constituents of the causal theory – objects, properties, and laws – must be integrated into a coherent whole. On our account, making sense of sensory input is a type of program synthesis, but it is *unsupervised* program synthesis. I will show how our system makes sense of a variety of sensory sequences, including rhythmic sequences, sequence induction IQ tasks, and occlusion tasks. It is noteworthy that our system is able to achieve human-level performance on these IQ tasks, even though it was not designed to solve those particular tasks. In the second half I will describe our neuro-symbolic framework for distilling interpretable theories out of streams of raw, unprocessed sensory experience. First, we extend the definition of the apperception task to include ambiguous (but still symbolic) input: sequences of sets of disjunctions. Next, we use a neural network to map raw sensory input to disjunctive input. Our binary neural network is encoded as a logic program, so the weights of the network and the rules of the theory can be solved jointly as a single SAT problem. This way, we are able to jointly learn how to perceive (mapping raw sensory information to concepts) and apperceive (combining concepts into declarative rules).

Program
13:15 - 13:30 Doors open
13.30 - 13.40 Introduction (Luc De Raedt)
13.40 - 14.30 Keynote “The Apperception Engine” (Richard Evans - DeepMind)
14.30 - 15.00 Presentation “RoadMap TAILOR” (Marc Schoenauer, Michela Milano)
15:00 - 15:15 Break
15.15 - 15.45 Breakout room
15.45 - 16.00 Wrapping up
16.00 - 16.30 Poster Spotlights
16.30 - 16.45 Virtual Coffee
16.45 - 18.00 Poster Session
18.00 - 18.15 Next Steps

Appendix 6: Connecting to Research Camp "Automating Data Science" WP 4 - WP 7 (ERC SYNTH project of Luc de Raedt) 2-4/2/2022

Abstract. Data analysis is a difficult process that requires a skilled data scientist. A typical analysis requires many different steps: Selecting the right subset of data, pre-processing the data into the right format (data-wrangling), determining the learning task, selecting the right algorithms, evaluating the result. The field of automated data science tries to democratise data analysis and make it more accessible to non-experts by automating these different steps as much as possible. This event is organised by the ERC AdG project SYNTH, which has been devoted to the goal of automating and democratising data science. The program spans three afternoons (February 2nd, 3rd and 4th). The first two days will feature invited keynote talks by Tijl De Bie, Holger Hoos, and Sumit Gulwani, on various aspects of automating data science. It will also feature several talks and demonstrations by the PI of SYNTH, Luc De Raedt, and team members on topics such as automating data-wrangling, learning constraints and inductive models (with probabilistic programs) to model data, as well as integrating these steps in one common framework to make predictions and find anomalies. The last half-day will consist of a hands-on workshop with the SYNTH software package as well as a poster session. We encourage all participants to submit posters of their recent work in the field of automating data science. The event is also highly relevant to related projects the Leuven ML Lab is involved in, in particular, TAILOR Network of Excellence (WP 7 on Automated AI), the Grand Challenge on "AI-Driven Data Science" of Flanders AI Program, and the iBOF Project on "Automating Data Science: the Next Frontiers". This research camp is of interest for PhD students and researchers whose research is related to automated data science. During the poster session they will also be able to present their work.

More info on website

Appendix 7: WP 4 What Are the Next Measurable Challenges in AI? 3/3/2022

Abstract. Building systems that can integrate learning, reasoning and optimization has long been a dream for artificial intelligence. One of the major challenges, within this context, is certainly to evaluate novel ideas and frameworks on appropriate benchmarks. Too often, in fact, the tasks and the datasets that are considered and proposed for experimental evaluation are tailored to some algorithms or methodologies, and limited to ad-hoc scenarios and application domains. More in general, they lack an open and wider perspective to test the considered approaches across a variety of different tasks and under different conditions, making experimental comparisons hard to obtain. Can we define a set of requirements for a challenge/benchmark that goes beyond those currently available? Can we do it with the goal of having a benchmark (or rather a benchmarking framework maybe) that meets these requirements and can still be implemented in a reasonable time? possibly building on top of existing ones?

