OpenDR
–
Open Deep Learning Toolkit for Robotics

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

Deliverable D9.6: 2nd report on dissemination, exploitation plans and project newsletter

Date of delivery: 31 Dec 2021

Contributing Partners: Aristotle University of Thessaloniki, 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:	**D9.6:2nd report on dissemination, exploitation plans and project newsletter**	
Project:	**OpenDR** (ICT-10-2019-2020 RIA)	
Nature:	Report	**Dissemination Level: PU**blic
Authors:	Anastasios Tefas, Nikos Nikolaidis, Eleni Chantzi (AUTH), Dias Daniel (CYB), Lukas Hedegaard Morsing, Halil Ibrahim Ugurlu, Illia Oleksiienko, Alexandros Iosifidis (AU), Moncef Gabbouj (TAU), Alea Scovill (AGI), Gizem Bozdemir, Thomas Peyrucain (PAL)	
Lead Beneficiary:	AUTH (Aristotle University of Thessaloniki)	
WP	9	
Doc ID:	OPENDR_D9.6.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	6/11/2021	Document outline created
1.0	15/12/2021	First complete draft, ready for internal review
2.0	23./12/2021	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

Executive Summary **4**

1. Introduction **5**

2. Dissemination and Communication **5**

2.1 Dissemination and Communication Activities *5*
2.1.1 Project Website 5
2.1.2 Social Media 7
2.1.3 OpenDR Webpages at Partners Websites 10
2.1.4 Dissemination Material (Brochures, Banners, Posters) 11
2.1.5 Publications 13
2.1.5.1 Journal Papers 13
2.1.5.2 Conference Papers 18
2.1.5.3 Preprints 24
2.1.5.4 Books 26
2.1.6 Special Sessions and Issues 26
2.1.7 Workshops, Invited / Keynote Talks 27
2.1.8 Participation to Tradeshows, Exhibitions, EU Events, Industry Workshops, Awards 31
2.1.9 Collaboration with Other Projects and DIHs 32

2.2 Dissemination Performance Metrics *33*
2.2.1 Project Website 33
2.2.2 Social Media 34
2.2.2.1 Twitter 34
2.2.2.2 Facebook, YouTube 34
2.2.2.3 LinkedIn 35
2.2.3 Publications 35
2.2.4 Promotional Material 35
2.2.5 Events 36

3. Current Exploitation Plan **37**

3.1 CYB Exploitation Plan *38*

3.2 PAL Exploitation Plan *39*

3.3 AGI Exploitation Plan *39*

3.4 OpenDR Participation in Innovation Radar *40*

3.5 Protection and Management of IPR *40*

4. Second Project Newsletter **41**

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 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, D9.6 "2nd report on dissemination, exploitation plans and project newsletter", is a document presenting the project dissemination and exploitation activities within the second year of its lifecycle. It contains information regarding scientific publications describing the project's technical achievements, project-related presentations at various events, keynote talks, workshops and special sessions, 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 OpenDR exploitation plan, for the SME partners. The 2nd project newsletter is also included in this deliverable. Overall, the consortium believes that the project progresses very 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. Similar to any collaborative project, dissemination and communication activities are very important for OpenDR, since they will enable 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 second year (Y2) of the project, namely activities in the interval M13-M24. It contains (Section 2) information regarding scientific publications, organization of a special session, a special issue 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 cite a few. Information regarding the performance of dissemination and communication activities in Y2, through a number of Key Performance Indicators (KPIs) and their target values defined in the DoA, is also presented.

Effective exploitation of the project results is also very important for a project, and especially for its industrial partners. The deliverable includes (Section 3) a description of the current exploitation plans of the three SMEs that are part of the project consortium. Inclusion of the foreseen project innovations in the EU Innovation Radar initiative is also described in the same section. Finally, the second project newsletter, an important instrument in the project's communication strategy, is also included in this deliverable (Section 4 and Appendix I).

The presented results and activities indicate that the project's progress in terms of both dissemination and exploitation planning is very good and make the consortium confident that the project outcomes will be both 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 Y2 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

The project website (https://opendr.eu/) has been kept up-to-date by AUTH, with contributions from consortium members. Indeed during this period the website has been updated with new content, such as new publications (Figure 1), the latest news and events (Figure 2), as well as brief articles related to the project that can be found in the website's new subpage called "OpenDR Briefs" (Figure 3). These briefs are regularly published (bi-monthly) by OpenDR researchers and are short and easy to understand articles regarding the research conducted in the

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

project or topics related to the project. Thus, they are targeted towards the general public. The project website counts until today more than 50 posts and more than 5600 visitors. A new page that describes and provides access to the project datasets has also been added.

20	Adaptive Inference for Face Recognition leveraging Deep Metric Learning-enabled Early Exit	N. Passalis and A. Tefas	29th European Signal Processing Conference, (EUSIPCO), 2021
21	Pseudo-Active Vision for Improving Deep Visual Perception	N. Passalis and A. Tefas	IEEE International Conference on Image Processing (ICIP), 2021
22	Temporal Difference Rewards for End-to-end Vision-based Active Robot Tracking using Deep Reinforcement Learning	P. Tiritiris, N. Passalis and A. Tefas	International Conference on Emerging Techniques in Computational Intelligence (ICETCI), 2021
23	Efficient Training of Lightweight Neural Networks Using Online Self-Acquired Knowledge Distillation	M. Tzelepi and A. Tefas	IEEE International Conference on Multimedia and Expo (ICME), 2021
24	Supervised Domain Adaptation using Graph Embedding	Morsing, Lukas Hedegaard; Sheikh-Omar, Omar Ali; Iosifidis, Alexandros	25th International Conference on Pattern Recognition
25	Speech Command Recognition in Computationally Constrained Environments with a Quadratic Self-organized Operational Layer	Mohammad Soltanian, Junaid Malik, Jenni Raitoharju, Alexandros Iosifidis, Serkan Kiranyaz, Moncef Gabbouj	International Joint Conference on Neural Networks
26	Analysis of Voxel-Based 3D Object Detection Methods Efficiency for Real-Time Embedded Systems	Illia Oleksiienko and Alexandros Iosifidis	IEEE International Conference on Emerging Techniques in Computational Intelligence
27	DeepKoCo: Efficient latent planning with a robust Koopman representation	van der Heijden, Bas; Ferranti, Laura; Kober, Jens; Babuska, Robert	IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)

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

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

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

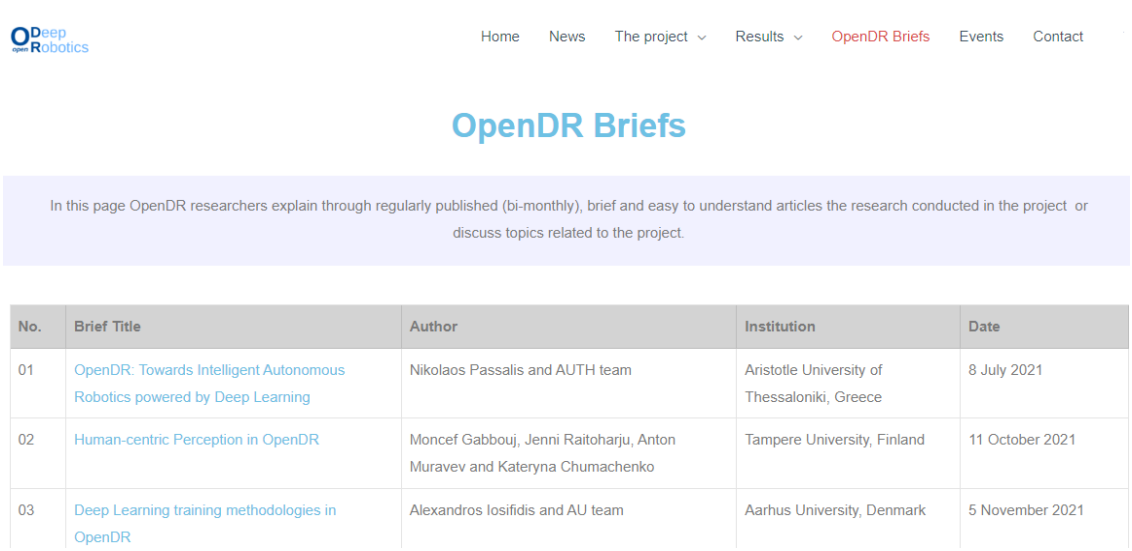

No.	Brief Title	Author	Institution	Date
01	OpenDR: Towards Intelligent Autonomous Robotics powered by Deep Learning	Nikolaos Passalis and AUTH team	Aristotle University of Thessaloniki, Greece	8 July 2021
02	Human-centric Perception in OpenDR	Moncef Gabbouj, Jenni Raitoharju, Anton Muravev and Kateryna Chumachenko	Tampere University, Finland	11 October 2021
03	Deep Learning training methodologies in OpenDR	Alexandros Iosifidis and AU team	Aarhus University, Denmark	5 November 2021

Figure 3: The new subpage "OpenDR Briefs"

2.1.2 Social Media

The project's promotion through its social media accounts continued to play a leading role among other dissemination activities during the period M13-M24. A significant number of posts and tweets have been published during this period. More details are provided below.

Twitter (@ OpenDR_EU)

The twitter account (https://twitter.com/OpenDR_EU) kept progressing very well both with respect to the number of tweets which were published, but also in terms of new followers and retweets by other twitter followers. Until now the project's twitter profile counts more than 230 followers and more than 150 tweets and retweets. Figure 4 depicts examples of tweets which were published during the period M13-M24.

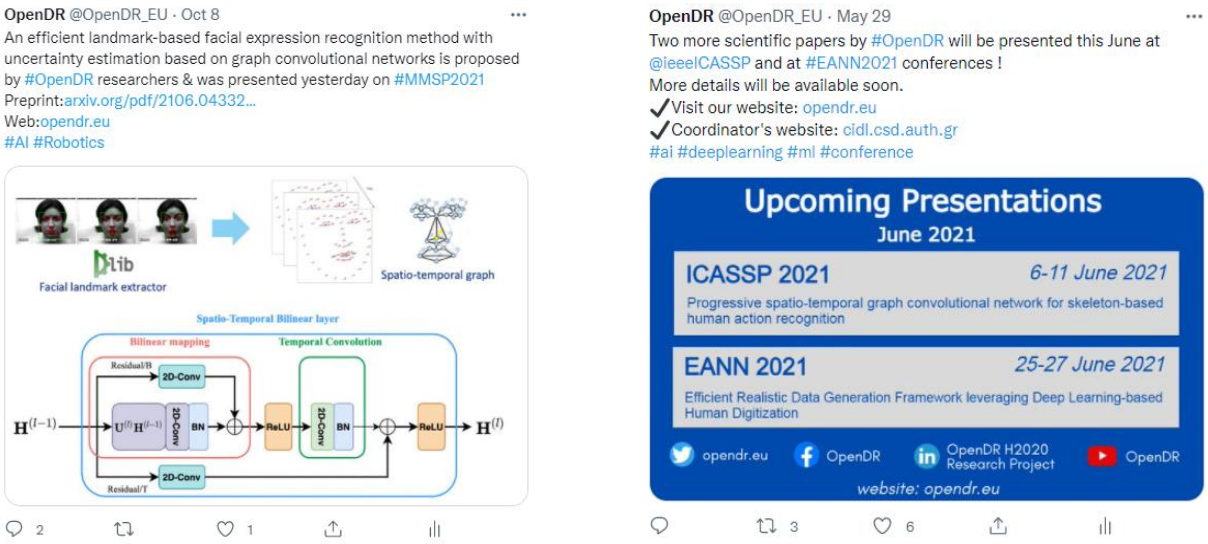

Figure 4: Examples of tweets published during the 2nd year

Facebook (@Opendr.eu)

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

The dissemination activity through the Facebook account (https://www.facebook.com/Opendr.eu) also continued in the second year of the project's life cycle. During this period the number of followers increased to 200, whereas close to 60 new posts were published. Most of the posts dealt with new journal and conference publications of the OpenDR team, but also with other project activities, for example the IROS Workshop (see the corresponding subsection) and the participation of OpenDR in the Google Summer of Code through the Deepbots supported project. (Figure 5)

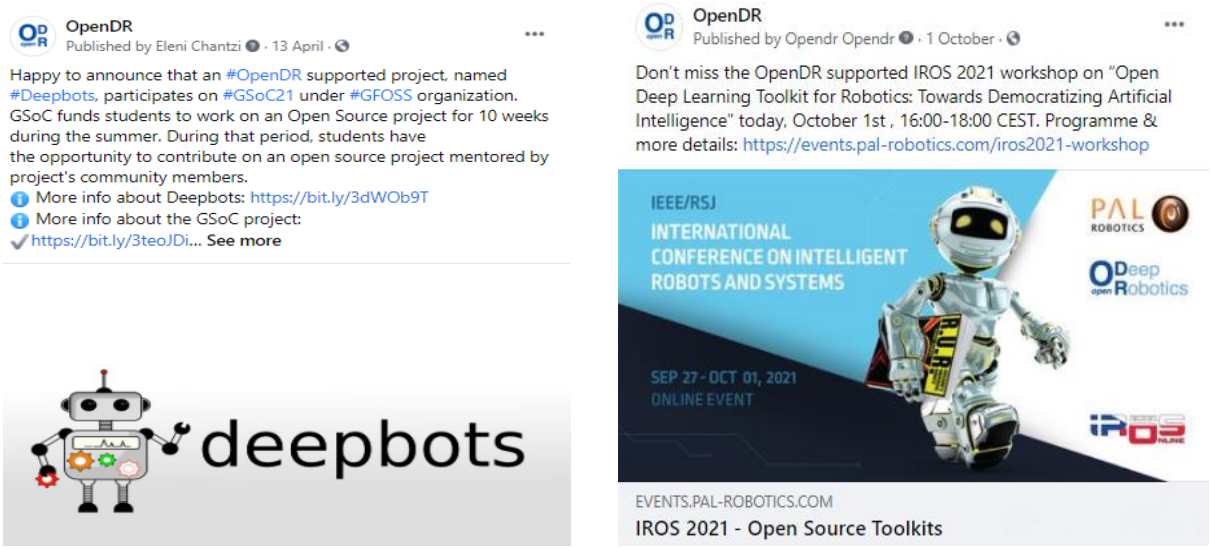

Figure 5: Examples of Facebook posts published during Y2

LinkedIn

The project continued to use the LinkedIn official page ("OpenDR Research Project") and LinkedIn group ("OpenDR H2020 Research project") to spread the news regarding its activities. The main content of the posts here was also related to the latest accepted journal and conference papers and the organised events. Until now the official page counts more than 245 followers, while the group members grew to 132. Examples of LinkedIn posts are displayed in Figure 6:

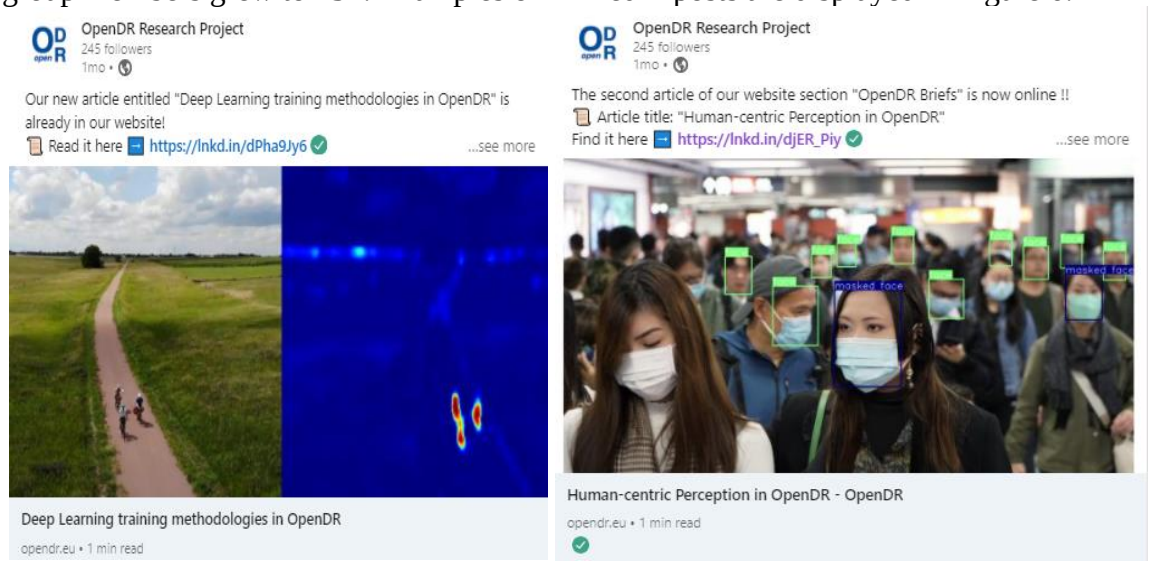

Figure 6: Examples of recent posts on LinkedIn page

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

YouTube

During the second year, OpenDR continued uploading new videos in its YouTube channel. So far, YouTube users can watch 12 videos regarding conference papers presentations as well as the project's promotional video. The channel counts 40 subscribers and more than 2000 views. It is worth mentioning that in the frameworks of the OpenDR supported IROS 2021 workshop "Open Deep Learning Toolkit for Robotics: Towards Democratising Artificial Intelligence" (see section 2.1.7), a new YouTube subchannel was created under the name "OpenDR-IROS Workshop", within the OpenDR main channel. Videos of the talks given by OpenDR partners in this workshop can be watched there (Figure 7).

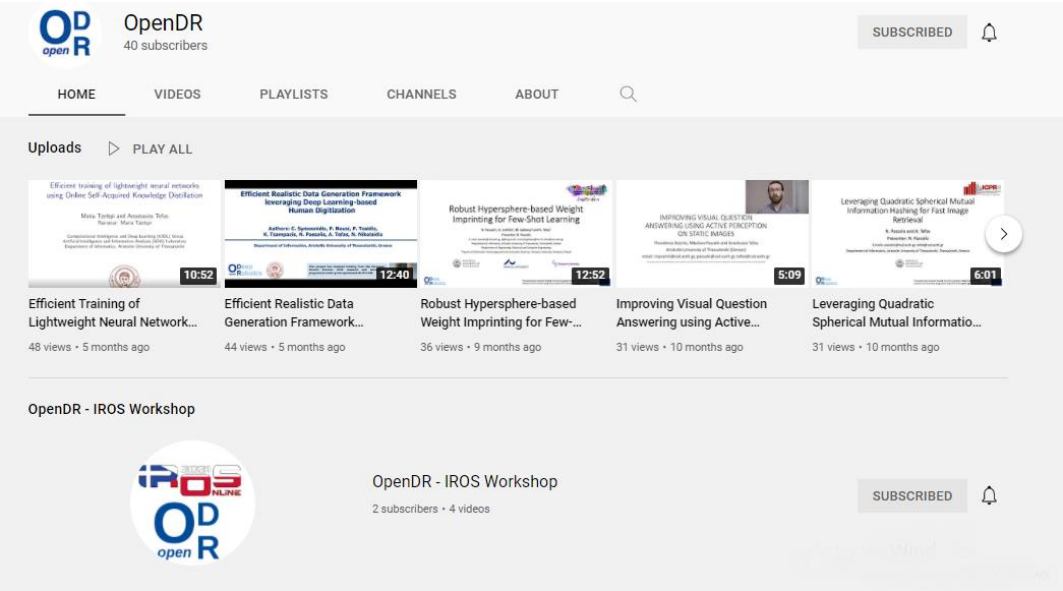

Figure 7: OpenDR and OpenDR-IROS Workshop YouTube channels

Partners social media activity

Apart from the activity in the "official" project channels mentioned above, OpenDR partners were also active in publicizing the project as well as related events and achievements. For example, partner PAL authored 14 OpenDR related posts in its Facebook account, 13 items in its LinkedIn account and posted 23 tweets in its Twitter corporate account from the start of the project (Figure 8).

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

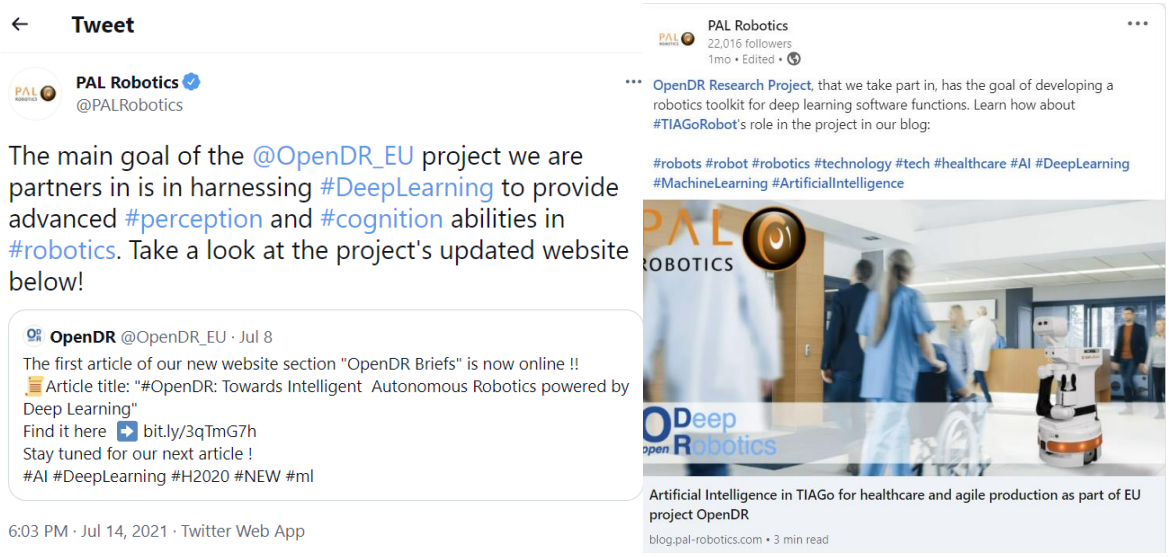

Figure 8: Examples of PAL Robotic posts regarding OpenDR

2.1.3 OpenDR Webpages at Partners Websites

During this period, OpenDR partners continued promoting the project through their websites. Indeed, PAL Robotics promoted the project on its website (Figure 9). The project is also mentioned in the Portfolio section of the CYB website, where a card is displayed describing the project and linking to the OpenDR website (Figure 10).

Figure 9: [OpenDR webpage] in PAL Robotics website

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

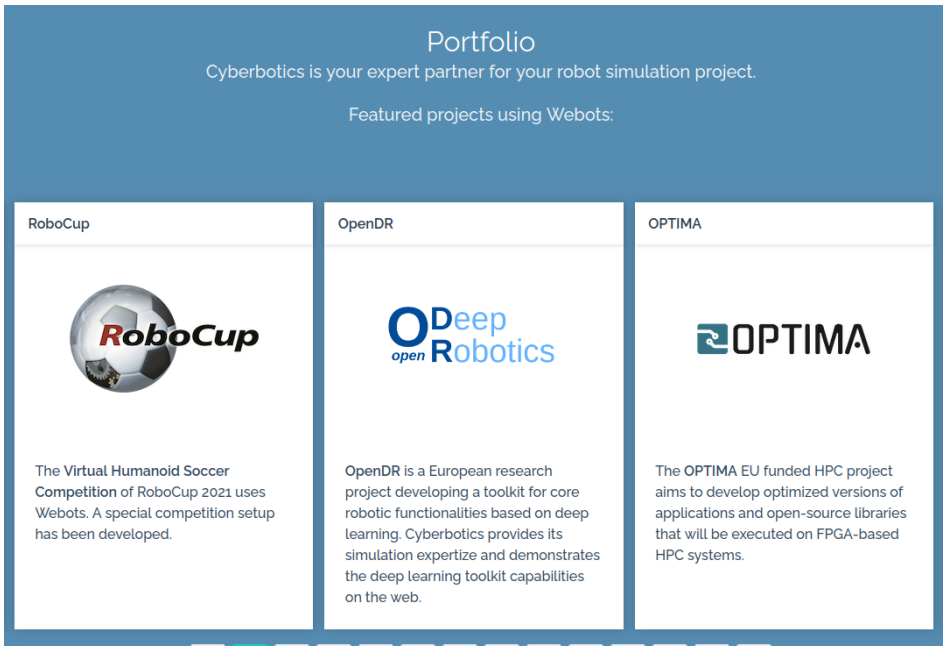

Fig 10: OpenDR description and link in the CYB website

2.1.4 Dissemination Material (Brochures, Banners, Posters)

In the framework of the project's dissemination activities, AUTH participated in the BEYOND 4.0, an Industry 4.0 exhibition/forum with participants from South Eastern Europe, the Mediterranean and the MENA region, which took place in Thessaloniki, Greece in October 2021. For this, a poster was created which contained all the fundamental information about the project, such as the objectives, the partners, the source of funding etc, along with infographics which made the poster more interesting (Figure 11).

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

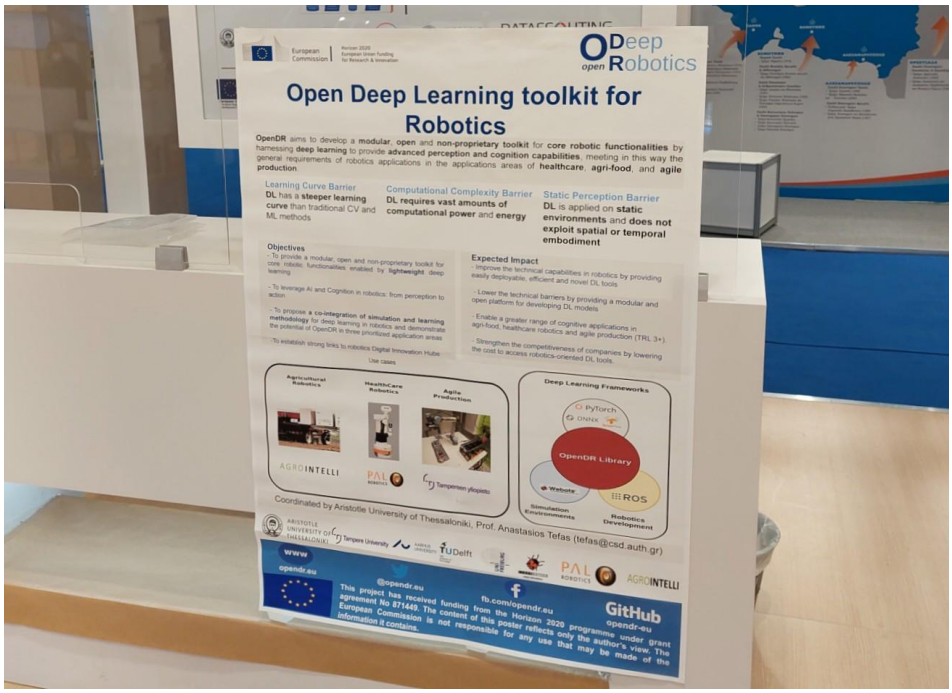

Figure 11: OpenDR poster for BEYOND 4.0 Expo

Moreover, in order to disseminate in the best way upcoming conference presentations, talks etc, OpenDR team often creates digital banners which are uploaded in the social media accounts as part of the relevant posts, in order to attract more users. Examples of such banners are shown in Figure 12.

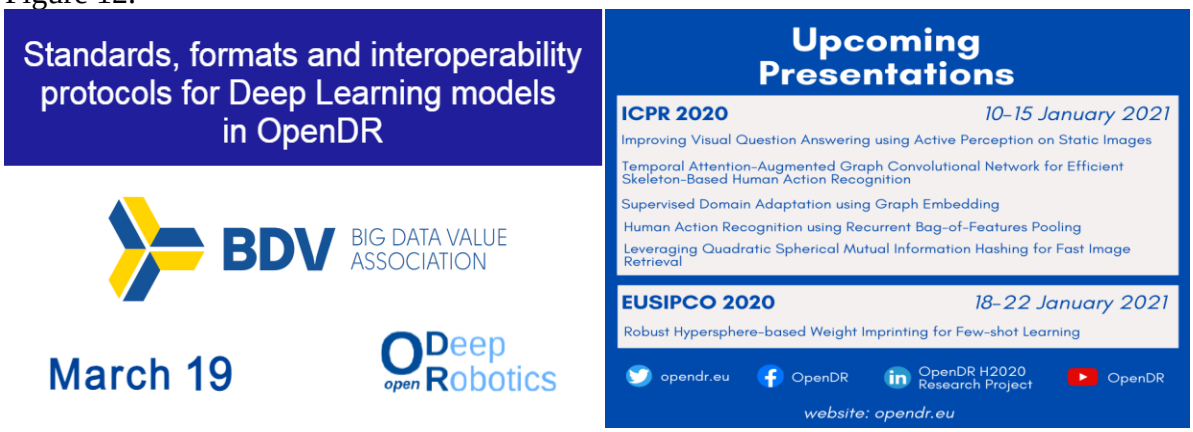

Figure 12: Banners for use in social media

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 twelve journal publications that have been accepted for publication or published. The full text of these publications can be found in the respective technical deliverables namely D3.2, D4.2, D5.2.

