



# BlueSky: How to Raise a Robot – A Case for Neuro-Symbolic AI in Constrained Task Planning for Humanoid Assistive Robots

Niklas Hemken  
niklas.hemken@student.kit.edu  
Karlsruhe Institute of Technology  
Karlsruhe, Germany

Florian Jacob  
florian.jacob@kit.edu  
Karlsruhe Institute of Technology  
Karlsruhe, Germany

Fabian Törnava  
fabian.peller-konrad@kit.edu  
Karlsruhe Institute of Technology  
Karlsruhe, Germany

Rainer Kartmann  
rainer.kartmann@kit.edu  
Karlsruhe Institute of Technology  
Karlsruhe, Germany

Tamim Asfour  
asfour@kit.edu  
Karlsruhe Institute of Technology  
Karlsruhe, Germany

Hannes Hartenstein  
hannes.hartenstein@kit.edu  
Karlsruhe Institute of Technology  
Karlsruhe, Germany

## ABSTRACT

Humanoid robots will be able to assist humans in their daily life, in particular due to their versatile action capabilities. However, while these robots need a certain degree of autonomy to learn and explore, they also should respect various constraints, for access control and beyond. We explore the novel field of incorporating privacy, security, and access control constraints with robot task planning approaches. We report preliminary results on the classical symbolic approach, deep-learned neural networks, and modern ideas using large language models as knowledge base. From analyzing their trade-offs, we conclude that a hybrid approach is necessary, and thereby present a new use case for the emerging field of neuro-symbolic artificial intelligence.

## CCS CONCEPTS

• Security and privacy → Access control; • Computer systems organization → Robotics; • Computing methodologies → Neural networks; Robotic planning.

## KEYWORDS

Humanoid Robots, Robot Task Planning, Activity-Centric Access Control, Deep Learning Based Access Control, Large Language Models, Neuro-Symbolic Access Control

## ACM Reference Format:

Niklas Hemken, Florian Jacob, Fabian Törnava, Rainer Kartmann, Tamim Asfour, and Hannes Hartenstein. 2024. BlueSky: How to Raise a Robot – A Case for Neuro-Symbolic AI in Constrained Task Planning for Humanoid Assistive Robots. In *Proceedings of the 29th ACM Symposium on Access Control Models and Technologies (SACMAT 2024)*, May 15–17, 2024, San Antonio, TX, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3649158.3657034>

## 1 INTRODUCTION

With an intensifying labor shortage in the care sector, assistive humanoid robots will most likely become a necessity in the future of our aging society. These robots will need to invade the most private spaces of humans they care for, thus safety, security and privacy issues in this field are of utmost importance. The potential of humanoid robots to assist humans lies in the ability to learn whatever is needed for assistance. However, ‘whatever is needed’ has to be restricted to safeguard safety, security, and privacy policies and preferences. The challenge for classical approaches, i.e., based on symbols and logical formulas, to ensure constraints lies in the task universality of humanoid robots: While standard robot designs like robotic arms or drones are limited to specific tasks, humanoid robots are intended to be on par with human beings regarding the variety of tasks they can solve. This wide range of tasks and deployments favors approaches in which the robot is granted increased sovereignty and should learn its authorizations in the field. Learning promises scalability of problem complexity up to task universality, while keeping manual specification complexity manageable for humans. However, learning changes the nature of access control from certain enforcement to probable observance.

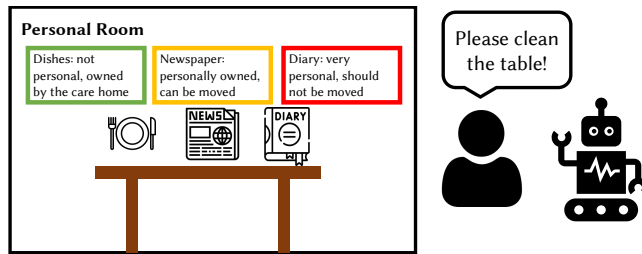
Future use cases of assistive robots in care home will require complex systems of requirements regarding safety, security and a lot of personalized requirements when directly interacting with patients. While some of these requirements are prescribed in policy documents like the official instructions for human care-home workers accessible to the robot via manual translation or machine learning, some requirements are purely based on “common sense” and a common understanding of social norms not compiled in a machine-readable way. An example of such norms is what objects a human considers as private belongings and what objects can inoffensively be interacted with.

In this Blue Sky paper, we investigate how symbolic and neural methods to robot task planning can ensure symbolic and neural constraints beyond access control, and how constraints from policy documents and social norms can be inferred and observed using large language models. We discuss the inherent trade-offs of symbolic and neural standard methods, which mark the endpoints of a neuro-symbolic spectrum. We propose neuro-symbolic hybrid methods as vision to bring assisted living with humanoid robots into practice.



This work is licensed under a Creative Commons Attribution International 4.0 License.

SACMAT 2024, May 15–17, 2024, San Antonio, TX, USA  
© 2024 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-0491-8/24/05.  
<https://doi.org/10.1145/3649158.3657034>



**Figure 1: An assistive robot is tasked to clean the table in the personal room of a care home resident. The challenge is to integrate an abstract “do not move private objects” policy as specific constraints into robot task planning. While privacy has a common understanding, it is a subjective preference that needs ad-hoc assessment by the robot. On the table are care home-owned dishes, thus not personal. The newspaper is personally-owned but contains no private information. The diary is personal and private.**

As toy example illustrated in Fig. 1, a robot is tasked to clean a table in a personal room, for which the robot now needs a plan. It considers whether it is allowed to move objects on the table, i.e., the authorization to move an object is a precondition of the action for moving that object. We assume constraints that formalize policies like “do not move personal items”, which allow, e.g., the removal of care home owned dishes, but disallow moving a diary. To put user-instructed policies like that in concrete terms, the robot needs a kind of policy information database or knowledge base to decide whether an action with an object is actually constrained by a policy.

In a first attempt to tackle this problem of balancing sovereignty and ‘obedience’, we proposed a joint modeling framework for robot task planning and access control [1]. By directly incorporating access control into task planning, robots are unable to even “think about” forbidden behavior. This paper is a continuation of our poster abstract [10]. While the fit between the notion of an activity in Activity-Centric Access Control [9] and task planning was shown in theory, we now report on preliminary practical insights. Recently, machine learning-based approaches to assist with or even perform access decisions have been proposed, as seen with Machine-Learning- and Deep-Learning-Based Access Control [20]. We refer to further related work in the following sections. Since the rate in which papers on deep learning are recently published is ‘breathtaking’, we cannot claim completeness, but focus on scientific works that substantiate our claims. Of course, the topic of our paper is also related to the wider topic of ethical artificial intelligence.

