OpenDR
–
Open Deep Learning Toolkit for Robotics

Project start date: 01.01.2020
Duration: 36 months
Lead contractor: Aristotle University of Thessaloniki

Deliverable D9.5: First report on dissemination, exploitation plans and project newsletter

Date of delivery: 31 Dec2020

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

Version: 4.0

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

Title:	**D9.5:First report on dissemination, exploitation plans and project newsletter**
Project:	**OpenDR** (ICT-10-2019-2020 RIA)
Nature:	Report **Dissemination Level: PU**blic
Authors:	Anastasios Tefas (AUTH), Nikos Nikolaidis (AUTH), Nikolaos Passalis (AUTH), Maria Tzelepi (AUTH), Paraskevi Nousi (AUTH), Pavlos Tosidis (AUTH), Eleni Chantzi (AUTH), Efstratios Kakaletsis (AUTH), Charalampos Symeonidis (AUTH) Alexandros Iosifidis (AU), Lukas Hedegaard Morsing (AU), Halil Ibrahim Ugurlu (AU), Stefania Pedrazzi (CYB), Moncef Gabbouj (TAU), Jenni Raitoharju (TAU), Alea Scovill (AGI), Alessandro Di Fava (PAL), Sarah Terreri (PAL), Carlos Vivas (PAL), Bas van der Heijden (TUD), Osama Mazhar (TUD), Laura Ferranti (TUD), Jens Kober (TUD), Robert Babuska (TUD), Tim Welschehold (ALU-FR), Abhinav Valada (ALU-FR)
Lead Beneficiary:	AUTH (Aristotle University of Thessaloniki)
WP	9
Doc ID:	OPENDR_D9.5.pdf

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

Document History

Version	Date	Reason of change
1.0	24/9/2020	First draft (document structure)
2.0	6/12/2020	First complete draft
3.0	20/12/2020	Document ready for internal review
4.0	28/12/2020	Final version, incorporating the reviewer comments, ready for submission

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

Table of contents

1 Contents

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

Executive Summary

Dissemination and communication activities 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 its impact. Exploitation is obviously also very important for a project. This deliverable, D9.5 “First report on dissemination, exploitation plans and project newsletter”, is a document that presents the project dissemination activities within the first 12 months of the project, as well as some dissemination performance indicators. It contains information regarding scientific publications, project-related presentations at various events, keynote talks and tutorials, as well as other means of communication and dissemination used by the project, such as the project website and social media channels or articles in magazines and newspapers. Moreover, the deliverable includes a description of the OpenDR initial exploitation plan. The 1st project newsletter is also included in this deliverable.

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 harnessing Deep Learning to provide advanced perception and cognition capabilities, meeting in this way the general requirements of robotics applications in the applications areas of healthcare, agri-food and agile production. As is natural for any collaborative project, dissemination and communication activities are very important for OpenDR, since they will make project results and findings known to the industry, academia and the general public, thus maximizing its impact. This deliverable is a document that presents the project dissemination activities within the first 12 months of the project. It contains (Section 2) information regarding scientific publications, organization of special sessions and a workshop, project-related presentations at various events, keynote talks and tutorials, as well as information regarding the project website and social media channels, dissemination material that has been created, articles in magazines and newspapers, etc. Information regarding the performance of the dissemination and communication activities so far, 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 obviously also very important for a project, and especially for its industrial partners. The deliverable includes (Section 3) a description of the OpenDR initial exploitation plan. Finally, the first project newsletter, a major element in the project's communication strategy, is also included in this deliverable.

2. Dissemination and Communication Activities

1.1 Short Description of Dissemination & Communication Plan

A detailed dissemination and communication plan for the project has been set out in the DoA. The basic elements of this plan are summarized in the following sections.

1.1.1 Dissemination plan

According to the DoA, the objectives of the dissemination activities, led by AUTH with contributions from all partners, will be to:

- **Plan**: Identify targets, messages, tools, and channels. Build an adequate and effective communication and dissemination plan to ensure the best impact of project results.
- **Design and produce dissemination tools**: Design a comprehensive set of communication material and tools to ensure an easy identification of the project and a major exposure.
- **Distribute and represent**: Use the dissemination channels. Organise project events and participate in workshops, conferences, and international/EC meetings.
- **Sustain:** Ensure a persistent and long-lasting visibility of the project activities and outcomes.

To maximize impact, the project consortium adopted a three-level strategy involving online and interactive tools and channels, non-electronic tools and channels, and physical interactive tools and channels:

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

A. Online dissemination and interaction:

- **Project website:** The project website plays the following roles:
  - Public image of the project and main online access point for the different target groups;
  - Information source: highlight project objectives, activities, outcomes and relevant updates
  - Repository of information: store and make available project resources and publications to targeted audiences.
  - Showcase project achievements in web-based interactive robot simulations where users can assess the robustness of the Deep Learning systems.
- **Software repository:** The software developed by OpenDR is to become publicly available in the well-known open-source repository, GitHub, the world's leading software development platform, ensuring a wide visibility. Links from the project website to the software repository (and vice versa) will ensure maximization of public visibility for both.
- **External channels:** Project results and activities are to be disseminated on external websites, such as consortium partners' websites or collaborating initiatives webpages, so as to increase awareness.

B. Offline dissemination:

- **Project public deliverables:** OpenDR is to produce (and already did so) a wide range of public deliverables, including the versions of the OpenDR toolkit. All of them will be made available on the project website in order to disseminate knowledge as widely as possible.
- **Project publications:** OpenDR actively and eagerly pursues publication of the project's scientific findings in prestigious international journals, conferences, and workshops. OpenDR's partners might also publish research articles containing breakthrough project results with simplified technical content in a way that is comprehensible to the general public.

C. Physical interactive dissemination:

- **Training activities:** OpenDR training and knowledge dissemination activities will be mostly focused on young researchers, postgraduate students and industry professionals.
- **Workshops/Symposia:** OpenDR's dissemination to researchers, research managers, company/industrial representatives and staff, is to be pursued through the organization of workshops and public fora.
- **Participation in exhibitions:** The OpenDR dissemination plan includes promoting the project and its findings by displaying prototypes in international exhibitions through the booths and other promoting activities of its consortium members.

These activities aim towards a number of dissemination target groups identified by OpenDR's partners and described in the DoA. These groups are presented in Table 1 along with the dissemination tools to be used per identified target group.

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

	MAIN TARGET GROUPS				
Dissemination tools and channels	**Policy makers**	**Research & Education Community**	**Private Sector**	**Related projects and initiatives**	**General public**
Website	✓	✓	✓	✓	✓
Publications		✓	✓	✓	
Open software & data	✓	✓	✓	✓	
Project public deliverables	✓	✓	✓	✓	
Training activities		✓	✓	✓	
Workshops	✓	✓	✓	✓	✓
Participation in exhibitions	✓	✓	✓	✓	✓

Table 1: Dissemination tools, channels and target groups

To maximize impact, constant monitoring of the effectiveness of the dissemination strategy through systematic collection of data was planned in the DoA. A number of dissemination key performance indicators (KPIs) against which to compare the measured performance were also set.

1.1.2 Communication plan

Apart from disseminating project results among the research community, policy-makers, and the private sector, the project vowed to make strong efforts towards communicating project information to the general public in an easy-to-understand, non-technical fashion including key facts, objectives and expected results, events and outputs.

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

Towards this direction, the project DoA set out a communication strategy whose main objectives are to:

- Raise awareness and ensure maximum visibility of the project key facts, objectives, activities and findings among targeted audience.
- Announce and promote OpenDR events and dissemination material.
- Support the dissemination objectives.

According to the communication plan, AUTH leads the project communication activities. The following elements and tools of the communication strategy were described in the DoA:

- **Logo and graphical identity:** The logo will give to the OpenDR project a unique independent identity that will be easily recognizable to communication-targeted groups. Based on the logo colour palette and fonts, a set of templates were set to be designed for project deliverables and presentations.
- **Media-kit:** A set of project promotional material is to be developed and distributed through various channels, including the project website, for publicity use. This media kit, including i) a project fact sheet, ii) a brochure, as well as iii) posters and iv) roll-up banners, will allow the project consortium to reach large audiences.
- **Video:** One short OpenDR video is to be produced to promote the project vision, objectives and challenges.
- **Website:** The project website, shortly described in the dissemination plan overview above, will serve as a portal through which continuous updates about the project progress and the results are presented and communicated to the general public.
- **Social media:** OpenDR communication plan includes the creation of project accounts on the most popular social media services such as Facebook, Twitter and LinkedIn. These will serve as a platform to disseminate, communicate, discuss, comment, consult and suggest research and policy topics with different target groups at different levels.
- **Newsletters:** Project newsletters are to be released on an annual basis and will enable the consortium to update the project community with latest project activities and results.
- **Press releases:** Press releases, news-bulletins, press conferences and interviews are to be organized. At least three press releases are foreseen in OpenDR's promotion strategy and will be provided by the industrial consortium members.

Table 2 shows a summary of the communication tools and the groups targeted by each tool (note that some tools are also used for dissemination purposes).

	MAIN TARGET GROUPS				
Communication tools and channels	**Policy makers**	**Research & Education Community**	**Private Sector**	**Related projects and initiatives**	**General public**
Visual identity	✓	✓	✓	✓	✓
Factsheet		✓	✓	✓	
Poster		✓	✓	✓	
Brochure		✓	✓	✓	

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

Website	✓	✓	✓	✓	✓
Social Networks	✓	✓	✓	✓	✓
Videos	✓	✓	✓	✓	✓
Newsletter	✓	✓	✓	✓	
Press releases	✓	✓	✓	✓	✓

Table 2: Communication tools, channels and target groups

Similar to dissemination activities, a series of Key Performance Indicators (KPIs) were identified for the communication activities and were detailed in the DoA.

1.2 Dissemination and Communication Activities M1-M12

A wealth of dissemination and communication activities took already place during the first 12 months of the project, a fact that denotes that the project takes these activities very seriously. Moreover, all WP9 deliverables scheduled for the first year, namely D9.1, D9.2. D9.3, D9.4 and this deliverable (D9.5), were delivered on time. In summary, the project undertook the following activities in this area:

- Creation and population of the project website.
- Establishment of the various social media accounts and authoring of news items for those channels.
- Creation of dissemination material including the project fact sheet, brochure, poster, promotional video, etc.
- Creation of a large amount of high-quality scientific publications namely 7 journal articles, 16 conference papers and one edited book (at the authoring stage).
- Organization of one training workshop, two special sessions in conferences and participation of a consortium member in the organizing committee of a major conference.
- Dissemination of OpenDR aims and findings through 2 invited/keynote talks and one tutorial.
- Participation of consortium members with talks into a number of conferences exhibitions, EU events, or industrial/professional workshops such as ERF 2020, IROS 2020, AI & Big Data Conference, etc.
- Collaboration with one DIH.
- Creation of the first newsletter.

Details are provided in the sections below. It should be noted here that the project GitHub page, which will be another major element of the project's dissemination strategy since it will host the OpenDR toolkit, will be available close to the toolkit release date, namely M18.

1.2.1 Project website

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

The website of the OpenDR project (available at https://opendr.eu/), was created by AUTH in the first month of the project and is the main tool for the project's dissemination activities. The website was officially delivered in M2, as part of Deliverable D9.3. It provides information about OpenDR's core ideas, partners involved, work structure, news, publications, etc. The website is maintained and kept up-to-date by AUTH, with contributions from the partners. Several pieces of information have been published since the beginning of the project (public deliverables, publications, events, news, etc.) while more than 3200 users have visited the website until now.

Through the top menu bar the user has the ability to choose different sections in order to be informed about the project. These sections include:

- news
- project information (consortium, work structure, communication material)
- results (deliverables, publications, videos)
- events
- contact information

At the right side of the menu bar, there are three clickable icons of the project's social media accounts (Facebook, Twitter, LinkedIn) that can lead the user to each account respectively.

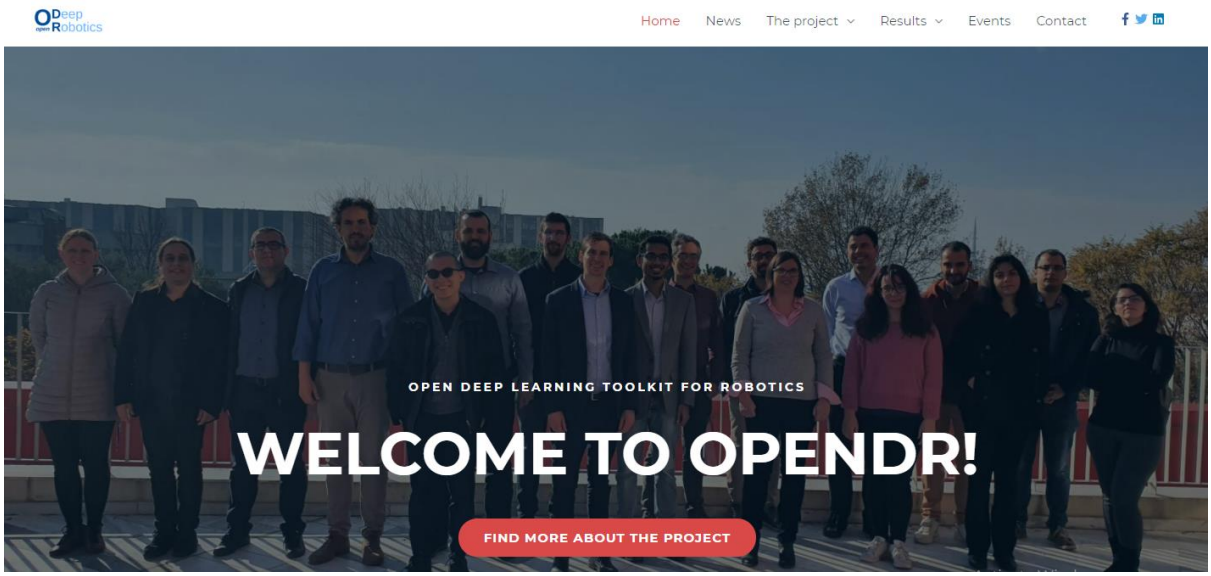

Figure 1: The header, menu and the landing page of the website

Below the main part of the home page (Fig. 1) there is a set of the latest most important news and events related to the project's activities, such as recent presentations, consortium meetings, organized events, etc. (Fig. 2).

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

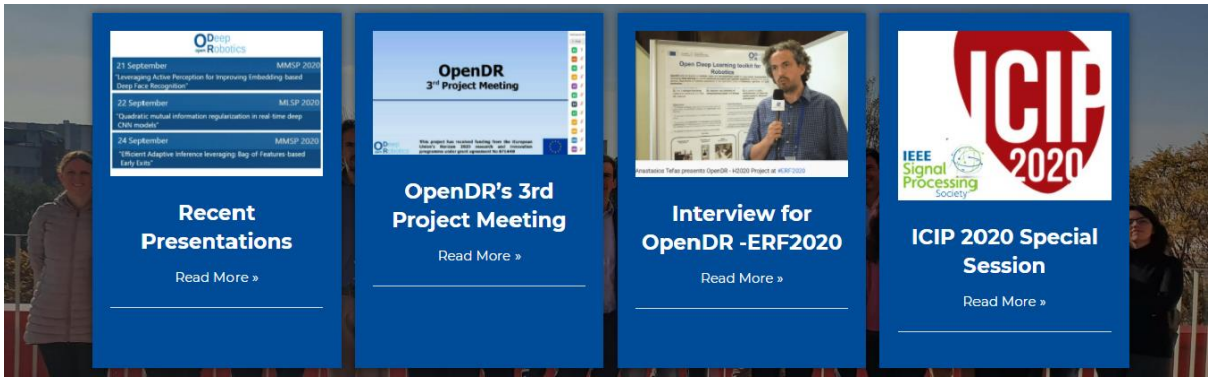

Figure 2: A set of the project's latest news and events

An overview of the latest tweets, extracted automatically from the project's Twitter account, is placed below the latest news panel, as can be seen in Fig. 3.

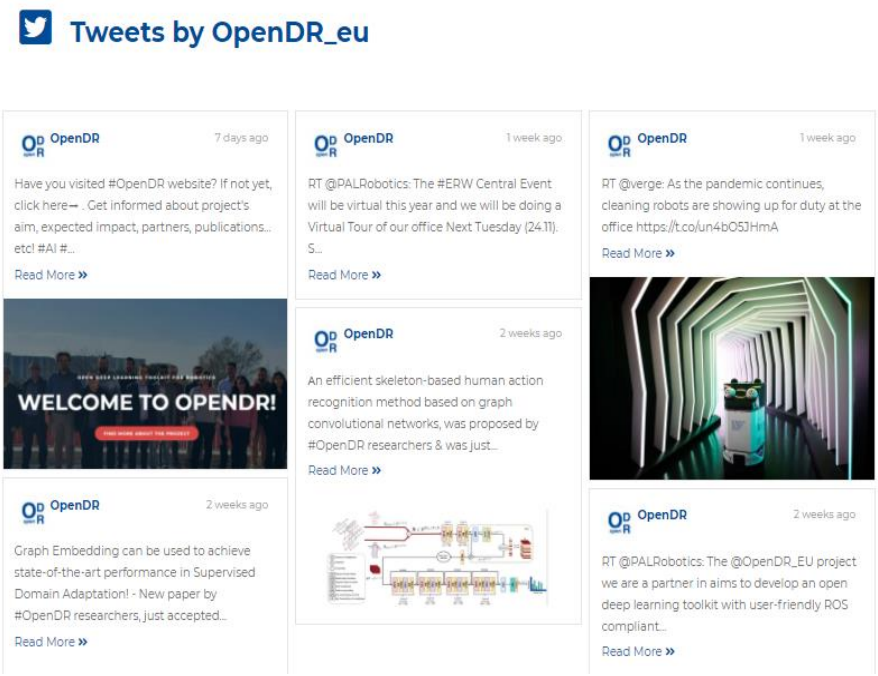

Figure 3: OpenDR's latest tweets feed

The footer of the website, available at the bottom of every page, provides general information about funding, coordination, as well as privacy policy and contact information (Fig. 4).

