Welcome to the 1st OpenDR Project Newsletter !

Newsletter 1.0: OpenDR challenges and current status

Almost everything we hear about artificial intelligence today is thanks to deep learning (DL). Deep learning has achieved tremendous performance jumps in the last decade in several Computer Vision (CV) and Machine Learning (ML) tasks, achieving in many cases super-human performance. However, DL cannot be currently fully exploited in robotics scenarios due to a number of barriers.

Learning Curve Barrier

DL has a **steeper learning curve** than traditional CV and ML methods

Computational Complexity Barrier

DL requires vast amounts of computational power and **energy**

Static Perception Barrier

DL is applied on **static environments** and **does not exploit spatial or temporal embodiment**

The need for an open deep learning toolkit that contains easy to train and deploy real-time, lightweight, Robot Operating System (ROS) compliant deep learning models for robotics is evident. This is where the **OpenDR** project enters.

What is OpenDR?

OpenDR “Open Deep Learning for Robotics Toolkit”, is a EU 2020 Project which was launched on January 2020 and aims to develop a modular, open and non-proprietary toolkit for core robotic functionalities by harnessing deep learning to provide advanced perception and cognition capabilities, meeting in this way the general requirements of robotics applications in the areas of healthcare, agri-food and agile production. The **OpenDR** project is coordinated by the Aristotle University of Thessaloniki, Greece and will be running throughout December 2022 with a total budget of 6.6 Million Euros.

OpenDR will enable real-time robotic visual perception on high-resolution data and enhance the robotic autonomy exploiting lightweight deep learning for deployment on robots and devices with limited computational resources. In addition, it aims to propose, design, train and deploy models that go beyond static computer vision and towards active robot perception, providing deep human-centric and environment active robot perception, as well as enhanced robot navigation, action and manipulation capabilities.

OpenDR’s expected impact is to improve the technical capabilities in robotics by providing easily deployable, efficient and novel Deep Learning tools, as well as to lower the technical barriers by providing a modular and open platform for developing Deep Learning for Robotics tools. Concerning industry, the project’s expected impact is to enable a greater range of applications in agri-food, healthcare robotics and agile production, as well as to strengthen the competitiveness of companies by lowering the cost to access robotics-oriented Deep Learning tools.

OpenDR Consortium

OpenDR consortium is a very good mix of 8 partners from 7 European Countries: 2 companies working in various fields of robotics, one company working in the field of robotics simulations, and 5 Universities that join the project with 4 robotics laboratories and 3 deep learning and computer/robot vision laboratories.

Aristotle University of Thessaloniki (AUTH) is the largest university in Greece, established in 1925. AUTH coordinates the project and leads the organization of dissemination activities. AUTH will focus its research on deep human centric active perception and cognition, where it will contribute on deep person/face/body part active detection/recognition and pose estimation, deep person/face/body part tracking, human activity recognition, as well as social signal analysis and recognition. AUTH will also lead the research in object detection/recognition and semantic scene segmentation and contribute to other areas such as evaluation and benchmarking activities of the project.

Tampere University (TAU) is Finland's second-largest university with 20.000 students and 330 professors. TAU participates with two labs/groups namely the Laboratory of Signal Processing at the Department of Computing Sciences and the Cognitive Robotics Group at the Department of Automation Technology and Mechanical Engineering. TAU will lead the research in deep human centric active perception and cognition, working mainly on deep speech and biosignals analysis and recognition, and will contribute to deep person/face/body part active detection/recognition and multi-modal human centric perception and cognition as well as in a number of other topics. TAU will also contribute on defining the agile production use case requirements and specifications and on the integration of OpenDR to this use case.

University of Freiburg (ALU-FR) is one of Germany's leading research institutions with an international reputation in many fields. ALU-FR will lead the research in deep environment active perception and cognition. ALU-FR will focus its research on Deep SLAM and 3D scene reconstruction, as well as on deep navigation. It will also contribute on developing methodologies for deep planning.

Aarhus University (AU), Denmark participates in OpenDR with two groups, namely the Data-Driven Analytics Group and the Artificial Intelligence in Robotics Group, both belonging to the Section of Electrical and Computer Engineering. AU will lead work on 2D/3D Object localization and tracking and will work on sensor information fusion, as well as object detection/recognition and semantic scene segmentation and understanding. AU will also contribute to a number of areas such as deep person/face/body part active detection/recognition, deep person/face/body part tracking, deep planning, etc.