This BlueSky paper is structured as follows: in Section 2 we introduce terminology for planning on a symbolic and signal-based level and formulate the problem statement of planning under constraints. While the problem formulation is generic, we focus in the following on constraints that result from access control, security, and privacy policies, and their integration with robotic planning. In Section 3 we discuss methods we consider as current methods, while in Section 4 we show what we expect for the future. We draft a roadmap in Section 5 that goes beyond any single research project and conclude in Section 6.

<sup>1</sup>In ancient Greece, “symbolons” were objects like bones broken and distributed between parties. The term “symballo” refers to later verifying whether the pieces, i.e., the distributed “sub-symbolons”, fit together for mutual authentication of parties. [16]

## 2 TERMINOLOGY & PROBLEM STATEMENT

A prompt is an instruction given by a user to perform a task to reach a certain goal. A signal-based prompt is, e.g., a user-written text or a microphone recording. A symbol-based prompt is formalized and refers to abstract symbols, like subjects, objects, and locations with attributes. A plan is a sequence of actions to be performed by the robot to fulfill the task and reach the goal based on its current initial state. A symbol-based plan is a sequence of abstract actions with their pre- and postconditions (effects) that is converted to a signal-based plan in the form of actuator commands needed for execution. Planning is the problem of finding such a sequence of actions. Symbolic planning does so by reshaping logical formulas and needs a symbol-based prompt as input to output a symbolic plan. Symbolic planning can be implemented either with classical search algorithms or symbolic artificial intelligence, but always uses a localist data representation where one symbol is represented in one discrete variable in information processing. Neural planning is the forward pass through a deep neural network that learned to plan and can either use a symbol- or signal-based prompt, to either create a symbol- or signal-based plan. End-to-end planning infers a signal-based plan from a signal-based prompt.

Neural information processing usually operates on a distributed data representation where one specific representation of a symbol in the respective feature space is mapped to not one single neuron, but a continuous activation pattern of many neurons<sup>1</sup>. The distributed data representation is the key that allowed deep neural networks to advance past symbolic methods in terms of generalizing from their learning material and flexibility.

We say that policies are abstract ideas of what is allowed and what is off limits, from individual user preferences over administrative instruction documents to non-written social norms. We define constraints as machine-executable representations of policies that have to be considered by the planner to produce plans that adhere to the constraints. A symbolic constraint is a constraint in form of, e.g., a first-order logical formula, as in classical access control. A neural constraint is a constraint in form of a deep neural network.

Solving planning problems under constraints is a classical topic. But when scaling up the complexity of tasks, constraints, and environments to task universality, classical methods might not deliver the performance required for real-world use cases like humanoid robots in assisted living scenarios. With task-specific robots, the available action options are typically rather limited and allow for manual specification of tasks, constraints, and environments. With task-universal robots, i.e., robots that are able to mimic human actions, the large space of potential actions leads to an most likely unacceptable overhead in manual specification of constraints as well as in runtime performance. As we will explicate in Section 3, however, neural planning struggles to actually ensure safety, security, and privacy constraints in its generated plans. Our thesis is that task-universal planning under safety, security, and privacy constraints requires a *hybrid* method, for which we see large language models, due to their proficiency in a broad range of tasks and internalized “common sense”, and approaches from neuro-symbolic artificial intelligence, due to their feature of bridging the gap between localist and distributed data representation, as promising research directions (c.f. Section 4).

### 3 ANALYSIS OF CURRENT METHODS

#### 3.1 Constraint-Ensuring Symbolic Planning

In symbolic planning, a planner finds a symbolic plan by searching through possible robot actions, trying to find transitions that transform the current state of the robot and its environment to the goal state inferred from the user prompt. In practice, the de-facto standard for this is the Planning Domain Definition Language (PDDL) [11], which we used in the following. The initial state is the current state of the robot and its environment, the goal state is the goal that the prompt described as target, and the transition rules are the actions the robot is capable of. The symbolic planner then tries to find a plan, i.e., a sequence of transitions that transforms the initial state into the goal state. Actions have preconditions required to be met before execution. Typical preconditions are that the robot has to be at a table and empty-handed to grab an item, or that an item has to be unobstructed to be grabbed. Transition preconditions for a symbolic planner are the equivalent of access control policies for a symbolic reference monitor. From an access control perspective, incorporating safety, security, and privacy constraints in form of preconditions is the most natural approach to generate plans that are aware of such constraints, exemplified in Listing 1.

```

1 (:action clean_from_table
2 :parameters
3   (?robot - robot
4    ?table - location
5    ?obj - on_table
6    ?remove - location)
7 :precondition (and
8   (non_personal ?obj)
9   (at ?robot ?table)
10  (at ?obj ?table)
11  (remove_loc ?remove)
12 )
13 :effect (and
14   (not (at ?obj ?table))
15   (at ?obj ?remove)
16 ))
17 ...
18 (:init
19   (at robot start)
20   (at newspaper table)
21   (at diary table)
22   (at dishes table)
23   (non_personal dishes)
24   (non_personal newspaper)
25   (personal diary)
26   (remove_loc remove))
27 (:goal
28   (forall (?obj - on_table) (or
29     (and (non_personal ?obj) (at ?obj remove))
30     (and (personal ?obj) (at ?obj table))
31   )))

```

Listing 1: Excerpt from our PDDL domain description

The abstraction of a PDDL transition corresponds to the abstraction of an activity in Activity-Centric Access Control (ACAC, [9]), whereby these concepts can be naturally combined. This workflow is illustrated in Fig. 2. We identify three different phases when designing such a planner. Only during runtime we automatically generate plans, while manual work is needed during development