Program
13:15 - 13:30 Doors open
13.30 - 13.40 Introduction (Luc De Raedt)

Introduction
13:15 - 13:30 Introduction & Expectations - Luc de Raedt
13:30 - 14:00 Invited Talk: Lessons Learned at NeurIPS 2021 Datasets and Benchmarks - Joaquin Vanschoren

PART I (grounding the discussion to literature)
14:00 - 14:15 Presentation Datasets/Systems Tables - Marco Lippi
14:15 - 15:30 Discussion on Tables - Working groups
15:30 - 15:45 Break

PART II (widening the perspective)
15:45 - 16:45 Panel on Limitations of Existing Benchmarks and New Challenges - Andrea Passerini

- Marco Gori
- Joaquin Vanschoren
- Kristian Kersting
- Michele Sebag
- Fosca Giannotti

16:45-18:00 Discussion on Panel - Working groups

Conclusions
18:00-18:15 What’s Next? - Luc de Raedt

Appendix 8: WP 4 Task 4.4 Learning and Reasoning for Perception, Spatial Reasoning, and Vision 30/5/2022

Abstract. The aim of the workshop is to have a global overview of all the work conducted in WP 4 so far, agree on the structure of the two deliverables and identify the future activities. The workshop opens with an invited talk by Bastian Leibe titled “Towards Sensing Human Actions at a Pixel Precision Level”.

Abstract of the talk. Computer Vision has made immense progress over the past decade, driven in large parts by major advances in (and a better understanding of) deep learning. In this talk, I will illustrate this progress by presenting examples of state-of-the-art approaches from our research in several areas of visual scene understanding, including object segmentation, tracking, human body pose estimation, and 3D semantic scene analysis. For each of those areas, I will show how deep learning approaches are currently being applied to solve visual scene understanding tasks. As the presented results will show, results of state-of-the-art vision methods are getting steadily closer to giving pixel accurate interpretations of visual scenes. This increased level of precision in delineating object boundaries has important implications on both the level of detail at which vision approaches are able to analyze a scene and on the trust one can potentially place in the results of this analysis. This is particularly relevant when considering combinations of sensing with reasoning about human actions, which we can explore together in a hopefully lively follow-up discussion.

Program

13:20 - 13:30 Intro & deliverables - Luc De Raedt

13:30 - 14:30 Invited Talk: Towards Sensing Human Actions at a Pixel Precision Level - Bastian Leibe - 30 mins + 30 mins Q/A

14:30 - 15:00 Reporting Scientific Activities for Different Tasks - Deliverable 1

- 3 minutes per task - task leaders (Task 4.1 - 4.4)
- Round table discussion around the deliverable (collection of material for the deliverable, definition of structure)

15:00 - 15:30 Reporting Scientific Activities from the Taskforce and Next Steps - Deliverable 2

- Summary on the work done by the Task force (15 mins) - Marco Lippi
- Round table discussion around the deliverable

15:30 - 15:45 Conclusions (Future activities) - Luc De Raedt

Appendix 9: Summer School on Artificial Intelligence 24-28/07/2023

About
1st European Summer School on Artificial Intelligence (ESSAI) & 20th Advanced Course on Artificial Intelligence (ACAI), organised jointly as 3rd TAILOR Summer School on Artificial Intelligence.

Tutorial Speaker: Roman Barták (CUNI)
Title: Foundations of Automated Planning
Location: Ljubljana, Slovenia
Web: http://ktiml.mff.cuni.cz/~bartak/ESSAI2023/

Course Description
Automating Planning is one of the core areas of Artificial Intelligence. It is a knowledge-based approach that exploits a formal model of the problem and uses search techniques to explore alternatives on how to achieve a goal. Thanks to formal background, automated planning contributes to trustworthy and explainable AI as the models can be used to explain and to verify the plans. Opposite to reactive techniques, automated planning reasons on the future and how an agent can influence it to achieve its own goals. Planning is hence an important capability of rational agents.

Planning is the reasoning part of acting. It deals with selecting and organising actions to achieve a given goal. The course introduces automated planning, starting with a formal logical model of planning tasks, through core planning techniques (state-space and plan-space search, planning graph and reduction-based techniques), till exploiting control knowledge and relation to scheduling and execution.

Appendix 10: Workshop Boolean Seminar Liblice 24-28/09/2023

About

The seminar is organised jointly by the Department of Theoretical Computer Science and Mathematical Logic of the Faculty of Mathematics and Physics of the Charles University in Prague and the Institute of Computer Science of the Academy of Sciences of the Czech Republic. From CUNI, Ondřej Čepek (CUNI) organised this with support from TAILOR.