Delft University of Technology (TUD) is the oldest and largest technical university in the Netherlands. TUD will lead/organize the research activities on deep action and control, deep planning, as well as deep navigation. Furthermore, TUD will also lead and undertake the research activities on human robot interaction. Finally, it will lead and organize the toolkit evaluation and benchmarking activities of the project.

Cyberbotics (CYB) is a Swiss spin-off company from EPFL, which has been developing the Webots robot simulator since 1998. CYB will lead efforts of defining the toolkit's requirements and specifications. CYB will also work on developing simulation environments and collecting data. Finally, it will also lead on toolkit integration by collecting and integrating all the OpenDR modules developed by the partners.

PAL Robotics (PAL) is a Spanish SME that provides robotic products and services. PAL will organize and coordinate the toolkit integration, as well as the use cases integration activities. PAL will also contribute on defining the healthcare robotics use case requirements and specifications and will work on the integration of OpenDR Toolkit to this use case, as well as on its evaluation.

Agro Intelligence APS (AGI), Denmark will organize and coordinate the toolkit evaluation, as well as the use cases specific toolkit evaluation activities. AGI will also contribute on defining the agri-food use case requirements and specifications and will work on the integration and evaluation of OpenDR Toolkit in this specific use case.

Work Performed so far

A lot has happened during the first year of the project. Read through the following sections to learn more!

Requirements and Specifications

At the beginning of the project, all partners actively collaborated for the definition of the requirements and specification of the OpenDR toolkit and the three application scenarios: agri-food, healthcare robotics and agile production. Indeed, a key part of the work has been the detailed description of these three robotics scenarios. Moreover, the consortium identified a set of objectives that the OpenDR DL tools should achieve with some specific hardware and device requirements. The OpenDR algorithms, to be developed and included in the first version of the toolkit have then been selected based on the use cases objectives. The OpenDR toolkit will provide multiple interfaces. A Python interface, will provide all the methods required for training and inference tasks. An additional C/C++ inference will be available for high performance applications.

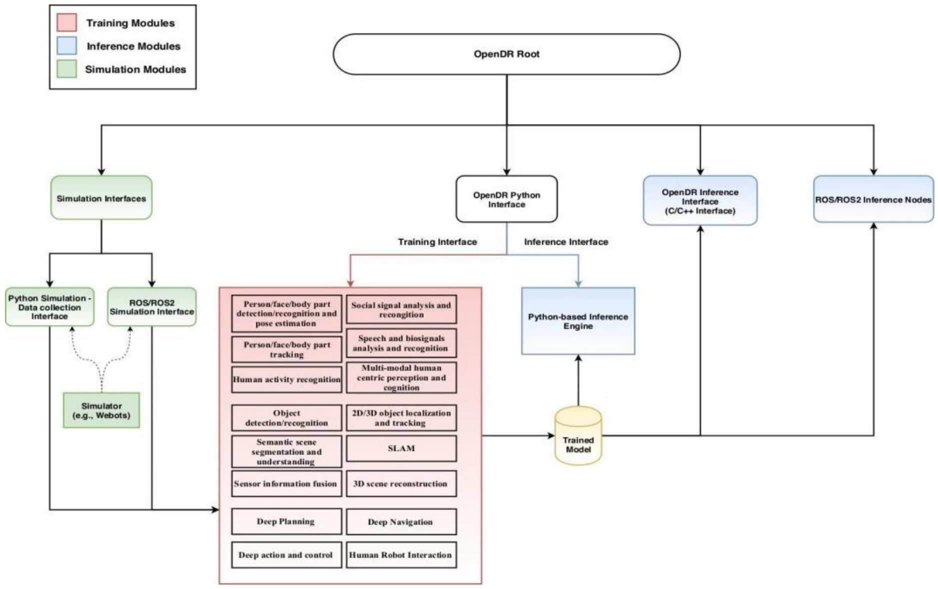

OpenDR Toolkit structure (better viewed in large magnification)

Finally, through a ROS/ROS2 interface the toolkit will be able to communicate with the robotics system or the simulation environment. Overall, the toolkit will be based on the most recent, established and widely used frameworks and standards in order to provide a state-of-the art efficient library.

Deep Human Centric Active Perception and Cognition

For the larger part of the year AUTH, AU and TAU have jointly worked on human centric tools and algorithms for the OpenDR project.