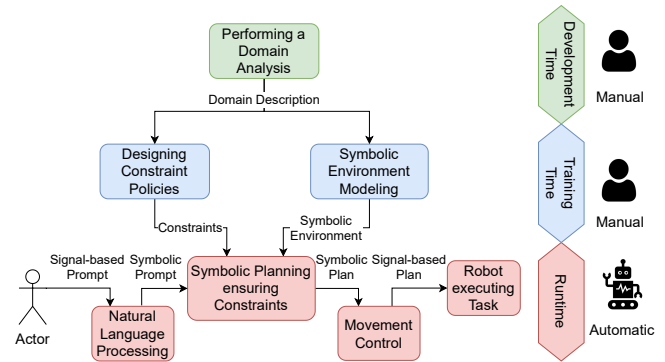


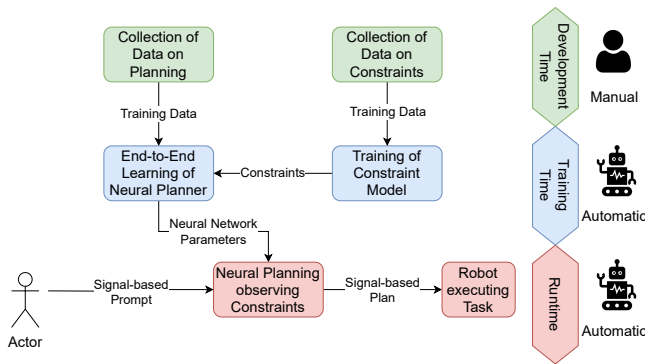
Figure 2: Workflow of designing a robot task planner using symbolic planning and symbolic constraints. We differentiate between three phases: In the development phase, the rules for the constraints are designed and the modeling for the symbolic planner is done. The training phase consists of the incorporation of such constraints into the planner. During runtime, the plan is generated and executed.

and training time. We were able to map<sup>2</sup> ACAC preconditions directly into PDDL preconditions. Resulting conditions of ACAC activities could be represented as *effect* in PDDL.

However, ACAC’s contextual conditions and current conditions, which are checked during the execution of an action, are not directly transferable. Contextual conditions require PDDL extensions to query external data. Current conditions require PDDL extensions that allow temporal planning and durative, interruptible actions. While it is possible to incorporate toy examples of ACAC into PDDL, we question the scalability up to the scenario of humanoid robots in a care home without requiring an equally complex manual specification. Every possible action, every possible object and every ACAC policy needs to be considered, inherently prohibiting such systems from being task-universal. While symbolic artificial intelligence can learn incrementally in the field and can thereby reduce manual specification complexity while still being understandable and tunable by humans, adapting symbolic artificial intelligence to a broad range of situations is still a ‘very manual’ process.

By relaxation from constraint-ensuring to constraint-observing planning, recent developments in the field of machine-learning based access control can be employed. Feasibility of access decisions performed by machine learning methods such as neural networks was, e.g., shown by Nobi et al. [21]. However, current systems show quite large error margins, reaching a precision of about 0.9 in an access control scenario for mobile phone apps. Nobi et al. propose the *Deep Learning Based Access Control* model, where a neural network directly performs access decisions. Such systems are mostly trained on access logs and further user and resource metadata. To combine symbolic planning with a learned access control model, the planner queries the learned model in addition or instead of symbolic constraint formulas, and uses a learned heuristic function that classifies branches of the state space as dead ends if a constraint is violated.

<sup>2</sup>Full artifacts available: <https://github.com/kitt-dsn/how-to-raise-a-robot-beyond-ac>



**Figure 3: Workflow of designing a neural planner that satisfies neural constraints. The development phase now consists only of data collection. The training phase includes the actual training of the deep neural network, while the neural constraints are being fed into the training of the planner. During runtime, a forward pass through the network is performed.**

### 3.2 Constraint-Observing Neural Planning

Deep-learned neural planners promise to be scalable to complex deployments, and are the natural counterpart to deep-learning based access control (DLBAC). When combining a neural constraint model with a neural planner, one gets a workflow as illustrated in Fig. 3. Again we identify three different phases. Compared to symbolic planning with symbolic constraints, i.e., classical access control, we can now perform the training phase automatically. The computational hardness of learning access control policies is already well-studied [15]: e.g., for typical access control models like access control matrices or role-based access control, there are problem formulations for which polynomial learning algorithms exist. Only the development phase needs manual work, which improves scalability.

A policy change that happens after initial policy definition necessitates incremental learning of the neural constraint model, i.e., in contrast to learning only from data all at once in the initial training phase, it has to learn new constraints from sequential data [17]. However, incremental learning is prone to catastrophic forgetting: On learning new data, neural networks have the tendency to disruptively forget previously learned knowledge [23]. Dealing with catastrophic forgetting has been an ongoing challenge for more than thirty years. While natural neural networks typically do not exhibit catastrophic forgetting, no general solution for artificial neural networks has been found yet [7]. In general, due to catastrophic forgetting, changes in the policy make retraining of the neural planner and neural constraints necessary. Still, ways to tackle policy changes in a manageable way have been shown regarding DLBAC administration [20]. The main idea is to retrain the networks in small steps, only using training data that represents the update that should be introduced. This approach intends to only change the neural networks to incorporate the new data from the update, but this cannot be guaranteed. One could also consider incorporating neural constraints into the reward calculation of reinforcement-learning based planners.

One can also combine neural planning with symbolic activity control. Analogously to the previous paragraph, but unconventionally for Generative Adversarial Network (GANs), one might

use symbolic constraints as discriminators to train the end-to-end learning-based neural planners. It is challenging to provide a distributed representation of constraints since symbolic constraints are localist representations. As approximation, a first step can be realized by using the differentiable loss function provided by existing deep learning based constraint systems, which can be used to train the network. Although this approach might not be as scalable as learning neural constraints in complex deployments due to manual policy engineering, it is more efficient to administrate since the symbolic constraints do not need to be retrained after a change. For our running example, we now need to incorporate policies on what are personal items again, which makes the whole process inherently more complex. Even though the neural planner learns its behavior, extensive policy engineering needs to be done. Once the training of the end-to-end learning-based planner is finished, this approach behaves similarly to using a deep-learning based activity control system: Given a prompt, the robot performs a forward pass through the neural planning network and receives a plan that probably observes the desired constraints.

## 4 ANALYSIS OF UPCOMING METHODS

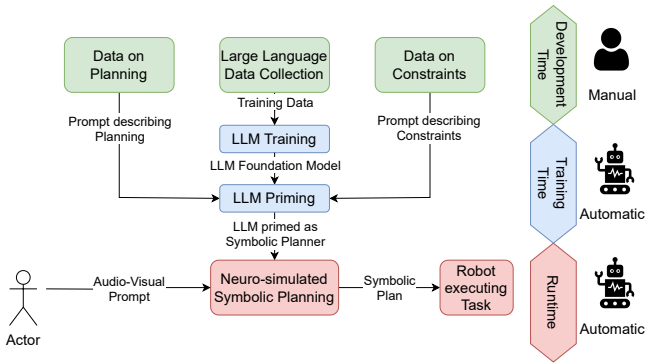
While symbolic and neural constraints are quite different compared to each other, the process of integrating any type of constraints in either symbolic or neural planning is similar. Therefore, one is not limited to choosing only one combination. We instead argue for a hybrid approach to combine the best of both worlds: one can use neural planning and neural constraints for versatility and universality, but ensure critical policies as symbolic constraints during runtime. For critical situations with limited universality, one can also fall back to symbolic planning to find an accurate solution slowly. Neural systems could also serve as recommendation engine for new policies that enhance the existing logical solution.