The first paper, a joint one by partners TAU and AU, was published in IEEE Internet of Things Journal:

Dat Thanh Tran, Moncef Gabbouj, Alexandros Iosifidis, "Remote Multilinear Compressive Learning with Adaptive Compression", IEEE Internet of Things Journal*, Early Access, 23 September 2021,* DOI 10.1109/JIOT.2021.3114743*, (impact factor 9.471).*

Multilinear Compressive Learning (MCL) is an efficient signal acquisition and learning paradigm for multidimensional signals. The level of signal compression affects the detection or classification performance of a MCL model, with higher compression rates often associated with lower inference accuracy. However, higher compression rates are more amenable to a wider range of applications, especially those that require low operating bandwidth and minimal energy consumption such as Internet-of-Things (IoT) applications. Many communication protocols provide support for adaptive data transmission to maximize the throughput and minimize energy consumption. By developing compressive sensing and learning models that can operate with an adaptive compression rate, we can maximize the informational content throughput of the whole application. In this paper, the authors proposed a novel optimization scheme that enables such a feature for MCL models. This proposal enables practical implementation of adaptive compressive signal acquisition and inference systems. Experimental results demonstrated that the proposed approach can significantly reduce the amount of computations required during the training phase of remote learning systems but also improve the informational content throughput via adaptive-rate sensing.

A paper describing work by AU researchers was published in IEEE Trans. on Image Processing:

Lukas Hedegaard, Omar Ali Sheikh-Omar and Alexandros Iosifidis, "Supervised Domain Adaptations: A Graph Embedding Perspective and a Rectified Experimental Protocol", IEEE Trans on Image Processing*, vol. 30, pp. 8619-8631,* DOI 10.1109/TIP.2021.3118978*, (impact factor 10.856).*

In this paper, the authors showed that Domain Adaptation methods using pairwise relationships between source and target domain data can be formulated as a Graph Embedding in which the domain labels are incorporated into the structure of the intrinsic and penalty graphs. Specifically, they analyse the loss functions of three existing state-of-the-art Supervised Domain Adaptation methods and demonstrate that they perform Graph Embedding. Moreover, the authors highlighted some generalisation and reproducibility issues related to the experimental setup commonly used to demonstrate the few-shot learning capabilities of these methods. To assess and compare Supervised Domain Adaptation methods accurately, they propose a rectified evaluation protocol, and report updated benchmarks on the standard datasets Office31 (Amazon, DSLR, and Webcam), Digits (MNIST, USPS, SVHN, and MNIST-M) and VisDA (Synthetic, Real).

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

A paper by AUTH was published in Pattern Recognition Letters:

Adamantios Zaras, Nikolaos Passalis, and Anastasios Tefas. "Improving knowledge distillation using unified ensembles of specialized teachers." Pattern Recognition Letters, vol 146, pp. 215-221, Jun. 2021, 10.1016/j.patrec.2021.03.014 (impact factor 3.756)

In this paper, the authors proposed a knowledge distillation method to enable transferring the knowledge between a very large network, called teacher and a smaller and faster one, called student. However, as recent evidence suggests, using powerful teachers often negatively impacts the effectiveness of the distillation process. The reasons behind this apparent limitation were studied and an approach that transfers the knowledge to smaller models more efficiently was proposed. To this end, multiple highly specialized teachers were employed, each one for a small set of skills, overcoming the aforementioned limitation, while also achieving high distillation efficiency by diversifying the ensemble. At the same time, the employed ensemble was formulated in a unified structure, making it possible to simultaneously train multiple models. The effectiveness of the proposed method was demonstrated using three different image datasets, leading to improved distillation performance, even when compared with powerful state-of-the-art ensemble-based distillation methods.

Another paper by AUTH was published in Signal Processing: Image Communication:

*Nikolaos Passalis, Anastasios Tefas "Deep supervised hashing using quadratic spherical mutual information for efficient image retrieval", S*ignal Processing: Image Communication, *vol 93, no 116146, pp. 1-10, April 2021, 10.1016/j.image.2021.116146 (impact factor 3.256)*

In this paper, AUTH proposed an efficient deep supervised hashing algorithm that optimizes the learned compact codes using an information-theoretic measure, the Quadratic Mutual Information (QMI). The proposed method is adapted to the needs of efficient image hashing and information retrieval leading to a novel information-theoretic measure, the Quadratic Spherical Mutual Information (QSMI), which can be especially useful in few-shot learning scenarios, as well as for mining data for training DL models. Apart from demonstrating the effectiveness of the proposed method under different scenarios and outperforming existing state-of-the-art image hashing techniques, this paper provides a structured way to model the process of information retrieval and develop novel methods adapted to the needs of different applications.

An Editorial of a Special Issue co-guest edited by E. Kayacan (AU) was published in IEEE Transactions on Artificial Intelligence:

W. He, Z. Li, E. Kayacan, C. Huang and M. V. Basin, "Guest Editorial: Special Issue on Artificial Intelligence for Robotics," in IEEE Transactions on Artificial Intelligence, vol. 2, no. 5, pp. 382-383, Oct. 2021, doi: 10.1109/TAI.2021.3119155. (impact factor N/A)

Topics covered in this special issue included but were not limited to AI for multirobot systems, AI in autonomous systems, AI in cognitive robotics, AI in human–robot interaction, AI-enabled robots to assist health professionals as in the case of COVID-19, applications of AI-enabled robots in industry, agriculture, logistics, medicine, augmented intelligence in robotics. The

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

editorial provided a concise foreword for the issue and included short descriptions of the 7 papers that were included in it, after peer review.

A paper by TUD was published in IEEE Robotics and Automation Letters:

Mazhar, O., Babuska, R., & Kober, J. (2021). GEM: Glare or Gloom, I Can Still See You - End-to-End Multi-Modal Object Detection. IEEE Robotics and Automation Letters, 6(4), 6321-6328. https://doi.org/10.1109/LRA.2021.3093871 (impact factor 6.54)

In this paper, the authors propose a multi-modal 2D object detector as well as deterministic and stochastic sensor-aware feature fusion strategies. The proposed fusion mechanisms are driven by the estimated sensor measurement reliability values/weights. Reliable object detection in harsh lighting conditions is essential for applications such as self-driving vehicles and human-robot interaction. The authors also propose a new 'r-blended' hybrid depth modality for RGB-D sensors. Through extensive experimentation, the authors show that the proposed strategies outperform the existing state-of-the-art methods on the FLIR-Thermal dataset, and obtain promising results on the SUNRGB-D dataset. Additionally, a new RGB-Infra indoor dataset is recorded, namely L515-Indoors, and it is demonstrated that the proposed object detection methodologies are highly effective for a variety of lighting conditions.

A paper by TUD was published in IEEE Robotics and Automation Letters:

J. Kulhánek, E. Derner, & R. Babuška (2021). Visual navigation in real-world indoor environments using end-to-end deep reinforcement learning. IEEE Robotics and Automation Letters, 6(3):4345–4352. https://doi.org/10.1109/LRA.2021.3068106, https://arxiv.org/pdf/2010.10903.pdf (impact factor 6.54)

Visual navigation is essential for many applications in robotics, from manipulation, through mobile robotics to automated driving. Deep reinforcement learning (DRL) provides an elegant map-free approach integrating image processing, localization, and planning in one module, which can be trained and therefore optimized for a given environment. However, to date, DRL-based visual navigation was validated exclusively in simulation, where the simulator provides information that is not available in the real world, e.g., the robot's position or segmentation masks. This precludes the use of the learned policy on a real robot. Therefore, the authors present a novel approach that enables a direct deployment of the trained policy on real robots. They designed a new powerful simulator capable of domain randomization. To facilitate the training, visual auxiliary tasks and a tailored reward scheme was proposed. The policy is fine-tuned on images collected from real-world environments. The method was evaluated on a mobile robot in a real office environment. The training took approximately 30 hours on a single GPU. In 30 navigation experiments, the robot reached a 0.3-meter neighbourhood of the goal in more than 86.7 % of cases. This result makes the proposed method directly applicable to tasks like mobile manipulation.

A paper by ALU-FR was published in IEEE Robotics and Automation Letters:

D. Honerkamp, T. Welschehold and A. Valada (2021). Learning Kinematic Feasibility for Mobile Manipulation Through Deep Reinforcement Learning. IEEE Robotics and Automation Letters, vol.

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

6, no. 4, pp. 6289-6296, Oct. 2021. https://doi.org/10.1109/LRA.2021.3092685, https://arxiv.org/abs/2101.05325 (impact factor 6.54)

Mobile manipulation tasks remain one of the critical challenges for the widespread adoption of autonomous robots in both service and industrial scenarios. While planning approaches are good at generating feasible whole-body robot trajectories, they struggle with dynamic environments as well as the incorporation of constraints given by the task and the environment. On the other hand, dynamic motion models in the action space struggle with generating kinematically feasible trajectories for mobile manipulation actions. In this paper ALU-FR proposed a deep reinforcement learning approach to learn feasible dynamic motions for a mobile base while the end-effector follows a trajectory in task space generated by an arbitrary system to fulfill the task at hand. This modular formulation has several benefits: it enables to readily transform a broad range of end-effector motions into mobile applications, it allows to use the kinematic feasibility of the end-effector trajectory as a dense reward signal and its modular formulation allows it to generalise to unseen end-effector motions at test time. The authors demonstrated the capabilities of our approach on multiple mobile robot platforms with different kinematic abilities and different types of wheeled platforms in extensive simulated as well as real-world experiments.

A paper previously reported as preprint by ALU-FR has now been published in International Journal of Computer Vision (IJCV):

Rohit Mohan and Abhinav Valada (2021). EfficientPS: Efficient Panoptic Segmentation. International Journal of Computer Vision (IJCV), vol. 129, no. 5, pp. 1551–1579, 2021. https://doi.org/10.1007/s11263-021-01445-z, https://arxiv.org/abs/2004.02307 (impact factor 11.541)

Understanding the scene in which an autonomous robot operates is critical for its competent functioning. Such scene comprehension necessitates recognizing instances of traffic participants along with general scene semantics which can be effectively addressed by the panoptic segmentation task. In this paper, the authors introduce the Efficient Panoptic Segmentation (EfficientPS) architecture that consists of a shared backbone which efficiently encodes and fuses semantically rich multi-scale features. They incorporate a new semantic head that aggregates fine and contextual features coherently and a new variant of Mask R-CNN as the instance head. They also propose a novel panoptic fusion module that congruously integrates the output logits from both the heads of our EfficientPS architecture to yield the final panoptic segmentation output. Additionally, they introduce the KITTI panoptic segmentation dataset that contains panoptic annotations for the popularly challenging KITTI benchmark. Extensive evaluations on Cityscapes, KITTI, Mapillary Vistas and Indian Driving Dataset demonstrate that the proposed architecture consistently sets the new state-of-the-art on all these four benchmarks while being the most efficient and fast panoptic segmentation architecture to date.

A paper by ALU-FR has was published in IEEE Transactions on Robotics (T-RO):

Kshitij Sirohi, Rohit Mohan, Daniel Büscher, Wolfram Burgard Abhinav Valada, EfficientLPS: Efficient LiDAR Panoptic Segmentation. IEEE Transactions on Robotics (T-RO), 2021. https://arxiv.org/abs/2102.08009 (impact factor 9.72)

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

Panoptic segmentation of point clouds is a crucial task that enables autonomous vehicles to comprehend their vicinity using their highly accurate and reliable LiDAR sensors. Existing top-down approaches tackle this problem by either combining independent task-specific networks or translating methods from the image domain ignoring the intricacies of LiDAR data and thus often resulting in sub-optimal performance. In this paper, ALU-FR authors present the novel top-down Efficient LiDAR Panoptic Segmentation (EfficientLPS) architecture that addresses multiple challenges in segmenting LiDAR point clouds including distance-dependent sparsity, severe occlusions, large scale-variations, and re-projection errors. EfficientLPS comprises a novel shared backbone that encodes with strengthened geometric transformation modeling capacity and aggregates semantically rich range-aware multi-scale features. It incorporates new scale-invariant semantic and instance segmentation heads along with the panoptic fusion module which is supervised by our proposed panoptic periphery loss function. Additionally, a regularized pseudo labeling framework is formulated to further improve the performance of EfficientLPS by training on unlabelled data. The proposed model is benchmarked on two large-scale LiDAR datasets: nuScenes, for which the authors also provide ground truth annotations, and SemanticKITTI. Notably, EfficientLPS sets the new state-of-the-art on both these datasets.

A paper by TAU has been published in Frontiers in Robotics and AI:

A. Angleraud, A. Mehman Sefat, M. Netzev and R. Pieters, Coordinating Shared Tasks in Human-Robot Collaboration by Commands, Frontiers in Robotics and AI, 2021. DOI: https://doi.org/10.3389/frobt.2021.734548 (impact factor: 4.331)

Human-robot collaboration is gaining more and more interest in industrial settings, as collaborative robots are considered safe and robot actions can be programmed easily by, for example, physical interaction. Despite this, robot programming mostly focuses on automated robot motions and interactive tasks or coordination between human and robot still requires additional developments. For example, the selection of which tasks or actions a robot should do next might not be known beforehand or might change at the last moment. Within a human-robot collaborative setting, the coordination of complex shared tasks, is therefore more suited to a human, where a robot would act upon requested commands. In this work, we explore the utilization of commands to coordinate a shared task between a human and a robot, in a shared work space. Based on a known set of higher-level actions (e.g., pick-and-placement, hand-over, kitting) and the commands that trigger them, both a speech-based and graphical command-based interface are developed to investigate its use. While speech-based interaction might be more intuitive for coordination, in industrial settings background sounds and noise might hinder its capabilities. The graphical command-based interface circumvents this, while still demonstrating the capabilities of coordination. The developed architecture follows a knowledge-based approach, where the actions available to the robot are checked at runtime whether they suit the task and the current state of the world. Experimental results on industrially relevant assembly, kitting and hand-over tasks in a laboratory setting demonstrate that graphical command-based and speech-based coordination with high-level commands is effective for collaboration between a human and a robot.