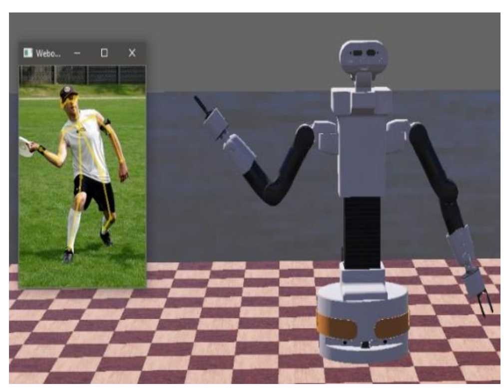

Demonstration of the TIAGo robot, implemented in the Webots simulator, mimicking a human's pose

Notable contributions were made in a number of research areas and the key project objectives were advanced, bringing the team closer to the realization of the powerful, flexible and efficient robotics toolkit. Deep learning models for robotics applications have to be lightweight due to computation power restrictions and achieve real-time performance. AUTH has worked on preparing such models for robotic vision tasks as well as models that adjust their runtime based on the available resources. AU, also made significant progress in skeleton-based human activity recognition, creating fast algorithms that can even outperform current state-of-the-art solutions. TAU proposed a framework for multilinear compressed learning, which is highly efficient with respect to memory and computation, achieving superior performance for face recognition tasks. TAU has also worked on efficient solutions for anomaly detection in heart signals and speech command recognition.

Deep Environment Active Perception and Cognition

ALU-FR made significant progress on deep environment perception and cognition through panoptic segmentation. Though relative methods in the domain are computationally intensive, ALU-FR proposed an architecture, called EfficientPS, that allows real-time inference on high-resolution input.

TUD has developed sensor fusion strategies for multi-modal object detection that efficiently exploit sensors redundancy in harsh lighting conditions. Moreover, they proposed a lightweight learning-free data augmentation method which creates random highlights and shadows to mimic such harsh conditions.

Panoptic segmentation results of the EfficientPS network, proposed by ALU-FR

Deep Robot Action and Decision Making

Partners AU, ALU-FR, TUD, TAU worked on the design of novel navigation, planning, and control algorithms. All the involved partners contributed to the state of the art and to the key project objectives.

ALU-FR developed a deep Reinforcement Learning (RL) method which enables a mobile robot to perform navigation for manipulation by generating kinetically feasible base motions given an end-effector motion following an unknown policy to fulfil a certain task. TUD investigated the design of an efficient model-based agent that can learn from images and be robust to possible distractions.

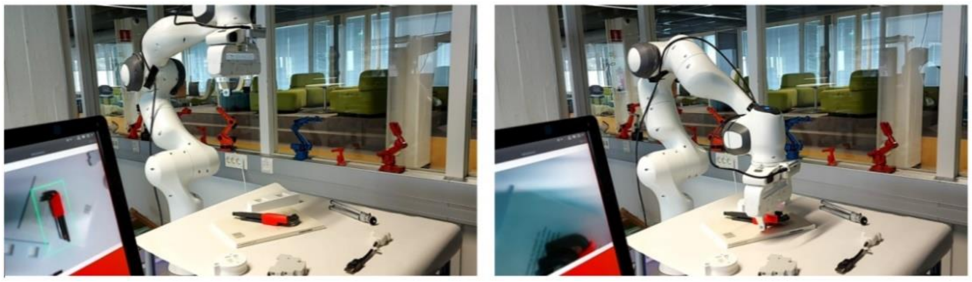

Single-Demonstration Grasping model at the hover pose (left) and the grasp pose (right)

AU has developed end-to-end planning methods with deep RL for autonomous drone racing and local replanning for the agricultural use-case. Finally, TAU has investigated robot grasping models for handling industrial objects in the Agile Production use case. A Single-Demonstration Grasp model that is light-weight and easy to train is under development and being tested with a collaborative robot.

