Neuro-Symbolic Planning and PUI Exam 2023

Stefan Edelkamp

Neural Networks in Planning

- Combination of Symbolic and Subsymbolic AI
- Alias of Data-Driven and Deductive AI
- Deep/Reinforcement Learning for AI
- Learning Motions for Sampling-Based Motion Planning
- Learning Task Decompositions with Large Language Models
- Modeling Graph Neural Networks in Logic
- Learning Plans in Grid Domains such as Sokoban
- Learning Plans with Improved Loss Functions: L*, L_GBFS, L_A*

Why?

Safe AI

- No harm to people or the environment

Robust AI

- Stability and sensitivity in all situations

Scalable AI

- Crucial for real-life applications

Explainable AI

- Understand the complexity to have confidence

Hot Topic

E.g., Workshop series: Bridging the Gap Between AI Planning and Reinforcement Learning (PRL)
PRL @ ICAPS 2023
•Workshop/Tutorial dates: Prague, Czech Republic, July 9-10, 2023.
Some of the accepted papers will be invited to be presented at the IJCAI edition of the workshop as well.
PRL @ IJCAI 2023
•Workshop/Tutorial dates: Macao, S.A.R, August 19-21, 2023

Previous Editions
•PRL @ IJCAI 2022
•PRL @ ICAPS 2022
•PRL @ ICAPS 2021
•PRL @ ICAPS 2020

AIPlan4EU

Making planning applicable for everyone

- **AIPlan4EU: Bringing AI Planning to the European On-Demand AI Platform**

- https://www.aiplan4eu-project.eu/

EU-Tuples

TUPLES (TrUstworthy Planning and scheduling with Learning and ExplanationS) aims to build trusted planning and scheduling systems that are safe, robust, explainable and efficient

- https://tuples.ai/project/

NEXT-AI [tbd]

Follow-UP RCI - Reaserch Center Informatics

… under Consideration of Czech Goverment Consortium led by CIIRC and AIC 4 Excellent research

Mission

From foundational AI research towards breakthrough applications in healthcare, energy and environment, logistics, trustworthy digital society and new AI-driven industry.

Vision

NEXT-AI will be the leading center of excellence for artificial intelligence and machine learning research in central and eastern Europe by 2030.

The center is proposed by a collective

of internationally recognized researchers (demonstrated by publication/citation record, awarded ERC projects, strong international experience and delivered impact outside of the scientific research) in various fields driven by Artificial Intelligence and Machine Learning at several public research organizations in the Czech Republic.

The Moments MCTS and NNs Married [DM]

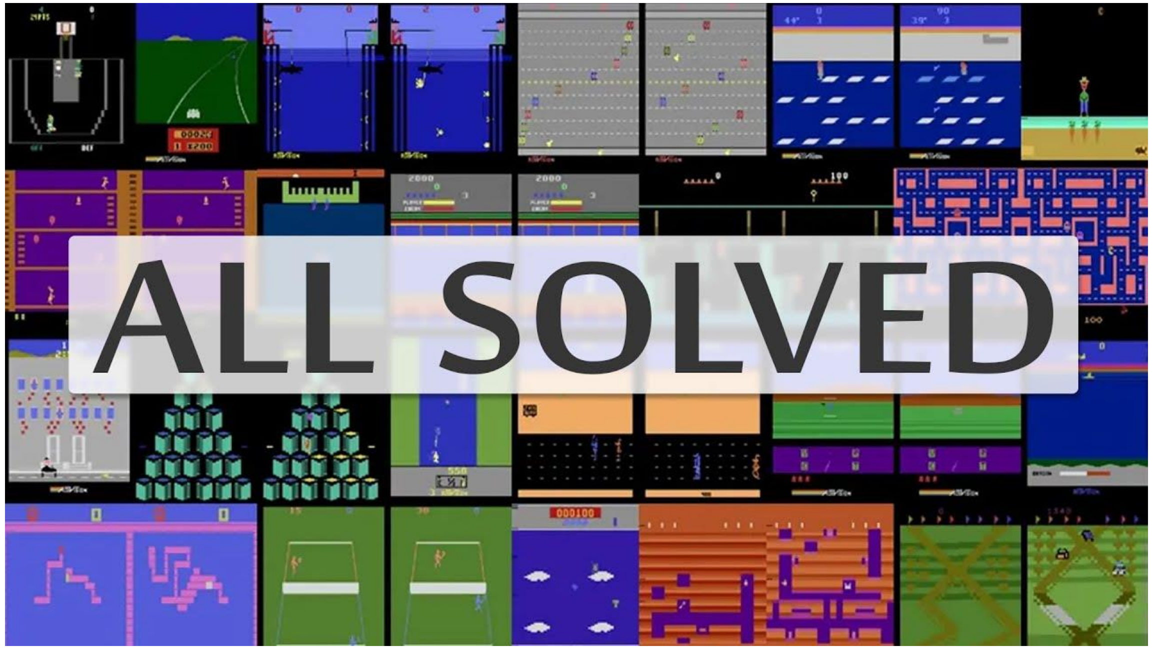

Atari Games

Computer Go

MCTS in Alignment, Routing, Packing, Robotics, Manufacturing, Graph Problems,… [Stefan, Yazz et al.]

Nested Rollout Policy Adaptation for Multiagent System Optimization in Manufacturing

