OpenDR
–
Open Deep Learning Toolkit for Robotics

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

Deliverable D9.7: 3rd report on dissemination, exploitation plans and project newsletter

Date of delivery: 31 Dec 2022

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

Version: 2.0

Title:	**D9.7:3rd report on dissemination, exploitation plans and project newsletter**	
Project:	**OpenDR** (ICT-10-2019-2020 RIA)	
Nature:	Report	**Dissemination Level: PU**blic
Authors:	Paraskevi Nousi (AUTH), Nikolaos Passalis (AUTH), Pavlos Tosidis (AUTH), Eleni Chantzi (AUTH), Anastasios Tefas (AUTH), Nikolaos Nikolaidis (AUTH), Daniel Dias (CYB), Roel Pieters (TAU), Halil Ibrahim Ugurlu (AU), Erdal Kayacan (AU), Lukas Hedegaard Morsing (AU), Illia Oleksiienko (AU), Alexandros Iosifidis (AU), Bas van der Heijden (TUD), Jelle Luijkx (TUD), Laura Ferranti (TUD), Jens Kober (TUD), Robert Babuśka (TUD), Niclas Vödisch (ALU-FR), Tim Welschehold (ALU-FR), Abhinav Valada (ALU-FR)	
Lead Beneficiary:	AUTH (Aristotle University of Thessaloniki)	
WP	9	
Doc ID:	OPENDR_D9.7.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
0.1	6/11/2022	Document outline created
1.0	10/12/2022	First complete draft, ready for internal review
2.0	22./12/2022	Final version, incorporating reviewer comments

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

Table of 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, D9.7 "3rd report on dissemination, exploitation plans and project newsletter", is a document presenting the project dissemination and exploitation activities within the third 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, workshops, tutorials 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 SME partners exploitation plans. The 3rd project newsletter is also included. Overall, the consortium believes that the project progresses 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 project results known to the industry, academia and the general public, therefore increasing its impact. This deliverable presents the project dissemination activities within the third year (Y3) of the project, namely activities in the period M25-M36. It contains (in Section 2) information regarding scientific publications, organization of the project summer school and a tutorial at IROS 2022, participation in two additional tutorials and a workshop, project-related presentations at various events, invited talks, as well as information regarding the project website and social media channels, dissemination material that has been created, collaboration with other projects and DIHs, to mention just a few. Information regarding the performance of dissemination and communication activities in Y3, 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 current exploitation plans of the three SMEs that are part of the project consortium are presented, whereas Section 4 includes a brief overview of the project's handling of IPR protection and management. Finally, the third project newsletter, an important element of the project's communication strategy, is briefly described in Section 5 and included in Appendix I.

The presented results and activities indicate that the project's progress in terms of both dissemination and exploitation planning is extremely good. The consortium is confident that the project outcomes will be widely disseminated and efficiently exploited.

2. Dissemination and Communication

The project dissemination and communication plan was described in the Description of Action and also summarized in D9.5 (M12). The dissemination and communication activities in this period as well as a KPI-based evaluation of the project performance in this respect are presented in the subsections below.

2.1 Dissemination and Communication Activities

2.1.1 Project Website

Similar to the previous periods, the project website has been used for publishing updates regarding the project activities and progress, maintaining this way the role of the main information hub for people interested in its activities. During this period (M25-M36) a significant number of posts such as publications (Figure 1), events (Figure 2), short and easy to understand articles (OpenDr briefs, Figure 3) regarding the conducted research, regularly published by OpenDR

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

researchers, have been posted, enriching the available content. The project website counts until today more than 65 posts (news, OpenDR briefs, videos etc) and more than 11000 visitors, almost doubling the number of visitors since M24. In the main menu bar, a new button named “OpenDR Toolkit” which leads to the project’s GitHub repository was added, providing easy access to the toolkit.

39	A UAV Video Data Generation Framework for Improved Robustness of UAV Detection Methods	C. Symeonidis, C. Anastasiadis, N. Nikolaidis	IEEE 24th International Workshop on Multimedia Signal Processing (MMSP), 2022
40	Continual SLAM: Beyond Lifelong Simultaneous Localization and Mapping through Continual Learning	N. Vödisch, D. Cattaneo, W. Burgard and A. Valada	International Symposium on Robotics Research (ISRR), 2022
41	Learning Long-Horizon Robot Exploration Strategies for Multi-Object Search in Continuous Action Spaces	F. Schmalstieg, D. Honerkamp, T. Welschehold, and A. Valada	International Symposium on Robotics Research (ISRR), 2022
42	OpenDR: An Open Toolkit for Enabling High Performance, Low Footprint Deep Learning for Robotics	N. Passalis, S. Pedrazzi, R. Babuska, W. Burgard, D. Dias, F. Ferro, M. Gabbouj, O. Green, A. Iosifidis, E. Kayacan, J. Kober, O. Michel, N. Nikolaidis, P. Nousi, R. Pieters, M. Tzelepi, A. Valada, and A. Tefas	2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2022
43	Continual 3D Convolutional Neural Networks for Real-time Processing of Videos	L. Hedegaard and A. Iosifidis	European Conference on Computer Vision (ECCV), 2022
44	Continual Inference: A Library for Efficient Online Inference with Deep Neural Networks in PyTorch	L. Hedegaard and A. Iosifidis	European Conference on Computer Vision (ECCV), 2022
45	Active Particle Filter Networks: Efficient Active Localization in Continuous Action Spaces and Large Maps	D. Honerkamp, S. Guttikonda, and A. Valada	IROS 2022 Workshop: Probabilistic Robotics in the Age of Deep Learning, 2022
46	Continual Transformers: Redundancy-Free Attention for Online Inference	L. Hedegaard, A. Bakhtiarnia A. Iosifidis	Neural Information Processing Systems (Neural IPS) Workshops, 2022

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

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

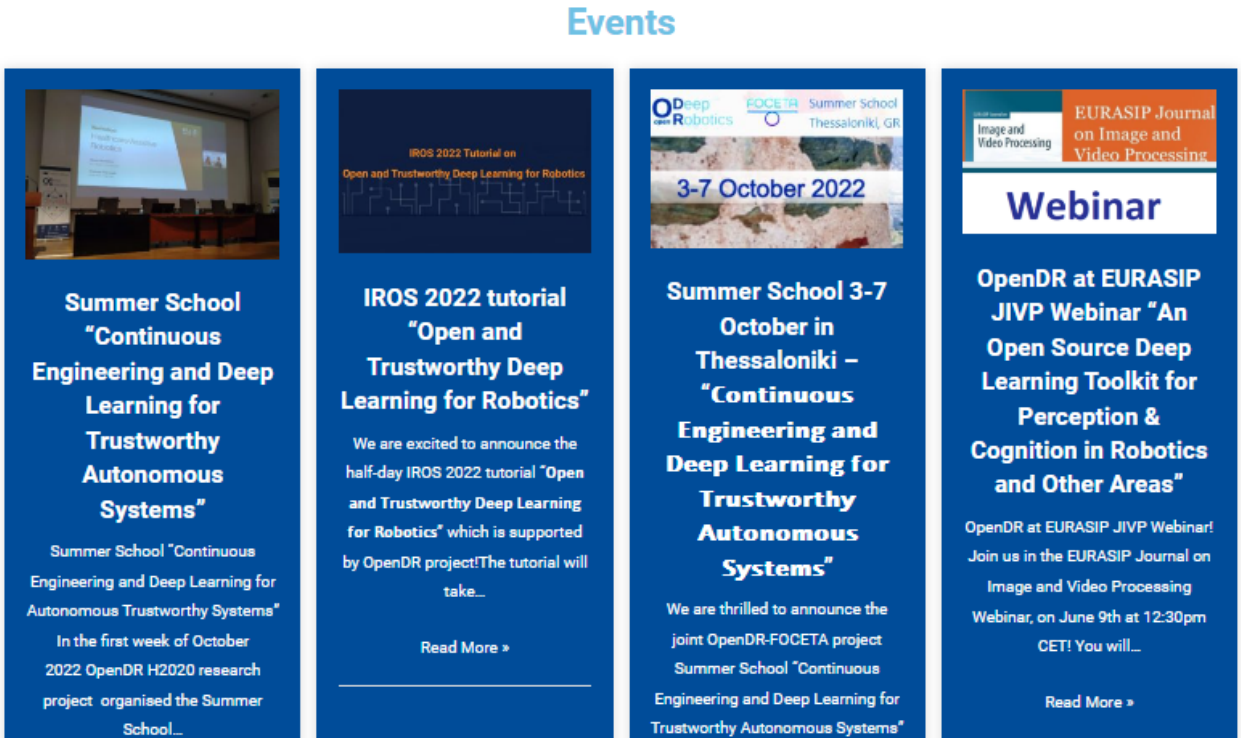

Figure 2: Project's latest events listed in the project website

04	EAGERx: An Engine Agnostic Gym Environment with Reactive extension	Bas van der Heijden, Jelle Luijkx, Laura Ferranti, Jens Kober, and Robert Babuska	Delft University of Technology - TU Delft, Netherlands	26 January 2022
05	Learning Kinematic Feasibility for Mobile Manipulation through Deep Reinforcement Learning	Tim Welschehold, Daniel Honerkamp, and Abhinav Valada	Albert-Ludwigs-University Freiburg, Germany	11 March 2022
06	Robotic simulations on the web: education, dissemination and competition organization	Daniel Dias & CYB team	Cyberbotics Ltd, Switzerland	5 May 2022
07	Pal Robotics Integration of Deep Learning tools on TIAGo robot for healthcare usecase	Lorna Mckinlay, Julia Atsu Romero, Thomas Peyrucain, Gizem Bozdemir	PAL Robotics, Spain	5 July 2022
08	Utilizing a Deep Learning human detection model with an agricultural field robot powered by Deep Learning	AgroIntelli team	Agro Intelligence ApS, Denmark	14 November 2022

Figure 3: Brief articles authored by OpenDR partners during the period M25-M36

2.1.2 Social Media

During this period OpenDR has been remarkably active in all its social media accounts. The content posted in these accounts was focused on the dissemination of journal and conference publications of the OpenDR team, as well as other project activities, like events, tutorials, toolkit releases etc. More information about the project's activity in LinkedIn, Twitter, YouTube and Facebook is provided below.

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

Twitter (@OpenDR_EU)

The project's Twitter account has made remarkable progress in terms of followers and tweets/retweets (Figure 4) during the third year of the project. More specifically, the twitter account has gathered until today more than 427 followers, while the number of tweets/retweets has exceeded 235. By comparing with M24 figures (230 followers, 150 tweets/retweets) it is obvious that the project is very active in Twitter and this has led to an increase of almost 100% in its followers.

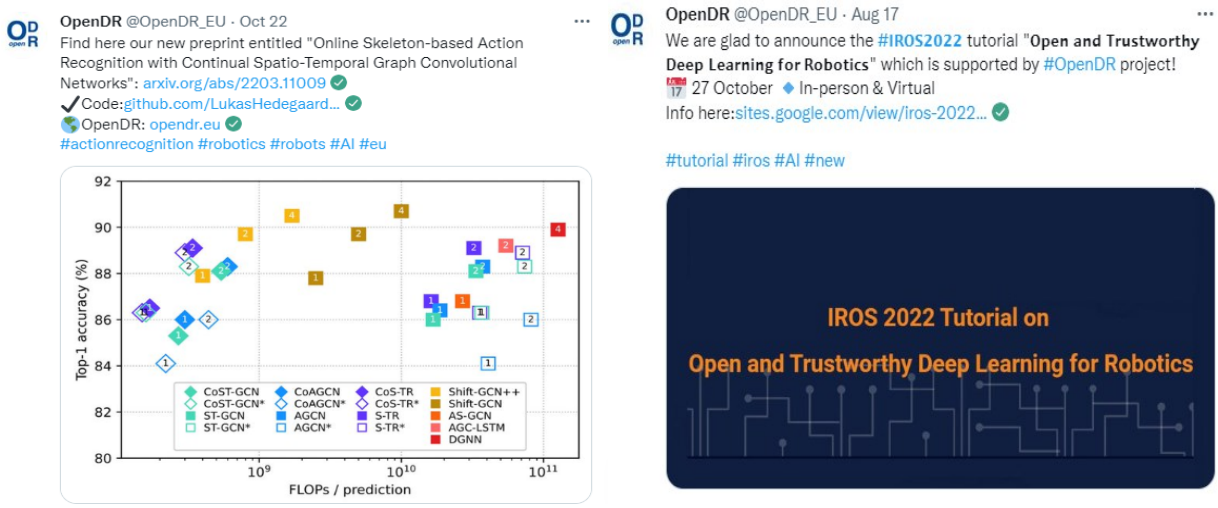

Figure 4: Examples of tweets published during the 3rd year

Facebook (@Opendr.eu)

The OpenDR Facebook account is also on a rising path, regarding the page likes/followers and the posts. More precisely, the account has until today more than 335 page likes/ 347 followers, while more than 200 posts (Figure 5) have been created, promoting the project's results, events, news, etc. and contributing at the same time to the increase of the website visitors. In comparison to the previous year's activity, there has been a significant increase of almost 90% in the page likes (190 on M24) as well as in the page posts (130 on M24).

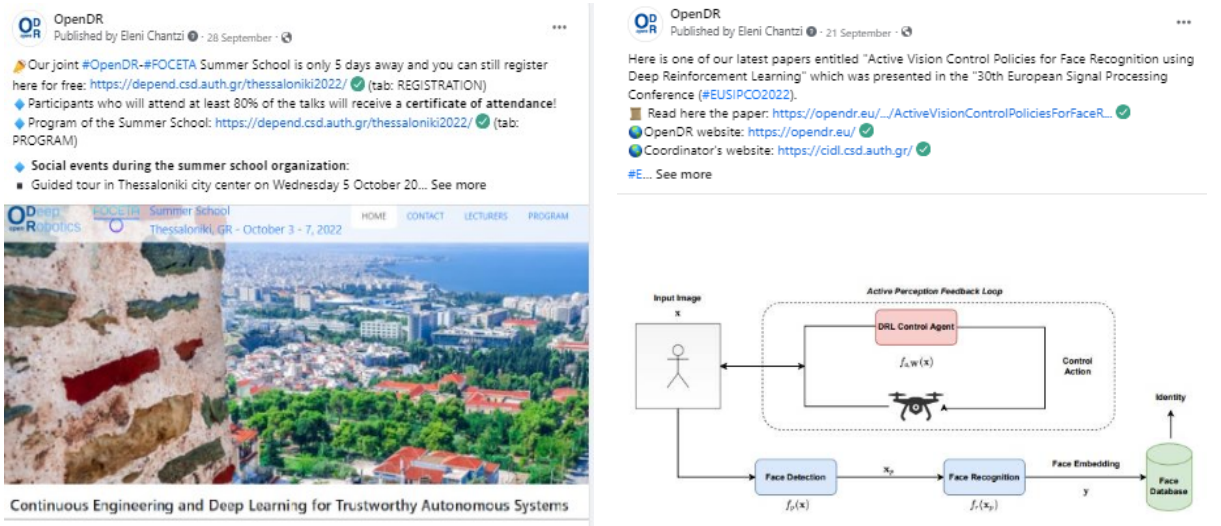

This project has received funding from the *European Union's Horizon 2020 research and innovation programme* under grant agreement No *871449*.

Figure 5: Examples of facebook posts published during the 3rd year

LinkedIn

LinkedIn plays a leading role among the project's social media accounts. Indeed the project has gathered until today more than 850 followers in its official page "OpenDR Research Project" and about 160 members in the group page "OpenDR H2020 Research Project". The importance in the use of this medium lies in the fact that the project results (for example conference and journal publications and toolkit tutorials) are disseminated to a more targeted community with scientific and industrial background, interested in OpenDR objectives and results. Compared to the metrics of M24, there has been a very large increase of the LinkedIn page followers (+605 followers), while the LinkedIn group has exceeded the goal of 150 members (+30 since M24).

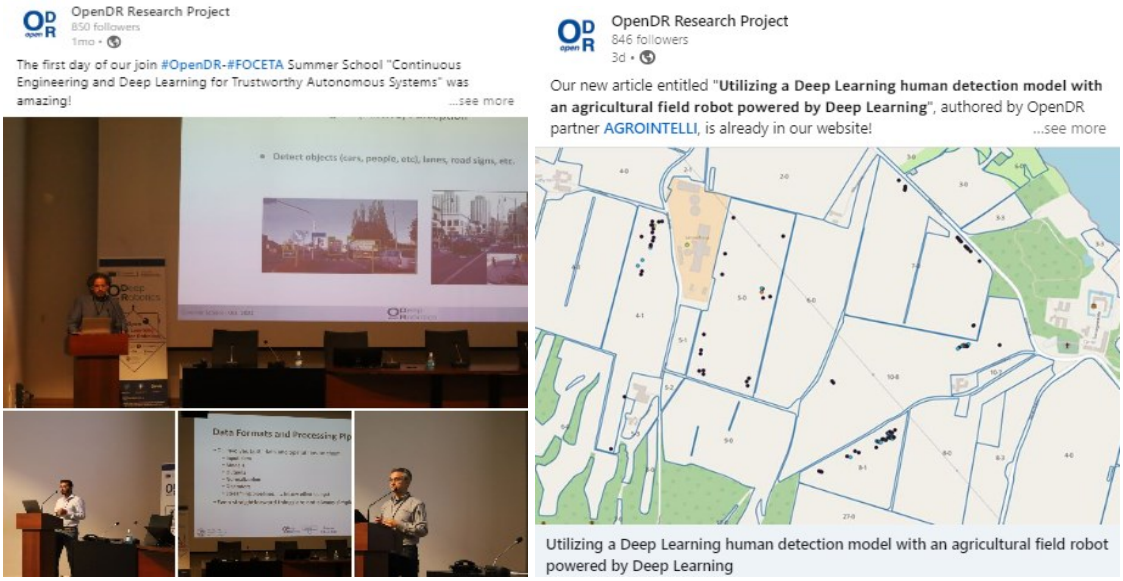

Figure 6: Examples of LinkedIn posts published during the 3rd year

YoutTube

The OpenDR consortium continued to use its YouTube channel in order to support the project's dissemination. So far, 16 videos have been uploaded, most of them from conference presentations, and the channel has gathered in total more than 3500 views (+1500 since M24) and 62 subscribers (+22 since M24).

