Deep-Learning and HPC to Boost Biomedical Applications for Health (DeepHealth)

Monica Caballero
EVERIS Spain
Barcelona, Spain
monica.caballero.galeote@everis.com

Jon Ander Gómez
Universitat Politècnica de València (UPV)
Valencia, Spain
jon@upv.es

Aimilia Bantouna
WINGS ICT Solutions
Athens, Greece
abantoun@wings-ict-solutions.eu

Abstract— **This document introduces the DeepHealth project: "Deep-Learning and HPC to Boost Biomedical Applications for Health". This project is funded by the European Commission under the H2020 framework program and aims to reduce the gap between the availability of mature enough AI-solutions and their deployment in real scenarios. Several existing software platforms provided by industrial partners will integrate state-of-the-art machine-learning algorithms and will be used for giving support to doctors in diagnosis, increasing their capabilities and efficiency. The DeepHealth consortium is composed by 21 partners from 9 European countries including hospitals, universities, large industry and SMEs.**

Keywords- High Performance Computing, Big Data, Pilot test cases, Very large databases, Deep Learning, Biomedical applications.

I. DESCRIPTION AND PROJECT GOALS

Healthcare is one of the key sectors of the global economy, especially in Europe where the EU28 total expenditure on health care reached an average of 10.1% of GDP in 2015 [1]. Global healthcare expenditure is currently estimated a USD7.6 trillion and its growth will continue greater than the GDP, reaching double the rate in some countries [2]. Any improvement in healthcare systems has a high impact on the welfare society, hence, the use of technologies in health is clearly a strong path to more efficient healthcare, benefitting to people and public budgets. On the other hand, European public health systems are generating large datasets of biomedical data, in particular images as most medical examinations use image-based processes. These datasets are in continuous growth and constitute a large unexploited knowledge database since most of its value comes from the interpretations of the experts. Nowadays, this process is still performed manually in most cases and global knowledge sharing is complex.

In the context of automating and accelerating the analysis of the health data and processes, health scientific discovery and innovation are expected to quickly move forward under the so-called "fourth paradigm of science", which relies on unifying the traditionally separated and heterogeneous high-performance computing and big data analytics environments. Under this paradigm, the DeepHealth project will provide HPC computing power at the service of biomedical applications by using Hybrid HPC + Big Data computer architectures; and apply Deep Learning (DL) and Computer Vision (CV) techniques on large and complex biomedical datasets to support new and more efficient ways of diagnosis, monitoring and treatment of diseases.

In particular, DeepHealth aims at:

- Filling, as fast as possible, the gap between the availability of new technologies and making extensive use of them;
- Reducing the time to design and develop end-user applications deployed in software platforms;
- Increasing the productivity of IT staff working in the health sector by allowing them to design, train and test many more predictive models in the same period of time; and
- Providing medical personnel with a friendly and individualized digital decision support.

The technical objectives of the DeepHealth project are:

- Developing the DeepHealth toolkit for training and testing predictive models based on Deep Neural Networks (DNNs).
- Developing the European Distributed Deep Learning Library (EDDLL), that will be part of the toolkit. This library will be of general purpose, but in this project will be applied to medical imaging.
- Developing the European Computer Vision Library (ECVL), that will also be part of the toolkit. This library will act as a wrapper of other existing libraries for transforming/processing images, and will include specific image processing algorithms not available in other libraries but needed for some of the 14 use cases of this project.
- Integrating both libraries into existing software platforms in order to use predictive models designed and trained by using the DeepHealth toolkit.

II. PARTICIPANTS

The DeepHealth project involves 21 partners from 9 European countries, gathering a multidisciplinary group from research organizations (10), health organizations (4) as well as (4) large and (3) SME industrial partners, with strong commitment towards innovation, exploitation and sustainability.