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

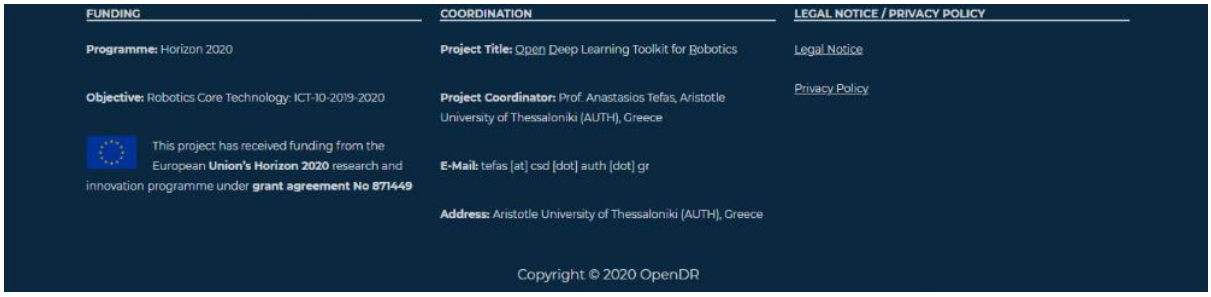

Figure 4: OpenDR's website footer

Users who wish to find out more information about the project, for example papers, deliverables, presentations, etc., can visit the relevant pages of the website (Fig. 5, Fig. 6) through the menu bar. Additional information regarding the project website can be found in Deliverable D9.3.

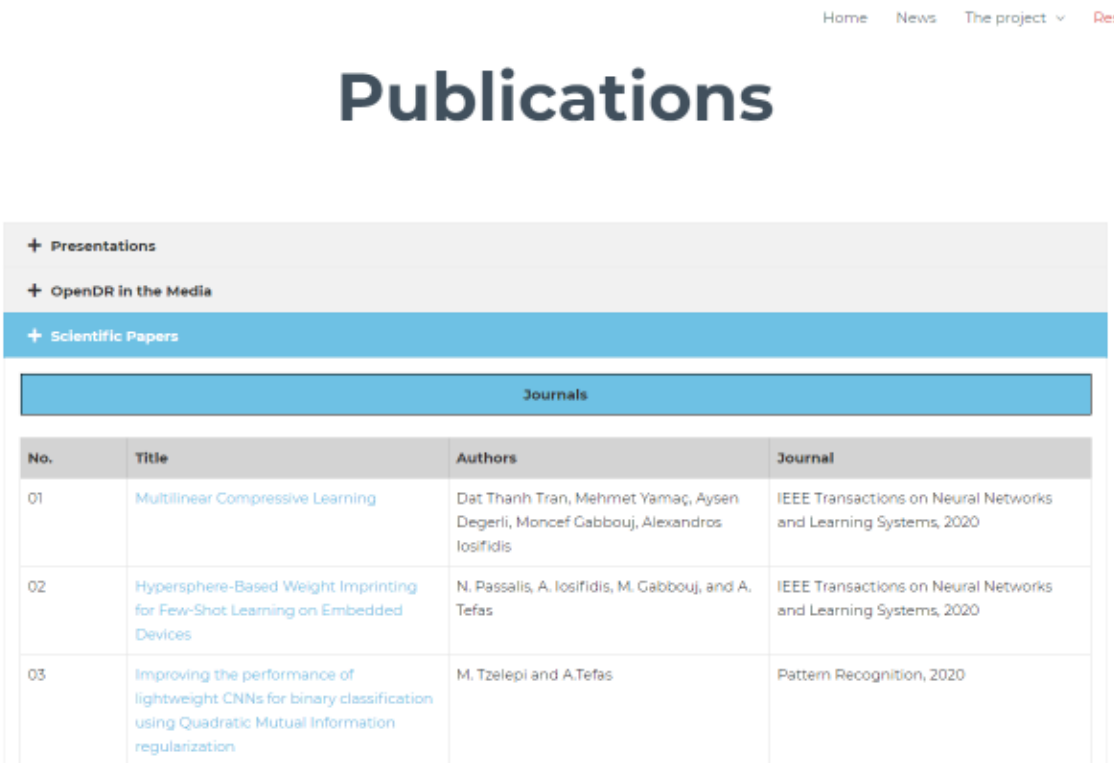

Figure 5: The Publications page of the project website

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

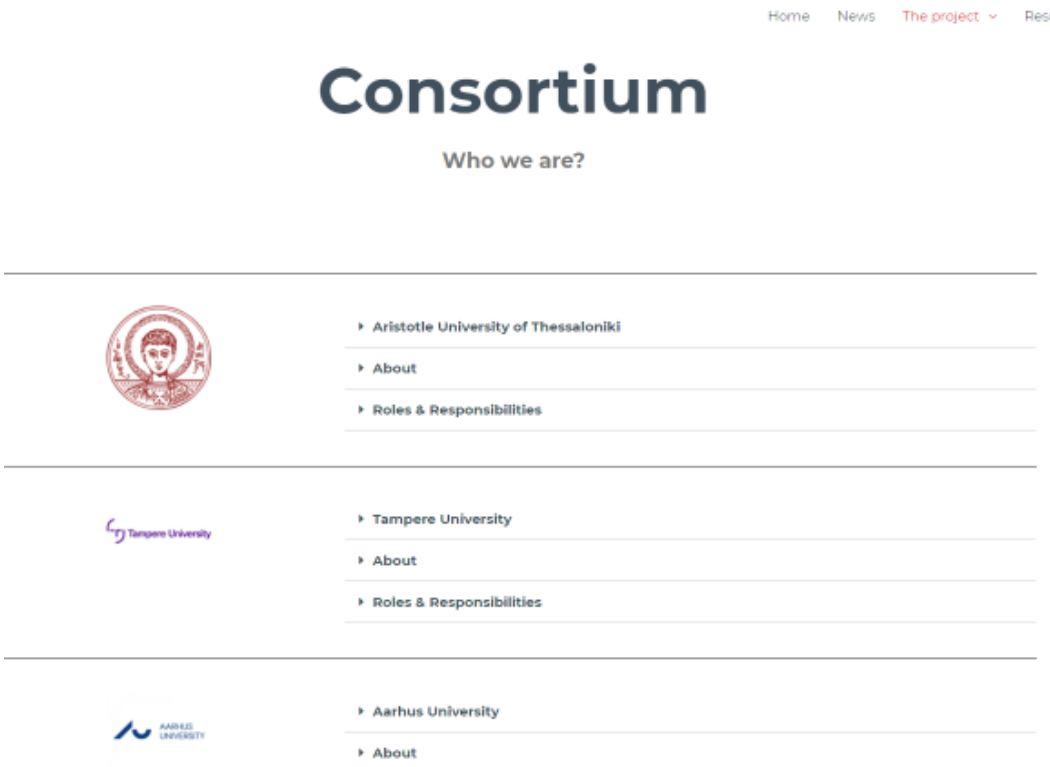

Figure 6: The Consortium page of the website

1.2.2 Social media

OpenDR has been very active in the most popular dissemination channels, namely LinkedIn, Facebook, Twitter and YouTube, where AUTH has created project accounts since the first month of the project. Social media accounts creation and setup is described in the respective Deliverable, namely D9.3, delivered on time in M2. AUTH is responsible for continuous and active social media content update, while the other partners are supportive through the provision of information related to dissemination activities. Furthermore, OpenDR partners that maintain social media accounts support the project's dissemination by releasing OpenDR-related content in their corporate accounts. The response of the audience until today could be considered as more than satisfactory, as shown by the number of followers of each account.

Twitter (@ OpenDR_EU)

The Twitter account of the project can be found in the following link: https://twitter.com/OpenDR_EU. The purpose of this social media account is the fast distribution of information about the project's actions and achievements (Fig. 7) A large amount (more than 90) tweets/retweets have been posted until today, on topics such as new project publications, presentations of project results at conferences and other events, organization of events (e.g. special sessions) from consortium members, etc. (Fig. 8). An important number of tweets deal with news and articles from around the web, related to the project. Through this activity, the Twitter account has attracted so far almost 150 followers, a very satisfactory figure given that the project is still in an early stage.

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

Figure 7: OpenDR's twitter account

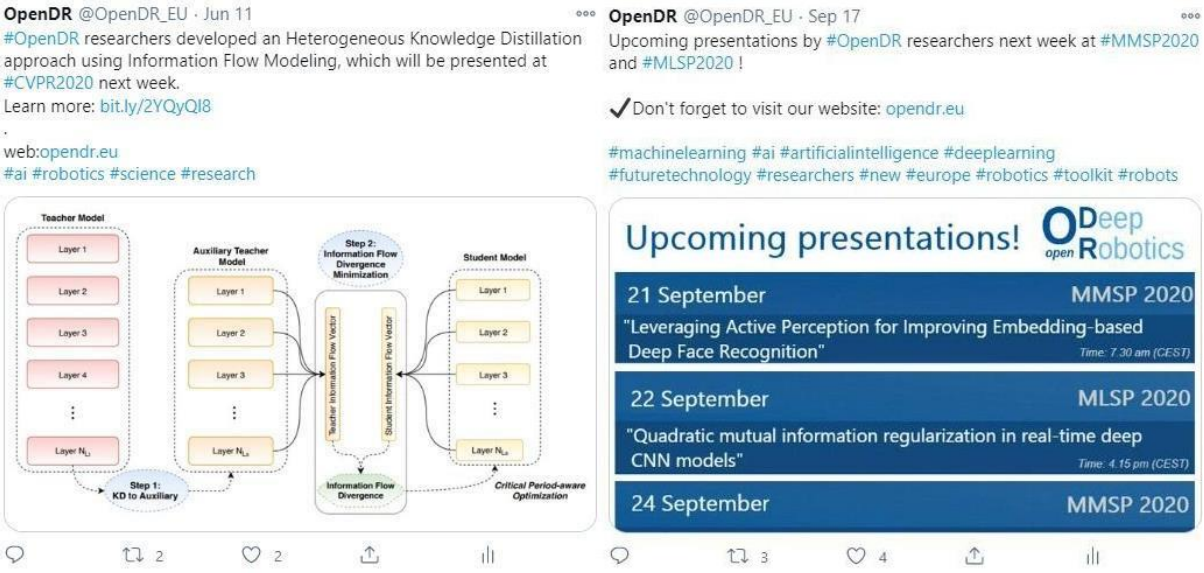

Figure 8: Examples of tweets

Facebook (@Opendr.eu)

In order to raise awareness and ensure maximum visibility of the project, a Facebook account was created under the name "**OpenDR**", which can be found at https://www.facebook.com/Opendr.eu/ (Fig. 9). Until today more than 70 posts have been

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

created, aiming to provide information about project activities/publications, as well as various developments in the fields of Robotics and Deep Learning. The current number of account followers is 175 while the page has attracted so far more than 170 likes. Both figures can be considered as more than satisfactory, given the fact that the project has not reached yet its major expected contribution, namely the OpenDR Toolkit. Examples of Facebook posts are presented in Fig. 10.

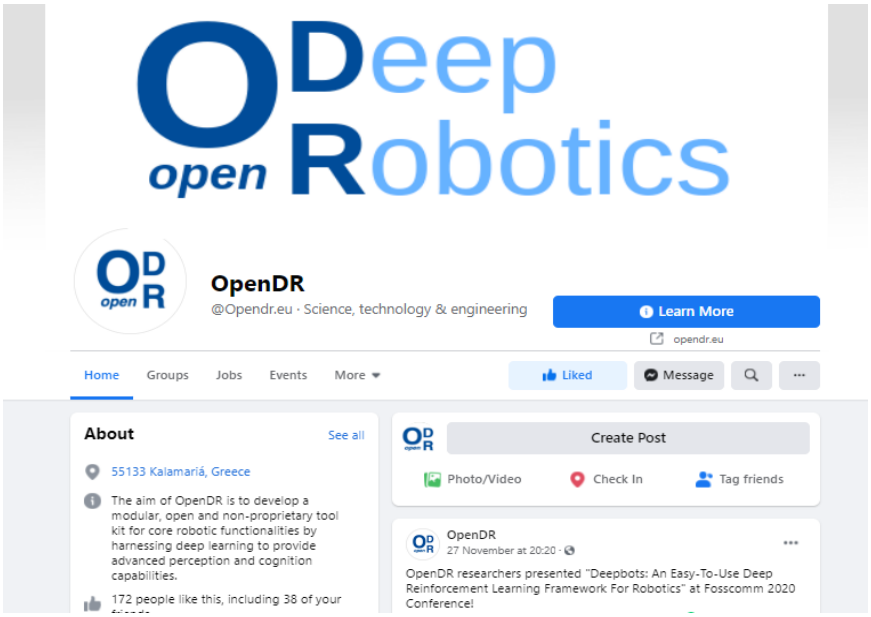

Figure 9: OpenDR's Facebook account

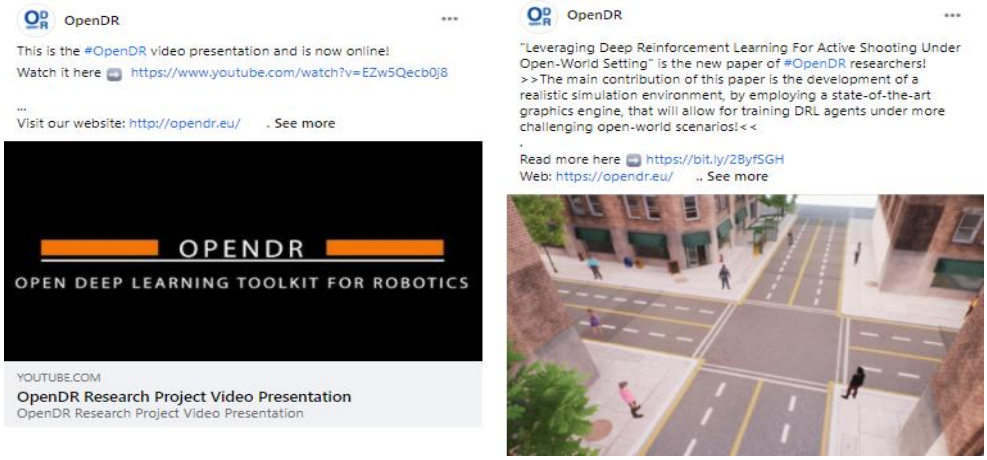

Figure 10: Examples of Facebook posts

LinkedIn

For further dissemination of the project activities and results, a LinkedIn group has been created under the name **"OpenDR H2020 Research project"** and is available at https://www.linkedin.com/groups/13807356/ (Fig. 11). The basic idea behind this group's

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

creation is the existence of an online team, whose members exchange ideas about the project and Deep Learning in general, thus maintaining their interest in the project. The way someone can become a group member is either by sending a "Request to join" or by receiving an invitation from another group member. The group has currently 112 members.

What is more, a LinkedIn page has been created, as it can attract a wide number of followers easily through the public posts (Fig. 13). The LinkedIn page of the project can be found at https://www.linkedin.com/company/opendr-research-project, under the name **"OpenDR Research Project"** (Fig. 12). Currently more than 185 people follow this page.

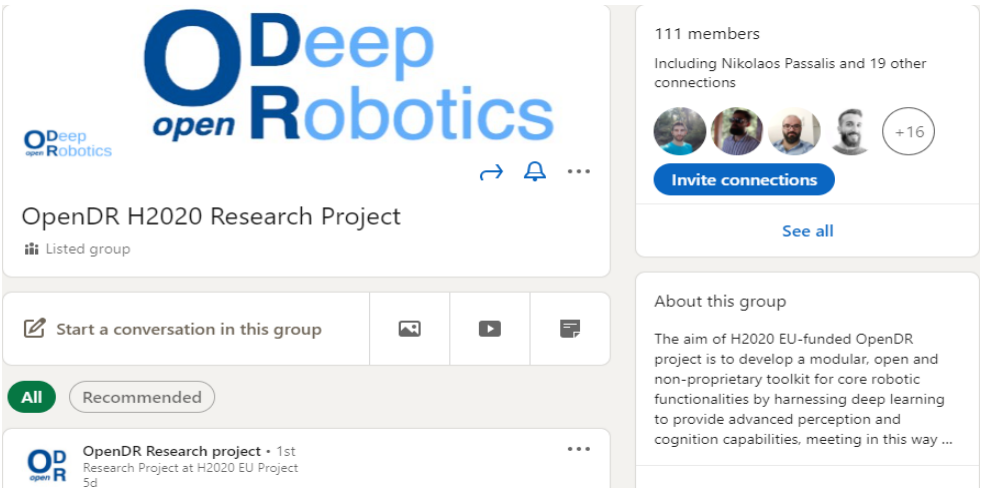

Figure 11: OpenDR's LinkedIn group

Figure 12: OpenDR's LinkedIn page

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

Until today, a large number of posts have been sent out through the LinkedIn group and page, some of which are presented in Fig. 17. Many of these posts disseminate the project publications as well as researchers' participation in important international conferences.