Partners social media activity

Apart from the activity in the "official" project channels mentioned above, OpenDR partners were also active in publicising the project as well as related events and achievements. For example, partner PAL authored 33 OpenDR related posts in its Facebook, LinkedIn and Twitter corporate accounts during M25-36 (Figure 7).

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

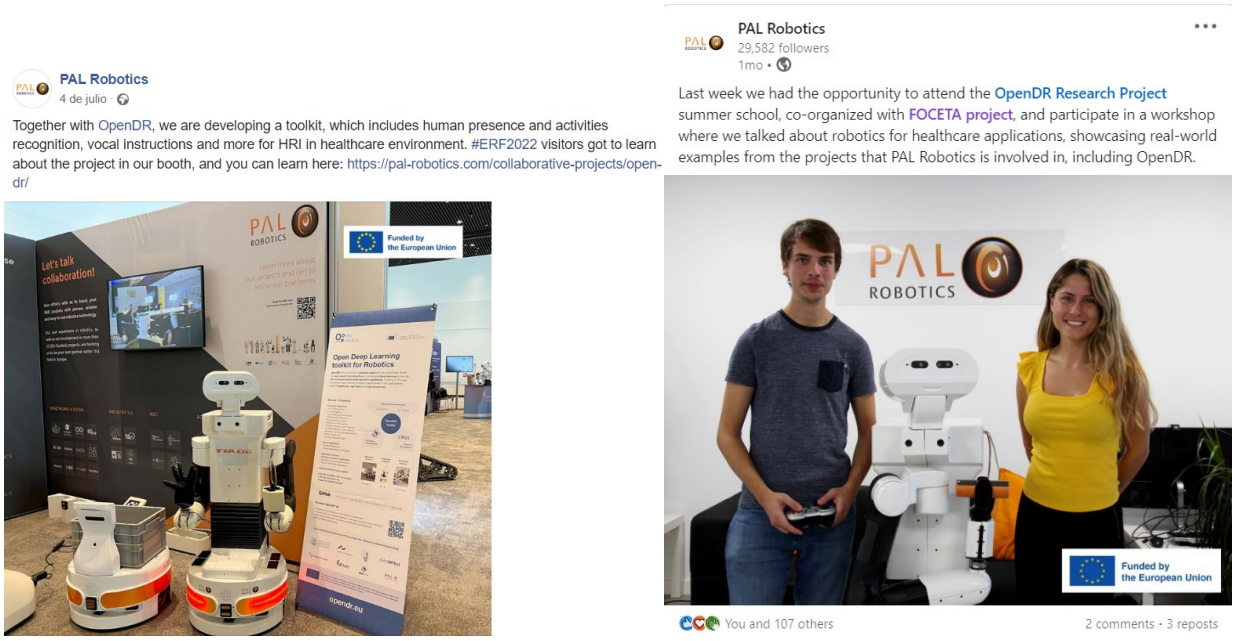

Figure 7: Examples of PAL Robotics posts regarding OpenDR

2.1.3 OpenDR Webpages at Partners' Websites

During this period, OpenDR partners continued promoting the project through their websites. Indeed, CYB website includes a project description and a link to the project website (Figure 8).

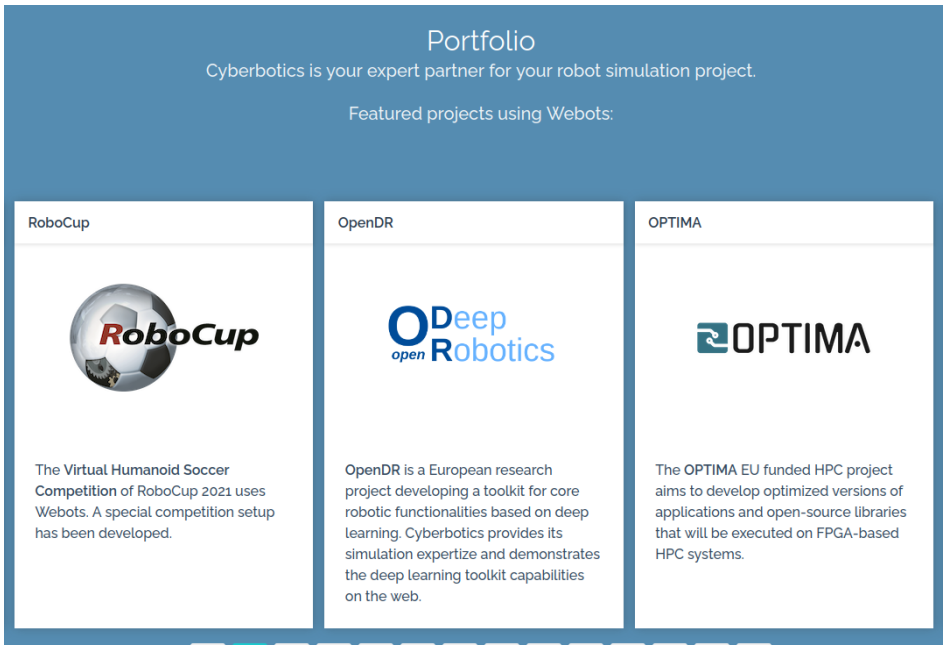

Figure 8: OpenDR description and link in the CYB website

The project is also mentioned in the Collaborative Projects section of PAL's website, where a card is displayed describing the project and linking to the OpenDR website (Figure 9).

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

Figure 9: PAL's website pages describing the project and linking to the OpenDR website

The Summer School "Continuous Engineering and Deep Learning for Trustworthy Autonomous Systems" (Section 2.1.6.2) was advertised through the website of the Aristotle University of Thessaloniki (AUTH) Department of Informatics (Figure 10), providing information about the program, the talks and links to the OpenDR website.

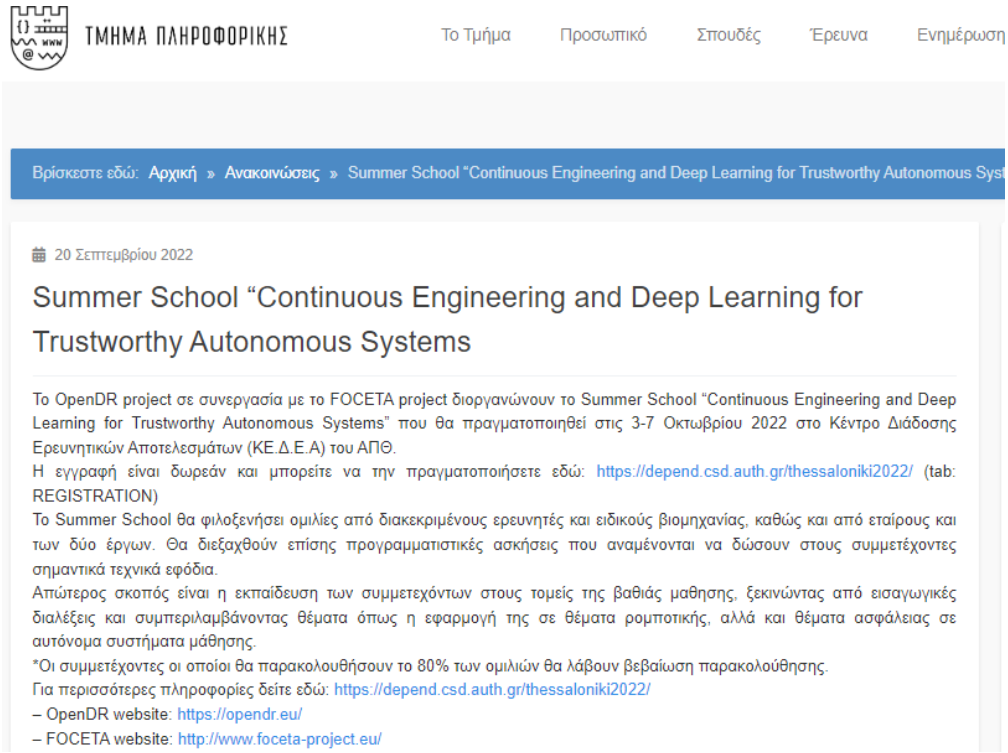

Figure 10: Post regarding the Summer School in the website of the Dept. of Informatics, AUTH

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

2.1.4 Dissemination Material (Brochures, Banners, Posters)

To further support the dissemination activities, a poster has been created in the current period. The poster (Figure 11) includes essential information about the project, such as the objectives and the consortium, as well as a brief reference to results related to version 1.1 of the OpenDR toolkit. Infographics, information about the project's social media accounts and a QR code which leads to the toolkit repository are also included. Furthermore, the first roll-up banner of the project was created (Figure 12). The banner also displays essential information about the project (logo, consortium, funding, social media accounts). Both the poster and the roll up banner were recently used in the Summer School "Continuous Engineering and Deep Learning for Trustworthy Autonomous Systems" and can be freely downloaded from the website (section "Communication Material"). Another version of the poster was displayed in PAL´s booth in ERF22 (Figure 13).

Figure 11: OpenDR poster

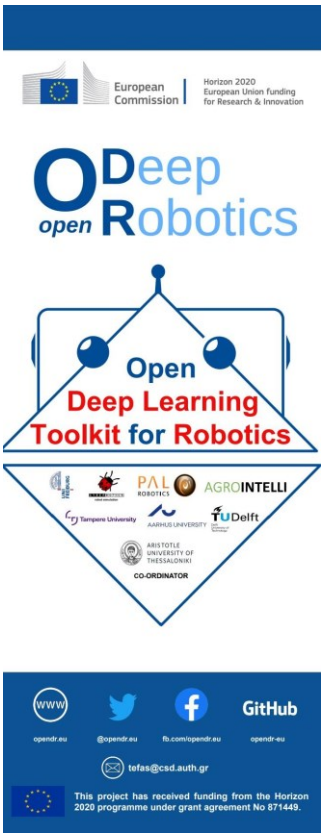

Figure 12: OpenDR roll up banner

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

Figure 13: OpenDR Poster

Updated versions of the OpenDR brochure (Figure 14) and the project factsheet were also created and can be found/downloaded in the project's website. This new edition of the brochure includes additional information regarding the features of the OpenDR Toolkit Version 1.1. Both the brochure and the project factsheet have been already distributed in recent events such as the European Robotics Forum 2022 and the Summer School "Continuous Engineering and Deep Learning for Trustworthy Autonomous Systems".

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

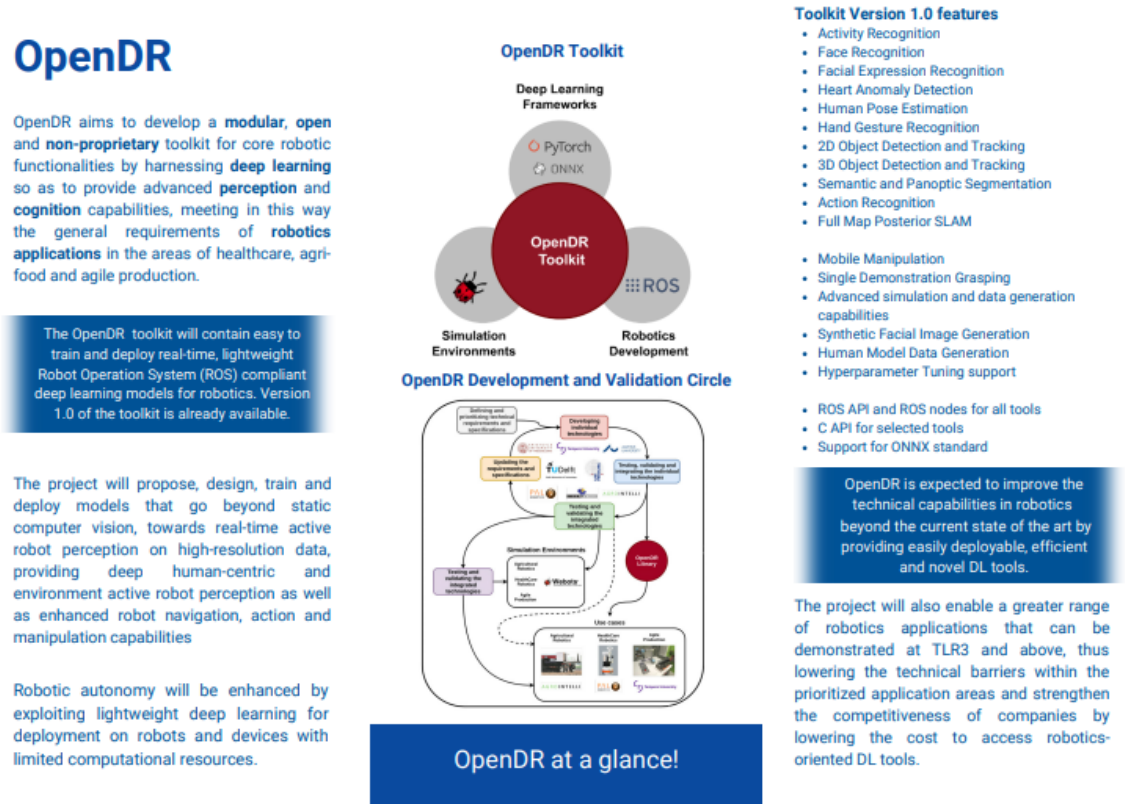

Figure 14: The updated version of the OpenDR brochure

Another brochure was created by PAL for the HRI2022 -EU meet up event in Barcelona, Spain sponsored by PAL Robotics and Eurecat (Figure 15).

Figure 15: Brochure created by PAL for the HRI2022 -EU meetup event

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 6 journal papers that have been accepted for publication or published.

A paper was published by AUTH in Neurocomputing journal:

- *M. Tzelepi, N. Passalis and A. Tefas, Probabilistic Online Self-Distillation, Neurocomputing, Vol. 493, pp. 592-604, 2022*

Deploying state-of-the-art deep learning models on devices with limited computational power imposes certain computation and storage restrictions. Knowledge Distillation, i.e. training compact models by transferring knowledge from more powerful models, constitutes a promising route to address this issue that has been followed during the recent years. A limitation of conventional knowledge distillation is that it is a long-lasting, computationally and memory demanding process, since it requires multiple stages of training process. To this end, a novel online probabilistic self-distillation method, namely Probabilistic Online Self-Distillation (POSD), aiming to improve the performance of any deep neural model in an online manner, was proposed. The authors argued that considering a classification problem, apart from the explicit concepts expressed with the hard labels, there are also implicit concepts expressed with the so-called latent labels. These implicit concepts reflect similarities among data, regardless of the classes. Then, the goal is to maximize the Mutual Information between the data samples and the latent labels. In this way, they were able to derive additional knowledge from the model itself, without the need of building multiple identical models or using multiple models to teach each other, like existing online distillation methods, rendering the POSD method more efficient. The experimental evaluation on six datasets showed that the proposed method improves the classification performance.

Furthermore, the AU team published a paper in Robotics:

- *Ugurlu, H.I.; Pham, X.H.; Kayacan, E. Sim-to-Real Deep Reinforcement Learning for Safe End-to-End Planning of Aerial Robots. Robotics 2022, 11, 109. https://doi.org/10.3390/robotics11050109*

In this study, a novel end-to-end path planning algorithm based on deep reinforcement learning is proposed for aerial robots deployed in dense environments. The learning agent finds an obstacle-free way around the provided rough, global path by only depending on the observations from a forward-facing depth camera. A novel deep reinforcement learning framework is proposed to train the end-to-end policy with the capability of safely avoiding obstacles. The Webots open-source robot simulator is utilized for training the policy, introducing highly randomized environmental configurations for better generalization. The training is performed without dynamics calculations through randomized position updates to minimize the amount of data processed. The trained policy is first comprehensively evaluated in simulations involving physical dynamics and software-in-the-loop flight control. The proposed method is proven to have a 38% and 50% higher success rate compared to both deep reinforcement learning-based and artificial potential field-based baselines, respectively. The generalization capability of the method is verified in simulation-to-real transfer without further training. Real-time experiments are conducted with several trials in two different scenarios, showing a 50% higher success rate of the proposed method compared to the deep reinforcement learning-based baseline. The full paper can be found in D5.3.

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

An AU-TAU joint work was published in Software Impacts journal:

- *I. Oleksiienko, D. T. Tran and A. Iosifidis, "Variational Neural Networks implementation in Pytorch and JAX", Software Impacts, vol. 14: 100431, 2022*

Bayesian Neural Networks consider a distribution over the network's weights, which provides a tool to estimate the uncertainty of a neural network by sampling different models for each input. Variational Neural Networks (VNNs) consider a probability distribution over each layer's outputs and generate parameters for it with the corresponding sub-layers. The authors provide two Python implementations of VNNs with PyTorch and JAX machine learning libraries that ensure reproducibility of the experimental results and allow implementing uncertainty estimation methods easily in other projects. The full paper can be found in D4.3.

In yet another joint effort, AU, TAU and AUTH published a paper in IEEE Access:

- *D.T. Tran, N. Passalis, A. Tefas, M. Gabbouj and A. Iosifidis, "Attention-based Neural Bag-of-Features Learning for Sequence Data", IEEE Access, vol. 10, pp. 45542-45552, 2022*

In this paper, the authors propose 2D-Attention (2DA), a generic attention formulation for sequence data, which acts as a complementary computation block that can detect and focus on relevant sources of information for the given learning objective. The proposed attention module is incorporated into the recently proposed Neural Bag of Feature (NBoF) model to enhance its learning capacity. Since 2DA acts as a plug-in layer, injecting it into different computation stages of the NBoF model results in different 2DA-NBoF architectures, each of which possesses a unique interpretation. Extensive experiments were conducted in financial forecasting, audio analysis as well as medical diagnosis problems to benchmark the proposed formulations in comparison with existing methods, including the widely used Gated Recurrent Units. The empirical analysis shows that the proposed attention formulations can not only improve performances of NBoF models but also make them resilient to noisy data. The full paper can be found in D3.3.