Another paper by TAU has been published in IEEE Access:

Anton Muravev, Jenni Raitoharju and Moncef Gabbouj, "Neural Architecture Search by Estimation of Network Structure Distributions," IEEE Access, vol. 9, pp. 15304-15319, 2021, DOI: 10.1109/ACCESS.2021.3052996 (impact factor 3.367)

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

The influence of deep learning is continuously expanding across different domains, and its new applications are ubiquitous. The question of neural network design thus increases in importance, as traditional empirical approaches are reaching their limits. Manual design of network architectures from scratch relies heavily on trial and error, while using existing pretrained models can introduce redundancies or vulnerabilities. Automated neural architecture design is able to overcome these problems, but the most successful algorithms operate on significantly constrained design spaces, assuming the target network to consist of identical repeating blocks. While such approach allows for faster search, it does so at the cost of expressivity. In this paper, the authors instead propose an alternative probabilistic representation of a whole neural network structure under the assumption of independence between layer types. Their matrix of probabilities is equivalent to the population of models, but allows for discovery of structural irregularities, while being simple to interpret and analyze. They construct an architecture search algorithm, inspired by the estimation of distribution algorithms, to take advantage of this representation. The probability matrix is tuned towards generating high-performance models by repeatedly sampling the architectures and evaluating the corresponding networks, while gradually increasing the model depth. The algorithm is shown to discover non-regular models which cannot be expressed via blocks, but are competitive both in accuracy and computational cost, while not utilizing complex dataflows or advanced training techniques, as well as remaining conceptually simple and highly extensible.

2.1.5.2 Conference Papers

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

A paper by AUTH was accepted and presented in the International Conference on Image Processing (ICIP):

Nikolaos Passalis, and Anastasios Tefas. "Pseudo-Active Vision For Improving Deep Visual Perception Through Neural Sensory Refinement." IEEE International Conference on Image Processing, Anchorage, USA, Sept. 2021, 10.1109/ICIP42928.2021.9506195

In this paper, the authors examine active vision approaches that hold the credentials for improving the accuracy of Deep Learning (DL) models for many challenging visual analysis tasks and varying environmental conditions. However, active vision approaches are typically closely tied to the underlying hardware, slowing down their adoption, while they typically increase the latency of perception systems, since sensory data must be recaptured. To this end, the authors propose a pseudo-active data refinement method that works by appropriately refining the sensory input, without having to reacquire the sensor data through traditional camera control approaches. The proposed method is fully differentiable and can be trained for the task at hand in an end-to-end fashion, while it can be directly deployed in a wide variety of systems, tasks and conditions. The effectiveness and robustness of the proposed method is demonstrated across a variety of tasks using two challenging datasets.

Another paper by AUTH was accepted and presented in the European Signal Processing Conference (EUSIPCO):

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

Nikolaos Passalis, Anastasios Tefas "Adaptive Inference for Face Recognition leveraging Deep Metric Learning-enabled Early Exits", European Signal Processing Conference (EUSIPCO), Dublin, Ireland, Aug. 2020, DOI N/A

In this paper, the authors worked on models that support adaptive computational graphs allowing for easily adapting the computations to the available resources by selecting the most appropriate computational path. However, such models are typically used in classification settings, e.g., using early exits, despite that DL models often aim at extracting representations (metric learning), e.g., face recognition. In this work, AUTH provided a metric learning oriented early exit methodology for DL models. As it was demonstrated, employing early exits in metric learning scenarios pose unique challenges compared to existing methodologies for classification-oriented early exits. To this end, the Bag-of-Features model was employed to efficiently extract compact representations from any layer of a DL model that is then combined with an efficient linear regressor to match the final representation of the model (without having to feedforward the whole computational graph). The proposed method is agile and can be directly used with any pre-trained DL model, while it is end-to-end differentiable, allowing for further fine-tuning the models towards having multiple early exits. The effectiveness of the proposed method was demonstrated using five face verification/recognition datasets.

Another work by AUTH was presented in the International Conference on Emerging Techniques in Computational Intelligence (ICETCI):

Pavlos Tiritiris, Nikolaos Passalis, Anastasios Tefas "Temporal Difference Rewards for End-to-end Vision-based Active Robot Tracking using Deep Reinforcement Learning", 2021 International Conference on Emerging Techniques in Computational Intelligence (ICETCI), Hyderabad, India, Aug. 2020, 10.1109/ICETCI51973.2021.9574071

In this paper, the authors worked on object tracking that allows for localizing moving objects in sequences of frames providing detailed information regarding the trajectory of objects that appear in a scene. In this paper, they studied active object tracking, where a tracker receives an input visual observation and directly outputs the most appropriate control actions in order to follow and keep the target in its field of view, unifying in this way the task of visual tracking and control. This is in contrast with conventional tracking approaches, as typically developed by the computer vision community, where the problem of detecting the tracked object in a frame is decoupled from the problem of controlling the camera and/or the robot to follow the object. Deep Reinforcement Learning (DLR) methods hold the credentials for overcoming these issues, since they allow for tackling both problems, i.e., detecting the tracked object and providing control commands, at the same time. However, DRL algorithms require a significantly different methodology for training compared to traditional computer vision models, e.g., they rely on dynamic simulations for training instead of static datasets, while they are often notoriously difficult to converge, often requiring reward shaping approaches for increasing convergence speed and stability. The main contribution of this paper is a DRL, vision-based active tracking method, along with an appropriately designed reward shaping approach for active tracking problems. The developed methods are evaluated using a state-of-the-art robotics simulator, demonstrating good generalization on various dynamic trajectories of moving objects under a wide range of different setups.

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

Another paper by AUTH was accepted and presented at the 2021 International Conference on Multimedia and Expo (ICME 2021):

Maria Tzelepi and Anastasios Tefas, "Efficient Training of Lightweight Neural Networks Using Online Self-Acquired Knowledge Distillation." 2021 IEEE International Conference on Multimedia and Expo (ICME). IEEE, 2021

In this paper, the authors proposed an online distillation method for efficiently training lightweight deep learning models for classification tasks. The proposed method overcomes the limitations of conventional distillation (i.e., complex, long-lasting, memory and computationally demanding process), by acquiring the additional knowledge about the similarities of the data samples with all the classes from the model itself and also in an online fashion. To achieve this goal, the knn non-parametric density estimation was employed for estimating the unknown probability distributions of the data samples in the output space of the network. This allowed for directly estimating the posterior class probabilities of the data samples and using them as soft labels that encode explicit information of the data samples with the classes. The experimental evaluation on four datasets validated the effectiveness of the proposed method.

Yet another paper by AUTH was presented at the 2021 International Conference on Information, Intelligence, Systems & Applications (IISA 2021):

Maria Tzelepi and Anastasios Tefas, "Semantic Scene Segmentation for Robotics Applications." 2021 12th International Conference on Information, Intelligence, Systems & Applications (IISA). IEEE, 2021.

In this work, the authors dealt with semantic scene segmentation for robotics applications. Semantic scene segmentation is a challenging step towards comprehensive scene understanding accompanied by many robotics applications, e.g., autonomous navigation. In such applications there are certain computational restrictions e.g., operation on low-power GPUs, at speed and also for high-resolution input. Existing semantic segmentation methods provide evaluation results under different setups and mainly considering high-power GPUs. In this paper, the behaviour of the most successful segmentation models in terms of deployment speed, under various setups (high-power and low-power GPUs, input sizes) for robotics applications was investigated. The target is to provide a comparative study of existing state-of-the-art models so as to select the most compliant with the robotics applications requirements.

Moreover, a paper by AUTH was presented at the 2021 International Conference on Engineering Applications of Neural Networks (EANN 2021):

Charalampos Symeonidis and Paraskevi Nousi and Pavlos Tosidis and Kostas Tsampazis and Nikolaos Passalis and Anastasios Tefas and Nikos Nikolaidis, "Efficient Realistic Data Generation Framework leveraging Deep Learning-based Human Digitization", International Conference on Engineering Applications of Neural Networks (EANN), 2021.

In this paper the authors presented a method capable of automatically generating realistic synthetic data with annotations, while leveraging existing images both for the creation of 3D models of humans, as well as for employing realistic backgrounds. The proposed method has minimal cost, since it does not require handcrafted 3D models or simulation environments and it

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

can be used for generating data for a variety of tasks, such as person detection, face recognition, and human pose estimation. The conducted experimental evaluation demonstrated that the generated data are suitable for training DL models in most cases. It is worth noting that for some tasks, such as face recognition, further improvements can be obtained. These results highlight the potential of the proposed data generation method both for generating large-scale datasets for data-scarce domains, as well as for minimizing the distribution shift that is experienced in many robotics applications, e.g., when transferring the models from a simulation environment in real deployment and vice versa.

Two papers by TUD were accepted and presented at the 2021 International Conference on Intelligent Robots and Systems (IROS 2021):

Bas van der Heijden, Laura Ferranti, Jens Kober, Robert Babuska, "DeepKoCo: Efficient latent planning with a task-relevant Koopman Representation", 2021 International Conference on Intelligent Robots and Systems (IROS 2021), Prague, Czech Republic/Online.

This paper presents DeepKoCo, a novel model-based agent that learns a latent Koopman representation from images. This representation allows DeepKoCo to plan efficiently using linear control methods, such as linear model predictive control. Compared to traditional agents, DeepKoCo learns task-relevant dynamics, thanks to the use of a tailored lossy autoen-coder network that allows DeepKoCo to learn latent dynamics that reconstruct and predict only observed costs, rather than all observed dynamics. As the results show, DeepKoCo achieves similar final performance as traditional model-free methods oncomplex control tasks while being considerably more robust todistractor dynamics, making the proposed agent more amenable for real-life applications

Jacob E. Kooi, Robert Babuska, "Inclined Quadrotor Landing using Deep Reinforcement Learning", 2021 International Conference on Intelligent Robots and Systems (IROS 2021), Prague, Czech Republic/Online.

This paper focuses on landing a quadrotor on an inclined surface, which is a challenging manoeuvre. The final state of any inclined landing trajectory is not an equilibrium, which precludes the use of most conventional control methods. The authors propose a deep reinforcement learning approach to design an autonomous landing controller for inclined surfaces. Using the proximal policy optimization (PPO) algorithm with sparse rewards and a tailored curriculum learning approach, a robust policy can be trained in simulation in less than 90 minutes on a standard laptop. The policy then directly runs on a real Crazyflie 2.1 quadrotor and successfully performs real inclined landings in a flying arena. A single policy evaluation takes approximately 2.5 ms, which makes it suitable for a future embedded implementation on the quadrotor.

A paper by AU was accepted and presented in the 2021 IEEE International Workshop on Multimedia Signal Processing (MMSP):

Negar Heidari and Alexandros Iosifidis. "Progressive Spatio-Temporal Bilinear Networks with Monte Carlo Dropout for Landmark-based Facial Expression Recognition with Uncertainty Estimation" IEEE International Workshop on Multimedia Signal Processing (MMSP), Tampere, Finland, Oct. 2021

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

Deep neural networks have been widely used for feature learning in facial expression recognition systems. However, small datasets and large intra-class variability can lead to overfitting. In this paper, the authors propose a method which learns an optimized compact network topology for real-time facial expression recognition utilizing localized facial landmark features. This method employs a spatio-temporal bilinear layer as backbone to capture the motion of facial landmarks during the execution of a facial expression effectively. Besides, it takes advantage of Monte Carlo Dropout to capture the model's uncertainty which is of great importance to analyze and treat uncertain cases. The performance of the method is evaluated on three widely used datasets and it is comparable to that of video-based state-of-the-art methods while it has much less complexity.

Another paper by AU was accepted and presented in the 2021 International Conference on Emerging Techniques in Computational Intelligence (ICETCI):

Illia Oleksiienko and Alexandros Iosifidis. "Analysis of voxel-based 3D object detection methods efficiency for real-time embedded systems" International Conference on Emerging Techniques in Computational Intelligence, Hyderabad, India, Aug. 2022

In this paper, the authors analyze the efficiency of two popular voxel-based 3D object detection methods providing a good compromise between high performance and speed based on two aspects, their ability to detect objects located at large distances from the agent and their ability to operate in real time on embedded platforms equipped with high-performance GPUs. Experiments show that these methods mostly fail to detect distant small objects due to the sparsity of the input point clouds at large distances. Moreover, models trained on near objects achieve similar or better performance compared to those trained on all objects in the scene. This means that the models learn object appearance representations mostly from near objects. The findings of the study suggest that a considerable part of the computations of existing methods is focused on locations of the scene that do not contribute to successful detection. This means that the methods can achieve a speed-up of 40-60% by restricting operation to near objects while not sacrificing much in performance.

A paper by TAU and AU was accepted and presented in the 2021 British Machine Vision Conference (BMVC):

Firas Laakom, Kateryna Chumachenko, Jenni Raitoharju, Alexandros Iosifidis and Moncef Gabbouj. "Learning to ignore: rethinking attention in CNNs" British Machine Vision Conference (BMVC), Virtual, Nov. 2021

Recently, there has been an increasing interest in applying attention mechanisms in Convolutional Neural Networks (CNNs) to solve computer vision tasks. Most of these methods learn to explicitly identify and highlight relevant parts of the scene and pass the attended image to further layers of the network. In this paper, the authors argue that such an approach might not be optimal. Arguably, explicitly learning which parts of the image are relevant is typically harder than learning which parts of the image are less relevant and, thus, should be ignored. In fact, in vision domain, there are many easy-to-identify patterns of irrelevant features. For example, image regions close to the borders are less likely to contain useful information for a classification task. Based on this idea, the authors propose to reformulate the attention mechanism in CNNs to learn to ignore

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

instead of learning to attend. Specifically, they propose to explicitly learn irrelevant information in the scene and suppress it in the produced representation, keeping only important attributes. This implicit attention scheme can be incorporated into any existing attention mechanism. In this work, this idea is validated using two recent attention methods Squeeze and Excitation (SE) block and Convolutional Block Attention Module (CBAM). Experimental results on different datasets and model architectures show that learning to ignore, i.e., implicit attention, yields superior performance compared to the standard approaches.

