OpenDR
–
Open Deep Learning Toolkit for Robotics

Project start date: 01.01.2020
Duration: 48 months
Lead contractor: Aristotle University of Thessaloniki

Deliverable D9.8: Final report on dissemination, exploitation plans and project newsletter

Date of delivery: 31 Dec 2023

Contributing Partners: Aristotle University of Thessaloniki, Cyberbotics S.à r.l (CYB). Tampere University (TAU), Aarhus University (AU), Delft University of Technology (TUD), Albert-Ludwigs-University of Freiburg (ALU-FR), PAL Robotics SL (PAL), Agro Intelligence APT (AGI)

Version: 3.0

Title:	**D9.8: Final report on dissemination, exploitation plans and project newsletter**	
Project:	**OpenDR** (ICT-10-2019-2020 RIA)	
Nature:	Report	**Dissemination Level: PU**blic
Authors:	Anastasios Tefas (AUTH), Nikos Nikolaidis (AUTH), Eleni Chantzi (AUTH), Pavlos Tosidis (AUTH), Theodoros Manousis (AUTH), Maria Tzelepi (AUTH), Paraskevi Nousi (AUTH), Avramelou Loukia (AUTH), Tsampazis Konstantinos (AUTH), Mpampis Emmanouil (AUTH), Efstratios Kakaletsis (AUTH), Kirtas Emmanouil (AUTH), Katsikas Dimitrios (AUTH), Moustakidis Vasileios (AUTH), Passalis Nikolaos (AUTH), Symeonidis Charalampos (AUTH), Spanos Dimitrios (AUTH) Olivier Michel (CYB), Niclas Vödisch (ALU-FR), Daniel Honerkamp (ALU-FR), Nicolai Dorka (ALU-FR), Tim Welschehold (ALU-FR), Abhinav Valada (ALU-FR), Gizem Bozdemir (PAL), Alea Scovill (AGI), Halil Ibrahim Ugurlu (AU), Amir Ramezani Dooraki (AU), Erdal Kayacan (AU), Alexandros Iosifidis (AU), Akif Ekrekli (TAU), Mikael Petäjä (TAU), Roel Pieters (TAU), Anton Muravev (TAU), Moncef Gabbouj (TAU)	
Lead Beneficiary:	AUTH (Aristotle University of Thessaloniki)	
WP	9	
Doc ID:	OPENDR_D9.8.pdf	

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

Document History

Version	Date	Reason of change
1.0	24/10/2023	Deliverable outline
2.0	15/12/2023	First complete version ready for internal review
3.0	28/12/2023	Final version incorporating the reviewer comments

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

Contents

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

Executive Summary

Dissemination and exploitation plans are very important for a collaborative project, since they make project results and findings known to the industry, academia and the general public, thus maximizing the project's impact. This deliverable, D9.8 "Final report on dissemination, exploitation plans and project newsletter", is a document presenting the project dissemination and exploitation activities within the final year of its lifecycle. It contains information regarding scientific publications describing the project's technical achievements, project-related presentations at various events, keynote talks, seasonal schools, as well as other means of dissemination used by the project, such as the project website and social media channels. In addition, the deliverable includes an update on the partners exploitation plans. The 4th and final project newsletter is also included. Finally, a short overview of all project achievements in terms of dissemination, throughout its duration, is provided. Overall, the consortium believes that the project performed extremely well and in accordance with the plan, in terms of both dissemination and exploitation planning.

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

1. Introduction

The OpenDR project aims at developing a modular, open and non-proprietary toolkit for core robotic functionalities by using Deep Learning to enable advanced perception and cognition capabilities, thus meeting the general requirements of robotics applications, but also focusing in the application areas of healthcare, agri-food and agile production. Obviously, dissemination and communication activities are very important for OpenDR, since they help in making the project results known to the industry, academia and the general public, therefore increasing its impact. This deliverable presents the project dissemination activities within the fourth and final year (Y4) of the project, namely activities in the period M37-M48. It contains (in Section 2) information regarding scientific publications, organization of the 2nd project summer school, plenary speeches, participation to events, as well as information regarding the project website and social media channels, dissemination material that has been created, collaboration with other projects, to mention just a few. Information regarding the performance of dissemination and communication activities in Y4, presented through several Key Performance Indicators (KPIs) and their respective target values defined in the DoA, is also presented.

Effective exploitation of the project results is also very important for a project, especially for its industrial partners. In Section 3, the final exploitation plans of the three SMEs that are part of the project consortium, as well as other partners are presented, whereas Section 4 includes an overview of the project's handling of IPR protection and management. The fourth project newsletter, an important element of the project's communication strategy, is briefly described in Section 5 and included in Appendix I. Finally, conclusions and an overview of the project achievements in terms of dissemination throughout its duration are included in Section 6

The presented results and activities indicate that the project's achievements in terms of both dissemination and exploitation planning in both this period and its entire duration were very important. The project outcomes, including the toolkit, were widely disseminated and efficiently exploited.

2. Dissemination and communication activities

2.1 Dissemination and communication activities

2.1.1 Project Website

The OpenDR website[1] has been the main information tool for the project's dissemination activities since M2,when it was created. Maintained and kept up-to-date by AUTH with contributions from all partners, the website reached by M48 213 posts and was visited more than 14900 times in total. Website visitors had the chance to read a variety of documents and posts, such as publications (Figure 1), information about events and the latest releases of the toolkit (Figures 2,3), short and easy to understand articles (OpenDR briefs, Figure 4) etc.

[1] https://opendr.eu/

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

59	ActiveFace: A Synthetic Active Perception Dataset for Face Recognition	C. Georgiadis, N. Passalis, N. Nikolaidis	IEEE 25th International Workshop on Multimedia Signal Processing (MMSP), 2023,*Disclaimer
60	Enabling High-Resolution Pose Estimation in Real Time Using Active Perception	T. Manousis, N. Passalis and A. Tefas	IEEE International Conference on Image Processing (ICIP), 2023,*Disclaimer
61	Deep Active Perception for Object Detection using Navigation Proposals	S. Ginargiros, N.Passalis and A. Tefas	IEEE Symposium Series on Computational Intelligence (SSCI), 2023,*Disclaimer
62	Deep Active Robotic Perception for Improving Face Recognition Under Occlusions	V. Dimaridou, N. Passalis and A. Tefas	IEEE Symposium Series on Computational Intelligence (SSCI), 2023,*Disclaimer
63	CoDEPS: Online Continual Learning for Depth Estimation and Panoptic Segmentation	N. Vödisch, K. Petek, W. Burgard, A. Valada	Robotics: Science and Systems (RSS), 2023,

Figure 1: Some of the latest publications (conferences) listed on the website

The final version of the OpenDR toolkit has been released!

The **final version (v3.0)** of the **Open Deep Learning Toolkit for Robotics** (OpenDR) has just been released! This version includes lots of new tools for face recognition, pose estimation, human detection, object detection, gesture recognition, fall and wave detection, demos etc ! More specifically it includes:

- Binary High Resolution Learner
- ROS2 node for EfficientLPS
- Fall and wave detection ROS nodes
- Continual SLAM: Adds a new SLAM tool for Continual SLAM

Figure 2: Announcement of the OpenDR toolkit release on the website

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

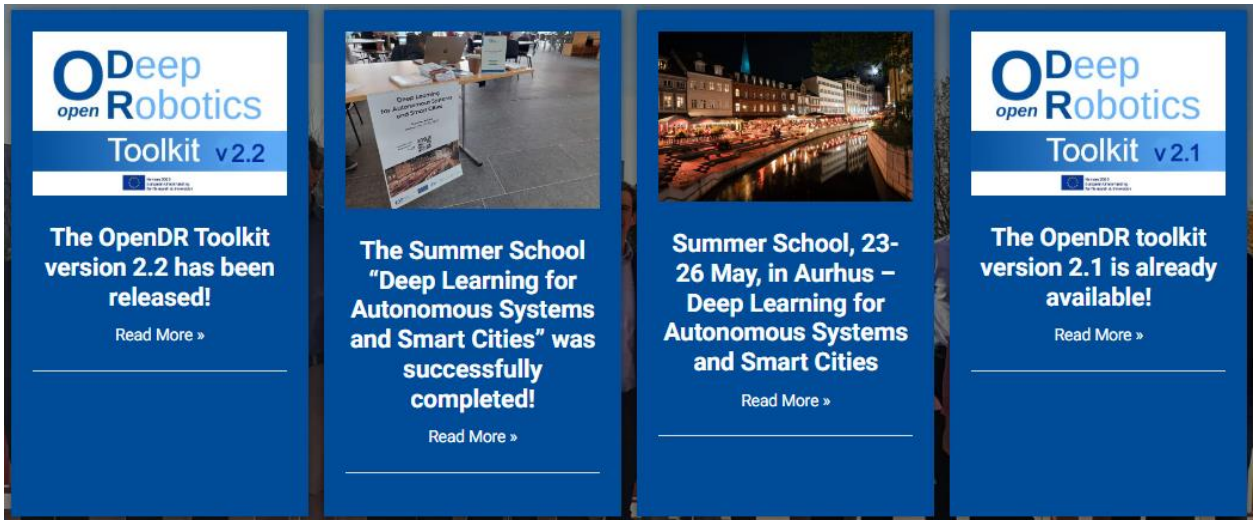

Figure 3: Some of the project's news listed in the website

09	Domain Shift Evaluation and Alleviation Methods in OpenDR powered by Deep Learning	Paraskevi Nousi and AUTH team	Aristotle University of Thessaloniki, Greece	13 January 2023
10	Deep learning for robotics in agile production	Roel Pieters and TAU team	Tampere University, Finland	10 March 2023
11	Continual Inference: A new paradigm for efficient online processing with deep neural networks	Lukas Hedegaard, Negar Heidari, and Alexandros Iosifidis	Aarhus University, Denmark	19 June 2023
12	PARTNR: Pick and place Ambiguity Resolving by Trustworthy iNteractive learning	Jelle Luijkx and TU Delft team	Delft University of Technology - TU Delft, Netherlands	21 July 2023
13	Continual Learning for Robotics	Niclas Vodich and Abhinav Valada	Albert-Ludwigs University Freihburg, Germany	4 September 2023
14	Simulated Humanoid Wrestling Competition in webots.cloud	Olivier Michel	Cyberbotics Ltd., Switzerland	5 October 2023

Figure 4: Brief articles authored by OpenDR partners during the period M37-M48

2.1.2 Social Media

Throughout the four-year period of its life cycle, OpenDR was very active in the most popular dissemination channels, namely LinkedIn, Facebook, Twitter and YouTube. Indeed plenty of posts have been published, the content of which dealt with the scientific papers that were produced, the OpenDR toolkit releases, the attendance of OpenDR members in international conferences and workshops, the events that were organized in the framework of the project etc. Examples and information regarding the social media accounts content are provided in the following sections:

Twitter (@OpenDR_EU)

The twitter account[2] has undoubtedly been on a rising path regarding the number of followers and the tweets/retweets (Figure 5) during the last year of the project. Indeed, the OpenDR twitter account counts today more than 600 followers and more than 285 tweets/retweets, whereas the numbers for M36 were 424 followers and 242 tweets/retweets respectively.

[2] https://twitter.com/OpenDR_EU

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

Figure 5: Examples of tweets posted during the period M37-M48

Facebook (@Opendr.eu)

The Facebook page[3] of the project, although it was not promised in the DoA, was created to support the dissemination activities and was used to a great extent for one more year. As a result, the page counts today (M48) more than 350 likes and more than 250 posts (Figure 6) the content of which is related to the project's publications, participation in conferences and workshops, organization of events, etc.

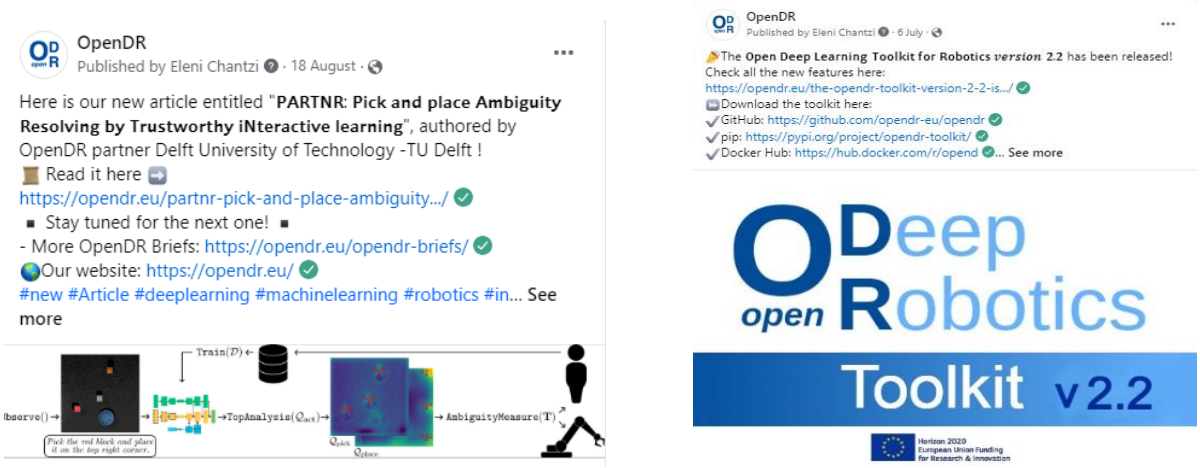

Figure 6: Examples of Facebook posts during the period M37-M48

LinkedIn

OpenDR had a notable presence in the LinkedIn platform, which was used extensively during the last year in order to disseminate the project's activities. In terms of numbers, the official LinkedIn page "OpenDR Research Project[4]" gathered throughout the four years of the project's life more than 1030 followers while today the LinkedIn group counts 172 members. The main content of the posts here was also related to the accepted journal and conference papers, organised events, brief articles, toolkit dissemination, etc. (Figure 7).

[3] https://www.facebook.com/Opendr.eu

[4] https://www.linkedin.com/company/opendr-research-project

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

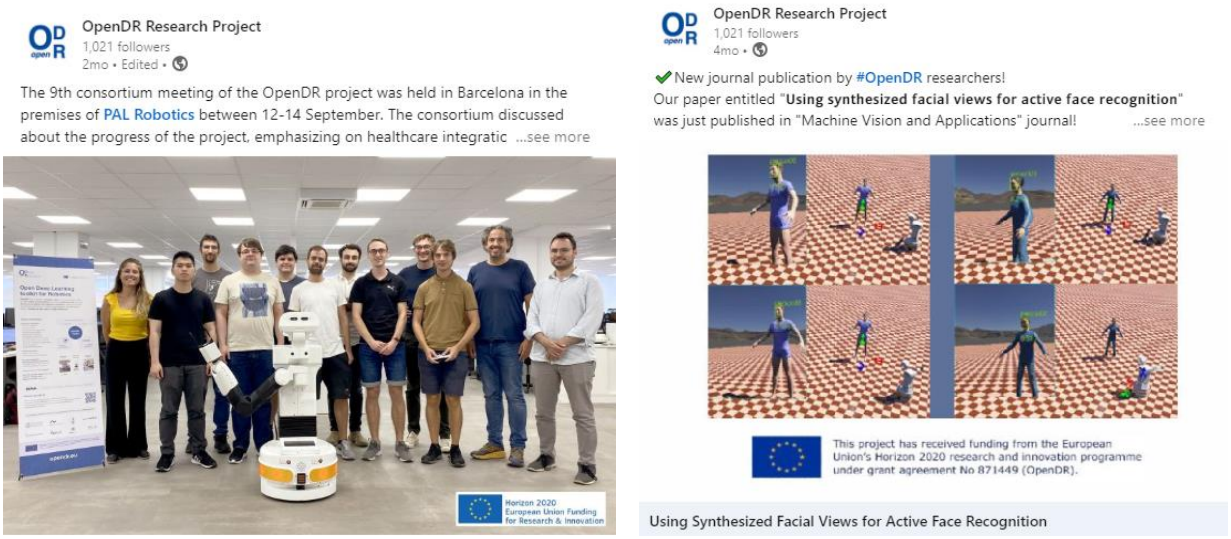

Figure 7: Examples of LinkedIn posts during the period M37-M48

YoutTube

During the period M37-48 the consortium continued to upload videos in the Youtube channel "OpenDR[5]". 18 videos, either promotional or related to conference presentations were posted, and have gathered in total 5425 views (+1925 since M36), while the Youtube channel counts 82 subscribers (+20 since M36).