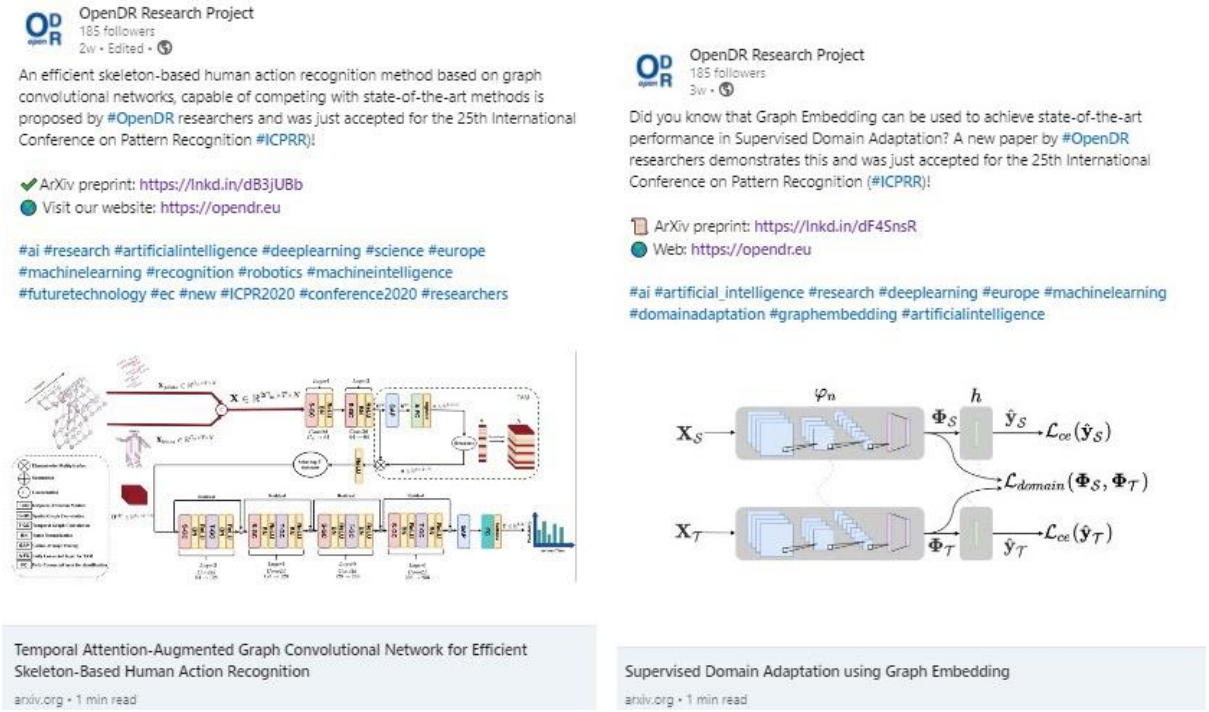

Figure 13: Examples of LinkedIn posts

YouTube channel

A YouTube channel has also been created, despite the fact that no such plan existed in the DoA. At its current stage, the channel contains video presentations of OpenDR's conference publications (including CVPR, MLSP, MMSP, AIAI and ICME) as well as the project's promotional video. Overall, the channel has 27 subscribers and more than 800 views. Its structure is depicted in Fig. 14 and an example of an uploaded video can be seen in Fig. 15.

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

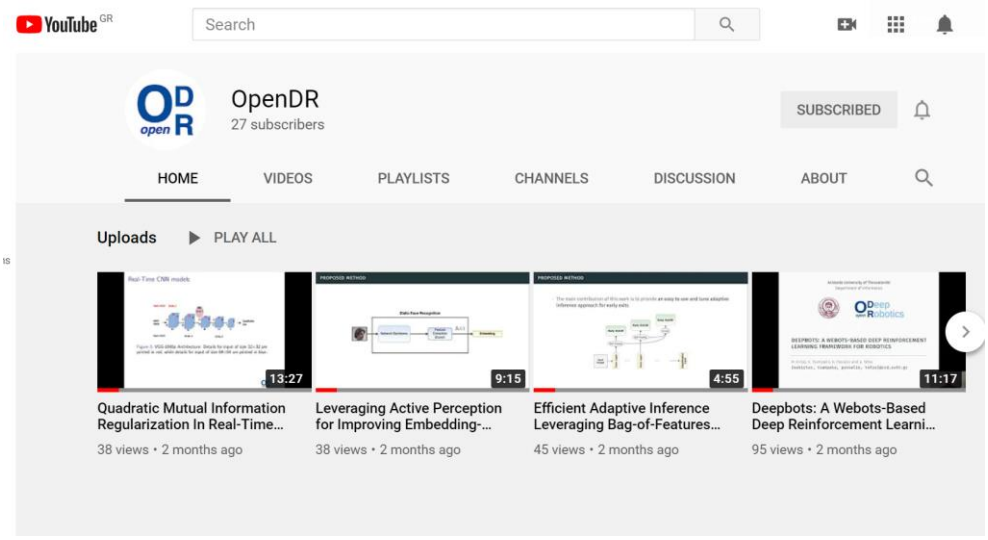

Figure 14: The YouTube channel of the OpenDR project

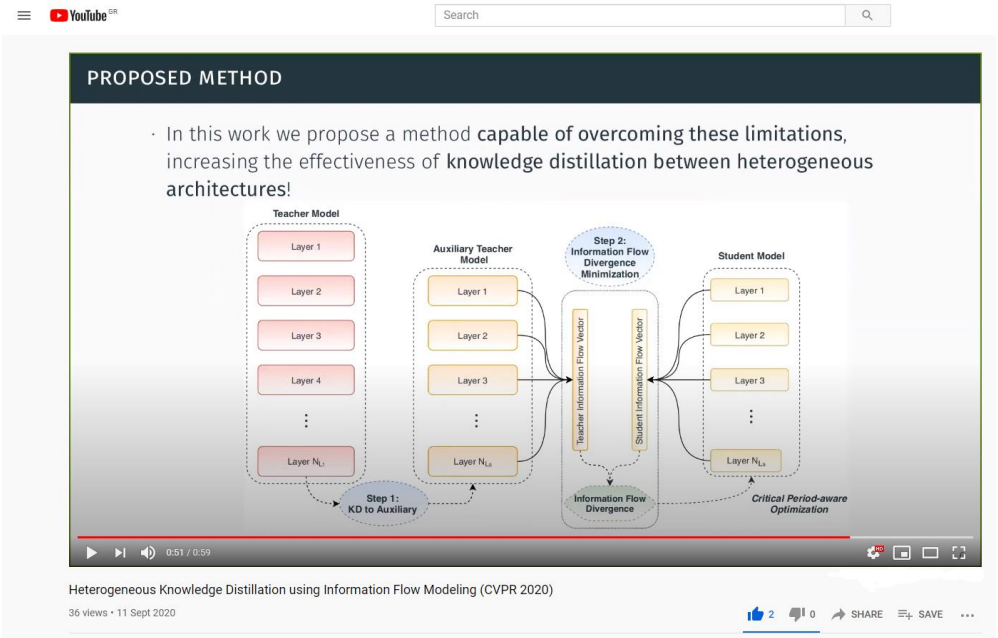

Figure 15: An example of an uploaded video in OpenDR's youtube channel

ResearchGate

In order to further support the dissemination activities of the project, an account at ResearchGate was created, through which, the scientific papers are disseminated to the users of the platform, aiming to cover all the possible dissemination channels and raise awareness for the project (Fig. 16). The project account can be found under the name "OpenDR H2020 Research Project".

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

Figure 16: The ReserchGate account of the OpenDR project

Partners social media activity

Apart from the activity in the "official" project channels mentioned above, OpenDR partners were also active into publicizing the project as well as related events and achievements. For example, partner PAL authored 6 OpenDR related posts in its Facebook account, 4 items in its LinkedIn account and posted 13 tweets, again in its Twitter corporate account. Similarly, members of the AU OpenDR team created posts in Aarhus University Engineering LinkedIn account.

Figure 17: Examples of PAL's Robotic posts

1.2.3 OpenDR webpages at partners websites

OpenDR partners are promoting the project through their websites. Indeed, the project is mentioned in the Portfolio section in the CYB website (Fig. 18), where a card is displayed describing the project and linking to the OpenDR website.

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

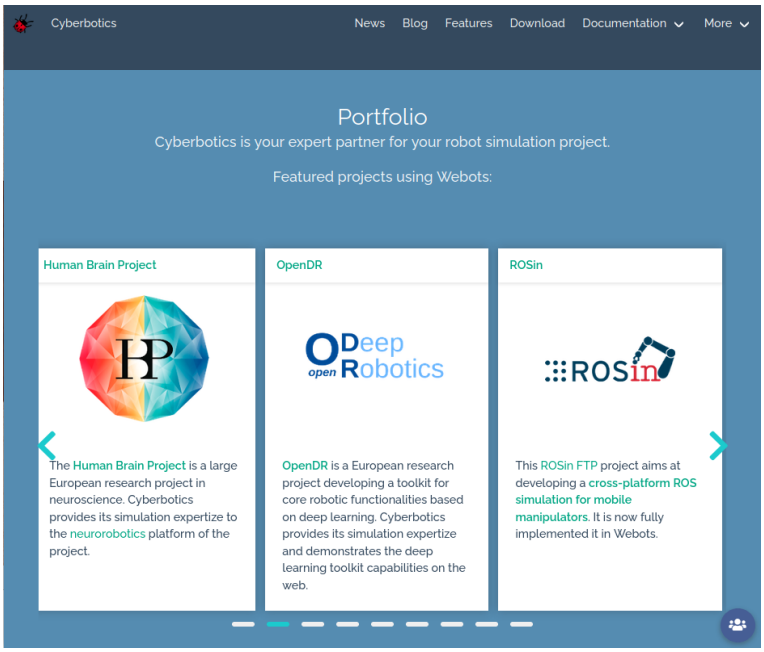

Figure 18: Portfolio section in the CYB website

Also, PAL Robotics included in its website the description of OpenDR project and posted a blog item (Fig. 19) regarding "Artificial Intelligence in TIAGo for healthcare and agile production as part of EU project OpenDR".

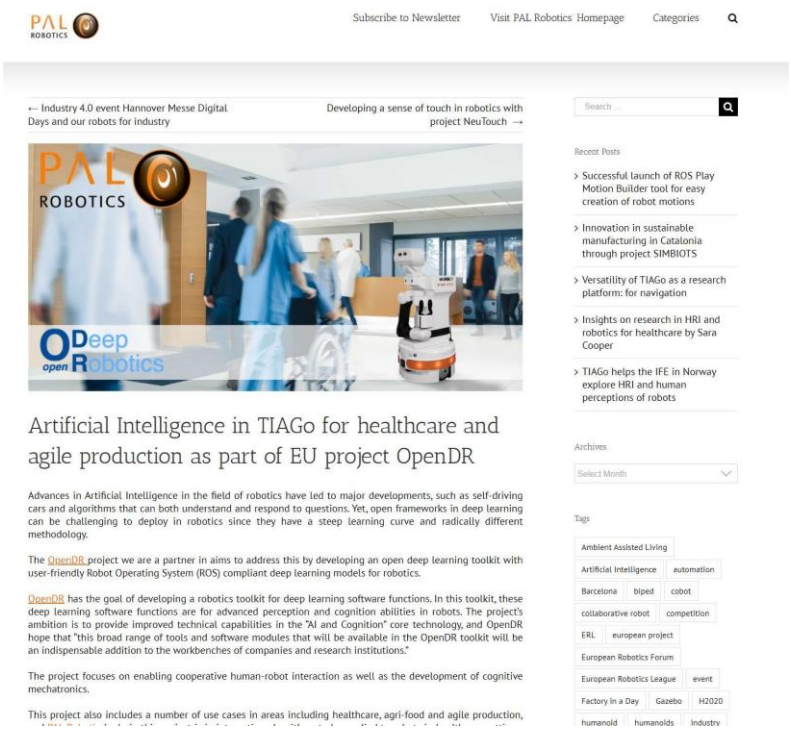

Figure 19: OpenDR post in the PAL Robotics website

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

Moreover, ALU-FR included brief descriptions of the OpenDR project on both the Autonomous Intelligent Systems and Robot Learning Lab Webpages (Fig. 20).

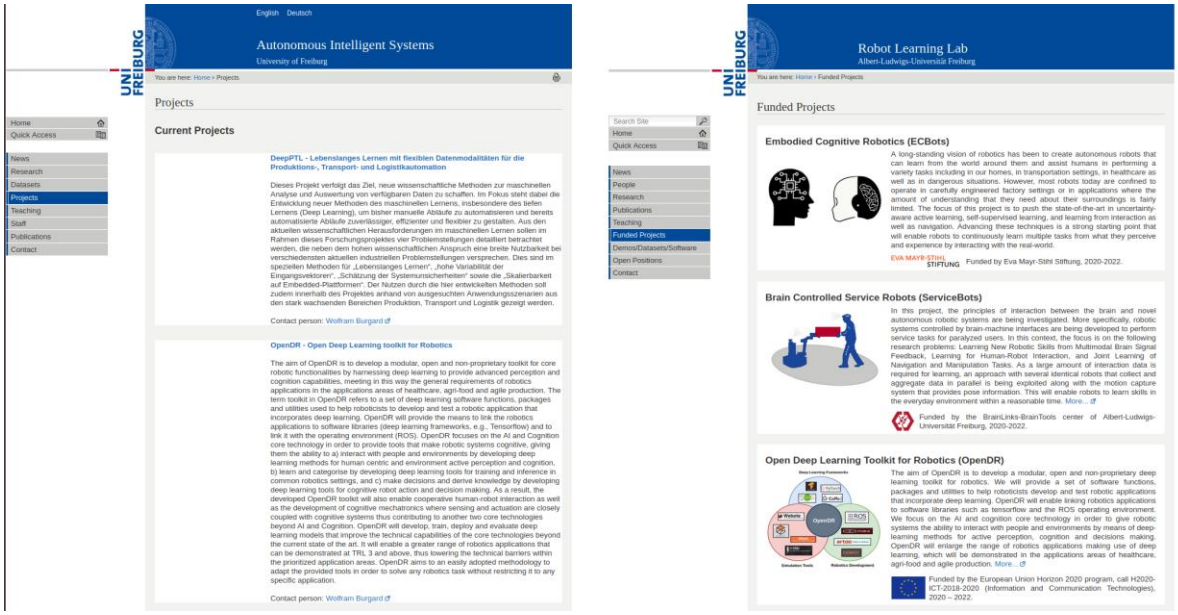

Figure 20: OpenDR description on ALU-FR's Webpages

1.2.4 Dissemination material

AUTH has created several pieces of dissemination material in order to be used by the partners in events where they participate. Unfortunately, due to COVID-19, most events were held in a virtual way and thus the use of such materials was so far rather limited.

Project logo

OpenDR's logo (Fig. 21) is the centerpiece of the project visual identity, facilitating many dissemination activities, ranging from presentations and exhibitions to the online presence of OpenDR to various social media and online source code repositories. OpenDR's consortium proposed several logo designs, among which four were shortlisted, and a voting procedure was used to select the final logo. The logo was presented as part of deliverable D9.1 The technical specifications that are related to OpenDR branding and should be followed when preparing dissemination material are also provided in this deliverable.

Figure 21: OpenDR's logo

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

Fact sheet and project presentation

Within the frameworks of the project's dissemination strategy and as promised in the DoA, a project fact sheet was created in the first month and delivered as part of deliverable D9.1. The fact sheet (Fig. 22) consists of four pages that include appealing info-graphics related to the project's objectives, consortium expected impact and social media accounts. This fact sheet has been already shared as a hard copy to the participants of the European Robotics Forum 2020 where OpenDR was present. Furthermore, it was shared to online users through the project's social media accounts and website. A project presentation (Fig. 23) has been also created to further assist the presentation of the project at various events. The presentation was also delivered as part of deliverable D9.1.

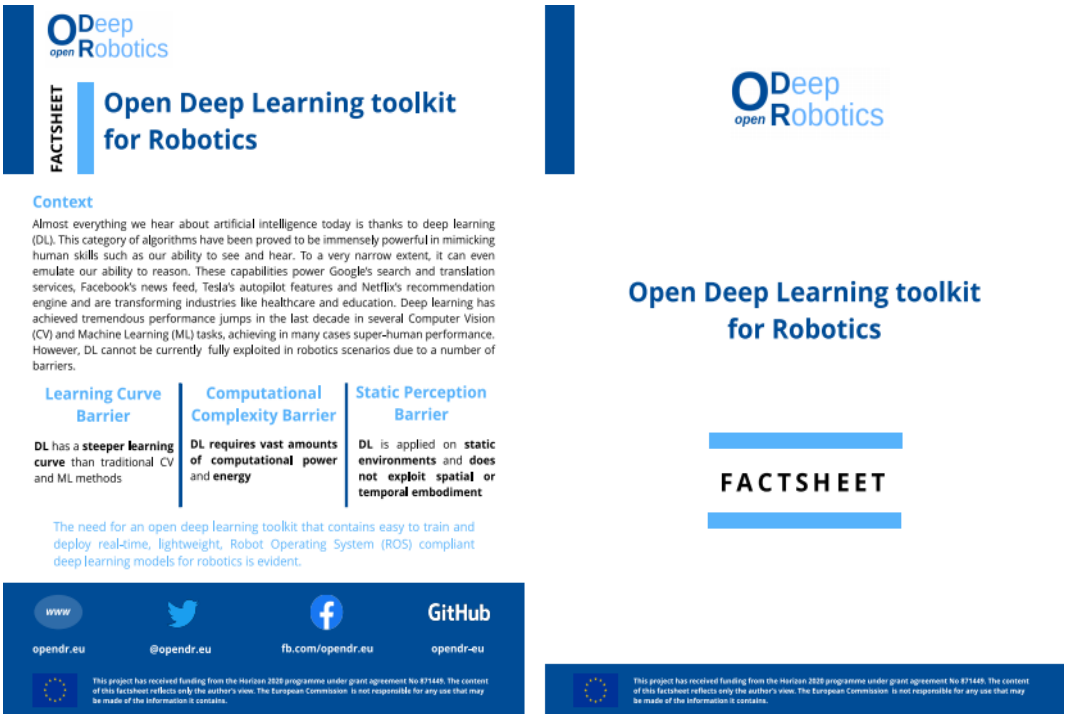

Figure 22: The first and the second page of the OpenDR fact sheet

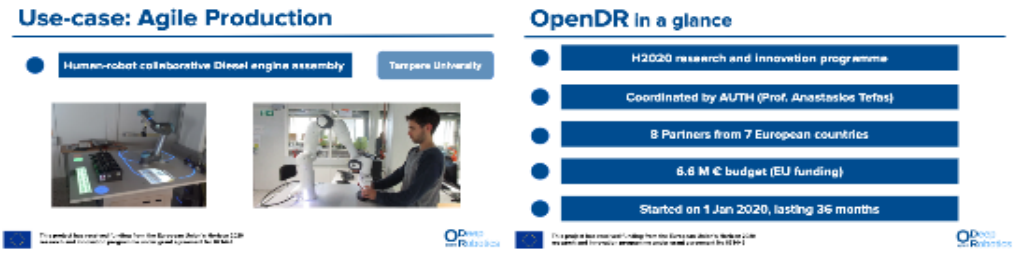

Figure 23: Slides from the OpenDR presentation

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

Poster

As an additional dissemination material, a promotional poster was designed in order to be used at events that the project will organize or contribute to. The poster provides information regarding the project's core idea, objectives, consortium, etc. Similar to the project fact sheet, the poster has been used to support the presentation of OpenDR at the European Robotics Forum 2020. The poster is available at the OpenDR website and has also been posted at the project's social media accounts (Fig. 24).