Finally, ALU-FR published two papers in IEEE Robotics and Automation Letters (RAL):

- *M. Büchner and A. Valada, "3D Multi-Object Tracking Using Graph Neural Networks With Cross-Edge Modality Attention," in IEEE Robotics and Automation Letters, vol. 7, no. 4, pp. 9707-9714, 2022*

Online 3D multi-object tracking (MOT) has witnessed significant research interest in recent years, largely driven by demand from the autonomous systems community. However, 3D offline MOT is relatively less explored. Labelling 3D trajectory scene data at a large scale while not relying on high-cost human experts is still an open research question. In this work, ALU-FR proposed Batch3DMOT which follows the tracking-by-detection paradigm and represents real-world scenes as directed, acyclic, and category-disjoint tracking graphs that are attributed using various modalities such as camera, LiDAR, and radar. A multi-modal graph neural network that uses a cross-edge attention mechanism mitigating modality intermittence, which translates into sparsity in the graph domain was presented. Additionally, the authors presented attention-weighted convolutions over frame-wise k-NN neighborhoods as suitable means to allow information exchange across disconnected graph components. The approach was evaluated using various sensor modalities and model configurations on the challenging nuScenes and KITTI datasets. Extensive experiments demonstrate that the proposed approach yields an overall improvement of 3.3% in the AMOTA score on nuScenes thereby setting the new state-of-the-art for 3D tracking and further enhancing false positive filtering. The full paper can be found in D4.3.

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

- *R. Mohan and A. Valada, "Perceiving the Invisible: Proposal-Free Amodal Panoptic Segmentation," in IEEE Robotics and Automation Letters, vol. 7, no. 4, pp. 9302-9309, 2022*

Amodal panoptic segmentation aims to connect the perception of the world to its cognitive understanding. It entails simultaneously predicting the semantic labels of visible scene regions and the entire shape of traffic participant instances, including regions that may be occluded. In this work, the authors formulate a proposal-free framework that tackles this task as a multi-label and multi-class problem by first assigning the amodal masks to different layers according to their relative occlusion order and then employing amodal instance regression on each layer independently while learning background semantics. For this, they propose the PAPS architecture that incorporates a shared backbone and an asymmetrical dual-decoder consisting of several modules to facilitate within-scale and cross-scale feature aggregations, bilateral feature propagation between decoders, and integration of global instance-level and local pixel-level occlusion reasoning. Further, they propose the amodal mask refiner that resolves the ambiguity in complex occlusion scenarios by explicitly leveraging the embedding of unoccluded instance masks. Extensive evaluation on the BDD100K-APS and KITTI-360-APS datasets demonstrate that the proposed approach sets the new state-of-the-art on both benchmarks. The full paper can be found in D4.3.

2.1.5.2 Conference Papers

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

The consortium published a joint paper in IEEE/RSJ International Conference on Intelligent Robots and Systems where an overview of the OpenDR toolkit is provided:

- *Passalis, N., Pedrazzi, S., Babuska, R., Burgard, W., Dias, D., Ferro, F., ... & Tefas, A. (2022). OpenDR: An Open Toolkit for Enabling High Performance, Low Footprint Deep Learning for Robotics. Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, 2022.*

Existing Deep Learning (DL) frameworks typically do not provide ready-to-use solutions for robotics, where very specific learning, reasoning, and embodiment problems exist. Their relatively steep learning curve and the different methodologies employed by DL compared to traditional approaches, along with the high complexity of DL models, which often leads to the need of employing specialized hardware accelerators, further increase the effort and cost needed to employ DL models in robotics. Also, most of the existing DL methods follow a static inference paradigm, as inherited by the traditional computer vision pipelines, ignoring active perception, which can be employed to actively interact with the environment in order to increase perception accuracy. In this paper, the authors presented the Open Deep Learning Toolkit for Robotics (OpenDR). OpenDR aims at developing an open, non-proprietary, efficient, and modular toolkit that can be easily used by robotics companies and research institutions to efficiently develop and deploy AI and cognition technologies to robotics applications, providing a solid step towards addressing the aforementioned challenges. The design choices, along with an abstract interface that was created to overcome these challenges are also detailed. This interface can describe various robotic tasks, spanning beyond traditional DL cognition and inference, as known by existing frameworks, incorporating openness, homogeneity and robotics-oriented perception e.g., through active perception, as its core design principles.

AUTH has presented 7 papers in conferences such as EUSIPCO 2022, ICIP 2022, ICPR 2022, MMSP 2022 etc:

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

- *P. Tosidis, N. Passalis and A. Tefas, "Active Vision Control Policies for Face Recognition using Deep Reinforcement Learning," 2022 30th European Signal Processing Conference (EUSIPCO), 2022, pp. 1087-1091.*

Robotic systems are capable of interacting with their environment in order to better sense their surroundings. This key ability of robotic systems is often ignored when developing Deep Learning models, since the later are usually trained using static datasets. This limits the ability of robots to perceive the environment in challenging scenarios. On the other hand, integrating perception and action in tightly coupled systems while operating on-the-edge, holds the credentials for deploying DL-enabled robots in such scenarios; Thus leading to more robust agents that can solve challenging tasks more accurately. In this work, the authors investigate whether active perception approaches can be employed and integrated into robotic systems in order to improve face recognition accuracy, as well as, study the effect of such an approach on the computational requirements for edge applications. To this end, they propose a DRL-based control approach for training agents that are able to identify task-relevant objects, as well as, issue the appropriate control commands to acquire better results. Through the conducted experimental evaluation, it is demonstrated that the proposed method leads to significant improvements in face recognition over the rest of the evaluated approaches by providing accurate control commands. The full paper can be found in D3.3.

- *C. Symeonidis, C. Anastasiadis, N. Nikolaidis, "A UAV Video Data Generation Framework for Improved Robustness of UAV Detection Methods," in Proceedings of the IEEE International Workshop on Multimedia Signal Processing (MMSP), 2022.*

Recent advances have facilitated the development and popularization of Unmanned Aerial Vehicles (UAVs) that can operate semi or fully autonomously. The real-time, accurate visual detection of UAVs is crucial for various tasks and applications including surveillance (e.g., detecting UAVs flying over restricted areas such as airports) or multi-robot systems (e.g., a swarm of UAVs that need to cooperate and avoid collisions between swarm members in GPS-denied environments). The small target-to-image ratio and large similarity with other flying objects makes the visual detection of UAVs a challenging task. In addition, data distribution shifts can have a major negative impact to UAV detection frameworks, often trained on a wide variety of datasets to achieve an adequate level of robustness. As an attempt to mitigate the effect of these issues, the authors present a method that can generate realistic annotated video data depicting flying UAVs, using as input real background videos and 3D UAV models. The conducted experimental evaluation showed that the synthetic data are both challenging and realistic and that detectors trained on a combination of real-world and synthetic data, exhibit an improved generalization performance, achieving better precision rates when evaluated on real datasets that are visually distinct from the corresponding real training data. The full paper can be found in D6.3.

- *C. Symeonidis, I. Mademlis, I. Pitas and N. Nikolaidis, "Auth-Persons: A Dataset for Detecting Humans in Crowds from Aerial Views," in Proceedings of the IEEE International Conference on Image Processing (ICIP), 2022, pp. 596-600.*

Recent advances in artificial intelligence, control and sensing technologies have facilitated the development of autonomous Unmanned Aerial Vehicles (UAVs). Detecting humans from video input captured on-the-fly from UAVs is a critical task for ensuring flight safety, mostly handled with lightweight Deep Neural Networks (DNNs). However, the detection of individual people in the case of dense crowds and/or distribution shifts (i.e., significant visual differences

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

between the training and the test sets) is still very challenging. This paper presents AUTH-Persons, a new, annotated, publicly available video dataset, that consists of both real and synthetic footage, suitable for training and evaluating aerial-view person detection algorithms. The synthetic data were collected from 8 visually distinct photorealistic outdoor environments, and they mostly contain scenes with crowded areas, where heavy occlusions and high person densities pose challenges to common detectors. This dataset is employed to evaluate the generalization performance of various state-of-the-art detection frameworks, by testing them on environments that are visually distinct from those they have been trained on. To achieve this, AUTH-Persons was carefully split into training and test sets that differ realistically in appearance. Finally, given that Non-Maximum Suppression (NMS) methods at the end of person detection pipelines typically suffer in crowded scenes, the performance of various NMS algorithms is also compared in AUTH-Persons. The full paper can be found in D3.3.

- *M. Tzelepi, N. Tragkas, and A. Tefas. "Improving binary semantic scene segmentation for robotics applications." International Conference on Engineering Applications of Neural Networks. Springer, Cham, 2022.*

Robotics applications are accompanied by particular computational restrictions, i.e., operation at sufficient speed, on embedded low power GPUs, and also for high-resolution input. Semantic scene segmentation performs an important role in a broad spectrum of robotics applications, e.g., autonomous driving. In this paper, AUTH focuses on binary segmentation problems, considering the specific requirements of the robotics applications. To this aim, they utilized the BiseNet model, which achieves significant performance considering the speed-segmentation accuracy trade-off. The target of this work is two-fold. Firstly, to propose a lightweight version of BiseNet model, providing significant speed improvements. Secondly, to explore different losses for enhancing the segmentation accuracy of the proposed lightweight version of BiseNet on binary segmentation problems. The experiments conducted on various high and low power GPUs, utilising two binary segmentation datasets validated the effectiveness of the proposed method.

- *M. Tzelepi, C. Symeonidis, N. Nikolaidis and A. Tefas, "Real-time synthetic-to-real human detection for robotics applications." 2022 13th International Conference on Information, Intelligence, Systems & Applications (IISA). IEEE, 2022.*

During the recent years, Deep Learning achieved exceptional performance in various computer vision tasks, paving auspicious research directions for its application in robotics. A key component for its exceptional performance is the availability of sufficient training data. However obtaining such amount of training data constitutes a challenging task, especially considering robotics applications. Thus, synthetic data have recently been regarded as a promising tool to overcoming the data availability problem. In this work the authors first build a synthetic human dataset, and then they train a lightweight model, capable of operating in real-time for high-resolution input on low-power GPUs, for discriminating between humans and non-humans. The target of the work is to assess the generalization of the model trained on synthetic data, to real data, and also to explore the effect of using (few) real images in the training phase. As it is shown through quantitative and qualitative results the use of only few real images can beneficially affect the performance of the synthetic-to-real real-time model. The full paper can be found in D3.3.

- *M. Tzelepi, C. Symeonidis, N. Nikolaidis and A. Tefas, "Multilayer Online Self-Acquired Knowledge Distillation" 26th International Conference on Pattern Recognition (ICPR), 2022*

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

Online knowledge distillation has been proposed as an auspicious approach for circumventing the flaws of the conventional offline distillation (i.e., complex, and computationally and memory demanding process). In this work, a novel online self-distillation method, named Multilayer Online SelfAcquired Knowledge Distillation (MOSAKD), is proposed, aiming to develop fast-to-execute and effective models that can comply with applications with memory and computational restrictions, e.g., robotics applications. The MOSAKD method is able to mine additional knowledge both from the intermediate and the output layers of a deep neural model in an online fashion. To achieve this goal, k-nn non-parametric density estimation for estimating the unknown probability distributions of the data samples in the feature space generated by any neural layer is used. This enables us to compute the soft labels that explicitly express the similarities of the data with the classes, by directly estimating the posterior class probabilities of the data samples. The experimental evaluation on four datasets, including a dataset of synthetic images, indicates the effectiveness of the MOSAKD method and the superiority over existing online distillation methods.

- *P. Nousi, E. Mpampis, N. Passalis, O. Green, & A. Tefas. (2022). A Novel Dataset for Evaluating and Alleviating Domain Shift for Human Detection in Agricultural Fields. IEEE Symposium Series On Computational Intelligence 2022.*

AUTH has evaluated the impact of domain shift on human detection models trained on well-known object detection datasets when deployed on data outside the distribution of the training set, as well as proposed methods to alleviate such phenomena based on the available annotations from the target domain. Specifically, the OpenDR Humans in Field dataset was introduced, collected in the context of agricultural robotics applications, using the Robotti platform, allowing for quantitatively measuring the impact of domain shift in such applications. Furthermore, the importance of manual annotation was examined by evaluating three distinct scenarios concerning the training data: a) only negative samples, i.e., no depicted humans, b) only positive samples, i.e., only images which contain humans, and c) both negative and positive samples. The results indicate that good performance can be achieved even when using only negative samples, if additional consideration is given to the training process. Furthermore, it was shown that positive samples increase performance especially in terms of better localization. The dataset is publicly available for download at https://github.com/opendr-eu/datasets. The full paper can be found in D6.3.

This year, AU members have published 3 conference papers, including one in the prestigious ECCV conference:

- *L. Hedegaard and A. Iosifidis, "Continual 3D Convolutional Neural Networks for Real-time Processing of Videos", European Conference on Computer Vision, Tel-Aviv, Israel, Oct. 23-27 2022*

The authors introduce Continual 3D Convolutional Neural Networks (Co3D CNNs), a new computational formulation of spatio-temporal 3D CNNs, in which videos are processed frame-by-frame rather than by clip. In online tasks demanding frame-wise predictions, Co3D CNNs dispense with the computational redundancies of regular 3D CNNs, namely the repeated convolutions over frames, which appear in overlapping clips. The authors show that Continual 3D CNNs can reuse pre-existing 3D-CNN weights to reduce the per-prediction floating point operations (FLOPs) in proportion to the temporal receptive field while retaining similar memory requirements and accuracy. This is validated with multiple models on Kinetics-400 and Charades with remarkable results: CoX3D models attain state-of-the-art complexity/accuracy trade-offs on Kinetics-400 with 12.1-15.3x reductions of FLOPs and 2.3-3.8% improvements in accuracy compared to regular

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

X3D models while reducing peak memory consumption by up to 48%. Moreover, we investigate the transient response of Co3D CNNs at start-up and perform extensive benchmarks of on-hardware processing characteristics for publicly available 3D CNNs. The full paper can be found in D3.3.

- *L. Hedegaard and A. Iosifidis, "Continual Inference: A Library for Efficient Online Inference with Deep Neural Networks in PyTorch", International Workshop on Computational Aspects of Deep Learning (ECCV Workshops), Tel-Aviv, Israel, Oct. 23-27 2022*

In this paper, the authors present Continual Inference, a Python library for implementing Continual Inference Networks (CINs) in PyTorch, a class of Neural Networks designed specifically for efficient inference in both online and batch processing scenarios. We offer a comprehensive introduction and guide to CINs and their implementation in practice, and provide best-practices and code examples for composing complex modules for modern Deep Learning. Continual Inference is readily downloadable via the Python Package Index and at http://www.github.com/lukashedegaard/continual-inference . The full paper can be found in D3.3.

- *L. M. Hedegaard, A. Bakhtiarnia and A. Iosifidis, "Continual Transformers: Redundancy-Free Attention for Online Inference", Advances in Neural Information Systems (NeurIPS) Workshop on Vision Transformers: Theory and Applications, 2022 (arXiv:2201.06268)*

Transformers in their common form are inherently limited to operate on whole token sequences rather than on one token at a time. Consequently, their use during online inference on time-series data entails considerable redundancy due to the overlap in successive token sequences. In this work, the authors propose novel formulations of the Scaled Dot-Product Attention, which enable Transformers to perform efficient online token-by-token inference on a continual input stream. Importantly, our modifications are purely to the order of computations, while the outputs and learned weights are identical to those of the original Transformer Encoder. We validate our Continual Transformer Encoder with experiments on the THUMOS14, TVSeries and GTZAN datasets with remarkable results: Our Continual one- and two-block architectures reduce the floating point operations per prediction by up to 63x and 2.6x, respectively, while retaining predictive performance. The full paper can be found in D3.3.

AU and TAU collaborated in two conference papers for MLSP and ICPR 2022:

- *K. Chumachenko, A. Iosifidis and M. Gabbouj, "Self-attention Neural Bag-of-Features", IEEE International Workshop on Machine Learning for Signal Processing, Xi'an, China, August 22-25, 2022*

In this work, the authors propose several attention formulations for multivariate sequence data. They build on top of the recently introduced 2D-Attention and reformulate the attention learning methodology by quantifying the relevance of feature/temporal dimensions through latent spaces based on self-attention rather than learning them directly. In addition, the authors propose a joint feature-temporal attention mechanism that learns a joint 2D attention mask highlighting relevant information without treating feature and temporal representations independently. The proposed approaches can be used in various architectures and we specifically evaluate their application together with the Neural Bag of Features feature extraction module. Experiments on several sequence data analysis tasks show the improved performance yielded by our approach compared to standard methods. The full paper can be found in D3.3.

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

- *K. Chumachenko, A. Iosifidis and M. Gabbouj, "Self-attention fusion for audiovisual emotion recognition with incomplete data", International Conference on Pattern Recognition, Montreal, Quebec, Canada, August 21-25, 2022*

In this paper, the authors consider the problem of multimodal data analysis with a use case of audiovisual emotion recognition. We propose an architecture capable of learning from raw data and describe three variants of it with distinct modality fusion mechanisms. While most of the previous works consider the ideal scenario of presence of both modalities at all times during inference, the authors evaluate the robustness of the model in the unconstrained settings where one modality is absent or noisy, and propose a method to mitigate these limitations in a form of modality dropout. Most importantly, they have found that following this approach not only improves performance drastically under the absence/noisy representations of one modality, but also improves the performance in a standard ideal setting, outperforming the competing methods. The full paper can be found in D3.3.

Moreover, TAU presented a paper in IEEE CASE 2022:

- *A. Mehman Sefat, A. Angleraud, E. Rahtu and R. Pieters, "SingleDemoGrasp: Learning to Grasp From a Single Image Demonstration", IEEE International Conference on Automation Science and Engineering (CASE), Mexico City, Mexico, August 20-24, 2022*

In this work TAU proposed a fast modelling approach for vision-based object grasp detection. Based on a single human object grasp annotation, an augmented dataset of RGB training images is generated, to be utilized for training a grasp detection model. Four different planar grasp detection models, each with different human annotations and grasp detection approach, are evaluated and implemented in simulation. All models are light-weight (below 0.5GB), enabling real-time inference. Best results were obtained with a keypoint-based model, which was further demonstrated with real robot grasping experiments. In all, from a human object grasp annotation, the augmented dataset and grasp model training, the approach enables the generation of a planar object grasp model in under 15 minutes. The full paper can be found in D5.3.