2.1.3 Dissemination Material (Press Release, Brochures, Factsheet)

The second Press Release of the OpenDR project was created by AUTH on M46. The Press Release (Figure 8) provides information about the toolkit, the aims, achievements and impact of the project so far, the successful dissemination activities of the project (conference and journal publications), etc. The second Press Release was uploaded on the project website and promoted through the project social media accounts and through emails. Subsequently, it was sent to OpenDR partners for further distribution through their dissemination channels (Figure 9). One more press release will be distributed just after the end of the project (early January 2024).

[5] https://www.youtube.com/@opendr7620

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

Thessaloniki (Greece) October 2023 – OpenDR "Open Deep Learning Toolkit for Robotics" is a research project funded by the **EU Horizon 2020** research and innovation program. **The project was launched in January 2020 and since then develops a modular, open and non-proprietary toolkit for core robotic functionalities by harnessing AI, more specifically deep learning, to provide advanced perception and cognition capabilities, meeting in this way the general requirements of robotics applications in numerous areas including healthcare, agri-food and agile production, which are the project's three use-case areas.** The project brings together 5 top-ranked academic and research institutes: Aristotle University of Thessaloniki (Greece), Tampere University (Finland), Aarhus University (Denmark), Delft University of Technology (Netherlands), University of Freiburg (Germany) and 3 leading industry partners: Cyberbotics (Switzerland), PAL Robotics (Spain) and AgroIntelli (Denmark). **The project is coordinated by Prof. Anastasios Tefas, Aristotle University of Thessaloniki.**

Five versions of the OpenDR toolkit have been publicly released so far. The first one was released in December 2021 providing more than 20 methods related to core robotic functionalities, an intuitive and easy-to-use Python interface, a C language API (Application Programming Interface) for selected tools, a wealth of usage examples and supporting tools, as well as ready-to-use ROS (Robot Operating System) nodes. Since then, 4 additional versions were released, expanding the number of available methods, as well as providing performance improvements to existing ones. In addition, several other improvements were made, including the addition of modular installation options, the support for newer CUDA versions, the implementation of a refined ROS/ROS2 interface, etc. The latest version was released on 3rd of July and provides, among others, tools for high resolution analysis, object detection and tracking, efficient LiDAR-based panoptic segmentation, continual inference transformers, facial emotion estimation and multi-object search, along with several bug fixes and a more rich C API for several tools. The toolkit is freely accessible in GitHub, the world's most popular software development and sharing platform.

Through the toolkit and the respective research, OpenDR aims, and so far succeeds, in having a significant impact by improving the technical capabilities in robotics through the provision of easily deployable, efficient and novel Deep Learning tools, as well as by lowering 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. Usage statistics seem to verify that these targets will be achieved. Indeed, the acceptance of the toolkit from the robotics, deep learning and computer vision communities has been extremely encouraging so far, the GitHub repository was awarded more than 540 stars from its users, whereas the toolkit as a whole or individual tools have been downloaded more than 17000 times since its first release.

In addition, **an impressive number of scientific journal and conference papers, more than 90 so far,** have been generated throughout this almost four-year period. Indeed, OpenDR has published the results of its groundbreaking research in prestigious robotics and AI journals, including IEEE (Institute of Electrical and Electronics Engineers), such as IEEE Transactions on Neural Networks and Learning Systems, IEEE Transactions on Artificial Intelligence, IEEE Transactions on Image processing, IEEE Robotics and Automation Letters and Neurocomputing, as well as highly influential conferences including IEEE Conference on Computer Vision and Pattern Recognition (CVPR), IEEE International Conference on Multimedia and Expo (ICME), IEEE International Conference on Image Processing (ICIP), IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) etc.

The project has also organized two very successful Summer Schools one on "Continuous Engineering and Deep Learning for Autonomous Trustworthy Systems" in Thessaloniki, Greece, in cooperation with FOCETA H2020 Project and one on "Deep Learning for Autonomous Systems and Smart Cities", in Aarhus, Denmark, in cooperation with Aarhus University. Both Summer Schools were very well attended. These events brought together AI experts and enthusiasts to share knowledge and foster learning in the areas of deep learning and robotics. **Other significant dissemination activities** included organization of a tutorial on Open And Trustworthy Deep Learning for Robotics in the highly esteemed International Conference on Intelligent Robots and Systems (IROS 2022), a workshop on Open Deep Learning Toolkit for Robotics: Towards Democratizing Artificial Intelligence at IROS 2021, organization of special sessions in conferences, invited talks etc.

The very good progress achieved by the consortium was verified in the two review meetings that have been conducted by the EU so far. In these meetings the consortium had the opportunity to present its results to external experts (reviewers) that praised the project's achievements and provided constructive comments for its forthcoming activities.

The project has received funding from the European Union's Horizon 2020 research and innovation program under grant agreement No 871449.

For more information visit **OpenDR's** website: www.opendr.eu

For updates follow us also in:
https://github.com/opendr-eu/opendr
https://www.linkedin.com/in/opendr-research-project-94a868199/
https://www.facebook.com/Opendr.eu
https://twitter.com/OpenDR_EU
https://www.youtube.com/@opendr7620

Figure 8: The 2nd Press Release of the OpenDR project

Figure 9: The promotion of the 2nd Press Release (in Greek language) through the Aristotle University of Thessaloniki (AUTH) social media accounts

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

Moreover, updated versions of the OpenDR brochure (Figure 10) and the project factsheet were created and placed on the project website. These new editions include important information about the features of the OpenDR Toolkit Version 2.2 and have been distributed as hard copies in several international conferences/workshops such as MMSP 2023, ICIP 2023 and SSCI 2023.

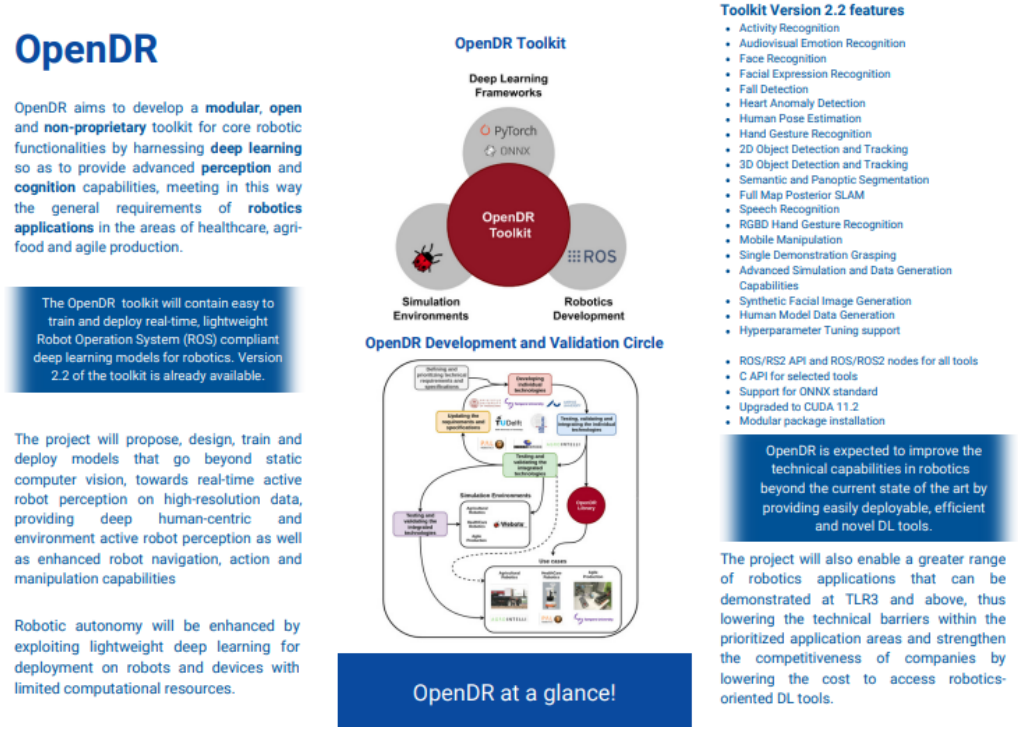

Figure 10: The updated version of the OpenDR brochure including information about the toolkit version 2.2

2.1.4 GitHub

Following the first release of the toolkit on M24, the OpenDR consortium released five more public versions of the toolkit until M48 (six public versions in total), through GitHub. The latest versions focused on expanding the number of available methods, as well as providing performance improvements to existing ones. In addition, several other improvements were made, including the addition of modular installation options, the support for newer CUDA versions, the implementation of a refined ROS/ROS2 interface, etc. The final version of the OpenDR toolkit, tagged 3.0.0, was released on December 2023 and includes several new tools and provides new functionalities, such as gesture recognition, FSeq2 non-maximum suppression, continual SLAM, intent recognition, object detection class filtering, object detection models for agricultural use cases, RL-based active perception, adaptive high-resolution pose estimation, and others. Furthermore, several existing tools have been enhanced by fixing bugs and including additional demos and ROS nodes. Also, the installation process is now more robust and efficient, while the development pipelines have been significantly improved allowing for more easily developing new tools. The new releases were announced and promoted actively through social media. The response of the public and especially of the scientific and robotics community to the OpenDR toolkit was extremely positive. Indeed, until M48 the toolkit gathered 567 stars, more than 17000 downloads and 86 forks in the GitHub platform.

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

2.1.5 Publications

2.1.5.1 Journal Papers

Research conducted within the project during this period led to eight (8) journal publications that have been accepted for publication or published.

AU published one paper in Procedia Computer Science:

- Illia Oleksiienko, Dat Thanh Tran, Alexandros Iosifidis, “Variational Neural Networks”, Procedia Computer Science, 222, pp. 104-113, 2023.

Bayesian Neural Networks provide a tool to estimate the uncertainty of a neural network by considering a distribution over weights and sampling different models for each input. In this paper, we propose a method for uncertainty estimation in neural networks which, instead of considering a distribution over weights, samples outputs of each layer from a corresponding Gaussian distribution, parametrized by the predictions of mean and variance sub-layers. In uncertainty quality estimation experiments, we show that the proposed method achieves better uncertainty quality than other single-bin Bayesian Model Averaging methods, such as Monte Carlo Dropout or Bayes By Backpropagation methods.

AUTH published one paper in Machine Vision and Applications journal:

- Efstratios Kakaletsis, Nikos Nikolaidis, “Using synthesized facial views for active face recognition”, Machine Vision and Applications, Volume 34, Issue 4, Jul 2023.

Active perception/vision exploits the ability of robots to interact with their environment, for example move in space, towards increasing the quantity or quality of information obtained through their sensors and, thus, improving their performance in various perception tasks. Active face recognition is largely understudied in recent literature. Attempting to tackle this situation, in this paper, we propose an active approach that utilizes facial views produced by photorealistic facial image rendering. Essentially, the robot that performs the recognition selects the best among a number of candidate movements around the person of interest by simulating their results through view synthesis. This is accomplished by feeding the robot’s face recognizer with a real-world facial image acquired in the current position, generating synthesized views that differ by from the current view and deciding, based on the confidence of the recognizer, whether to stay in place or move to the position that corresponds to one of the two synthesized views, in order to acquire a new real image with its sensor. Experimental results in three datasets verify the superior performance of the proposed method compared to the respective “static” approach, approaches based on the same face recognizer that involve synthetic face frontalization and synthesized views, random direction robot movement, robot movement towards a frontal location based on view angle estimation, as well as a state of the art active method. Results from a proof of concept simulation in a robotic simulator are also provided.

ALU-FR published four papers in the IEEE Robotics and Automation Letters (RA-L):

- J. Arce, N. Vödisch, D. Cattaneo, W. Burgard, and A. Valada, "PADLoC: LiDAR-Based Deep Loop Closure Detection and Registration using Panoptic Attention", IEEE Robotics and Automation Letters (RA-L), vol. 8, issue 3, pp. 1319-1326, March 2023.

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

A key component of graph-based SLAM systems is the ability to detect loop closures in a trajectory to reduce the drift accumulated over time from the odometry. Most LiDAR-based methods achieve this goal by using only the geometric information, disregarding the semantics of the scene. In this work, the authors introduce PADLoC, a LiDAR-based loop closure detection and registration architecture comprising a shared 3D convolutional feature extraction backbone, a global descriptor head for loop closure detection, and a novel transformer-based head for point cloud matching and registration. They present multiple methods for estimating the point-wise matching confidence based on diversity indices. Additionally, to improve forward-backward consistency, they propose the use of two shared matching and registration heads with their source and target inputs swapped by exploiting that the estimated relative transformations must be inverse of each other. Furthermore, they leverage panoptic information during training in the form of a novel loss function that reframes the matching problem as a classification task in the case of the semantic labels and as a graph connectivity assignment for the instance labels. They perform extensive evaluations of PADLoC on multiple real-world datasets demonstrating that it achieves state-of-the-art performance.

- Fabian Schmalstieg, Daniel Honerkamp, Tim Welschehold, Abhinav Valada, Learning Hierarchical Interactive Multi-Object Search for Mobile Manipulation IEEE Robotics and Automation Letters (RA-L), 2023, 2023.

Existing object-search approaches enable robots to search through free pathways, however, robots operating in unstructured human-centered environments frequently also have to manipulate the environment to their needs. In this work, we introduce a novel interactive multi-object search task in which a robot has to open doors to navigate rooms and search inside cabinets and drawers to find target objects. These new challenges require combining manipulation and navigation skills in unexplored environments. We present HIMOS, a hierarchical reinforcement learning approach that learns to compose exploration, navigation, and manipulation skills. To achieve this, we design an abstract high-level action space around a semantic map memory and leverage the explored environment as instance navigation points. We perform extensive experiments in simulation and the real world that demonstrate that, with accurate perception, the decision making of HIMOS effectively transfers to new environments in a zero-shot manner. It shows robustness to unseen subpolicies, failures in their execution, and different robot kinematics. These capabilities open the door to a wide range of downstream tasks across embodied AI and real-world use cases. Code and videos are publicly available at http://himos.cs.uni-freiburg.de/.

- A. Younes, D. Honerkamp, T. Welschehold and A. Valada, “Catch Me If You Hear Me: Audio-Visual Navigation in Complex Unmapped Environments with Moving Sounds”, Robotics and Automation Letters (RA-L), 8(2), 2023

Audio-visual navigation combines sight and hearing to navigate to a sound-emitting source in an unmapped environment. While recent approaches have demonstrated the benefits of audio input to detect and find the goal, they focus on clean and static sound sources and struggle to generalize to unheard sounds. In this work, the authors propose the novel dynamic audio-visual navigation benchmark which requires catching a moving sound source in an environment with noisy and distracting sounds, posing a range of new challenges. They introduce a reinforcement learning approach that learns a robust navigation policy for these complex settings. To achieve this, they propose an architecture that fuses audio-visual information in the spatial feature space to learn correlations of geometric information inherent in both local maps and audio signals. They demonstrate that the proposed approach consistently outperforms the current state-of-the-art by a

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

large margin across all tasks of moving sounds, unheard sounds, and noisy environments, on two challenging 3D scanned real-world environments, namely Matterport3D and Replica. The benchmark is available at http://dav-nav.cs.uni-freiburg.de/ .

