TAILOR

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.1 Report

Document type (nature)	Report
Deliverable No	4.1
Work package number(s)	4
Date	26 June 2022
Responsible Beneficiary	P 5 - KU Leuven
Editor(s)	Luc De Raedt
Publicity level	Public
Short description (Please insert the text in the Description of Deliverables in the Appendix 1.)	Foundations, techniques, algorithms and tools for integrating learning, reasoning and optimisation v.1

History			
Revision	**Date**	**Modification**	**Author**
1.0	30 June 2022	Minor	Luc de Raedt

Document Review		
Reviewer	**Partner ID / Acronym**	**Date of report approval**
Umberto Straccia	CNR	28 June, 2022
Ana Paiva	IST	26 June, 2022

Table of Contents

Summary of the report

Objectives: 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. It will also study applications in two different domains. The WP is divided into four main scientific Tasks, and is connected to other WPs by two tasks.

Introduction to the Deliverable

There are two deliverables for WP 4, that are both divided into an intermediate report (v1 M22) and a final report (v2 due at the end of the project).

Deliverables

D4.1: Foundations, techniques, algorithms and tools for integrating learning, reasoning and optimisation. (report) Report on the scientific challenges tasks T4.1 & T4.2.

D4.3: Integrated learning, reasoning and optimisation in practice (report). Report on the scientific challenges tasks T4.3 & T4.4.

This TAILOR WP has largely focused on two types of meetings and workshops. In the first type, there has been an emphasis on foundations, techniques, and challenges for integrating learning, reasoning and optimization. In this type of workshop, the four scientific topics that characterize the first four tasks of WP4 within TAILOR have been covered. This has not only provided us with insight into the foundations and challenges connected to this WP, it has also delivered a number of interesting tutorials and survey papers, that have partly or fully been inspired by TAILOR and that led to novel insights and often also collaborations. The current Deliverable 4.1 starts with these results, and then outlines the other results obtained within the WP. The second type of meeting was connected to the important taskforce of WP4 around benchmarks, datasets and systems. Given the plethora of different systems, representations and datasets, it is not easy to see the general picture in this diverse landscape. Therefore, we decided to start up a taskforce that would collect existing data, systems and study and compare the abilities and performance of different systems. This forms the basis for Deliverable 4.3 and promises to result in useful insights (and publications) about this.

Thus rather than dividing the deliverables along the task dimensions T4.1 / 2 vs T4.3 / 4 we found it more appropriate to report on the foundational issues in D4.1 and focus on the results of the taskforce in D4.3 as this is related to the potential and practice of WP 4 techniques.

Organisation

The following people are responsible for the Deliverable:

Partner ID / Acronym	**Name**	**Role**	**Other**
KU Leuven	Luc De Raedt	Leader of Task 4.1	Emanuele Sansone
University of Bologna	Michela Milano	Leader of Task 4.2	-
University of Siena/CINI	Marco Gori	Leader of Task 4.3	-
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 confrontational 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 cannot really be used for complex reasoning; 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 awareness 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. 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 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. W.r.t 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 specialized workshops:

- **Neuro-Symbolic Learning Computation (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 regularize 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. Although it has been shown that embeddings can be used for simple inference tasks, it is still unclear how to combine them to support multi-step reasoning and to use them for explainable AI.
- **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 optimize, 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 view on 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, and start with Tasks 4.1. and 4.3, followed by Task 4.2.

We also list, for completeness, also the workshops organized for this WP in the Appendices.

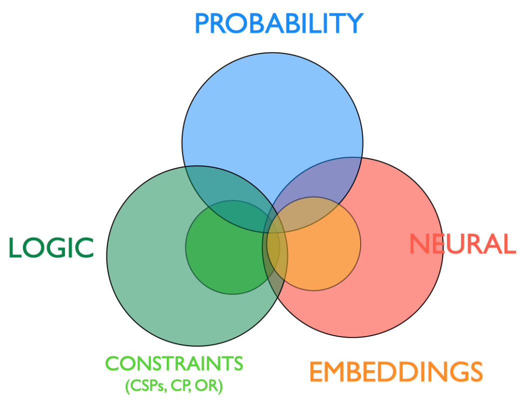

Task 4.1: Integrating Representations for Learning and Reasoning

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 in [De Raedt et al. 2020]. This survey was also the basis for tutorials at AAAI 2021, IJCAI 2021, and the TAILOR Summer School in 2021 (see link).