The involved research organizations are:

a) *Universitat Politècnica de València (UPV)*

b) *Commissariat à l'Énergie Atomique et aux Énergies Alternatives (CEA)*

c) *Barcelona Supercomputing Center (BSC)*

d) *Università degli Studi di Torino (UNITO)*

e) *Università degli Studi di Modena e Reggio Emilia (UNIMORE)*

f) *Centro di Ricerca Sviluppo e Studi Superiori in Sardegna (CRS4)*

g) *Otto von Guericke University Magdeburg (OVGU)*

h) *Ecole Polytechnique Federale De Lausanne (EPFL)*

i) *Stelar Security Technology Law Research UG (STELAR)*

j) *Karolinska Institutet (KI)*

The health organizations are:

a) *Azienda Ospedaliera Citta Della Salute E Della Scienza Di Torino (CDSS)*

b) *Centre Hospitalier Universitaire Vaudois, Lausanne (CHUV)*

c) *Spitalul Clinic "Prof. Dr. Th. Burghele" (SCHTB)*

d) *Fundación para el fomento de la Investigación Sanitaria y Biomédica de la Comunitat Valenciana (FISABIO)*

The large industrial partners are:

a) *Everis Spain SLU (EVR)*

b) *Philips Medical Systems Nederland B.V. (PHILIPS)*

c) *Thales Six Gts France SAS (THALES SIX)*

d) *SIVECO Romania SA (SIVECO)*

The SME industrial partners are:

a) *WINGS ICT Solutions Information & Communication Technologies IKE (WINGS)*

b) *TREE Technology SA (TREE)*

c) *PRO DESIGN Electronic GmbH (PRO DESIGN)*

Several partners are members of private public partnerships. BSC, CEA, CRS4, EVR, PHILIPS, THALES, CDSS, SIVECO, UNITO, UPV and WINGS are members of the Big Data Value Association (http://bdva.eu). PRO DESIGN and UPV are members of the European Technology Platform for High Performance Computing (https://www.etp4hpc.eu).

III. EXPECTED RESULTS

DeepHealth is a three-year project that started on January, 1st 2019 and will end on December, 31st 2021.

The DeepHealth project will develop a flexible and scalable framework for running on HPC and Big Data environments. The DeepHealth framework will be based on two new core technology libraries to be developed within the project and integrated into existing (biomedical) software platforms: the European Distributed Deep Learning Library (EDDLL) and the European Computer Vision Library (ECVL). These libraries will take full advantage of the current and coming development of HPC systems deployed over Europe, and will provide a transparent use of heterogeneous hardware accelerators to optimize the training of predictive models, while considering performance and accuracy trade-offs. The DeepHealth framework will also include a front-end to simplify the use of the core technology libraries, both for the manipulation of images and for the training of Deep Neural Networks (DNNs). In summary, what from now on we are going to reference as the DeepHealth toolkit will integrate three components: the EDDLL, the ECVL and the front-end.

Both training and predictive algorithms will be able to run on Hybrid HPC + Big Data architectures with heterogeneous hardware independently of the existing (biomedical) software platforms. Then, these platforms will benefit from the DeepHealth framework since they will be able to directly use the models trained by using the toolkit.

The libraries will be integrated into seven existing software platforms: three commercial platforms: **Open Innovation** (PHILIPS), **Lumen** (EVR) and **PIAF** (THALES), and four research-oriented platforms: **MigraineNet** (WINGS), **ExpressIF™** (CEA), **Open-DeepHealth** (UNITO) and **Digital Pathology** (CRS4).

The DeepHealth framework will be validated in 14 pilot test-beds following real world use cases and datasets in three main areas: (i) neurological diseases, (ii) tumor detection and early cancer prediction; and (iii) digital pathology and automated image annotation. In all the pilots, the privacy requirements will be strictly satisfied, since no data will be transferred from the servers of the use case providers to any computer of any other partner.

UC1: Migraine & Seizures prediction (WINGS). Personalized migraine predictions have been already achieved by the powerful mechanisms of the MigraineNet platform by combining user's information, AI algorithms and DL techniques. User's information includes the migraineur's user profile, migraine attacks characteristics and the patient's daily routine. A main target of this use case is the reduction of the required model training and evaluation times through the parallelization and use of HPC technologies to be offered by the project. Additionally, the

quickly processing data thanks to simultaneously developing/training and evaluating prediction models for more migraineurs is expected to enhance the scalability of the MigraineNet system. Last but not least, expansions of platform's implementation in other health issues such as epileptic seizures are going to robust platform's potential, because the DL techniques will also be applied to meet the requirements for seizure prediction using Electrocardiogram (ECG) and Electroencephalogram (EEG) data.