Another paper by TAU and AU was accepted and presented in the ICML 2021 Workshop on Overparameterization: Pitfalls & Opportunities:

Kateryna Chumachenko, Firas Laako, Jenni Raitoharju, Alexandros Iosifidis and Moncef Gabbouj. "Rethinking compactness in deep neural networks" ICML 2021 Workshop on Overparameterization: Pitfalls & Opportunities, Virtual, July 2021

Deep neural networks are a type of over-parameterized models which are able to achieve high performance despite having typically more parameters than training samples. Recently, there has been an increasing interest in uncovering and understanding the different phenomena that occur in the over-parameterized regime induced by such neural networks. In this paper, the authors aim to shed light on the relationship between class compactness of the learned feature representations and the model performance. Surprisingly, the authors find that models that learn more class-invariant features do not necessarily perform better. Moreover, they show that during training, class-wise variance increases and the models learn a less compact and more outspread representation of the classes.

Yet another paper by TAU and AU was accepted and presented in the International Joint Conference on Neural Networks (IJCNN) 2021:

Mohammad Soltanian, Junaid Malik, Jenni Raitoharju, Alexandros Iosifidis, Serkan Kiranyaz and Moncef Gabbouj, "Speech Command Recognition in Computationally Constrained Environments with a Quadratic Self-Organized Operational Layer," International Joint Conference on Neural Networks (IJCNN), 2021.

Automatic classification of speech commands has revolutionized human computer interactions in robotic applications. However, employed recognition models usually follow the methodology of deep learning with complicated networks which are memory and energy hungry. So, there is a need to either squeeze these complicated models or use more efficient lightweight models in order to be able to implement the resulting classifiers on embedded devices. In this paper, the authors pick the second approach and propose a network layer to enhance the speech command recognition capability of a lightweight network and demonstrate the result via experiments. The employed method borrows the ideas of Taylor expansion and quadratic forms to construct a better representation of features in both input and hidden layers. This richer representation results in recognition accuracy improvement as shown by extensive experiments on Google speech commands (GSC) and synthetic speech commands (SSC) datasets.

A paper by TAU was accepted and presented in the 2021 20th International Conference on Advanced Robotics (ICAR):

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

Saad Ahmad, Kulunu Samarawickrama, Esa Rahtu and Roel Pieters. "Automatic Dataset Generation From CAD for Vision-Based Grasping", 20th International Conference on Advanced Robotics (ICAR), Virtual, Dec. 2021.

Recent developments in robotics and deep learning enable the training of models for a wide variety of tasks, from large amounts of collected data. Visual and robotic tasks, such as pose estimation or grasping, are trained from image data (RGB-D) or point clouds that need to be representative for the actual objects, to acquire accurate and robust results. This implies either generalized object models or large datasets that include all object and environment variability, for training. However, data collection is often a bottleneck in the fast development of learning-based models. In fact, data collection might be impossible or even undesirable, as physical objects are unavailable or the physical recording of data is too time-consuming and expensive. For example, when building a data recording setup with cameras and robotic hardware. CAD tools, in combination with robot simulation, offer a solution for the
generation of training data that can be easily automated and that can be just as realistic as real world data. In this work, TAU proposed a data generation pipeline that takes as input a CAD model of an object and automatically generates the required training data for object pose estimation and object grasp detection. The object data generated are: RGB and depth image, object binary mask, class label and ground truth pose in camera- and world frame. They demonstrated the dataset generation of several sets of industrial object assemblies and evaluated the trained models on state of the art pose estimation and grasp detection approaches. Code and video are available at: https://github.com/KulunuOS/gazebo dataset generation

2.1.5.3 Preprints

Six preprints that describe results obtained within the project were published during this period. A preprint by TUD is available on ArXiv:

Osama Mazhar, Jens Kober, "Random Shadows and Highlights: A new data augmentation method for extreme lighting conditions", arXiv:2101.05361, 2021.

This paper proposes a new data augmentation method, Random Shadows and Highlights (RSH) to acquire robustness against lighting perturbations. The method creates random shadows and highlights on images, thus challenging the neural network during the learning process such that it acquires immunity against such input corruptions in real world applications. It is a parameter-learning free method which can be integrated into most vision related learning applications effortlessly. With extensive experimentation, it is demonstrated that RSH not only increases the robustness of the models against lighting perturbations, but also reduces overfitting significantly. Thus RSH should be considered essential for all vision related learning systems.

A preprint by ALU-FR is available on ArXiv:

Borna Bešić and Abhinav Valada. "Dynamic Object Removal and Spatio-Temporal RGB-D Inpainting via Geometry-Aware Adversarial Learning", arXiv:2008.05058v3, 2021.

This paper introduces DynaFill, a method to remove dynamic objects from dense 3D maps and efficiently fill the gaps by synthesizing plausible color, texture and geometry in regions occluded by dynamic objects. The proposed architecture follows a coarse-to-fine topology and

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

incorporates our gated recurrent feedback mechanism to adaptively fuse information from previous timesteps. The architecture is optimized using adversarial training to synthesize fine realistic textures which enables it to hallucinate color and depth structure in occluded regions online in a spatially and temporally coherent manner, without relying on future frame information. Casting the inpainting problem as an image-to-image translation task, the model also corrects regions correlated with the presence of dynamic objects in the scene, such as shadows or reflections. The authors introduce a large-scale hyperrealistic dataset with RGB-D images, semantic segmentation labels, camera poses as well as groundtruth RGB-D information of occluded regions. Extensive quantitative and qualitative evaluations show that the proposed approach achieves state-of-the-art performance, even in challenging weather conditions. Furthermore, the authors present results for retrieval-based visual localization with the synthesized images that demonstrate the utility of our approach.

A preprint by AU is available on ArXiv:

Lukas Hedegaard and Alexandros Iosifidis. "Continual 3D Convolutional Neural Networks for Real-time Processing of Videos", arXiv:2106.00050, 2021

This paper introduces 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 processing 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 multiple clips. While yielding an order of magnitude in computational savings, Co3D CNNs have memory requirements comparable with that of corresponding regular 3D CNNs and are less affected by changes in the size of the temporal receptive field. The authors show that Continual 3D CNNs initialised on the weights from pre-existing state-of-the-art video recognition models reduce the floating point operations for frame-wise computations by 10.0-12.4x while improving accuracy on Kinetics-400 by 2.3-3.8. Moreover, they investigate the transient start-up response of Co3D CNNs and perform an extensive benchmark of online processing speed as well as accuracy for publicly available state-of-the-art 3D CNNs on modern hardware.

A preprint by TAU and AU is available on ArXiv:

Dat Thanh Tran, Moncef Gabbouj and Alexandros Iosifidis. "Knowledge Distillation by Sparse Representation Matching", arXiv:2103.17012, 2021

In this paper, the authors propose Sparse Representation Matching (SRM), a method to transfer intermediate knowledge obtained from one Convolutional Neural Network (CNN) to another by utilizing sparse representation learning. SRM first extracts sparse representations of the hidden features of the teacher CNN, which are then used to generate both pixel-level and image-level labels for training intermediate feature maps of the student network. The authors formulate SRM as a neural processing block, which can be efficiently optimized using stochastic gradient descent and integrated into any CNN in a plug-and-play manner. Experiments demonstrate that SRM is robust to architectural differences between the teacher and student networks, and outperforms other KD techniques across several datasets.

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

A preprint by TAU, AUTH and AU is available on ArXiv:

Dat Thanh Tran, Nikolaos Passalis, Anastasios Tefas, Moncef Gabbouj and Alexandros Iosifidis. "Attention-based Neural Bag-of-Features Learning for Sequence Data", arXiv:2005.12250, 2021

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. The authors conducted extensive experiments 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.

A preprint by AUTH is available on TechRxiv:

C. Symeonidis, I. Mademlis, I. Pitas and N. Nikolaidis, "Neural attention-driven Non-Maximum Suppression for person detection", TechRxiv preprint, 10.36227/techrxiv.16940275, 2021.

Non-Maximum Suppression (NMS) is the last step in a typical object detection system. Heavy occlusions when detecting humans in images of crowded areas impose great challenges to most NMS methods, despite the importance of such a task for human safety-centric applications. This paper presents Seq2Seq-NMS, a novel deep neural architecture for performing NMS in similar hard cases, which relies on reformulating NMS as a sequence-to-sequence problem. The proposed method relies on the Multihead Scaled Dot-Product Attention mechanism in order to efficiently capture interrelations across the sequence of candidate detections, while also jointly exploiting visual appearance and geometric properties of the input ROIs in order to better represent them. Quantitative evaluation on three public person detection datasets, each one using a different detector, showed that Seq2Seq-NMS can provide state-of-the-art results at the IoU threshold used for annotating its training dataset, with acceptable inference runtime requirements and good behaviour for large numbers of raw candidate ROIs per image.

2.1.5.4 Books

As mentioned in D9.5, following an invitation by Elsevier, Profs. Alexandros Iosifidis (AU) and Anastasios Tefas (AUTH) started editing in the previous period a book focusing on the topics of OpenDR. The title of the book is "Deep Learning for Robot Perception and Cognition". During this period, the book was finalized (the great majority of chapters were written by project partners) and will be published in January 2022.

2.1.6 Special Sessions and Issues

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

During the second year of the project, one Special Session and one Special Issue on topics related to the project were organized by consortium members

Indeed, a Special Issue co-organized by E. Kayacan, AU, was published in IEEE Transactions on Artificial Intelligence:

W. He, Z. Li, E. Kayacan, C. Huang and M. V. Basin, "Guest Editorial: Special Issue on Artificial Intelligence for Robotics," in IEEE Transactions on Artificial Intelligence, vol. 2, no. 5, pp. 382-383, Oct. 2021, doi: 10.1109/TAI.2021.3119155.

The special issue was dedicated to AI for robotics, including AI to support human–robot interaction, AI for multirobot systems, AI learning algorithms in robotics, ethics of AI in the context of robotics, and applications of AI-enabled robots. Topics covered in this special issue included, but were not limited to AI for multirobot systems, AI in autonomous systems, AI in cognitive robotics, AI in human–robot interaction, AI-enabled robots to assist health professionals as in the case of COVID-19, applications of AI-enabled robots in industry, agriculture, logistics, medicine, augmented intelligence in robotics. The special issue received 17 articles on AI for robotics from different countries, of which seven articles contributed significantly and were accepted for publication. These papers offered innovative ideas and concepts, discoveries and improvements, and novel applications to the field of AI for robotics.

A Special Session on Artificial Neural Networks for Computer and Robot Vision was organized by Prof. Alexandros Iosifidis (AU) in the International Conference on Emerging Techniques in Computational Intelligence 2021, technically co-sponsored by the IEEE Computational Intelligence Society. The Conference was organized in a hybrid mode. The session included presentations of five papers. A screenshot from the Special Session, where a contributing paper from AU is presented is shown in Figure 12.

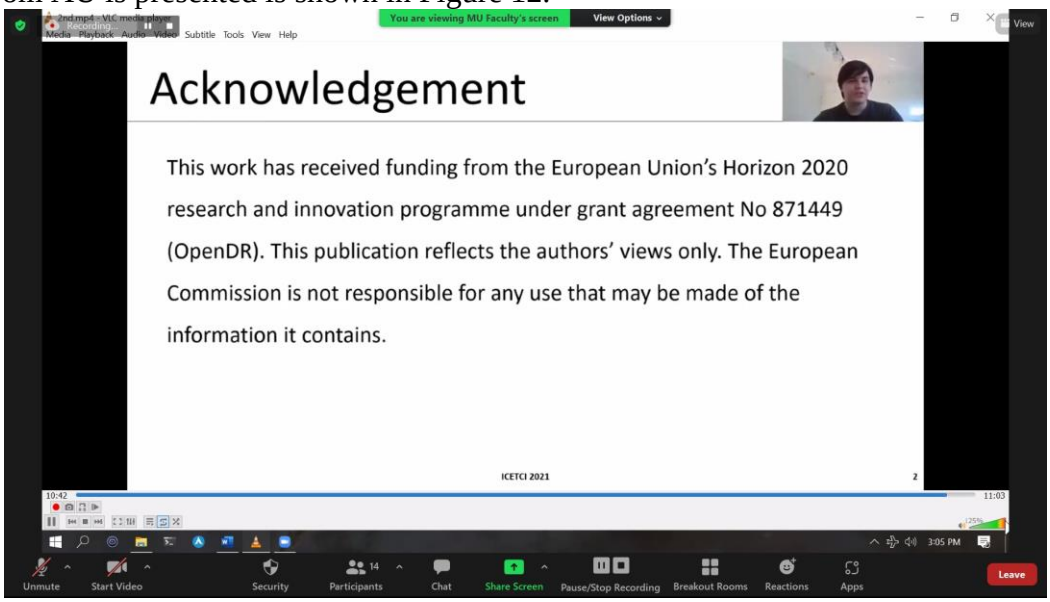

Figure 12: Screenshot from the presentation of AU-contributed paper in the ICETCI 2021 Special Session on Artificial Neural Networks for Computer and Robot Vision

2.1.7 Workshops, Invited / Keynote Talks

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

Dissemination of OpenDR information and findings through workshops or invited/keynote talks has obviously a high impact since, usually, the audience consists of researchers or industrial representatives, interested in the specific topic of the talk/tutorial. The following activities took place during the second year of the project:

IROS 2021 Workshop

OpenDR successfully organized on October 1st 2021, under the leadership of PAL Robotics, an IROS 2021 workshop on "Open Deep Learning Toolkit for Robotics: Towards Democratizing Artificial Intelligence". The workshop, which was held virtually (Figure 13) and was well attended, was structured in two sessions, each followed by a lively panel discussion involving the session speakers and a moderator. Members of the OpenDR consortium contributed to the workshop with interesting talks on a variety of topics:

- Deep Learning and AI Applications in the Webots Robot Simulator (Stefania Pedrazzi)
- Robotic grasping in Agile Production (Roel S. Pieters)
- Deep Learning models for efficient video-based action recognition (Alexandros Iosifidis)
- Learning Efficient Representations for Perception, Tracking, and Localization (Abhinav Valada)
- How to benefit from learning in predictive control (Laura Ferranti)
- Business perspectives and applications of OpenDR for AGROINTELLI (Ole Green)

Moreover, the workshop included talks from prominent roboticists namely:

- Using Open Deep Learning Toolkits in ROS/ROS2 (Francisco Martin Rico, Rey Juan Carlos University Spain)
- European Training Network on Personalized Robotics as Service Oriented applications (Silvia Rossi, University of Naples "Federico II", Italy)

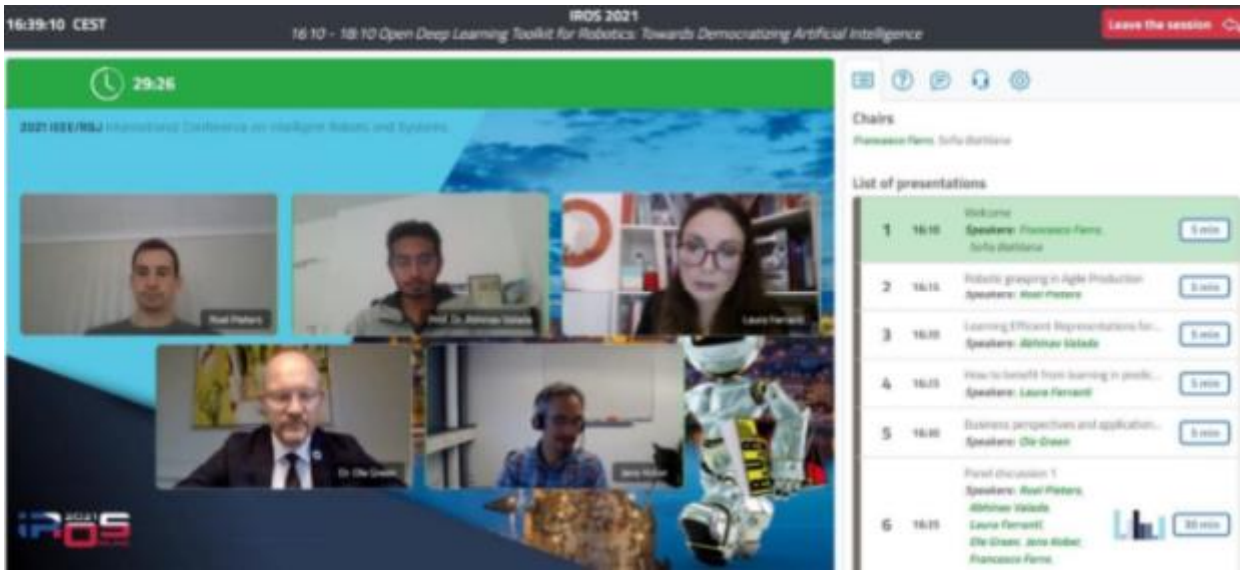

This project has received funding from the *European Union's Horizon 2020 research and innovation programme* under grant agreement No *871449*.

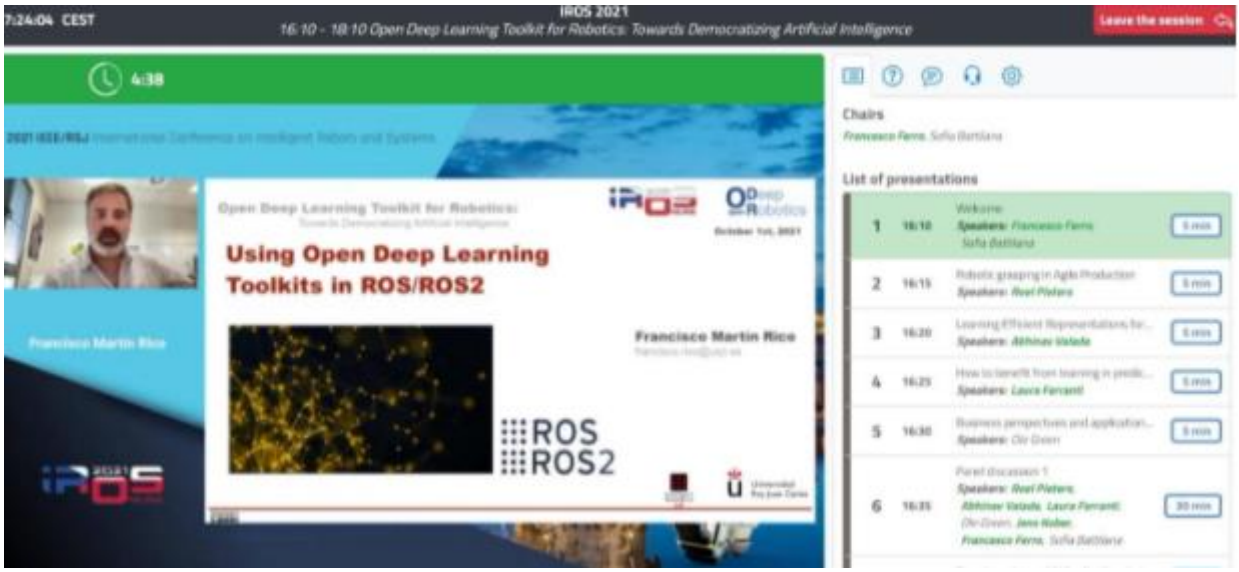

Figure 13: Screenshots from the virtual IROS workshop

Apart from their live talks, speakers made available longer recorded versions of their talks (some of which are available on the OpenDR YouTube channel) as well as the talk slides. Francesco Ferro and Sofia Battilana (PAL) played a key role in the workshop, by coordinating the flow and providing opening and closing remarks.

In order to disseminate the workshop "Open Deep Learning Toolkit for Robotics: Towards Democratising Artificial Intelligence" a new webpage was created at PAL: https://events.pal-robotics.com/iros2021-workshop (Figure 14). The page provided information regarding the workshop's objectives, the topics of interest and the program, as well as the invited speakers and the organisers. A registration button leading to the IROS webpage was also included.

Figure 14: IROS Workshop official website at PAL

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

An additional webpage was created in the OpenDR website ("Events" section) to maximise awareness (Figure 15).

Figure 15: IROS workshop webpage at OpenDR website

Furthermore, as already mentioned, a new YouTube "OpenDR-IROS Workshop" sub-channel was created under the OpenDR account to include recorded talks given by OpenDR partners.

Invited talk at IEEE CASE workshop

Dr. Roel Pieters (TAU) was invited to give a talk in the workshop 'Novel robot technologies for agile manufacturing', as part of the 17th IEEE International Conference on Automation Science and Engineering (CASE), to discuss robotic object grasping and manipulation in Agile Production, which finds overlap with topics in OpenDR. The OpenDR project and toolkit were presented as well.

Invited talk at Harbin Institute of Technology

TUD gave an invited talk at Harbin Institute of Technology in July 2021. The content of the talk included TUD OpenDR work on harsh lighting conditions and multimodal object detection.

Webinar at PROSA

In April 2021 A. Iosifidis (AU) provided a webinar organized by PROSA, a Union of IT professionals in Denmark with about 16.000 members, where he introduced OpenDR's objectives and the work of AU.

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

2.1.8 Participation to Tradeshows, Exhibitions, EU Events, Industry Workshops, Awards

Trade shows, exhibitions, events organized by the EU or similar venues are important for the dissemination of OpenDR's aims and results. During this period, the project was present at the following two fora:

Participation in Beyond 4.0 Expo

The AUTH OpenDR team presented the project at the Beyond 4.0 Expo, an Industry 4.0 international exhibition/forum, which took place in Thessaloniki, Greece in October 2021. Among the main aims of Beyond 4.0 was to discuss the impact and opportunities of Industry 4.0 for South East Europe, the Mediterranean & the MENA region and accelerate the potential of key economic sectors. AUTH team presented a poster in the University booth during the whole duration of the exhibition (3 days) Beyond 4.0 attracted a diverse audience from different countries, including entrepreneurs, developers, researchers and students.

Google Summer of Code

OpenDR participated through the "Deepbots" supported project at the Google Summer of Code 2021 (GSoC2021) under the Open Technologies Alliance (GFOSS) organisation. GSoC funded two students in order to develop an infrastructure related to Deepbots which would support evolutionary algorithms and a set of testing environments with humanoid robots in Webots respectively.

Participation in BDVA Workshop

On March 2021 a presentation entitled "Standards, formats and interoperability protocols for Deep Learning models in OpenDR" was given by Dr N. Passalis (AUTH) on "Big Data Value Association (BDVA) workshop, where the main aspects of the project were pointed out, while among the topics discussed was the use of standards that ensure the interoperability in the fragmented landscape of Deep Learning.

Participation in the European Robotics Forum (ERF) 2021

The project and its objective were also mentioned in the ERF 2021 virtual event (13-15th of April), under the workshop entitled "How to make an impact in EU Projects and boost innovation" by the PAL team (Figure 16).

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

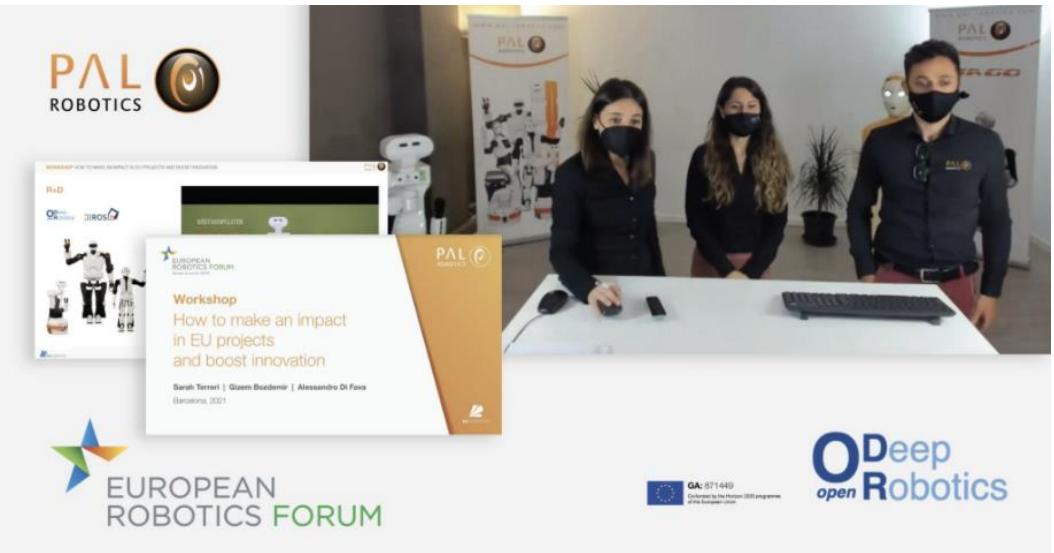

Figure 16: Virtual workshop at ERF 2021 by PAL Robotics

PAL activities

PAL was particularly active in disseminating the project in a number of events. Indeed PAL Robotics talked about the project in the company´s virtual Webinar "Advancing robotics through EU projects" (July 1st 2021) along with its other collaborative projects. The target group of the webinar was entrepreneurs, developers, researchers who are interested in exploring collaborations in European robotics. Moreover project work and objectives were disseminated by PAL in a Webinar at the Cleveland State University, Ohio, USA (April 2021) and two workshops in Spain (Jornada de inteligencia artificial aplicada al turismo, April 2021/ Jornada de Mecatrónica May 2021).

AGI award at Agritechnica 2022

AGI won a silver star at Agritechnica 2022 for Robotti and RoboVeg selective broccoli harvesting combination. Although the award does not concern innovations obtained within the project, it is important that the Robotti will be the testbed for the agri-food use case.

2.1.9 Collaboration with Other Projects and DIHs

Collaboration with two projects and one DIH took place during the second year of the project. Moreover PAL plans to introduce OpenDR technology in other projects it participates in.

FELICE

Project coordinator A. Tefas has been contacted in March 2021 by the coordinator of FELICE H2020 EU project (www.felice-project.eu, the project started in January 2021 and targets the application priority area of agile production) with a request for collaboration between the two projects. After discussions in the Steering Committee and consultation with the PO, the following collaboration framework has been established:

- OpenDR disseminated Deliverable D2.1 on requirements and specifications to FELICE.

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

- FELICE will consider using the tools of OpenDR toolkit whenever possible (e.g., in a project showcase), appropriately acknowledging the OpenDR toolkit and citing the utilized OpenDR methods in the corresponding publications.
- FELICE will consider giving feedback to OpenDR about the Requirements and Specifications of OpenDR, and OpenDR will take into account this feedback.
- FELICE will consider contributing to the OpenDR toolkit when mature results will be produced by FELICE.

The consortium believes that interesting results can come out of this collaboration in the near future.

SqM Farm

AGI will be combining the work from the SqM Farm project (Danish national project, GUDP fund) with the OpenDR project. In the SqM Farm project, AGI is creating automatic mapping services of various DL models and collected Robotti data. Combining these two, AGI can map and show the user plant maps from DL algorithms developed in OpenDR**.**

Trinity DIH

During the second year (M13-M24) OpenDR's continued active engagement with the DIH TRINITY was managed by TAU on the industrial relevance of the OpenDR toolkit in the field of Agile Production. The activities of OpenDR, the development of the OpenDR toolkit and its individual tools have been communicated to DIH TRINITY and its coordination team. Moreover, the plan of a first open release of the toolkit at M24 has been communicated and dissemination through the network of the Digital Innovation Hub has been requested.

PAL plans to use OpenDR technology in other projects

For PAL, the OpenDR project opens various possibilities of collaboration with projects it participates in, like the SHAPES and SPRING projects but also SANDRO and ALMI projects where tools from the OpenDR toolkit are planned to be integrated in a healthcare use case of these projects.