Recently, Large Language Models (LLMs) became a popular method for artificial intelligence. Especially the model ChatGPT by OpenAI [22] caught the attention of the general public, due to its seemingly close resemblance of human conversation while also being able to give symbolic answers, like solving a task in a programming language or even PDDL. Due to their proficiency in processing and generating both signal- and symbol-based input and output, they are a natural building block for hybrid systems.

Neuro-symbolic artificial intelligence stands for the combination of symbolic and neural approaches to solve problems, and promises the “best of both worlds”, i.e., combining the advantages of both symbolic and neural methods. Thus, the popularity of neuro-symbolic methods has risen in recent years [24], and those methods promise to bring together the flexibility of neural planning with guaranteeing constraint satisfaction of symbolic planning. We discuss LLMs using our toy example next, and then generalize to discuss neuro-symbolic methods to constraint-ensuring universal task planning in the following.

### 4.1 LLMs as Planner and Knowledge Base

While LLMs represent a break-through for dialogues with humans, LLMs are capable of solving many kinds of problems given the right prompt. Due to this versatility, LLMs are foundational models, from which a problem-specific model is derived through priming. The



**Figure 4: Workflow of designing a robot task planner using a multi-modal large language model as foundation model. In the development phase, the prompts for priming are manually engineered.**

LLM can act on the priming because of the breadth and sheer size of its training data, which includes descriptions of planning problems, but also allows the models to generalize knowledge between domains – for an LLM, solving a planning problem like the classical wolf, goat and cabbage river crossing problem in English is of no fundamental difference than solving it in another natural language, or in a formal planning language like PDDL.

We now describe an experiment with priming LLMs to behave like symbolic planners to solve constrained planning problems. We can see such a workflow illustrated in Fig. 4. With recent advances in multimodal large language model, we envision the model to operate not only on textual information, but also on audio-visual user prompts. We tested the primed LLM approach<sup>1</sup> by formulating our running example as natural-language prompt, and used ChatGPT 4 as constraint-observing neural planner to generate natural-language plans, shown in Listing 2. Specifically, lines 1 and 12 reduce ChatGPT to a constraint-observing neural planner by making it emulate a constraint-ensuring symbolic planner. Lines 2-8 simulate sensors and memory, lines 9-11 instruct ChatGPT with a goal and constraints. We obtain a plan that reaches the goal of cleaning the table, while not interacting with the most personal item on the table, the diary. This experiment shows that we can use a GPT-based LLM to do neural planning while observing neural constraints inferred from natural-language prompts.

One interesting effect of using LLMs for constraint-observing planning is their performance on problems that require knowledge of *common sense*. Our toy example delivered an insight into such behavior. We tasked the LLM-based planner with only handling objects that are not considered private. From experience and intuition, humans know that a diary is a very private object and dirty dishes are not, but machines need to encode such information somehow. When using LLMs as planner, such information does not need to be given explicitly to the planner, because the LLM can infer it from its training data. Since it is trained on extensive textual data that encodes a lot of human interaction, such social norms are known. When asked whether a diary is private, a *yes* has a higher probability than a *no*. The LLM therefore not only acts as a symbolic planner, but also as knowledge base for common sense and object attribution for the simulated planner. This integrated knowledge is of tremendous value for systems deployed in real

- 1 Suppose you are a symbolic planning system that generates plans for task-universal assistive and humanoid robots based on a domain description that is provided below.
- 2 Domain:
- 3 – Care home that uses robots to care for their residents. The care home has lots of personal rooms for the residents, common spaces like a dining area, and administrative spaces like a kitchen or offices
- 4 – Personal room of a resident, that has a table inside
- 5 – on the table are the following things:
- 6 – newspaper that was paid for by the resident
- 7 – dirty dishes that are owned by the care home and are cleaned by the care home
- 8 – a diary that most certainly contains very private information
- 9 Goal: Remove as many objects from the table as possible
- 10 Constraints: Do not interact with objects that you would consider personal.
- 11 Possible actions: Move, analyze an object, take an object, place an object
- 12 Put out the plan in terms of a sequence of actions the robot performs. Put out the action you used and the object you interacted with or where the object was placed. You can only use the actions from the list above. Put out a list of actions and objects that the plan consists of. Summarize the reasoning for the overall plan as a single text paragraph. The plan must be conform with all given constraints.
- 13 # Exemplary output from ChatGPT 4
- 14 Plan Sequence:
- 15 1. Analyze object: newspaper
- 16 2. Analyze object: dirty dishes
- 17 3. Analyze object: diary
- 18 4. Take object: dirty dishes
- 19 5. Move to: kitchen or designated cleaning area
- 20 6. Place object: dirty dishes in cleaning area or dishwashing station
- 21 Reasoning:
- 22 Given the constraints and the goal, the robot needs to first distinguish between objects that are personal and objects that belong to the care home. Analyzing each object ensures the robot can determine its nature and importance. The newspaper, although bought by the resident, could be seen as potentially non-intrusive. However, it was paid for by the resident which could imply some personal value, so it's left untouched. The diary is most certainly private, so it's not touched. The dirty dishes belong to the care home and must be cleaned, so they are the only objects removed from the table, in line with the constraints set.

**Listing 2: A prompt to simulate a symbolic planner for a basic domain description, and the ChatGPT 4 answer. Shared chat transcript at: <https://chat.openai.com/share/77d31c8a-239c-4d57-a96c-f70078f88503>**

world applications, because such social norms will no longer need to be encoded directly, if a system should obey to them. Especially humanoid assistive robots act in very sensitive environments and will need to follow a very complex set of rules. While some of these rules will be available in the form of certain documents, some rules will only exist in the form of social norms, for which no single “ISO norm” document exists. LLMs provide a solution to encoding such norms, without extensive work on building a symbolic version of them. Thus, LLMs serve as knowledge bases for such systems.

For any kind of constrained task planner, quality properties of generated plans are crucial for their usefulness in practice, like whether generating plans are syntactically valid, semantically correct, consistent between multiple consecutive planner executions, and stable, in the sense that similar situations should lead to similar plans. Symbolic planners scrape the possibility space until they find a valid and correct plan, and if a planner returns a solution, it can guarantee its quality properties. However, solution quality of LLMs is an open problem [19], and for LLMs instructed to behave like a symbolic planner, one can find examples where neither of

1 A farmer with a pig, a goat, and a cabbage must cross a river by boat. The boat can carry only the farmer and a single item. If left unattended together, the goat or the pig would eat the cabbage. How can they cross the river without anything being eaten?