UC2: UNITOPath – Classification of whole-slide histological images of colorectal biopsy samples (UNITO). Histo-pathological characterization of colorectal polyps is the major instrument for the pathologist in order to examine and infer the actual risk for cancer. The whole-slide are scans from H&E stains, in which *haematoxylin* and *eosin* are chemical substances used to achieve visible color contrast in biological tissues also for microscopic pictures. A total of six different classes have been chosen: (1) Hyperplastic polyp (HP), (2) Normal tissue (NORM), (3) Tubular Adenoma, high-grade dysplasia (TA.HG), (4) Tubular Adenoma, low-grade dysplasia (TA.LG), (5) Tubulovillous Adenoma, high-grade dysplasia (TVA.HG) and (6) Tubulovillous Adenoma, low-grade dysplasia (TVA.LG).

UC3: Brain (UNITO and CDSS). This UC is based on data provided by the Interventional Neuroradiology Unit of Città della Salute University Hospital. The task will be to synthetize CT Perfusion maps of the brain from reduced data sets, maps that are informative enough to be used for ordinary clinical purposes. Occlusion of a cerebral vessel causes a sudden decrease of the blood perfusion of the corresponding vascular territory. Blood flow deficit is greater at the centre of this area, being the most peripheral regions furnished by collateral circulation. Since brain resistance to ischemia is limited, a lesion develops in a few minutes or hours, starting from the centre of the ischemic territory, and progressively enlarging up to the most peripheral territory. The lesion brain at a given time is defined "ischemic core", the surrounding at risk brain, not death yet, is defined "penumbra". Distinction between core and penumbra is critical for efficiently cure the patient. CT perfusion has been demonstrated to allow to identify both regions. CT perfusion measure, with a sample time of roughly 1 Hz, the passage of a contrast media bolus into the brain, on a pixel-by-pixel basis. Serial low-dose scans are acquired; time-density curves, corresponding to the contrast media passage in brain tissue, are calculated; parametric maps are calculated too. The most relevant parameters used in clinical practice are Cerebral Blood Volume (CBV) and Cerebral Blood Flow (CBF), the integral and the tangent of the curve respectively. The aim is to demonstrate that DNNs may generate synthetic CBV and CBF maps with reduced and/or noisy sets of data, since they perform well in such conditions.

UC4: Chest (UNITO and CDSS). Lung nodules are small focal lesions in lung parenchyma, can be solitary or multiple and in many cases are accidentally found in CT scans: their identification is time consuming in current clinical activity

for the radiologist and, since these small lesions are difficult to be "labelled", patients often need to perform follow-up CT scans in order to assess their benignity/malignancy, resulting in increased radiation exposure and anxiety for the patient and increased work load for doctors. The goal of UC4 will be to train AI-based systems in order to recognize lung nodules using chest CT scans, providing radiologists with an efficient tool for daily activity.

UC5: Deep Image Annotation (UNITO and CDSS). Automatic annotation of image-based medical examinations with text in natural language is desirable and necessary. This use case aims at implementing and evaluating machine learning models for the automatic generation of medical reports, learning from datasets composed by image examinations and related reports. Recent developments in Deep Learning models allowed to automatically generate natural text from features extracted with a multimodal approach, where images and related text are combined in a joint semantic description [3,4].

UC6: Promort: A vertical application to support digital pathology in the context of prostate tumor diagnosis (KI and CRS4). A vertical application will be created that will use models based on the EDDLL to automatically suggest annotations for prostate tumor slides. The application will present these suggestions to a pathologist for validation and, if necessary, correction. The DeepHealth libraries will be used to exploit the Promort image dataset to train effective models and then to integrate the prediction functionality into the new vertical application. The Karolinska Institute will test this application in a clinical research setting, where it will be tested in a continuous model improvement cycle.