- N. Dorka, T. Welschehold, J.Boedecker, and W. Burgard, "Adaptively Calibrated Critic Estimates for Deep Reinforcement Learning", Robotics and Automation Letters (RA-L), 2023

Accurate value estimates are important for off-policy reinforcement learning. Algorithms based on temporal difference learning typically are prone to an over- or underestimation bias building up over time. In this paper, the authors propose a general method called Adaptively Calibrated Critics (ACC) that uses the most recent high variance but unbiased on-policy rollouts to alleviate the bias of the low variance temporal difference targets. They apply ACC to Truncated Quantile Critics, which is an algorithm for continuous control that allows regulation of the bias with a hyperparameter tuned per environment. The resulting algorithm adaptively adjusts the parameter during training rendering hyperparameter search unnecessary and sets a new state of the art on the OpenAI gym continuous control benchmark among all algorithms that do not tune hyperparameters for each environment. ACC further achieves improved results on different tasks from the Meta-World robot benchmark. Additionally, the authors demonstrate the generality of ACC by applying it to TD3 and showing an improved performance also in this setting.

Further, ALU-FR published one paper in the IEEE Transactions on Robotics Journal (T-RO):

- D. Honerkamp, T. Welschehold and A. Valada, “N2M2: Learning Navigation for Arbitrary Mobile Manipulation Motions in Unseen and Dynamic Environments”, IEEE Transactions on Robotics (T-RO), vol. 39, no. 5, pp. 3601-3619, Oct. 2023, doi: 10.1109/TRO.2023.3284346.

Despite its importance in both industrial and service robotics, mobile manipulation remains a significant challenge as it requires a seamless integration of end-effector trajectory generation with navigation skills as well as reasoning over long-horizons. Existing methods struggle to control the large configuration space, and to navigate dynamic and unknown environments. In previous work, the authors proposed to decompose mobile manipulation tasks into a simplified motion generator for the end-effector in task space and a trained reinforcement learning agent for the mobile base to account for kinematic feasibility of the motion. In this work, they introduce Neural Navigation for Mobile Manipulation (N2M2) which extends this decomposition to complex obstacle environments and enables it to tackle a broad range of tasks in real world settings. The resulting approach can perform unseen, long-horizon tasks in unexplored environments while instantly reacting to dynamic obstacles and environmental changes. At the same time, it provides a simple way to define new mobile manipulation tasks. The authors demonstrate the capabilities of the proposed approach in extensive simulation and real-world experiments on multiple kinematically diverse mobile manipulators. Code and videos are publicly available at http://mobile-rl.cs.uni-freiburg.de.

TAU published one paper in Robotics and Computer Integrated Manufacturing:

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

- A. Angleraud, A. Ekrekli, K. Samarawickrama, G. Sharma and R. Pieters, “Sensor-based human-robot collaboration for Industrial Tasks”, Robotics and Computer Integrated Manufacturing (RCIM), vol. 86, p. 102663, April 2024, DOI: 10.1016/j.rcim.2023.102663

Collaboration between human and robot requires interaction modalities that suit the context of the shared tasks and the environment in which it takes place. While an industrial environment can be tailored to favour certain conditions (e.g., lighting), some limitations cannot so easily be addressed (e.g., noise, dirt). In addition, operators are typically continuously active and cannot spare long time instances away from their tasks engaging with physical user interfaces. Sensor-based approaches that recognize humans and their actions to interact with a robot have therefore great potential. This work demonstrates how human–robot collaboration can be supported by visual perception models, for the detection of objects, targets, humans and their actions. For each model we present details with respect to the required data, the training of a model and its inference on real images. Moreover, we provide all developments for the integration of the models to an industrially relevant use case, in terms of software for training data generation and human–robot collaboration experiments. These are available open-source in the OpenDR toolkit at https://github.com/opendr-eu/opendr. Results are discussed in terms of performance and robustness of the models, and their limitations. Although the results are promising, learning-based models are not trivial to apply to new situations or tasks. Therefore, we discuss the challenges identified, when integrating them into an industrially relevant environment.

2.1.5.2 Conference Papers

Eighteen (18) conference papers that describe results obtained within the project were published or accepted for presentation during this period.

TUD presented one paper at the IEEE 21st International Conference on Industrial Informatics (INDIN) and another one at the Sixteenth European Workshop on Reinforcement Learning:

- F. Sibona, J. Luijkx, B. van der Heijden, L. Ferranti, and M. Indri “EValueAction: a proposal for policy evaluation in simulation to support interactive imitation learning”, IEEE 21st International Conference on Industrial Informatics (INDIN) (accepted), 2023.

The up-and-coming concept of Industry 5.0 foresees human-centric flexible production lines, where collaborative robots support human workforce. In order to allow a seamless collaboration between intelligent robots and human workers, designing solutions for non-expert users is crucial. Learning from demonstration emerged as the enabling approach to address such a problem. However, more focus should be put on finding safe solutions which optimize the cost associated with the demonstrations collection process. This paper introduces a preliminary outline of a system, namely EValueAction (EVA), designed to assist the human in the process of collecting interactive demonstrations taking advantage of simulation to safely avoid failures. A policy is pre-trained with human-demonstrations and, where needed, new informative data are interactively gathered and aggregated to iteratively improve the initial policy. A trial case study further reinforces the relevance of the work by demonstrating the crucial role of informative demonstrations for generalization.

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

- A. Keijzer, B. van der Heijden, and J. Kober “Prioritizing States with Action Sensitive Return in Experience Replay”, Sixteenth European Workshop on Reinforcement Learning, 2023.

Experience replay for off-policy reinforcement learning has been shown to improve sample efficiency and stabilize training. However, typical uniformly sampled replay includes many irrelevant samples for the agent to reach good performance. We introduce Action Sensitive Experience Replay (ASER), a method to prioritize samples in the replay buffer and selectively model parts of the state-space more accurately where choosing sub-optimal actions has a larger effect on the return. We experimentally show that this can make training more sample efficient and that this allows smaller parametric function approximators – like neural networks with few neurons – to achieve good performance in environments where they would otherwise struggle.

ALU-FR presented three papers at the Robotics: Science and Systems (RSS) conference, the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshop for Continual Learning and the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS):

- N. Vödisch, K. Petek, W. Burgard, and A. Valada, “CoDEPS: Online Continual Learning for Depth Estimation and Panoptic Segmentation”, Robotics: Science and Systems (RSS), 2023.

Operating a robot in the open world requires a high level of robustness with respect to previously unseen environments. Optimally, the robot is able to adapt by itself to new conditions without human supervision, e.g., automatically adjusting its perception system to changing lighting conditions. In this work, we address the task of continual learning for deep learning-based monocular depth estimation and panoptic segmentation in new environments in an online manner. We introduce CoDEPS to perform continual learning involving multiple real-world domains while mitigating catastrophic forgetting by leveraging experience replay. In particular, we propose a novel domain-mixing strategy to generate pseudo-labels to adapt panoptic segmentation. Furthermore, we explicitly address the limited storage capacity of robotic systems by leveraging sampling strategies for constructing a fixed-size replay buffer based on rare semantic class sampling and image diversity. We perform extensive evaluations of CoDEPS on various real-world datasets demonstrating that it successfully adapts to unseen environments without sacrificing performance on previous domains while achieving state-of-the-art results.

- N. Vödisch, D. Cattaneo, W. Burgard, and A. Valada, “CoVIO: Online Continual Learning for Visual-Inertial Odometry”, IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops, 2023.

Visual odometry is a fundamental task for many applications on mobile devices and robotic platforms. Since such applications are oftentimes not limited to predefined target domains and learning-based vision systems are known to generalize poorly to unseen environments, methods for continual adaptation during inference time are of significant interest. In this work, we introduce CoVIO for online continual learning of visual-inertial odometry. CoVIO effectively adapts to new domains while mitigating catastrophic forgetting by exploiting experience replay. In particular, we propose a novel sampling strategy to maximize image diversity in a fixed-size replay buffer that

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

targets the limited storage capacity of embedded devices. We further provide an asynchronous version that decouples the odometry estimation from the network weight update step enabling continuous inference in real time. We extensively evaluate CoVIO on various real-world datasets demonstrating that it successfully adapts to new domains while outperforming previous methods.

- J. Vertens, N. Dorka, T. Welschehold, M. Thompson, and W. Burgard "Improving Deep Dynamics Models for Autonomous Vehicles with Multimodal Latent Mapping of Surfaces", IEEE/RSJ International Conference on Intelligent Robots and Systems, 2023.

The safe deployment of autonomous vehicles relies on their ability to effectively react to environmental changes. This can require maneuvering on varying surfaces which is still a difficult problem, especially for slippery terrains. To address this issue, we propose a new approach that learns a surface-aware dynamics model by conditioning it on a latent variable vector storing surface information about the current location. A latent mapper is trained to update these latent variables during inference from multiple modalities on every traversal of the corresponding locations and stores them in a map. By training everything end-to-end with the loss of the dynamics model, we enforce the latent mapper to learn an update rule for the latent map that is useful for the subsequent dynamics model. We implement and evaluate our approach on a real miniature electric car. The results show that the latent map is updated to allow more accurate predictions of the dynamics model compared to a model without this information. We further show that by using this model, the driving performance can be improved on varying and challenging surfaces.

AUTH members published in total 9 conference papers during this period:

- Manousis, Theodoros, Nikolaos Passalis, and Anastasios Tefas. "Enabling High-Resolution Pose Estimation in Real Time Using Active Perception." IEEE International Conference on Image Processing (ICIP), 2023.

Deep Learning (DL) models have enabled very accurate pose estimation. However, most of the existing approaches require images of relatively high resolution, since locating body parts and joints accurately is challenging, which increases the computational cost of these approaches. To overcome this limitation in this paper we propose an active perception method for high-resolution pose estimation that enables efficiently selecting the most appropriate image region for analysis and then employing a bottom-up pose estimator on the corresponding region. This allows for significantly improving the efficiency of pose estimation by selectively analyzing in high resolution only the parts of the image that contain humans. To ensure the computational efficiency of the proposed method we propose using low-resolution heat maps extracted using the same pose estimation model in order to guide the active perception process. The proposed method is model agnostic since it can be combined with any bottom-up pose estimation model in order to enable high-resolution analysis. We have experimentally evaluated the proposed method using a well-known pose estimation model, Lightweight OpenPose, demonstrating its effectiveness on three high-resolution variants of the COCO2017 dataset.

- Konstantinos Tsampazis, Manos Kirtas, Pavlos Tosidis, Nikolaos Passalis, and Anastasios Tefas. "Deep Reinforcement Learning with Action Masking for Differential-Drive Robot Navigation using Low-cost Sensors." IEEE International Workshop on Machine Learning for Signal Processing, 2023.

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

Driving a wheeled differential-drive robot to a target can be a complicated matter when trying to also avoid obstacles. Usually, such robots employ a variety of sensors, such as LiDAR, depth cameras, and others, that can be quite expensive. To this end, in this paper, we focus on a simple differential-drive wheeled robot that uses only inexpensive ultrasonic distance sensors and touch sensors. We propose a method for training a Reinforcement Learning (RL) agent to perform robot navigation to a target while avoiding obstacles. In order to increase the efficiency of the proposed approach we design appropriate action masks that can significantly increase the learning speed and effectiveness of the learned policy. As we experimentally demonstrated, the proposed agent can robustly navigate to a given target even in unknown procedurally generated environments, or even when denying part of its sensor input. Finally, we show a practical use-case using object detection to dynamically search for, and move to objects within unknown environments.

- Manos Kirtas, Konstantinos Tsampazis, Loukia Avramelou, Nikolaos Passalis, and Anastasios Tefas. "Using Part-based Representations for Explainable Deep Reinforcement Learning." ECML PKDD – Workshop: Uncertainty meets Explainability, 2023.

Utilizing deep learning models to learn part-based representations holds significant potential for interpretable-by-design approaches, as these models incorporate latent causes obtained from feature representations through simple addition. However, training a part-based learning model presents challenges, particularly in enforcing non-negative constraints on the model's parameters, which can result in training difficulties such as instability and convergence issues. Moreover, applying such approaches in Deep Reinforcement Learning (RL) is even more demanding due to the inherent instabilities that impact many optimization methods. In this paper, we propose a non-negative training approach for actor models in RL, enabling the extraction of part-based representations that enhance interpretability while adhering to non-negative constraints. To this end, we employ a non-negative initialization technique, as well as a modified sign-preserving training method, which can ensure better gradient flow compared to existing approaches. We demonstrate the effectiveness of the proposed approach using the well-known Cartpole benchmark.

- Charalampos Georgiadis, Nikolaos Passalis and Nikos Nikolaidis. "ActiveFace: A Synthetic Active Perception Dataset for Face Recognition." IEEE International Workshop on Multimedia Signal Processing, 2023.

Active vision aims to enhance the efficiency of computer vision methods by enabling the capturing sensor, usually placed on a robot or, more generally, an autonomous system, to dynamically adjust its viewpoint, position or parameters in real-time. This capability allows for more precise decision-making by the model. However, training and evaluating an active vision model often necessitates a substantial number of annotated images, which must be captured under various sensor and environmental settings. These diverse images enable the model to learn the underlying dynamics of the active perception process. Unfortunately, collecting and annotating such datasets is a challenging and expensive task. It involves not only providing hand-crafted ground truth annotations but also ensuring that actions, such as moving around / towards / away from a person, are properly "imitated" to enable active vision approaches to model the corresponding active perception dynamics. To address these limitations, in this paper we propose a synthetic facial image generation pipeline specifically designed to support active face recognition, developed using a highly realistic simulation framework based on Unity. The developed pipeline allows for the generation of facial images for a set of persons at various view angles, distances, illumination conditions, and backgrounds. We demonstrate the effectiveness of our approach by training and

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

evaluating a recently proposed embedding-based active face recognizer, as well as extending it to perform 2 axis control, leveraging the additional information provided by the generated dataset. To facilitate replication and encourage the use of the generated dataset for training and evaluating other active vision approaches, we also provide the associated assets and the developed dataset generation pipeline.

- Stefanos Ginargyros, Nikolaos Passalis, Anastasios Tefas. "Deep Active Perception for Object Detection using Navigation Proposals." IEEE Symposium Series on Computational Intelligence, 2023.

Deep Learning (DL) has brought significant advances to robotics vision tasks. However, most existing DL methods have a major shortcoming - they rely on a static inference paradigm inherent in traditional computer vision pipelines. On the other hand, recent studies have found that active perception improves the perception abilities of various models by going beyond these static paradigms. Despite the significant potential of active perception, it poses several challenges, primarily involving significant changes in training pipelines for deep learning models. To overcome these limitations, in this work, we propose a generic supervised active perception pipeline for object detection that can be trained using existing off-the-shelf object detectors, while also leveraging advances in simulation environments. To this end, the proposed method employs an additional neural network architecture that estimates better viewpoints in cases where the object detector confidence is insufficient. The proposed method was evaluated on synthetic datasets - constructed within the Webots robotics simulator -, showcasing its effectiveness in two object detection cases.

- Valia Dimaridou, Nikolaos Passalis, Anastasios Tefas. "Deep Active Robotic Perception for Improving Face Recognition Under Occlusions." IEEE Symposium Series on Computational Intelligence, 2023.

Recent studies have demonstrated that active perception can improve the perception abilities of deep learning (DL) models. However, there are challenges associated with using active perception in DL models, including the need for datasets and/or realistic simulations that can support the training process, along with the difficulty of predicting the final target position, which reduces planning efficiency. To address these challenges, this work presents a methodology for enhancing the perception abilities of DL models through active perception. The methodology proposes a way to create datasets for active perception by fusing existing large-scale datasets and decomposing the active perception problem into three sub-tasks for face recognition. The sub-tasks aim to determine the appropriateness of the current view for face recognition, the direction in which the robot should move for a better viewpoint, and the expected amount of movement required. A novel trial-based approach is introduced to estimate the final target position, making the method platform-agnostic and easily applicable to different robots. The proposed methodology is validated through experiments on two well-known face verification datasets that have been augmented with occlusions, demonstrating its effectiveness in enhancing the perception abilities of DL models through active perception.

- Dimitrios Tsiakmakis, Nikolaos Passalis, and Anastasios Tefas. “Improving Inertial-based UAV Localization using Data-efficient Deep Reinforcement Learning”, European Signal Processing Conference, 2023.