Figure 24: The project poster

Brochure

A promotional brochure (a triptych) was created in month 10 in support of the OpenDR communication activities. It includes general information about the project's main idea, expected impact, consortium and social media accounts, with special QR codes for immediate

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

access. This brochure can be shared to participants of events, talks, workshops, etc., as a direct and easy way of informing them about the project (Fig. 25, Fig. 26). Unfortunately, due to the pandemic, and similar to the other materials that have been created, its distribution has been very limited.

Figure 25: The external page of the OpenDR promotional brochure

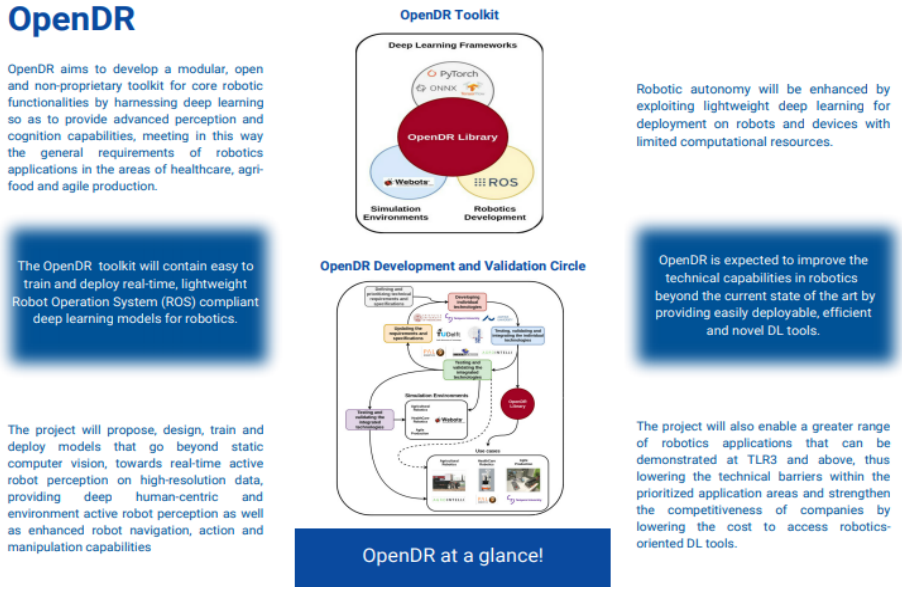

Figure 26: The internal page of the OpenDR promotional brochure

Video Presentation

A promotional video was created on the eighth month of the project, as project deliverable D9.4. The 7-minute video, which was created by AUTH with contributions (short video interviews, videos showcasing partners' technologies, etc.) from all partners, aims at providing a brief overview and introducing the viewer to the main objectives and elements of

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

the project. The video was uploaded on OpenDR's YouTube channel (created under the name "OpenDR"). It was also posted at the project's Twitter, Facebook and LinkedIn accounts and placed at the OpenDR website (Fig. 27). The video has so far attracted more than 500 views. Additional details can be found in D9.4.

Figure 27: Frames from the OpenDR video presentation

Press release

A press release, which provides an overview of the project aims and expected impact, as well as brief information regarding the project partners and the project social media channels, was created and delivered as deliverable D9.2 in M2 (Fig. 28). This press release was submitted (in Greek) to the Press Office of Aristotle University of Thessaloniki for dissemination through its communication channels and was also disseminated through OpenDR's web site (https://opendr.eu/press-release) and its social media channels.

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

Thessaloniki (Greece) February 2020 – OpenDR "Open Deep Learning for Robotics Toolkit", is a new EU 2020 Project which was launched at January 1st 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 Prof. Anastasios Tefas at Aristotle University of Thessaloniki in Greece, it will be running throughout the period of January 2020 to December 2022 and in total there are 8 partners from 7 different countries participating in the project.

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, 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, **OpenDR'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.

The **OpenDR** consortium consists of 5 top-ranked academic and research institutes: Aristotle University of Thessaloniki (Greece), Tampere University (Finland), Aarhus University (Denmark), Delft University of Technology (Netherlands), University of Freiburg (Germany) and 3 leading industry partners: CYBERBOTICS (Switzerland), PAL Robotics (Spain) and AgroIntelli (Denmark).

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

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

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

Figure 28: The first press release of the OpenDR project

1.2.5 Publications

During its first year, the project has produced a large number of high quality scientific publications, as detailed in the following subsections.

1.2.5.1 Books

Following an invitation by Elsevier, Profs. Alexandros Iosifidis (AU) and Anastasios Tefas (AUTH) initiated the process for editing a book focusing on the topics of OpenDR in summer 2020. The title of the book is "Deep Learning for Robot Perception and Cognition". With this book, the editors hope to contribute to the literature with a complete collection of topics for understanding the conceptual and mathematical background needed for approaching a large number of robot perception and cognition tasks from an end-to-end learning point of view. The book will serve as a textbook useful for academics, industry, and practitioners in Robot Vision, Intelligent Control, Mechatronics and Deep Learning focusing on Robotic Perception and Cognition Tasks. Partners of OpenDR will contribute in this effort by contributing chapters falling within their expertise and including, among others, research results from the project. The current plan is that the book will be formed by 22 chapters (20 of

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

them authored by OpenDR members) and will have a length of 425-500 pages, and that it will be published in autumn/winter 2021. The (working) Table of Contents, along with the contributing partners of OpenDR, is the following:

Chapter 1: Introduction (AU and AUTH)
Chapter 2: Neural Networks and the Backpropagation (AUTH)
Chapter 3: Convolutional Neural Networks (TAU)
Chapter 4: Graph Convolutional Neural Networks (AU)
Chapter 5: Recurrent Neural Networks (AUTH)
Chapter 6: Deep Reinforcement Learning (AUTH)
Chapter 7: Lightweight Deep Learning (AUTH)
Chapter 8: Knowledge Transfer (AUTH)
Chapter 9: Progressive and Compressive Deep Learning (TAU and AU)
Chapter 10: Representation Learning and Retrieval (AUTH)
Chapter 11: Object Detection and Tracking (TAU and AU)
Chapter 12: Semantic Scene Segmentation for Robotics (ALU-FR)
Chapter 13: 3D Object Detection and Tracking (AU)
Chapter 14: Activity Recognition (AU)
Chapter 15: Autonomous Navigation and Planning in Drone Racing (AU)
Chapter 16: Tactile and Haptic Perception for Robots
Chapter 17: Deep Learning for Collaborative Robots in Agile Production (TAU)
Chapter 18: Deep Learning in Multi-agent Systems
Chapter 19: Simulation Environments (AUTH)
Chapter 20: Biosignal time-series analysis (TAU)
Chapter 21: Medical Image Analysis (AU)
Chapter 22: Deep Learning for Robotics examples using OpenDR (AUTH)

1.2.5.2 Journal papers

Research conducted within the project during this period led to 7 (peer reviewed) journal articles that have been accepted for publication or were published. The full text of these publications can be found in the relevant technical deliverables.

A joint paper by partners TAU and AU, was published in IEEE Transactions on Neural Networks and Learning Systems:

- Dat Thanh Tran, Mehmet Yamaç, Aysen Degerli, Moncef Gabbouj, Alexandros Iosifidis, “Multilinear Compressive Learning”, IEEE Transactions on Neural Networks and Learning Systems, (Early Access), pp. 1-13, doi: 10.1109/TNNLS.2020.2984831, 2020 (Impact Factor: 8.793).

In this paper, the authors propose Multilinear Compressive Learning, a framework that takes into account the tensorial nature of multi-dimensional signals in the acquisition step and builds the subsequent inference model on the structurally sensed measurements. Theoretical complexity analysis shows that the proposed framework is more efficient compared to its vector-based counterpart in both memory and computation requirements. With extensive experiments, it is empirically shown that the Multilinear Compressive Learning framework

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

outperforms the vector-based framework in object classification and face recognition tasks, and scales favorably when the dimensionalities of the original signals increase, making it highly efficient for high-dimensional multi-dimensional signals.

Prof. A. Iosifidis (AU) was the sole author of two papers published in IEEE Access and Pattern Recognition Letters:

- Alexandros Iosifidis, "Probabilistic Class-Specific Discriminant Analysis", IEEE Access, vol. 8, pp. 183847-183855, doi: 10.1109/ACCESS.2020.3029514, 2020.

In this paper, the author formulates a probabilistic model for class-specific discriminant learning. The proposed model can naturally incorporate the multi-modal structure of the negative class, which is neglected by existing class-specific methods. Moreover, it can be directly used to define a probabilistic classification rule in the discriminant subspace. It is also shown that existing class-specific discriminant analysis methods are special cases of the proposed probabilistic model and, by casting them as probabilistic models, they can be extended to class-specific classifiers. Experimentally, it is shown that when combined with data representations obtained by using deep neural networks, i.e. under a transfer learning setting, the proposed probabilistic classifier leads to competitive performance compared to related methods.

- Alexandros Iosifidis, "Class Mean Vector Component and Discriminant Analysis", Pattern Recognition Letters, vol. 140, pp. 207-213, doi: 10.1016/j.patrec.2020.10.014, 2020.

The kernel matrix used in kernel methods encodes all the information required for solving complex nonlinear problems defined on data representations in the input space using simple, but implicitly defined, solutions. Spectral analysis on the kernel matrix defines an explicit nonlinear mapping of the input data representations to a subspace of the kernel space, which can be used for directly applying linear methods. However, the selection of the kernel subspace is crucial for the performance of the proceeding processing steps. In this paper, the author proposes a component analysis method for kernel-based dimensionality reduction that optimally preserves the pairwise distances of the class means in the feature space. An extensive analysis on the connection of the proposed criterion to those used in kernel principal component analysis and kernel discriminant analysis is provided, leading to a discriminant analysis version of the proposed method. The conducted analysis also provides more insights on the properties of the feature spaces obtained by applying these methods. Experimentally, it is shown that when combined with data representations obtained by using deep neural networks, i.e. under a transfer learning setting, the proposed component analysis method leads to competitive performance compared to related methods.

A joint paper by AU, TAU and AUTH was published in the highly influential IEEE TNNLS.

- N. Passalis, A. Iosifidis, M. Gabbouj, and A. Tefas, "Hypersphere-Based Weight Imprinting for Few-Shot Learning on Embedded Devices", IEEE Transactions on Neural Networks and Learning Systems, (Early Access), pp. 1-6, doi: 10.1109/TNNLS.2020.2979745, 2020.

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

Weight imprinting (WI) was recently introduced as a way to perform gradient descent-free few-shot learning. Due to this, WI was almost immediately adapted for performing few-shot learning on embedded neural network accelerators that do not support back-propagation, e.g., edge tensor processing units. However, WI suffers from many limitations, e.g., it cannot handle novel categories with multimodal distributions and special care should be given to avoid overfitting the learned embeddings on the training classes since this can have a devastating effect on classification accuracy (for the novel categories). In this article, the authors propose a novel hypersphere-based WI approach that is capable of training neural networks in a regularized, imprinting-aware way effectively overcoming the aforementioned limitations. The effectiveness of the proposed method is demonstrated using extensive experiments on three image data sets.

Finally, three papers (including two joint ones) were published in the highly cited Pattern Recognition journal.

- M. Tzelepi and A.Tefas, "Improving the performance of lightweight CNNs for binary classification using Quadratic Mutual Information regularization", Pattern Recognition, vol: 106, pp. 107407, doi: 10.1016/j.patcog.2020.107407, 2020.

In this paper, the authors propose regularized lightweight deep convolutional neural network models, capable of effectively operating in real-time on-drone for high-resolution video input. Furthermore, they study the impact of hinge loss against the cross entropy loss on the classification performance, mainly in binary classification problems. Finally, they propose a novel regularization method motivated by the Quadratic Mutual Information, in order to improve the generalization ability of the utilized models. Extensive experiments on various binary classification problems involved in autonomous systems are performed, indicating the effectiveness of the proposed models. The experimental evaluation on four datasets indicates that hinge loss is the optimal choice for binary classification problems, considering lightweight deep models. Finally, the effectiveness of the proposed regularizer in enhancing the generalization ability of the proposed models is also validated.

- M. Krestenitis, N. Passalis, A. Iosifidis, M. Gabbouj and A. Tefas, "Recurrent bag-of-features for visual information analysis", Pattern Recognition, vol. 106, pp. 107380, doi: 10.1016/j.patcog.2020.107380, 2020.

Deep Learning (DL) has provided powerful tools for visual information analysis. For example, Convolutional Neural Networks (CNNs) are excelling in complex and challenging image analysis tasks by extracting meaningful feature vectors with high discriminative power. However, these powerful feature vectors are crushed through the pooling layers of the network, that usually implement the pooling operation in a less sophisticated manner. This can lead to significant information loss, especially in cases where the informative content of the data is sequentially distributed over the spatial or temporal dimension, e.g., videos, which often require extracting fine-grained temporal information. A novel stateful recurrent pooling approach, that can overcome the aforementioned limitations, is proposed in this paper. The proposed method is inspired by the well-known Bag-of-Features (BoF) model, but employs a stateful trainable recurrent quantizer, instead of plain static quantization, allowing for efficiently processing sequential data and encoding both their temporal, as well as their spatial aspects. The effectiveness of the proposed Recurrent BoF model to enclose spatio-temporal information compared to other competitive methods is demonstrated using six different datasets and two different tasks.

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

- N. Passalis, J. Raitoharju, A. Tefas and M. Gabbouj, “Efficient adaptive inference for deep convolutional neural networks using hierarchical early exits”, Pattern Recognition, vol. 105, pp. 107346, doi: 10.1016/j.patcog.2020.107346, 2020.

Early exits are capable of providing Deep Learning models with adaptive computational graphs that can readily adapt on-the-fly to the available resources. Despite their advantages, existing early exit methods suffer from many limitations which limit their performance, e.g., they ignore the information extracted from previous exit layers, they are unable to efficiently handle feature maps with large sizes, etc. To overcome these limitations the authors propose a Bag-of-Features (BoF)-based method that is capable of constructing efficient hierarchical early exit layers with minimal computational overhead, while also providing an adaptive inference method that allows for early stopping the inference process when the network is confident enough for its output, leading to significant performance benefits. To this end, the BoF model is extended and adapted to the needs of early exits by constructing additive shared histogram spaces that gradually refine the information extracted from the various layers of a network, in a hierarchical manner, while also employing a classification layer reuse strategy to further reduce the number of parameters needed per exit layer. Note that the proposed method is generic and can be readily combined with any neural network architecture. The effectiveness of the proposed method is demonstrated using five different image datasets, proving that early exits can be readily transformed into a practical tool, which can be effectively used in various real-world embedded applications.

1.2.5.3 Conference papers

16 conference papers that describe results obtained within the project were published or accepted for presentation during this period. The full text of these publications can be found in the relevant technical deliverables.

One paper was published in IEEE International Conference on Image Processing (ICIP) 2020.

- Dat Thanh Tran, Moncef Gabbouj and Alexandros Iosifidis, “Subset Sampling For Progressive Neural Network Learning”, IEEE International Conference on Image Processing, 2020

Progressive Neural Network Learning is a class of algorithms that incrementally construct the network's topology and optimize its parameters based on the training data. While this approach exempts the users from the manual task of designing and validating multiple network topologies, it often requires an enormous number of computations. In this paper, the authors propose to speed up this process by exploiting subsets of training data at each incremental training step. Three different sampling strategies for selecting the training samples according to different criteria are proposed and evaluated. They also propose to perform online hyperparameter selection during the network progression, which further reduces the overall training time. Experimental results in object, scene and face recognition problems demonstrate that the proposed approach speeds up the optimization procedure considerably while operating on par with the baseline approach exploiting the entire training set throughout the training process.

A joint AU-TAU paper was published in IEEE Symposium Series on Computational Intelligence (IEEE SSCI) 2020.

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

- Dat Thanh Tran, Moncef Gabbouj and Alexandros Iosifidis, "Performance Indicator in Multilinear Compressive Learning", IEEE Symposium Series on Computational Intelligence (IEEE SSCI), 2020

In Compressive Learning in general, and in Multilinear Compressive Learning (MCL) in particular, the number of compressed measurements captured by a compressive sensing device characterizes the storage requirement or the bandwidth requirement for transmission. This number, however, does not completely characterize the learning performance of a MCL system. In this paper, the authors analyse the relationship between the input signal resolution, the number of compressed measurements and the learning performance of MCL. Empirical analysis shows that the reconstruction error obtained at the initialization step of MCL strongly correlates with the learning performance, thus can act as a good indicator to efficiently characterize learning performances obtained from different sensor configurations without optimizing the entire system.

Five papers were accepted in International Conference on Pattern Recognition (ICPR) 2020 (to be held virtually in January 2021 due to Covid-19).