UC7: Major Depression (OVGU). Major Depressive Disorder (MDD), with a lifetime prevalence of 15%, affects nearly one in six people. Many of them are not in adequate treatment because the medical system does not have enough capacity or they are unaware of their disorder. In the future, decision support systems will be able to help doctors in diagnosing and monitoring disorders in cases where specialist services are not available. This use case will use clinical, longitudinal and omics data from patients as well as physical activity and biological data during the disease course in order to identify patterns that are relevant to the treatment of the patients. The main target of this use case is to perform classification of disordered vs. healthy states. To this end, state-of-the-art machine learning based algorithms will be used to detect disorders over big amounts of neuroimaging and biological marker data. A secondary aim is to establish predictors for a change in clinical state (e.g. prediction of therapy outcome under treatment) or to recommend the k treatments with a high probability to be successful.

UC8: Dementia (OVGU). Dementia is an age-related disorder affecting 46 Million people that includes deficits in cognitive, emotional and social abilities with both degenerative (e.g. Alzheimer's Disease) and nondegenerative

causes. In the second case, there are applicable treatment options. Therapy and enhancing of cognitive reserve can be most effective when accurate diagnosis and very early diagnosis can be achieved. The diagnostic approaches being involved are interdisciplinary with psychiatry, neurology, neuroradiology, and clinical chemistry. Decision support systems will be most effective when they are running on fast platforms. These systems could be used by all general practitioners in addition to neurologists, psychiatrists and psychologists. This use case will exploit clinical, longitudinal, imaging, laboratory and omics data from patients in order to create models, i.e., to identify disease specific patterns, that can better predict the clinical course of the patient and can lead to treatment options. The main target of this use case is to classify healthy states vs. several forms of dementia (Alzheimer, vascular, fronto-temporal, mixed) and mild cognitive impairment. Also, an aim is to find predictors for evaluating the disease progression, what can be helpful for clinical guidance or in the decision for the necessary support that the subject needs in the near future. To this end, state-of-the-art machine learning based algorithms will be used for classifying subjects by using biological markers and neuroimaging data.

UC9: Study of structural changes in lumbar spine pathology. (FISABIO). The analysis of 8,597 patients, with a total of 42,436 MRI Scans (Axial-T2, Sagittal-T1, Sagittal-T2 and Sagittal-STIR) from the Public Health Service of the Valencian Region will be used to create a population database of medical imaging associated with Electronic Health Records, anonymized and interoperable with the Medical Imaging Databank of the Valencian Region (BIMCV). This will constitute the MIDAS dataset: Massive Image Data Anatomy of the Spine. Lumbar spine use case aims to develop an infrastructure with massive storage capacity and intensive computational modelling of lumbar magnetic resonance images together with clinical data of patients with low back pain so as to investigate possible correlations between pathology and structural alterations.

UC10: Predictive and Population Model for Alzheimer's Disease (AD) using Structural Neuroimaging (FISABIO). The 10K Big Data in Brain Imaging of Valencia Region aims at exploiting the rich information of MRI images from more than 10,000 individuals. This dataset provides detailed anatomical and morphological data of the brain, and it is composed by 2D clinical images corresponding to an average of 24 slices per session, and 6 sessions per patient in average, so the total number of 2D images is greater than 1.4 Million. When reconstructed into 3D images more than 60,000 samples will be available. This use case aims to correlate brain identified surface areas according to different criteria (age, gender, pathologies, and geographical area) in the population of the Valencian region. For achieving this purpose, the abovementioned MRI images will be used. This kind of images make possible to quantify all the structures of interest, in order to compare which areas of the brain produce the possible variability between healthy and pathological subjects, specifically in relation to Alzheimer's disease.

UC11: Image Analysis and prediction for Urology (SIVECO and SCHTB). The image analysis is of utmost importance for a correct diagnosis in urology. Urologic Radiology includes several ways to carefully look inside organs like the kidneys, ureters, bladder, and reproductive organs. Imaging tools range from the basic ones like x-rays to more complex tools. Generally, these imaging tests are quick and offer a high level of detail, most are pain-free. The various existing imaging technologies have meant greater diagnostic power and lead to an important increase in the number and quality of images obtained and viewed by practicing clinicians, mainly in the field of urology, where static and dynamic images are fundamental to the diagnosis and treatment of almost all conditions. Our focus will be on the diagnosis from images of kidney tumors (benign and malignant) and also adrenal tumors, for which the use of AI for diagnosis purposes is rarely described in the specialized literature. A main target of this use case is to develop a catalog of datasets with anonymized medical images and its related metadata obtained from the patients with renal and adrenal tumors, together with their radiological description and diagnosis. Using the DeepHealth toolkit, the objectives of the use case will be to growth in the volume of imaging services, improving the diagnostic capabilities with the help of AI-based solutions, and leading to changes in the standard of care for these conditions and expansions in the coverage for new diagnostic imaging modalities.