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

Precise localization is a critical task for many Unmanned Aerial Vehicle (UAV)-based applications. Inertial-based navigation, which relies on Inertial Measurement Units (IMUs), is extensively used to this end, due to its low-cost and small footprint. However, IMU-based localization leads to accumulating significant localization errors. To overcome this limitation, in this paper we propose a data-efficient Deep Reinforcement Learning (DRL) method that enables learning how to correct localization errors from IMUs leading to more precise localization. In contrast with supervised approaches, the proposed method employs a novel data augmentation and regularization approach, which requires collecting a minimal number of real examples, while it is also platform-agnostic and can account for manufacturing impressions. The effectiveness of the proposed method is demonstrated both in a simulation environment, as well as using a real UAV.

- Angelos Nalmpantis, Nikolaos Passalis, and Anastasios Tefas. "End-to-end trainable gaussian filtering for electrocardiogram signal classification using deep learning." European Signal Processing Conference, 2023.

Deep Learning (DL) is increasingly used in electrocardiograms (ECGs)-based signal analysis allowing for diagnosing a variety of cardiac disorders and enabling the development of automated diagnosis and prognosis systems. However, ECG signals are typically noisy, requiring a series of prepossessing steps before feeding them into DL models for analysis. The main contribution of this paper is a simple, yet very effective end-to-end hybrid trainable filtering and feature extraction pipeline, that is formulated as a set of differentiable layers that can be incorporated into any DL model and employ Gaussian filters as a prior model. Indeed, it is experimentally demonstrated that the proposed method can increase the performance of ECG-based classification both for baseline architectures (e.g., MLPs, CNNs, LSTMs) as well as for state-of-the-art architectures proposed in the literature for ECG signal analysis. Given the simplicity and effectiveness of the proposed method, we believe that it constitutes a practical tool that can readily be incorporated into any DL pipeline that is used for time-series analysis, potentially further increasing its performance with minimal computational overhead.

- Efstratios Kakaletsis, Nikos Nikolaidis,"Active Face Recognition through View Synthesis", European Signal Processing Conference, 2023,

Active vision exploits the ability of robots to interact with their environment, towards increasing the quantity / quality of information obtained through their sensors and, therefore, improving their performance in perception tasks. Active face recognition is largely understudied in recent literature. In this paper, we propose an active approach that utilizes facial views produced by facial image rendering. The robot that performs face recognition selects the best candidate rotation around the person of interest by simulating the results of such movements through view synthesis. This is achieved by passing to the robot's face recognizer a real world facial image acquired in the current position, generating synthesized views that differ by $\pm\theta_\circ$ from the current view. Then, it decides, on the basis of the confidence of the recognizer, whether to stay in place or move to the position that corresponds to one of the two synthesized views, so as to to acquire a new real image. Experimental results in two datasets verify the superior performance of the proposed method compared to the respective static approach and an approach based on the same face recognizer that involves face frontalization with synthesized views.

AU presented one paper at IEEE International Workshop on Machine Learning for Signal Processing and another one at ISR Europe 2023; 56th International Symposium on Robotics:

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

- Illia Oleksiienko, Alexandros Iosifidis, “Layer Ensembles”, IEEE International Workshop on Machine Learning for Signal Processing, pp1-6, 2023.

Deep Ensembles, as a type of Bayesian Neural Networks, can be used to estimate uncertainty on the prediction of multiple neural networks by collecting votes from each network and computing the difference in those predictions. In this paper, we introduce a method for uncertainty estimation that considers a set of independent categorical distributions for each layer of the network, giving many more possible samples with overlapped layers than in the regular Deep Ensembles. We further introduce an optimized inference procedure that reuses common layer outputs, achieving up to 19x speed up and reducing memory usage quadratically. We also show that the method can be further improved by ranking samples, resulting in models that require less memory and time to run while achieving higher uncertainty quality than Deep Ensembles.

- Ugurlu, H. I., Bardakci, D., Pham, H. X., & Kayacan, E., (2023) "Guidance of Agricultural Ground Robots Team with an Aerial Vehicle: A Cost-Effective Solution," ISR Europe 2023; 56th International Symposium on Robotics, Stuttgart, Germany.

Increasing the operational efficiency of agricultural machines is essential by the use of artificial intelligence (AI)-based navigation, planning, and control algorithms to handle the increasing demand for food production without compromising sustainability. In this study, an end-to-end path planning algorithm (AgroRL) is proposed for aerial-ground robots team collaboration. In the proposed solution, while main operations in the field are handled by the ground vehicle, the aerial robot is responsible for re-planning a collision-free trajectory for the ground robot when the robot faces an obstacle. Deep reinforcement learning is used for training the end-to-end policy for local re-planning of the aerial robot. The agent, informed by the global trajectory, generates local plans based on depth images. Variational autoencoders are also investigated for dimension reduction of the depth images in obstacle avoidance context to speed up deep reinforcement learning and alleviate the computational complexity of the policy network. The agriculture environment is developed in the Webots open-source robot simulator for training and testing purposes. The efficiency and efficacy of the end-to-end planner are evaluated over a number of cluttered field scenarios. The simulation experiments demonstrate a single aerial vehicle guiding multiple ground robots in agricultural operations.%

Two papers were presented at the IEEE International Conference on Robotic Computing by TAU:

- A. Ekrekli, A. Angleraud, G. Sharma and R. Pieters, “Co-speech gestures for human-robot collaboration”, IEEE International Conference on Robotic Computing, 2023.

Collaboration between human and robot requires effective modes of communication to assign robot tasks and coordinate activities. As communication can utilize different modalities, a multi-modal approach can be more expressive than single modal models alone. In this work we propose a co-speech gesture model that can assign robot tasks for human-robot collaboration. Human gestures and speech, detected by computer vision and speech recognition, can thus refer to objects in the scene and apply robot actions to them. We present an experimental evaluation of the multi-modal co-speech model with a real-world industrial use case. Results demonstrate that multimodal

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

communication is easy to achieve and can provide benefits for collaboration with respect to single modal tools.

- G. Sharma, A. Angleraud, R. Pieters, “Multi-label Annotation for Visual Multi-Task Learning Models”, IEEE International Conference on Robotic Computing, 2023.

Deep learning requires large amounts of data, and a well-defined pipeline for labeling and augmentation. Current solutions support numerous computer vision tasks with dedicated annotation types and formats, such as bounding boxes, polygons, and key points. These annotations can be combined into a single data format to benefit approaches such as multi-task models. However, to our knowledge, no available labeling tool supports the export functionality for a combined benchmark format, and no augmentation library supports transformations for the combination of all. In this work, these functionalities are presented, with visual data annotation and augmentation to train a multi-task model (object detection, segmentation, and key point extraction). The tools are demonstrated in two robot perception use cases.

2.1.5.3 PhD Theses

Six PhD theses which were associated and were, at least partially, funded by the project were completed during its 4-year duration. These theses are listed below:

In June 2023, the PhD thesis entitled “Efficient Online Processing with Deep Neural Networks” was submitted by Lukas Hedegaard to the Graduate School of the Faculty of Technical Sciences at Aarhus University. This thesis is connected to the OpenDR project and includes improvements of neural networks for increasing their efficiency. Specifically, a core contribution addresses the efficiency aspects during online inference, where the system must make an updated prediction swiftly after new input information is received, without waiting for subsequent data. Here, the concept of Continual Inference Networks (CINs) is proposed and explored across four publications. CINs extend prior state-of-the-art methods developed for offline processing of spatio-temporal data and reuse their pre-trained weights, improving their online processing efficiency by an order of magnitude. These advances are attained through a bottom-up reorganization of computational components and judicious architectural modifications. The benefit to online inference is demonstrated by reformulating several widely used network architectures into CINs, including 3D Convolutional Neural Networks, Spatio-temporal Graph Convolutional Networks, and Transformer Encoders. An orthogonal contribution tackles the concurrent adaptation and computational acceleration of a large source model into multiple lightweight derived models. Drawing on fusible adapter networks and structured pruning, Structured Pruning Adapters achieve superior predictive accuracy under aggressive pruning using significantly fewer learned weights compared to fine-tuning with pruning.

In January 2023, the PhD thesis entitled “Uncertainty Estimation for 3D Object Detection and Tracking” was submitted by Illia Oleksiienko to the Graduate School of the Faculty of Technical Sciences at Aarhus University. This thesis has been supported by the OpenDR project and includes improvements of 3D Object Detection and Tracking methods. It first provides an overview of 3D object detection and tracking, uncertainty estimation methods and their applications, then describes the technical contributions to improve the real-world applicability of the state-of-the-art methods by improving the voxel-based 3D object detection inference for embedded devices, proposing a novel method for real-time embedded 3D single object tracking, improving the uncertainty estimation methods and applying the proposed uncertainty estimation methods to 3D single and multiple object tracking, improving the performance of those methods.

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

These contributions have made progress in the fields of 3D object detection, tracking and uncertainty estimation and provided the directions for future research.

In November 2021, the PhD thesis entitled "Representation Learning for Graph-Structured Data Analysis" was submitted by Negar Heidari to the Graduate School of the Faculty of Technical Sciences at Aarhus University. This thesis has been partially supported by the OpenDR project and it includes improvements of Spatio-Temporal Graph Convolutional Networks which were proposed for processing a sequence of graphs through time, i.e. spatio-temporal graphs, to improve the state-of-the-art performance in terms of classification accuracy and computational complexity. These include novel efficient solutions for spatio-temporal graph classification in computer vision applications, i.e. human action recognition and facial expression recognition. These contributions have made progress in the field of graph-based learning and provided intuitions for future research.

In March 2021 Maria Tzelepi was awarded a PhD degree by the Department of Informatics, Aristotle University of Thessaloniki, for her PhD thesis entitled "Deep learning techniques in digital media". This thesis has been partially supported by the OpenDR project for the period January 2020-March 2021. During this period and within the scope of the project Dr Tzelepi conducted research on a number of topics that included semantic scene segmentation, knowledge distillation, quadratic mutual information regularization and lightweight deep learning methods.

In November 2021 Paraskevi Nousi was awarded a PhD degree by the Department of Informatics, Aristotle University of Thessaloniki, for her PhD thesis entitled "Supervised and Unsupervised Deep Learning Methodologies". This thesis has been partially supported by the OpenDR project for the period January 2020-November 2021. During this period and within the scope of the project Dr Nousi conducted research on a number of topics that included active pose estimation, dataset generation for the agriculture use-case and 3D tracking using voxel pseudo images.

In June 2023 Efstratios Kakaletsis finished his research on "Human-Centric Computer Vision Techniques for Autonomous Systems" as part of his PhD studies at the Department of Informatics, Aristotle University of Thessaloniki and obtained permission to author the dissertation. In this research, which has been partially supported by the OpenDR project, Mr Kakaletsis conducted research on active face recognition as well as on using view synthesis towards generating multi-view facial datasets. The thesis is in the final stages of proofreading and will be submitted to the Department of Informatics in the next few weeks.

2.1.5.4 Preprints

8 preprints that describe results obtained within the project were published during this period.

AU published two preprints:

- Amir Ramezani Dooraki, Alexandros Iosifidis, "Curiosity-Driven Reinforcement Learning based Low-Level Flight Control", arXiv:2307.15724, 2023.

Curiosity is one of the main motives in many of the natural creatures with measurable levels of intelligence for exploration and, as a result, more efficient learning. It makes it possible for humans and many animals to explore efficiently by searching for being in states that make them surprised with the goal of learning more about what they do not know. As a result, while being curious, they learn better. In machine learning literature, curiosity is mostly combined with reinforcement learning-based algorithms as an intrinsic reward. This work proposes an algorithm based on the drive of curiosity for autonomous learning to control by generating proper motor speeds from odometry data. The quadcopter controlled by our proposed algorithm can pass through obstacles while controlling the Yaw direction of the quad-copter toward the desired location. To achieve

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

that, we also propose a new curiosity approach based on prediction error. We ran tests using on-policy, off-policy, on-policy plus curiosity, and the proposed algorithm and visualized the effect of curiosity in evolving exploration patterns. Results show the capability of the proposed algorithm to learn optimal policy and maximize reward where other algorithms fail to do so.

- Illia Oleksiienko, Alexandros Iosifidis, "Uncertainty-Aware AB3DMOT by Variational 3D Object Detection", arXiv:2302.05923, 2023.

Autonomous driving needs to rely on high-quality 3D object detection to ensure safe navigation in the world. Uncertainty estimation is an effective tool to provide statistically accurate predictions, while the associated detection uncertainty can be used to implement a more safe navigation protocol or include the user in the loop. In this paper, we propose a Variational Neural Network-based TANet 3D object detector to generate 3D object detections with uncertainty and introduce these detections to an uncertainty-aware AB3DMOT tracker. This is done by applying a linear transformation to the estimated uncertainty matrix, which is subsequently used as a measurement noise for the adopted Kalman filter. We implement two ways to estimate output uncertainty, i.e., internally, by computing the variance of the CNNs outputs and then propagating the uncertainty through the post-processing, and externally, by associating the final predictions of different samples and computing the covariance of each predicted box. In experiments, we show that the external uncertainty estimation leads to better results, outperforming both internal uncertainty estimation and classical tracking approaches. Furthermore, we propose a method to initialize the Variational 3D object detector with a pretrained TANet model, which leads to the best performing models.

Moreover, AU and AUTH published the following preprint:

- Illia Oleksiienko, Paraskevi Nousi, Nikolaos Passalis, Anastasios Tefas, Alexandros Iosifidis, "Variational Voxel Pseudo Image Tracking", arXiv:2302.05914.

Uncertainty estimation is an important task for critical problems, such as robotics and autonomous driving, because it allows creating statistically better perception models and signaling the model's certainty in its predictions to the decision method or a human supervisor. In this paper, we propose a Variational Neural Network-based version of a Voxel Pseudo Image Tracking (VPIT) method for 3D Single Object Tracking. The Variational Feature Generation Network of the proposed Variational VPIT computes features for target and search regions and the corresponding uncertainties, which are later combined using an uncertainty-aware cross-correlation module in one of two ways: by computing similarity between the corresponding uncertainties and adding it to the regular cross-correlation values, or by penalizing the uncertain feature channels to increase influence of the certain features. In experiments, we show that both methods improve tracking performance, while penalization of uncertain features provides the best uncertainty quality.

TUD published the following two preprints:

- B. van der Heijden, J. Luijkx, L. Ferranti, J. Kober and R. Babuska "EAGERx: Graph-Based Framework for Sim2real Robot Learning", technical report, 2023.

Sim2real, that is, the transfer of learned control policies from simulation to real world, is an area of growing interest in robotics due to its potential to efficiently handle complex tasks. The sim2real approach, however, is hampered by discrepancies between simulation and reality, inaccuracies in

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

physical phenomena modeling, and asynchronous control, among others. To this end, we introduce EAGERx, a framework with a unified software pipeline for both real and simulated robot learning. It can support various simulators and aids in integrating state, action and time-scale abstractions to facilitate learning. EAGERx's integrated delay simulation, domain randomization features, and proposed synchronization algorithm contribute to narrowing the sim2real gap. We demonstrate the efficacy of EAGERx in accommodating diverse robotic systems and maintaining consistent simulation behavior. EAGERx is open source and its code is available at https://eagerx.readthedocs.io.

- D.-J. Boonstra, L. Willemet, J. D. Luijkx and M. Wiertlewski, "Learning to estimate incipient slip with tactile sensing to gently grasp objects," under review for the IEEE International Conference on Robotics and Automation, 2023.