Program http://clp.mff.cuni.cz/booleanseminar/index.php?page=program

Sunday, September 24

19:00 Dinner

Monday, September 25

09:00-10:30 Session 1 – 3 Speakers, Chaired by Ondřej Čepek
10:30-11:00 Coffee break
11:00-12:00 Session 2 – 1 Speaker, Chaired by Martin Milanic
12:30-13:30 Lunch
14:00-18:00 Discussions
18:30-19:00 Dinner

Tuesday, September 26

09:00-10:30 Session 3 – 3 Speakers, Chaired by Endre Boros
10:30-11:00 Coffee break
11:00-12:00 Session 4 – 2 Speakers, Chaired by Yves Crama
12:00-13:00 Lunch
13:00-18:00 Trip
18:30 Dinner

Wednesday, September 27

09:00-10:30 Session 5 – 3 Speakers, Chaired by Thomas Eiter
10:30-11:00 Coffee break
11:00-12:00 Session 6 – 2 Speakers, Chaired by Stefan Szeider
12:30-13:30 Lunch
14:00-18:00 Discussions
18:00 Dinner

Thursday, September 28

09:00 Departure

Appendix 11: WP 4 Workshop on "WebCrow Project" 24/02/2023

First Workshop about WebCrow Benchmark in TAILOR.
Organised by CINI and Expert.ai.
Link: https://sailab.diism.unisi.it/tailor-webcrow-workshop/

Objective
WebCrow is an AI software developed to solve multilingual crossword puzzles using NLP technology. NLP approaches may exploit different sources of knowledge like ontologies, knowledge graphs and the web. Then the system attempts to provide the correct answer for each textual input clue, by means of a sequence of logic and linguistic steps.

The focus of the workshop will be on highlighting some research lines connecting neural-symbolic approaches and crossword puzzle solvers and generators. In particular, we aim at presenting the WebCrow as a candidate for defining a clear benchmark for neural-symbolic systems.

Expected output:

- Discussion on new ideas for improving the Webcrow solver performances by exploiting ontologies, especially in multi-language frameworks.
- Promising directions for the construction and evaluation of new benchmarks for neural-symbolic AI on crossword puzzles.
- Setting up new collaborations to work on open research directions.

Program
10:00 - 10:10 – Welcome and Introduction (Francesco Giannini, Marco Gori, CINI)
10:10 - 10:25 – Introducing WebCrow Project (Marco Ernandes, Expert.ai)
10:25 - 11:00 – Crossword Competition (Giovanni Angelini, Expert.ai)
11:00 - 11:15 – WebCrow Architecture (Andrea Zugarini, Expert.ai)
11:15 - 11:30 – Research Directions (Kamyar Zeinalipour, Unisi, Andrea Zugarini, Expert.ai)
11:30 - 11:45 – WebCrow Benchmark: Tasks and Challenges (Marco Ernandes, Expert.ai)
11:45 - 12:00 – Open Discussion and Feedbacks (Francesco Giannini, Marco Gori, CINI)

Appendix 12: WP 4 Workshop on “Benchmarks for Neural-Symbolic AI” 03/07/2023

About

- Seven talks on Neural-Symbolic benchmarks within different domains.
- The workshop has been collocated at the NeSy 2023 Conference.
- https://sailab.diism.unisi.it/tailor-wp4-workshop-at-nesy/

Abstract
The study of Neural-Symbolic (NeSy) approaches has been a longstanding goal in the field of AI. However, the evaluation of novel ideas and frameworks in this context poses a significant challenge due to the lack of appropriate benchmarks. Existing tasks and datasets tend to exhibit a myopic approach, tailored exclusively to specific algorithms or methodologies, and limited to narrow scenarios and application domains. This limitation hinders the ability to obtain meaningful experimental comparisons and constrains the overall progress of NeSy AI. Consequently, there is an urgent need for more comprehensive and inclusive benchmarks that facilitate the testing of diverse approaches across a wide range of tasks and under varying conditions. The workshop will present some prominent examples of existing benchmarks in different sub-fields, like continual learning, NLP, knowledge graphs, and so forth, from the perspective of NeSy-based approaches. In addition, we will discuss the importance of developing new benchmarks that address some shortcomings of existing ones.