UC12: Skin cancer melanoma detection (UNIMORE). Skin cancer is the most common cancer in humans and it has an increasing trend. It affects mainly Caucasians and accounts for a very high probability in life, especially in relation to increased life-span. Even if its mortality rate remains low, its morbidity is relevant in the population, demanding for surgery frequently occurring on the face with disfiguring consequences. This use case is based on data available from a public online archive and an internal database which will be organized and made public within the ethical requirements. The aim is to provide a useful tool from the segmentation of medical images and the automated melanoma detection from dermoscopy images. To this end, clinical, dermoscopy and confocal microscopy images, annotated with conclusive diagnosis (histologic or clinically confirmed), and relevant patient's data will be used so as to facilitated the interpretation of a) magnified skin images of the tumor and visualization of its sub-surface structures (coming from dermoscopy), as well as b) in vivo confocal microscopy images, which provide nearly histologic information without tissue excision.

UC13: Epileptic seizures detection (EPFL and CHUV). The main target of this use case is to perform detection of epileptic seizures on ECG and EEG signals. To this end, state-of-the-art machine learning based algorithms will be trained and will need to be constantly running the epilepsy detection algorithms over a very large amount of data.
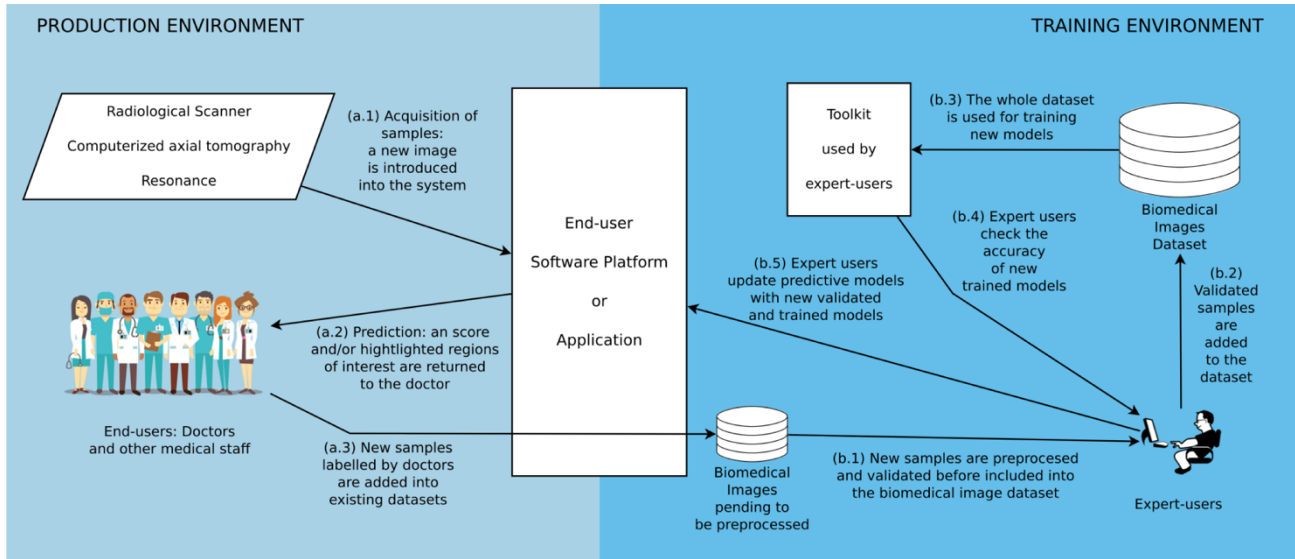
Figure 1. A typical scenario where tasks related to biomedical image processing are carried out. Steps a.1, a.2 and a.3 are done in production environments on a daily basis. Steps b.1 to b.5 in training environments are carried out when necessary.

