ECAI 2024
U. Endriss et al. (Eds.)

doi:10.3233/FAIA240468

Contributions to the Journal Track

Gregor Behnke

Chair of the ECAI-2024 Journal Track
University of Amsterdam

Abstract. The journal track of the 27th European Conference on Artificial Intelligence (ECAI-2024) offered the authors of papers recently accepted for publication by either one of the two leading discipline-wide journals in AI, *Artificial Intelligence* (AIJ) and the *Journal of Artificial Intelligence Research* (JAIR), the opportunity to present their work at the conference without undergoing an additional round of reviewing. Papers were eligible only if no part had previously been presented at a conference with archival proceedings. Traditionally, the authors of such papers would have missed out on the opportunity to present their work to a broader research audience. This limitation tends to discourage the submission of original work to journals without prior conference publications on the same topic. The intention of the journal track is to encourage a "journal-first" publication strategy–by giving authors the option to present their work at a suitable conference venue such as ECAI. On the following pages, for each paper presented at the journal track, we list bibliographic information and the abstract of the original publication.

Dealing with expert bias in collective decision-making

Axel Abels, Tom Lenaerts, Vito Trianni, and Ann Nowé. Dealing with expert bias in collective decision-making. *Artificial Intelligence*, 320:103921, 2023. DOI: 10.1016/j.artint.2023.103921

Quite some real-world problems can be formulated as decision-making problems wherein one must repeatedly make an appropriate choice from a set of alternatives. Multiple expert judgments, whether human or artificial, can help in taking correct decisions, especially when exploration of alternative solutions is costly. As expert opinions might deviate, the problem of finding the right alternative can be approached as a collective decision making problem (CDM) via aggregation of independent judgments. Current state-of-the-art approaches focus on efficiently finding the optimal expert, and thus perform poorly if all experts are not qualified or if they display consistent biases, thereby potentially derailing the decision-making process. In this paper, we propose a new algorithmic approach based on contextual multi-armed bandit problems (CMAB) to identify and counteract such biased expertise. We explore homogeneous, heterogeneous and polarized expert groups and show that this approach is able to effectively exploit the collective expertise, outperforming state-of-the-art methods, especially when the quality of the provided expertise degrades. Our novel CMAB-inspired approach achieves a higher final performance and does so while converging more rapidly than previous adaptive algorithms.

DeepSym: Deep Symbol Generation and Rule Learning for Planning from Unsupervised Robot Interaction

Alper Ahmetoglu, M. Yunus Seker, Justus Piater, Erhan Oztop, and Emre Ugur. DeepSym: Deep Symbol Generation and Rule Learning for Planning from Unsupervised Robot Interaction. *Journal of Artificial Intelligence Research*, 75:709–745, 2022. DOI: 10.1613/jair.1.13754

Symbolic planning and reasoning are powerful tools for robots tackling complex tasks. However, the need to manually design the symbols restrict their applicability, especially for robots that are expected to act in open-ended environments. Therefore symbol formation and rule extraction should be considered part of robot learning, which, when done properly, will offer scalability, flexibility, and robustness. Towards this goal, we propose a novel general method that finds action-grounded, discrete object and effect categories and builds probabilistic rules over them for non-trivial action planning. Our robot interacts with objects using an initial action repertoire that is assumed to be acquired earlier and observes the effects it can create in the environment. To form action-grounded object, effect, and relational categories, we employ a binary bottleneck layer in a predictive, deep encoder-decoder network that takes the image of the scene and the action applied as input, and generates the resulting effects in the scene in pixel coordinates. After learning, the binary latent vector represents action-driven object categories based on the interaction experience of the robot. To distill the knowledge represented by the neural network into rules useful for symbolic reasoning, a decision tree is trained to reproduce its decoder function. Probabilistic rules are extracted from the decision paths of the tree and are represented in the Probabilistic Planning Domain Definition Language (PPDDL), allowing off-the-shelf planners to operate on the knowledge extracted from the sensorimotor experience of the robot. The deployment of the proposed approach for a simulated robotic manipulator enabled the discovery of discrete representations of object properties such as 'rollable' and 'insertable'. In turn, the use of these representations as symbols allowed the generation of effective plans for achieving goals, such as building towers of the desired height, demonstrating the effectiveness of the approach for multi-step object manipulation. Finally, we demonstrate that the system is not only restricted to the robotics domain by assessing its applicability to the MNIST 8-puzzle domain in which learned symbols allow for the generation of plans that move the empty tile into any given position.