Program (chair: Francesco Giannini)

11:00 - 11:05 Opening
11:05 - 11:20 Towards a survey and unified benchmark for neural-symbolic AI by Luc De Raedt and Robin Manhaeve
11:20 - 11:30 Neuro-Symbolic Continual Learning: A Novel Task with Novel Benchmarks by Emanuele Marconato
11:30 - 11:40 Entity Linking Benchmark in Computational Food: Combining Text Mining, Logical Reasoning, and LLM Prompting by Agnieszka Lawrynowicz
11:40 - 11:50 Challenges for Neuro-Symbolic Approaches: Case Study on Legal Analytics and Argument Mining by Federico Ruggeri
11:50 - 12:00 A neuro symbolic setup for link prediction in KGEs by Michelangelo Diligenti
12:00 - 12:10 Three New Benchmarks for Neuro-Symbolic Systems by Fabrizio Riguzzi
12:10 - 12:20 V-LoL: A Diagnostic Dataset for Visual Logical Learning by Devendra Singh Dhami
12:20 - 13:00 Round Table Discussion by Michelangelo Diligenti and Francesco Giannini

Appendix 13: Workshop on Neuro-Symbolic Benchmarks, 11/03/2024

About
A workshop focused on neuro-symbolic benchmarks: insightful talks on new neuro-symbolic benchmarks, followed by a discussion on the current state of benchmarking and experimental evaluation in NeSy. The workshop was held online.
https://tailor-network.eu/events/workshop-on-neuro-symbolic-benchmarks-march-11-1200-1400-cet/

Program
12:00 - 12:05 Opening statements
12:05 - 12:25 Pietro Barbieri: Interpretable Graph Networks Formulate Universal Algebra Conjectures
12:25 - 12:45 Luca Lorello: The KANDY Benchmark: Incremental Neuro-Symbolic Learning and Reasoning with Kandinsky Patterns
12:45 - 13:05 Samuele Bortolotti: Benchmarks for Neuro-Symbolic Reasoning Shortcuts
13:05 - 13:25 Elisabetta Gentili: New Benchmarks for Neuro-Symbolic Systems
13:25 - 13:45 Eleonora Giunchiglia: ROAD-R: the Autonomous Driving Dataset with Logical Requirements
13:45 - 14:00 Discussions

Organisers
Fabrizio Riguzzi (University of Ferrara)
Luc De Raedt (KU Leuven)
Marco Lippi (University of Modena and Reggio Emilia)
Neil Yorke-Smith (TU Delft)
Alice Bizzarri (University of Ferrara)
Damiano Azzolini (University of Ferrara)
Elisabetta Gentili (University of Ferrara)
Luca Salvatore Lorello (University of Pisa)
Robin Manhaeve (KU Leuven)

Appendix 14: Workshop on Neuro-Symbolic AI, 20/03/2024

An online workshop on neuro-symbolic AI, hosting an invited talk by Floris Geerts from the Adrem Data Lab at the University of Antwerp. The talk, titled "The Expressive Power of Graph Learning," was a distinguished paper at ICLR 2022 (link).
https://tailor-network.eu/events/workshop-on-neuro-symbolic-ai-march-20-1300-1600-cet/

Program

13:00 - 13:10 Opening Remarks
13:10 - 13:30 DeepSeaProbLog by Lennert De Smet
13:30 - 14:20 Invited talk: The Expressive Power of Graph Learning by Floris Geerts
14:20 - 14:40 Interpretable Neural-Symbolic Concept Reasoning by Francesco Giannini
14:40 - 15:00 The expressive power of pooling in graph neural networks by Veronica Lachi
15:00 - 15:20 DeepSoftLog by Jaron Maene
15:20 - 15:40 Reasoning Shortcuts in Neuro-Symbolic AI: Characterization, Mitigations, and Awareness by Emanuele Marconato
15:40 Closing remarks

Abstract Invited talk (The Expressive Power of Graph Learning by Floris Geerts)
A key component of graph and relational learning methods is the computation of vector representations of the input graphs or relations. In this talk, we model this computation as queries, mapping discrete relational objects into the realm of real vectors. We then revisit the expressive power of graph learning methods from this unifying query language perspective. Here, we consider the expressive power related to separability of inputs and to the approximation power of functions.