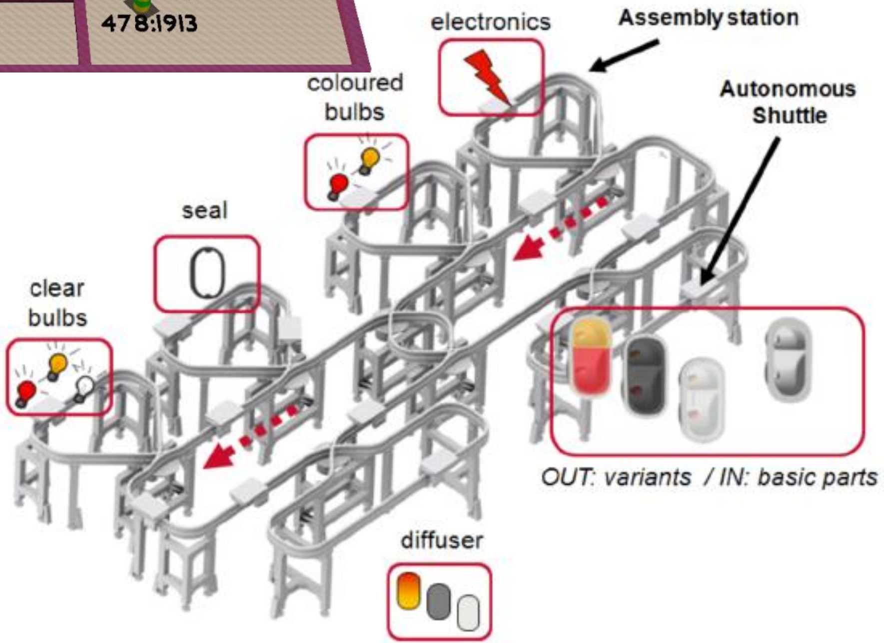

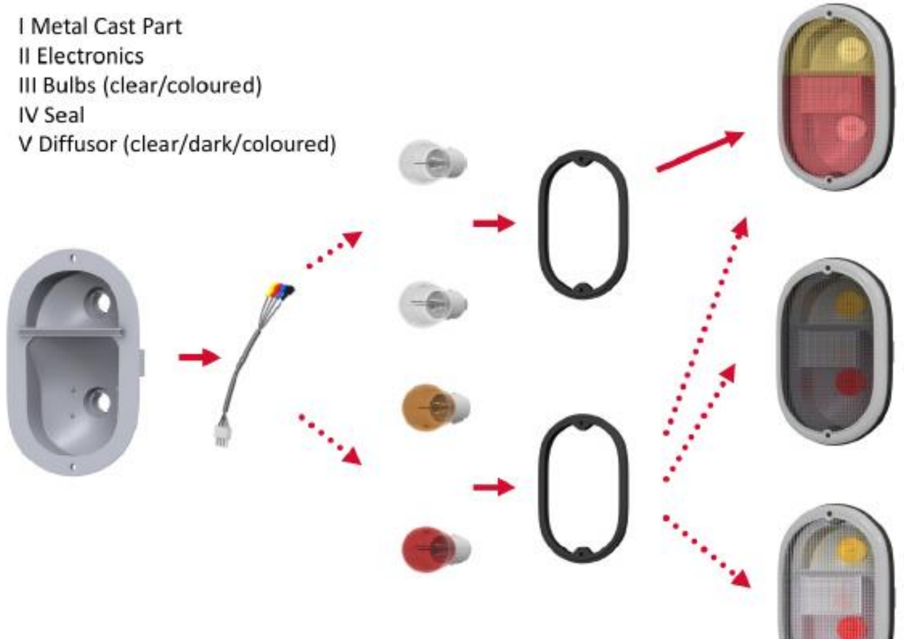

Monte-Carlo and Reinforcement Search [Jan]

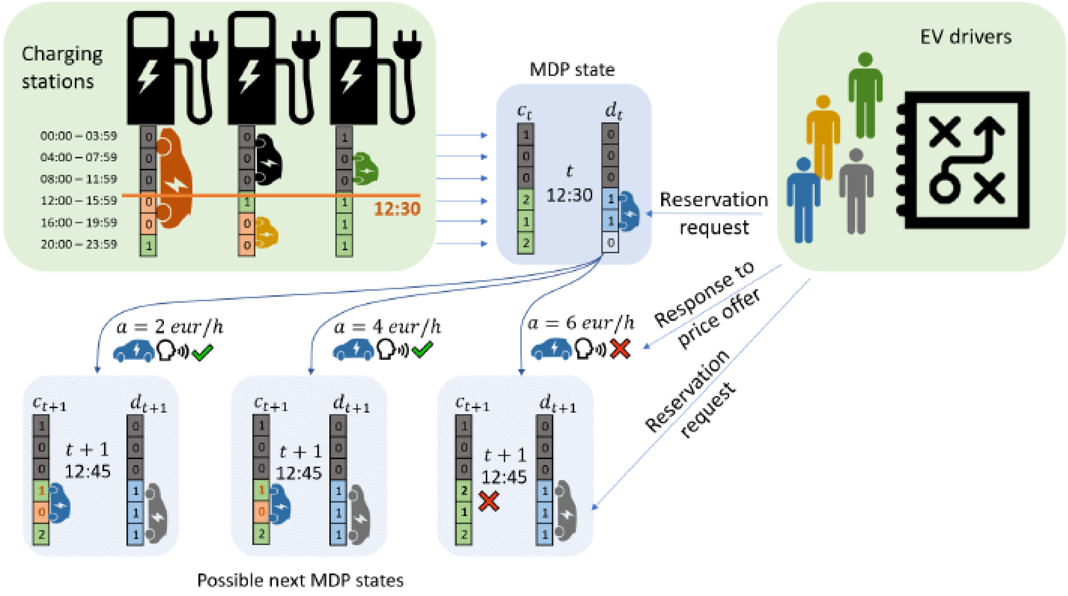

Figure 1. Illustration of the MDP states. The blue squares represent the MDP states. At timestep t, the capacity of the charging station is expressed by the capacity vector c_t. Elements of the vector represent available charging capacity in corresponding timeslots (time ranges in the green square).

Dynamic pricing is a proven technique to increase revenue in markets with heterogeneous demand. This work proposes a Markov Decision Process (MDP)-based approach to revenue- or utilization-maximizing dynamic pricing for charging station operators.

We implement the method using a Monte Carlo Tree Search (MCTS) algorithm and evaluate it in simulation using a range of problem instances based on a real-world dataset of EV charging sessions

Neuro-Symbolic AI methods [Rostia, Ondrej, Felip]

Develop new methods specifically for the setting of planning and scheduling systems.

Focus on their explainability in terms of the underlying domain language.

Resulting methods will contribute to build AI planning and scheduling systems that are transparent, robust, safe, and scalable.

Common requisite **is first-order logic** (FOL), which is a sufficiently expressive formal language to capture relational interdependencies in data processed by AI planning and scheduling systems.

Deal with **graph neural networks** (GNNs), which are nowadays one of the machine learningmodels of choice for complex data; they are used, among others, in several state-of-the-art planning and scheduling systems.

