Sustainable computer vision for autonomous machines

The main aim of this grant is to develop sustainable computer vision methods for autonomous machines, which assume device limitations and a variety of sensors. We examine issues such as adapting machine learning models to new types of data obtained, for example, with the help of event cameras, and we are introducing new architectures suitable for limited computing resources, i.e. GPU, CPU, and RAM. We focus on the theory of machine learning and not on the hardware aspect, taking into account environmental issues, which are an extremely important aspect of the development of machine learning. Therefore, priority is given to methods that not only improve efficiency and effectiveness of models but also reduce the carbon footprint. This approach translates into a potentially larger number of applications of various machines, while reducing the adverse impact of the technology on the environment.

Animal protection, fire detection or support for safety services. Our solutions could potentially be used in drones as a tool supporting the protection of national parks, including animals against poaching. They allow for fast and efficient monitoring of large land areas in remote locations thanks to panoramic vision and specialized data from laser or thermal scanners, for example. As a result, it is possible to monitor movements of animals or to early detect forest fires. We already have drones that can perform such missions, but their operation requires specially trained personnel. There are also limitations related to equipment. In order to operate autonomously, with minimal human support, a drone must be able to detect and identify animals, or notice forest fires, without outside help, and then take appropriate action. This application perfectly fits into the goals of our group, among which we should mention the introduction of innovative methods of active visual exploration.

Robots that can see… like people. One of the issues we are interested in is active visual exploration. Machines do not observe their surroundings like humans. This is due to the fact that the limited capabilities of sensors, or the size of the battery, do not allow the machine to analyze the entire environment at the same time. Therefore, the challenge for an autonomous vehicle is to analyze the environment, as well as to decide what to focus attention on and what next steps to take. In other words, active exploration addresses the problem of limited sensor capabilities in real-world scenarios, where subsequent observations are actively selected based on the environment. For example, robot sensors have a limited field of view, the environment is constantly changing, and computational costs are high, which complicates obtaining complete information about the environment. Therefore, in order to infer about the entire environment, the agent must most efficiently sample new observations. We solve this problem by introducing new techniques based on well-generalized transformers, for example.

Fast adaptation is the key to success. An important aspect of our research is the introduction of models capable of generalization, so that a model trained for one environment (such as the prairie) can be quickly adapted to a new one (such as a jungle). To this end, we are introducing novel approaches to self-supervised learning that trains the model without the use of labels, which in turn benefits applications in a variety of fields. As a result, we kind of recycle models, significantly reducing the amount of energy required for training.