

Integrating Perception, Learning and Verification in Interactive Autonomous Systems

Danica Kragic

Michael Felsberg

Alexandre Proutiere

Kalle Åström

Vision for the project

Perception is the basis of **any** interactive autonomous systems. Similarly to humans, artificial systems base their actions on various sensory modalities such as vision, force and torque sensing, range sensing, tactile sensing. The choice of the employed sensory modality is usually based on the set of tasks a system is supposed to perform. For example, visual scene analysis is essential to enable navigation in dynamic environments shared with humans and containing visual symbols and sign. Human-robot physical collaboration requires in addition tactile and force-torque sensing. There are several scientific challenges in relation to developing autonomous, interactive systems with a goal to achieve robustness and ability to adapt to new knowledge and generalize it to new situations. Data representation and fusion of sensory information have to be combined with online learning of visual models from weakly annotated data with minimal supervision. Leveraging and learning from increasingly large amount of data critically requires the development of new theoretical tools for data analytics and learning, considering inputs from several sensory modalities. These tools need to be generic and pervasive, including mathematical models, learning and inference algorithms, as well as appropriate optimization techniques. In this project, we will define and develop novel perception capabilities, acquired through learning, for the use in interactive and autonomous systems. The application areas will be service robotics settings, industrial assembly lines (through collaboration with ABB), public safety (through collaboration with Ericsson and SAAB), and Collaborative Automated Transport Systems (through collaboration with Autoliv, Volvo, SAAB and Scania). Besides progress in the areas of machine learning using large amounts of perceptual data from unstructured environments, this will also comprise novel verification methods to guarantee correctness, reliability, and robustness of the resulting systems, especially also in collaboration with other WASP subprojects.

Research Challenges

The envisioned systems will be able to interact with and adapt to their environment, as well as collect and learn from data to take informed decisions. The scientific challenges are:

1. The design of systems learning from very large amounts of data require classical machine learning techniques such as classification, clustering, detection and sketching algorithms. The focus here will be on deep learning and Bayesian non-parametric inference along with its associated sampling techniques. Both tools are expected to bring breakthrough improvements into critical tasks in robotics and human-machine interaction, including visual classification, speech recognition, and natural language processing.
2. When, on the contrary, the training data is relatively sparse, it may be crucial to transfer knowledge from other domains where large training data has been available. We plan to develop novel transfer learning techniques using e.g. generative probabilistic and topological models. We expect these techniques to constitute fundamental building blocks of autonomous systems whose sensing and learning capabilities must be complemented by processing information from other sources, e.g. self-driving vehicles.
3. Autonomous systems critically need to interact with and learn from evolving data in an online manner. For example, knowledge originating from visual perception and influencing representations on the higher level (and vice versa) will interact with probabilistic planning. The recognition of human actions is essential to modeling, analysis and syntheses of human-in-the-loop collaborative systems, which will enable bootstrapping of autonomous agents in completely unknown environments with a minimum of cognitive load for the human operator. The research focus will be on online weakly supervised, reinforcement, and Hebbian learning, Gaussian process methods and bandit and expert optimization.

4. The final class of systems consists of those learning autonomously, i.e., deciding on what they need to learn, and gathering the appropriate training data towards this aim. This calls for the development of decision methods and tools to seek and acquire training data to improve system performance in a fully automated way. For example, online perception and real-time sensing, requires data to be acquired in an explorative manner, and future work will focus on exploiting feedback mechanisms and self-assessment during learning.

Industrial Challenges

Industrial standards require certification and standardization of methods, which is difficult to achieve for learning based systems in unstructured environments. New approaches to data augmentation and simulation for the purpose of evaluation on large scales are required, along with safety by design principles. Another challenge is fast programming of assembly lines directly from human demonstration as well as both physical and non-physical human-robot collaboration. One specific example is the generation and sharing of map and 3D model updates between agents, which is a very useful functionality in a highly dynamic world and in particular in catastrophe cases. Potential customers are companies building planning or navigation systems, assembly lines as well as blue-light forces. The generic extension of object detection and learning capabilities is probably of even higher industrial impact. Many companies struggle currently with present deep learning techniques, in particular if their specific use-cases are not well covered by training on ImageNet. For instance traffic safety systems need much more specific datasets for training than what ImageNet provides and they need to be updated continuously whenever new object types occur (e.g. hover boards). The aspect of object interaction and manipulation is of special industrial interest and we already have two PhD students from ABB starting on this project. Finally, software verification is of interest for most of the companies involved in WASP and there are already several collaborations ongoing.

Sub-projects

Semantic structure from motion for autonomous systems

David Gillsjö, academic PhD, MIG/LU

Object recognition and 3D scene reconstruction have so far largely been studied independently. An example is multiple view geometry where a major success in recent years has been the ability to automatically reconstruct large-scale 3D models from collections of 2D images. The approach is based on purely geometric concepts and it is mostly passive, utilizing no semantic scene understanding. Limitations are apparent: certain scene elements cannot be reconstructed because geometry is under-constrained, mid-level gestalts and category specific priors cannot be easily leveraged, and the model ultimately provides a point cloud and a texture map, not a semantic representation that enables effective navigation or interaction. The goal here is to develop an integrated framework capable of recognizing, navigating and mapping based on geometric computer vision and deep learning techniques.