Also, TUD presented a paper at a NeurIPS workshop:

- *J. Luijkx, Z. Ajanović, L. Ferranti and J. Kober, "PARTNR: Pick and place Ambiguity Resolving by Trustworthy iNteractive leaRning", NeurIPS Robot Learning Workshop: Trustworthy Robotics, 2022*

Several recent works show impressive results in mapping language-based human commands and image scene observations to direct robot executable policies (e.g., pick and place poses). However, these approaches do not consider the uncertainty of the trained policy and simply always execute actions suggested by the current policy as the most probable ones. This makes them vulnerable to domain shift and inefficient in the number of required demonstrations. The authors extend previous works and present the PARTNR algorithm that can detect ambiguities in the trained policy by analyzing multiple modalities in the pick and place poses using topological analysis. PARTNR employs an adaptive, sensitivity-based, gating function that decides if additional user demonstrations are required. User demonstrations are aggregated to the dataset and used for subsequent training. In this way, the policy can adapt promptly to domain shift and it can minimize the number of required demonstrations for a well-trained policy. The adaptive threshold enables users to achieve the user-acceptable level of ambiguity to execute the policy autonomously and in turn, increase the trustworthiness of our system. The performance of PARTNR is demonstrated in a table-top pick and place task. The full paper can be found in D5.3.

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

Finally, ALU-FR presented two papers in ISRR 2022 and a paper at an ICML workshop:

- *F. Schmalstieg, D. Honerkamp, T. Welschehold and A. Valada, "Learning Long-Horizon Robot Exploration Strategies for Multi-Object Search in Continuous Action Spaces", Proceedings of the International Symposium on Robotics Research (ISRR), 2022*

Recent advances in vision-based navigation and exploration have shown impressive capabilities in photorealistic indoor environments. However, these methods still struggle with long-horizon tasks and require large amounts of data to generalize to unseen environments. In this work, ALU-FR presented a novel reinforcement learning approach for multi-object search that combines short-term and long-term reasoning in a single model while avoiding the complexities arising from hierarchical structures. In contrast to existing multi-object search methods that act in granular discrete action spaces, the proposed approach achieves exceptional performance in continuous action spaces. The authors performed extensive experiments and showed that it generalizes to unseen apartment environments with limited data. Furthermore, zero-shot transfer of the learned policies to an office environment in real world experiments were demonstrated. The full paper can be found in D5.3.

- *N. Vödisch, D. Cattaneo, W. Burgard, and A. Valada, "Continual SLAM: Beyond Lifelong Simultaneous Localization and Mapping through Continual Learning", in Proceedings of the International Symposium on Robotics Research (ISRR), 2022*

Robots operating in the open world encounter various different environments that can substantially differ from each other. This domain gap also poses a challenge for Simultaneous Localization and Mapping (SLAM) being one of the fundamental tasks for navigation. In particular, learning-based SLAM methods are known to generalise poorly to unseen environments hindering their general adoption. In this work, the authors introduce the novel task of continual SLAM extending the concept of lifelong SLAM from a single dynamically changing environment to sequential deployments in several drastically differing environments. To address this task, they propose CL-SLAM leveraging a dual-network architecture to both adapt to new environments and retain knowledge with respect to previously visited environments. CL-SLAM is compared to learning-based as well as classical SLAM methods and the advantages of leveraging online data are shown. They also extensively evaluate CL-SLAM on three different datasets and demonstrate that it outperforms several baselines inspired by existing continual learning-based visual odometry methods. The full paper can be found in D4.3.

- *N. Dorka, T. Welschehold, and W. Burgard, "Dynamic Update-to-Data Ratio: Minimizing World Model Overfitting", in Decision Awareness in Reinforcement Learning Workshop at ICML, 2022*

Early stopping based on the validation set performance is a popular approach to find the right balance between under- and overfitting in the context of supervised learning. However, in reinforcement learning, even for supervised sub-problems such as world model learning, early stopping is not applicable as the datase is continually evolving. As a solution, the authors propose a new general method that dynamically adjusts the update to data (UTD) ratio during training based on underand overfitting detection on a small subset of the continuously collected experience not used for training. They apply the method to DreamerV2, a state-of-the-art modelbased reinforcement learning algorithm, and evaluate it on the DeepMind Control Suite and the Atari 100k benchmark. The results demonstrate that one can better balance under- and overestimation by adjusting the UTD ratio with the proposed approach compared to the default setting in

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

DreamerV2 and that it is competitive with an extensive hyperparameter search which is not feasible for many applications. The method eliminates the need to set the UTD hyperparameter by hand and even leads to a higher robustness with regard to other learning-related hyperparameters further reducing the amount of necessary tuning. The full paper can be found in D5.3.

2.1.5.3 Preprints

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

AUTH and AU members published the following preprint:

- *I. Oleksiienko, P. Nousi, N. Passalis, A. Tefas and A. Iosifidis, "VPIT: Real-time Embedded Single Object 3D Tracking Using Voxel Pseudo Images", arXiv:2206.02619, 2022*

In this paper, the authors propose a novel voxel-based 3D single object tracking (3D SOT) method called Voxel Pseudo Image Tracking (VPIT). VPIT is the first method that uses voxel pseudo images for 3D SOT. The input point cloud is structured by pillar-based voxelization, and the resulting pseudo image is used as an input to a 2D-like Siamese SOT method. The pseudo image is created in the Bird's-eye View (BEV) coordinates, and therefore the objects in it have constant size. Thus, only the object rotation can change in the new coordinate system and not the object scale. For this reason, the authors replace multi-scale search with a multi-rotation search, where differently rotated search regions are compared against a single target representation to predict both position and rotation of the object. Experiments on KITTI Tracking dataset show that VPIT is the fastest 3D SOT method and maintains competitive Success and Precision values. Application of a SOT method in a real-world scenario meets with limitations such as lower computational capabilities of embedded devices and a latency-unforgiving environment, where the method is forced to skip certain data frames if the inference speed is not high enough. The authors implement a real-time evaluation protocol and show that other methods lose most of their performance on embedded devices, while VPIT maintains its ability to track the object. The full paper can be found in D4.3.

In addition, AU published the following four preprints in the current period:

- *N. Heidari and A. Iosifidis, "Learning Diversified Feature Representations for Facial Expression Recognition in the Wild", arXiv:2210.09381, 2022*

Diversity of the features extracted by deep neural networks is important for enhancing the model generalization ability and accordingly its performance in different learning tasks. Facial expression recognition in the wild has attracted interest in recent years due to the challenges existing in this area for extracting discriminative and informative features from occluded images in real-world scenarios. In this paper, the authors propose a mechanism to diversify the features extracted by CNN layers of state-of-the-art facial expression recognition architectures for enhancing the model capacity in learning discriminative features. To evaluate the effectiveness of the proposed approach, they incorporate this mechanism in two state-of-the-art models to (i) diversify local/global features in an attention-based model and (ii) diversify features extracted by different learners in an ensemble-based model. Experimental results on three well-known facial expression recognition in-the-wild datasets, AffectNet, FER+, and RAF-DB, show the effectiveness of our method, achieving the state-of-the-art performance of 89.99% on RAF-DB, 89.34% on FER+ and the competitive accuracy of 60.02% on AffectNet dataset.

- *I. Oleksiienko and A. Iosifidis, "Layer Ensembles", arXiv:2210:04882, 2022*

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

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, the authors introduce a novel method for uncertainty estimation called Layer Ensembles 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. They further introduce Optimized Layer Ensembles with an inference procedure that reuses common layer outputs, achieving up to 19x speed up and quadratically reducing memory usage. They also show that Layer Ensembles 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. The full paper can be found in D4.3.

- *L. Hedegaard, N. Heidari and A. Iosifidis, "Online Skeleton-based Action Recognition with Continual Spatio-Temporal Graph Convolutional Networks", arXiv:2203.11009, 2022*

Graph-based reasoning over skeleton data has emerged as a promising approach for human action recognition. However, the application of prior graph-based methods, which predominantly employ whole temporal sequences as their input, to the setting of online inference entails considerable computational redundancy. In this paper, the authors tackle this issue by reformulating the Spatio-Temporal Graph Convolutional Neural Network as a Continual Inference Network, which can perform step-by-step predictions in time without repeat frame processing. To evaluate the method, they create a continual version of STGCN, CoST-GCN, alongside two derived methods with different self-attention mechanisms, CoAGCN and CoS-TR. They investigate weight transfer strategies and architectural modifications for inference acceleration, and perform experiments on the NTU RGB+D 60, NTU RGB+D 120, and Kinetics Skeleton 400 datasets. Retaining similar predictive accuracy, they observe up to 109× reduction in time complexity, on-hardware accelerations of 26×, and reductions in maximum allocated memory of 52% during online inference.

- *L. Hedegaard, A. Alok, J. Jose, A. Iosifidis, Structured Pruning Adapters, arXiv: 2211.10155, 2022*

The authors propose Structured Pruning Adapters (SPAs), a family of compressing, task-switching network adapters, that accelerate and specialize networks using tiny parameter sets. Specifically, they propose a channel- and a block-based SPA and evaluate them with a suite of pruning methods on both computer vision and natural language processing benchmarks. Compared to regular structured pruning with fine-tuning, our channel-SPA improves accuracy by 6.9% on average while using half the parameters at 90% pruned weights. Alternatively, it can learn adaptations with 17×fewer parameters at 70% pruning with 1.6% lower accuracy. Similarly, the proposed block-SPA requires far fewer parameters than pruning with fine-tuning.

Finally, ALU-FR managed to publish five preprints:

- *D. Honerkamp, T. Welschehold and A. Valada, "N2M2: Learning Navigation for Arbitrary Mobile Manipulation Motions in Unseen and Dynamic Environments", arXiv:2206.08737, 2022*

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

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

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. The full paper can be found in D5.3.

- *D. Honerkamp, S. Guttikonda and A. Valada, "Active Particle Filter Networks: Efficient Active Localization in Continuous Action Spaces and Large Maps", arXiv:2209.09646, 2022*

Accurate localization is a critical requirement for most robotic tasks. The main body of existing work is focused on passive localization in which the motions of the robot are assumed given, abstracting from their influence on sampling informative observations. While recent work has shown the benefits of learning motions to disambiguate the robot's poses, these methods are restricted to granular discrete actions and directly depend on the size of the global map. The authors propose Active Particle Filter Networks (APFN), an approach that only relies on local information for both the likelihood evaluation as well as the decision making. To do so, they couple differentiable particle filters with a reinforcement learning agent that attends to the most relevant parts of the map. The resulting approach inherits the computational benefits of particle filters and can directly act in continuous action spaces while remaining fully differentiable and thereby end-to-end optimizable as well as agnostic to the input modality. They demonstrate the benefits of the proposed approach with extensive experiments in photorealistic indoor environments built from real-world 3D scanned apartments. Videos and code are available at http://apfn.cs.uni-freiburg.de/. The full paper can be found in D5.3.

- *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", arXiv:2111.14843, 2022*

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 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/.

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

- *J. Arce, N. Vödisch, D. Cattaneo, W. Burgard, and A. Valada, "PADLoC: LiDAR-Based Deep Loop Closure Detection and Registration using Panoptic Attention", in arXiv preprint arXiv:2209.09699, 2022*

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. The code of the work is publicly available at http://padloc.cs.uni-freiburg.de/. The full paper can be found in D4.3.

- *N. Dorka, T. Welschehold, J.Boedecker, and W. Burgard, "Adaptively Calibrated Critic Estimates for Deep Reinforcement Learning", in arXiv preprint arXiv:2111.12673, 2022*

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. The full paper can be found in D5.3.

2.1.5.4 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. Five new briefs have been published during this period, on topics that include "Robotic simulations on the web: education, dissemination and competition organization", "Learning Kinematic Feasibility for Mobile Manipulation through Deep Reinforcement Learning" and "PAL Robotics Integration of Deep Learning tools on TIAGo robot for healthcare use case".

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

2.1.6 Workshops, Schools Invited / Keynote Talks and Tutorials

Dissemination of OpenDR information and findings through workshops, invited/keynote talks, tutorials and seasonal schools has obviously a high impact since, usually, the audience consists of researchers or industrial representatives interested in the specific topic of the talk/tutorial. During this period, the project organized or participated in a significant number of such events, as described below.

OpenDR was present at the tutorial "Deep Learning Methods for Robot Control" which was organised by Robert Babuska (TUD) on behalf of OpenDR in August 2022 in the Conference on Control Technology and Applications (CCTA 2022). In this tutorial, the participants had the chance to attend several talks by OpenDR partners on a variety of OpenDR related topics, such as:

- Interactive Imitation Learning: Shaping policies through demonstrations, corrections, and rewards (Jens Kober, TUD)
- End to end planning for aerial-ground robots team collaboration for green transition in agriculture (Halil Ibrahim and Erdal Kayacan, AU)
- Deep learning-based perception (Nikolaos Passalis, AUTH)
- EAGERx: a robot learning toolkit for effective transfer from simulation to real-world applications (Jelle Luijkx, TUD)
- Vision-based navigation of aerial robots - using (un)conventional sensors (AU)

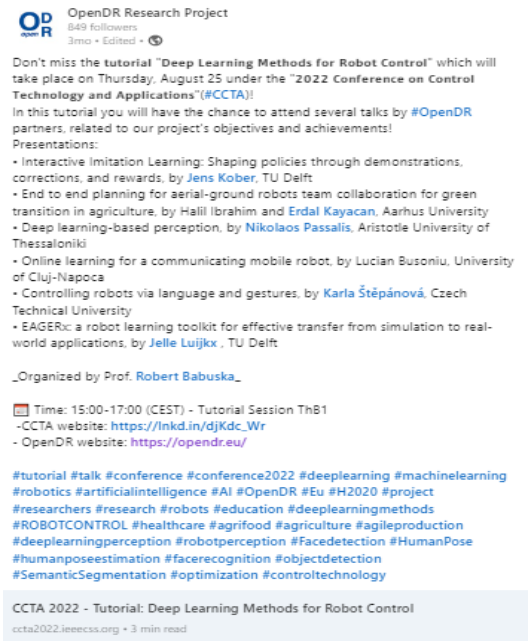

Figure 16: A LinkedIn post promoting the participation of OpenDR in the CCTA tutorial

OpenDR was also present at the tutorial "Tools for Robotic Reinforcement Learning" which was organised in May 2022 within the International Conference on Robotics and Automation (ICRA). The tutorial included a hands-on session on OpenDR tool EAGERx. Among the organisers of the tutorial were Bas van der Heijden, Jelle Luijkx and Jens Kober (TUD). The tutorial was supported by the EU H2020 projects VERtical Innovation in the Domain of Robotics Enabled by Artificial intelligence Methods and OpenDR.

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

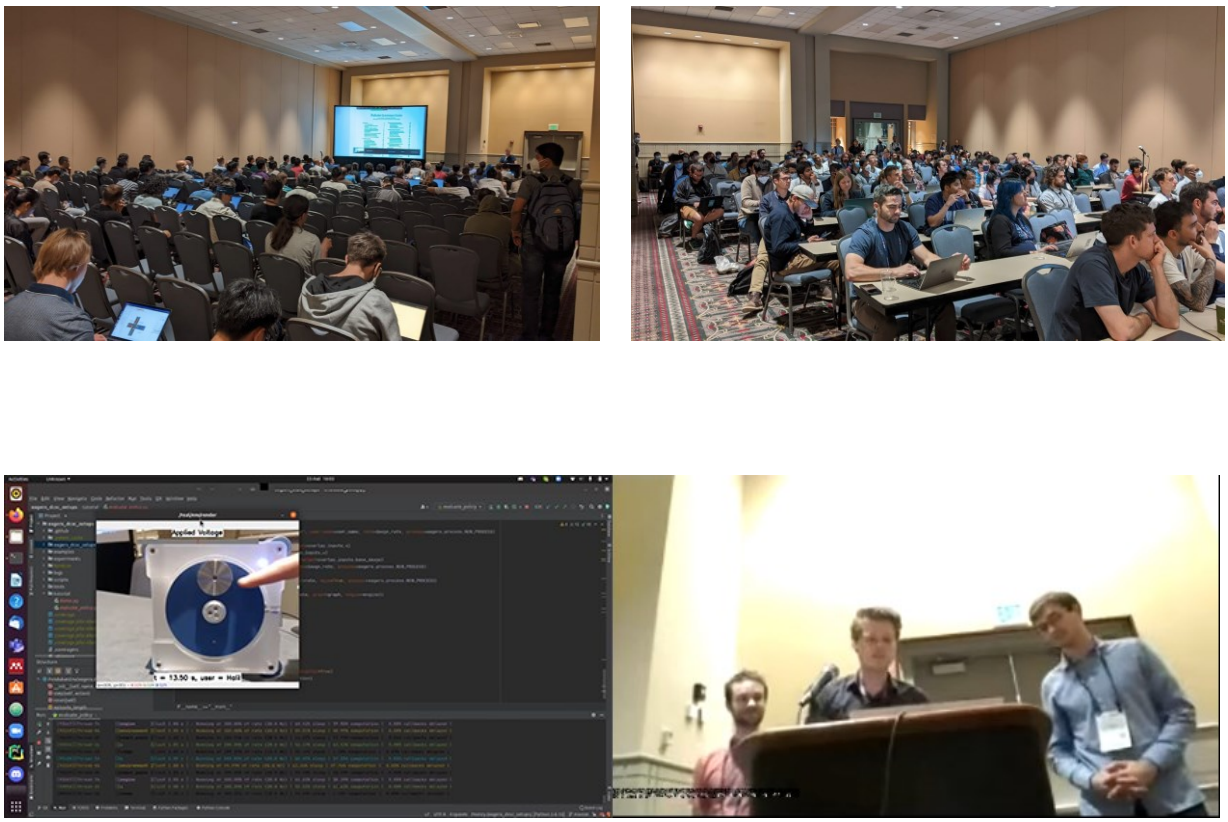

Figure 17: Impressions from the ICRA tutorial

In addition, on March 25th, 2022 Prof. Alexandros Iosifidis gave an invited talk on "Human Action Recognition in Videos" in the context of the Computer Science Colloquium Talks at the University of Louisiana Lafayette (ULL) where, among others, he presented the OpenDR work of AU in human action recognition based on video classification (Continual 3D CNNs) and human body skeletons (PST-GCN). The talk was attended by students at the ULL interested in topics of Computer/Robot Vision and Machine Learning.