2.2 Dissemination Performance Metrics

In the second year of the project's life cycle, OpenDR consortium focused on maintaining and improving the key performance indicators (KPIs) which had been defined in the DoA, by producing significant scientific work and disseminating these results in the best way. An important number of scientific papers was published, new posts were created on the website, numerous announcements were made in the social media accounts etc. These contributed in keeping the dissemination performance metrics at a high level. In the following sections we will present the obtained KPI values and compare them to targets that had been set in the DoA.

2.2.1 Project Website

An important number of new posts and documents was created/uploaded on the project website throughout this second year. Through these posts, all the new journal and conference

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

papers, announcements about events as well as participation in workshops and conferences were made known to the public. Thus, the project managed to already achieve the targets which were set in the DoA (even those set for M48) and also record important progress in relation to the period M1-M12.

KPI	Target (for M48, unless stated otherwise)	M12	M24
Number of visitors	1000 (by M12)	More than 3200	More than 5600
Average duration of visits:	2:00	1:36	1:46
Number of posts published (including news):	>50	15	56
Number of documents available in the website repository:	>20	18	41

2.2.2 Social Media

The OpenDR project continued its presence in the most popular social media channels such as Twitter, Facebook and LinkedIn, by updating the corresponding content at a very satisfactory pace, which contributed in already exceeding the M48 target values that had been defined in the DoA, for many key performance indicators.

2.2.2.1 Twitter

During the second year, the Twitter account gathered even more followers than those mentioned in the deliverable 9.5, reaching the number 230, while the tweets and retweets published were more than 150 in total. Both of these metrics indicate that the goals which were set at the beginning of the project were achieved although we are still on M24.

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

2.2.2.2 Facebook, YouTube

The Facebook account of the project which was created to support the dissemination activities, despite not being mentioned in the DoA, counts more than 190 likes and more than 130 posts (+60 since M12) from the beginning of the project. As far as the YouTube channel is concerned, so far 12 videos of conference paper presentations account for more than 2020 views

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

in total and 40 subscribers. In addition, four videos were posted in the new IROS 2021 workshop sub-channel. The metrics show a very satisfactory increase compared to M12 (800 views and 27 subscribers in M12).

2.2.2.3 LinkedIn

OpenDR tried to increase awareness and disseminate all the second year's publications and news through the LinkedIn group and the LinkedIn official page. According to the target that has been set, group members should be more than 150 by M48, while currently it counts 132 members, a value very close to the target. Furthermore, so far the group counts 105 posts by group members. The LinkedIn official page has gathered until now more than 245 followers, a fact which, in combination with the number of the Linkedin group members, shows the very satisfactory presence of OpenDR in this platform.

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

2.2.3 Publications

According to the DoA the number of scientific papers should be more than 8 per year, but, as shown in the table below, the goal has been again surpassed by far in the second year. From M13 until M24, 14 additional conference papers were accepted/presented and 12 more journal papers were published (26 in total). Thus, so far the project has published 25 papers per year. These numbers show dedication to the ultimate project goal which is the production of significant scientific work.

KPI	Target	M12	M24
Number of publications	>8 per year	24 (16 Conferences, 7 Journals, 1 Book)	50 (14 new Conferences, 12 new Journals)

2.2.4 Promotional Material

From M13 to M24 there has been a considerable increase in two of the KPIs related to the promotional material, namely the size of the dissemination list and the number of views of the promotional video. For both these KPIs the M48 target has already been achieved. Regarding the production of new promotional material (newsletter, press releases, promotional videos, brochure) the consortium is confident that the targets that have been defined in the DoA are going to be achieved by M48. However, there is considerable uncertainty about achieving the goals regarding the use of the factsheet, poster and brochures in live events, as the COVID19 pandemic proved to be a major obstacle in the organization of live events and most continue to be virtual. However the OpenDR consortium will seize all the opportunities regarding the use of communication material in any future physical event. After all, the consortium believes that the effect of these shortcomings

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

in the dissemination campaign is rather minimal and is by far compensated by the extensive online campaigns (newsletters, website, social media etc) undertaken.

KPI	Target (M48)	M12	M24
Number of e-newsletters published	**4**	**1**	**2**
Size of the dissemination list	**> 500**	**>1000 (estimate)**	**>3000 (estimate)**
Number of press releases	**3**	**1**	**1**
Number of videos to be produced	**2**	**1**	**1**
Number of video views	**500**	**529**	**911**
Number of events where the factsheet is used	**10**	**1**	**1**
Number of events where the poster is used	**10**	**1**	**2**
Number of brochures	**2**	**1**	**1**
Number of events where the brochure is used	**10**	**-**	**-**
Number of copies distributed (brochures)	**800 hardcopies + eBrochures**	**-**	**-**
Number of events where a roll-up banner is used	**4**	**0**	**0**

2.2.5 Events

Being in the second year of the project's duration and given the pandemic situation as mentioned above, OpenDR has made important progress concerning the organisation of new training activities (Google Summer of Code), workshops (IROS Workshop), and use of exhibition booths (BEYOND 4.0). Some of the M48 targets in this area (number of training activities and their participants, number of workshops/tutorials) have already been achieved in M24, however this does not mean that the consortium will stop being active in these areas. Regarding the other KPIs, the consortium is optimistic that the relevant targets are going to be achieved by M48.

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

KPI	Target (M48)	M12	M24
Number of training activities	**2**	**1**	**2**
Number of participants per activity	**20**	**130**	~70
Number of web-based training activities	**1**	**-**	**-**
Number of Workshops / Symposia / Tutorials	**2**	**1 (tutorial)**	**2 (1 tutorial, 1 workshop)**
Number of participants per workshop	**50**	**-**	**~30**
Number of exhibition booths:	**2**	**-**	**1**

2.3 Conclusions and Future Plan

The large number of dissemination activities (especially in what concerns publications, 50 since the project start, without counting preprints) that materialized within the second year of the project, as well as the values of the KPIs and the fact that, in a significant number of them, the M48 targets have already been reached, signifies that OpenDR has made important progress and continues to be in a very good shape, despite the negative effect of COVID-19. As expected, the effects of the pandemic are more prominent in KPIs that involve physical events (e.g., number of distributed brochures or number of events where the poster has been used). However, 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.

The release of the first public version of the toolkit in M24 will signal the start of a new series of activities that will focus on attracting relevant parties in getting to know and use it. These activities might include training events such as webinars or online courses and tutorials. The consortium will also strengthen its activities towards liaising with relevant DIHs so as to make them and their members aware of the toolkit and the opportunities that it offers. The consortium will obviously closely monitor the adoption and visibility of the toolkit in GitHub by using metrics such as number of stars, forks, unique visitors etc.

3. Current Exploitation Plan

The OpenDR toolkit will provide a modular, generic, and open library that can bring state-of-the art deep learning approaches that can be used to common robotic scenarios, both simulated

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

and real. The industrial partners have plans to exploit the potential of this library, PAL and AGI to enhance their own robotic platforms and CYB to provide 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 and safety of the Robotti as well as the objects that surround it.

PAL's exploitation plan focuses on the commercialization of assistive robots, in this context OpenDR is expected to bridge the difficulties intrinsic to an unpredictable healthcare environment.

CYB's plan is to demonstrate the capabilities of the Webots robot simulator 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. Moreover the participation of the project in the EU Innovation Radar initiative and the current status of IPR protection and management are described in the relevant sections.

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 potential, but also to increase visibility of Webots in this field, which is currently lacking. By introducing tutorials, sample applications and more in depth scenarios that showcase OpenDR in different use-cases it will help CYB acquire more users and therefore expand the Webots community. 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, or for the creation of analogous scenarios and to organize competitions in the field. This is expected to yield additional requests for paid services and boost visibility.

With the first public release of the OpenDR toolkit set for December 2021, and to allow for the tools to gain some maturity, CYB plans to include the first demos and tutorials with the June release of Webots simulator, in order to ensure a high quality introduction of these tools and its potential to its user base.

Additionally, CYB plans to turn some project's results into web services and charge for its usage. For instance, users will be able to edit, run and view their simulations using the OpenDR toolkit directly from the web without needing to install anything, and could be charged for running their simulations on the provided host machines. This simplified approach lowers the barrier of entry for new users while at the same time boosting the visibility of the simulator and the OpenDR toolkit that powers it. It can be an effective approach in the academic field to introduce students to robotics and deep learning, to organize international competitions with minimal overhead, but also for researchers to easily disseminate their results as it does not require any setup in order to play and interact with the simulation.

For this purpose, work continues at CYB in order to increase its presence on the cloud by developing these tools and to prepare for the future exploitation that will be applied in a later project phase when a first version of the toolkit and scenarios will be available.

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

3.2 PAL Exploitation Plan

As already mentioned in D9.5, the advances in PAL are allowing more improved skills to become available for our robots that can bring an added value to our future robotics applications. An open source toolkit, like the one developed in OpenDR, is key to having more accessible AI skills and functionalities to a wider group of people and not only to the experts in the field.

The PAL-led use case will allow 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. The OpenDR toolkit will provide useful tools in order to create complex and meaningful applications in healthcare, making the power of deep learning available to people without significant expertise. For example, to create a receptionist robot, it is possible to use the output of the face detection algorithm to crop images and import them in the database to train the face recognition and provide input of the names of the people in order to implement custom welcoming messages for recognized persons.

We continue spreading the progress of the project to other PAL Robotics business units and managers to show the case studies and the methodology applied within the project, highlighting the possible benefits coming from the use of the OpenDR toolkit and discussing the possibilities of using the developed toolkit on the other platforms such as ARI, Talos, Kangaroo in the future. Once the tools are available and mature enough, PAL is planning to integrate them first to the Tiago platforms which will help us to develop new functionalities and new services for our clients.

We would like to put forward the effort that we started in the field of assistive robots and focus on commercializing our robots making their use easier in an unpredictable healthcare environment. In order to achieve this, we foresee a market analysis that will define, characterize and segment the potential opportunities for different solutions developed in the project, as well as describe the value chains serving each technology. Preliminary business models will be identified taking into account the outcomes of the market analysis and the identified new value chains. We will also study the possibility to extend the tests to real world environments with different scenarios in the healthcare field as well as other potential applications so as to ensure full functionality of all features of the technology.

The OpenDR project will open various possibilities of collaboration as discussed before in D9.5 like the SHAPES and SPRING projects but also for SANDRO and ALMI projects where tools from the OpenDR toolkit are planned to be integrated in a healthcare use case of these projects. With the first public release of the OpenDR toolkit set for December 2021, PAL plans to integrate the ROS ready software in these projects.

3.3 AGI Exploitation Plan

Through the OpenDR project, AGI aims to develop and potentially commercialize and exploit the following technologies: intra-row weeder system, increasing safety on Robotti, plant row guidance system, and plant identification models.

For the intra-row weeder system, the necessary technologic building blocks are currently being built in the OpenDR project and it is unknown when it will be mature enough for exploitation.

AGI will be using several of the OpenDR tools to increase the safety for Robotti and the objects in the surrounding environment. These tools include human detection, human pose detection, animal detection, and object detection. When the toolkit and tools are integrated, tested

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

and evaluated, it is currently the plan to allow pilot users to use the tools. Depending on the robustness and feedback from the pilot users, AGI will start exploiting the technology.

For the plant row guidance system, AGI is currently developing the system. The system is getting close to being able to test and this should happen in 2022 or 2023. When the system has been tested and validated, AGI will exploit and commercialize it.

The plant identification DL algorithms can be exploited in many ways. Through another project, AGI is building an automatic data mapping service from Robotti's cameras. Using this platform, the output of the DL algorithms will be mapped and displayed on a website. Combining this with the plant identification DL algorithms, AGI will be able to show the farmers where their crop is and to give a current count of their crop. In addition, the plant identification algorithms can also be used in the intra-row weeder, spot sprayer, band sprayer, and other precision farming implements that will be able to give a reduction of pesticide use, helping to protect the environment.

When the exploitation process starts, AGI will demonstrate and showcase these technologies. AGI has found it is important that the farmers are able to see and experience the technologies first hand before they adopt them.

3.4 OpenDR Participation in Innovation Radar

OpenDR has been selected to participate in the EU Innovation Radar, an initiative to collect structured data about the innovation profile of EU funded projects and their outputs. The goal of the initiative is to identify high-potential innovators and innovations and their specific 'go to market' needs. Following this invitation the consortium filled and submitted the respective questionnaire, as required. In this, 9 innovations were identified and described, one of them, actually the major one, being the OpenDR toolkit itself. The remaining 8 innovations came from the three SMEs that participate in the project, namely 2 from CYB (ROS2 Interface for Webots robot simulator, Framework for running Webots simulations on the Web), 3 from AGI (Plant row guidance system for Robotti agricultural robot, Intra-row weeder system for Robotti agricultural robot, Extra safety for Robotti agricultural robot using DL obstacle detection) and 3 from PAL (Improved human robot interaction for healthcare and other collaborative use cases, Improved awareness of the robot environment for in house/ hospital working, Make OpenDR toolkit available for PAL Robotics robots).

3.5 Protection and Management of IPR

The consortium agreed on developing the OpenDR toolkit under the Apache 2.0 license. With this permissive open source software license, the tools could be used and sold by the industrial partners as part of their commercial products. However, Apache 2.0 license is not copyleft, and thus the developers and researchers will retain the ownership of the code they develop, and if the code is redistributed or modified, the original author as well as the modifications have to be explicitly mentioned.

This Intellectual Property (IP) management approach is flexible enough to protect the copyright and allow commercial exploitations. Additionally, publicly releasing the toolkit will guarantee a much higher 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

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

contributions from OpenDR partners or external contributors that will maintain and extend the library.

4. Second Project Newsletter

In accordance with the DoA, the consortium has created, under the coordination of AUTH, the second project Newsletter. The document, which gives a brief introduction to the project, its consortium and illustrates its major results so far (including the first public release of the toolkit), can be found in Appendix I. The newsletter will be placed 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 or 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

2nd 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 2nd OpenDR Project Newsletter !