Fusion of visual tracking approaches and machine learning of object detection and recognition

Gustav Häger, academic PhD, CVL/LiU

State-of-the-art visual object detection and recognition methods are based on deep learning approaches making use of data from ImageNet. Once learned, these modules remain static and applications in new problem domains require additional off-line learning with specific datasets typically much smaller than ImageNet. On the contrary, state-of-the-art visual object tracking is initialized with a single patch and remains adaptive to new data throughout its operation. The goal of this sub-project is to use this adaptive visual modelling from visual object tracking and to extend it to multiple-aspect generative modelling. This generative model can then be used to train the detectors and classifiers with problem specific data, leading to a fusion of visual perception modalities.

Online learning for visual navigation of UAVs

Bertil Grelsson, industrial PhD, SAAB Dynamics AB + CVL/LiU

An unmanned aerial vehicle (UAV) is conducting a mission in an area where a detailed 3D map is available. The UAV carries an onboard fisheye camera for surveillance and reconnaissance purposes. Current conditions demand GPS-free ego-localization and navigation in the area. A high quality 3D-map

with visual texture will be used to train a convolutional neural network (CNN) to enable online coarse ego-localization from aerial fisheye images captured by the UAV. We will also focus on questions related to combining machine learning methods with geometric approaches, e.g. for navigating in partially changed or destroyed environments while updating their local maps on the fly. The updated map is particularly useful to be shared with other agents in the same system.

Deep learning for visual tracking

Martin Danelljan, affiliated PhD, CVL/LiU

For advanced visual perception, objects and features in the environment need to be detected, classified and tracked. The tracking of objects provides situation awareness, while feature tracking can be used for mapping and localization. Unlike many related computer vision problems, deep learning has only achieved partial success for visual tracking due to two fundamental challenges: (1) the online nature of the learning problem and (2) the lack of training data. We will address these challenges by investigating novel deep learning architectures that combine offline learning of generic visual features suitable for the tracking task with flexible online learning approaches.

Learning for task based grasping

Mia Kovic, academic PhD, KTH

The specific goal of the work in this thesis will be on the machine learning approaches for multisensory, task based grasping. We will address problems of data representation and data fusion given multiple of sensory modalities such as vision, tactile and force-torque. We will develop online learning via crowdsourcing and incremental learning using weakly annotated data. Leveraging and learning from increasingly large amount of data requires development of new theoretical tools and we will look into state free representations such as Probabilistic State Representation. Task relation will be ensured through the use of Bayesian models where we will address structure learning in networks with discrete and continuous nodes.

Autonomous skill acquisition for robot assembly tasks

Shahbaz Khader, industrial PhD, KTH+ABB

The goal is to develop learning mechanisms for bi-manual assembly skills. This includes learning low-level sensorimotor signals and also learning fine manipulation planning functions. The objective is to make the entire process more "autonomous" with robot learning from its own experience in addition to human specified goals in small parts assembly. We will develop scenarios to define the requirements in terms of perception (What are the appropriate/available sensory modalities? How are these integrated?), learning (What can we measure and how complex are data? What can be learned in an unsupervised manner?) and control (How do we adopt model predictive control for both low and high level tasks?)

Planning and learning for robot assembly tasks

Johan Wessen, industrial PhD, KTH+ABB

The main goal of the project is to develop methodologies for autonomous, interactive systems to achieve robustness and ability to adapt to new knowledge and to generalize to new situations, especially with the focus on small parts assembly lines. The project will enable a robot to learn how to perform assembly tasks from own interaction and when interacting with a human. The special goal of this work is to develop learning mechanisms that can learn from sensory data to classify a work scene in a manner that a collision avoiding path planner can adapt to the classification. This will likely have at least three aspects: learning to classify the objects in a given scene, develop strategies for a path planner to react to the different classifications, and third on how to use the increased perception to make the robot more autonomous.

Reinforcement learning in Markov Decision Processes and Model Training

Daniel Wrang, academic PhD, KTH

Reinforcement learning constitutes a versatile tool for decision-making and optimization in uncertain environments. We will here develop novel reinforcement learning algorithms to tackle problems in adaptive control problems in large-scale dynamical systems modeled as Markov Decision Processes (MDPs), and in online optimization problems related to model training in machine learning. The first

objective is to devise algorithms learning as fast as possible the optimal policy in MDPs whose parameters are initially unknown. The second objective of this work is to propose online and adaptive learning methods to speed up the training of a model using data. Usually, these methods leverage stochastic optimization techniques, i.e., using in each iteration, a sample chosen randomly from the data to improve our knowledge of the model parameters. We plan here to apply reinforcement learning tools to devise faster training algorithms.