Dynamic Controllability of Temporal Plans in Uncertain and Partially Observable Environments

Arthur Bit-Monnot and Paul Morris. Dynamic Controllability of Temporal Plans in Uncertain and Partially Observable Environments. *Journal of Artificial Intelligence Research*, 77:1311–1369, 2023.

DOI: 10.1613/jair.1.13065

The formalism of Simple Temporal Networks (STNs) provides methods for evaluating the feasibility of temporal plans. The basic formalism deals with the consistency of quantitative temporal requirements on scheduled events. This implicitly assumes a single agent has full control over the timing of events. The extension of Simple Temporal Networks with Uncertainty (STNU) introduces uncertainty into the timing of some events. Two main approaches to the feasibility of STNUs involve (1) where a single schedule works irrespective of the duration outcomes, called Strong Controllability, and (2) whether a strategy exists to schedule future events based on the outcomes of past events, called Dynamic Controllability. Case (1) essentially assumes the timing of uncertain events cannot be observed by the agent while case (2) assumes full observability.

The formalism of Partially Observable Simple Temporal Networks with Uncertainty (POSTNU) provides an intermediate stance between these two extremes, where a known subset of the uncertain events can be observed when they occur. A sound and complete polynomial algorithm to determining the Dynamic Controllability of POSTNUs has not previously been known; we present one in this paper. This answers an open problem that has been posed in the literature.

The approach we take factors the problem into Strong Controllability micro-problems in an overall Dynamic Controllability macro-problem framework. It generalizes the notion of labeled distance graph from STNUs. The generalized labels are expressed as max/min expressions involving the observables. The paper introduces sound generalized reduction rules that act on the generalized labels. These incorporate tightenings based on observability that preserve dynamic viable strategies. It is shown that if the generalized reduction rules reach quiescence without exposing an inconsistency, then the POSTNU is Dynamically Controllable (DC). The paper also presents algorithms that apply the reduction rules in an organized way and reach quiescence in a polynomial number of steps if the POSTNU is Dynamically Controllable.

Remarkably, the generalized perspective leads to a simpler and more uniform framework that applies also to the STNU special case. It helps illuminate the previous methods inasmuch as the max/min label representation is more semantically clear than the ad-hoc upper/lower case labels previously used.

Multi-Objective Reinforcement Learning Based on Decomposition: A Taxonomy and Framework

Florian Felten, El-Ghazali Talbi, and Grégoire Danoy. Multi-Objective Reinforcement Learning Based on Decomposition: A Taxonomy and Framework. *Journal of Artificial Intelligence Research*, 79:679–723, 2024. DOI: 10.1613/jair.1.15702

Multi-objective reinforcement learning (MORL) extends traditional RL by seeking policies making different compromises among conflicting objectives. The recent surge of interest in MORL has led to diverse studies and solving methods, often drawing from existing knowledge in multi-objective optimization based on decomposition (MOO/D). Yet, a clear categorization based on both RL and MOO/D is lacking in the existing literature. Consequently, MORL researchers face difficulties when trying to classify contributions within a broader context due to the absence of a standardized taxonomy.

To tackle such an issue, this paper introduces multi-objective reinforcement learning based on decomposition (MORL/D), a novel methodology bridging the literature of RL and MOO. A comprehensive taxonomy for MORL/D is presented, providing a structured foundation for categorizing existing and potential MORL works. The introduced taxonomy is then used to scrutinize MORL research, enhancing clarity and conciseness through well-defined categorization. Moreover, a flexible framework derived from the taxonomy is introduced. This framework accommodates diverse instantiations using tools from both RL and MOO/D. Its versatility is demonstrated by implementing it in different configurations and assessing it on contrasting benchmark problems. Results indicate MORL/D instantiations achieve comparable performance to current state-of-the-art approaches on the studied problems. By presenting the taxonomy and framework, this paper offers a comprehensive perspective and a unified vocabulary for MORL. This not only facilitates the identification of algorithmic contributions but also lays the groundwork for novel research avenues in MORL.