Newsletter 2.0: First release of the toolkit and current status

Almost everything we hear about 0artificial intelligence today is thanks to deep learning (DL). Deep learning has achieved tremendous performance jumps in the last decade in several Computer Vision (CV) and Machine Learning (ML) tasks, achieving in many cases super-human performance. However, DL cannot be currently fully exploited in robotics scenarios due to a number of barriers.

Learning Curve Barrier	Computational Complexity Barrier	Static Perception Barrier
DL has a **steeper learning curve** than traditional CV and ML methods	**DL requires vast amounts of computational power** and **energy**	**DL** is applied on **static environments** and **does not exploit spatial or temporal embodiment**

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

What is OpenDR?

OpenDR "Open Deep Learning for Robotics Toolkit", is a EU 2020 Project which was launched on January 2020 and aims to develop a modular, open and non-proprietary toolkit for core robotic functionalities by harnessing deep learning to provide advanced perception and cognition capabilities, meeting in this way the general requirements of robotics applications in the areas of healthcare, agri-food and agile production. The **OpenDR** project is coordinated by the Aristotle University of Thessaloniki, Greece and will be running throughout December 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 will focus its research on deep human centric active perception and cognition, where it will contribute on deep person/face/body part active detection/recognition and pose estimation, deep person/face/body part tracking, human activity recognition, as well as social signal analysis and recognition. AUTH will also lead the research in object detection/recognition and semantic scene segmentation and contribute to other areas such as evaluation and benchmarking activities of the project.

Tampere University (TAU) is Finland's second-largest university with 20.000 students and 330 professors. TAU participates with two labs/groups namely the Laboratory of Signal Processing at the Department of Computing Sciences and the Cognitive Robotics Group at the Department of Automation Technology and Mechanical Engineering. TAU will lead the research in deep human centric active perception and cognition, working mainly on deep speech and biosignals analysis and recognition, and will contribute to deep person/face/body

part active detection/recognition and multi-modal human centric perception and cognition as well as in a number of other topics. TAU will also contribute on defining the agile production use case requirements and specifications and on the integration of OpenDR to this use case.

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

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

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

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

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

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

Work Performed in the 2nd Year

1st Public Release of the Toolkit

Following months of development, integration and debugging as well as countless videoconferencing sessions, the **first official public release of OpenDR** is finally accessible in GitHub, as well as through pip and Docker Hub! OpenDR provides an intuitive and easy-to-use **Python interface**, a **C API** for selected tools, **a wealth of usage examples and supporting tools**, as well as **ready-to-use ROS nodes**. The toolkit provides more than 20 methods, for human pose estimation, face detection, recognition, facial expression recognition, semantic and panoptic segmentation, video and skeleton based action recognition, image, multimodal and point cloud based object detection, 2D and 3D object tracking, speech command recognition, heart anomaly detection, navigation for wheeled robots, and grasping. A set of data generation utilities, a hyperparameter tuning tool and a framework to easily apply RL both in simulation and real robotics applications are also included. All methods and their parameters are thoroughly documented, demonstration examples are available to showcase their functionality, and continuous integration tests ensure both consistency of the code and that no conflicts arise between the different tools. At the same time, OpenDR is built to

support Webots Open Source Robot Simulator, while it also extensively follows industry standards, such as ONNX model format and OpenAI Gym Interface. We look forward to receiving your feedback, bug reports and suggestions for improvements!

Deep Human Centric Active Perception and Cognition

Throughout the second year of the project AUTH, AU and TAU have continued their work on the human centric tools and algorithms for the OpenDR project, bringing the team closer to the realization of the powerful, flexible and efficient robotics toolkit.

AUTH has made additional contributions to the development of active perception in OpenDR. Specifically, a controller based on deep reinforcement learning, capable of issuing movement commands such that the face recognition confidence from an on-board camera image is maximized. Another work tackled active perception by predicting the best view angles for face recognition via synthetic image generation.

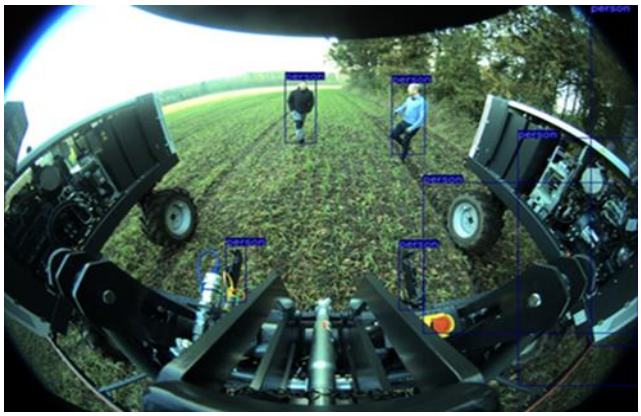

Human detection example on an image collected from the Robotti platform. Both true and false detections are evident, highlighting the need to adapt existing models to the particularities of agriculture use case.

AUTH also investigated artificial data generation to expand on the existing datasets, demonstrating significant performance improvements when both real and synthetic data are used for training DNN models. In addition, AUTH proposed a novel DNN-based Non-Maximum Suppression (NMS) method for the person detection task. The method reformulates NMS for object detection as a sequence-to-sequence problem and proposes a novel DNN-based architecture for classifying a ROI as "correct" or "potentially suppressed". Finally, AUTH jointly with AGI evaluated the performance of standard human detection models in the practical setting of the OpenDR agriculture use case, using visual data collected from

Robotti agricultural robot, highlighting the need to tackle the domain adaptation problem within the project.

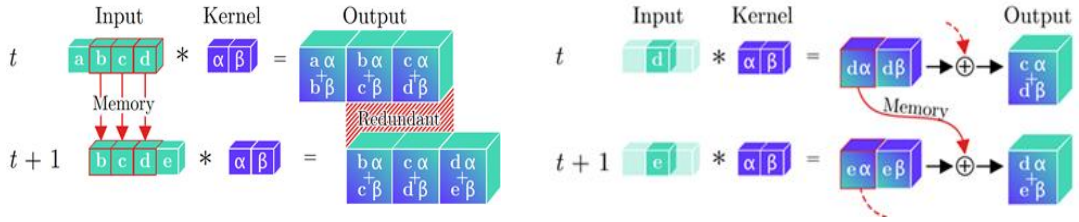

Standard 3D convolution operation (on the left) and proposed continual convolution (on the right). The proposed operation avoids redundant computations and requires less memory

To facilitate efficient video processing, AU has designed continual convolution – a replacement for the standard 3D convolution that allows fast per-frame predictions, while being able to reuse the exact same weights. Continual convolutions can thus be incorporated into existing models without fine tuning, unlike other similar operations from the literature. AU applied continual convolution in its own graph convolutional models for human activity recognition, yielding an increase in throughput from 7.0 times to 19.9 times with the same prediction accuracy. AU also proposed to extend its spatio-temporal graph models from skeleton-based human action recognition to landmark-based facial expression recognition, achieving accuracy comparable with state-of-the-art approaches, while using models of lower complexity.

TAU has started working on human centric multimodal perception and cognition, proposing a novel attention-based methodology where the neural models learn to ignore less relevant portions of the inputs. The developed attention mechanism developed was incorporated into TAU methods, allowing them to learn a joint spatio-temporal mask instead of independent separate masks. A tool for multimodal hand gesture recognition from RGBD data input was also developed.

Deep Environment Active Perception and Cognition

Partners AU, ALU-FR, AUTH, TAU and TUD worked towards providing a multitude of lightweight deep learning methods for deep environment active perception and cognition in various fields.

In the field of object detection/recognition and semantic scene segmentation, AUTH proposed an end-to-end trainable pseudo-active sensory refinement method

that works by applying a number of neural transformation layers on the sensor data. This allows for refining the sensory input, without having to reacquire the sensor data. Moreover, AUTH developed a single-stage self-distillation method for improving the performance of any deep neural model, in an online fashion, allowing for training compact yet effective classification models. Furthermore, AUTH worked on improving object recognition methods, by incorporating label embedding criteria into the learning objective of lightweight neural networks, capturing both general class similarities as well as instance specific resemblances between samples. AU proposed a type of neural networks called Variational Neural Networks, which can be used for the estimation of the deep learning model uncertainty. TAU proposed Sparse Representation Matching, a method to transfer intermediate knowledge obtained from one Convolutional Neural Network to another one by using sparse representation learning. TAU also conducted research towards better understanding of learned feature spaces in image classification tasks in deep learning. It is expected that models achieving good performance converge to feature spaces with highly-separable compact classes. TAU found that better compactness is not necessarily associated with better performance.

In the field of object localization and tracking, AU worked towards 3D object detection based on spherical projection images of point cloud data. The proposed method is able to regress the 3D size, rotation and position offset of an object. AU proposed methods for 3D object detection and tracking, based on spherical projection images and Siamese PointPillars, respectively. AUTH developed a 2D tracking methodology that focuses on active object tracking. The tracker receives an input visual observation and directly outputs the most appropriate control actions to follow and keep the target in its field of view. ALU-FR integrated their panoptic segmentation approach for images, EfficientPS, into the toolkit and extended it to LiDAR-based point clouds. Additionally, ALU-FR improved their work towards monocular vision-based localization in 3D maps.

In the field of sensor information fusion TAU has continued its work on the multimodal sensor fusion framework for manipulation tasks. The formulation of the representation learning task was extended with additional objective functions to support input reconstruction and cross-modal compensation, enabling the system to react to corrupted inputs. Additionally, TAU is working on optimizing the structures of the fusion modules and the feature encoder as a whole, using neural architecture search. Finally, TUD has developed efficient sensor fusion strategies in the context of object detection for RGB and IR/depth sensors for harsh lighting conditions. The proposed fusion strategies learn to exploit sensor redundancy in extreme lighting conditions by intelligently determining the scalar/mask weights

for the dominant sensor modality. A novel data augmentation scheme has been proposed to mimic such extreme conditions.

Deep Robot Action and Decision Making

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

AU proposed a novel end-to-end path planning algorithm for a team of multiple aerial-ground robots, targeting the agriculture use case. The main operations in the field are handled by the ground vehicle whereas the aerial robot is responsible for re-planning a collision-free trajectory for the ground robot in case of obstacles. The method was successfully tested in simulations and real-world experiments.

ALU-FR developed a novel approach for mobile navigation and manipulation. Given an arbitrary end-effector motion, reinforcement learning is used to control the robot base to output actions that ensure that the end-effector motions remain kinematically feasible. The method allows the use of a novel dense reward signal for long horizon tasks and shows strong zero-shot generalisation to unseen tasks and motions.

TUD finalized their work on the design of a novel and efficient model-based agent that learns a latent Koopman representation from images and is robust to possible distractions. This representation allows the method to plan efficiently using linear control methods, such as linear model predictive control. In addition, TUD developed EAGER, an engine agnostic gym environment for robotics tasks that will be integrated in the OpenDR toolkit.

TAU investigated robot grasping models for handling industrial objects in the agile production use case. Their single demonstration grasping has been integrated in Webots and in the OpenDR toolkit. In addition, a human-robot interaction scenario was developed that enables the collaboration between human and robot based on commands given by a person.

Different stages of the hand-over scenario by speech command-based coordination.

Simulation Environments and Data

AUTH finalized an annotated mixed (real along with synthetic) image dataset for human-centric perception tasks. The dataset was generated by using deep learning-generated 3D human models to populate real background images with humans in various poses and positions. In addition, starting from hand-crafted 3D human models it generated a dataset of textured SMPL-D parametric statistical body models, and made it available for direct use in simulations in Webots.

Animatable 3D human models based on the parametric SMPL+D body model.

Simulation is an important factor in the development of the OpenDR toolkit. It allows for the rapid setup of scenarios in which to test the algorithms being developed, to easily generate the data necessary for the training and testing of the models as well as to easily disseminate the results of this work.

Desktop rendering, old web rendering engine and current web rendering engine.

CYB extended the capabilities of the Webots simulator in a number of ways: ameliorating the rendering of simulations run on the cloud through porting its web rendering engine to the web using WebAssembly, converting its entire library to the default ROS coordinate system for better compatibility, adding functionalities needed in DL applications, and improving ROS/ROS2 compatibility, providing now 16 ROS2 packages.

Toolkit Integration for the Use Cases

Integration of toolkit functionalities in robotics platforms that will be used in the three project use cases has already started. PAL collaborated with AUTH in order to integrate tools such as face and mask detection and active or static face recognition inside the Jetson TX2 embedded in the TIAGo robot.

OpenDR toolkit code running in the Jetson TX2 embedded in TIAGo.

TAU has worked to integrate several tools developed in the OpenDR toolkit to the agile production use case, including a single demonstration grasping model and the speech recognition model. A human-robot collaboration scenario was developed that concerns the assembly of a Diesel engine, including tasks for both a robot (Franka Emika) and a human and the handling of various components.

Agile production use case scenario combining the single demonstration object grasping model and the speech recognition model.

The integration of the OpenDR tools for the agriculture use case will start at the beginning of 2022 and field tests will be conducted later in the same year. AGI has been in the process of improving and integrating the SLAM plant row guidance into the vision system of its Robotti agricultural robot and also building the intra row prototype and testing it.

Dissemination

Dissemination and communication are taken very seriously in OpenDR. Indeed, the consortium undertook numerous efforts in various directions to attract interest in the project findings and results. The project website and its social media accounts (Facebook, Twitter, LinkedIn, YouTube) 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 official website.

Similar to the first year, the consortium managed to generate a high volume of publications. Fourteen papers were presented or accepted in high quality, well-established international conferences (including ICME, IROS, ICIP, EUSIPCO etc.) and twelve papers were published or accepted in scientific journals (including the highly influential IEEE TIP, IEEE TRO and IJCV). In total, project partners have published so far 50 conference and journal papers, one edited book and numerous preprints.

A particularly important dissemination activity was a workshop on "Open Deep Learning Toolkit for Robotics: Towards Democratizing Artificial Intelligence" that was organized at IROS 2021. The workshop was well attended and included talks from consortium members and prominent roboticists. Other significant dissemination activities included organization of a conference special session, co-organization of an IEEE TAI special issue on AI for robotics, invited talks, participation in events like ERF 2021 or GSoC 2021 and so on. The release of the first public version of the toolkit in M24 will signal the start of a new series of activities that will focus on attracting relevant parties in getting to know and use it.

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