2 # Exemplary output from ChatGPT 4

3 The puzzle you've described is a classic river crossing problem, though it's typically presented with a wolf, a goat, and a cabbage. The solution to this problem, where the pig or the goat would eat the cabbage if left unattended, is as follows:

4 1. The farmer takes the goat (or pig, it doesn't matter which) across the river first and leaves it on the other side.

5 2. The farmer then goes back alone to the original side.

6 3. He takes the cabbage across the river next.

7 4. After leaving the cabbage on the other side, he takes the goat (or pig) that was left there during the first trip back with him to the original side.

8 5. He leaves the goat (or pig) on the original side and takes the pig (or goat, the one he hasn't taken across yet) to the other side.

9 6. The farmer leaves the pig (or goat) with the cabbage on the other side and goes back to the original side to get the goat (or pig).

10 7. Finally, he takes the goat (or pig) across the river.

11 Now, all three items – the pig, the goat, and the cabbage – are safely on the other side of the river. At no point were the cabbage and the animals left unattended together, so nothing was eaten.

**Listing 3: A river crossing puzzle as stated in Wikipedia [26], but with the wolf replaced by a pig. The ChatGPT 4 solution leaves either of the herbivores alone with the cabbage. Shared chat transcript: <https://chat.openai.com/share/b9eef797-a4ad-4b50-a9d1-3cc294e744f6>**

those properties are fulfilled. In our experiments, the dishes were consistently moved to the kitchen, while the newspaper and diary were always considered too personal to interact with. However, with ChatGPT 3.5, the newspaper was sometimes interacted with and sometimes not, showing probabilistic inconsistencies not exhibited by a (deterministic) symbolic planner. These inconsistencies stem from the fact that the LLM only simulated a symbolic planner using neural methods. The LLM chose among the predicted text continuations with the highest probability, and if two alternatives have no clear winner, the solution will vary, leading to inconsistent plans. Planning-primed LLMs produce outputs that *sound like* a plan, but that are not necessarily a correct or even valid plan. As a well-known example, Listing 3 shows a prompt for a variation of the river-crossing problem usually formulated using a wolf, a goat, and a cabbage, which is a common way to test the limits of LLMs (c.f., e.g., Gerhard Paaß in [8]). Here, the wolf is replaced with a pig, and the constraints are stated accordingly that neither of the animals can be left unattended with the cabbage. While ChatGPT 4 confirms this fact, its plan does not ensure this constraint, as it leaves one of the animals alone with the cabbage, due to missing symbolic understanding of the problem. Instead, it provides a solution similar to the regular solution with the wolf, due to its prevalence in the training data. When compared to standard approaches of Section 3, likelihood-based text completion in the form of LLMs exhibit a form of task universality that has high potential for meeting practical task and policy complexity needs of humanoid robots [4, 5]. However, ?? demonstrates that inspiring results in terms of complexity are offset by disillusioning solution quality properties, when compared to quality standards set by standard symbolic planning and access control methods that can guarantee for their results.

Lvl	Description based on Garcez et al.[6]	Example for Constrained Robot Task Planning
1	Standard neural networks deep-learned for using symbolic input and output	Learned neural network planner that produces symbolic plans from symbolic prompts
2	A neural network is loosely coupled with a symbolic problem solver	A LLM is used as knowledge base for a symbolic constraint system
3	A neural network focusing on one task is interacting with a symbolic system performing a complementary task	A neural planner produces symbolic plans, which are verified post-planning to ensure symbolic security and privacy constraints through a symbolic constraint system
4	Symbolic knowledge is given into the training set of a neural network	The training data is generated by a simulator based on symbolic domain and constraint representations
5	Symbolic logic is mapped onto an embedding which acts as a soft-constraint on the network's loss function	The loss function of a neural planner has an embedded symbolic constraint system, penalizing illegal plans and emphasizing legal ones
6	True symbolic reasoning inside a neural engine	A neural planner runs completely on a symbolic representation of the domain and its constraints

**Table 1: Levels of neuro-symbolic integration according to Kautz [13] and their respective realization for robot-task planning.**

While LLMs worked well for showcasing versatile planning behavior based on the knowledge of social norms in our toy example, the river-crossing problem example confirms the need for symbolic reasoning to get correct solutions. The emerging field of neuro-symbolic artificial intelligence acts as a promising direction to combine versatility and quality properties like correctness and consistency [18], which we look into next.

## 4.2 Integration of Neural and Symbolic AI

The term *neuro-symbolic AI* describes systems that combine symbolic with neural approaches. The above-mentioned hybrid systems are in this category. Neuro-symbolic AI systems are differentiated on the extent of the integration of symbolic data representation in the neural process. Kautz [13] differentiates between six different types of neuro-symbolic integration (c.f. Table 1). This classification ranges from standard deep learning or large language models, where the input and output can be symbolic, but the inner workings are not, to type six integration, where full symbolic reasoning is possible within the neural engine. We especially want to emphasize how the integration of symbolic reasoning into the training process happens on each level. For our analysis, we consider each level and give an example for our constrained task planning scenario.

Our first experiment using a symbolic planner and PDDL from Section 3.1 has no neural reasoning at all, i.e., provides no neuro-symbolic integration, and therefore it is a “**level zero**” system. The same is true for our end-to-end learned neural planner architecture from Section 3.2, as no symbolic representation is involved.

Kautz’ categorization starts with **level one**, which is classical deep learning trained on symbolic input to produce symbolic output. Since there is no symbolic reasoning inside the system, data representation is purely distributed inside the neural engine. The integration of constraints may be possible by only using valid plans as training data, hoping to not generate invalid ones. The LLM system from Section 4.1 would fall in this category if prompted with symbolic inputs to produce symbolic outputs, e.g. in PDDL.

**Level two** consists of loosely coupled systems. For example, a symbolic problem solver can use a neural system as some kind of sub-routine, as stated by Kautz [13]. In our constrained task planning case, this level could be reached with a symbolic planner like in our first experiment, enhanced to use a LLM as knowledge base. The sub-routine would allow the symbolic planner to check an item for its level of privacy and the societal knowledge encoded in the LLM. It is important to note that the training process of both parts is completely decoupled, i.e., the symbolic part has no influence on the training of the neural part and vice-versa.