Appendix 15: Joint WP 4-5 Workshop on the Integration of Large Language Models and Reasoning 19/04/2024

Hosted by Work Packages 4 and 5, this online workshop explores the fusion of large language models and reasoning. We were thrilled to welcome two esteemed speakers: Guy Van den Broeck (UCLA) and Scott Sanner (University of Toronto).
https://tailor-network.eu/events/workshop-on-the-integration-of-large-language-models-and-reasoning-april-19th-1300-1700-cet/

Program

13.00-13.30 Opening remarks
13.10-13.25 Refining and Improving the Reasoning Capabilities of LLMs by Paul Debjit
13.25-13.40 Reasoning over Description Logic-based Contexts with Transformers by Angelos Poulis, Eleni Tsalapati, Manolis Koubarakis
13.40-13.55 Comparing Evolutionary Methods and LLMs for Program Synthesis by Leonardo Lucio Custode, Chiara Camilla Rambaldi Migliore
13.55-14.00 Break
14.00-14.50 Symbolic AI 3.0 (S3): Rise of the LLMs by Scott Sanner (University of Toronto)
14.50-15.05 Break
15.05-15.20 Enriching interactive explanations with fuzzy temporal constraint networks by Alberto Bugarín-Diz
15.20-15.35 SayCanPay: Heuristic Planning with LLMs using Learnable Domain Knowledge by Rishi Harzra
15.35-15.50 Chatbots and LLMs for Constraint Programming: Opportunities and Challenges by Tias Guns
15.50-16.00 Break
16.00-16.50 Symbolic Reasoning for Large Language Model by Guy Van den Broeck
16.50-17.00 Closing remarks

Invited talks

Title: Symbolic AI 3.0 (S3): Rise of the LLMs, by Scot Sanner
Large Language Models (LLMs) such as ChatGPT, GPT-4, and Gemini have emerged as a revolutionary technology for natural language and visual reasoning and numerous related AI applications. I'll discuss some of my group's own work on abstract reasoning and interactive conversational systems leveraging LLMs and the game-changing realisations that I have taken away from these investigations. This talk will then discuss some general implications of the LLM era and my conjectures as to how it will shift (and has already shifted) research foci in the near future and enable levels of user-facing AI deployment that were unthinkable just two years ago.

Title: Symbolic Reasoning for Large Language Models by Guy Van den Broeck
Many expect that AI will solve society's problems by simply being more intelligent than we are. Implicit in this bullish perspective is the assumption that AI will naturally learn to reason from data: that it can form trains of thought that "make sense", similar to how a human expert might reason about a case, or more formally, how a mathematician might prove a theorem. This talk will investigate the question whether this behaviour can be learned from data, and

how we can design the next generation of AI techniques that can achieve such capabilities. It will focus on neurosymbolic reasoning for large language models, both at training and generation time, using probabilistic circuits as the architecture that bridges learning and reasoning

Appendix 16: Workshop on Human-Interpretable AI at KDD conference 26/08/2024

Workshop on Human-Interpretable AI at KDD conference, to be held on August 26, 2024. The workshop is organised as a part of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD) that will be held in Barcelona, Spain. The speakers are yet to be announced. https://human-interpretable-ai.github.io/

Introduction and Goals

As deep neural networks have become a fundamental building block in recent Artificial Intelligence (AI) systems, there has been a sharp increase in attempts to explain these notoriously opaque models. Human-Interpretable AI (HI-AI), a sub-field in explainable AI (XAI), is emerging as a promising direction where methods aim to construct explanations using representations that are aligned to high-level human-understandable concepts or symbols rather than low-level representations (e.g., pixels/saliency maps). This workshop aims to spearhead research on topics within HI-AI by:

- providing a general overview of the key aspects of HI-AI to equip all researchers with the necessary background and definitions.
- running a call for papers for researchers in fields related to HI-AI to present their works as part of our poster session and have the opportunity to include these works in workshop proceedings (details can be found in our call for papers).
- creating a space for active researchers in HI-AI to share and discuss novel ideas through invited keynote talks and contributing talks from a selected number of accepted papers.

We welcome contributions covering novel post-hoc explainability or interpretable-by-design approaches and theoretical analysis of existing works. Additionally, position contributions speculating on the future potential of this field are highly encouraged. Finally, we welcome contributions from related fields such as ethical AI, knowledge-driven machine learning, human-machine interaction, applications in medicine and industry, and analyses from regulatory experts.