Planning via **C2 Logic**s & **Lifted Relational Neural Networks**

Neural networks for model-free and scale-free automated planning [Misa]

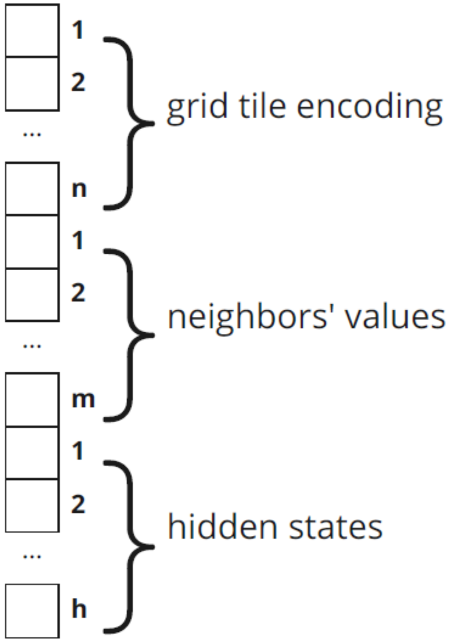

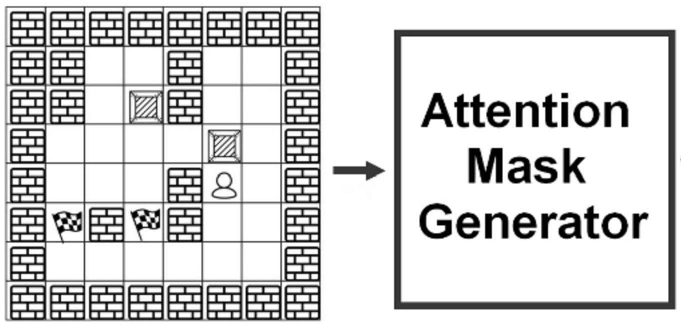

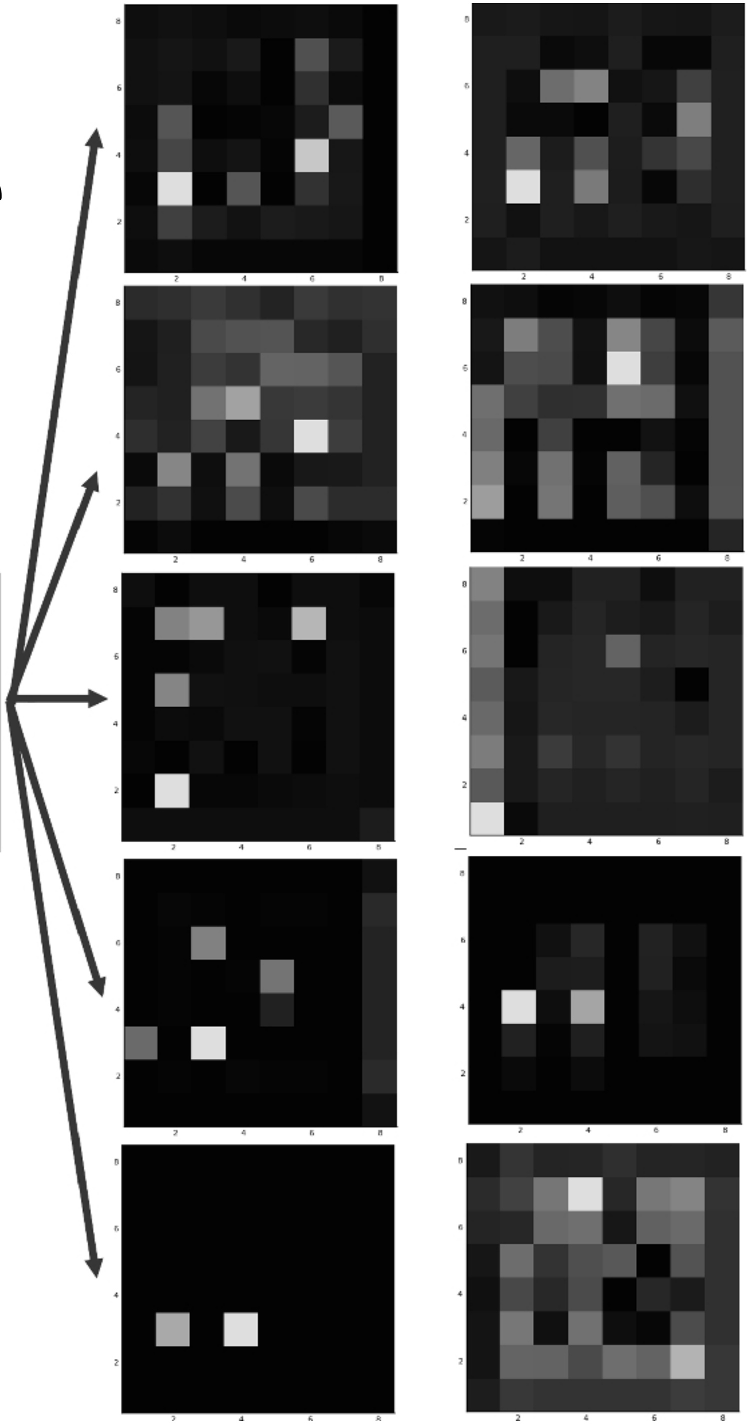

Example of generated attention masks for *Sokoban* instance

Learning Heuristic Estimates for Planning in Grid Domains by Cellular Simultaneous Recurrent Networks

Topics: Deep Learning; Machine Learning; Neural Networks; Planning and Scheduling