To gently grasp objects, robots need to generate enough friction without creating damage by applying the right amount of force. In practice, implementing this force regulation is challenging since it requires knowledge of the friction coefficient, which can vary from object to object and even from grasp to grasp. Fortunately, tactile sensing can provide information about friction notably by detecting the moment when the object slips away from the grasp. These tactile sensors capture distributed information about the deformation of the artificial skin in the normal and tangential direction, from which slippage can be detected. However, current approaches only react to slip, which leads to significant object movement. The movement can in turn induce a failure of the grasp and damage. In this study, we introduce a machine-learning method that anticipates slip by computing the so-called safety margin of grasp. It represents the margin of an additional lateral force that the frictional contact can bear. To find this value, we use a high-density camera-based tactile sensor to measure the 3D deformation of the surface over 82 points. We trained a Convolutional Neural Network (CNN) to obtain a frictional safety margin estimate from the tactile images. The safety margin offers a powerful metric for regulation and therefore a simple proportional controller was enough to robustly grasp a wide collection of objects. The results show that this control method outperforms slip detection methods, by reducing regrasp reaction times while decreasing the maximum applied grasping force to the low-value range of 1-3 N.

During this period TAU released one preprint:

- Kateryna Chumachenko, Alexandros Iosifidis, Moncef Gabbouj, "Improving Unimodal Inference with Multimodal Transformers", arXiv preprint arXiv:2311.10170, 2023.

This paper proposes an approach for improving performance of unimodal models with multimodal training. Our approach involves a multi-branch architecture that incorporates unimodal models with a multimodal transformer-based branch. By co-training these branches, the stronger multimodal branch can transfer its knowledge to the weaker unimodal branches through a multi-task objective, thereby improving the performance of the resulting unimodal models. We evaluate our approach on tasks of dynamic hand gesture recognition based on RGB and Depth, audiovisual emotion recognition based on speech and facial video, and audio-video-text based sentiment analysis. Our approach outperforms the conventionally trained unimodal counterparts. Interestingly, we also observe that optimization of the unimodal branches improves the multimodal branch, compared to a similar multimodal model trained from scratch.

ALU-FR released two preprints:

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

- M. Käppeler, K. Petek, N. Vödisch, W. Burgard, and A. Valada, “Few-Shot Panoptic Segmentation With Foundation Models”, arXiv preprint arXiv:2309.10726, 2023.

Current state-of-the-art methods for panoptic segmentation require an immense amount of annotated training data that is both arduous and expensive to obtain posing a significant challenge for their widespread adoption. Concurrently, recent breakthroughs in visual representation learning have sparked a paradigm shift leading to the advent of large foundation models that can be trained with completely unlabeled images. In this work, we propose to leverage such task-agnostic image features to enable few-shot panoptic segmentation by presenting Segmenting Panoptic Information with Nearly 0 labels (SPINO). In detail, our method combines a DINOv2 backbone with lightweight network heads for semantic segmentation and boundary estimation. We show that our approach, albeit being trained with only ten annotated images, predicts high-quality pseudo-labels that can be used with any existing panoptic segmentation method. Notably, we demonstrate that SPINO achieves competitive results compared to fully supervised baselines while using less than 0.3% of the ground truth labels, paving the way for learning complex visual recognition tasks leveraging foundation models. To illustrate its general applicability, we further deploy SPINO on real-world robotic vision systems for both outdoor and indoor environments

- E. Greve, M. Büchner, N. Vödisch, W. Burgard, and A. Valada, “Collaborative Dynamic 3D Scene Graphs for Automated Driving”, arXiv preprint arXiv:2309.06635, 2023.

Maps have played an indispensable role in enabling safe and automated driving. Although there have been many advances on different fronts ranging from SLAM to semantics, building an actionable hierarchical semantic representation of urban dynamic scenes from multiple agents is still a challenging problem. In this work, we present Collaborative URBan Scene Graphs (CURB-SG) that enable higher-order reasoning and efficient querying for many functions of automated driving. CURB-SG leverages panoptic LiDAR data from multiple agents to build large-scale maps using an effective graph-based collaborative SLAM approach that detects inter-agent loop closures. To semantically decompose the obtained 3D map, we build a lane graph from the paths of ego agents and their panoptic observations of other vehicles. Based on the connectivity of the lane graph, we segregate the environment into intersecting and non-intersecting road areas. Subsequently, we construct a multi-layered scene graph that includes lane information, the position of static landmarks and their assignment to certain map sections, other vehicles observed by the ego agents, and the pose graph from SLAM including 3D panoptic point clouds. We extensively evaluate CURB-SG in urban scenarios using a photorealistic simulator.

2.1.5.5 OpenDR Briefs

During this period, the consortium continued to publish its OpenDR Briefs in the project’s website, aimed primarily at the general public. Briefs are regularly published (bi-monthly) by OpenDR researchers and are short and easy to understand articles regarding the research conducted in the project or topics related to the project. Six (6) new briefs have been published during this period, on topics that include “Domain Shift Evaluation and Alleviation Methods in OpenDR powered by Deep Learning”, “Deep learning for robotics in agile production”, “Continual Inference: A new paradigm for efficient online processing with deep neural networks”, “PARTNR: Pick and place Ambiguity Resolving by Trustworthy iNteractive learning”, “Continual Learning for Robotics”,”Simulated Humanoid Wrestling Competition in webots.cloud”. In total 14 such articles were published.

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

2.1.6 Summer School, Invited / Keynote Talks

Dissemination of OpenDR information and findings through seasonal schools or invited/keynote talks has obviously a high impact, since, usually, the audience consists of students and researchers, interested in the specific topic of the school/talk. One plenary and one invited talk were delivered during the fourth year of the project:

- Prof. Anastasios Tefas (AUTH) delivered a plenary talk on "Deep Learning for Active Robotic Perception" at the "15th International Joint Conference on Computational Intelligence" (IJCCI 2023) which took place in November 2023, in Rome, Italy. The presentation focused on OpenDR achievements.
- In July 2023 OpenDR was present at the "2nd Workshop on "Neuromorphic photonics and applications" which was organized by three EU funded projects (NEoteRIC H2020, PROMETHEUS Horizon Europe, PlasmoniAC H2020) in the premises of the University of West Attica in Athens, Greece. During this event, Prof. Anastasios Tefas, gave an invited talk entitled "Open deep learning tools for neuromorphic computing" through which the OpenDR project was disseminated.

An additional important dissemination event that took place during this period was the Summer School that was organized in May 2023 in Aarhus, Denmark. Details are provided in the subsection below.

2.1.6.1 Summer School on "Deep Learning for Autonomous Systems and Smart Cities"

In May 2023 OpenDR organized the "Deep Learning for Autonomous Systems and Smart Cities" Summer School in collaboration with MARVEL H2020 research project. The School was held in the premises of Aarhus University from the 23rd to the 26th of May and was attended by 36 persons from various countries (Figure 11). OpenDR topics were presented on the first two days.

Talks on various topics relevant to OpenDR's research fields were provided by members of the consortium. Topics included "Human action recognition", "Deep Reinforcement Learning" (in 2 parts), "Deep Learning for 2D Object Detection", "Deep Learning for Object Tracking", "Robotic Grasping for Agile Production", "Continual Learning for Robotics", "Efficient deep learning-based monocular 3D perception for agile robot navigation", "Robotic Simulation Environments", "Deep Learning for Agriculture", "Geometric Deep Learning: GNNs", "Deep Learning for 3D Object Detection/Tracking", "Deep Learning for Semantic Segmentation", and a hands-on session focusing on the OpenDR toolkit. MARVEL project members also provided talks on topics related to OpenDR such as "Audio data for AI applications", "Deep Learning methods for audio AI", "Compact and dynamic models for AI on edge devices", "Federated Learning", "Mondel Compression/DNN Pruning", and "Audio-visual data analysis for UAV-based monitoring", while invited Professors/Researchers from Aarhus University also gave talks on "Enabling Agile Edge Intelligence in Edge-Cloud Continuum" (in 2 parts) and "Deep Learning in multiagent systems", and "Bayesian Deep Learning".

As mentioned above, a session of the School was dedicated to a hands-on introduction to the OpenDR toolkit. During this session, the AUTH team showcased the use of different tools of the OpenDR toolkit and the attendees had the opportunity to follow programming examples illustrating the functionalities of the OpenDR tools.

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

Figure 11: Photos from the Summer School's presentations.

To support the School dissemination and to allow for a smooth registration of the participants, a website was created[6] (Figure 12) and a webpage was created in the website of OpenDR (Figure 13). Moreover, information about the OpenDR project was disseminated to the participants of the School in the form of flyers and factsheets (Figure 14).

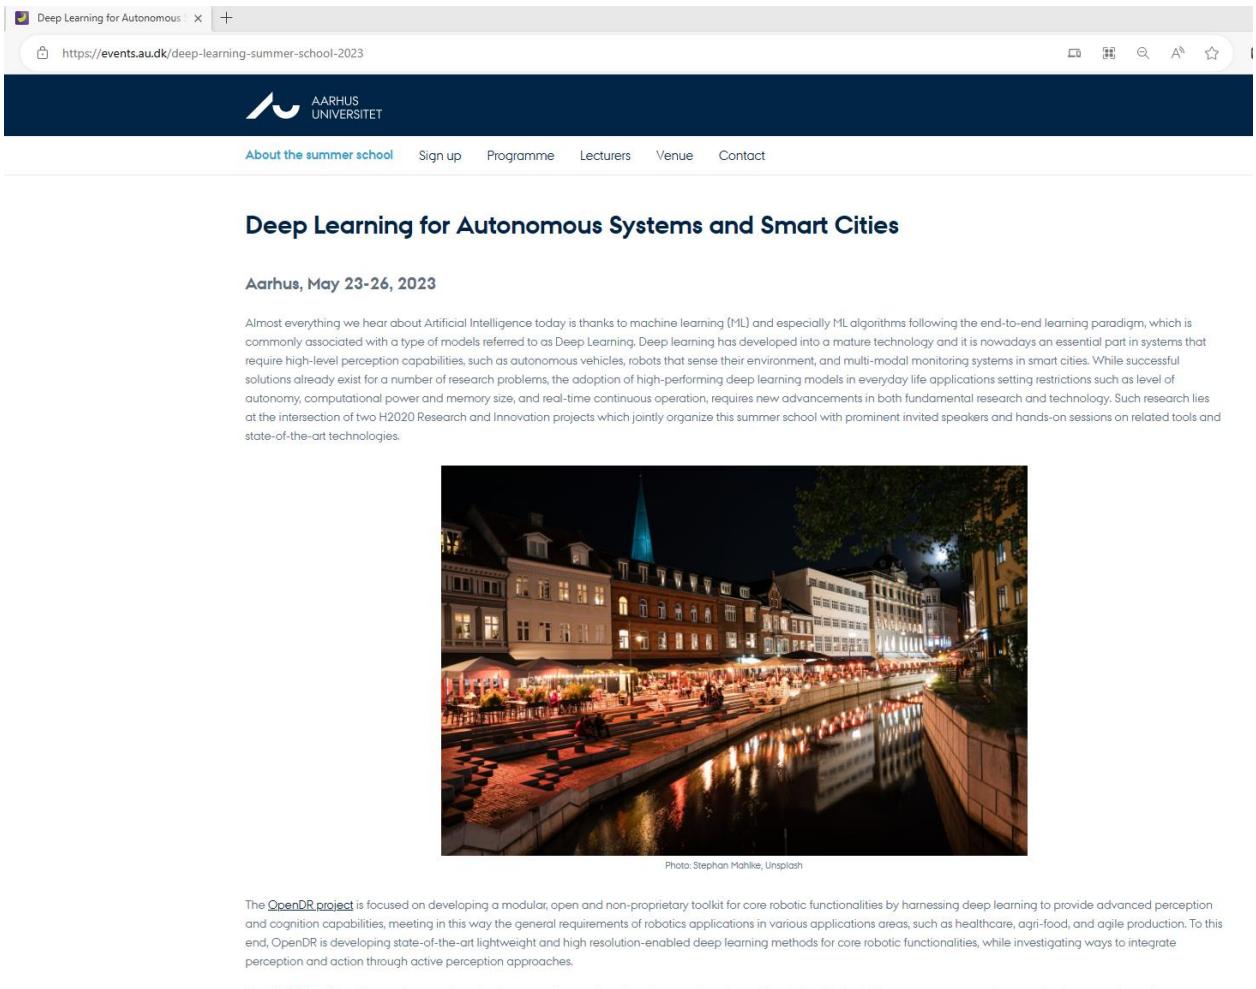

Figure 12: A screenshot of the Summer School website hosted by Aarhus University.

[6] https://events.au.dk/deep-learning-summer-school-2023

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

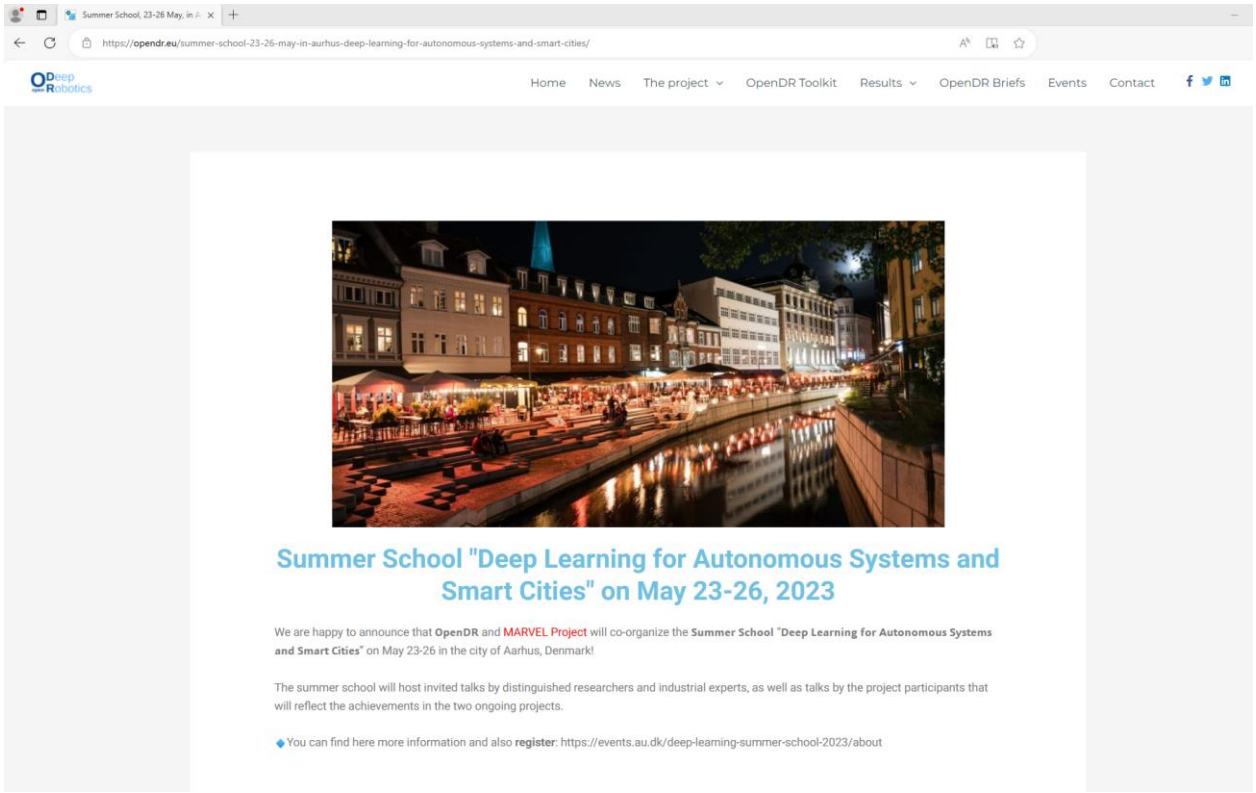

Figure 13: A screenshot of the Summer School webpage in the website of OpenDR.

Figure 14: Flyers and information material about OpenDR distributed to the participants of the Summer School.

2.1.7 Participation to Tradeshows, Exhibitions, EU Events, Industry Workshops

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

Events such as trade shows, exhibitions, events organized by the EU, etc., are obvious venues for the dissemination of OpenDR aims and results. During this period, the project was present at the following fora:

- PAL Robotics was present this year at the European Robotics Forum (ERF2023) in Odense, Denmark (14-16 March) and disseminated OpenDR objectives and results in a Workshop entitled “How can AI and Robotics be used in the care context? The sustainability of care for older adults” and organized by PAL Robotics and Poznan University of Medical Sciences. Within this workshop, Dr. Paraskevi Nousi from AUTH also disseminated the OpenDR project through the presentation "Open Deep Learning Methods for Robotics in Healthcare" (Figure 15).