Also, in June 2022 N. Passalis and N. Nikolaidis from the AUTH team were invited to present the topic "An Open Source Deep Learning Toolkit for Perception & Cognition in Robotics and Other Areas" in a one hour webinar within the EURASIP Journal and Video Processing series of monthly webinars. The webinar was promoted through the OpenDR website as well as the project's social media accounts (Figure 18).

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

Figure 18: A website post from the "News" section (left) and a LinkedIn post (right) for the Webinar

OpenDR was also present in the workshop "The challenges of the extreme-scale multi-modal analytics applications" which was organised in May 2022 by the MARVEL H2020 research project. In this workshop an overview of the OpenDR project was given by the AUTH team. OpenDR participation was announced in OpenDR website (Figure 19) as well as in the project's social media accounts.

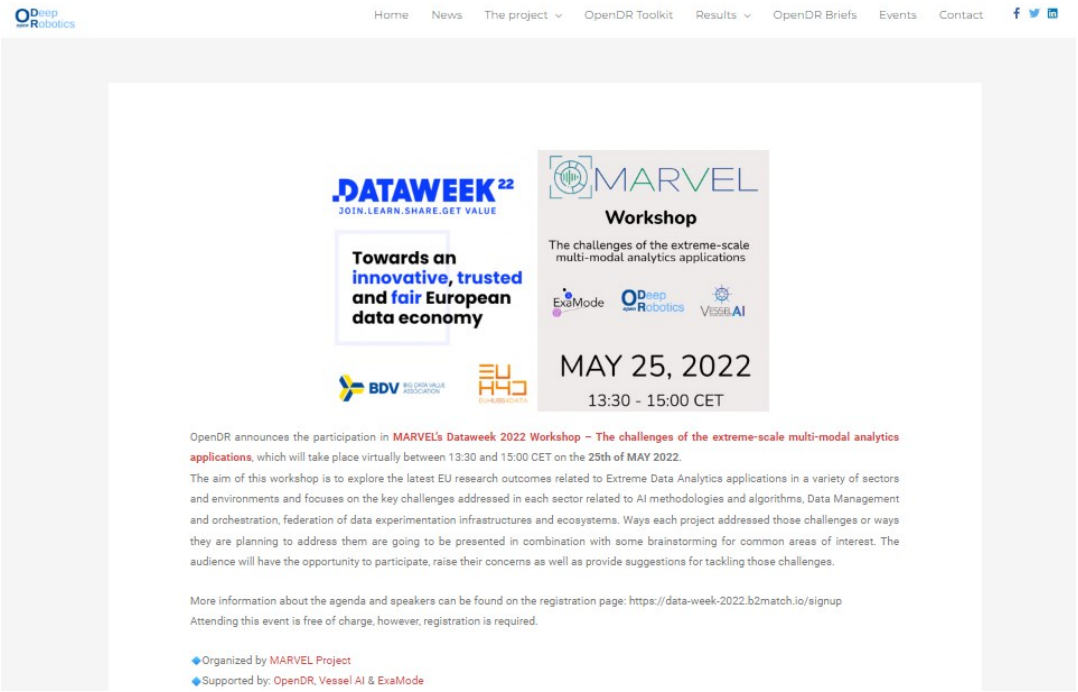

Figure 19: OpenDR's participation in MARVEL Workshop announced in the project's website

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

Moreover, AU presented their work entitled “AgroRL: end-to-end planning of air-ground multi-robot teams for green digital farming” as a poster (Figure 20) at the Workshop on Agricultural Robotics and Automation within ICRA22.

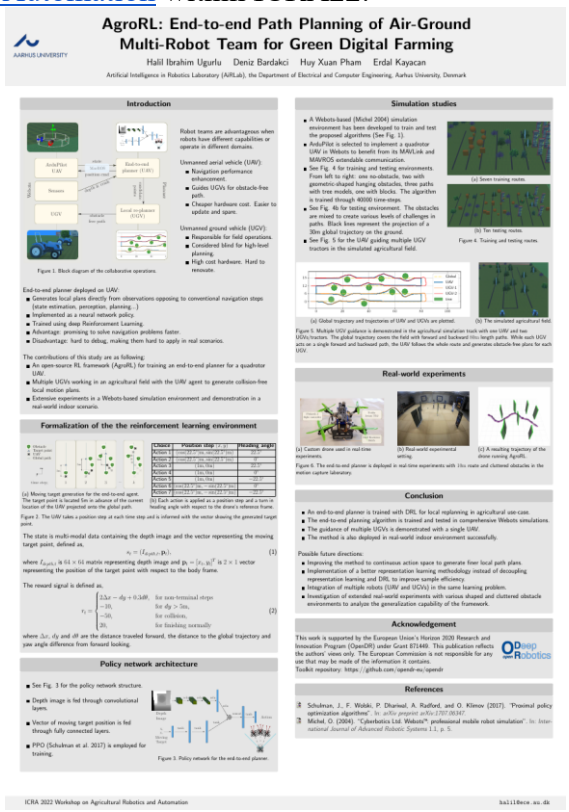

Figure 20: A poster at Workshop on Agricultural Robotics and Automation ICRA22

Last but not least, during the third year of the project, consortium members successfully organised a tutorial in the prestigious IROS 2022 and a summer school as described in the following subsections.

2.1.6.1 IROS 2022 Tutorial

On October 27th 2022 OpenDR organised the tutorial “Open and Trustworthy Deep Learning for Robotics” in IROS 2022, Kyoto, Japan, under the leadership of Aarhus University (AU). The dissemination of the tutorial was supported with the creation of a new webpage (Figure 21) by AU.

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

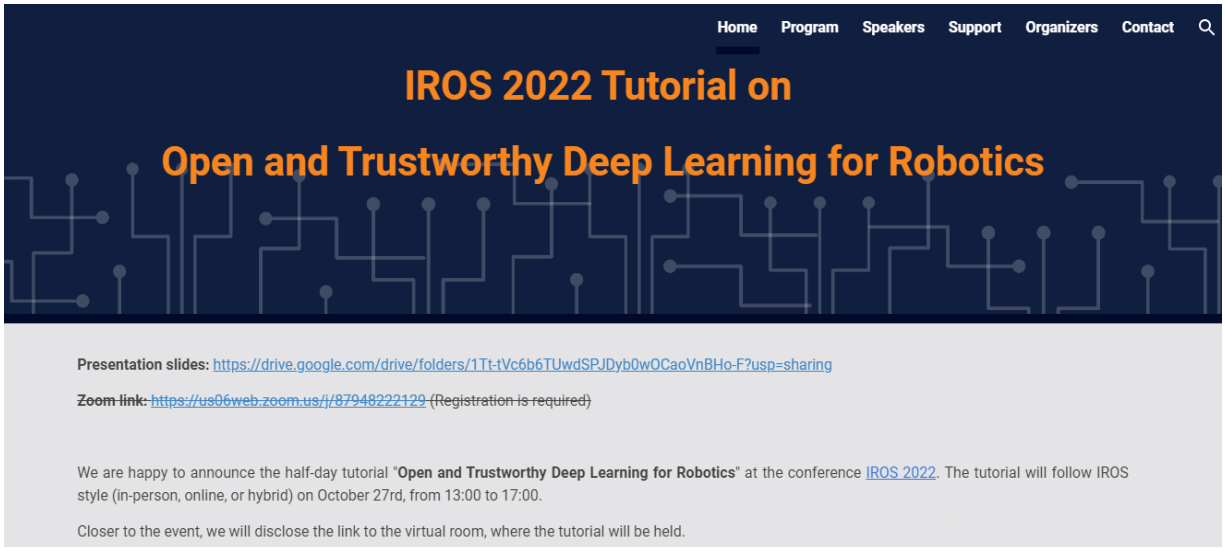

Figure 21: The landing page of the tutorial's website

The tutorial was also disseminated through the OpenDR website, in sections "News" and "Events" (Figure 22) as well as through the project's social media accounts (Figure 23).

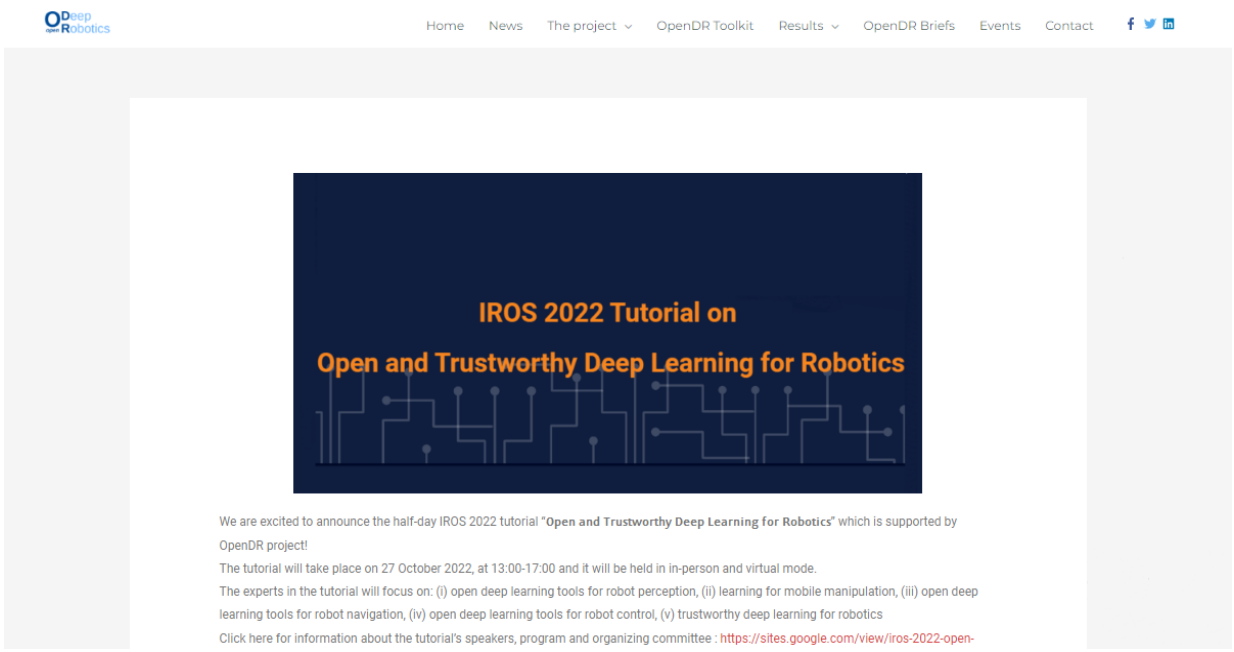

Figure 22: Announcement of the IROS 2022 tutorial in the OpenDR website

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

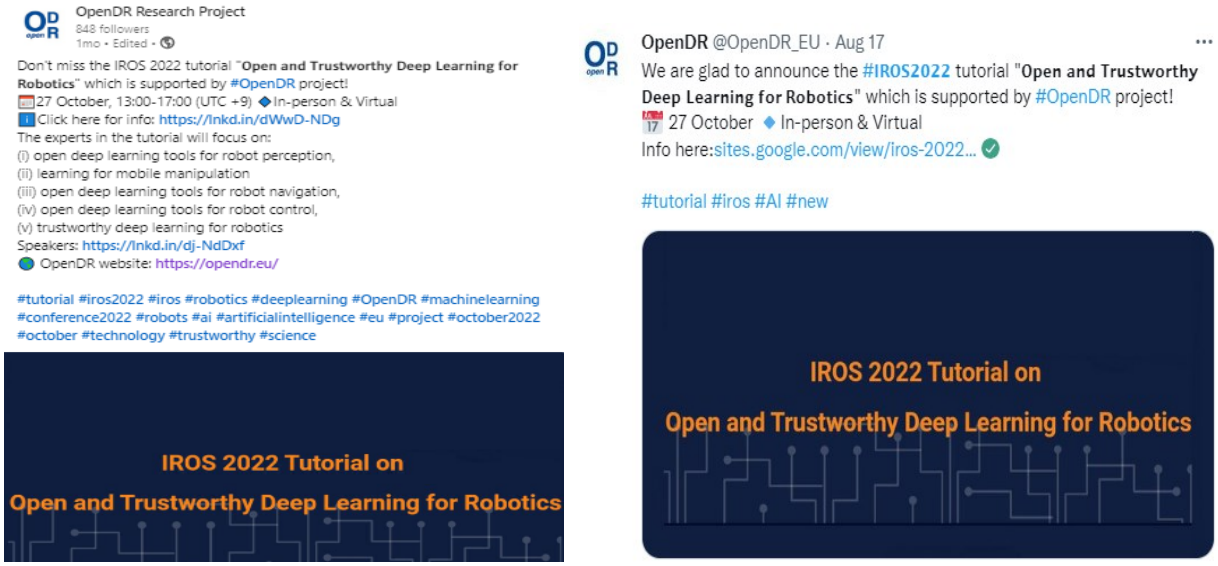

Figure 23: Announcement of the tutorial in OpenDR LinkedIn (left) and Twitter accounts (right)

The tutorial, which was attended by approximately 25 persons (Figure 24), was held in a dual (in-person and virtual) mode, and included, an introductory talk, four talks by OpenDR partners, one invited talk and a section with concluding remarks and discussion:

- Open Deep Learning Tools for Robot Perception, A. Tefas, A. Iosifidis and N. Passalis (AUTH & AU)
- Trustworthy Deep Learning for Robotics, X. Huang, University of Liverpool, UK
- Learning for Mobile Manipulation, A. Valada and D. Honerkamp (ALU-FR)
- Open Deep Learning Tools for Robot Navigation, E. Kayacan (AU)
- Open Deep Learning Tools for Robot Control, R. Babuska and J. Kober (TUD)

After the end of the tutorial, the presentation slides were uploaded on the tutorial's webpage and further disseminated through the project's website and social media accounts.

Figure 24: A snapshot from the IROS 2022 tutorial

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

2.1.6.2 Summer School on Continuous Engineering and Deep Learning for Trustworthy Autonomous Systems

In October 2022 OpenDR organised the "Continuous Engineering and Deep Learning for Trustworthy Autonomous Systems" Summer School in collaboration with FOCETA H2020 research project. The School was held in the premises of Aristotle University of Thessaloniki from the 3rd to the 7th of October and was attended by more than 75 persons from various countries (Figure 25). School attendees were provided with a certificate of attendance after the end of the event. OpenDR topics were presented on the first two days, whereas an OpenDR hands-on workshop was organised on the last day.

Various interesting talks relevant to OpenDR's research fields were provided by members of the consortium. Topics included "Deep Learning" (in 3 parts), "DL for Object Detection/Tracking", "Robotic Grasping for Agile Production", "Robotic Simulation Environments" etc. FOCETA project members provided talks on topics such as "Continuous Engineering of Trustworthy Autonomy", "DNN Robustness and Resiliency Approaches", "Formal Methods for Safe and Accountable AI", among others.

As already mentioned, on the last day of the School, a hands-on workshop took place. During this workshop, the AUTH team showcased the use of different tools of the OpenDR toolkit. The workshop attendees had the opportunity to participate by following the tutorials using their laptops. During the School, a number of social events, such as a guided tour to Thessaloniki city centre and a guided visit to the Archaeological Museum of Thessaloniki, were organised.

Figure 25: Snapshots from the Summer School's presentations

To better support the School dissemination, a website was created (Figure 26) . Furthermore, several items were posted in the social media accounts of OpenDR and FOCETA projects (Figure 27). Last but not least, a promotional poster presenting important information about the Summer School was also created (Figure 28).

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

Figure 26: A screenshot of the the landing page of the Summer School website

Figure 27: Examples of a Facebook (left) and a LinkedIn post published before and during the Summer School.

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

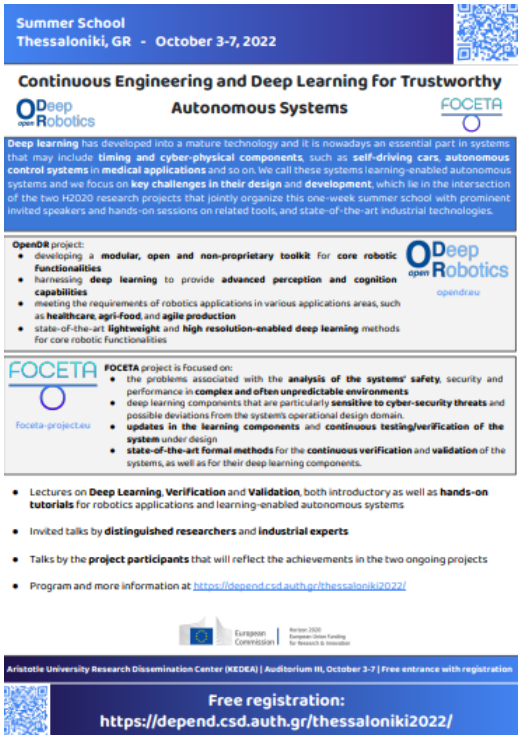

Figure 28: The promotional poster for the Summer School

As already mentioned in Section 2.1.6.2, the Summer School was also promoted through the website of the Department of Informatics, AUTH. Posts in the Department's Facebook account were also created (Figure 29).

