




Contents lists available at ScienceDirect

Computer Methods and Programs in Biomedicine

journal homepage: www.elsevier.com/locate/cmpb

IODeep: An IOD for the introduction of deep learning in the DICOM standard

Salvatore Contino¹, Luca Cruciatà¹, Orazio Gambino^{*}, Roberto Pirrone

Department of Engineering, University of Palermo, Palermo, 90128, Sicily, Italy

ARTICLE INFO

MSC:
00-01
99-00

Keywords:

Medical image segmentation
Decision making in medical diagnosis
DICOM
Information Object Definition
Deep Neural Networks
Artificial Intelligence

ABSTRACT

Background and objective: In recent years, Artificial Intelligence (AI) and in particular Deep Neural Networks (DNN) became a relevant research topic in biomedical image segmentation due to the availability of more and more data sets along with the establishment of well known competitions. Despite the popularity of DNN based segmentation on the research side, these techniques are almost unused in the daily clinical practice even if they could support effectively the physician during the diagnostic process. Apart from the issues related to the explainability of the predictions of a neural model, such systems are not integrated in the diagnostic workflow, and a standardization of their use is needed to achieve this goal.

Methods: This paper presents *IODeep* a new DICOM Information Object Definition (IOD) aimed at storing both the weights and the architecture of a DNN already trained on a particular image dataset that is labeled as regards the acquisition modality, the anatomical region, and the disease under investigation.

Results: The IOD architecture is presented along with a DNN selection algorithm from the PACS server based on the labels outlined above, and a simple PACS viewer purposely designed for demonstrating the effectiveness of the DICOM integration, while no modifications are required on the PACS server side. Also a service based architecture in support of the entire workflow has been implemented.

Conclusion: *IODeep* ensures full integration of a trained AI model in a DICOM infrastructure, and it is also enables a scenario where a trained model can be either fine-tuned with hospital data or trained in a federated learning scheme shared by different hospitals. In this way AI models can be tailored to the real data produced by a Radiology ward thus improving the physician decision making process.

Source code is freely available at <https://github.com/CHILab1/IODeep.git>.

1. Introduction

As reported in the most recent scientific literature, state-of-the-art techniques for segmenting biomedical images make use of AI, and in particular Deep Learning (DL) models that provide an average accuracy higher than 90% on most tasks. Nevertheless, DL models are almost no present in the daily diagnostic practice, due to three main problems stated below:

- Statement 1: DL models are not integrated in the medical images diagnostic platforms.
- Statement 2: DL models can not be trained on huge data sets coming from radiologic wards due to issues related to both the patient's

data privacy and the security issues implied in building a computing infrastructure distributed across several hospitals.

- Statement 3: Physicians do not trust the predictions by a neural model as they do not understand the “clinical way” used by the model to segment regions that are relevant for the diagnosis.

Despite the problems stated above, using DL models in biomedical imaging platforms would bring multiple benefits:

- Improve efficiency: DL models can automate routine tasks such as image segmentation and analysis, freeing up radiologists' time and improving workflow efficiency.

* Corresponding author.

E-mail address: orazio.gambino@unipa.it (O. Gambino).

¹ These authors contributed equally.

<https://doi.org/10.1016/j.cmpb.2024.108113>

Received 15 January 2024; Received in revised form 22 February 2024; Accepted 1 March 2024

Available online 8 March 2024

0169-2607/© 2024 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

- Cognitive and visual fatigue decreasing: The radiologists' cognitive and visual fatigue entail the wrong diagnosis, whereas an AI algorithm is not affected by this. As a consequence, AI is a valid support during the image analysis by the radiologist [1].
- Increase accuracy: DL models can help to identify very subtle patterns in medical images that may be missed by the human eye, leading to more accurate and reliable diagnoses [2].
- Enhance the quality of care: AI can provide additional information to support clinical decision-making, helping healthcare providers to make more informed and effective treatment decisions [3].
- Facilitate research: AI based techniques can help to automate the analysis of large data sets, enabling researchers to identify new patterns and insights that may not be apparent through traditional methods.
- Mass screening: the ability to analyze large data sets provided by AI algorithms can be used to perform mass screening on the occasion of campaigns to raise awareness of a particular illness [4].
- Improve patient outcomes: by improving the speed and accuracy of diagnoses, DL models can help to improve patient outcomes and reduce the need for follow-up testing.