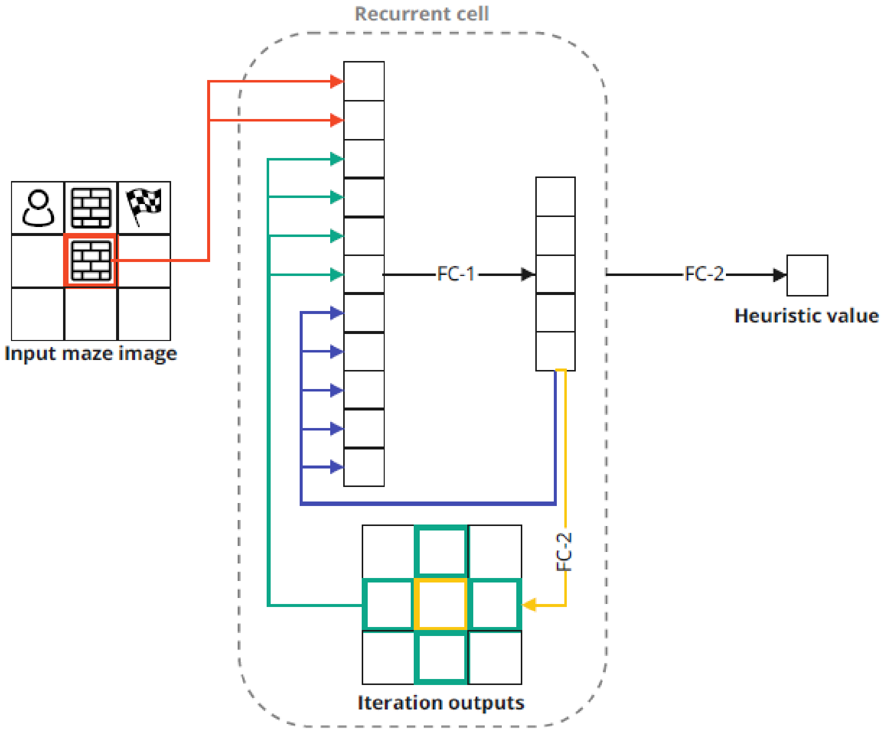

Learning heuristic functions for classical planning algorithms has been a great challenge in the past years. The biggest bottleneck of this technique is the choice of an appropriate description of the planning problem suitable for machine learning. Various approaches were recently suggested in the literature, namely grid-based, image-like, and graph-based. In this work, we extend the latest grid-based representation with layered architecture capturing the semantics of the related planning problem. Such an approach can be used as a domain-independent model for further heuristic learning. This representation keeps the advantages of the grid-structured input and provides further semantics about the problem we can learn from. Together with the representation, we also propose a new network architecture based on the Cellular Simultaneous Recurrent Networks (CSRN) that is capable of learning from such data and can be used instead of a heuristic function in the state-space search algorithms.

Semantically Layered Representation for Planning Problems and Its Usage for Heuristic Computation Using Cellular Simultaneous Recurrent Neural Networks

Topics: Deep Learning; Machine Learning; Neural Networks; Planning

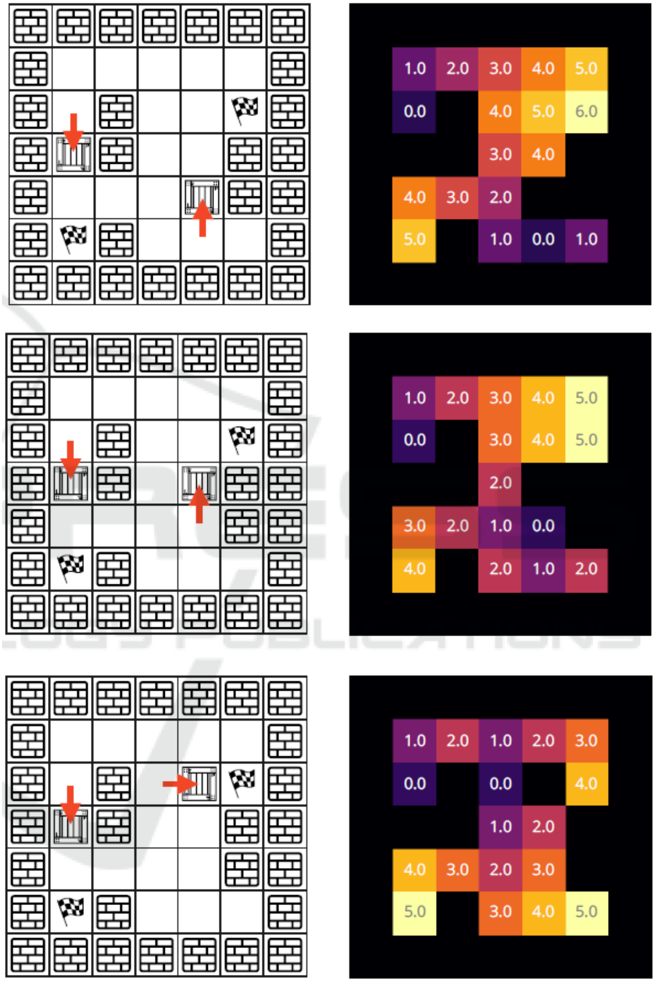

Automated planning provides a powerful general problem solving tool, however, its need for a model creates a bottleneck that can be an obstacle for using it in real-world settings. In this work we propose to use neural networks, namely Cellular Simultaneous Recurrent Networks (CSRN), to process a planning problem and provide a heuristic value estimate that can more efficiently steer the automated planning algorithms to a solution. Using this particular architecture provides us with a scale-free solution that can be used on any problem domain represented by a planar grid. We train the CSRN architecture on two benchmark domains, provide analysis of its generalizing and scaling abilities. We also integrate the trained network into a planner and compare its performance against commonly used heuristic functions.

Heuristic search planning with deep neural networks using imitation, attention and curriculum learning [Leah]