- Lukas Hedegaard Morsing, Omar Ali Sheikh-Omar and Alexandros Iosifidis, "Supervised Domain Adaptation using Graph Embedding", International Conference on Pattern Recognition (ICPR), 2020

Getting Deep Convolutional neural networks to perform well requires a large amount of training data. When the available labelled data is small, it is often beneficial to use transfer learning to leverage a related larger dataset (source) in order to improve the performance on the small dataset (target). Among the transfer learning approaches, domain adaptation methods assume that distributions between the two domains are shifted and attempt to realign them. In this paper, the authors consider the domain adaptation problem from the perspective of dimensionality reduction and propose a generic framework based on graph embedding. Instead of solving the generalised eigenvalue problem, they formulate the graph-preserving criterion as a loss in the neural network and learn a domain-invariant feature transformation in an end-to-end fashion. The proposed approach leads to a powerful Domain Adaptation framework; a simple LDA-inspired instantiation of the framework leads to state-of-the-art performance on two of the most widely used Domain Adaptation benchmarks, Office31 and MNIST to USPS datasets.

- Negar Heidari and Alexandros Iosifidis, "Temporal Attention-Augmented Graph Convolutional Network for Efficient Skeleton-Based Human Action Recognition", International Conference on Pattern Recognition (ICPR), 2020

Graph Convolutional Networks (GCNs) have been very successful in modelling non-Euclidean data structures, like sequences of body skeletons forming actions modelled as spatio-temporal graphs. Most GCN-based action recognition methods use deep feed-forward networks with high computational complexity to process all skeletons in an action. This leads to a high number of floating point operations (ranging from 16G to 100G FLOPs) to process a single sample, making their adoption in restricted computation application scenarios infeasible.

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

In this paper, the authors propose a Temporal Attention Module (TAM) for increasing the efficiency in skeleton-based action recognition by selecting the most informative skeletons of an action at the early layers of the network. They incorporate the TAM in a light-weight GCN topology to further reduce the overall number of computations. Experimental results on two benchmark datasets show that the proposed method outperforms with a large margin the baseline GCN-based method while having x2.9 less number of computations. Moreover, it performs on par with the state-of-the-art with up to x9.6 less number of computations.

- T. Bozinis, N. Passalis and A. Tefas, "Improving Visual Question Answering using Active Perception on Static Images", 25th International Conference on Pattern Recognition (ICPR), 2020

Visual Question Answering (VQA) is one of the most challenging emerging applications of Deep Learning. Providing powerful attention mechanisms is crucial for VQA, since the model must correctly identify the region of an image that is relevant to the question at hand. However, existing models analyse the input images at a fixed and typically small resolution, often leading to discarding valuable fine-grained details. To overcome this limitation, in this work the authors propose a reinforcement learning-based active perception approach that works by applying a series of transformation operations on the images (translation, zoom) in order to facilitate answering the question at hand. This allows for performing fine-grained analysis, effectively increasing the resolution at which the models process information. The proposed method is orthogonal to existing attention mechanisms and it can be combined with most existing VQA methods. The effectiveness of the proposed method is experimentally demonstrated on a challenging VQA dataset.

- Marios Krestenitis, Nikolaos Passalis, Alexandros Iosifidis, Moncef Gabbouj and Anastasios Tefas, "Human Action Recognition using Recurrent Bag-of-Features Pooling", 25th International Conference on Pattern Recognition (ICPR), 2020

Bag-of-Features (BoF)-based models have been traditionally used for various computer vision tasks, due to their ability to provide compact semantic representations of complex objects, e.g., images, videos, etc. Indeed, BoF has been successfully combined with various feature extractions methods, ranging from handcrafted feature extractors to powerful Deep Learning models. However, BoF, along with most of the pooling approaches employed in Deep Learning, fails to capture the temporal dynamics of the input sequences. This leads to significant information loss, especially when the informative content of the data is sequentially distributed over the temporal dimension, e.g., videos. In this paper the authors propose a novel stateful recurrent quantization and aggregation approach in order to overcome the aforementioned limitation. The proposed method is inspired by the well-known Bag-of-Features (BoF) model, but employs a stateful trainable recurrent quantizer, instead of plain static quantization, allowing for effectively encoding the temporal dimension of the data. The effectiveness of the proposed approach is demonstrated using three video action recognition datasets.

- N. Passalis and Anastasios Tefas, "Leveraging Quadratic Spherical Mutual Information Hashing for Fast Image Retrieval", 25th International Conference on Pattern Recognition (ICPR), 2020

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

Several deep supervised hashing techniques have been proposed to allow for querying large image databases. However, it is often overlooked that the process of information retrieval can be modelled using information-theoretic metrics, leading to optimizing various proxies for the problem at hand instead. Contrary to this, the authors propose a deep supervised hashing algorithm that optimizes the learned codes using an information-theoretic measure, the Quadratic Mutual Information (QMI). The proposed method is adapted to the needs of large-scale hashing and information retrieval leading to a novel information-theoretic measure, the Quadratic Spherical Mutual Information (QSMI), that is inspired by QMI, but leads to significant better retrieval precision. Indeed, the effectiveness of the proposed method is demonstrated under several different scenarios, using different datasets and network architectures, outperforming existing deep supervised image hashing techniques.

One paper was also published in the prestigious IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2020. Another paper was also included in CVPR's Workshop on Scalability in Autonomous Driving.

- N. Passalis, M. Tzelepi, and A. Tefas, "Heterogeneous Knowledge Distillation using Information Flow Modeling", IEEE Conference on Computer Vision and Pattern Recognition, (CVPR) 2020

Knowledge Distillation (KD) methods are capable of transferring the knowledge encoded in a large and complex teacher into a smaller and faster student. Early methods were usually limited to transferring the knowledge only between the last layers of the networks, while latter approaches were capable of performing multi-layer KD, further increasing the accuracy of the student. However, despite their improved performance, these methods still suffer from several limitations that restrict both their efficiency and flexibility. First, existing KD methods typically ignore that neural networks undergo through different learning phases during the training process, which often requires different types of supervision for each one. Furthermore, existing multi-layer KD methods are usually unable to effectively handle networks with significantly different architectures (heterogeneous KD). In this paper the authors propose a novel KD method that works by modelling the information flow through the various layers of the teacher model and then train a student model to mimic this information flow. The proposed method is capable of overcoming the aforementioned limitations by using an appropriate supervision scheme during the different phases of the training process, as well as by designing and training an appropriate auxiliary teacher model that acts as a proxy model capable of "explaining" the way the teacher works to the student. The effectiveness of the proposed method is demonstrated using four image datasets and several different evaluation setups.

- Juana Valeria Hurtado, Rohit Mohan, Wolfram Burgard, Abhinav Valada. "MOPT: Multi-Object Panoptic Tracking". The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshop on Scalability in Autonomous Driving, 2020

Comprehensive understanding of dynamic scenes is a critical prerequisite for intelligent robots to autonomously operate in their environment. Research in this domain, which

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

encompasses diverse perception problems, has primarily been focused on addressing specific tasks individually rather than modelling the ability to understand dynamic scenes holistically. In this paper, the authors introduce a novel perception task denoted as multi-object panoptic tracking (MOPT), which unifies the conventionally disjoint tasks of semantic segmentation, instance segmentation, and multi-object tracking. MOPT allows for exploiting pixel-level semantic information of 'thing' and 'stuff' classes, temporal coherence, and pixel-level associations over time, for the mutual benefit of each of the individual sub-problems. To facilitate quantitative evaluations of MOPT in a unified manner, the authors propose the soft panoptic tracking quality (sPTQ) metric. As a first step towards addressing this task, they propose the novel PanopticTrackNet architecture that builds upon the state-of-the-art top-down panoptic segmentation network EfficientPS by adding a new tracking head to simultaneously learn all sub-tasks in an end-to-end manner. Additionally, several strong baselines are presented that combine predictions from state-of-the-art panoptic segmentation and multi-object tracking models for comparison. Finally, the authors present extensive quantitative and qualitative evaluations of both vision-based and LiDAR-based MOPT that demonstrate encouraging results.

Members of the consortium also published papers in the International Conference on Artificial Intelligence Applications and Innovations (AIIA) 2020 the flagship IEEE International Conference on Multimedia and Expo (ICME), 2020, as well as in European Signal Processing Conference (EUSIPCO) 2020 (to be held virtually in January 2021 due to Covid-19).

- M. Kirtas, K. Tsampazis, N. Passalis, and A. Tefas, “Deepbots: A Webots-Based Deep Reinforcement Learning Framework for Robotics”, 16th International Conference on Artificial Intelligence Applications and Innovations (AIAI), 2020

Deep Reinforcement Learning (DRL) is increasingly used to train robots to perform complex and delicate tasks, while the development of realistic simulators contributes to the acceleration of research on DRL for robotics. However, it is still not straightforward to employ such simulators in the typical DRL pipeline, since their steep learning curve and the enormous amount of development required to interface with DRL methods significantly restrict their use by researchers. To overcome these limitations, in this work the authors present an open-source framework that combines an established interface used by DRL researchers, the OpenAI Gym interface, with the state-of-the-art Webots robot simulator in order to provide a standardized way to employ DRL in various robotics scenarios. Deepbots aims to enable researchers to easily develop DRL methods in Webots by handling all the low-level details and reducing the required development effort. The effectiveness of the proposed framework is demonstrated through code examples, as well as using three use cases of varying difficulty.

- A. Tzimas, N. Passalis, and A. Tefas, “Leveraging Deep Reinforcement Learning For Active Shooting Under Open-World Setting”, IEEE International Conference on Multimedia and Expo (ICME), 2020

Recent advances in Deep Reinforcement Learning (DRL) led to the development of powerful agents that can learn how to perform complicated tasks in an end-to-end fashion operating directly on raw unstructured data, e.g., images. However, the real world performance of such methods critically relies on the quality of the simulation environments used for training

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

them. The main contribution of this paper is the development of a realistic simulation environment, by employing a state-of-the-art graphics engine, for training DRL agents that are able to control a drone for performing active shooting. In contrast with previous approaches, that solely relied on simplistic constrained datasets, the environment employed in this work supports a challenging open-world setting, providing a solid step towards developing effective RL methods for various drone control tasks. An appropriate reward shaping approach is also introduced in this work, ensuring that the agent will behave as expected, avoiding erratic movements, as demonstrated through the conducted experiments.

- N. Passalis, A. Iosifidis, M. Gabbouj, and A. Tefas, "Robust Hypersphere-based Weight Imprinting for Few-shot Learning", European Signal Processing Conference (EUSIPCO), 2020

Performing fast few-shot learning is increasingly important in a number of embedded applications. Among them, a form of gradient-descent free learning known as Weight Imprinting was recently established as an efficient way to perform few-shot learning on Deep Learning (DL) accelerators that do no support back-propagation, such as Edge Tensor Processing Units (Edge TPUs). Despite its efficiency, WI comes with a number of critical limitations. For example, WI cannot effectively handle multimodal novel categories, while it is especially prone to overfitting that can have devastating effects on the accuracy of the models on novel categorizes. To overcome these limitations, in this paper the authors propose a robust hypersphere-based WI approach that allows for regularizing the training process in an imprinting-aware way. At the same time, the proposed formulation provides a natural way to handle multimodal novel categories. Indeed, as demonstrated through the conducted experiments, the proposed method leads to significant improvements over the baseline WI approach.

Yet another paper was published in IEEE International Workshop on Machine Learning for Signal Processing (MLSP) 2020.

- M. Tzelepi and A.Tefas, "Quadratic Mutual Information Regularization in real-time deep CNN models", IEEE International Workshop on Machine Learning for Signal Processing (MLSP), 2020

In this paper, regularized lightweight deep convolutional neural network models, capable of effectively operating in real-time on devices with restricted computational power for high-resolution video input are proposed. Furthermore, a novel regularization method motivated by the Quadratic Mutual Information, in order to improve the generalization ability of the utilized models is proposed. Extensive experiments on various binary classification problems involved in autonomous systems are performed, indicating the effectiveness of the proposed models as well as of the proposed regularizer.

Two papers were published in IEEE International Workshop on Multimedia Signal Processing (MMSP) 2020.

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

- N.Passalis, J. Raitoharju, M. Gabbuj, A. Tefas, "Efficient Adaptive Inference leveraging Bag-of-Features-based Early Exits", IEEE International Workshop on Multimedia Signal Processing, 2020

Early exits provide an effective way of implementing adaptive computational graphs over Deep Learning models. In this way it is possible to adapt them on-the-fly to the available computational resources or even to the difficulty of each input sample, reducing the energy and computational power requirements in many embedded and mobile applications. However, performing this kind of adaptive inference also comes with several challenges, since the difficulty of each sample must be estimated, and the most appropriate early exit must be selected. It is worth noting that existing approaches often lead to highly unbalanced distributions over the selected early exits, reducing the efficiency of the adaptive inference process. At the same time, only a few resources can be devoted to the aforementioned process, in order to ensure that an adequate speedup will be obtained. The main contribution of this work is to provide an easy to use and tune adaptive inference approach for early exits that can overcome some of these limitations. In this way, the proposed method allows for a) obtaining a more balanced inference distribution among the early exits, b) relying on a single and interpretable hyper-parameter for tuning its behaviour (ranging from faster inference to higher accuracy), and c) improving the performance of the networks (increasing the accuracy and reducing the time needed for inference). Indeed, the effectiveness of the proposed method over existing approaches is demonstrated using four different image datasets.

- N. Passalis, and A. Tefas, "Leveraging Active Perception for Improving Embedding-based Deep Face Recognition", IEEE International Workshop on Multimedia Signal Processing, 2020

Even though recent advances in Deep Learning (DL) led to tremendous improvements for various computer and robotic vision tasks, existing DL approaches suffer from a significant limitation: they typically ignore that robots and cyber-physical systems are capable of interacting with the environment in order to better sense their surroundings. In this work the authors argue that perceiving the world through physical interaction, i.e., employing active perception, allows for both increasing the accuracy of DL models, as well as for deploying smaller and faster models. To this end, an active perception-based face recognition approach is proposed, which is capable of simultaneously extracting discriminative embeddings, as well as predicting in which direction the robot must move in order to get a more discriminative view. Based on the authors, this is the first embedding-based active perception method for deep face

- Daniele Cattaneo, Domenico Giorgio Sorrenti, Abhinav Valada. "CMRNet++: Map and Camera Agnostic Monocular Visual Localization in LiDAR Maps". The IEEE International Conference on Robotics and Automation (ICRA) Workshop on Emerging Learning and Algorithmic Methods for Data Association in Robotics, 2020. (Spotlight)

Localization is a critically essential and crucial enabler of autonomous robots. While Deep Learning has made significant strides in many computer vision tasks, it is still yet to make a sizeable impact on improving capabilities of metric visual localization. One of the major hindrances has been the inability of existing Convolutional Neural Network (CNN)-based pose

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

regression methods to generalize to previously unseen places. The recently introduced CMRNet effectively addresses this limitation by enabling map independent monocular localization in LiDAR-maps. In this paper, the authors now take it a step further by introducing CMRNet++, which is a significantly more robust model that not only generalizes to new places effectively, but is also independent of the camera parameters. They enable this capability by combining Deep Learning with geometric techniques, and by moving the metric reasoning outside the learning process. In this way, the weights of the network are not tied to a specific camera. Extensive evaluations of CMRNet++ on three challenging autonomous driving datasets, i.e., KITTI, Argoverse, and Lyft5, show that CMRNet++ outperforms CMRNet as well as other baselines by a large margin. More importantly, for the first-time, the authors demonstrate the ability of a Deep Learning approach to accurately localize without any retraining or fine-tuning in a completely new environment and independent of the camera parameters.

Until now, 10 of the above papers have already been presented in eight conferences and workshops (ICIP, SSCI, CVPR, AIAI, ICME, MLSP, MMSP, ICRA) while another 6 papers are to be presented in ICPR 2020 and EUSIPCO 2020 conferences. All the conferences/workshops until now were held virtually due to the coronavirus pandemic situation and the relevant presentations of the scientific papers were done with the use of video recordings. Some frames of the presentations that have been done so far are depicted in Fig. 29.

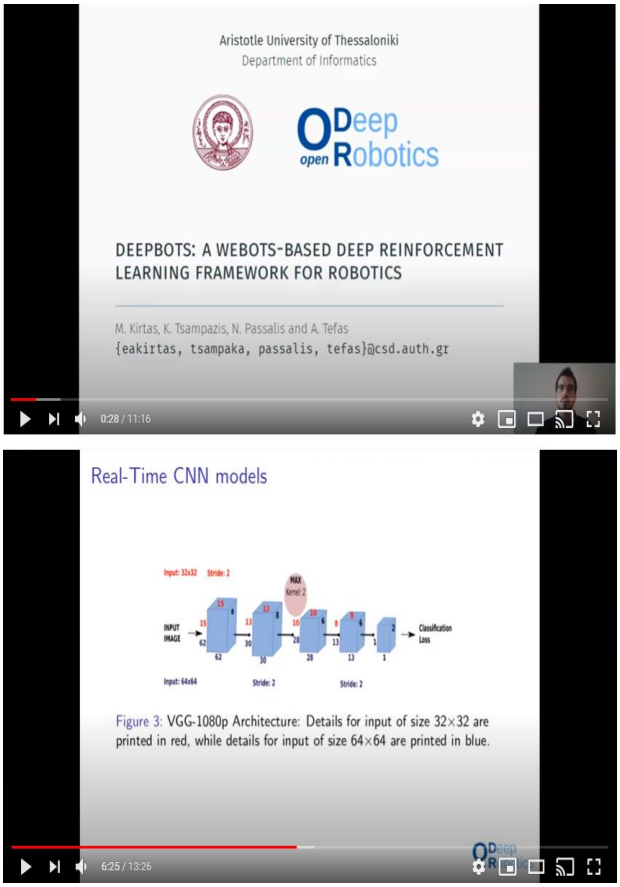

Figure 29: Frames from video presentations in conferences

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

1.2.5.4 Preprints

In addition to the publications in peer-reviewed journals and conferences mentioned in the previous subsections, a number of papers have been posted in the open access archive arXiv. These papers will soon be submitted (or are already submitted and under review) to appropriate forums. The full text of these papers can be found in the relevant technical deliverables.

