Externally-Funded, Collaborative Projects

Current Projects

PURSUIT - Purposeful Signal-symbol Relations for Manipulation Planning (Austrian Science Fund (FWF), 2023-2026): Artificial intelligence (AI) task planning approaches permits projecting robotic applications outside industrial settings by automatically generating the required instructions for task execution. However, the abstract representation used by AI planning methods makes it complicated to encode physical constraints that are critical to successfully execute a task: What specific movements are necessary to remove a cup from a shelf without collisions? At which precise point should a bottle be grasped for a stable pour afterwards? These physical constraints are normally evaluated outside AI planning using computationally expensive trial-and-error strategies. PURSUIT focuses on a new task and motion planning (TAMP) approach where the evaluation of physical constraints starts at perception stage and propagates through planning and execution using a single heuristic search. The approach is based on a common signal-symbol representation that encodes physical constraints in terms of the “purpose” of object relations in the context of a task: Is the hand-bottle relation adequate for picking up the bottle for a stable pouring afterwards? This is a fundamental difference with respect to the purpose-agnostic descriptions of object relations used by current TAMP approaches. Our TAMP approach aims to quickly render task plans that are physically feasible, avoiding the intensive computations of trial-and-error approaches.

ELSA - Effective Learning of Social Affordances for Human-Robot Interaction (ANR/FWF AAPG, 2022-2026): Affordances are action opportunities directly perceived by an agent to interact with its environment. The concept is gaining interest in robotics, where it offers a rich description of the objects and the environment, focusing on the potential interactions rather than the sole physical properties. In this project, we extend this notion to social affordances. The goal is for robots to autonomously learn not only the physical effects of interactive actions with humans, but also the humans’ reactions they produce (emotion, speech, movement). For instance, pointing and gazing in the same direction make humans orient towards the pointed direction, while pointing and looking at the finger make humans look at the finger. Besides, scratching the robot’s chin makes some but not all humans smile. The project will investigate how learning human- general and human-specific social affordances can enrich a robot’s action repertoire for human-aware task planning and efficient human-robot interaction.

SEAROCO - Seamless Levels of Abstraction for Robot Cognition (Austrian Science Fund (FWF) - Lise Meitner Project, 2019-2023): The project seeks to develop a robotic cognitive architecture that overcomes the difficulties found when integrating different levels of abstractions (e.g. AI and robotic techniques) for task plan and execution in unstructured scenarios. The backbone of the project is a unified approach that permits searching for feasible solutions for new tasks execution at all the levels of abstractions simultaneously, where symbolic descriptions are no longer disentangled from the physical aspects they represent.

Completed Projects (Selection)

OLIVER - Open-Ended Learning for Interactive Robots (EUREGIO IPN, 2019-2022): We would like to be able to teach robots to perform a great variety of tasks, including collaborative tasks, and tasks not specifically foreseen by its designers. Thus, the space of potentially-important aspects of perception and action is by necessity extremely large, since every aspect may become important at some point in time. Conventional machine learning methods cannot be directly applied in such unconstrained circumstances, as the training demands increase with the sizes of the input and output spaces. Thus, a central problem for the robot is to understand which aspects of a demonstrated action are crucial. Such understanding allows a robot to perform robustly even if the scenario and context change, to adapt its strategy, and to judge its success. Moreover, it allows the robot to infer the human intent and task progress with respect to the goal, enabling it to share the task with humans, offer help or ask for help, resulting in natural human-robot cooperative behavior.

IMAGINE - Robots Understanding Their Actions by Imagining Their Effects (EU H2020, 2017-2021): seeks to enable robots to understand the structure of their environment and how it is affected by its actions. “Understanding” here means the ability of the robot (a) to determine the applicability of an action along with parameters to achieve the desired effect, and (b) to discern to what extent an action succeeded, and to infer possible causes of failure and generate recovery actions.

FlexRoP - Flexible, assistive robot for customized production (FFG (Austria) ICT of the Future, 2016-2019): Production of mass customized products is not easy to automate since objects and object positions remain more uncertain compared to mass production scenarios. Uncertainty handling motivates the application of advanced sensor-based control strategies which increases system complexity of robot applications dramatically. A possible solution to this conflict is the concept of task level or skill based programming that will render modern robot systems. Such systems can be applied without safety fence, are easier to program, more applicable and transformable into capable robot assistants. The project will implement a skill based programming framework and will apply it on selected industrial demo scenarios and evaluate research results. The main focus of the project is the application of methods to acquire process information by worker monitoring and thus make the robot assistants self-learning.