The **third level** is, e.g., the “do not cross the red line” approach. We again have a distinct neural and symbolic system. On this level, the output of one is given as input to the other one. While on level two, the data flow was primarily handled by one of the systems, with the other one being just a sub-routine, it is now handled by both systems. A typical example is the pipeline of a neural computer vision system which produces the symbolic input for a symbolic planner, whose symbolic output is translated to motor signals by another neural system. In our scenario, an example would be a pipeline that feeds the symbolic plans of an LLM-based planner to a symbolic constraint enforcement system, as seen with Yang et al.’s LLM coupled with a linear temporal logic constraint enforcement [28]. An important distinction of level three in contrast to level two is that now both systems take part in a joint training procedure. Kautz mentions a feedback loop that flows back from the symbolic system to optimize the training of the neural system.

**Level four** represents systems where symbolic knowledge is put into the neural system by encoding it into synthetic training data. This approach uses symbolic knowledge on the problem domain, but also general symbolic knowledge of physics or mathematics applied to the domain, to generate valid and invalid solutions for many problem instances. Note that finding a problem instance for a given, randomly-generated solution, i.e., to generate a domain in which a given plan does not violate any constraints, might be much faster than the other way around, i.e., actual symbolic planning. An exemplary system on this level would be a neural constraint enforcement machine trained on synthetic decision logs. The training data would consist of access requests and corresponding access decisions, generated from a set of symbolic constraints that formalize the access policy. Analogously, a neural task planner could be trained with synthetic pairs of tasks and suitable plans. It is important to note the difference to training on real world access logs here: All synthetic training samples have symbolic knowledge inherently encoded. This is not necessarily the case on real world

data, where the inherent logic might be overlaid by perturbations, or not be present at all. Symbolically generated, synthetic training data further deepens the integration of the symbolic reasoning into the training process of the neural system. On level three, it was a feedback loop back into the training, now it is directly embedded.

On **level five**, one looks at systems where symbolic logic is directly mapped onto the loss function of the neural system. This level is especially challenging because the distributed data representation of a symbolic system is typically not differentiable. In our planning example, this approach would mean that for example a symbolic constraint enforcement is directly mapped into the loss function of the neural planner. Plans that are valid would be emphasized by it and plans that are not valid would be penalized. Since the symbolic logic is directly in the loss function of the neural system, we cannot distinguish the symbolic part in the training process anymore, it has become part of it.

The last step is further strengthening this integration of symbolic logic into the neural system. Such systems are described by **level six** neuro-symbolic systems, representing the deepest neuro-symbolic integration. On this level, a real distinction of neural and symbolic parts is almost impossible: The symbolic knowledge is deeply encoded in the whole neural system. Systems on this level might be compared to human brains, bridging the gap between automatic and instinctive part of the brain and the slow and logical part, according to Kahnemanns notion of Systems 1 and 2 [12]. State of the art neural systems correspond to the System 1 fast and instinctive part, and are not able to bridge the gap to also think slow and logical. This behavior is exemplarily demonstrated in Listing 3, where ChatGPT “reads over” the peculiarities of the specific problem instance, and instead provides a solution similar to the one for the common instance presumably more prevalent in its training data. Our example for this level is a neural planner that runs completely on a symbolic representation and its constraints. The vagueness of this example is because such system do not exist currently, and it is still unclear how they could look like in practice, or whether they are feasible at all. However, such level six systems would be learned symbolic reasoning machines. Due to their capability of “thinking, fast and slow” both in the neural and symbolic world, they would be a game-changer for constrained task planning. One would be able to completely explain the machines and guarantee their behavior, neither of which is currently possible using black-box deep-learned neural methods.

## 5 ROADMAP: CHALLENGES AHEAD

### 5.1 Neuro-Symbolic Integration of Constraint Satisfaction and Task Planning

On current systems, we observe that task planning for real-world robots scales best to complex situations using neural planners, while in the access control community, symbolic approaches are most prevalent to strictly enforce security and privacy constraints. Therefore, we see the challenge of neuro-symbolic integration for constrained robot task planning in integrating a neural planner with a symbolic constraint system. While deepening the integration of the symbolic and neural spheres is still an open research problem, we present goals and challenges for deep neuro-symbolic integration for constrained robot task planning in this section.

Deep-learned neural planning requires a distributed representation of plans instead of a localist representation as a list of task actions, in order to be able to learn in the planning space embedding. With end-to-end learning-based robot task planners, we get this distributed representation of plans for free. With sensor data as feature space, we can directly perform neural operations on plan-level, which is encoded by the internal weights of the neural network that represents the planner. To interact with a neural planner, constraints need a distributed representation as well. Recent developments in deep learning based access control promise analogous approaches to the distributed representation of end-to-end learning based robot planners [21]. The deep neural networks that perform access decisions encode distributed constraints in their internal weights, and provide a feature space on which such constraints can be learned. However, satisfaction of constraints in distributed representation can still not be guaranteed (c.f. Section 3.2).

Thus the integration target of constraint-ensuring neural task planning comes down to the following challenges that can be considered variations of the well-known signal-to-symbol gap challenges:

- How can a distributed representation of planning work together with a distributed representation of constraints?
- How can symbolic algorithms reason on suitable distributed data representations?
- How to learn symbolic knowledge on distributed feature spaces?

In this paper, we have described approaches up to level two in the previous sections. As a research community, we can currently build neuro-symbolic systems up to level three and four. For level three, we can use existing systems and plug them together, level four can be realized using synthetic training data generated from symbolic knowledge on planning and constraints. Obviously, the leap to level five is the challenge, however, it would also bring humanoid assistive robots on a new level that might be necessary for practical task-universal usage in the real world. Existing neural constraint enforcement systems such as DLBAC [21] show that the integration of such systems into the loss function of planners appears on the horizon of current research. Controlled text generation of LLMs is also a current research topic, addressing the question of how to make a LLM produce text of a specific kind, as seen with Dathari et al. [2] or Keskar et al. [14]. Such imposing of constraints is a promising fit for imposing constraints on planners to use LLMs in some form, as knowledge bases or even directly as planners. While these controllers are also neural networks as of today and therefore are level four systems at most, an approach with symbolic controllers would be a level five neuro-symbolic AI.

A type six neuro-symbolic system that utilizes truly symbolic reasoning inside a neural engine would make many problems of constraining task-universal robots disappear, and thereby would represent a major break-through for secure, private, but task-universal humanoid assistive robots in the real world. Such a highly-integrated neuro-symbolic system would be as task-universal, flexible and fast as today’s neural planners, but still able to guarantee that learned constraints would be met as today’s symbolic planners. A truly neuro-symbolic system would easily be able to scale in the dimension of task and constraint complexity, because those could be learned and would not need to be manually designed. However,

while we seem to have major parts for functioning level five neuro-symbolic systems at hand, level six still presents a grand challenge that is possibly much farther off in the future, if achievable at all.