Figure 15: ERF23

- PAL also had the chance to talk about the project in a workshop about “TIAGo and ROS2” organized in ROSCon Madrid 2023 (29 September), Spain (Figure 16)

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

Figure 16: Noel Jimenez Garcia (PAL)'s presentation at ROSCon Madrid

- PAL disseminated the project in several additional events such as EuroShop 2023, Dusseldorf, Germany (26 Feb - 2 March); Advanced Factories Expo & Congress 2023, Barcelona, Spain (18-20 April); ICRA23, London, UK (29 May- 2 Jun); Automatica 2023, Munich, Germany (27-30 June); IROS23, Detroit, USA (1-5 October) ESN Annual Seminar 2023, Barcelona, Spain (9-10 October).
- Furthermore, the CTO of PAL Robotics had to chance to discuss the project results at the ROS Community meeting in ROSCon23, New Orleans, USA (18-20 October) (Figure 17)

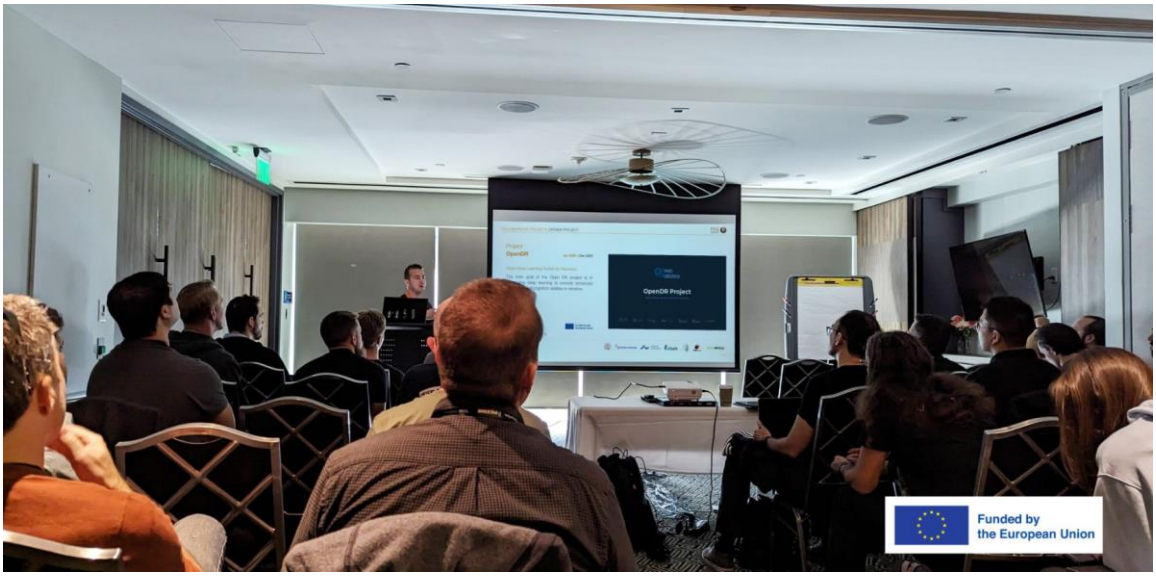

Figure 17: Luca Lash (CTO in PAL) presentation in ROSCon23

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

2.1.8 Collaboration with Other Projects

During the previous period, TAU has obtained, through TRINITY DIH, a list of projects and DIHs that were related to Open DR topics. In October 2023 AUTH went through this list to identify the 25 most relevant and still active projects and their contact points. To maximize the project's recognition as well as the possibility of collaboration, emails regarding the OpenDR objectives and achievements (with focus in the toolkit) as well as pointers to the project's social media channels and the website were sent to those projects.

Moreover, in November 2023 the EU project MOZART[7] contacted OpenDR towards clustering/joint dissemination activities such as including OpenDR in the list of MOZART relevant projects, exchanging news items for posting them in the two projects social media or MOZART newsletter, etc. The OpenDR consortium happily agreed, and the cooperation started.

Also, OpenDR collaborated with the H2020 MARVEL project where AU is participating. The two projects co-organized the Summer School on "Deep Learning for Autonomous Systems and Smart Cities" in Aarhus Denmark (see Section 2.1.6.1).

In addition, CYB has been working on several projects for industrial customers which partially benefited from some achievements of the OpenDR project, including Webots, webots.cloud and deep learning interfaces. These included the following:

- KHG GmbH: development of a simulation of remote-controlled robots that should intervene in case of nuclear disaster. The simulation is used for training the pilots of these robots.
- Code Institute: development of online robot simulations on the web used for an online robotics curriculum aiming at teaching robotics to high school students in the UK
- Harvard University: development of a simulation of a wearable robot and deployment of this simulation in the AWS cloud to run deep learning algorithms.
- RoboCup: development of a humanoid robot soccer simulation to be used by the RoboCup organization for setting-up a realistic simulation of the humanoid league.
- OpenRoberta: development of an online simulation of NAO robots for the OpenRoberta online robot graphics programming tools based on Blockly.

Finally, for PAL the OpenDR project opened various possibilities of collaboration with projects it participates in, like SPRING and ALMI, where tools from the OpenDR toolkit have been or will be integrated in healthcare-related use cases of these projects.

2.1.9 Organization of Contests and Challenges

During this period CYB organized two simulated humanoid robot wrestling programming contests respectively for the ICRA 2023[8] and IROS 2023[9] robotics conferences. These programming contests were based on the Webots and webots.cloud technologies developed in the OpenDR project. CYB also provided competitors with the ROS 2 Vulcanexus distribution that includes potentially all the ROS 2 ecosystem to which OpenDR has been contributing. Everything was running automatically in the cloud (GitHub action), launching games on a worker machine equipped with an GPU for accelerated machine learning. Each competition gathered more than 32

[7] https://mozart-robotics.eu/resources/

[8] https://github.com/cyberbotics/wrestling/blob/main/result-icra2023.md

[9] https://github.com/cyberbotics/wrestling/blob/main/result-iros2023.md

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

teams of compe0titors worldwide. An article[10] published at opensource.com introduces the ICRA 2023 competition.

2.2 Dissemination Performance Metrics

2.2.1 Project Website

During this fourth year a significant number of new posts and documents were created and uploaded on the project website. Through these posts, all the new journal and conference papers, announcements about events as well as participation in workshops and conferences were made known to the public. The project achieved, by a large margin, all the goals which were defined in the DoA (with the exception of the average duration of visits metric) and also made significant progress, in relation to the previous period (M24-36).

KPI	Target (for M48, unless stated otherwise)	M12	M24	M36	M48
Number of visitors	1000 (by M12)	More than 3200	More than 5600	More than 10800	More than 14900
Average duration of visits:	2:00	1:36	1:46	1:15	1: 22
Number of posts published (including news):	>50	15,	56	66	79
Number of documents available in the website repository:	>50	18	41	101	134

2.2.2 Social Media

As already stated, OpenDR tried to increase awareness and disseminate all the fourth year's publications and news by using the most popular social media, such as LinkedIn, Twitter and Facebook, while the Youtube channel was used for publishing videos that were created in the frameworks of conference presentations and the project's dissemination activities in general. The social media accounts have been used to a great extent not only during the period M37-M48, but also throughout all the four-year project's lifecycle, and, as a result, all target values of the key performance indicators which were set in DoA have been accomplished by a large margin.

[10] https://opensource.com/article/23/2/ai-robot-wrestling-open-source

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

2.2.2.1 Twitter

KPI	Target (M48)	M12	M24	M36	M48
Twitter followers	>100	151	230	424	610
Number of tweets/retweets	>150	89	151	242	285

2.2.2.2 Linkedin

The target of exceeding 150 group members was achieved during the previous period. Currently, the LinkedIn group has 172 members and counts 245 posts. As far as the LinkedIn official page is concerned (no target values were set in the DoA), it has gathered until now more than 1030 followers, a fact which, in combination with the number of the Linkedin group members, shows the very satisfactory presence of OpenDR in this platform during these four years.

KPI	Target (M48)	M12	M24	M36	M48
LinkedIn group members	>150	111	132	158	172

2.2.2.3 Publications

Like all the previous years, the number of publications during the period M37-M48 exceeded the target which was defined in the DoA. Indeed, while the goal that had been set was 8 publications per year, the OpenDR consortium produced during the last year 8 journal and 18 conference papers. The number of publications that the project generated throughout all the four years of its life, reached 100, i.e., 25 per year, which can be considered as impressive.

KPI	Target	M12	M24	M36	M48
Number of publications	>8 per year	24 (16 Conferences, 7 Journals, 1 Book)	50 (14 new Conferences, 12 new Journals)	74 (18 new Conferences, 6 new Journals)	100 (18 new Conferences, 8 new journals)

2.2.2.4 GitHub

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

Following the first release of the toolkit on M24, the OpenDR consortium released five more public versions of the toolkit until M48 (six public versions in total), through GitHub. The response of the public and especially of the scientific and robotics community to the OpenDR toolkit was extremely positive, taking into consideration the numbers that have been achieved regarding the KPIs which were stated to the DoA. Indeed, until M48 the toolkit gathered 567 stars, more than 17000 downloads and 86 forks in the GitHub platform.

KPI	Target (M48)	M36	M48
Number of stars	>100	301	567
Number of Downloads	>500	>7500	>17000
Number of forks	>30	54	86

2.2.2.5 Dissemination Material

Throughout the period M37-M48 OpenDR consortium tried to increase as much as possible the use of dissemination material in order to maximize the project's awareness and impact to the public. As a result, almost all KPI targets which were stated in the DoA have been achieved. However, despite the fact that OpenDR consortium tried to do its best, COVID19 pandemic prevented the project from achieving the KPIs regarding the number of events where the poster and the roll up banner were used, due to the fact that throughout the period M3-M24 no physical events could be organized by the OpenDR project. However this can be considered a minor pitfall.

KPI	Target (M48)	M12	M24	M36	M48
Number of e-newsletters published	4	1	2	3	4
Size of the dissemination list	>500	>1000 (estimate)	>3000 (estimate)	>5000 (estimate)	>8000 (estimate)
Number of press releases	3	1	1	1	3
Number of videos to be produced	2	1	1	1	2
Number of video views	500	529	911	1526	2800
Number of events where the	10	1	1	4	11

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

factsheet is used					
Number of events where the poster is used	10	1	2	6	8
Number of brochures	2	1	1	2	2
Number of events where the brochure is used	10	-	-	4	11
Number of copies distributed (brochures)	800 hardcopies + eBrochures	-	-	550	850
Number of events where a roll-up banner is used	4	0	0	1	2

3. Final Exploitation Plan

The exploitation plans of the outcomes of the OpenDR project differ depending on the type of partner. Industrial partners are generally re-using the technologies, tools, datasets, documentation and know-how developed in OpenDR to apply them to business projects they have with their customers and partners. Academic partners are mainly exploiting the outcomes of the project to publish high quality research papers in prestigious journals and conferences. Indeed, the technical and scientific achievements of the OpenDR project allow more efficient research progress. The tools developed in OpenDR are also exploited by academic partners for educational purposes. This section details the exploitation plans for each partner of the project.

3.1 CYB Exploitation Plan

CYB plans to exploit the outcomes of the OpenDR project by continuing to use the technologies developed in OpenDR for industrial and academic applications in a similar way to what was achieved through various collaborations during the project (see section 2.1.8). This includes in particular the Webots open-source software simulator for which industrial and academic partners of CYB need some adaptations, improvements, interface with other software, simulation scenarios, simulation models, training, documentation, etc. This also includes the sales of cloud computing services for robot simulations based on the webots.cloud technology developed throughout the OpenDR project.

3.2 AGI Exploitation Plan

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

AGI has already exploited the Crop & Weed tool by providing crop and weed heat maps to those who have Robotti with CropEye. Currently, the heat mapping service is free. AGI plans to enhance their remote supervision webpage by adding the human detection tool to the live streaming front from the robot's front and rear cameras and further integrating the human detection tool as a helping safety function. AGI will continue to develop the plant row guidance system.

3.3 PAL Exploitation Plan

In healthcare scenarios PAL plans to use the various tools provided by the OpenDR toolkit to create applications that are better aware of the environment and are able to interact in a seamless manner with the patient, doctor, nurses and persons visiting the hospital or care houses.

For example, following the public release of the OpenDR toolkit in December 2022, PAL successfully integrated the OpenDR fall detection ROS node into another company's social robot, ARI. These results are being exploited in the EU project SPRING, where ARI is utilized for the elderly use case in a real hospital in Paris, France. The OpenDR toolkit is also applied in the ALMI project, showcasing fall detection in healthcare scenarios. Furthermore, OpenDR tools are being used in SESAME for human and object detection and they have adapted the Yolo OpenDR tool to be able to detect low light and full light in the use case scenario.

Continuing PAL´s collaboration with healthcare providers not only validates and refines project solutions but also plays an important role in advancing the TRL for product finalization, crucial for broader adoption and market penetration. As part of PAL's ongoing efforts in the realm of assistive robots, the company has focused on streamlining the commercialization of robotic technology for enhanced usability in unpredictable healthcare environments.

To pave the way to a deployable robot in the healthcare sector, Tiago Pro, the latest iteration of the Tiago robot, was created with new arms that integrate force sensing, torque sensing and brakes in each joint to improve the safety of the robot around humans. Tools from the OpenDR toolkit are planned to be incorporated in this new platform, elevating its performance and bringing it closer to meeting the demands of robotics in healthcare.

In general, PAL aims to revolutionize care environments by offering enhanced social interaction, continuous monitoring, and cognitive engagement support for end-users. This long-term vision shows the company's commitment to making a meaningful impact on the quality of the provided care. This acquired expertise not only contributes to advancements in the social robotics sector but also strengthens the company foothold in the Assistive Robotics industry.

3.4 AUTH Exploitation Plan

Being a Higher Education Institution, AUTH exploits, promotes and disseminates research results mainly by publishing in well-known and peer reviewed international scientific journals and established international scientific conferences, while providing open access to these publications using mainly the "green " open-access model. In this sense, AUTH "academic exploitation" has been very successful, having published within the project duration 10 journal and 35 conference papers, some of them in cooperation with other partners. These activities will continue after the end of the project, since there are papers currently in the making and also because a number of research activities that started within the project will continue after its conclusion. Contacts with other related research projects, research/academic institutions and some industrial players have

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

also been established, in an effort to highlight the importance of the achieved results and promote their use.

Moreover, the developed tools and authored publications have already been exploited for improving education quality at undergraduate and postgraduate level in related state-of-the-art topics. Indeed such material was used for teaching purposes, whereas numerous student projects in courses on neural networks and machine / deep learning were based on or utilized tools from the OpenDR Toolkit. In addition, a significant number of pre-graduate and MSc Theses dealt with OpenDR topics and produced interesting and, in many cases, publishable results. Finally, 8 PhD students worked, as part of their theses, on OpenDR topics and were partially funded by it. 3 of these theses were completed within the project duration. Obviously, training young researchers on cutting-edge technologies is of utmost importance for improving the competitiveness of Europe in the scientific and technological arena.

OpenDR has also motivated the AUTH team that worked on the project to partially shift its research directions towards computer vision / deep learning / perception /decision-making for robotics. It is obvious that the accumulated know-how in this area will influence the research directions and topics of interest of its members long after the end of the project. AUTH will attempt to use OpenDR R&D outcomes to address other related areas and markets, such as media analysis, fintech, game development etc.

Finally, another important exploitation avenue for AUTH is to use OpenDR derived know-how to strengthen its position as a key player in the European/international scientific arena. Such know-how will help AUTH in securing positions in consortia competing for EU R&D projects. Thus, AUTH will participate in several Horizon Europe R&D proposals on related topics, thanks to its involvement in OpenDR.