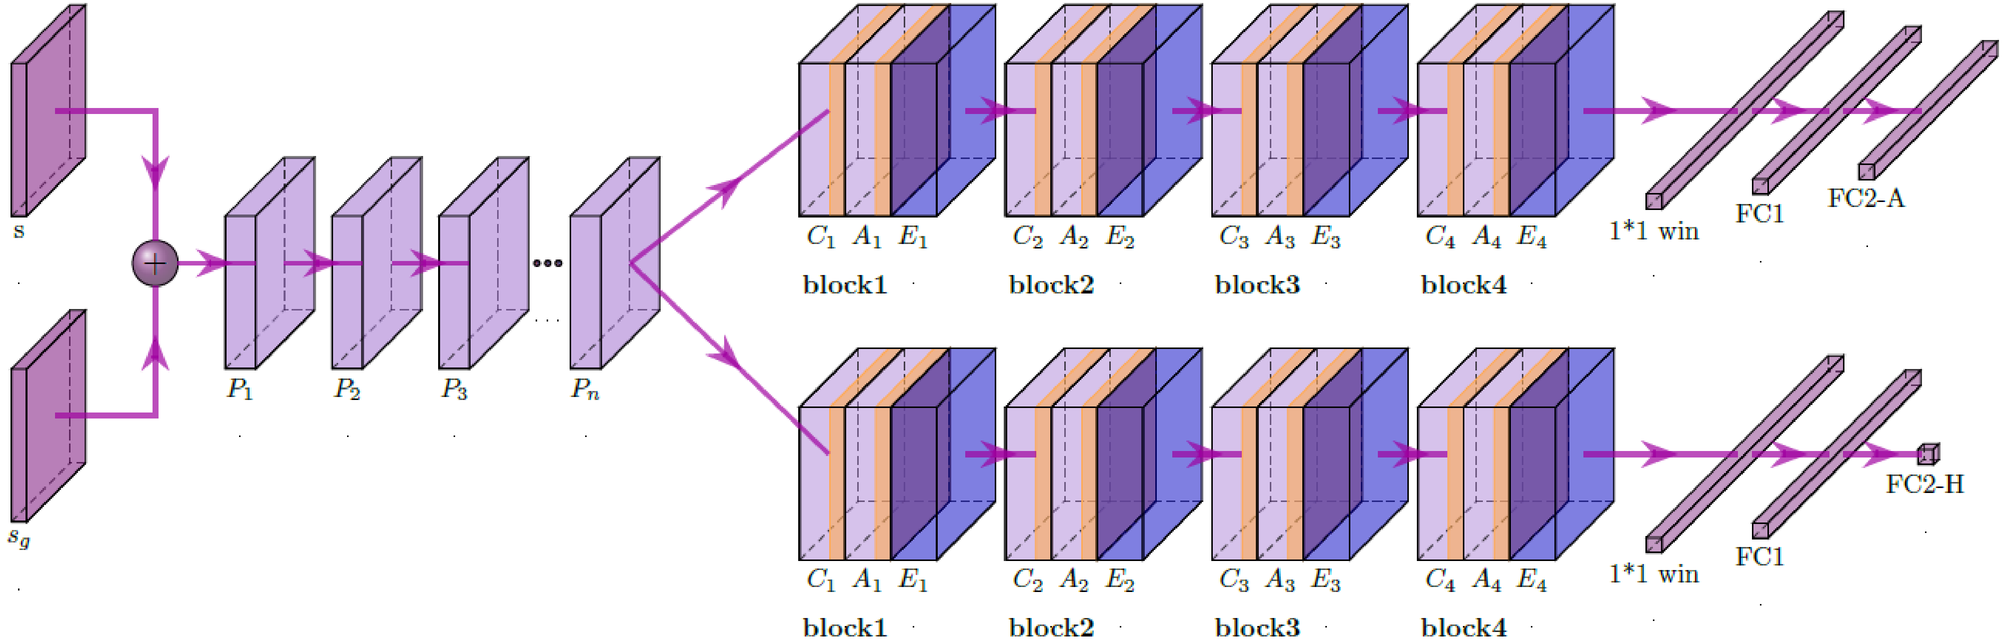

This work presents a network model to learn a heuristic capable of relating distant parts of the state space via optimal plan imitation using the attention mechanism, which drastically improves the learning of a good heuristic function. To counter the limitation of the method in the creation of problems of increasing difficulty, we demonstrate the use of curriculum learning, where newly solved problem instances are added to the training set, which, in turn, helps to solve problems of higher complexities and far exceeds the performances of all existing baselines including classical planning heuristics. We demonstrate its effectiveness for grid-type PDDL domains.

A Differentiable Loss Function for Learning Heuristics in A*

This paper argues that this does not necessarily lead to a faster search of A* algorithm since its execution relies on relative values instead of absolute ones. As a mitigation, we propose a L* loss, which upper-bounds the number of excessively expanded states inside the A* search. The L* loss, when used in the optimization of state-of-the-art deep neural networks for automated planning in maze domains like Sokoban and maze with teleports, significantly improves the fraction of solved problems, the quality of founded plans, and reduces the number of expanded states to approximately 50%

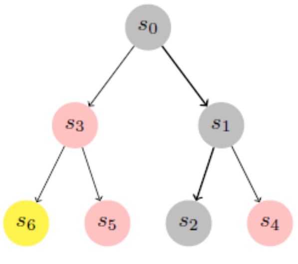

(a)

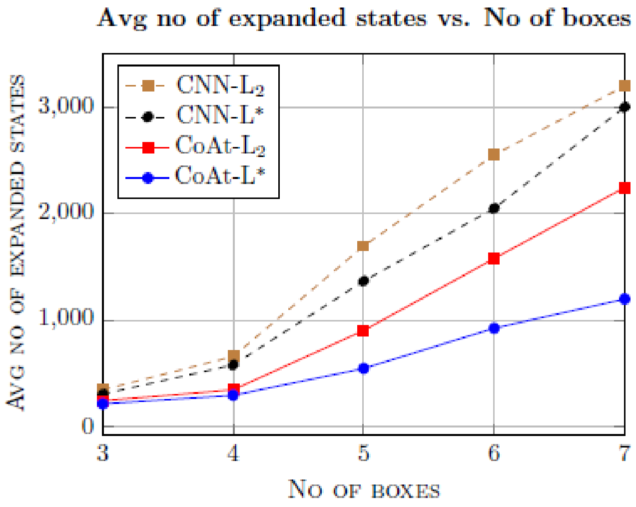

(a) Sokoban

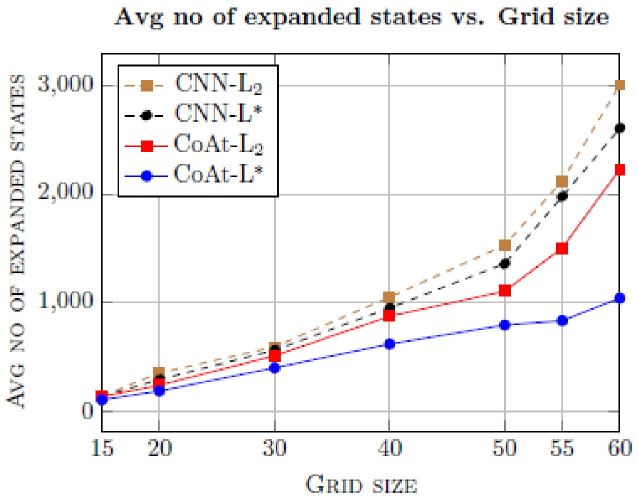

(b) Maze with teleports

Deep RRT* [Dang]