However, a deep level six neuro-symbolic integration might not be necessary for acceptable trade-offs regarding constrained task planning. Pure symbolic problem solving always comes with the advantage of guaranteeing certain behavior, which neural approaches cannot provide. For example, one could think of high-level symbolic planners that provide a general idea, and only fine-grained planning is done using neural systems. The same idea can be used for constraints that are imposed or observed by such planning systems. The integration of symbolic reasoning might also be done by imposing symbolic *red lines* to the neural planning system. The neural planner would then be constrained by a symbolic system that encodes rules of high importance. Aspects with a lower security level can however be constrained by a neural planning system, because errors do not have such an impact there. We conclude that the goal is not necessarily maximum neuro-symbolic integration, but rather to find the adequate integration level for the task at hand.

## 5.2 Next Steps

In this paper, we argued that neuro-symbolic hybrid approaches are necessary to bring constrained task-universal planning to real-world problem complexity. We listed many variations of possible hybrid approaches, which still need to prove their feasibility in real robot experiments. Feasibility especially means how well the different approaches cope with practical task and policy complexity in e.g. our assistive humanoid robot in care homes scenario. A practical feasibility analysis could also determine necessary and sufficient levels of neuro-symbolic integration for the scenario, finding out what type of neuro-symbolic integration would fit best. Deploying neuro-symbolic systems on real robots is an extensive challenge, but simulating such situations with the necessary degree of realism using information gathered from care homes seems feasible.

In the space of possible neuro-symbolic approaches, we see LLMs as promising building block due to their inherent knowledge on social norms and ability to work with multimodal, layered instructions – we imagine a LLM that is primed to behave like a humanoid assistive robot in a care home using documents and norms for human trainees, that use input from the robot’s camera, other sensors and memory to perform a task given as voice command in natural language. We see advances in this direction in the form of multimodal LLM-based planners that take vision and user command as input and are able to respect basic physical constraints (c.f., e.g. [4, 5]). These approaches directly embed images and other sensor data from robots into the language-embedding space and use these during training and fine-tuning stages. It can be observed that the models gain an implicit understanding of their surroundings. Similar to learned physics reasoning based on scene images that looks like an actual physics simulation based on a scene graph [3, 27], these developments need to be continued to provide learned access control reasoning based on natural language policies that looks like an actual reference monitor based on symbolic policies.

However, further research into guaranteeing quality properties of generated plans like validity, correctness, consistency, and stability when using LLMs is necessary – we see a promising direction in



building upon advances in controlling LLM output as a way to make LLM-based planners adhere to constraints (c.f., e.g. [2, 14, 18, 25]). In addition, performance of LLM-based planning in comparison to standard symbolic and neural planning systems in terms of speed and correctness is yet to be evaluated. Specific to LLMs used as knowledge base for planners, the general performance of inferring social norms also has to be evaluated.

Orthogonal to testing the feasibility of current ideas for hybrid approaches, neuro-symbolic AI advances as a field of itself. Any progress in this field, especially in the integration of a neural engine and a symbolic constraint system, would be directly applicable in constrained task planning for humanoid assistive robots.

## 6 CONCLUSION

We discussed combinations of symbolic and neural planning approaches under access control constraints and beyond. We highlighted their trade-offs and showed their benefits and shortcomings. While neural planners scale better to complex deployments in terms of required manual specification, they introduce an error margin and hinder administrability. The usage of neural systems such as LLMs as knowledge base for access decisions presents a direct way of incorporating social norms into such access decisions. Designing symbolic planning systems, however, requires more manual tasks than neural planning systems. While neural systems seem to have an advantage when scaling problem complexity to task universality, to actually teach them security, they do not only challenge us to quantify what we mean with security, but also to make that notion differentiable. In order to create deployable systems one should, therefore, strive for hybrid designs to combine best of both worlds: neural planning with neural constraints leads to probable observance of policies, critical policies will be safeguarded by symbolic constraints. Neuro-symbolic AI promises to advance the field of constrained task planning of assistive humanoid robots, and of access and activity control as well.

## ACKNOWLEDGMENTS

We like to thank Jan Niehues for his valuable comments, in particular for his hints on controlled text generation and controller architectures. This work has been supported by the project “Stay young with robots” (JuBot). The JuBot project was made possible by funding from the Carl Zeiss Foundation.