The survey bridges the gap between neuro-symbolic AI (NeSy), where the goal is to incorporate symbolic reasoning into neural networks with statistical relational AI (StarAI). NeSy already has a long tradition, and it has recently attracted a lot of attention from various communities. Many approaches to NeSy aim at extending neural networks with logical reasoning.

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 centered 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] to point out the similarities between these two endeavors and in this way stimulate more cross fertilization. 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 categorize 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 in 2016 [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 fastly grown in the last years relies on exploiting low-level distributional semantics for nodes and edges, that is the fundamental idea behind Knowledge Graph Embeddings (KGE). Among others, a recent survey on the state of the art of KGE is considered in [Dai et al. 2020], that extends a previous well-established survey of Wang et al. [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. split the review of existing KGE models by considering at first the embedding models only leveraging the information coming from observed triplets in the KG, and then considering classes of advanced models that utilize additional semantic information to improve the performance of the original methods. The general class of embedding models just relying on true-facts can be categorized 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 semantics information, and/or other kinds of approaches providing a different reasoning scheme, is currently one of the fundamental research direction that it is widely under investigation, and some of the main advancements in this area have been collected in some recent surveys like [Li et al. 2020] and [Zhang et al. 2021]. 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] focuses on neuro-symbolic reasoning techniques.

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 in [Maggini et al. 2022][1], 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.

StarAI. Logic and Probability, where it is useful to also further distinguish propositional from relational or first-order logical representations. Work w.r.t propositional representations has focussed on probabilistic circuits (and knowledge compilation), which are important for efficient inference and which typically are differentiable, and hence, also relevant from a neural network perspective.

There are several contributions related to probabilistic circuits. For instance, iSPN [Zecevic et al. 2021] considers the problem of learning interventional distributions (i.e., answering causal queries) with tractable probabilistic models (gated SPN); [Fadja et al. 2021] introduce hierarchical probabilistic logic programs, a restricted probabilistic logic programming language that 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, obtaining tractable inference from discriminative rule models while operating on the relational domain; and [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.

Other contributions to StarAI (and its subfield probabilistic logic programming - PLP) include [Azzolini et al. 2022] who study the problem of reasoning about the existence of a path between two nodes that can route a payment of a given size leveraging multiple models based on PLP for the lightning network; PCTL-REBEL [Yang et al. 2021] contributes a StarAI

[1]To appear in "*AI in the Age of Neural Networks and Brain Computing*", Eds. R. Kozma, C. Alippi, Y. Choe, F. Morabito 2nd edition,

approach to model checking for relational MDPs in a lifted manner, which is also connected to planning and WP 5; and [Pellegrini et al. 2021] introduce Learnable Aggregation Functions, a fully differentiable layer for set aggregation that can can approximate several extensively used aggregators (such as average, sum, maximum) as well as more complex functions like variance and skewness and smooth interpolations between aggregators.

StarAI approaches to NeSy. As indicated in the survey [De Raedt et al. 2020] already described under Task 4.1, many NeSy approaches take inspiration from StarAI or can be interpreted in these terms. For instance, Relational Neural Machines (RNM) [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] proposes a new approach for reasoning tasks taking advantage of differentiable forward-chaining using first-order logic. DeepStochLog [Winters et al. 2022] is an extension of the StarAI inspired stochastic logic programs (which are based on probabilistic definite clause grammars, ako probabilistic unification based grammar) towards neural networks. The extension is based on the notion of a neural predicate (introduced in DeepProbLog). SLASH [Skryagin et al. 2022] consists of Neural-Probabilistic Predicates (NPPs) and logic programs which are united via answer set programming. The probability estimates resulting from NPPs 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. VAEL [Misino et al. 2022] bridges the gap between variational autoencoders and probabilistic logic programs. This NeSy system offers capabilities which go beyond the traditional properties of deep generative models. Indeed, the reasoning component provides an inductive bias for the latent space of a VAE, which allows to structure the representation so that the model can generalize to previously unseen tasks, the so called task generalization. Furthermore, this hybrid generative model offers the capability to learn from smaller amounts of training data compared to purely neural-based approaches. Finally, 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.