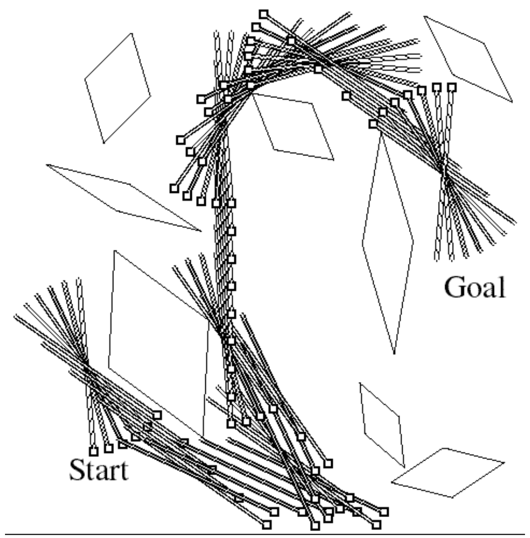

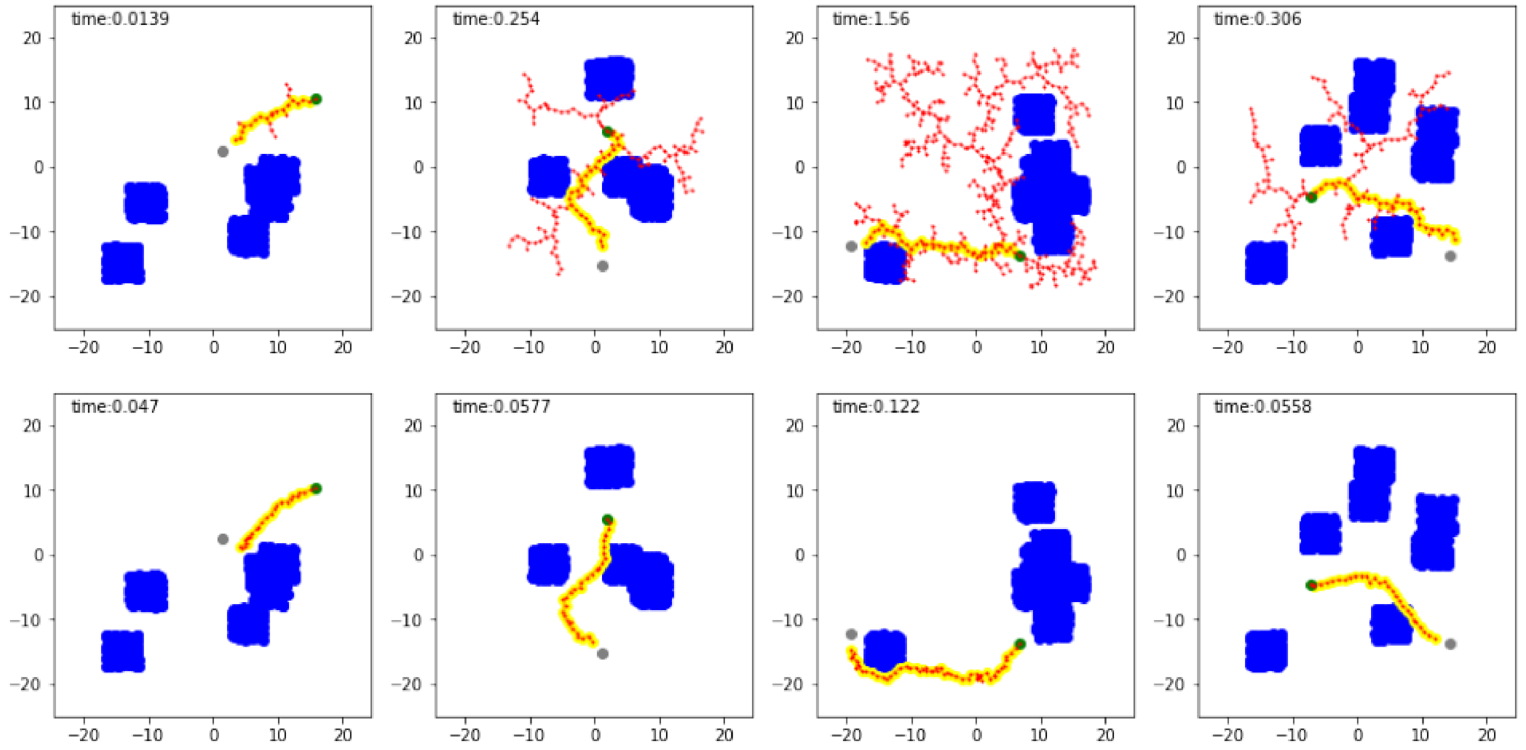

Figure 1: The first row is the results of RRT*. The second row is results of our method. The blue points are the obstacles cloud points. The red nodes are the tree nodes. The yellow lines are path. The green nodes are the start states, and the gray nodes are the goal states.