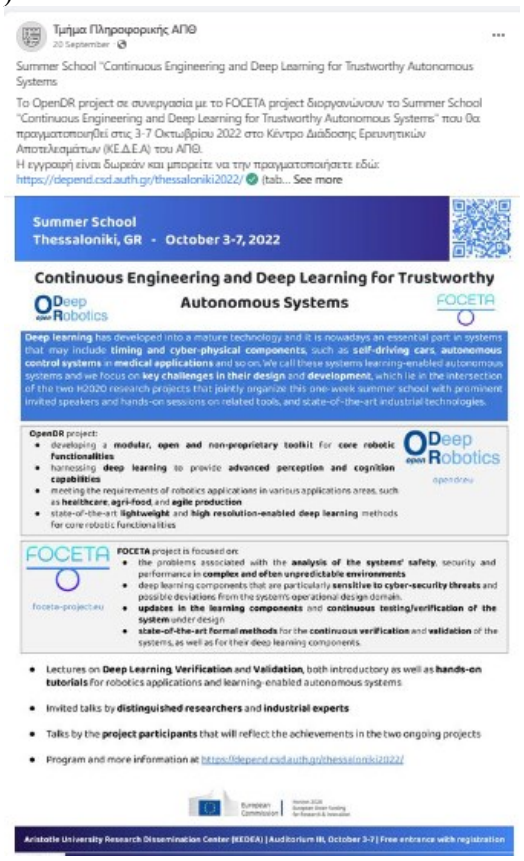

Figure 29: The post at the Department of Informatics (AUTH) Facebook account regarding the Summer School

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

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

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 in a number of fora, mainly thanks to PAL efforts:

PAL Robotics disseminated the project and its objectives in the European Robotics Forum (ERF2022) in Rotterdam, Netherlands (28-30rd of June). More specifically, the project activities were disseminated in the following events (Figure 30):

- in the workshop entitled "Robotics and AI & Data, The need of AI and Robotics" by Carlos Vivas, Business Manager.
- in the workshop entitled "Agile Production, Mobile Manipulators – the Gap between Reality and Expectation" by Sarah Terreri, Project Manager.
- in PAL´s booth.

TUD also disseminated the project and its objectives at ERF2022. TUD collected video contributions from the partners and displayed the generated video at its booth. The video can be found at this link.

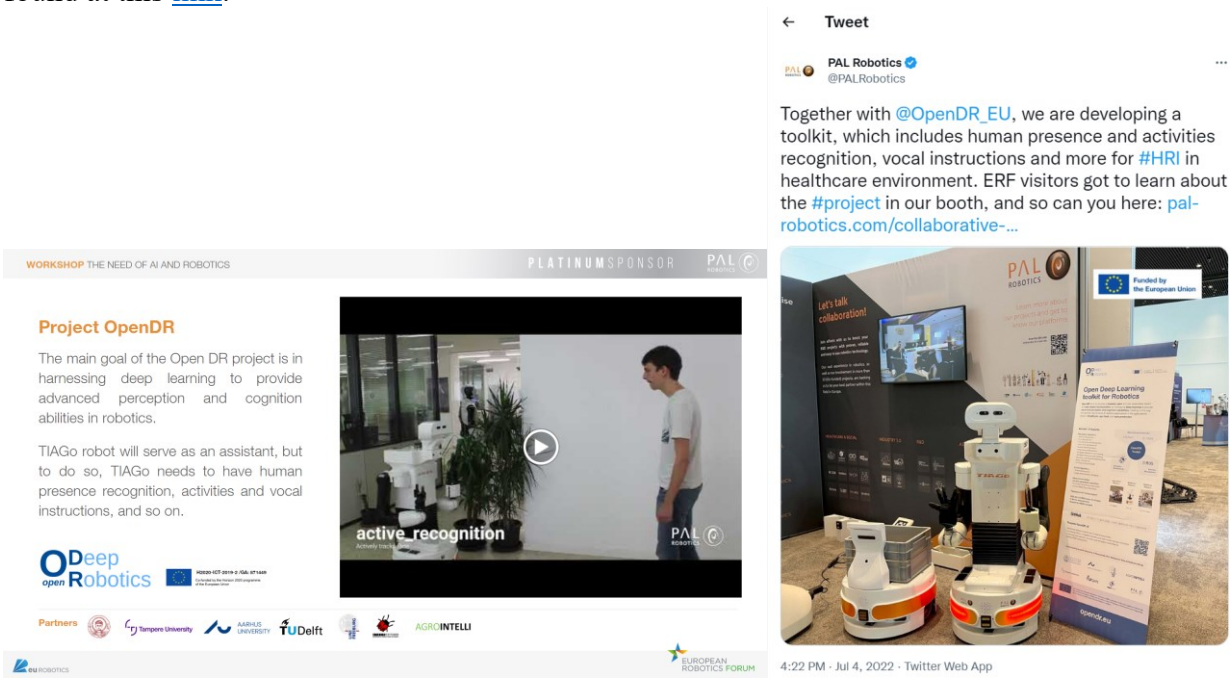

Figure 30: PAL booth and "Workshop The need of AI and Robotics" in ERF22

Moreover, PAL Robotics sponsored this year's (2022) HRI event, which took place in Barcelona, Spain in March and an OpenDR poster, created for the event, was displayed alongside posters for other collaborative projects where PAL participates (Figure 31).

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

Figure 31: Reception in the HRI22 event

Also, PAL disseminated the project in several additional events (Figure 32) such as ICRA2022, Philadelphia, USA (23-27 May); Automatica 2022, Munich Germany (21-24 June); XLIII Jornada Automatica 2022, La Rioja, Logroño, Spain (7-9 September); IROS2022, Kyoto, Japan (23-27 October).

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

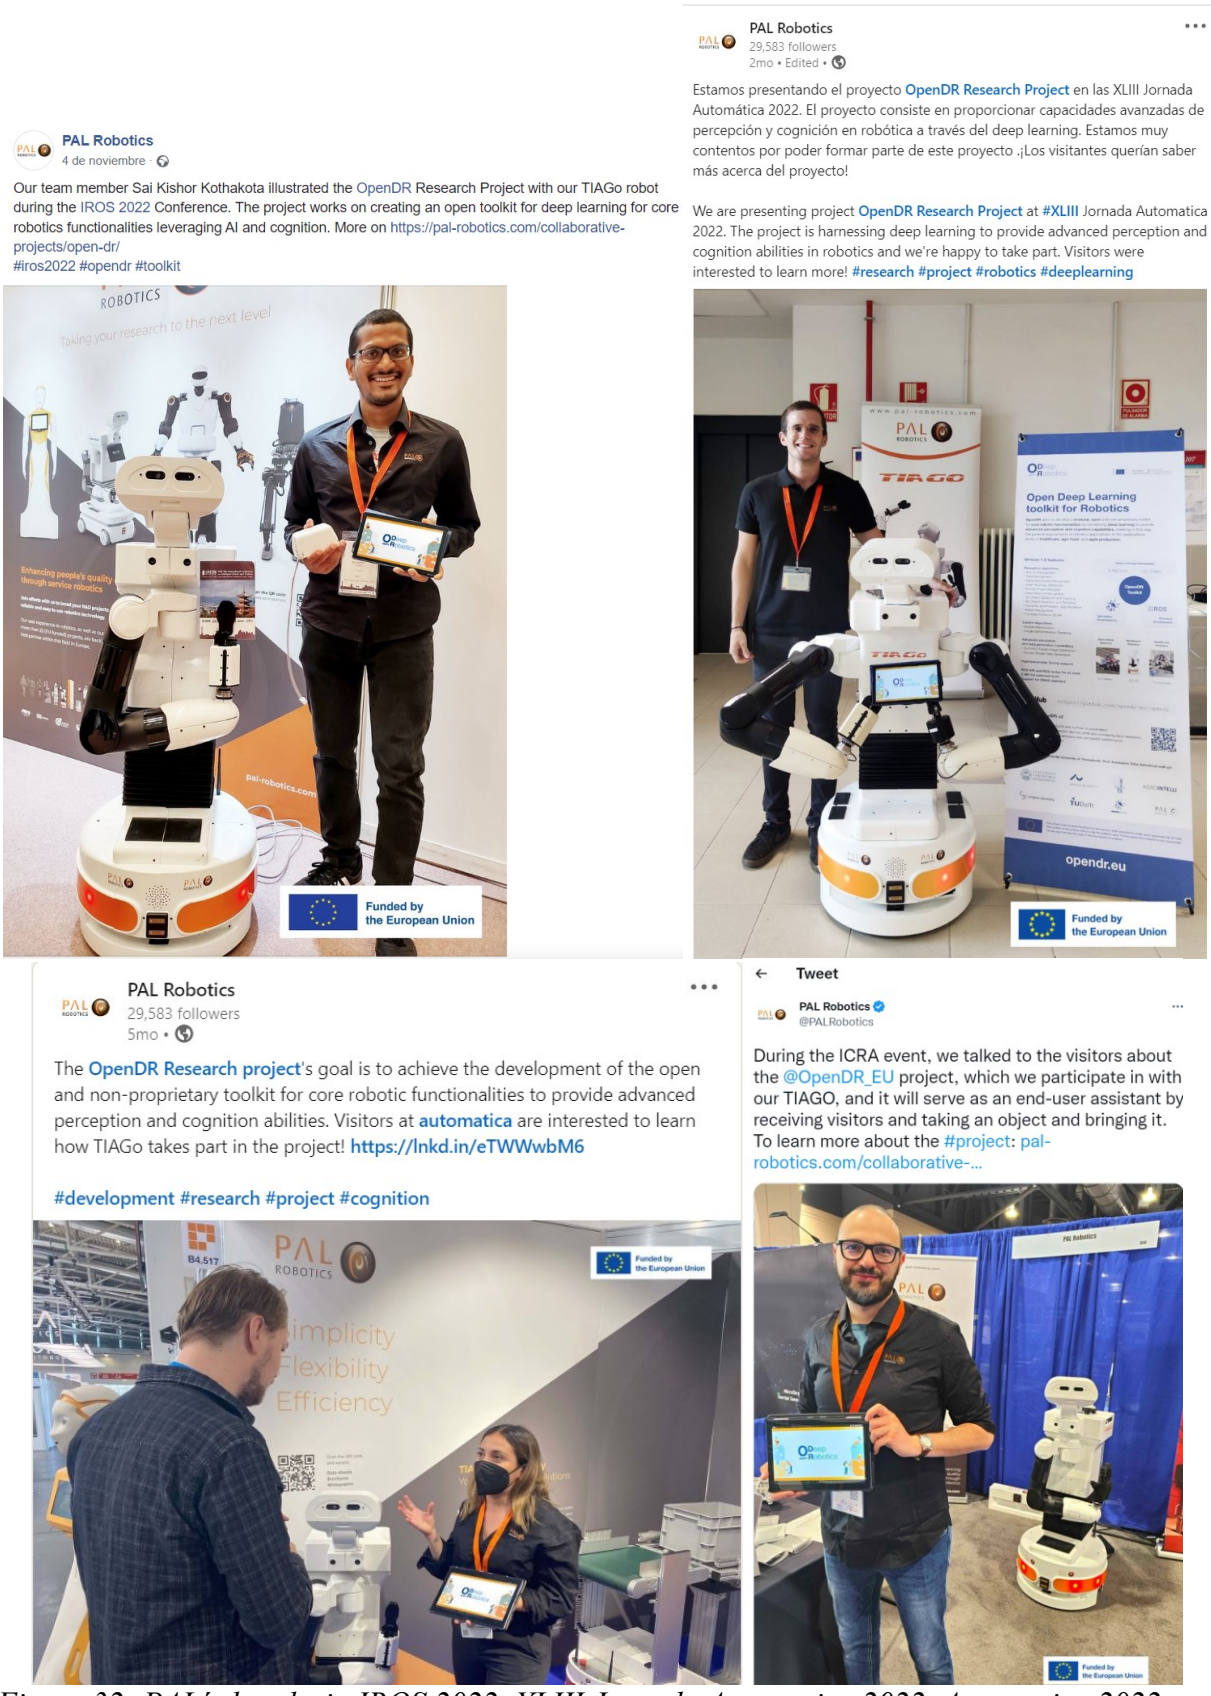

Figure 32: PAL´s booths in IROS 2022, XLIII Jornada Automatica 2022, Automatica 2022 and ICRA 2022

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

In addition, during the Humanoids 2022 conference, held in Okinawa, Japan in November, the CEO Francesco Ferro gave a presentation (Figure 33) in the workshop “Can we build Baymax?” and talked about OpenDR.

Figure 33: Talk of PAL’s CEO Francesco Ferro on Humanoids 2022 Conference

Finally, OpenDR was present in two events that focused on promoting the participation of women in science and engineering, since the consortium believes that women shall play a leading role in these fields and AI / DL in particular. Indeed Dr. Negar Heidari (AU) presented her OpenDR work “Attention-Augmented ST-GCN for Efficient Skeleton-based Human Action Recognition” at the NeurIPS 2022 Workshop Women in Machine Learning, after submitting an abstract which was peer reviewed prior to acceptance. Moreover, AUTH participated in the “4th Summit on Gender Equality in Computing” (GEC’22) which took place in June 2022 in Thessaloniki, Greece. The AUTH team presented two works, the first one entitled “Lightweight Deep Learning Methods for Robotics Applications Using the OpenDR Toolkit” and the second “Multilayer Online Self-Acquired Knowledge Distillation”. The presentations were supported by two posters (Figure 34).

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

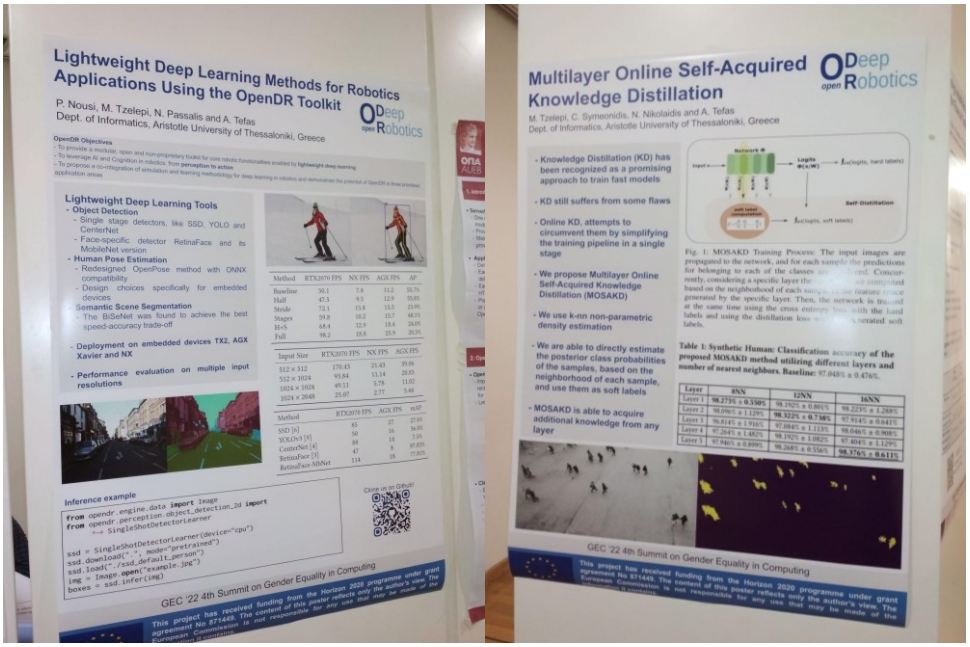

Figure 34: OpenDR at GEC'22

2.1.8 Collaboration with Other Projects and DIHs

Active engagement with DIH TRINITY on agile production was maintained in this period by regular communication and dissemination actions. OpenDR communicated to the DIH TRINITY network the existence of the toolkit and its capabilities and the plan of a second release by the end of year 3. In addition, collaboration between OpenDR and TRINITY included the plan (end of 2022) to upload the OpenDR Agile Production use case as a use case (Sensor-based Human-Robot Collaboration) and modules (Human detection, Human action recognition, object detection and robot grasping) to the Digital Access Point of TRINITY. OpenDR has also been invited to speak at a TRINITY proposed workshop at ERF2023.

For PAL the OpenDR project opened various possibilities of collaboration with projects it participates in, like the SHAPES, SPRING, SANDRO and ALMI projects where tools from the OpenDR toolkit are planned to be integrated in healthcare-related use cases of these projects. OpenDR research results are already used in an ALMI project demonstration in a real world scenario where the TIAGo robot assists an Alzheimer patient in the kitchen. The demo video will be publicly released at the beginning of 2023.

Finally, OpenDR collaborated with the H2020 MARVEL project where AU is participating. Indeed, as already mentioned, OpenDR was present in the workshop "The challenges of the extreme-scale multi-modal analytics applications" which was organised by MARVEL. Furthermore, the two projects will join forces for the co-organization of a Summer School on "Deep Learning for Autonomous Systems and Smart Cities" that will take place in Aarhus Denmark, May 23-26, 2023. OpenDR topics will be presented in the first two days of the School while MARVEL speakers will cover the remaining two days.

2.1.9 Participation into Contests and Challenges

Tools of the OpenDR toolkit have been utilized by a student team from Tampere University, participating in the ADAPT 1st Field Campaign, organized by the H2020 project METRICS. During the campaign, participants were able to evaluate the performance of their robot system and algorithms in the field of agile production, with respect to assembly and human-robot

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

collaboration. ADAPT benchmarks evaluate robot perception and actions (e.g., object detection, pose estimation, grasping and manipulation and collaborative robot programming). Objects used were 3D printed gear assemblies to be manipulated by any robot and any gripper. The ADAPT 1st field campaign competition took place from May 6 to June 22 in the form of a video submission track and a live demo track. The challenge consisted of the assembly of a helical gear by collaborative programming with human and robot. The team achieved 2nd place for the video submission track and received the open-source award for their contributions.

2.1.10 OpenDR Toolkit Dissemination

The first public release of the OpenDR toolkit (V1.0, December 2021) through GitHub (Figures 36, 37) was disseminated through the project website, the 2nd newsletter, as well as through the most popular social media platforms, namely LinkedIn, Twitter and Facebook. Also, as mentioned in Section 2.2.1, an "OpenDR Toolkit" button that leads directly to the toolkit was added in the main menu bar of the project website. Apart from the first announcement, a significant number of posts and tweets have been published in the aforementioned accounts (Figure 35). Furthermore, in order to perform targeted dissemination of the toolkit release to the scientific community, AUTH announced the release of the toolkit through email lists, such as euRobotics and CVML. The toolkit was also announced through social media accounts maintained by the partners. Release of version 1.1.1 of the toolkit (June 2022) was also publicised through social media.