An Attention Model for the Formation of Collectives in Real-World Domains

Adrià Fenoy, Filippo Bistaffa, and Alessandro Farinelli. An Attention Model for the Formation of Collectives in Real-World Domains. *Artificial Intelligence*, 328:104064, 2024. DOI: 10.1016/j.artint.2023.104064

We consider the problem of forming collectives of agents inherent in application domains aligned with Sustainable Development Goals 4 and 11 (i.e., team formation and ridesharing, respectively). We propose a general solution approach based on a novel combination of an attention model and an integer linear program (ILP). In more detail, we propose an attention encoder-decoder model that transforms a collective formation instance to a weighted set packing problem, which is then solved by an ILP. Results on collective formation problems inherent in the ridesharing and team formation domains show that our approach provides comparable solutions (in terms of quality) to the ones produced by state-of-the-art approaches specific to each domain. Moreover, our solution outperforms the most recent general approach for forming collectives based on Monte Carlo tree search.

Levi and Harper identities for non-prioritized belief base change

Marco Garapa, Eduardo Fermé, and Maurício Reis. Levi and Harper identities for non-prioritized belief base change. *Artificial Intelligence*, 319:103907, 2023. DOI: 10.1016/j.artint.2023.103907

In this paper, we investigate the relation between shielded base contraction postulates and credibility-limited (CL) base revision postulates. More precisely, we identify (i) the relation between the postulates satisfied by a shielded base contraction operator and the postulates satisfied by the CL base revision operator that is defined from it by means of the consistency-preserving Levi identity and (ii) the relation between the postulates satisfied by a CL base revision operator and the postulates satisfied by the shielded base contraction operator that is defined from it by means of the Harper identity. Furthermore, we show that the consistency-preserving Levi identity and the Harper identity establish a one-to-one correspondence between the twenty classes of shielded base contractions presented in [1] and the twenty classes of credibility-limited base revisions presented in [2].

[1] Marco Garapa, Eduardo Fermé, Maurício D.L. Reis. Shielded base contraction. *Artificial Intelligence*, 259:186–216, 2018.

[2] Marco Garapa, Eduardo Fermé, Maurício D.L. Reis. Credibility-limited base revision: new classes and their characterizations. *Journal of Artificial Intelligence Research*, 69:1023–1075, 2020.

Polarized Message-Passing in Graph Neural Networks

Tiantian He, Yang Liu, Yew-Soon Ong, Xiaohu Wu, and Xin Luo. Polarized Message-Passing in Graph Neural Networks. *Artificial Intelligence*, 331:104126, 2024. DOI: 10.1016/j.artint.2024.104129

In this paper, we present Polarized message-passing (PMP), a novel paradigm to revolutionize the design of message-passing graph neural networks (GNNs). In contrast to existing methods, PMP captures the power of node-node similarity and dissimilarity to acquire dual sources of messages from neighbors. The messages are then coalesced to enable GNNs to learn expressive representations from sparse but strongly correlated neighbors. Three novel GNNs based on the PMP paradigm, namely PMP graph convolutional network (PMP-GCN), PMP graph attention network (PMP-GAT), and PMP graph PageRank network (PMP-GPN) are proposed to perform various downstream tasks. Theoretical analysis is also conducted to verify the high expressiveness of the proposed PMP-based GNNs. In addition, an empirical study of five learning tasks based on 12 real-world datasets is conducted to validate the performances of PMP-GCN, PMP-GAT, and PMP-GPN. The proposed PMP-GCN, PMP-GAT, and PMP-GPN outperform numerous strong message-passing GNNs across all five learning tasks, demonstrating the effectiveness of the proposed PMP paradigm.

Iterative Train Scheduling under Disruption with Maximum Satisfiability

Alexandre Lemos, Filipe Gouveia, Pedro T. Monteiro, and Inês Lynce. Iterative Train Scheduling under Disruption with Maximum Satisfiability. *Journal of Artificial Intelligence Research*, 79:1047–1090, 2024. DOI: 10.1613/jair.1.14924

This paper proposes an iterative Maximum Satisfiability (MaxSAT) approach designed to solve train scheduling optimization problems. The generation of railway timetables is known to be intractable for a single track. We consider hundreds of trains on interconnected multi-track railway networks with complex connections between trains. Furthermore, the proposed algorithm is incremental to reduce the impact of time discretization.

The performance of our approach is evaluated with the real-world Swiss Federal Railway (SBB) Crowd Sourcing Challenge benchmark and Periodic Event Scheduling Problems benchmark (PESPLib). The execution time of the proposed approach is shown to be, on average, twice as fast as the best existing solution for the SBB instances. In addition, we achieve a significant improvement over SAT-based solutions for solving the PESPLib instances.