Framework

→Unity
→ROS

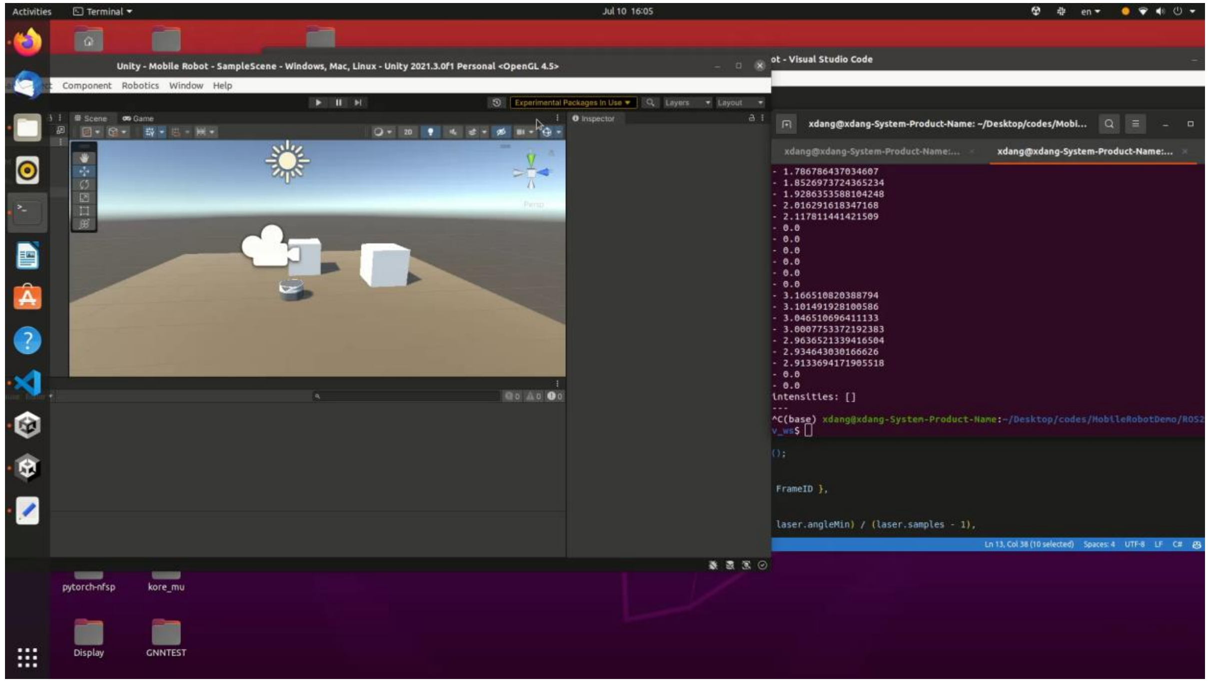

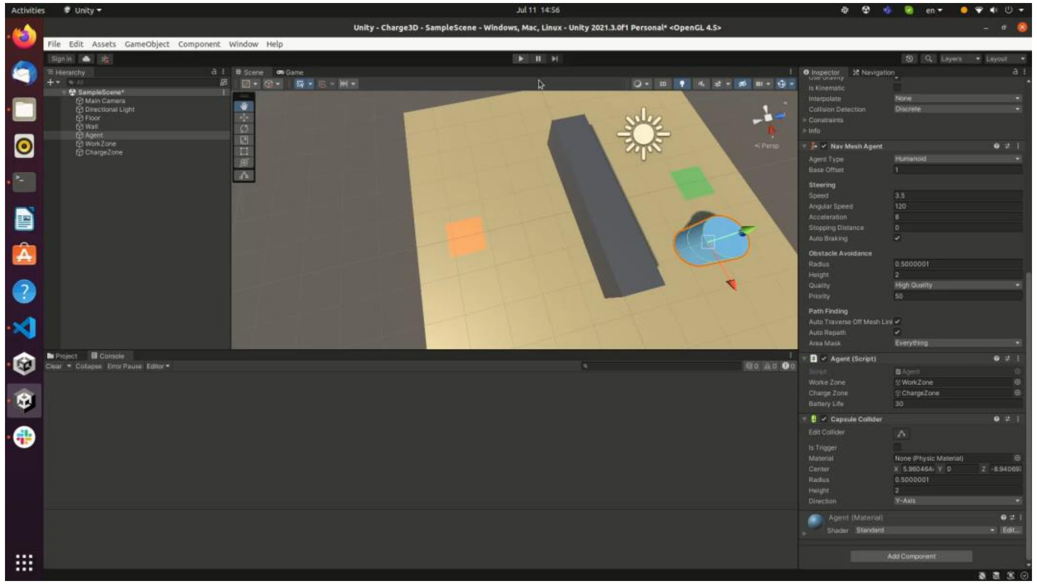

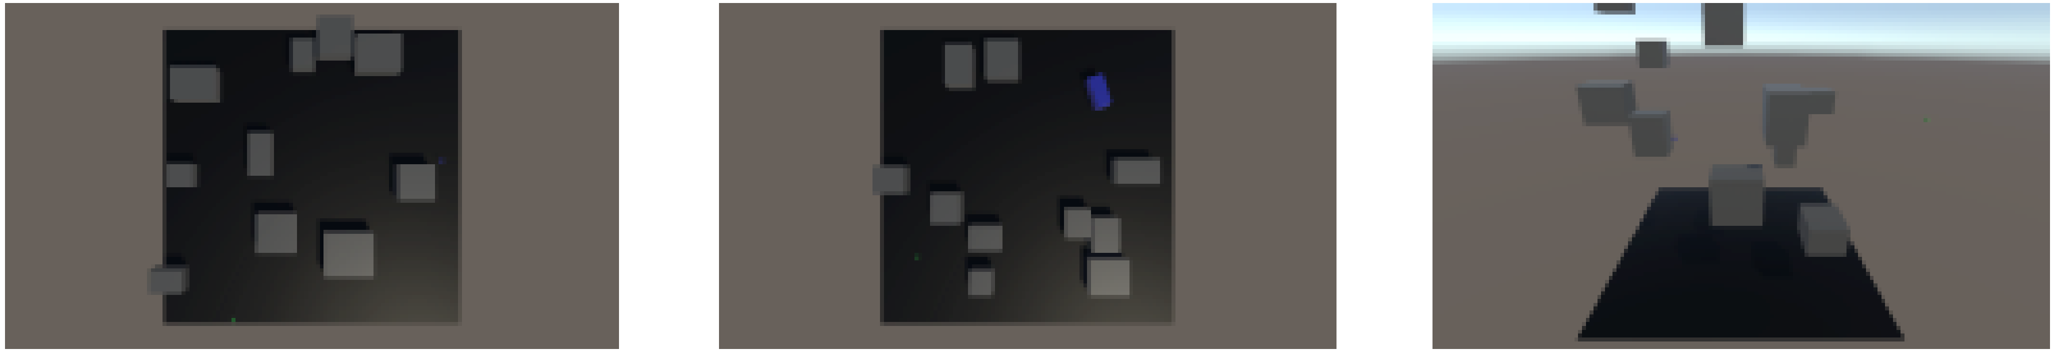

Deep RRT* in Unity

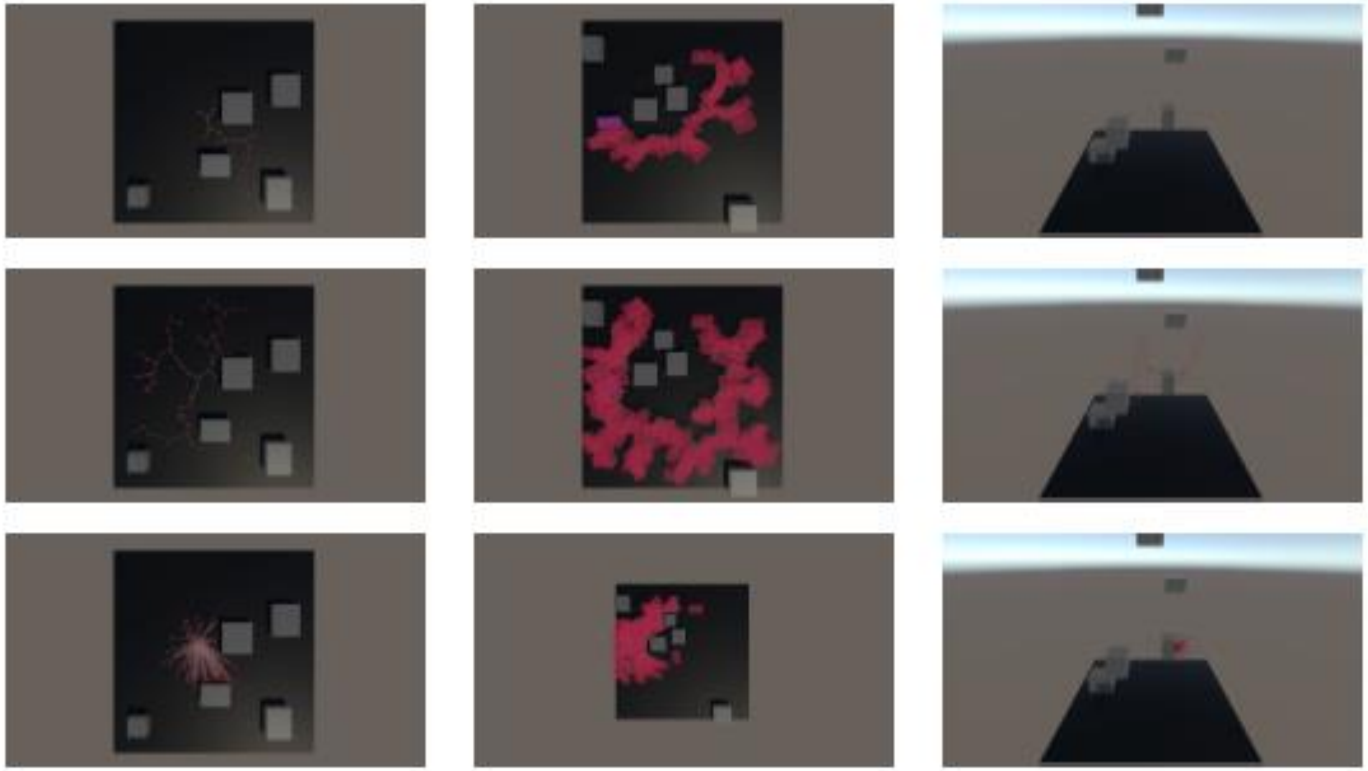

Critic/Actor and CLIP: Connecting text and images

Goal of Research: Steer a Robot via Natural Language

- Decompose Instructions by Large Language Model
- Merge the Modes for Language and Execution

- E.g. Move red ball to/from gray area

Uses Critic/Actor model to decompose and OpenAIs CLIP ot connect text with images.

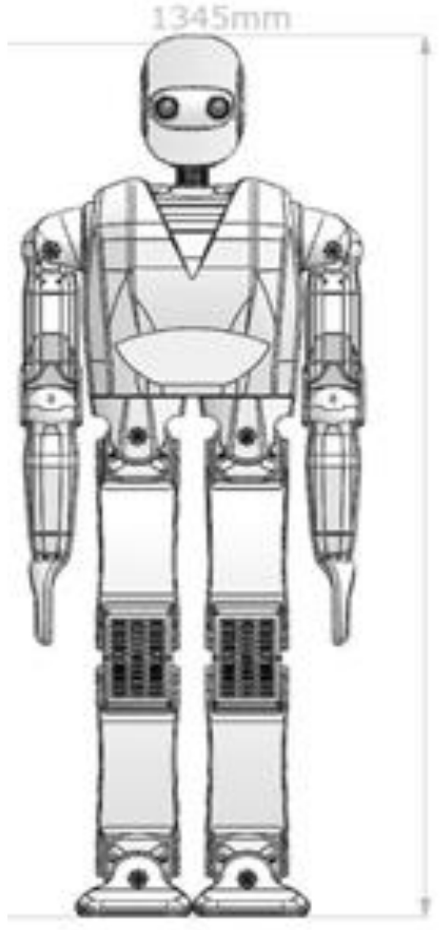

Towards the Exam…

Exams Dates / Room

- 1st Exam May 31st / **June 1st**
  - **10:00-noon (12:30)**
- 2nd Exam **June 7th**/ June 14th
  - **10:00-noon (12:30)**

- Room: KNE-301 for 1.6. and KNE-126 for 7.6.
- Building E KN
- Signing in for the exam available in KOS System

Instructions also available in the PUI courseware Intranet

https://cw.fel.cvut.cz/wiki/courses/pui/lectures#lecture_plan_20222023

Requirements

Requirements, Points, and Language

- For attending an exam, you need to fulfill requirements for **Zápočet** (Check your KOS).
- Total number of points possible to get from the exam is **50.**
- You can get maximum of **50** points from the **written part**

Language

- The questions in the written exam are in **English** but you can ask the examiner for an explanation in Czech. You can answer them either in **Czech or English.**

Exam Layout

- Cover Page, name on every page Full 2h (Theoretical Part: approx 15+15 = 30 min / Practical Part: approx 90min)
- Example (Heuristics, STL/VI/MCTS)
- Scoring: 50 Points in Assignment, 30 Points in Practical Part (20/Misa+10/Jan) +
- 20 Points in Theoretical Part (10/Tonda+10/Stefan)
- Grades A >= 90 Points, B >= 80, C >= 70, D >= 60, E >= 50, otherwise F
- All Results will be put and available to into BRUTE
- 28th 2nd Assignment.
- Resits (if fail, likely a third exam)

Content Theoretical Part: Q/A

Classical Planning [Tonda]

- Intro, Formalisms, PDDL
- State-Space Search and Heuristics
- Relaxation Heuristics
- Landmarks in Planning
- Potential Heuristics and Abstractions

Extended Planning [Stefan]

- Nondeterministic Planning
- Probabilistic Planning
- Monte-Carlo Search (UCT / NRPA)