- Dat Thanh Tran, Nikolaos Passalis, Anastasios Tefas, Moncef Gabbouj and Alexandros Iosifidis, "Attention-based Neural Bag-of-Features Learning for Sequence Data", arXiv preprint, arXiv:2005.12250, 2020

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. With extensive experiments on sequential data analysis problems, including audio data analysis and medical data analysis problems, it is empirically shown that the proposed 2DA-NBoFs model performs favourably compared to existing methods, including the widely used Gated Recurrent Units models. 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.

- Dat Thanh Tran, Moncef Gabbouj and Alexandros Iosifidis, "Multilinear Compressive Learning with Prior Knowledge", arXiv preprint, arXiv:2002.07203

The recently proposed Multilinear Compressive Learning (MCL) framework combines Multilinear Compressive Sensing and Machine Learning into an end-to-end system that takes into account the multidimensional structure of the signals when designing the sensing and feature synthesis components. The key idea behind MCL is the assumption of the existence of a tensor subspace which can capture the essential features from the signal for the downstream learning task. Thus, the ability to find such a discriminative tensor subspace and optimize the system to project the signals onto that data manifold plays an important role in Multilinear Compressive Learning. In this paper, the authors propose a novel solution to address both of the aforementioned requirements, i.e., How to find those tensor subspaces in which the signals of interest are highly separable? and How to optimize the sensing and feature synthesis components to transform the original signals to the data manifold found in the first question? In the proposed approach, the discovery of a high-quality data manifold is conducted by training a nonlinear compressive learning system on the inference task. Its knowledge of the data manifold of interest is then progressively transferred to the MCL components via multi-stage supervised training with the supervisory information encoding how the compressed measurements, the synthesized features, and the predictions should be like. The proposed knowledge transfer algorithm also comes with a semi-supervised adaptation that enables compressive learning models to utilize unlabelled data effectively. Extensive experiments

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

demonstrate that the proposed knowledge transfer method can effectively train MCL models to compressively sense and synthesize better features for the learning tasks with improved performances, especially when the complexity of the learning task increases.

- Lukas Hedegaard Morsing, Omar Ali Sheikh-Omar and Alexandros Iosifidis, "Supervised Domain Adaptation: A Graph Embedding Perspective and a Rectified Experimental Protocol", arXiv preprint, arXiv:2004.11262

The performance of machine learning models tends to suffer when the distributions of the training and test data differ. Domain Adaptation is the process of closing the distribution gap between datasets. In this paper, it is shown 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. In this paper, the authors analyse the loss functions of existing state-of-the-art Supervised Domain Adaptation methods and demonstrate that they perform Graph Embedding. Moreover, they highlight some generalisation and reproducibility issues related to the experimental setup commonly used to demonstrate the few-shot learning capabilities of these methods. A rectified evaluation setup is proposed for more accurately assessing and comparing Supervised Domain Adaptation methods, and experiments on the standard benchmark datasets Office31 and MNIST-USPS are reported.

- Negar Heidari and Alexandros Iosifidis, "Progressive Spatio-Temporal Graph Convolutional Network for Skeleton-based Human Action Recognition", arXiv preprint, arXiv:2011.05668

Graph convolutional networks (GCNs) have been very successful in skeleton-based human action recognition where the sequence of skeletons is modelled as a graph. However, most of the GCN-based methods in this area train a deep feed-forward network with a fixed topology that leads to high computational complexity and restricts their application in low computation scenarios. In this paper, the authors propose a method to automatically find a compact and problem-specific topology for spatio-temporal graph convolutional networks in a progressive manner. Experimental results on two widely used datasets for skeleton-based human action recognition indicate that the proposed method has competitive or even better classification performance compared to the state-of-the-art methods with much lower computational complexity.

- Negar Heidari and Alexandros Iosifidis, "On the spatial attention in Spatio-Temporal Graph Convolutional Networks for Skeleton-based Human Action Recognition", arXiv preprint, arXiv:2011.05668

Graph convolutional networks (GCNs) achieved promising performance in skeleton-based human action recognition by modelling a sequence of skeletons as a spatio-temporal graph. Most of the recently proposed GCN-based methods improve the performance by learning the graph structure at each layer of the network using a spatial attention applied on a predefined graph Adjacency matrix that is optimized jointly with model's parameters in an end-

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

to-end manner. In this paper, the authors analyse the spatial attention used in spatio-temporal GCN layers and propose a symmetric spatial attention for better reflecting the symmetric property of the relative positions of the human body joints when executing actions. They also highlight the connection of spatio-temporal GCN layers employing additive spatial attention to bilinear layers, and they propose the spatio-temporal bilinear network (ST-BLN) which does not require the use of predefined Adjacency matrices and allows for more flexible design of the model. Experimental results show that the three models lead to effectively the same performance. Moreover, by exploiting the flexibility provided by the proposed ST-BLN, one can increase the efficiency of the model.

- Anton Muravev, Jenni Raitoharju and Moncef Gabbouj, "Neural Architecture Search by Estimation of Network Structure Distributions", arXiv preprint arXiv:1908.06886

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 an approach allows for faster search, it does so at the cost of expressivity. The authors instead propose an alternative probabilistic representation of a whole neural network structure under the assumption of independence between layer types. The matrix of probabilities is equivalent to the population of models, but allows for discovery of structural irregularities, while being simple to interpret and analyse. 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 proposed 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.

- Mohammad Soltanian, Junaid Malik, Jenni Raitoharju, Alexandros Iosifidis, Serkan Kiranyaz, Moncef Gabbouj. "Speech Command Recognition in Computationally Constrained Environments with a Quadratic Self-organized Operational Layer", arXiv preprint arXiv:2011.11436 (under review in IEEE ICASSP 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 light-weight 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

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

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

- Kulhánek, J., Derner, E., & Babuška, R. (2020). Visual Navigation in Real-World Indoor Environments Using End-to-End Deep Reinforcement Learning. *arXiv preprint arXiv:2010.10903*.

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 image segmentation masks. This precludes the use of the learned policy on a real robot. Therefore, the authors propose a novel approach that enables a direct deployment of the trained policy on real robots. They have designed visual auxiliary tasks, a tailored reward scheme, and a new powerful simulator to facilitate domain randomization. The policy is fine-tuned on images collected from real-world environments. The method has been evaluated on a mobile robot in a real office environment. The training took ~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.

- Rohit Mohan and Abhinav Valada. “EfficientPS: Efficient Panoptic Segmentation”, arXiv preprint arXiv:2004.02307, 2020.

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. The authors 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 IndianDriving 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.

1.2.5.5 Other Publications

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

During this period, a number of press items regarding the project or its results were authored/posted by either consortium members or parties external to the project. Indeed AU OpenDR team members posted an article in the website of the Department of Engineering, Aarhus University, entitled "New robotic AI will improve sustainability in agriculture" (27.03.2020)

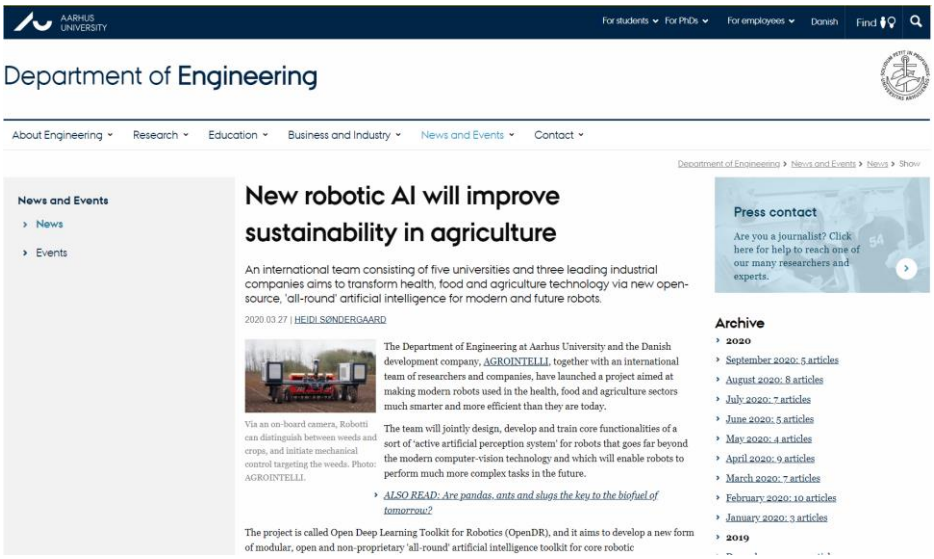

Figure 30: OpenDR post in the Department of Engineering, Aarhus University, website

Moreover, due to its success in the benchmarks and the Robust Vision challenge 2020 at ECCV2020, the EfficientPS network from ALU-FR was covered in numerous media articles including *Tech Xplore* (Fig. 31) and *Springer Professional* (Fig. 32).

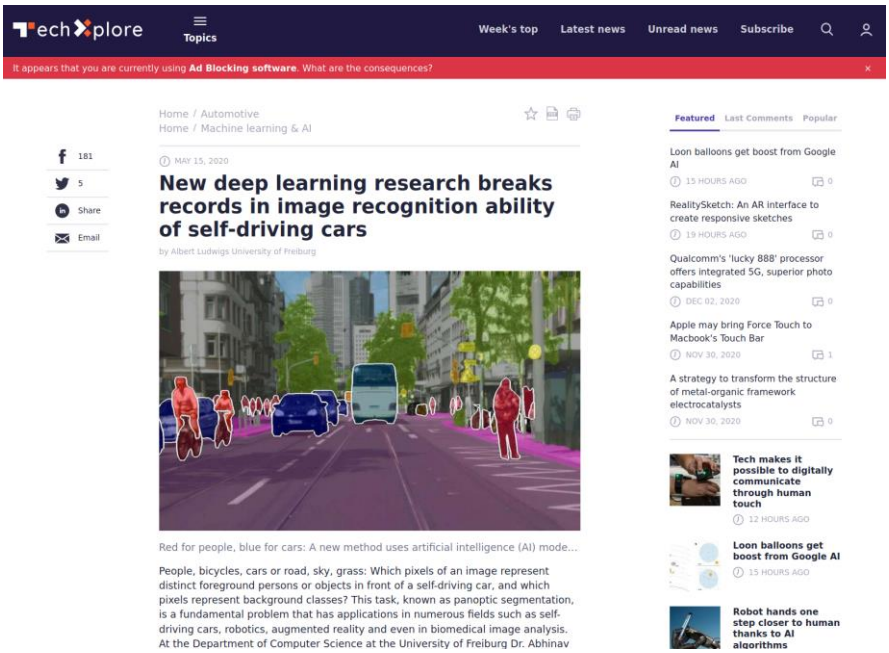

Figure 31: Tech Xplore article for EfficientPS network from ALU-FR

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

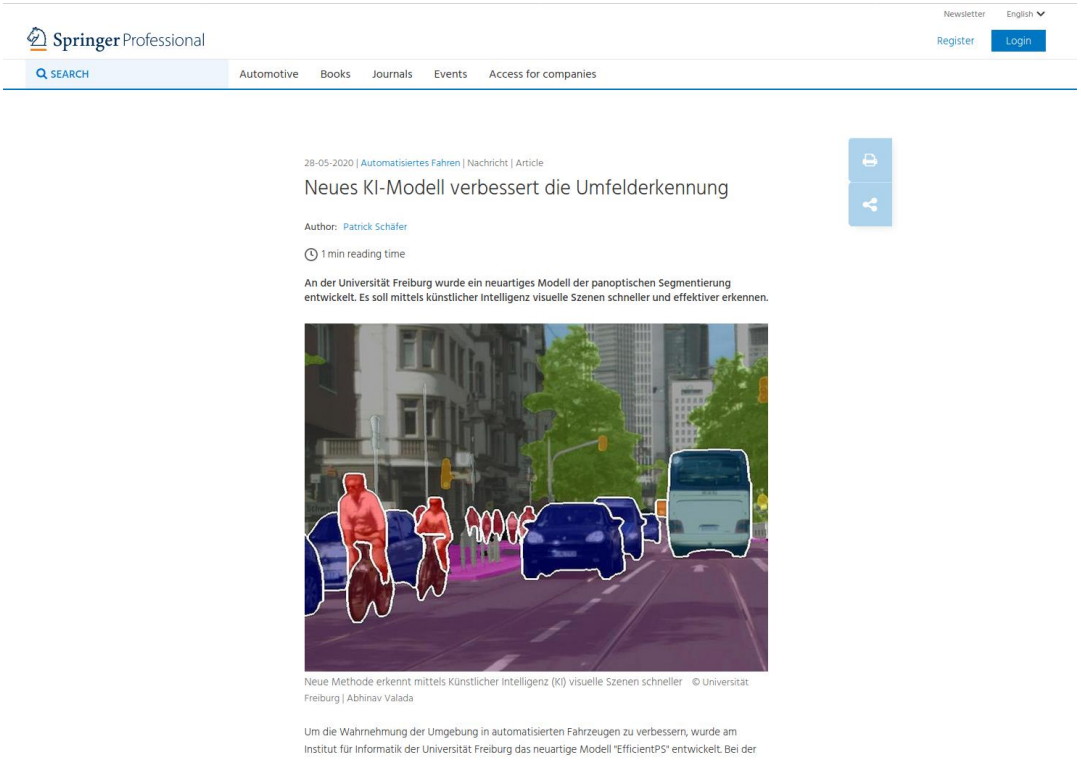

Figure 32: Springer Professional article for EfficientPS network from ALU-FR

Finally, following the distribution of the first press release through the Press Office of Aristotle University of Thessaloniki to its communication channels, a number of articles appeared in more than 15 Greek media websites such as thestival.gr, inewsgr.com, makthes.gr, etc. (Fig. 33).

Figure 33: News item regarding the project in a Greek media website.

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

1.2.6 Workshops, conferences, special sessions

During the first year of the project, consortium members took important initiatives related to the organization of a workshop, scientific/organizational support of a major conference and organization of two conference special sessions on topics related to the project.

On January 13th 2020, along with OpenDR's first project meeting, an open training workshop entitled "Deep Learning for Robotics" was organized by Prof. Anastasios Tefas and held in the premises of AUTH in Thessaloniki, Greece. In this workshop the OpenDR research project was introduced to the audience for the first time with a speech by Prof. Tefas on the main aspects and goals of the project. Furthermore, several lectures were given by the members of the consortium regarding the progress and state of the art in areas related to Deep Learning and robotics, as well as their applications, mainly in agriculture, healthcare and agile production (Fig. 34, Fig. 35). In more detail the workshop included lectures in the following topics:

- OpenDR: Open Deep Learning for Robotics Toolkit
- Deep Learning for Robot Perception and Cognition
- Deep Learning for BioSignal Analysis
- Deep Learning for Robot Navigation
- Deep Learning for Robot Planning, Action and Decision Making
- Webots: Open-Source Robots Simulation
- Robotics in Agriculture
- Robotics in Healthcare
- Cognitive human-robot interaction in Agile Production

The workshop was attended by approximately 130 participants, making it a big success.

Figure 34: The first workshop of the OpenDR project organized by AUTH on January 2020

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

Figure 35: The poster of the OpenDR first workshop

Moreover Moncef Gabbouj (TAU) served as the General Chair of the IEEE Signal Processing Flagship Conference, the International Conference on Image Processing, ICIP 2020, United Arab Emirates (fully virtual due to COVID19), 25-28 October, 2020 (Fig. 36, Fig. 37). The conference attracted over 7500 attendees from 105 countries, 3 keynotes, 14 tutorials (one of which was organized by members of OpenDR), 13 special sessions (one was organized by OpenDR), 55 regular sessions and 4 industry workshops. 698 papers were presented at ICIP, including several from OpenDR researchers. Tampere University (TAU), an OpenDR partner, was one of the University sponsors of ICIP 2020.

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

Figure 36: Moncef Gabbouj speaking at the Opening Ceremony of IEEE International Conference on Image Processing, ICIP 2020.

Figure 37: Moncef Gabbouj (TAU) served as the General Chair of IEEE ICIP 2020.

In addition, a Special Session on "Deep Learning for Robotic Perception and Cognition" was co-organized by Alexandros Iosifidis (AU) and Anastasios Tefas (AUTH) in the same conference (IEEE ICIP) (Fig. 38) and was supported by OpenDR. The purpose of the Special Session was to provide a forum to exchange ideas and to discuss developments in Deep Learning models for Robotic Perception and Navigation, address the challenges of bringing Deep Learning into robotic platforms, including low power Unmanned Aerial Vehicles (drones), Collaborative Robots, and Autonomous Vehicles. A hybrid Call for Papers process was followed as, according to the Conference rules, a list of (six) invited papers was included in the Session proposal. After the acceptance of the Special Session, an open call was made and the Session was advertised in the conference website and the organizers professional networks. The peer review process was handled centrally by the conference organization committees. Seven papers were presented in the Session, proposing methodologies for 6-DOF pose estimation, Deep Learning-based segmentation for autonomous marine operations, human action recognition in drone videos, novelty detection, generalized critic policy optimization for

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

deep reinforcement learning, Lidar-based camera calibration, and path planning for obstacle detection. Information on the Session and links to the papers can be found here.