SQUIRREL (EU FP7-ICT-STREP, 2014-2018): Clutter in an open world is a challenge for many aspects of robotic systems, especially for autonomous robots deployed in unstructured domestic settings, affecting navigation, manipulation, vision, human robot interaction and planning. SQUIRREL addresses these issues by actively controlling clutter and incrementally learning to extend the robot's capabilities while doing so. We term this the B3 (bit by bit) approach, as the robot tackles clutter one bit at a time and also extends its knowledge continuously as new bits of information become available. SQUIRREL is inspired by a user driven scenario, that exhibits all the rich complexity required to convincingly drive research, but allows tractable solutions with high potential for exploitation. We propose a toy cleaning scenario, where a robot learns to collect toys scattered in loose clumps or tangled heaps on the floor in a child's room, and to stow them in designated target locations.

3rdHand (EU FP7-ICT-STREP, 2013-2017) develops a semi-autonomous robot assistant that acts as a third hand of a human worker. It will be straightforward to instruct even by an untrained layman worker, allow for efficient knowledge transfer between tasks, and enable effective collaboration between a human worker with a robot third hand. The main contributions of this project will be the scientific principles of semi-autonomous human-robot collaboration, a new semi-autonomous robotic system that is able to (i) learn cooperative tasks from demonstration, (ii) learn from instruction, and (iii) transfer knowledge between tasks and environments.

PaCMan - Probabilistic and Compositional Representations for Object Manipulation (EU FP7-ICT-STREP, 2013-2016) advances methods for object perception, representation and manipulation so that a robot is able to robustly manipulate objects even when those objects are unfamiliar, and even though the robot has unreliable perception and action. The proposal is founded on two assumptions. The first of these is that the representation of the object's shape in particular and of other properties in general will benefit from being compositional (or very loosely hierarchical and part based). The second is that manipulation planning and execution benefits from explicitly reasoning about uncertainty in object pose, shape etcetera; how it changes under the robot's actions, and the robot should plan actions that not only achieve the task, but gather information to make task achievement more reliable.

Xperience (EU FP7-ICT-IP, 2011-2015) pursues two principal objectives. The first goal is to show that the state of the art enactive embodied cognition systems can be significantly enhanced by using structural bootstrapping - a concept taken from language learning. The second goal is to implement a complete robot system for automating introspective, predictive, and interactive understanding of actions and dynamic situations.

IntellAct (EU FP7-ICT-STREP, 2011-2014) addresses the problem of understanding and exploiting the meaning (semantics) of manipulations in terms of objects, actions and their consequences for reproducing human actions with machines. This is in particular required for the interaction between humans and robots in which the robot has to understand the human action and then to transfer it to its own embodiment.

LearnBiP (EU FP7-ICT ECHORD Experiment, 2011-2012) has two main aims. First it utilizes the huge amount of data generated in industrial bin-picking for the introduction of grasp learning. Second it evaluates the potential of the SCHUNK dexterous hand SDH-2 for its application in industrial bin-picking.

SignSpeak (EU FP7-ICT-STREP, 2009-2012) focused on scientific understanding and vision-based technological development for continuous sign language recognition and translation. The aim was to increase the linguistic understanding of sign languages and to create methods for transcribing sign language into text.

PACO-PLUS (“perception, action and cognition through learning of object-action complexes”, EU FP6-IST-IP, 2006-2010) brings together an interdisciplinary research team to design and build cognitive robots capable of developing perceptual, behavioural and cognitive categories that can be used, communicated and shared with other humans and artificial agents. To demonstrate our approach we are building robot systems that will display increasingly advanced cognitive capabilities over the course of the programme. They will learn to operate in the real world and to interact and communicate with humans. To do this they must model and reflectively reason about their perceptions and actions in order to learn, act and react appropriately. We hypothesize that such understanding can only be attained by embodied agents and requires the simultaneous consideration of perception and action.

TRICTRAC (2003-2006), directed by J. Piater, aimed at the development of algorithms for real-time object tracking in one or more live video streams. It was a joint project between the Université de Liège and the Université Catholique de Louvain funded by the Walloon Region. Some results are summarized in a video.