In view of the previous considerations, this work provides a complete solution to the problem formulated in Statement 1, and it also represents an enabling technology for answering both Statement 2 and Statement 3. The system proposed in our work leverages the well known DICOM (Digital Imaging and Communications in Medicine)² standard to achieve the solution, by integrating in the normal DICOM communication between a PACS server and a viewer, the storage and retrieval of a DNN trained purposely on a data set that is related to a particular acquisition modality, anatomical region, and disease. We achieve our stated goal through the creation of a proper Information Object Definition (IOD) that we called *IODeep*. *IODeep* contains both the architecture and the weights of a neural network purposely trained on a medical image data set; in our approach, several DNNs are stored in the PACS server, each in a separate *IODeep* instance, and a novel algorithm is proposed to select the best matching DNN for the images under investigation, based on the information mentioned above that *IODeep* stores in some proper DICOM tags. Network selection is transparent for the physician who only asks the system to start AI based Regions of Interest (ROI) proposal. The selected network is used to perform client-side prediction of the ROI on the slice that the physician is currently looking at. On the PACS viewer side, ROIs are visually proposed to the physician who can explicitly validate them as relevant for the diagnosis. The PACS client embeds the set of validated ROIs in an RT Structure Set IOD, which is stored on the server. The entire workflow is conducted in accordance with the DICOM protocol, preserving the security required by the standard, and no modification is required on the PACS server side. The main contributions of the paper are the following:

- The complete information architecture of *IODeep* along with the DICOM compliant workflow for both ROI prediction and storage on the PACS server.
- A novel algorithm based on the *IODeep* DICOM tags to select the correct DNN for the data under investigation in a transparent way for the physician.
- A purposely designed back-end independent JSON format to describe DNN architectures to be used as a possible DNN description inside *IODeep*.
- A purposely designed lightweight PACS viewer that implements all the client side communication involved in using *IODeep*, while no modification is required on the server side.

- A service based implementation of the entire ROI prediction scenario where both network selection and prediction run as external REST services that communicate with the PACS infrastructure

The rest of the paper is arranged as follows. Section 2 reports some theoretical background about DL models for medical image analysis. The motivations of our work are outlined in Section 3, while the *IODeep* architecture is described in Section 4 along with the detailed ROI prediction workflow, and a comparison with the current abstract model devised by the DICOM community for AI integration. Final remarks and conclusions are reported in Section 5.

2. Theoretical background

A relevant amount of recent papers in the scientific literature on biomedical image processing make large use of deep-learning approaches. Such types of algorithms are applied without any regard for image modality, pathology and body part with the aim to either segment or simply classify both 2D and 3D data. DNNs are used mainly in three kinds of task: classification, semantic segmentation, and region proposal.