We also analyzed real schedule data from Switzerland and the Netherlands to create a disruption generator based on probability distributions. The novel incremental algorithm allows solving the train scheduling problem under disruptions with better performance than traditional algorithms.

Functional Relation Field: A Model-Agnostic Framework for Multivariate Time Series Forecasting

Ting Li, Bing Yu, Jianguo Li, Guoli Yang, and Zhanxing Zhu. Functional Relation Field: A Model-Agnostic Framework for Multivariate Time Series Forecasting. *Artificial Intelligence*, 334:104158, 2024. DOI: 10.1016/j.artint.2024.104158

In multivariate time series forecasting, the most popular strategy for modeling the relationship between multiple time series is the construction of graph, where each time series is represented as a node and related nodes are connected by edges. However, the relationship between multiple time series is typically complicated, e.g. the sum of outflows from upstream nodes may be equal to the inflows of downstream nodes. Such relations widely exist in many real-world scenarios for multivariate time series forecasting, yet are far from well studied. In these cases, graph might be insufficient for modeling the complex dependency between nodes. To this end, we explore a new framework to model the inter-node relationship in a more precise way based our proposed inductive bias, *Functional Relation Field*, where a group of functions parameterized by neural networks are learned to characterize the dependency between multiple time series. Essentially, these learned functions then form a "field", i.e. a particular set of constraints, to regularize the training loss of the backbone prediction network and enforce the inference process to satisfy these constraints. Since our framework introduces the relationship bias in a data-driven manner, it is flexible and model-agnostic such that it can be applied to any existing multivariate time series prediction networks for boosting performance. The experiment is conducted on one toy dataset to show our approach can well recover the true constraint relationship between nodes. And various real-world datasets are also considered with different backbone prediction networks. Results show that the prediction error can be reduced remarkably with the aid of the proposed framework.

A k-additive Choquet integral-based approach to approximate the SHAP values for local interpretability in machine learning

Guilherme Pelegrina, Leonardo Duarte, and Michel Grabisch. A k-additive Choquet integral-based approach to approximate the SHAP values for local interpretability in machine learning. *Artificial Intelligence*, 325:104014, 2023. DOI: 10.1016/j.artint.2023.104014

Besides accuracy, recent studies on machine learning models have been addressing the question on how the obtained results can be interpreted. Indeed, while complex machine learning models are able to provide very good results in terms of accuracy even in challenging applications, it is difficult to interpret them. Aiming at providing some interpretability for such models, one of the most famous methods, called SHAP, borrows the Shapley value concept from game theory in order to locally explain the predicted outcome of an instance of interest. As the SHAP values calculation needs previous computations on all possible coalitions of attributes, its computational cost can be very high. Therefore, a SHAP-based method called Kernel SHAP adopts a strategy that approximates such values with less computational effort. However, we see two weaknesses in Kernel SHAP: its formulation is difficult to understand and it does not consider further game theory assumptions that could reduce the computational cost. Therefore, in this paper, we propose a novel approach that addresses such weaknesses. Firstly, we provide a straightforward formulation of a SHAP-based method for local interpretability by using the Choquet integral, which leads to both Shapley values and Shapley interaction indices. Thereafter, we propose to adopt the concept of k-additive games from game theory, which contributes to reduce the computational effort when estimating the SHAP values. The obtained results attest that our proposal needs less computations on coalitions of attributes to approximate the SHAP values.

Right Place, Right Time: Proactive Multi-Robot Task Allocation Under Spatiotemporal Uncertainty

Charlie Street, Bruno Lacerda, Manuel Mühlig, and Nick Hawes.

Right Place, Right Time: Proactive Multi-Robot Task Allocation Under Spatiotemporal Uncertainty. *Journal of Artificial Intelligence Research*, 79:137–171, 2024. DOI: 10.1613/jair.1.15057

For many multi-robot problems, tasks are announced during execution, where task announcement times and locations are *uncertain*. To synthesise multi-robot behaviour that is robust to early announcements and unexpected delays, multi-robot task allocation methods must explicitly model the stochastic processes that govern task announcement. In this paper, we model task announcement using continuous-time Markov chains which predict when and where tasks will be announced. We then present a task allocation framework which uses the continuous-time Markov chains to allocate tasks *proactively*, such that robots are near or at the task location upon its announcement. Our method seeks to minimise the *expected total waiting duration* for each task, i.e. the duration between task announcement and a robot beginning to service the task. Our framework can be applied to any multi-robot task allocation problem where robots complete spatiotemporal tasks which are announced stochastically. We demonstrate the efficacy of our approach in simulation, where we outperform baselines which do not allocate tasks proactively, or do not fully exploit our task announcement models.