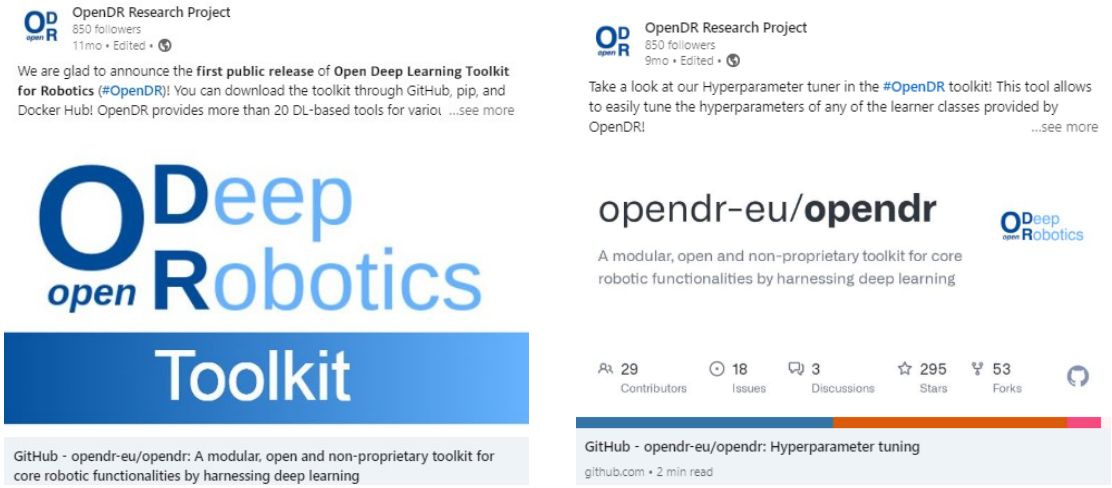

Figure 35: LinkedIn posts for the OpenDR toolkit release (left) and the Hyperparameter Tuning (right)

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

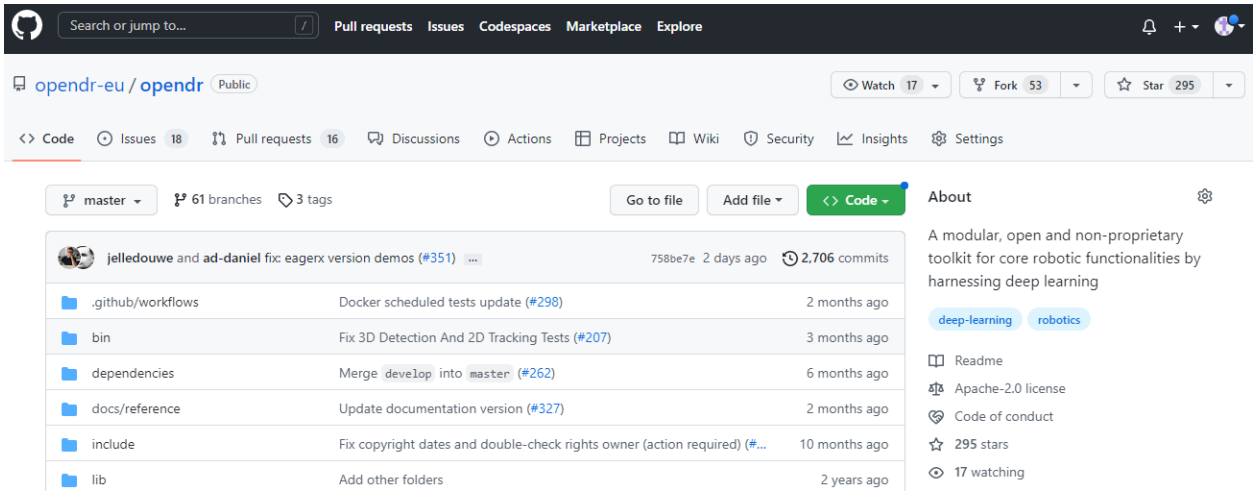

Figure 36: The OpenDR toolkit GitHub repository

Figure 37: A screenshot from the GitHub repository

Apart from online dissemination, the toolkit was also publicised in a series of events and talks (see previous Sections) such as the IROS tutorial, the Summer School, ERF 2022, the EURASIP JIVP webinar etc. The dissemination efforts were effective: the toolkit repository has currently more than 300 stars and is watched by 17 persons. More details are provided in Section 2.2.4.

2.2 Dissemination Performance Metrics

AUTH is continuously monitoring the performance of dissemination activities. So far, the results are extremely good, as detailed in the following subsections.

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

2.2.1 Project Website

Until today, almost all goals that had been set at the beginning of the project regarding the website, have been achieved. This fact indicates that all the dissemination actions that are accomplished through this communication channel have a very effective impact. For example, the number of website visitors up to M36 exceeds 10800, a very good figure indeed.

KPI	Target (for M48, unless stated otherwise)	M12	M24	M36
Number of visitors	1000 (by M12)	More than 3200	More than 5600	More than 11000
Average duration of visits:	2:00	1:36	1:46	1:15
Number of posts published (including news):	>50	15	56	66
Number of documents available in the website repository:	>50	18	41	103

2.2.2 Social Media

The presence of OpenDR in the most popular social media accounts could be characterised as extremely successful, as all target values of the key performance indicators defined in DoA have been already accomplished.

2.2.2.1 Twitter

Regarding Twitter, the tweets and retweets are already 235, while the account's followers are more than 420, up by almost 90% compared to M24 (the M48 target was set to 100). These metrics indicate the successful use of this social medium.

KPI	Target (M48)	M12	M24	M36
Twitter followers	>100	151	230	427
Number of tweets/retweets	>150	89	151	235

2.2.2.2 Facebook, YouTube

The Facebook account of the project was created in order to support the dissemination effort and present the project's results to an even wider target group. At the moment, it counts more than 345 likes/followers and more than 200 posts (+75 since M24). The YouTube account of the project has gathered so far almost 3569 views (+1569 since M24) and 62 subscribers, while

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

the target of the promotional video views according to the KPIs (500 views) has been achieved (almost 1600 views until M36).

2.2.2.3 LinkedIn

The LinkedIn group is in a very satisfactory status, as the target of the 150 group members (by the end of the project) has already been achieved. Indeed, the group counts so far 160 members and more than 180 posts. The LinkedIn page which has been created in order to maximise the awareness of the project and its results, has gathered until today more than 850 followers and counts more than 190 posts, related mainly to the project's results and news. The combination of these two information hubs definitely increases awareness for the project within the professional community.

KPI	Target (M48)	M12	M24	M36
LinkedIn group members	>150	111	132	160

2.2.3 Publications

As is the previous year, during the period M24-M36 the number of the publications have exceeded the expected target. Indeed, 18 new conference papers and 6 new journal papers have been added to the publications' list, thus, the project counts 74 publications in total, a number that shows its high level of scientific research.

KPI	Target	M12	M24	M36
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)

2.2.4 GitHub

The OpenDR toolkit was released for the first time on M24 through a GitHub repository. Since then, a considerable number of stars (>300), downloads and forks (>50) have been observed, which shows significant interest by the scientific community for the content of the toolkit. Focusing on downloads, the toolkit recorded an impressive 1100 unique GitHub clones (originated from users who wanted to have full access to all the capabilities provided by the toolkit) in just one year. Indeed, the total number of clones exceeded 45000, which however includes those from the continuous integration. Furthermore, an estimated 1700 pulls have been performed for the ready to use docker images (which mainly targets less experienced developers who want to directly try the toolkit) and an estimated 5000 pip downloads (originated from users who use selectively parts of the Python API of the toolkit, e.g., individual tools without installing the whole toolkit). The larger number of pip downloads indicates that the target group is familiar with Python and prefers to install only the parts of the toolkit that are relevant to their needs. Note that the statistics for docker and pip have been adjusted to remove traffic that might have been automatically generated, e.g. by the CI system. Based on these statistics, OpenDR consortium estimates that at least 7500 downloads from developers have been performed in this period, exceeding by far the corresponding M48 KPI target (a total of 500 downloads for the toolkit).

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

KPI	Target (M48)	M36
Number of stars	>100	303
Number of Downloads	>500	>7500
Number of forks	>30	54

2.2.5 Dissemination Material

During the period M24-M36 there was a significant increase in most of the KPIs which are related to the promotional (dissemination) material, namely the number e-newsletter dissemination list recipients, the promotional video views, events where the poster, factsheet, brochure and roll up banner were used, the number of brochures and the distributed copies. Obviously, the improvement of the COVID19 situation has enabled participation in more in-person events, such as conferences, workshops, summer schools etc. Thus, despite the hurdles imposed by the pandemic, it is possible that by the end of the project (M48), the consortium will have achieved most of the targets which have been set in the DoA with respect to dissemination material. In any case the consortium believes that dissemination through electronic means (website, social media etc) is much more effective than "traditional" dissemination methods (posters, brochures etc).

KPI	Target (M48)	M12	M24	M36
Number of e-newsletters published	4	1	2	3
Size of the dissemination list	>500	>1000 (estimate)	>3000 (estimate)	>5000 (estimate)
Number of press releases	3	1	1	1
Number of videos to be produced	2	1	1	1
Number of video views	500	529	911	1622
Number of events where the factsheet is used	10	1	1	4
Number of events where the poster is used	10	1	2	6
Number of brochures	2	1	1	2
Number of events	10	-	-	4

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

where the brochure is used				
Number of copies distributed (brochures)	800 hardcopies + eBrochures	-	-	550
Number of events where a roll-up banner is used	4	0	0	1

2.3 Conclusions and Future Plans

The very large number of successfully organized dissemination activities (especially in what concerns publications, 74 since the project start excluding preprints) that took place within the third year of the project, the very good values that were achieved for the KPIs and the fact that, for the great majority of them, the M48 targets have already been reached and exceeded, signifies that OpenDR has made important progress and continues to be in a very good shape, despite the negative effect of COVID-19 especially in Y1 and Y2. As expected, the effects of the pandemic are still visible in KPIs that involve physical events (e.g., number of distributed brochures or number of events where the poster has been used). However, the project is catching up in this aspect of dissemination. Furthermore, the impact of these shortcomings in the effectiveness of the overall dissemination strategy is rather minimal. Based on the above, the consortium feels that there is no need to adapt/modify its dissemination strategy. It will however continue to monitor its performance and implement corrective measures whenever needed, in the final year.

The release of the first public version of the toolkit in M24 gave rise to a new series of activities that focused on attracting relevant parties in getting to know and use it. The relevant figures show that this has been achieved in a very satisfactory level, but efforts will continue in the final year with e.g., additional training events such as webinars, online courses and tutorials, including the seasonal school that will be organized in Denmark in May 2023, in collaboration with H2020 project MARVEL. The consortium will also continue to monitor the adoption and visibility of the toolkit in GitHub.

3. Current Exploitation Plan

The OpenDR toolkit is a modular, generic, and open library that includes state-of-the art deep learning approaches that can be used to common robotic scenarios, both simulated and real. The industrial partners (SME's) have plans to exploit the potential of this library, PAL and AGI towards enhancing their robotic platforms and CYB towards providing an ecosystem in which to develop and disseminate robotic solutions relying on deep learning.

AGI intends to rely on the OpenDR tools to ameliorate reliability, efficiency and safety of the Robotti. Indeed, it plans to exploit a system for increased Robotti safety, the plant row guidance system and a plant identification tool, all of which utilize OpenDR tools.

PAL's exploitation plan focuses mainly on the commercialization of assistive robots. In this context, OpenDR tools are expected to help coping with the difficulties intrinsic to an unpredictable healthcare environment.

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

CYB's plan is to demonstrate the capabilities of the Webots robot simulator (including the web-based version) in the deep learning field, with the expectation of increasing visibility to its product and therefore bring requests for additional paid services.

In the following sections, each SME partner describes and refines their exploitation plan, as it stands in the current phase of the project.

3.1. CYB Exploitation Plan

CYB's exploitation plans remain unchanged. CYB plans to rely on the OpenDR toolkit and the scenarios developed in the project in order to demonstrate the capabilities of Webots robotics simulator in the Deep Learning field. Interest in this field remains high both in the industrial and academic spheres, making it very appealing as a market both for its future potential, but also to increase the visibility of Webots as a platform in this field.

In the last year special effort was undertaken in order to have simulations running on the web, this being a domain in which CYB was previously absent. Webots and all other technologies developed by CYB around it are open source so the primary source of revenue comes from the services built on top of these products. CYB believes the web as a platform is necessary for its growth as it has an important potential to expose Webots and its capabilities to a wider audience, while at the same time growing a community around its product. A larger user base and increased visibility is expected to yield more opportunities for collaborations with industrial or academic partners that might require support in the implementation of these tools for their own use-case. More specifically, requests for the following services are expected to increase:

- Requests of new features in the simulator.
- Support in the creation of educational curricula centred around robotics and deep learning.
- Organisation of international competitions or benchmarks.
- Requests to host servers where web simulations can be run on.

The necessary tools and infrastructure to achieve these objectives have now been implemented and therefore the primary focus in the coming year will be centred heavily around the dissemination of the results CYB and the other partners in the consortium have achieved.

3.2. PAL Exploitation Plan

The use of robots in non-industrial environments is set to continue increasing dramatically over the coming years. Creating assistive robots not only means combining several subsystems but also combining them into robots that can be adapted to follow the specific physical constraints that especially the older population is facing and the requirements of the care environment.

The PAL-led use case allows the execution of more tests for assistive robots like TIAGo in complex healthcare-like environments including the presence of humans, that is the basis of the targeted PAL roadmap for healthcare robots. Thanks to the first release of the toolkit in December 2021, PAL Robotics:

- Integrated the toolkit into TIAGo via Jetson-NY which has more computing power and allows integrating multiple tools at the same time and running multiple AI solutions, in order to use them in the healthcare use case.
- Carried out several tests to ensure the toolkit is working properly.

PAL Robotics continued spreading the progress of the project to other business units and managers within the company so as to show the case studies and the methodology applied within the project. The company has already discussed the adaptation of the OpenDR toolkit to match the

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

format of the ROS4HRI framework in order to facilitate its use for human-robot interaction (HRI) in ROS. By doing so, the community of AI and Robotics could further explore the project's results.

On the other hand, as discussed in D9.6, with the first public release of the OpenDR toolkit set in December 2021, PAL had planned to integrate the ROS ready software in other collaborative projects. Indeed the OpenDR toolkit was integrated in TIAGo robot in another collaborative project called ALMI. In this project, PAL has demonstrated the fall detection tool in order for TIAGo to detect a person on the floor and call for help in a real world healthcare scenario. Furthermore, PAL tested the toolkit in another social robot of the company named ARI. The results are planned to be exploited in the project SPRING where ARI is used for the elderly use case. This will happen by the end of the project.

Due to its vast network of European contacts, PAL can further collaborate with healthcare providers and care centres that can act as early adopters of the project's solutions and provide opportunities to extend the tests to different scenarios within the healthcare field or other potential fields of applications, so as to ensure full functionality of all technology features. This will contribute to the advancing of the TRL towards product finalisation and, later on, the broader adoption and market penetration of the solution.

Moreover, PAL would also like to continue the effort that it started in the field of assistive robots and focus on commercialising its robots, making their use easier in an unpredictable healthcare environment.

3.3. AGI Exploitation Plan

As the end of the project is approaching, the plans for exploitation are becoming clearer. AGI plans on exploiting:

1. Increased safety system
2. Plant row guidance system
3. Plant identification tool

For increasing safety, AGI will combine the OpenDR toolkit with its current remote viewing system. The OpenDR toolkit will be used as a safety helping function; when a human is detected, the remote viewer will be informed that there is a human seen by the FrontEye camera. As the technology is explored, it is expected that additional uses will emerge that can also be investigated and potentially exploited. In addition, the human pose detection and animal detection tools will be considered. After the systems are combined, AGI will be able to map where humans have been seen at a farm level. This will allow AGI to be able to make better safety system decisions. It is also possible to combine the lidar and the OpenDR toolkit to remove false positives, but at this time, it is out of the scope of the project.

The plant row guidance system will be tested in 2023 and improved for robustness. There is still very much a need for this technology, as farmers wish to seed with their 6m wide precision seeders and then have the robot follow the crop rows. As soon as the technology is robust, validated and proven, AGI will exploit the plant row guidance system (Figure 38).

For the plant identification algorithms, as mentioned in the previous technical report, AGI was building a mapping system. We are commercialising the CropEye camera system in Jan 2023 (Project: SqM Farm, Fund: GUDP), capable of acquiring between 2 and 4 downward facing images of the farmer's field, then uploading or saving the images to disk. The farmer is then able to view the images using the Robotti portal. The image locations are shown, and the farmer can virtually walk through their field (named Virtual Field Walk). The CropEye and Virtual Field Walk have been shown at SIMA and Agromek agricultural trade shows in 2022.
The CropEye system and Virtual Field Walk is the base technology for the next step, in which the farmer will be able to view a map of their field that shows the crop and weed densities and

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

coverages. From these maps, they can gain insight into how well the individual plants are performing compared to their neighbours. In addition, the farmer will be able to see how many monocot/dicot weeds they have in the images.

Through exploration, we have learned that we can also use the plant identification algorithms on the front and implement cameras. By comparing the crop count in the front image with the crop count in the rear image at the same approximate position, we are able to analyse whether or not the mechanical weeder is destroying the crop. We are also able to analyse the images for weed density and count. If the count has not changed significantly then the weeder tines or shares are not adjusted correctly. In this way, we can monitor the performance of the mechanical weeder and inform the farmer when the performance drops below the threshold.

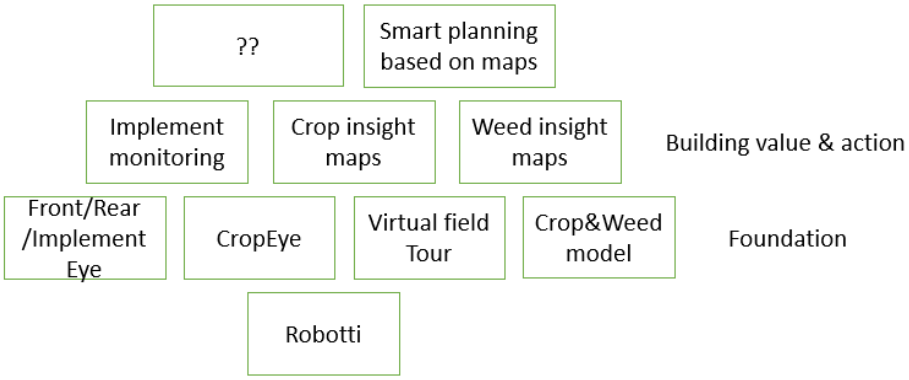

Figure 38: The technology building blocks that will exploit the plant identification tool.