- Metric and Temporal Planning, Preferences

Sample Questionaire Non-Deterministic & Probabilistic Planning

What is an MDP?

How to Solve MDPs (Name three approaches)

What are weak, strong, strong cyclic plans?

What is VI and what does it have in common with SSSP?

How efficient to solve?

State the Bellmann Equation of VI?

Sample Questionaire Monte-Carlo

- Which kind of Machine Learning, where successful?
- Name 4 stages of MCTS
- What is the Function of Exploration-Exploitation in UCT?
- Theory behind UCT
- What is NRPA and when is it useful?
- Number of rollouts for k Iterations and depth d

Sample Questionaire: Temporal Planning

Temporal Planning in PDDL, Snap Action Concept

What is a Simple Temporal Network?

How do the Constraints look like?

How do STN apply to Temporal Planning?

How To determine Consistency, Earliest Starting Time?

Which Algorithm to Choose to solve STNs and How Efficient?

Content Practical Part: Hands-On Examples

1st Part [Misa]

- State Spaces in Planning
- Heuristics:

  h_add, h_ff, h_lm, h_pot, h_abs, h_pdb

2nd Part [Jan]

- MDPs, Value Iteration and Variants
- Simple Temporal Networks
- Monte-Carlo Search

What is Allowed/Requested to Bring to the Exam

- Calculators. No Computer, No phones.
- Pen/Pencil (Mandatory!)
- Additional Pages provided from the tutors

- Jan and Misa will be supervising the exams
- Grade for the exams and final course marks, reporting date

→ Soon (around 1-2 week after final exam)

Good Luck alias
Deserved Success

Announcement: May 25, 2023 2:30pm – 3:30pm

- **Card games - the future of research**
  Many board games like Checkers, Nine-Men-Morris have been solved, or, as in Chess, Shogi or Go, computers outperform humans. Therefore, research attention has shifted to card games. After some variants of Poker have also been solved to a satisfying degree, trick-taking games like Skat or Bridge have been identified as current AI game playing challenges. One obstacle is that ---facing the larger number of tricks and degree of uncertainty--- the application of deep and reinforcement learning appears to be less apparent than in perfect-information games. In Spring 2022, we won an online tournament with 20 series of 36 games against a top player.