StarAI, NeSy, Embeddings and Ontologies. The StarAI and NeSy methods are often also integrated with embeddings, knowledge graphs and/or graph neural networks. For instance, Relational Reasoning Networks (R2N) [Marra et al. 2021] perform learning and reasoning in latent spaces, by means of embedding representations of logic atoms and FOL formulas; Logic Explained Networks (LEN) [Ciravegna et al. 2021, Barbiero et al. 2022] are explainable-by-design neural networks that can be used to both learn from and learn of logical constraints, by exploiting concept embedding models; [Galassi et al. 2021] proposes a neural-symbolic approach to mine argument components and relations from textual corpora; and LGE [Dhami et al. 2021] proposes a structure learning algorithm for learning embeddings for relational data (knowledge graphs) taking into account the local information in the graph. The method makes use of Gaifman locality theorem to obtain these embeddings.

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 characterizing 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] presents 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 [Petkovic et al. 2021], and semi-supervised learning [Petkovic et al. 2022].

Task 4.2: Integrating Representations for Learning and Optimization

State of the art. One recent survey 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 only partially 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.

Achievements w.r.t Tasks 4.2.

There are some key contributions that fit the regret minimisation viewpoint. First, [Mulamba et al. 2021] introduces 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 HW dimensioning and configuration for online algorithms in the energy system domain, under an heterogeneous set of constraints. The ML models are used to predict the online algorithms performance on different HW 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 learning approach for acquiring mixed-integer linear programming models from historical data that leverages both gradient-based and combinatorial search for learning.

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] contribute 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) contribute an approach to anomaly detection in multichannel data using sparse representation in RADWT frames and (Lombardi et al. 2020) analyse regularized approaches for constrained machine learning.

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 or whether they just regurgitate pieces of similar answers seen in their vast amounts of training data is currently still subject of an intense scientific debate [Marcus & Davis, 2020]. 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], or more recently OpenAI’s DALL-E model for image generation from text captions [Ramesh et al. 2021; Ramesh et al. 2022].

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, we 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 regularization, 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. We thus believe continual learning of VQA systems is also worthy to be further investigated.

Achievements w.r.t Tasks 4.4

Important achievements with regard 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 categorized by their representative techniques including regularization, 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 categorized based on the input data modality.

Other contributions for Task 4.3 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 maximizes 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.

Possible Future Work

In the following, we highlight a selection of the future work single partners are planning to explore and collaborations that are fueled by the TAILOR project.

- UNIPI is performing joint work with the University of Edinburgh on assessing the reasoning capabilities of deep learning models in chemical applications: results show evidence of algebraic reasoning capabilities being needed to reliably tackle chemical reaction prediction tasks.
- UNIPI and TU Delft are developing an approach to control complex physical systems on a learned compressed representation where the Hamiltonian/Lagrangian structure of the system is preserved.
- TU Delft is looking at facilitating certain chemical engineering tasks with exact training of graph neural networks.
- UNIBO is working on generalizations of Decision Focused Learning that rely on Reinforcement Learning and bi-level optimization. A joint effort with UNITN and TUE will also research the use of EML with iterative sampling (similarly to what is done in Bayesian Optimization) to handle constrained optimization problems with blackbox optimization functions. UniBo is also working on improving how ML models are encoded within optimization models in EML, especially to expand the pool of supported ML models (e.g., convolutional neural networks) and to improve scalability (e.g., NN distillation or quantization techniques).
- CNR is working on learning constraints in binary classification problems considering the possible presence of outliers.
- CINI is performing joint work with KU Leuven on integrating logical knowledge into embedding representation of knowledge graphs while guaranteeing logical reasoning semantics.
- CINI is collaborating with the University of Cambridge in order to extract ontologies and first-order logic rules from graph-structured data.

- CINI is collaborating with other WP3 partners at the coordinated action (CA2) on Explainable Malware Detection of Task 3.1 to apply newly devised KGE systems to real-world problems.
- KU Leuven is learning constraints in collaboration with UNITN
- KU Leuven is further investigating the integration of StarAI and NeSy, in collaboration with other partners such as CINI (Uni Sienna and TU/Delft)

Importantly, the work conducted by the Taskforce, and covered more in detail in Deliverable 4.3, is strategic for WP 4, because it lays the basis for a common understanding of the challenges to solve and to drive future research activities and collaborations.