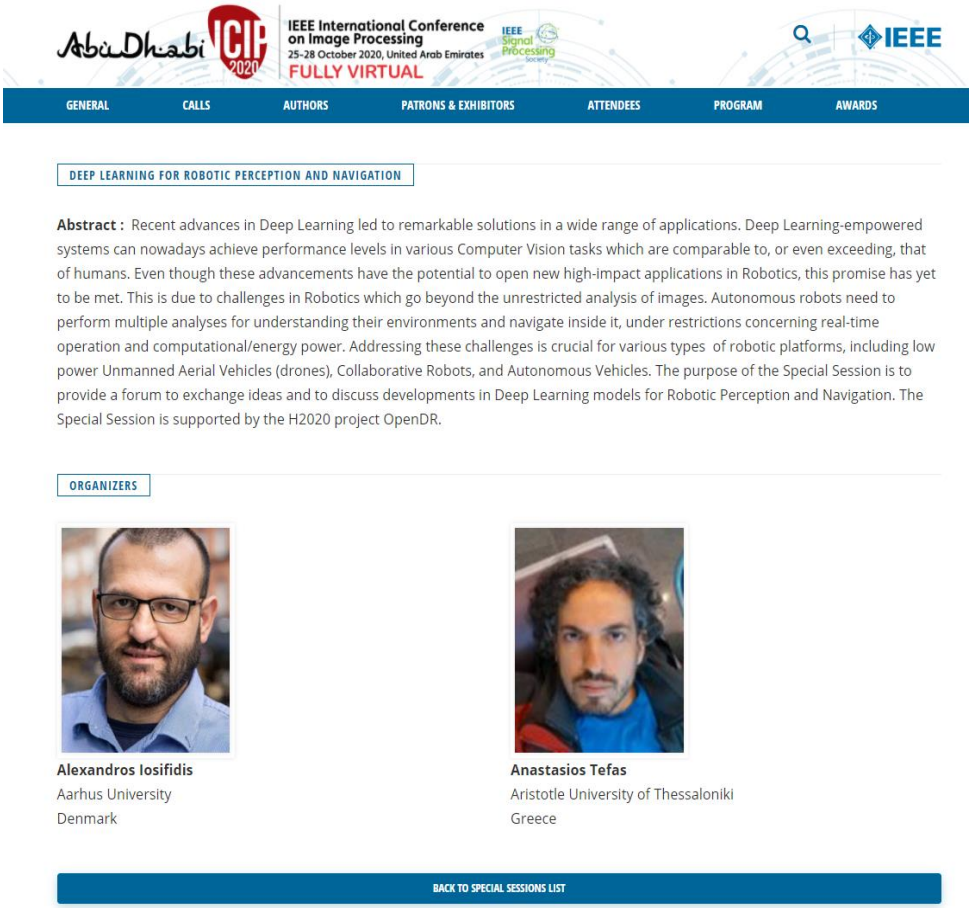

Figure 38: A Special Session on "Deep Learning for Robotic Perception and Cognition" was co-organized by Alexandros Iosifidis (AU) and Anastasios Tefas (AUTH) in IEEE ICIP

Finally, a Special Session on "Artificial Neural Networks for Computer and Robot Vision" will be organized by Alexandros Iosifidis (AU) in the International Conference on Emerging Techniques in Computational Intelligence (ICETCI), Hyderabad, India, 25-27 August, 2021. The session is supported by the IEEE Computational Intelligence Society Neural Networks Technical Committee and OpenDR. The purpose of the Special Session is to provide a forum to exchange ideas and to discuss developments in Artificial Neural Networks (Convolutional Neural Networks, Recurrent Neural Networks and Graph Neural Networks) with applications in Computer and Robot Vision as well as current issues such as those related to the adoption of high-performing difficult to train ANN-based solutions in real-life Computer and Robotic Vision problems.

1.2.7 Invited / keynote talks and tutorials

Dissemination of OpenDR information and findings through invited/keynote talks and tutorials has obviously high impact, since, usually, the audience consists of researchers or

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

industrial representatives, interested in the specific topic of the talk/tutorial. The following such activities took place during the first year of the project:

- A tutorial (Fig. 39, Fig. 40) on "Generalized Operational Neural Networks" was co-organized by Dat Thanh Tran (AU), Alexandros Iosifidis (AU), Junaid Malik (TAU), Serkan Kiranyaz and Moncef Gabbouj (TAU) in IEEE International Conference on Image Processing (IEEE ICIP), United Arab Emirates (fully virtual due to COVID19), 25-28 October, 2020. The tutorial presented several advanced ML techniques, many of which have been developed by OpenDR. OpenDR was showcased as a testbed for many topics presented at the tutorial. The project was acknowledged in the tutorial.

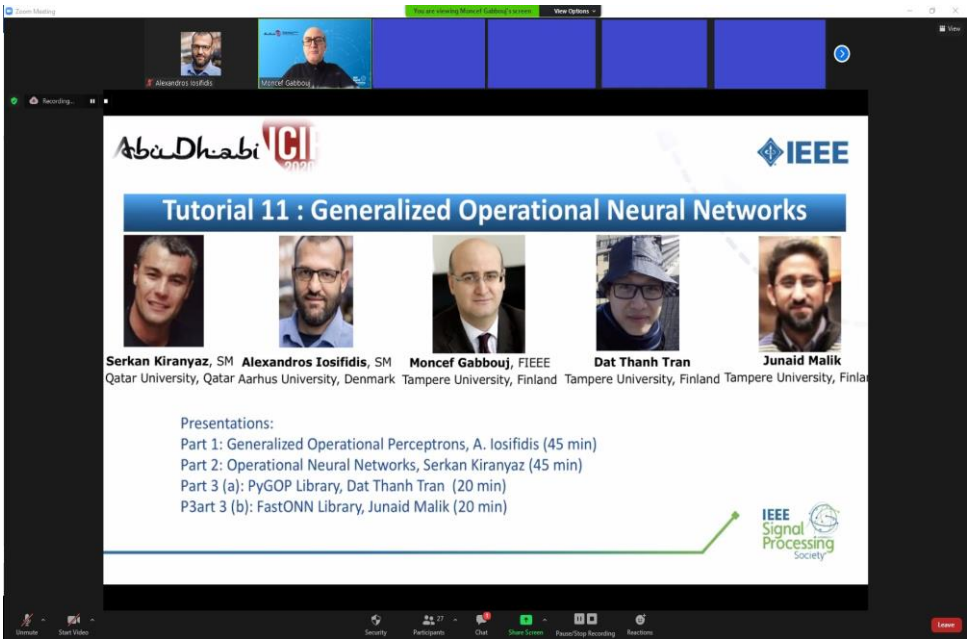

Figure 39: Presenters of the tutorial on "Generalized Operational Neural Networks", IEEE ICIP 2020

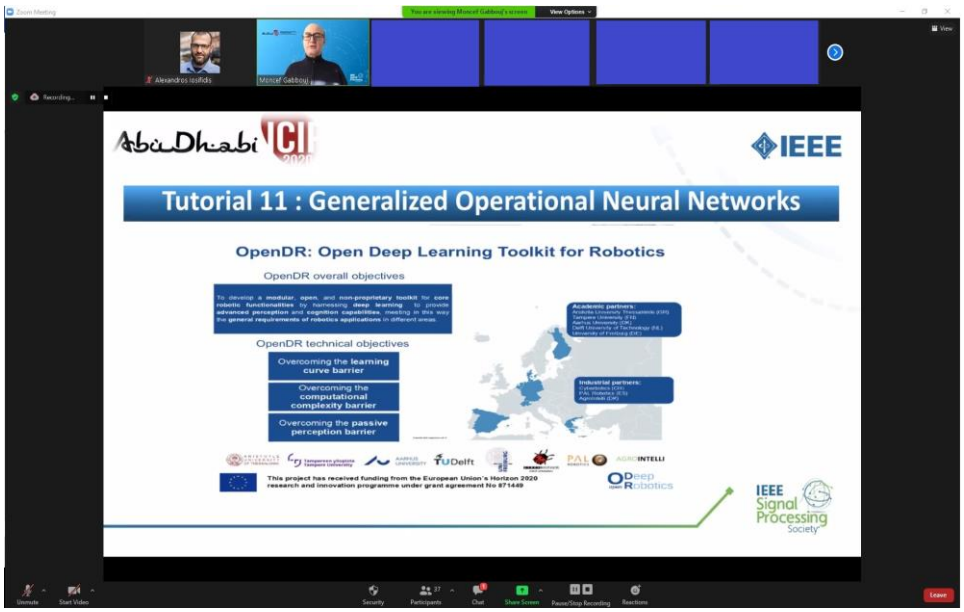

Figure 40: Tutorial on "Generalized Operational Neural Networks" co-organized by D.

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

Thanh Tran (AU), A. Iosifidis (AU), J. Malik (TAU), S. Kiranyaz and M. Gabbouj (TAU) in IEEE ICIP 2020

- Moncef Gabbouj (TAU) presented a keynote at the International Conference on Signal Processing and Information Security. The talk was entitled "The Super Neuron Model — Signal Processing Applications." The Conference was fully virtual and was held on 25-26 Nov 2020 (originally planned to be held in Dubai).
- OpenDR, its consortium and objectives were described in the invited talk of Alexandros Iosifidis (AU) (Fig. 41) at the Environmental Institute of Finland, Helsinki entitled "Data-Driven Analytics: from Shallow to Deep, from manual to automatic" (30.01.2020). The talk was open to all employees of the Environmental Institute of Finland and had approximately 15 attendees.

Figure 41: Invited talk of Alexandros Iosifidis at the Environmental Institute of Finland, Helsinki

1.2.8 Participation to tradeshows, exhibitions, conferences, EU events, industry workshops

Events such as tradeshows, exhibitions, events organized by the EU etc., are obvious venues for the dissemination of OpenDR aims and results. Due to the pandemic, a number of such events were cancelled or were held in a virtual way. This had an effect to the project members participation in such activities. For example, AGI was planning to participate in the Agromek trade show in February 2021, however the event was cancelled. During this period, the project was present at the following fora:

- OpenDR, its consortium and objectives were described in the European Robotics Forum (ERF) Malaga (03-06.03.2020) both as a poster and as invited talks in workshops (Fig. 42, Fig. 43). The talks were delivered by Francesco Ferro (PAL) (title: "Breakthroughs in humanoid and service robots") and Lorna McKinlay (PAL) (title: "ARI: the new generation of AI powered robots").

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

- Prof A. Tefas (AUTH) also presented the project in two talks during ERF 2020: one in a session on “New H2020 projects”, organized by the EU and another one regarding “Social Robotics for HealthCare” in a session on "AI in Social Robotics" organized by PAL. Moreover he gave a short interview (Fig. 44) presenting the aim and the partners of OpenDR, as well as the objectives and the expected impact of the project.

Figure 42: Invited talk by Lorna McKinlay (PAL) at ERF 2020 workshop

Figure 43: A poster and invited talk in ERF 2020 workshops by Francesco Ferro (PAL)

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

Figure 44: Interview given by Professor Anastasios Tefas at ERF 2020

- OpenDR, its consortium and objectives were described in the AiRo Webinar, virtual event (02.06.2020), in a talk, entitled "Service Robotics to improve people's quality of life", delivered by Francesco Ferro (PAL) . The target group of the webinar was entrepreneurs, developers, researchers and authorities who are interested in networking in Finland and internationally.
- Similarly, OpenDR, was described in the Super Minds Virtual Conference, virtual event (02.07.2020) in a talk by Francesco Ferro titled "Humanization of the robots". The virtual congress Super Minds wanted to align to the new tendencies on the talent and innovation in the international workplace.
- The project and its objectives were also described in the AI & Big Data Conference, virtual event (14-15.10.2020). The talk was delivered by Francesco Ferro (PAL) (Fig. 45) and was entitled "Open Deep Learning toolkit for Robotics in the healthcare sector". AI & Big Data Congress is the leading congress on Artificial Intelligence & Big Data and the meeting point for professionals, suppliers and companies that want to develop or are carrying out projects in the field of AI & Data Analytics.

Figure 45: Francesco Ferro's (PAL) talk in AI & Big Data Conference

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

- OpenDR, its consortium and objectives were described in the International Conference on Intelligent Robots and Systems (IROS) workshop on "Robots Building Robots" virtual event (25-30.10.2020) The talk was delivered by Francesco Ferro (PAL) (Fig. 46) and was titled "Robots building Robots - when users build their own robots".

Figure 46: Talk by Francesco Ferro (PAL) in IROS2020

- The *deepbots framework* (Fig. 47) included in the OpenDR Toolkit to ease the development and deployment of Reinforcement Learning algorithms in the Webots robot simulator, was presented in FOSSCOMM 2020 (Free and Open Source Software Communities Meeting). This conference aims to bring together the opensource communities, researchers and development teams and promote the usefulness of free and opensource software.

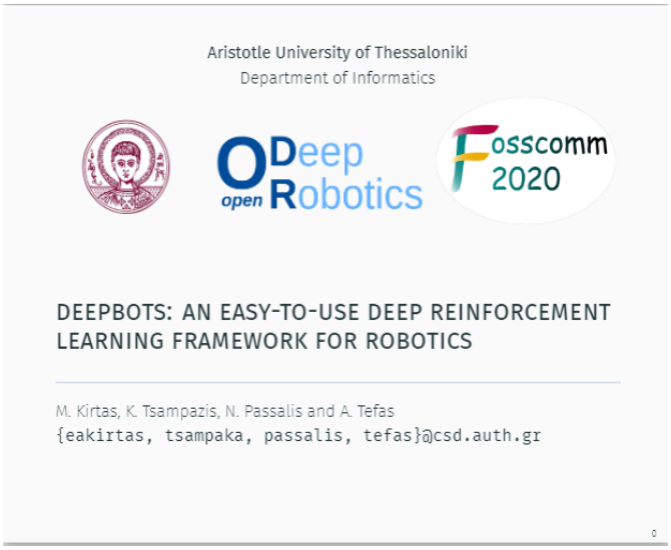

Figure 47: The deepbots framework presented in FOSSCOM 2020

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

1.2.9 Collaboration with other projects and DIHs

During this period there was an active engagement with the Digital Innovation Hub (DIH) TRINITY managed by TAU on the industrial relevance of the OpenDR toolkit in the field of Agile Production.

Indeed TRINITY held several workshops and meetings with SMEs, academics and other stakeholders in agile production and produced several reports detailing their results and outcomes, that also include the OpenDR agile production use case. In addition, TRINITY documented the needs and insights of SMEs, including the need of a toolkit that extends the capabilities of robots. The relevance of the OpenDR toolkit is therefore being addressed within the DIH TRINITY.

Furthermore, OpenDR partners have contacted a number of additional Digital Innovation Hubs, such as RODIN (RObotics Digital Innovation Network), AgROBOfood (Business-Oriented Support to the European Robotics and Agri-food Sector, towards a network of Digital Innovation Hubs in Robotics), and others, towards further extending and strengthening its links to DIHs.

1.2.10 Awards, participation into contests and challenges

Despite its early stage, project results already started to draw attention and distinctions. Indeed, the lightweight EfficientPS algorithm from ALU-FR won the Robust Vision Challenge 2020 Panoptic Segmentation Track, that was held in conjunction with ECCV2020.

Furthermore, Agro Intelligence (AGI) received the Solar Impulse award for its Robotti agricultural robot in October 2020. Although the award application process was not part of this project, the Robotti will be the testbed for the agri-food use case.

1.3 Dissemination and Communication Performance Evaluation

In order to assess the Project impact in the relevant communities it is important to measure the effects of the dissemination endeavours. The project has already defined its target audiences and started to monitor the degree of engagement and the effect of its strategies. Appropriate tools to collect and analyse user data (whenever possible and legally compliant) were used.

To measure the performance of the dissemination and communication activities, a number of performance indicators have been defined by the consortium in Part B of the DoA. Their aim is to help the project get a feeling of how well the dissemination is going and whether adjustments are needed. Target values were also set. In the following paragraphs we will take a look at the Key Performance Indicators (KPIs) and what already has been accomplished within the first year of the project. Obviously, the project will continue to monitor its dissemination performance through these KPIs throughout its duration.

Moreover, in the DoA the consortium has described a number of target values for the dissemination and communication materials that it will create and the events that it will organize during the project lifecycle. The progress towards reaching these goals is also reported in sections below.

1.3.1 Project website

KPIs that refer to the project website, along with their target values (all but one set for M36) and the figures obtained so far are presented in the table below. It is evident that a large number of internet users have visited the OpenDR website (more than 3200 until today) and

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

the respective target has been easily achieved. As a matter of fact, the website visitors count was three times the set target (1000 visitors by M12). However, the average duration of visit lags somewhat with respect to the target (M36) value, but with more interesting content becoming available as the project progresses, it is expected that the goal will be reached by the end of the project. Regarding the internal goals set in terms of website content (number of published posts and documents) the figures show that the number of posts is on track whereas the M36 target for documents (e.g., project publications) has almost been reached in the first year.

KPI	Target (for M36, unless stated otherwise)	Reached until M12
Number of visitors	1000 (by M12)	More than 3200
Average duration of visits:	2:00	1:36
Number of posts published (including news):	>50	15
Number of documents available in the website repository:	>20	18

1.3.2 Social media

OpenDR has been notably active in the most popular dissemination channels (Twitter, Facebook, LinkedIn) with its project accounts which were created right from the beginning. These communication efforts were received very well from the respective audiences, as presented below.

1.3.2.1 Twitter

As shown below, a large number of tweets/retweets have been posted until today, towards continuously feeding the public with information about the project's actions and achievements, as well as Deep Learning in general, so as to increase awareness for OpenDR's and improve its visibility**.** It is expected that the target value (150) will be easily surpassed by M36. Until now, the Twitter account has gathered more than 150 followers and the M36 target has been already surpassed, showing a strong interest for the project and its results.

KPI	Expected until M36	Reached until M12
Twitter followers	>100	151
Number of tweets/retweets	>150	89

1.3.2.2 LinkedIn

According to the DoA, the members of the LinkedIn group should be more than 150 at the end of the project. Until today, more than 110 users have become members of the OpenDR group and this fact shows that until M36 the goal will be achieved. Moreover, a LinkedIn page has been created for the project. Currently more than 185 people follow this page. Both figures reflect the fact that the group and the page are constantly fed with interesting content.

KPI	Expected until M36	Reached until M12
LinkedIn group members	>150	111

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

1.3.2.3 Facebook, YouTube

Although no KPIs and target values were set for Facebook in the DoA, the consortium has monitored its performance in this channel. Until today, more than 70 posts have been created, and the current number of account followers is 175 while the page has attracted so far more than 170 likes. These figures clearly show that the project is very active in communicating through this channel and this effort has paid off. Moreover, although not planned initially, the project has created a YouTube channel which has 27 subscribers and more than 800 views. Both figures are satisfactory, given that the channel was created in August 2020.