A private database provided by CHUV will be used. The algorithm used for preprocessing and classifying will initially be the most relevant one in the state-of-the-art that use Deep Learning techniques for population-wide analysis [5,6,7].

UC14: Objective fatigue assessment for Multiple Sclerosis patients (EPFL/CHUV). This use case can be divided into two main objectives: 1) the segmentation of MRI images using state-of-the-art Deep Learning techniques to extract the lesion and set an initial baseline for the patient, and 2) the association between the parameters obtained from different biosignals captured by wearable devices with a fatigue index for the patient, with the ultimate aim of obtaining objective indicators of fatigue. A prospective dataset will be acquired for this use case with data from 20 to 50 patients. While the first objective is fully aligned with the DeepHealth project goals and will take advantage of the new framework to improve the state-of-the-art techniques [8,9,10], the second objective has not yet been explored and there are no yet Deep Learning algorithms validated for this purpose.

IV.   KEY PERFORMANCE INDICATORS

Use cases pilots will allow to train models and evaluate the performance of the proposed solutions in terms of the following KPIs: **time-of-pre-processing-images**, **time-to-model-in-production** and **time-to-train-models**. These three KPIs will be measured in hours. Additionally, the performance of the training and predicting algorithms will be measured by the **speedup** and the **efficiency of parallelism**, two standard KPIs widely used in parallel and distributed computing works. Concerning the accuracy of predictive

models, standard KPIs as **error rate**, **sensitivity** vs **specificity**, or just the **accuracy** will be used.

V.   EXPECTED IMPACT

As commented in Section 1, one of the goals of the DeepHealth project is to increase the productivity of IT staff working in the health sector by allowing them to design, train and test many more predictive models in the same period of time. Here is where it is expected the highest impact of this project.

Figure 1 shows a typical scenario where image processing is needed for diagnosis. Doctors (mainly radiologists) and other medical personnel are the end users. IT staff, typically computer scientists, are the expert users. Doctors have the knowledge to label images, define objectives and provide metadata in relation to images. IT staff are in charge of processing the images provided by end users, organize the datasets, detect anomalies, perform image transformations when required, and train predictive models. These models are used by doctors via software platforms in order to get clues that can help them in the diagnosis. For doctors this process is transparent, they provide images to the system and get responses (e.g. value of an indicator or bio-marker, segmentation of the image for identifying different elements as tissues, bones, nerves, blood vessels, etc.).

The training environment in Figure 1 represents where IT staff work with datasets of images for training predictive models. The production environment in the same figure represents where the medical personnel use a software platform or application to get knowledge from a single image or from a set of images coming from a scan session that generate multiple slices. Trained predictive models are integrated into software platforms by expert users.

Designing and training predictive models based on DNNs requires a high technological background in Deep Learning. Getting this background requires several years of study and great effort. Not all the IT staff working in hospitals have available the time needed to acquire such strong background. Typically, they can be experts in Image Processing, but it is not realistic to expect that all computer scientists working in the health sector become experts on Deep Learning.

Therefore, DeepHealth is expected to impact on facilitating the daily work of computer scientists working in the health sector. The DeepHealth toolkit will allow IT staff to (1) design DNNs without a profound knowledge in Deep Learning, and (2) run the training algorithms over Hybrid HPC + Big Data architectures without a profound knowledge of HPC or Big Data or distributed computing or Cloud computing.

Beyond that, the DeepHealth project tackles three major areas of great interest and impact to the health community, due to their high economic costs and high societal impact: neurological diseases, early tumor detection and digital pathology. Expected impacts in these three areas and the respective use cases include, but are not limited to: a) improving well-being, b) increasing early diagnosis and improving treatments, c) increasing diseases and pathologies knowledge, and d) saving direct and indirect healthcare costs.

Finally, we would like to remark that the outcomes of this project will be useful to other sectors besides health because the EDDLL has been designed to be a general-purpose library for training any type of deep neural network, and the ECVL functionalities will also apply to image processing in general, not only for biomedical images. Hence, the impact of the DeepHealth project will be on two directions. Vertically on the health application domain. Horizontally on three technology areas, HPC, Deep Learning and Computer Vision, which are closely related because Deep Learning techniques are widely used in Computer Vision solutions, and both, DL and CV techniques will benefit from HPC regarding performance.