At the more global WP 4 level, further discussions and joint activities are planned in particular with WP 3, possibly with WP 5, and possibly also with the ELLIS Program on Semantic, Symbolic and Interpretable Machine Learning.

References

[Agrawal et al. 2015] VQA: Visual Question Answering. ICCV 2015

[Azzolini et al. 2022] Probabilistic Logic Models for the Lightning Network. Cryptography. 2022 [link] **- CINI**

[Barbiero et al. 2021] Entropy-based Logic Explanations of Neural Networks. AAAI 2022 [link] **- CINI**

[Brown et al. 2020] Language Models are Few-Shot Learners. arXiv 2020

[Cardillo and Straccia 2022] Fuzzy OWL-Boost: Learning fuzzy concept inclusions via real-valued boosting. Fuzzy Sets and Systems 2022 [link] **- CNR**

[Cima et al. 2021] Abstraction in Data Integration. LICS 2021 [link] **- La Sapienza**

[Cima et al. 2021b] Query Definability and Its Approximations in Ontology-based Data Management. CIKM 2021 [link] **- La Sapienza**

[Ciravegna et al. 2021] Logic explained networks. arXiv 2021 [link] **- CINI**

[Console et al. 2021] Model-theoretic Characterizations of Rule-based Ontologies. PODS 2021 [link] **- La Sapienza**

[Dai et al. 2021] A survey on knowledge graph embedding: Approaches, applications and benchmarks." Electronics 2020

[De Canditiis and De Feis 2021] Anomaly Detection in Multichannel Data Using Sparse Representation in RADWT Frames. Mathematics 2021 [link] **- CNR**

[De Filippo et al. 2022] HADA: an Automated Tool for Hardware Dimensioning of AI Applications. Knowledge-Based Systems 2022 [link] **- UNIBO**

[Devlin et al. 2019] BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. NAACL-HLT 2019

[Dhami et al. 2021] Non-Parametric Learning of Gaifman Models. ILP 2021 [link] **- TUDA**

[Dhami et al. 2022] A Statistical Relational Approach to Learning Distance-based GCNs. AAAI workshop 2022 [link] **- TUDA**

[Di Liello et al. 2020] Efficient Generation of Structured Objects with Constrained Adversarial Networks. NeurIPS 2020 [link] **- UNITN**

[Diligenti et al. 2022] A Constraint-Based Approach to Learning and Reasoning. AAAI 2020 [link] **- CINI**

[Fadja et al. 2021] Learning hierarchical probabilistic logic programs. Machine Learning 2021 [link] **- CINI**

[Galassi et al. 2021] Investigating Logic Tensor Networks for Neural-Symbolic Argument Mining. StarAI workshop 2021 [link] **- CINI**

[Gong et al, 2022] Meta agent teaming active learning for pose estimation. CVPR 2022 [link] **- LU**

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

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

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

[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] Hybrid reasoning in knowledge graphs: Combining symbolic reasoning and statistical reasoning. Semantic Web 2020

[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**

[Maggini et al. 2022] A Lagrangian Framework for Learning in Graph Neural Networks. AI in the Age of Neural Networks and Brain Computing 2022 **- CINI**

[Marcus & Davis, 2020] GPT-3: Bloviator", Technology Review 2020

[Marra et al. 2020] From Statistical Relational Learning to Neural-Symbolic Artificial Intelligence. Early version at IJCAI 2020 and long version at arxiv [link] **- KUL & TU/e**

[Marra et al. 2021] Learning Representations for Sub-Symbolic Reasoning. arXiv 2021 [link] **- KUL, CINI**

[Misino et al. 2022] VAEL: Bridging Variational Autoencoders and Probabilistic Logic Programming. arXiv 2022 [link] **- KUL & UNIBO**