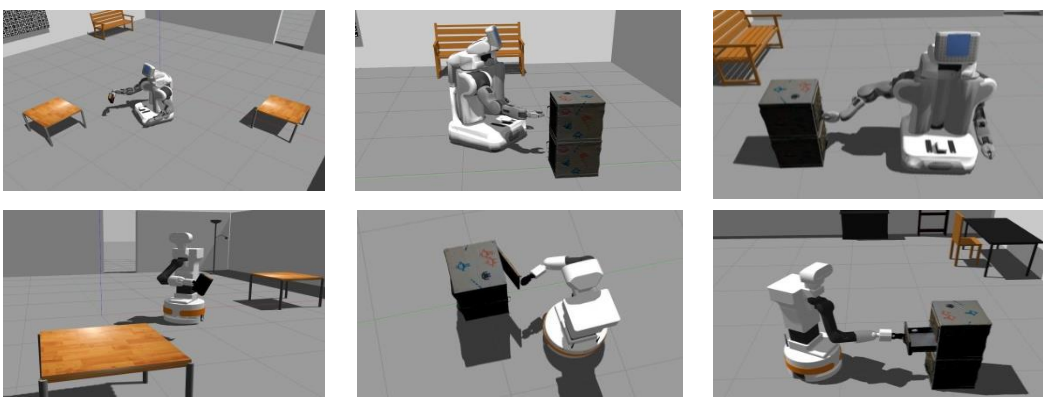

Example trajectories in the Gazebo physics simulator on different tasks for the PR2 (top) and TIAGo (bottom) mobile platforms

Simulation Environments and Data

Simulation has an important role in the development of the OpenDR toolkit: it helps in the generation of the data needed to train and test the developed DL tools and provides an easily accessible environment where to test the integration with the robotics system.

With these objectives in mind, AUTH worked on the generation of three human-related datasets: multi-view facial datasets, annotated image datasets depicting humans in various environments and a dataset of 3D human animatable models. These datasets will then be used for person detection, person recognition and pose estimation whereas 3D human models will also be used directly in the Webots simulator to create realistic robotics scenarios including human-robot interaction.

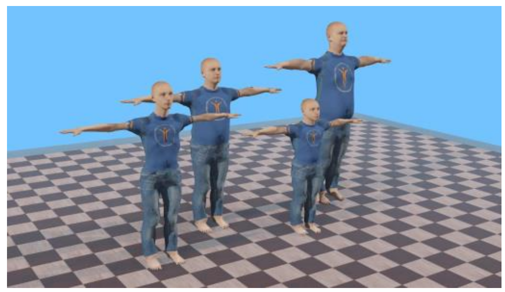

SMPL human body models in Webots generated with different shape parameters (right), reposed using existing animations (left)

Similarly, CYB worked on extending Webots to make it more robust and suitable in DL applications and on improving the compatibility with ROS/ROS2 by adding an export to URDF, switching to the ROS default coordinate system, and automatically generating a ROS2 compatible interface by parsing the Webots robot model. Finally, CYB started developing a platform to run simulations on the web and disseminate the results of the OpenDR project.

Dissemination

Dissemination and communication activities are a very important part of OpenDR. Although COVID-19 pandemic had a negative impact on the overall dissemination, all partners have made numerous and diverse efforts to attract interest in the project and its findings so far.

Right in the beginning, AUTH organized a workshop on Deep Learning for Robotics, in parallel to the project kick-off meeting. Members of the consortium were able to deliver detailed lectures about the progress of deep learning and robotics to a wide audience. To disseminate the project though the web, AUTH set up the official project website along with social media accounts (Facebook, Twitter, LinkedIn). Frequent updates about the project are provided through all these channels. In addition, a promotional video was created and uploaded on the project's YouTube channel, providing a brief but comprehensive overview of the project.

Despite being at an early stage of the project lifecycle, the consortium managed to generate a high volume of publications. A total of 16 papers were presented or accepted in high quality, well-established international scientific conferences (including CVPR2020, ICME2020, ICPR2020, ICRA2020 etc.) and 7 papers appeared in scientific journals (including the highly influential IEEE TNNLS and Elsevier's Pattern Recognition). OpenDR has also been actively present at prestigious scientific conferences and industrial events, where its members have served in organization committees, gave invited lectures, organized special sessions etc. A typical example was the strong presence of the consortium in IEEE ICIP2020, where Prof. Moncef Gabbouj (TAU) served as the General Co-chair and Profs. Alexandros Iosifidis (AU) and Anastasios Tefas (AUTH) co-organized a Special Session on "Deep Learning for Robotic Perception and Cognition". Overall, the consortium was present at numerous events (most of them virtual due to the pandemic), presenting the aims and results of OpenDR, while also exchanging ideas in domains relative to the project.

opendr.eu | **@OpenDR_EU** | **OpenDR** | **OpenDR H2020 Research Project**

This project has received funding from the *European Union's Horizon 2020 research and innovation programme* under Grant agreement No *871449*.