Lifted Reasoning for Combinatorial Counting

Pietro Totis, Jesse Davis, Luc de Raedt, and Angelika Kimmig. Lifted Reasoning for Combinatorial Counting. *Journal of Artificial Intelligence Research*, 76:1–58, 2023. DOI: 10.1613/jair.1.14062

Combinatorics math problems are often used as a benchmark to test human cognitive and logical problem-solving skills. These problems are concerned with counting the number of solutions that exist in a specific scenario that is sketched in natural language. Humans are adept at solving such problems as they can identify commonly occurring structures in the questions for which a closed-form formula exists for computing the answer. These formulas exploit the exchangeability of objects and symmetries to avoid a brute-force enumeration of all possible solutions. Unfortunately, current AI approaches are still unable to solve combinatorial problems in this way. This paper aims to fill this gap by developing novel AI techniques for representing and solving such problems. It makes the following five contributions. First, we identify a class of combinatorics math problems which traditional lifted counting techniques fail to model or solve efficiently. Second, we propose a novel declarative language for this class of problems. Third, we propose novel lifted solving algorithms bridging probabilistic inference techniques and constraint programming. Fourth, we implement them in a lifted solver that solves efficiently the class of problems under investigation. Finally, we evaluate our contributions on a real-world combinatorics math problems dataset and synthetic benchmarks.

Object-agnostic Affordance Categorization via Unsupervised Learning of Graph Embeddings

Alexia Toumpa, and Anthony G. Cohn. Object-agnostic Affordance Categorization via Unsupervised Learning of Graph Embeddings. *Journal of Artificial Intelligence Research*, 77:1–38, 2023. DOI: 10.1613/jair.1.13253

Acquiring knowledge about object interactions and affordances can facilitate scene understanding and human-robot collaboration tasks. As humans tend to use objects in many different ways depending on the scene and the objects' availability, learning object affordances in everyday-life scenarios is a challenging task, particularly in the presence of an open set of interactions and objects. We address the problem of affordance categorization for class-agnostic objects with an open set of interactions; we achieve this by learning similarities between object interactions in an unsupervised way and thus inducing clusters of object affordances. A novel depth-informed qualitative spatial representation is proposed for the construction of Activity Graphs (AGs), which abstract from the continuous representation of spatio-temporal interactions in RGB-D videos. These AGs are clustered to obtain groups of objects with similar affordances. Our experiments in a real-world scenario demonstrate that our method learns to create object affordance clusters with a high V-measure even in cluttered scenes. The proposed approach handles object occlusions by capturing effectively possible interactions and without imposing any object or scene constraints.

Undesirable Biases in NLP: Addressing Challenges of Measurement

Oskar van der Wal, Dominik Bachmann, Alina Leidinger, Leendert van Maanen, Willem Zuidema, and Katrin Schulz. Undesirable Biases in NLP: Addressing Challenges of Measurement. *Journal of Artificial Intelligence Research*, 79:1–40, 2024. DOI: 10.1613/jair.1.15195

As Large Language Models and Natural Language Processing (NLP) technology rapidly develop and spread into daily life, it becomes crucial to anticipate how their use could harm people. One problem that has received a lot of attention in recent years is that this technology has displayed harmful biases, from generating derogatory stereotypes to producing disparate outcomes for different social groups. Although a lot of effort has been invested in assessing and mitigating these biases, our methods of measuring the biases of NLP models have serious problems and it is often unclear what they actually measure. In this paper, we provide an interdisciplinary approach to discussing the issue of NLP model bias by adopting the lens of psychometrics – a field specialized in the measurement of concepts like bias that are not directly observable. In particular, we will explore two central notions from psychometrics, the construct validity and the reliability of measurement tools, and discuss how they can be applied in the context of measuring model bias. Our goal is to provide NLP practitioners with methodological tools for designing better bias measures, and to inspire them more generally to explore tools from psychometrics when working on bias measurement tools.