3.5 AU Exploitation Plan

In a similar manner, as a Higher Education Institution, AU published 7 articles in high-quality international journals and 11 papers in top-quality conferences, while it has also conducted research reported in 7 preprints which are considered for publication in conferences and journals. AU's "academic exploitation" has been very successful, and it will continue further work on the topics of OpenDR after the end of the project. AU established connections with other projects (both funded by EU and National agencies) in which the team is participating, that will exploit the usage of the tools developed in OpenDR for further methodology development and illustration of their capabilities.

Topics of OpenDR have been included as group projects in the Computer Vision course at the MSc study program of the Electrical and Computer Engineering at Aarhus University, coordinated by Prof. Iosifidis. Finally, 4 PhD students worked, as part of their theses, on OpenDR topics and were partially funded by it.

3.6 ALU-FR Exploitation Plan

Exploitation and dissemination are also viewed in a similar fashion by ALU-FR. ALU-FR published 8 articles in renowned international journals as well as 6 papers in top-tier international conferences in the fields of robotics, machine learning, and computer vision. The methods and tools developed during the duration of the project have already served as a foundation for new endeavors and will continue to do so after the end of the project. The project resulted in a strengthened collaboration among the project partners resulting in joint publications independent of the project as well as new joint project proposals. With regards to teaching and education,

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

OpenDR had a major impact. 4 PhD students were partially funded by OpenDR, making the conducted research work part of their theses. Furthermore, several MSc theses were written on project-related topics, in a few cases leading to published articles. We expect many coming topics for master theses and projects to build upon the tools and methods developed within the OpenDR project.

3.7 TUD Exploitation Plan

Exploitation and dissemination are viewed in a similar fashion to the other academic partners. TUD published 9 articles in renowned international journals and in top-tier international conferences in the fields of robotics, machine learning, and computer vision. The tools developed in the project will be the foundation to develop new technologies and methods within the robotics and machine learning community. EAGERx was presented in a tutorial session organized at one of the top robotics conferences (i.e., the 2022 International Conference on Robotics and Automation) and invited for presentation at the 2022 IEEE Control Technology and Applications (CCTA) Conference. With regards to teaching and education, OpenDR had a major impact. A workshop for MSc and PhD students was organized to introduce them to EAGERx. In addition, 2 PhD students were partially funded by OpenDR, making the conducted research work part of their theses. Furthermore, several MSc theses were written on project-related topics, in two cases leading to published articles.

3.8 TAU Exploitation Plan

In TAU, exploitation and dissemination are viewed in a similar manner to the other academic partners. TAU published 9 articles in high-quality international journals and 16 papers in top-quality conferences, while it has also conducted research reported in 2 preprints which are considered for publication in conferences and journals. TAU's "academic exploitation" has been very successful, and it will continue further work on the topics of OpenDR after the end of the project. TAU established connections with other projects (both funded by EU and National agencies) in which the team is participating, that will exploit the usage of the tools developed in OpenDR for further methodology development and illustration of their capabilities. Topics of OpenDR have been included in BSc and MSc courses, at the Computing Sciences unit and the Automation Technology and Mechanical Engineering unit at Tampere University, and as part of MSc theses. Finally, 4 PhD students worked, as part of their theses, on OpenDR topics and were partially funded by it.

4. Protection and Management of IPR

The Consortium Agreement (CA) of OpenDR provides the general framework for IPR protection and management. The CA states that any result generated before the effective date of the CA (i.e. background) shall remain with the respective party bringing such background to the project. Any result generated by a party after said date, during and within the scope of the project (i.e. Result) whether or not it qualifies for Intellectual Property Right (IPR) protection, shall vest

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

in the party that generated such Results. Partners are free to use the IP they developed within the scope of the project for applications outside the scope of the project without the need of consent from the other partners.

The CA of OpenDR also describes the IP brought into the project by the participants at the start of the project. The IP remains property of the party that brought it into the project, while project partners will obtain non-exclusive access to the IP for the purpose and duration of the project.

Early in the OpenDR lifecycle, it was decided that the project will release the OpenDR toolkit under the open-source Apache 2.0 license. With this permissive non-contaminating open-source software license, the tools could be used and sold by the industrial partners as part of their commercial products. Indeed, the license allows customers of industrial partners to reuse the licensed software within their own proprietary software without having to disclose their proprietary software. It also guarantees that the distributed software license will remain open-source forever. This is very much appreciated by industrial customers who don't want to be dependent on the possibly changing license conditions, including pricing, sub-licensing, maintenance, etc. When they decide to use some software dependencies for their own development, they need to be sure they can rely on these dependencies in the long run. However, Apache 2.0 license is not copyleft, and thus the developers and researchers will retain the ownership of the code they develop, and, if the code is redistributed or modified, the original author as well as the modifications must be explicitly mentioned.

In short, this Intellectual Property (IP) management approach is flexible enough to protect the copyright and allow commercial exploitations. Additionally, the public release of the toolkit guarantees a high visibility of the OpenDR project results, the research partners' works, and the industrial partners' robotics platforms. Then, the post-project commercial exploitation will also benefit from the open-source format since it would render it possible to continue receiving contributions from OpenDR partners or external contributors that will maintain and extend the library.

In accordance with the above, CYB decided to release all their software contributions under the Apache 2.0 license.

Regarding AGI, it has released a partial dataset for the Crop & Weed tool. However, the larger dataset will be kept confidential and unpublished. All other tools and datasets used in the agricultural use case will remain open-source.

Regarding PAL, its background as indicated in CA is the following:

- All the hardware technologies related to actuation systems, electronics boards, cable harnesses, sensors, network devices and onboard computers that are integrated in PAL Robotics' robots constitute background information and expertise of PAL Robotics
- All the relevant low level, middleware and high-level software modules and the Hardware Abstraction Control layer (ROS Control) developed for PAL Robotics' robots are prior knowledge of PAL Robotics.
- The Kinematics and Dynamics description of PAL Robotics' robots as well as its simulation environment in Gazebo are background knowledge of PAL robotics.
- The software modules like the autonomous navigation, motion planning, face detection/tracking, fiducial markers recognition, text-to-speech, are prior knowledge of PAL Robotics, developed for its large family of humanoid and service robots.

All other tools and datasets used in the healthcare use case will remain open-source.

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

It shall be also mentioned that the project has created a significant number of datasets (or software tools that can be used to create datasets). These are also publicly available and can be found in the respective GitHub page[11] and in the OpenDR community in Zenodo[12]

5. Fourth Project Newsletter

In accordance with the DoA, the consortium has created, under the coordination of AUTH, the fourth project Newsletter. The document, which provides a brief introduction to the project, its consortium and illustrates its major results during the fourth year (including the toolkit) can be found in Appendix I. The newsletter will be posted at the project website, and announcements will be posted on the project social channels. The newsletter will also be distributed through relevant email lists such as euRobotics and CVML and email lists maintained by the partners.

6 Overview and Conclusions

By considering OpenDR activities and achievements within the four years of its duration, one can conclude that the project has been very successful in terms of dissemination and exploitation. Indeed, the project has to present the following achievements:

- Publication of 66 conference papers in high caliber conferences, 33 papers in established scientific journals, numerous preprints (all available in the website) and 1 edited book.
- Organization of 2 summer schools, 1 training workshop, 1 conference workshop and 2 conference tutorials. Participation in 2 conference tutorials.
- Organization of 2 robot programming competitions (ICRA 2023 & IROS 2023).
- Organization of 2 special sessions in conferences and 1 special issue in a journal.
- Delivery of a significant number of plenary/invited talks by its members.
- Participation in a wealth of scientific and industrial events: ERF, IROS, ICRA, etc. to promote the project.
- Continuous update of the project website and intense activity in social media
    - More than 200 documents and posts on the website, ~15000 visitors
    - More than 280 tweets and retweets, more than 600 followers in Twitter/X
    - More than 1000 followers in the LinkedIn official page
    - 18 YouTube videos (2 promotional), more than 5400 views
- Dissemination of the OpenDR toolkit releases that contributed to its success (more than 560 stars, more than 17000 downloads)
- Creation of 4 newsletters, press releases, brochures, factsheets, posters etc.
- Collaboration with DIH TRINITY, and other projects (MARVEL, FOCETA, FELICE)
- Establishment of exploitation plans and directions.

[11] https://github.com/opendr-eu/datasets
[12] https://zenodo.org/communities/opendr

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

Through all these activities, the project managed to achieve and, in most cases, significantly surpass essentially all the targets (in terms of KPI) that it has set.

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

Appendix I

4th Newsletter

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

Welcome to the final OpenDR Project Newsletter

Newsletter 4.0: Final release of the OpenDR toolkit and other project achievements

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.

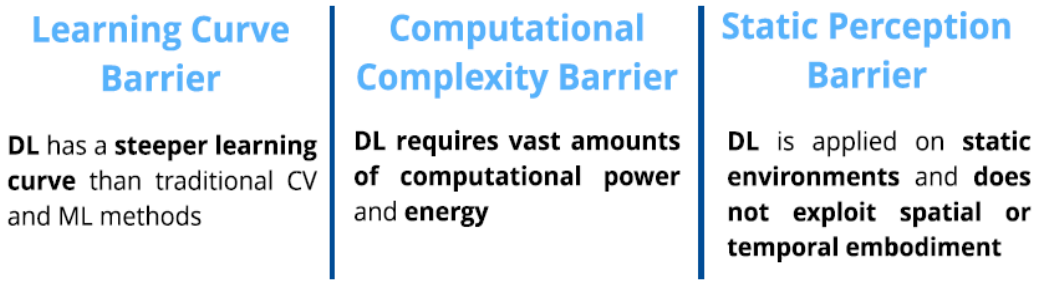

The need for an open deep learning toolkit that contains easy to train and deploy real-time, lightweight, Robot Operating System (ROS) compliant deep learning models for robotics is evident. This is where the **OpenDR** project enters.

What is OpenDR?

OpenDR "Open Deep Learning for Robotics Toolkit", is a EU 2020 Project which was launched on January 2020 and aimed 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 was coordinated by the Aristotle University of Thessaloniki, Greece and ran through December 2023 with a total budget of 6.6 Million Euros.

OpenDR aimed to 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 aimed 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 was 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 coordinated the project and led the organization of dissemination activities. AUTH focused its research on deep human centric active perception and cognition, where it contributed 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 also led the research in object detection/recognition and semantic scene segmentation and contributed 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 participated 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 led the research in deep human centric active perception and cognition, working mainly on deep speech and biosignals analysis and recognition, and contributed 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 also contributed 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 led the research in deep environment active perception and cognition. ALU-FR focused its research on Deep SLAM and 3D scene reconstruction, as well as on deep navigation. It also contributed on developing methodologies for deep planning.

Aarhus University (AU), Denmark participated 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 led work on 2D/3D Object localization and tracking and worked on sensor information fusion, as well as object detection/recognition and semantic scene segmentation and understanding. AU also contributed 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 led/organized the research activities on deep action and control, deep planning, as well as deep navigation. Furthermore, TUD led and undertook the research activities on human robot interaction. Finally, it led and organized 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 led efforts of defining the toolkit's requirements and specifications. CYB also worked on developing simulation environments and collecting data. Finally, it also led 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 organized and coordinated the toolkit integration, as well as the use cases integration activities. PAL also contributed on defining the healthcare robotics use case requirements and specifications and worked on the integration of OpenDR Toolkit to this use case, as well as on its evaluation.

Agro Intelligence APS (AGI), Denmark organized and coordinated the toolkit evaluation, as well as the use cases specific toolkit evaluation activities. AGI also contributed on defining the agri-food use case requirements and specifications and worked on the integration and evaluation of OpenDR Toolkit in this specific use case.

Work Performed in the 4th Year

Toolkit Integration & Evaluation

Following the first release of the toolkit on M24, the OpenDR consortium released five more public versions of the toolkit until M48 (six public versions in total), through its GitHub repository. The latest versions focused on expanding the number of available methods, as well as providing performance improvements to existing ones. In addition, several other improvements were made, including the addition of modular installation options, the support for newer CUDA versions, the implementation of a refined ROS/ROS2 interface, etc. The final version of the OpenDR toolkit was released on December 2023, includes several new tools and provides new functionalities, such as gesture recognition, FSeq2 non-maximum suppression, continual SLAM, intent recognition, object detection class filtering, object detection models for agricultural use cases, RL-based active perception, adaptive high-resolution pose estimation, and others. Furthermore, several existing tools have been enhanced by fixing bugs and including additional demos and ROS nodes. Also, the installation process is now more robust and efficient, while the development pipelines have been significantly improved allowing for more easily

developing new tools. The OpenDR team will continue supporting the tool in the years to come. We do believe that this major outcome of the project has fulfilled its goal of providing easy to train and deploy, real-time, lightweight, Robot Operating System (ROS) compliant deep learning models for robotics. Indeed, reception of the toolkit from the robotics / deep learning /computer vision community is already very encouraging: so far, the GitHub repository was awarded more than 550 stars, was forked 86 times and the toolkit (as a whole or individual tools) has been downloaded more than 17000 times since its first release in December 2021.

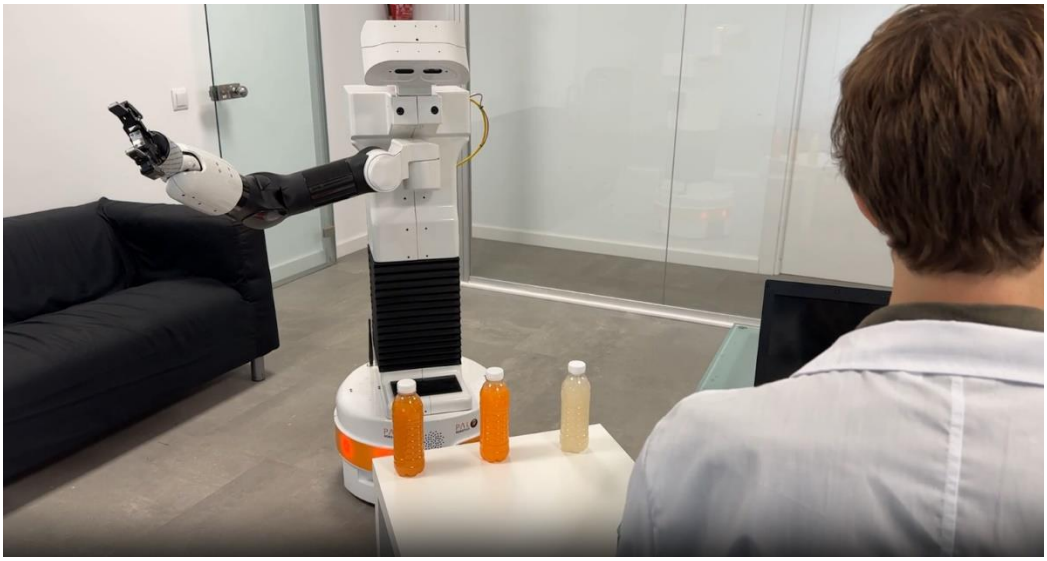

Figure 1: Human robot interaction in the healthcare scenario.

Moreover, the consortium successfully concluded the integration of tools in the three use cases. In the Healthcare use case, specifications were refined to enhance human-robot interaction and multi-modal information fusion, accompanied by a more intricate storyline. In addition, PAL seamlessly incorporated new tools into the healthcare scenario, achieving milestones like dynamic face recognition and improved speech recognition. The integration fostered a seamless interaction between natural speech recognition and PAL's internal chatbot which has resulted in a user-friendly interface for human-robot interaction with the TIAGo robot.

Furthermore, AGI has integrated and tested several tools, meeting the requirements for the Agriculture use case. The Crop & Weed tool was integrated into the OpenDR toolkit and into AGI’s Robotti and so did tools for detecting people and tractors and the plant row guidance system. Finally, the Agile Production use case saw updates integrating multimodal tasks, refining simulation-to-real requirements, and introducing natural language recognition.

In terms of evaluation, the consortium rigorously tested and evaluated in diverse environments the performance of a multitude of tools developed in the research work packages. At the same time, tools used in the three use cases were evaluated to check whether they fulfil the respective requirements.