[Mokady et al. 2021] ClipCap: CLIP Prefix for Image Captioning. arXiv 2021

[Morettin et al. 2021] Paolo Morettin, Pedro Zuidberg Dos Martires, Samuel Kolb, and Andrea Passerini. Hybrid probabilistic inference with logical and algebraic constraints: a survey. In IJCAI Survey Track 2021. [link] **- KUL, UNITN**

[Morettin et al. 2021] Hybrid probabilistic inference with logical and algebraic constraints: a survey. IJCAI 2021 [link] **- KUL, UNITN**

[Mulamba et al. 2021] Contrastive Losses and Solution Caching for Predict-and-Optimize. IJCAI 2021 [link] **- KUL**

[Nickel et al. 2016] A Review of Relational Machine Learning for Knowledge Graphs.Proceedings of the IEEE 2016

[Pellegrini et al. 2021] Giovanni Pellegrini, Alessandro Tibo, Paolo Frasconi, Andrea Passerini, and Manfred Jaeger. Learning Aggregation Functions. In IJCAI 2021. [link] **- UNITN**

[Peng et al. 2021] The Multi-Modal Video Reasoning and Analyzing Competition. ICCV Workshop 2021 [link] **- LU**

[Petkovic et al. 2021] Ensemble- and Distance- Based Feature Ranking for Unsupervised Learning. Int. J. Intell. Syst. 2021 [link] **- JSI**

[Petkovic et al. 2022] Feature ranking for semi-supervised learning MLJ 2022 **[**link**] - JSI**

[Qu et al, 2021] Recent Advances of Continual Learning in Computer Vision: An Overview. arXiv 2021 [link] **- LU**

[Ramesh et al. 2021] Zero-Shot Text-to-Image Generation. ICML 2021

[Ramesh et al. 2022] Hierarchical Text-Conditional Image Generation with CLIP Latents. arXiv 2022

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

[Scarselli et al. 2009] The graph neural network model. IEEE Transactions on Neural Networks 2009

[Shindo et al. 2022] Neuro-Symbolic Forward Reasoning. arXiv 2022 [link] **- TUDA**

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

[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] Efficient and Reliable Probabilistic Interactive Learning with Structured Outputs. AAAI Workshop 2022 [link] **- UNITN**

[Teso et al. 2022] Stefano Teso, Laurens Bliek, Andrea Borghesi, Michele Lombardi, Neil Yorke-Smith, Tias Guns, Andrea Passerini. Machine Learning for Combinatorial Optimisation of Partially-Specified Problems: Regret Minimisation as a Unifying Lens. arXiv 2022 [link] **- UNITN, KUL, TU Delft, UNIBO**

[van Driel, 2021] Learning Variable Activity Initialisation for Lazy Clause Generation Solvers. CPAIOR 2021 [link] **- TU Delft**

[Ventola et al. ILP 2021] Generative Clausal Networks: Relational Decision Trees as Probabilistic Circuits. ILP 2021 [link] **- TUDA**

[Vergari et al. 2021] A Compositional Atlas of Tractable Circuit Operations: From Simple Transformations to Complex Information-theoretic Queries. NeurIPS 2021 [link] **- UNITN**

[Winters et al. AAAI 2022] DeepStochLog: Neural Stochastic Logic Programming. AAAI 2022 [link] **- KUL**

[Yang et al. 2021] Lifted Model Checking for Relational MDPs. MLJ 2021 [link] **- KUL**

[Wang et al. 2017] Knowledge graph embedding: A survey of approaches and applications. IEEE Transactions on Knowledge and Data Engineering 2017

[Zecevic et al. 2021] Interventional Sum-Product Networks: Causal Inference with Tractable Probabilistic Models. NeurIPS 2021 [link] **- TUDA**

[Zhang et al. 2021] Neural, symbolic and neural-symbolic reasoning on knowledge graphs AI Open 2 2021

Appendix 1: Workshops and Activities Organized by WP 4

We provide a list of the workshop and activities organized 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 Gianotti 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

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 organization 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 behavior of a system, and 2) embedding such 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. Specic 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 direction has been so far constrained by limited resource 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 kind 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 partecipation of Task 7.4 of 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 formalization, 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 democratize data analysis and make it more accessible to non-experts by automating these different steps as much as possible. This event is organized by the ERC AdG project SYNTH, which has been devoted to the goal of automating and democratizing 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