## REFERENCES

- [1] Saskia Bayreuther, Florian Jacob, Markus Grotz, Rainer Kartmann, Fabian Peller-Konrad, Fabian Paus, Hannes Hartenstein, and Tamim Asfour. 2022. BlueSky: Combining Task Planning and Activity-Centric Access Control for Assistive Humanoid Robots. In *Proceedings of the 27th ACM Symposium on Access Control Models and Technologies (SACMAT '22)*. ACM, New York, 185–194. <https://doi.org/10.1145/3532105.3535018>
- [2] Sumanth Dathathri, Andrea Madotto, Janice Lan, Jane Hung, Eric Frank, Piero Molino, Jason Yosinski, and Rosanne Liu. 2020. Plug and Play Language Models: A Simple Approach to Controlled Text Generation. arXiv:1912.02164 [cs.CL]
- [3] Danny Driess, Jung-Su Ha, Russ Tedrake, and Marc Toussaint. 2021. Learning Geometric Reasoning and Control for Long-Horizon Tasks from Visual Input. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*. 14298–14305. <https://doi.org/10.1109/ICRA48506.2021.9560934>
- [4] Danny Driess, Fei Xia, Mehdi S. M. Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, Wenlong Huang, Yevgen Chebotar, Pierre Sermanet, Daniel Duckworth, Sergey Levine, Vincent Vanhoucke, Karol Hausman, Marc Toussaint, Klaus Greff, Andy Zeng, Igor Mordatch, and Pete Florence. 2023. PaLM-E: An Embodied Multimodal Language Model. arXiv:2303.03378 [cs.LG]
- [5] Anthony Brohan et al. 2023. RT-2: Vision-Language-Action Models Transfer Web Knowledge to Robotic Control. arXiv:2307.15818 [cs.RO]
- [6] Artur d’Avila Garcez and Luis C. Lamb. 2023. Neurosymbolic AI: The 3rd wave. *Artificial Intelligence Review* (2023). <https://doi.org/10.1007/s10462-023-10448-w>
- [7] Oscar C González, Yury Sokolov, Giri P Krishnan, Jean Erik Delanois, and Maxim Bazhenov. 2020. Can Sleep Protect Memories from Catastrophic Forgetting? *eLife* 9 (aug 2020), e51005. <https://doi.org/10.7554/eLife.51005>
- [8] Arne Grävemeyer. 2023. Chatbots reinlegen – Wie man KI-Sprachgeneratoren entlarvt. *c’t* 16 (2023), 116–119. <https://www.heise.de/select/ct/2023/16/2313010292321745772> (“Fooling Chatbots - How to Expose AI Speech Generators.”)
- [9] Maanak Gupta and Ravi Sandhu. 2021. Towards Activity-Centric Access Control for Smart Collaborative Ecosystems. In *Proceedings of the 26th ACM Symposium on Access Control Models and Technologies (SACMAT '21)*. ACM, New York, 155–164. <https://doi.org/10.1145/3450569.3463559>
- [10] Niklas Hemken, Florian Jacob, Fabian Peller-Konrad, Rainer Kartmann, Tamim Asfour, and Hannes Hartenstein. 2023. Poster: How to Raise a Robot - Beyond Access Control Constraints in Assistive Humanoid Robots. In *Proceedings of the 28th ACM Symposium on Access Control Models and Technologies (Trento, Italy) (SACMAT '23)*. Association for Computing Machinery, New York, NY, USA, 55–57. <https://doi.org/10.1145/3589608.3595078>
- [11] Adele Howe, Craig Knoblock, Drew McDermott, Ashwin Ram, Manuela Veloso, Daniel Weld, and David Wilkins. 1998. PDDL – *The Planning Domain Definition Language*. Technical Report. Yale Center for Computational Vision and Control.
- [12] Daniel Kahneman. 2011. *Thinking, Fast and Slow*. Farrar, Straus and Giroux, NY.
- [13] Henry A. Kautz. 2022. The third AI summer: AAAI Robert S. Engelmore Memorial Lecture. *AI Magazine* 43, 1 (2022), 105–125. <https://doi.org/10.1002/aaai.12036>
- [14] Nitish Shirish Keskar, Bryan McCann, Lav R. Varshney, Caiming Xiong, and Richard Socher. 2019. CTRL: A Conditional Transformer Language Model for Controllable Generation. arXiv:1909.05858 [cs.CL]
- [15] Xiaomeng Lei and Mahesh Tripunitara. 2023. The Hardness of Learning Access Control Policies. In *Proceedings of the 28th ACM Symposium on Access Control Models and Technologies (Trento, Italy) (SACMAT '23)*. Association for Computing Machinery, New York, NY, USA, 133–144. <https://doi.org/10.1145/3589608.3593840>
- [16] Henry George Lidell and Robert Scott. [n. d.]. *A Greek–English Lexicon*. Clarendon Press, Oxford. <https://www.perseus.tufts.edu/hopper/text?doc=Perseus:text:1999.04.0057:entry=su/mbolon>
- [17] Michael McCloskey and Neal J. Cohen. 1989. Catastrophic Interference in Connectionist Networks: The Sequential Learning Problem. *Psychology of Learning and Motivation*, Vol. 24. Academic Press, 109–165. [https://doi.org/10.1016/S0079-7421\(08\)60536-8](https://doi.org/10.1016/S0079-7421(08)60536-8)
- [18] Jasmin Mousavi. 2024. Towards Consistent Language Models Using Declarative Constraints. In *AAAI workshop on Neuro-Symbolic Learning and Reasoning in the Era of Large Language Models*. <https://openreview.net/forum?id=UQTGlr30Y>
- [19] Humza Naveed, Asad Ullah Khan, Shi Qiu, Muhammad Saqib, Saeed Anwar, Muhammad Usman, Naveed Akhtar, Nick Barnes, and Ajmal Mian. 2024. A Comprehensive Overview of Large Language Models. arXiv:2307.06435 [cs.CL]
- [20] Mohammad Nur Nobi, Ram Krishnan, Yufei Huang, and Ravi Sandhu. 2022. Administration of Machine Learning Based Access Control. In *ESORICS 2022*. Springer, Cham, 189–210. [https://doi.org/10.1007/978-3-031-17146-8\\_10](https://doi.org/10.1007/978-3-031-17146-8_10)
- [21] Mohammad Nur Nobi, Ram Krishnan, Yufei Huang, Mehrnoosh Shakarami, and Ravi Sandhu. 2022. Toward Deep Learning Based Access Control. In *Proceedings of the Twelfth ACM Conference on Data and Application Security and Privacy*. ACM, New York, 143–154. <https://doi.org/10.1145/3508398.3511497>
- [22] OpenAI. 2023. GPT-4 Technical Report. arXiv:2303.08774 [cs.CL]
- [23] Roger Ratcliff. 1990. Connectionist Models of Recognition Memory: Constraints Imposed by Learning and Forgetting Functions. *Psychological Review* 97, 2 (1990), 285. <https://doi.org/10.1037/0033-295X.97.2.285>
- [24] Amit Sheth, Kaushik Roy, and Manas Gaur. 2023. Neurosymbolic AI – Why, What, and How. arXiv:2305.00813 [cs.AI]
- [25] S. M Towhidul Islam Tonmoy, S M Mehedi Zaman, Vinija Jain, Anku Rani, Vipula Rawte, Aman Chadha, and Amitava Das. 2024. A Comprehensive Survey of Hallucination Mitigation Techniques in Large Language Models. arXiv:2401.01313 [cs.CL]
- [26] Wikipedia contributors. 2023. *River crossing puzzle*. [https://en.wikipedia.org/w/index.php?title=River\\_crossing\\_puzzle&oldid=1150277959](https://en.wikipedia.org/w/index.php?title=River_crossing_puzzle&oldid=1150277959)
- [27] Haotian Xue, Antonio Torralba, Josh Tenenbaum, Dan Yamins, Yunzhu Li, and Hsiao-Yu Tung. 2023. 3D-IntPhys: Towards More Generalized 3D-grounded Visual Intuitive Physics under Challenging Scenes. In *Advances in Neural Information Processing Systems*, A. Oh, T. Neumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (Eds.), Vol. 36. Curran Associates, Inc., 7116–7136. [https://proceedings.neurips.cc/paper\\_files/paper/2023/file/164687cb815daae754d33364716e65e6-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/164687cb815daae754d33364716e65e6-Paper-Conference.pdf)
- [28] Ziyi Yang, Shreyas S. Raman, Ankit Shah, and Stefanie Tellex. 2023. Plug in the Safety Chip: Enforcing Constraints for LLM-driven Robot Agents. arXiv:2309.09919 [cs.RO]