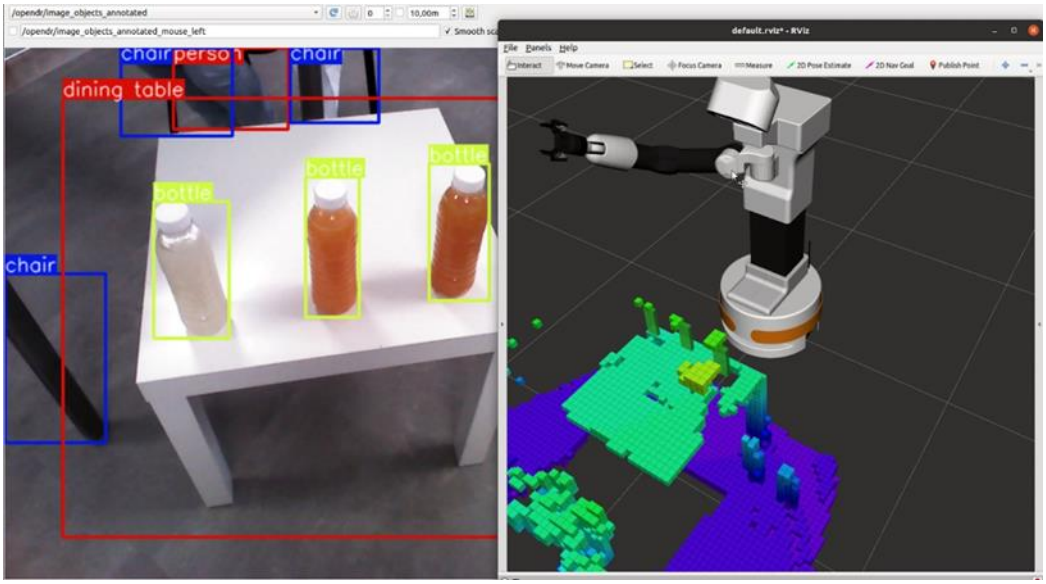

Figure 2: Yolov5 OpenDR 2D object detector (left) used in PAL Advanced grasping pipeline.

As a result, feedback and crucial insights were provided to the partners that developed them. Overall, the integration and testing of the tools, both in the toolkit, and in the three use cases was very successful.

Deep Human Centric Active Perception and Cognition

Work on human-centric active perception and cognition during this last period led to significant achievements. For example, AU worked on incorporating uncertainty estimation in skeleton-based human action recognition by proposing variational

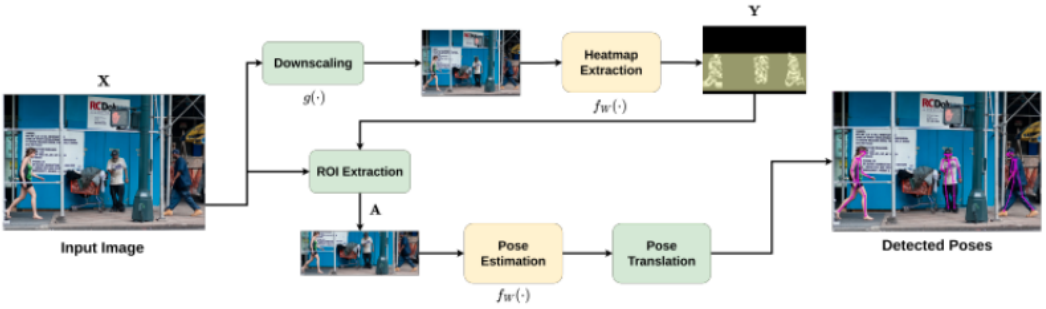

Figure 3: The active perception for high resolution pose estimation approach by AUTH

versions of Spatio-Temporal Graph Convolutional Networks and Adaptive Graph Convolutional Network models.

The proposed variational perception methods improve the quality of human action recognition and estimate model uncertainties. AUTH worked towards further extending a high resolution pose estimation approach developed in the project in order to more efficiently handle cases where multiple humans appear. Furthermore, AUTH further developed an embedding-based active perception approach for face recognition by leveraging a new dataset it has generated. The improved method allows for multi-axes control of the robot.

Finally, TAU developed a method to allow for better performance in a variety of unimodal deep learning tasks (such as hand gesture recognition, audiovisual emotion recognition, language sentiment analysis, etc.) by training the models in a multimodal fashion.

Deep Environment Active Perception and Cognition

In the final year of the project, partners AU, AUTH, ALU-FR, TAU, and AGI proposed various deep learning-based methods for environment active perception and cognition. In detail, AU introduced multiple methods for estimating and exploiting uncertainty in 3D perception tasks, e.g., by utilizing variational neural networks.

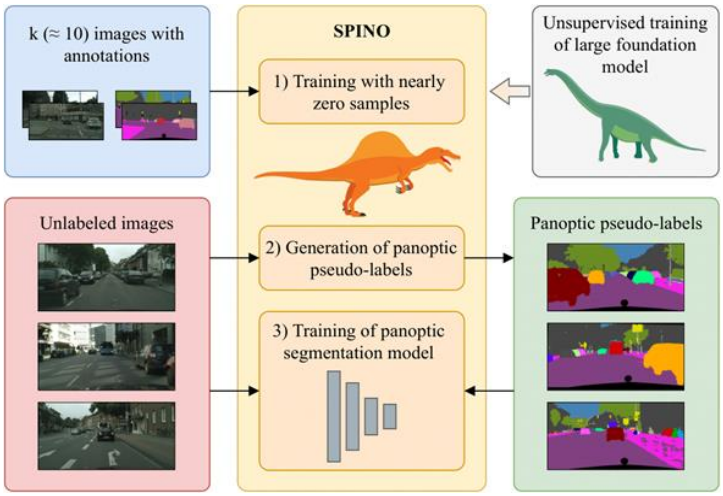

Figure 4: SPINO is a novel method for vision-based panoptic segmentation that enables few-shot learning by exploiting rich image features from an unsupervised foundation model.

Their work significantly improved the tracking of objects. Furthermore, AUTH developed lightweight object detection algorithms and models for agricultural applications while focusing on high-resolution processing. Moreover, AUTH continued working on methods to tackle one-hot encoding shortcomings by introducing a framework for learning soft label embeddings. Partner

ALU-FR proposed several methods for applying online continual learning to visual odometry, monocular depth estimation, and panoptic segmentation. These approaches enable seamless domain adaptation and are self-improving during the test time. ALU-FR further developed a novel 3D scene graph for urban driving that is collectively assembled by multiple agents and provides a high-level structure of large environments.

Additionally, ALU-FR proposed a novel approach for few-shot panoptic segmentation that requires as few as ten annotated images for training. Their method, SPINO, leverages rich image descriptions from a visual foundation model for label-efficient training. Next, TAU prepared benchmark evaluations for a multi-modal feature fusion framework that addresses recent data collection issues in both real and simulated domains. Finally, AGI developed and evaluated a new method for the mapping of agricultural fields and utilized this map for the autonomous navigation of a robot. Using in-house tools, the crop rows are automatically found along with the cross-track error of a row.

Deep Robot Action and Decision Making

In the field of deep robot action and decision making, partners AU, ALU-FR, TUD and AUTH continued and finalized their work towards the design of novel, state of the art navigation, planning, and control algorithms. AU presented a novel Deep Reinforcement Learning (DRL) training strategy for addressing robot navigation problems by leveraging the principles of the Lyapunov theory and a novel approach.

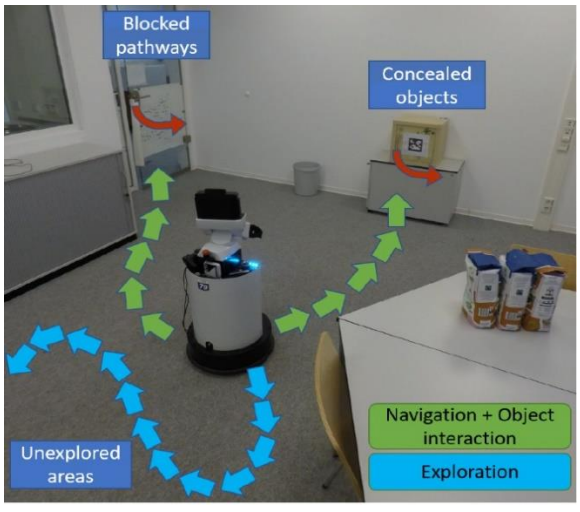

Figure 5: TUD Evaluation of PARTNR in a table-top pick and place task

for learning the low-level flight control of UAV robots. ALU-FR introduced an approach that learns to coordinate and combine previous OpenDR works on exploration and mobile manipulation to autonomously solve a novel interactive multi-object search task. Moreover, AUTH developed a DRL-based end-to-end trainable agent for differential-drive wheeled robot navigation, while also developing the appropriate techniques to improve learning efficiency. Furthermore, AUTH finalized

their work on a data-efficient DRL approach for robust inertial-based UAV localization. TUD focused on mitigating the sim2real gap and enhancing sample efficiency in off-policy reinforcement learning. In addition, TUD worked on simplifying and making safer the collaboration between humans and robots in an industrial context, whereas TAU's main objective was to utilize perception for human-robot collaboration, either by individual perception tools or by combining input from multiple tools into a fused output. Human speech, gestures and object perception provided the input for commanding robot actions and enabling collaboration in shared industrial tasks.

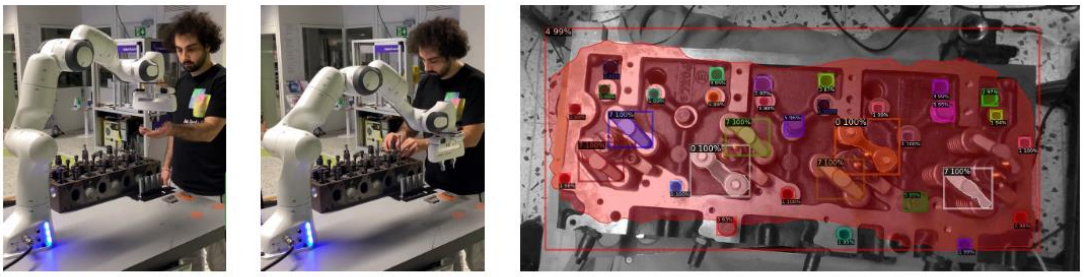

Figure 6: Results of robot-human hand-over and assembly tracking experiments proposed by TAU. (a) hand-over of a rocker arm from robot to human. (b) human assembly action of the rocker arm, while the robot fetches another rocker arm. (c) assembly tracking results, with several objects and their locations detected inside the Diesel engine bounding box.

Simulation Environments and Data

During this year, several models for simulating a robotics farming environment were developed by AU. These include models of robots: AGI's Robotti, tractor implements, agricultural fields, agricultural buildings, farm animals, humans, etc. as well as scenarios based mainly on the use of lidar and camera sensors. The animals, humans, buildings, fences, etc. are used mainly as obstacles that can be perceived by the robot sensors and avoided or causing the robot to stop to avoid a collision. Such an environment can be used to generate large volumes of simulation data that could be processed on-line (during the simulation) or off-line (handy for deep learning purposes and repeatability.

Moreover, AUTH dealt with the creation of a synthetic multimodal (audio, video and 3D models) dataset suitable for active facial expression recognition. In addition, AUTH worked towards a tool that uses Augmented Reality (AR) technology and allows the easy creation of mixed (real and synthetic) data depicting realistic 3D human models in various real environments, captured from different camera positions.

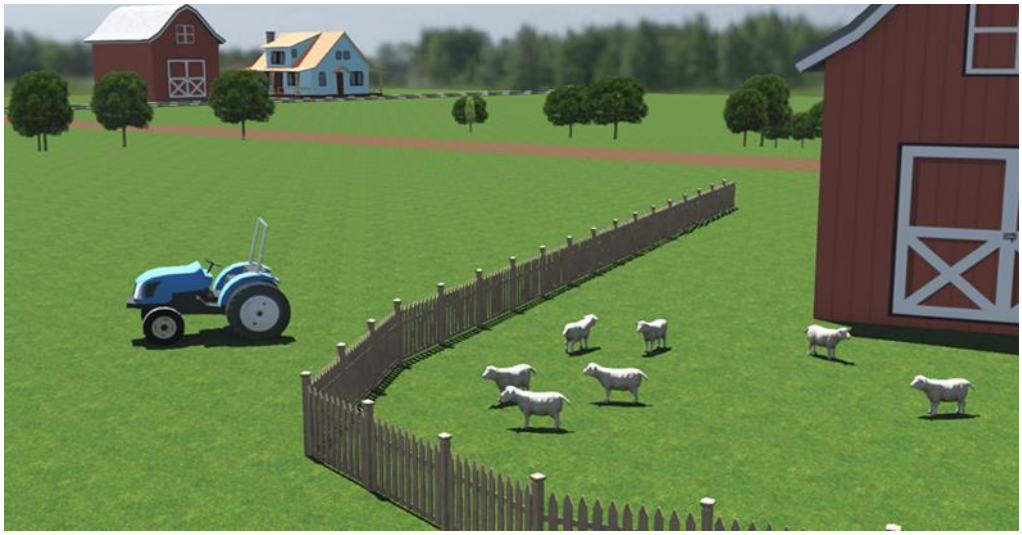

Figure 7: Simulated farm environment, including tractors, animals, barns, etc

AUTH also finalized its work on the Unity-based synthetic annotated data generation pipeline for active vision tasks and publicly released both the software and two datasets (a face datasets called ActiveFace and a full body dataset called ActiveHuman).

Figure 8: Robotti robot operating on a field and detecting the presence of humans.

Finally, CYB worked on further developing the Webots simulator and a new version was released: R2023b. The improvements include the adjunction of new primitive nodes, the integration of the TIAgo robot from PAL in Webots including new device models such as the Robotiq grippers 2f 85, 2f 140 etc. Moreover new versions of the Webots ROS 2 interface with several improvements were released whereas the

Webots web interface was improved with the ability to set up competitions running in webots.cloud.

Dissemination

During the last year of the project, a multitude of efforts, in various directions, were undertaken by the consortium in order to attract interest in the project findings and results. The project website and its social media accounts (Facebook, Twitter, LinkedIn, YouTube) continued to receive frequent updates and posts regarding new publications, project news, forthcoming events etc. Virtually all project publications but also videos and slides from conference presentations are available on the website.

Moreover, the consortium succeeded in generating, for yet another year, a high volume of publications. Eighteen papers were presented or accepted in highly esteemed international conferences (IROS, ICIP, EUSIPCO etc.) and eight papers were published or accepted in scientific journals (including IEEE T-RO and IEEE RA-L). In total, project partners have published 99 conference and journal papers, one edited book and numerous preprints. Not bad at all!

A particularly important dissemination activity was the organization, in collaboration with the MARVEL H2020 research project, of the 4-day "Deep Learning for Autonomous Systems and Smart Cities" Summer School, in Aarhus, Denmark, which was attended by more than 35 persons and included a hands-on introduction to the OpenDR toolkit. Other significant dissemination activities included 2 invited/plenary talks and participation, through talks or posters in partners' booths, in events like ERF 2023, IROS 2023 and ROSCon23.

Epilogue

A very productive four-year journey has come to an end! The project achieved, with hard work and effective partners collaboration, its goal of creating a modular, open and non-proprietary toolkit for core robotic functionalities by harnessing deep learning to provide advanced perception and cognition capabilities. The very good reception of the toolkit (more than 550 stars and 17000 downloads so far) shows that the toolkit has already had a significant impact in the robotics and deep learning communities. Apart from this, the consortium believes that the scientific outcomes of OpenDR, documented in 100 publications, have pushed the state of

the art in various areas such as deep environment and human-centric active perception and cognition, deep robot action and decision making, simulation environments and data generation etc.

It was a very interesting journey indeed! We hope that our contributions will be a small but important step towards smarter and more efficient robots of all kinds.

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