VI. SCOPE OF THE USE OF THE DEEPHEALTH FRAMEWORK

During the development of this three-year project, the two libraries of the DeepHealth toolkit (EDDLL and ECVL) will be integrated into seven software platforms. In a similar way, other commercial or research software platforms will have the possibility of integrating the two libraries with no additional costs than the person hours needed for the integration, given that the DeepHealth toolkit will be delivered as free and open-source software.

Integrating the libraries into software applications is an option, but the DeepHealth toolkit will be useful by itself, so that any biomedical imaging practitioner will have the chance to use it for any purpose. Additionally, European large industry companies and SMEs will have the chance of using the DeepHealth toolkit for training predictive models on Hybrid HPC + Big Data architectures, because the format for storing the DNNs will be the standard ONNX [11], so, if

their software applications already integrated another DL library able to read the ONNX format, all work is done.

In addition to its use in commercial applications, the DeepHealth toolkit will be ready to be used in future research and innovation projects under both the current H2020 and the new Horizon Europe (2021-2027) [12] research and innovation frameworks.

REFERENCES

[1] 2018 Ageing Report, European Commission, https://ec.europa.eu/info/sites/info/files/economy-finance/ip079_en.pdf

[2] World Economic Forum, Value in Healthcare, Insight Report April 2017.

[3] Shin, H. C., Roberts, K., Lu, L., Demner-Fushman, D., Yao, J., & Summers, R. M. (2016). Learning to read chest x-rays: Recurrent neural cascade model for automated image annotation. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 2497-2506).

[4] Jing, B., Xie, P., & Xing, E. (2018). On the automatic generation of medical imaging reports. In *Proceedings of the 56ᵗʰ* Annual *Meeting of the Association for Computational Linguistics (Long Papers)*, pages 2577–2586 Association for Computational Linguistics. *arXiv preprint arXiv:1711.08195*.

[5] Acharya, U. R., Oh, S. L., Hagiwara, Y., Tan, J. H., & Adeli, H. (2018). Deep convolutional neural network for the automated detection and diagnosis of seizure using EEG signals. Computers in Biology and Medicine, 100, 270–278. https://doi.org/10.1016/J.COMPBIOMED.2017.09.017

[6] Ullah, I., Hussain, M., Qazi, E., Aboalsamh, H. (2018). An Automated system for epilepsy detection using EEG brain signals based on Deep Learning Approach. https://arxiv.org/abs/1801.05412

[7] Yuan, Y., Xun, G., Jia, K., Zhang, A. (2019). A Multi-View Deep Learning Framework for EEG Seizure Detection. IEEE Journal of **Biomedical** and Health Informatics, 23 (1), 83-94.

[8] Brosch, T., Tang, L. Y. W., Yoo, Y., Li, D. K. B., Traboulsee, A., & Tam, R. (2016). Deep 3D Convolutional Encoder Networks With Shortcuts for Multiscale Feature Integration Applied to Multiple Sclerosis Lesion Segmentation. IEEE Transactions on Medical Imaging, 35(5), 1229–1239. https://doi.org/10.1109/TMI.2016.2528821

[9] Valverde, S., Cabezas, M., Roura, E., González-Villà, S., Pareto, D., Vilanova, J. C., Lladó, X. (2017). Improving automated multiple sclerosis lesion segmentation with a cascaded 3D convolutional neural network approach. NeuroImage, 155, 159–168. https://doi.org/10.1016/J.NEUROIMAGE.2017.04.034

[10] Roy, S., Butman, J. A., Reich, D. S., Calabresi, P. A., & Pham, D. L. (2018). Multiple Sclerosis Lesion Segmentation from Brain MRI via Fully Convolutional Neural Networks. ArXiv preprint http://arxiv.org/abs/1803.09172

[11] ONNX: Open Neural Network Exchange format. The new open ecosystem for interchangeable AI models. https://onnx.ai/supported-tools

[12] Horizon Europe – the next research and innovation framework programme: https://ec.europa.eu/info/designing-next-research-and-innovation-framework-programme/what-shapes-next-framework-programme_en