4. Protection and Management of IPR

There are no changes in IPR protection and management policies during this year. The consortium has agreed previously on developing the OpenDR toolkit under the Apache 2.0 licence. With this permissive open source software licence, the tools could be used and sold by the industrial partners as part of their commercial products. However, Apache 2.0 licence 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.

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.

5. Third Project Newsletter

In accordance with the DoA, the consortium has created, under the coordination of AUTH, the third project Newsletter. The document, which provides a brief introduction to the project, its consortium and illustrates its major results during the third year can be found in Appendix I. The

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

newsletter will be posted at the project website, and announcements will be posted in the project social channels. The newsletter will be also distributed through relevant email lists such as euRobotics and CVML and email lists maintained by the partners.

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

Appendix I

3rd Project 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 3rd OpenDR Project Newsletter !

Newsletter 3.0: Toolkit progress, integration and evaluation and current status

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

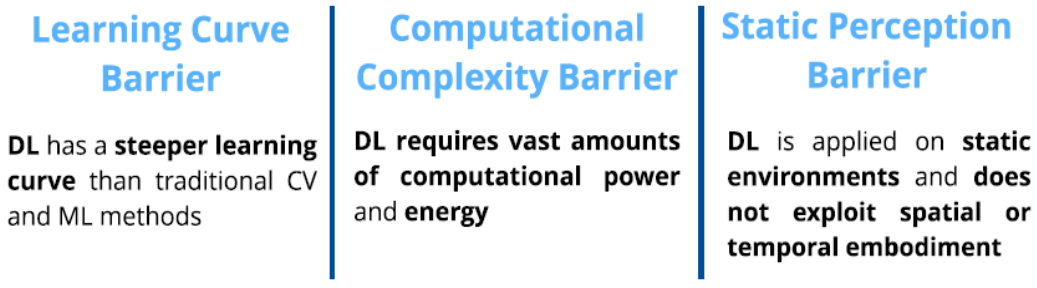

The need for an open deep learning toolkit that contains easy to train and deploy real-time, lightweight, Robot Operating System (ROS) compliant deep learning models for robotics is evident. This is where the **OpenDR** project enters.

What is OpenDR?

OpenDR "Open Deep Learning for Robotics Toolkit", is a EU 2020 Project which was launched on January 2020 and aims to develop a modular, open and non-proprietary toolkit for core robotic functionalities by harnessing deep learning to provide advanced perception and cognition capabilities, meeting in this way the general requirements of robotics applications in the areas of healthcare, agri-food and agile production. The **OpenDR** project is coordinated by the Aristotle University of Thessaloniki, Greece and will be running throughout December 2023 with a total budget of 6.6 Million Euros.

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

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

OpenDR Consortium

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

Aristotle University of Thessaloniki (AUTH) is the largest university in Greece, established in 1925. AUTH coordinates the project and leads the organization of dissemination activities. AUTH focuses its research on deep human centric active perception and cognition, where it contributes 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 leads the research in object detection/recognition and semantic scene segmentation and contributes to other areas such as evaluation and benchmarking activities of the project.

Tampere University (TAU) is Finland's second-largest university with 20.000 students and 330 professors. TAU participates with two labs/groups namely the Laboratory of Signal Processing at the Department of Computing Sciences and the Cognitive Robotics Group at the Department of Automation Technology and Mechanical Engineering. TAU leads the research in deep human centric active perception and cognition, working mainly on deep speech and biosignals analysis and recognition, and contributes 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 contributes 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 leads the research in deep environment active perception and cognition. ALU-FR focuses its research on Deep SLAM and 3D scene reconstruction, as well as on deep navigation. It also contributea on developing methodologies for deep planning.

Aarhus University (AU), Denmark participates in OpenDR with two groups, namely the Data-Driven Analytics Group and the Artificial Intelligence in Robotics Group, both belonging to the Section of Electrical and Computer Engineering. AU leads work on 2D/3D Object localization and tracking and works on sensor information fusion, as well as object detection/recognition and semantic scene segmentation and understanding. AU also contributes 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 leads/organizes the research activities on deep action and control, deep planning, as well as deep navigation. Furthermore, TUD leads and undertakes the research activities on human robot interaction. Finally, it leads and organizes 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 leads efforts of defining the toolkit's requirements and specifications. CYB also works on developing simulation environments and collecting data. Finally, it also leads 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 organizes and coordinates the toolkit integration, as well as the use cases integration activities. PAL also contributes on defining the healthcare robotics use case requirements and specifications and works on the integration of OpenDR Toolkit to this use case, as well as on its evaluation.

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

Work Performed in the 3rd Year

Toolkit Progress, Integration & Evaluation

The second version of the OpenDR toolkit has been officially released at the end of December 2022. **Version 2.0** of toolkit is more modular, allowing the users to install only the tools they need and includes several new tools like fall detection, SiamRPN 2D tracking, Continual Transformer Encoders, to cite some of them. Additionally, a ROS/ROS2 interface has been added for all tools and the documentation was enriched with instructions and benchmarks. The toolkit is accessible in **GitHub**, as well as through **pip** and **Docker Hub**. Reception of the toolkit from the robotics / deep learning /computer vision community was very good: so far, the GitHub repository was awarded more than 300 stars from its users, whereas the toolkit as a whole or individual tools have been downloaded more than 7500 times since its first release in December 2021. The OpenDR consortium is eager to receive your feedback, bug reports and suggestions for improvements!

Regarding integration, partners AUTH and PAL worked on the integration of the OpenDR toolkit in the different supported embedded platforms (NVIDIA Jetson TX2, AGX, and NX) focusing on the testing and integration of the toolkit in docker for all

the platforms to make it easy to use and install. PAL also focused on testing the ROS nodes and provided feedback to partners involved in tools creation towards improving them and making them user-friendly for ROS users. Furthermore, all partners developed and contributed ROS2 nodes, making the OpenDR toolkit ROS2-compliant and extending the range of supported configurations. AGI focused on evaluating the performance of the tools related to the agriculture use case and gathered developer's feedback for improving the interface and documentation of the toolkit.

Deep Human Centric Active Perception and Cognition

Throughout the third year of the project partners AUTH, AU and TAU have continued their work on human centric tools and algorithms, bringing the consortium even closer to the realization of the powerful, flexible and efficient robotics toolkit. Indeed, AUTH worked towards developing active perception models using a multitude of methods, including active perception for face recognition, high resolution pose estimation, as well non-maximum suppression suitable for person detection methods.

Heatmaps generated as part of the active efficient high resolution pose estimation approach proposed by AUTH.

The developed approaches include a DRL-based control approach for training agents that are able to identify and focus on task-relevant objects, i.e., humans, as well as issue appropriate control commands so as to acquire better results. AU has contributed novel methods for accelerating the online inference of both video- and skeleton-based human activity recognition networks, and also proposed and implemented Continual Inference Networks (CINs), which perform efficient step-by-step online processing. In the area of facial expression recognition, AU researchers proposed a mechanism for improving the generalization ability of the state-of-the-art models on unseen samples by learning diversified facial feature representations and encouraging the

learner to extract diverse spatial and channel-wise features. Finally, in multimodal perception, TAU has contributed a method for robust audiovisual emotion recognition. The method includes a new modality fusion approach based on self-attention. Additionally, a training approach, referred to as modality dropout, that increases robustness to incomplete or noisy data in one modality was introduced.

Deep Environment Active Perception and Cognition

During 2022, partners AU, AUTH, ALU-FR, TUD, AGI and TAU proposed a variety of methods that can be used for environment active perception and cognition. AU worked on designing approaches for deep learning models uncertainty estimation, introducing the Layer Ensembles method. AUTH developed a novel deep active object detection pipeline that provides active perception capabilities to existing object detectors by employing a separate planning network that regresses the rotation and translation that a robot should follow in order to increase the object detection confidence. Moreover, AU and AUTH joined forces to introduce a 3D single object tracking method called Voxel Pseudo Image Tracking (VPIT) that is suitable for real-time object tracking using embedded devices. Partner ALU-FR introduced a novel method for 3D multi-object tracking using a graph neural network called Batch3DMOT, which leverages multiple modalities including camera, LiDAR, and radar, and provides 3D tracks of multiple objects in a given scene, as well as a deep learning-based loop closure detection and point cloud registration approach for LiDAR-based SLAM systems. They also dealt with a new task called continual SLAM that combines lifelong SLAM with online domain adaptation to reflect challenges that occur when deploying a SLAM system to the open world and proposed a dual network architecture called CL-SLAM for this task.

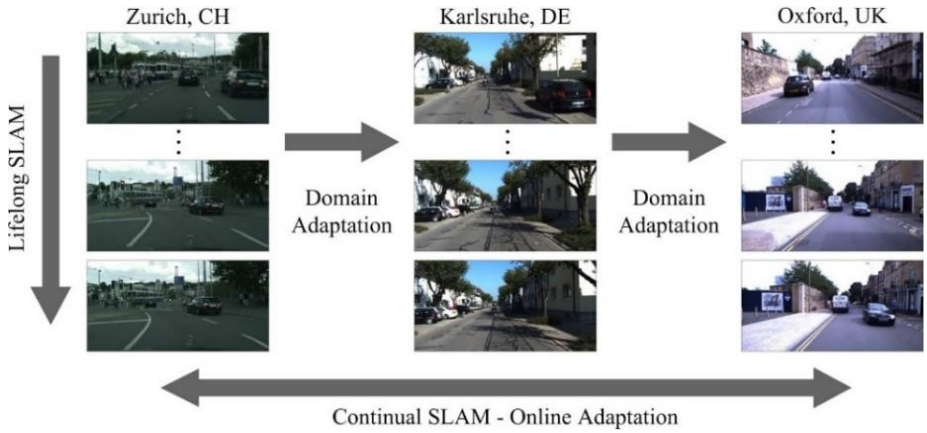

Continual SLAM: A new task that extends both lifelong SLAM and domain adaptation techniques.

Finally, TUD extended its previously proposed RSH data augmentation scheme for object detection in harsh lighting conditions towards polygon-shaped masked data augmentation in order to further improve robustness against a larger variety of lighting perturbations.

Deep Robot Action and Decision Making

In the area of deep robot action and decision making, partners AU, ALU-FR, TUD, TAU continued to work towards the design of novel navigation, planning, and control algorithms, thus contributing to the state of the art and to the key project objectives. AU presented an end-to-end planner trained with DRL for safe navigation in cluttered obstacle environments. The end-to-end planning algorithm is trained and tested in comprehensive simulations developed in Webots. The method was also successfully deployed in real-world indoor environments successfully and the corresponding experiments demonstrated that the proposed UAV planner trained solely with simulation can directly work in a real environment. On another front, ALU-FR developed a novel multi-object search approach that unifies short- and long-term reasoning in a single model. In addition, ALU-FR has developed an active localization method that combines differentiable particle filters with reinforcement learning to scale to large maps and continuous action spaces.

ALU-FR real world experiments on the PR2 (left) and HSR (right) robots.

TUD further improved its EAGERx toolkit, part of the OpenDR toolkit. The framework has now a consistent interface with an interactive GUI, unit tests with code coverage > 95%, and is accompanied by extensive documentation including a set of 10 interactive tutorials to make it easy for new users to get started. Moreover, TUD presented a delay simulation framework that allows delays to be accurately simulated in simulators that run faster than real- time. Furthermore, a novel communication protocol was proposed that reduces the effect of jitter on the sim2real performance. Finally, TUD started working on the PARTNR human-robot interaction algorithm, which learns to solve ambiguities in pick and place problems through interactive learning.

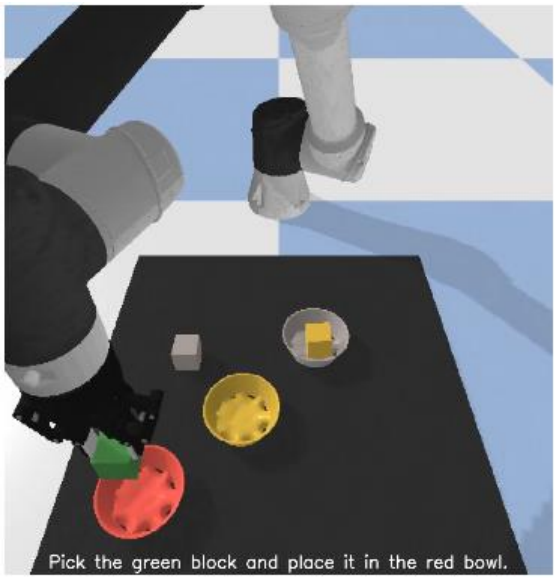

TUD Evaluation of PARTNR in a table-top pick and place task

TAU has made improvements to the SingleDemoGrasp tool. The tool now includes data annotation and augmentation functionalities as part of the toolkit, and the functionality to utilize different visual detection modules from the Detection 2D tool. Furthermore, TAU has developed a human-robot collaborative scenario that is inspired by the agile production use case. Within this scenario, perception tools (human skeleton detection, human action recognition detection & pose estimation of objects and targets) from the OpenDR toolkit are utilized to enable the robot to act as assistant to the human, by functi-onalities such as automated pick-and-place and robot to human hand-overs. TAU additionally provided basic HRC software templates, which can be utilized to replicate their experiments.

Simulation Environments and Data

Within the previous year, CYB extended the features of the Webots simulator by adding four new robots, several new assets to model industrial, healthcare and rural environments. The Camera node now supports either spherical or cylindrical projections for more realistic simulations and a new CadShape node that allows to

easily import 3D models in Collada (.dae) or Wavefront (.obj) format was introduced.

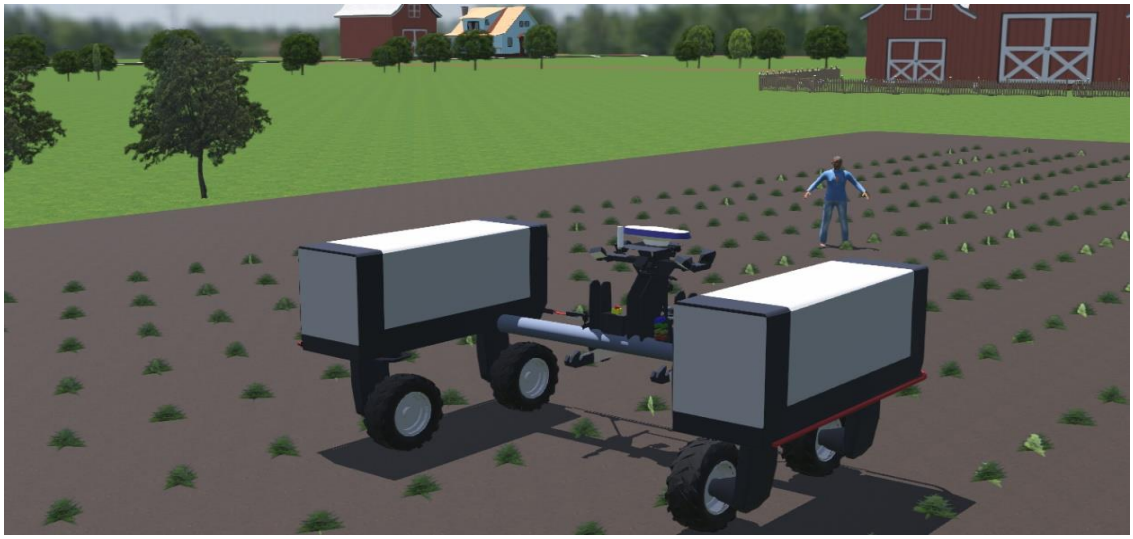

Robotti on a field

Improvements were also made in the way <extern> controllers communicate with the simulator allowing controllers to run over TCP or in isolated docker containers, decoupling the need to run everything on the same machine. Furthermore, a new service https://webots.cloud was announced which allows to easily disseminate results in the form of immersive 3D animations or simulations without needing to install any software. Furthermore, AUTH continued working towards the generation of datasets and simulation environments related to OpenDR use-cases and tools. This work included the generation of 3D facial models, a simulation environment and an image sequences dataset suitable for active facial expression recognition methods, an annotated synthetic full-body /face dataset for human-centric active vision tasks, and a video data generation framework for UAV detection methods. In addition AUTH, along with AGI, created a dataset designed for the training and evaluation of the performance of person detection methods in settings encountered in the agriculture use case. Finally, AU continued to improve the capability of its planning algorithms in its Webots R2021a implemented pipeline by incorporating OpenDR perception tools (2D object detection algorithms). As a result, the safety of algorithms increased by detecting the human presence and stopping the operation in such case.

Dissemination

Numerous efforts, in various directions, were undertaken by the OpenDR 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.

The consortium managed to generate, for yet another year, a high volume of publications. Eighteen papers were presented or accepted in well-established international conferences (including ECCV, ICIP, ICPR, EUSIPCO etc.) and six papers were published or accepted in scientific journals (including IEEE RA-L and Neurocomputing). In total, project partners have published so far 74 conference and journal papers, one edited book and numerous preprints.

A particularly important dissemination activity was a tutorial on "Open and Trustworthy Deep Learning for Robotics" that was organized by the project at IROS 2022, Kyoto, Japan. The tutorial was well attended and included talks on open Deep Learning tools for robot perception navigation and control, as well as other topics, presented mainly from consortium members. Of equal importance was the organization, in collaboration with the FOCETA H2020 research project, of the 5-day "Continuous Engineering and Deep Learning for Trustworthy Autonomous Systems" Summer School, which was attended by more than 75 persons, Other significant dissemination activities included invited talks, talks in two tutorials (at ICRA and CCTA 2022) and participation, through workshop talks or posters / videos in partners' booths, in events like ERF 2022, IROS 2022 or ICRA 2022.

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