1.3.3 Publications

As described in the DoA, the project set a goal of publishing more than 8 scientific papers per year. Until today, OpenDR has already reached 24 publications (one of them being an edited book, a project currently in progress), a fact that shows that the project has achieved in the best way its annual goal, despite the fact that this was its first year.

KPI	Target	Reached until M12
Number of publications	>8 per year	24 (16 Conferences, 7 Journals, 1 Book)

1.3.4 Promotional material

In the framework of its dissemination activities, the consortium set certain M36 goals, concerning different types of promotional material and their usage. As shown in the table below, the project is perfectly on track with respect to the promotional materials that it has to produce, having created so far the first newsletter and press release, a promotional video and a brochure, as well as additional materials not mentioned in this table. The newsletter is to be disseminated through tweets and Facebook/LinkedIn posts as well as through emailing to relevant lists. Thus, despite the fact that no specific dissemination list might be created, the first newsletter is expected to reach more persons than the ones anticipated for the end of the project.

As far as the distribution of printed promotional material is concerned, the pandemic has had a major negative impact, as most of the events have been converted to virtual ones. Indeed, the project's factsheet and poster have been used only in one event (ERF 2020), while the promotional brochure has not been used yet in any event. Depending on the developments in the COVID-19 front, these targets might or may not be reached by the end of the project. However this is not to be considered as a major issue since, if it occurs, it will be compensated by the project's extremely good performance in other dissemination activities.

KPI	Expected until M36	Reached until M12
Number of e-newsletters published	3	1
Size of the dissemination list	> 500	>1000 (estimate)

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

Number of press releases	3	1
Number of videos to be produced	2	1
Number of video views	500	529
Number of events where the factsheet is used	10	1
Number of events where the poster is used	10	1
Number of brochures	2	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

1.3.5 Events

Despite the pandemic and the early stage in the project's lifecycle, the consortium had organized two training/ dissemination events. Indeed, as mentioned in the corresponding sections, the consortium has successfully organized one training workshop, in conjunction with its kick-off meeting, that attracted a large number of participants (pre- and post graduate students, PhD students, researchers) far above the set target value. Moreover, members of the consortium organized one tutorial in a major conference (ICIP 2020). Thus, the project, despite the hurdles, can be considered to be on track with respect to events organization. It is expected that, with the ease of restrictions regarding COVID-19 and as the project will generate more and more substantial results, more importantly, the first version of the toolkit, it will easily manage to achieve the set goals by the end of its lifecycle.

It shall be noted that the DoA included a number of KPI's for the project software repository. However, as this will be available in M18, the respective figures will be reported in the next version of this document, in M24.

KPI	Expected until M36	Reached until M12
Number of training activities until M36	2	1
Number of participants per activity	20	130
Number of web-based training activities	1	-
Number of Workshops / Symposia / Tutorials	2	1 (tutorial)

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

Number of participants per workshop	50	-
Number of exhibition booths:	2	-

1.4 Conclusions and Future Plans

The large volume of dissemination and communication activities (especially in terms of publications) that occurred within the first year of the project, as well as the values of the KPIs, as detailed in the sections above, show that OpenDR is in a very good shape, despite the negative effect of COVID-19 and the fact that it is still at an early stage of its lifecycle. Obviously, as the work progresses and more results materialize, especially when the first version of the toolkit will become available, the dissemination efforts will intensify towards reaching the full set of KPI targets defined in DoA. The consortium will continue to monitor its performance and implement corrective measures if necessary.

2 Initial Exploitation Plan

2.1 Short Description of Exploitation Plan

The toolkit developed within the OpenDR project will provide a modular and open library that could be easily applied to common robotics scenarios and robotics hardware.
The industrial partners of the project have plans to exploit the resulting Deep Learning library to enhance their own robotic platforms (PAL and AGI) and robotic simulator sample scenarios (CYB) with additional AI and cognition capabilities and improved human-robot interaction.

In this early stage of the project, the available results are still fragmented and they cannot be applied to the robotics platforms and simulation yet. For this reason, all the industrial partners focused on preparing the infrastructure to integrate and exploit the OpenDR toolkit and collect feedback from customers and make sure that it meets the customer's demands. In M13 the consortium will start working on the integration of the different available tools in a common toolkit library and this will produce mature results that could be applied on the robotics platforms.

In the following sections, the three industrial partners of the OpenDR project describe in detail their current exploitation plan and the actions taken during this first year.

2.1.1 CYB exploitation plan

CYB plans to use the OpenDR toolkit and developed robotics scenarios to demonstrate the capabilities of the Webots robotics simulator in the Deep Learning area. An increasing number of Webots users is interested in DL and including additional simulation developed within the OpenDR project in the Webots built-it library will help acquiring more users and increasing the Webots community. New Webots releases are published regularly and CYB will be able to integrate the new OpenDR scenarios at the latest in the major Webots release published each year at the end of December and June. By extending the Webots library and community CYB expects to gain more visibility inside the robotics community and to receive more requests for paid services. In this early stage of the project, the consortium is still

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

developing the initial tools and simulations, so more concrete actions are planned for a later stage of the project.

CYB also plans to turn some project's results into web services and charge for the usage. Users will be able to run their own simulations using the OpenDR toolkit on the web without having to install any tools, but they could be charged for running their simulations on the provided host machine. For this purpose, CYB started to work on the development of the web platform to run users' simulations on the cloud. In this way, as soon as some simulations will be ready, they can be easily shown and disseminated on the web. Besides giving more visibility to the OpenDR project, services offering robotics simulation based on DL on the web are increasingly in demand and CYB will take this opportunity to jump into this market and extend the range of its services. Also, for this objective, at the moment CYB can only take actions to prepare the future exploitation that will be applied in a later project phase when a first version of the toolkit and scenarios will be available.

2.1.2 PAL exploitation plan

PAL Robotics is a leading provider of service robots (Fig. 48), especially involved in the future of the Socially Assistive Robots (SAR) in healthcare and in robots for the so called Industry 4.0. SAR shares with assistive robotics the goal to provide assistance to human users, specifying that the assistance is through social interaction.

Figure 48: PAL Robotics service robots

But the environments where such SAR operate are often the same environment in which humans operate, so there is a lot of uncertainty and unpredictability. Following this goal, robots such as the TIAGo manipulator for healthcare, that is involved in the OpenDR project, that will operate in these environments need more skills - making enhanced Deep Learning essential for successful deployment.

The use of robots in non-industrial environments is set to continue increasing dramatically over the coming years. Creating assistive robots not only means combining several subsystems but also combining them into robots that can be adapted to follow the specific physical constraints that the older person is facing and the requirements of the environment where the person lives. The advances in AI are allowing more improved skills available for robots that can bring an added value to the future of robotics applications. An open source toolkit, like the one presented in OpenDR, are key to have more accessible AI skills and functionalities to a wider group of people and not only to the experts on the field. The integration with Deep Learning could bring out the cognitive capabilities within the robot.

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

Deep learning could assist the robot to manoeuvre, manipulate and understand the surrounding environment. In such a complex environment AI and Deep Learning help robots to deal with uncertainty. In the following some benefits Deep Learning can bring to healthcare robots are listed: facial recognition, emotion recognition, people tracking, enhanced fall detection, speech recognition, environment understanding, object recognition, object tracking, control of manipulation of objects in case of uncertainty.

Internal exploitation activities for PAL will include regular meetings with the 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. The PAL use case will be useful to get the possibility to execute 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. We would like to put forward the effort that we started in the field of assistive robots and focus on commercializing our robots making easier their use in an unpredictable healthcare environment. In order to achieve this, we foresee a market analysis that will be done to 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 also will 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 different possibilities of collaboration on singular aspects of the project with any specific partner and also the liaison with other projects of healthcare where PAL is involved, like the SHAPES and SPRING projects, to check the availability to test and use the OpenDR toolkit outside the consortium, when a first version of the toolkit will be available.

Moreover, work will be done in OpenDR to improve the configuration of the Nvidia Jetson used to extend the robotic platform performances to perform very consuming tasks like the ones using the Deep Learning tools. The study of the use of extra devices like the Nvidia Xavier will be also conducted, to improve the execution performances of the Deep Learning algorithms.

2.1.3 AGI exploitation plan

Robotti is a field robot, capable of performing various different operations in a farmer's field such as harrowing, rotovating, seeding, mechanical weeding, spraying, etc. A farmer's field is an unstructured environment with the potential of many dangers both for the robot and its surroundings. Currently, Robotti is fitted with several safety features, but Deep Learning can help provide an extra safety layer by finding obstacles at a greater distance than the laser sensor and bumpers are able to. In addition, Deep Learning is able to identify the crop, enabling various different technologies to utilize this algorithm, such as intra-row mechanical weeding.

In this project, AGI focuses on two main Deep Learning algorithms: obstacle detection and crop plant detection. In addition, an intra-row weeder making use of the crop detection algorithm will be developed for the Robotti, enabling state-of-the-art organic weed control. With current intra-row weeding systems, they are not able to detect the plant when the weed reaches a higher level of density or a certain size. The Deep Learning crop detection algorithm is expected to solve this problem. Assuming the price per hectare and capacity is at a high

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

enough level, it will also be attractive for conventional farmers to use the intra-row weeder, significantly decreasing their herbicide use.

A new version of AGI's field robot, Robotti 2021 will be announced at a public release show in February or March 2021 before the agriculture season starts. The forward-facing camera with obstacle detection algorithm will be a standard component on the Robotti 2021 version. Feedback from customers and from the market will be collected and feed back into engineering development for continual updates. This release show will include press, social media, etc.

During the 2021 season, development on the obstacle detection camera system will be continued. Data will be collected from the onboard cameras, increasing the dataset and improving the performance of the obstacle detection algorithm.

For the plant detection algorithm, there are currently no exploitation plans. These plans will be solidified after the in-field testing that will be performed in 2021. There are no exploitation plans for the intra-row weeder for 2021. There is a high level of demand for this product, so it is the intention of AGI to develop and commercialize the product as soon as it is ready.

3 First Project Newsletter

In accordance with the DoA, the consortium has created, under the coordination of AUTH, the first project Newsletter. The document, which gives a brief introduction to the project, its consortium and illustrates its major results so far, 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. Efforts will be also made to distribute the newsletter through relevant email lists such as euRobotics, 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

First Project Newsletter

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

Welcome to the 1st OpenDR Project Newsletter !

Newsletter 1.0: OpenDR challenges and current status

OpenDR Deep Robotics

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

Learning Curve Barrier	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 2022 with a total budget of 6.6 Million Euros.

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

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

OpenDR Consortium

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

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

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

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

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

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

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

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

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

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

Work Performed so far

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

Requirements and Specifications

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

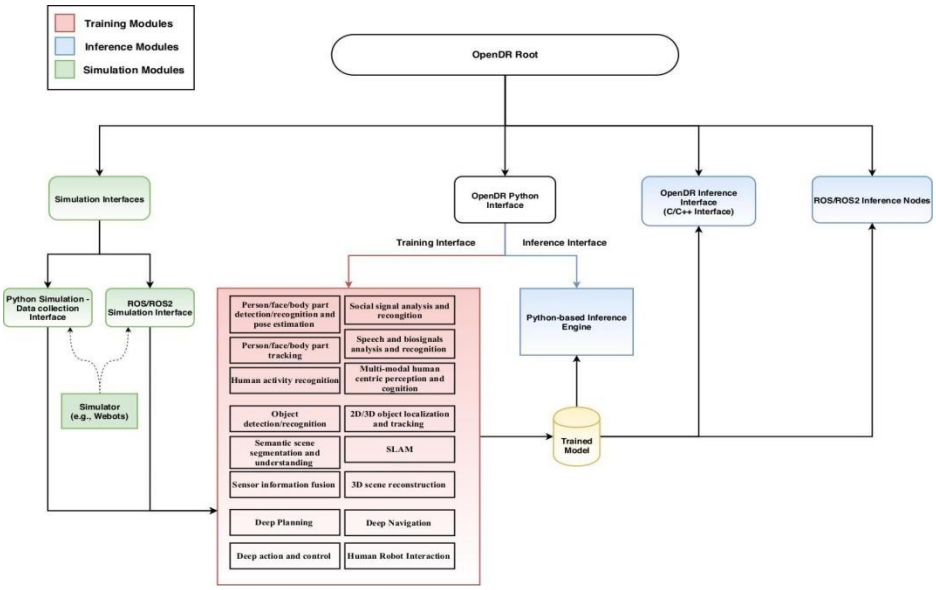

OpenDR Toolkit structure (better viewed in large magnification)

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

Deep Human Centric Active Perception and Cognition

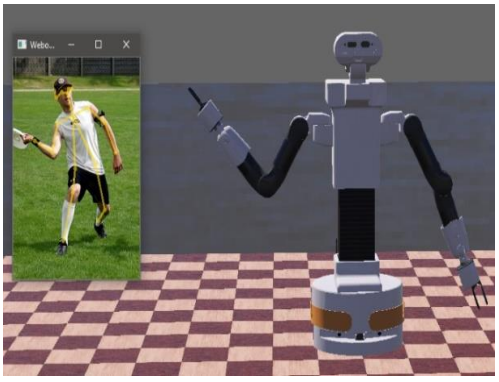

Demonstration of the TIAGo robot, implemented in the Webots simulator, mimicking a human's pose

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

Notable contributions were made in a number of research areas and the key project objectives were advanced, bringing the team closer to the realization of the powerful, flexible and efficient robotics toolkit. Deep learning models for robotics applications have to be lightweight due to computation power restrictions and achieve real-time performance. AUTH has worked on preparing such models for robotic vision tasks as well as models that adjust their runtime based on the available resources. AU, also made significant progress in skeleton-based human activity recognition, creating fast algorithms that can even outperform current state-of-

the-art solutions. TAU proposed a framework for multilinear compressed learning, which is highly efficient with respect to memory and computation, achieving superior performance for face recognition tasks. TAU has also worked on efficient solutions for anomaly detection in heart signals and speech command recognition.

Deep Environment Active Perception and Cognition

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

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

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

Deep Robot Action and Decision Making

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

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

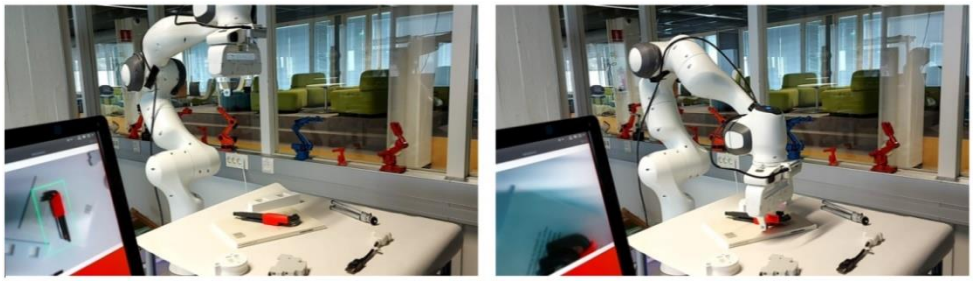

Single-Demonstration Grasping model at the hover pose (left) and the grasp pose (right)

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

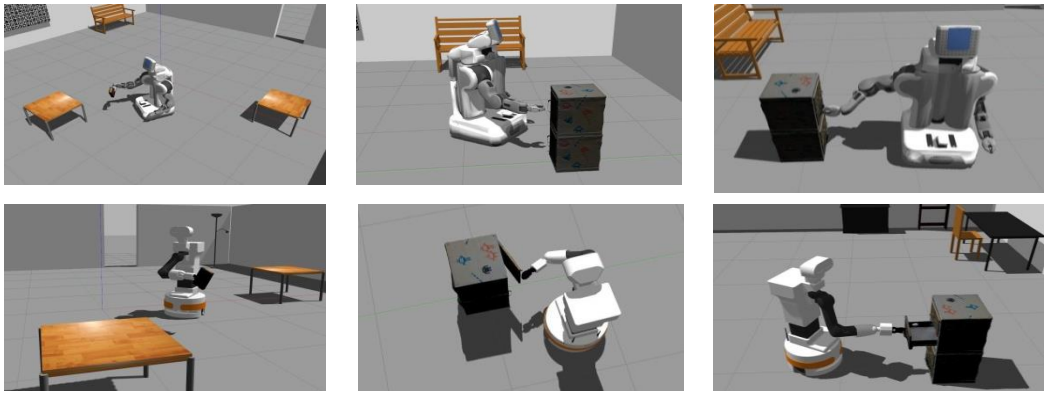

Example trajectories in the Gazebo physics simulator on different tasks for the PR2 (top) and TIAGo (bottom) mobile platforms

Simulation Environments and Data

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

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

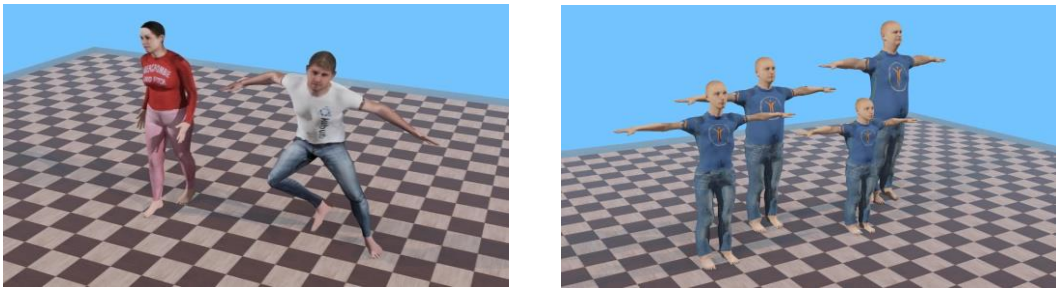

SMPL human body models in Webots generated with different shape parameters (right), reposed using existing animations (left)

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

Dissemination

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

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

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

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