Classification tasks aim at placing a (multi-)class label to the whole image, as in the case of histopathology where one is interested in assessing the presence of cancer tissues. Conversely, semantic segmentation is the task of labeling each pixel in the image as belonging to a certain class. In general, labels discriminate different tissues: a normal organ, a pathological one, a tumor lesion, and so on. Zhou et al. [5] propose a nested U-Net architecture for medical image segmentation. It soon achieved state-of-the-art results on several data sets including lung and liver CT scans. Schlemper et al. [6] introduce an attention mechanism to the U-Net architecture for pancreas segmentation in CT scans. The attention mechanism helps the model to focus on the relevant regions of the image, leading to improved segmentation accuracy. The approach proposed by Fu et al. [7] proposes a dual attention mechanism that has been adapted for medical image segmentation tasks with good results. The mechanism consists of a spatial attention module and a channel attention module, which capture spatial and channel-wise dependencies respectively. Dou et al. [8] propose a 3D anisotropic hybrid network for medical image segmentation, which combines 2D and 3D convolutional layers to better handle anisotropic medical images. Milletari et al. [9] introduce a family of V-Net architectures for volumetric medical image segmentation, which achieve state-of-the-art results on several data sets including liver, lung, and brain CT scans. The V-Net architectures are based on 3D convolutional layers and incorporate skip connections for improved segmentation accuracy. Christ et al. [10] propose a cascaded fully convolutional neural network (FCN) approach for automatic segmentation of liver and tumor from CT and MRI volumes. The proposed method consists of two stages, where the first stage segments the liver, and the second stage segments the tumor within the liver. The approach achieved state-of-the-art results on two liver tumor segmentation datasets. A review on the state of the art in deep learning-based methods for brain MRI segmentation is presented in [11]. The authors provide an overview of different types of deep learning models used for brain MRI segmentation, including fully convolutional networks, U-Nets, and deep residual networks. The paper also discusses current challenges and future directions for brain MRI segmentation. Qu et al. [12] present a deep learning-based method to segment pelvic bone tumors in MRI that is based on a multi-view fusion network to extract pseudo-3D information from two scans in different directions. Wang et al. [13] propose a novel deep neural architecture for segmenting multiple organs in abdominal CT scans, that is called Organ-Attention Network (OAN), and it is specifically designed to both classify and segment each individual organ with high accuracy. OAN uses a combination of convolutional and attention-based modules to selectively focus on each organ of interest while ignoring irrelevant regions in the image. Additionally, the paper proposes a statistical fusion

² <https://www.dicomstandard.org/>.

strategy to combine the individual organ segmentations generated by OAN into a complete multi-organ segmentation. Yan and colleagues [14] propose a deep learning approach for the automatic segmentation of pancreas from computed tomography (CT) scans. The proposed method uses a multi-scale U-Net architecture, tested on NIH pancreas segmentation data set [15], which includes an attention mechanism to highlight important regions in the input image. A very recent attention based CNN is a module based on Cortical Spiking models presented by Zhou et al. [16], that reduces the number of trainable parameters compared to well-known state-of-the-art architectures, while improving performance in terms of accuracy. In very recent times, the advent of Transformers networks, and self-attention mechanisms originated in the Natural Language Processing (NLP) community [17] have pushed also the research in Computer Vision towards the use of networks integrating this kind of technology. Vision transformers (ViTs) have become attractive due to their ability to encode long-range dependencies and to learn highly effective feature representations. Right now, ViTs are the new state of the art, surpassing CNN-based architectures in both classification and semantic segmentation tasks [18], adding more and more trainable parameters but decreasing the possibility of gaining information from eXplainable AI (XAI) techniques.

CNN-based networks are often used as a backbone for architectures trained to solve region proposal tasks. In fact, in contrast to semantic segmentation that needs pixel precision, region proposal architectures return predictions about ROIs, i.e. bounding boxes around the truly interesting pattern. Pun and Agarwal [19] in their 2022 survey analyzed the state-of-the-art of all U-Net variants used in region proposal tasks; the survey's emphasis is on the efficiency of the various variants in tricky tasks such as the diagnosis of severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2). Zhang et al. in 2023 [20] propose the MDU-Net, an architecture consisting of a multi-scale dense connections (MDC) encoder and a U-shaped decoder for biomedical image segmentation. The multi-scale dense connections, allowing shorter connections between layers close to the input and output, enable a much deeper U-net that reduces potential overfitting and further improves segmentation performance. Finally, Song et al. [21] in their work propose the Outlined Attention U-Network, (OAU-Net) based on a bypass branching strategy to solve biomedical image segmentation tasks, capable of detecting surface and deep features. The particularity of this architecture lies in the encoders that are based on residual convolution.

One of the main disadvantages of region proposal architectures is the high computational cost deriving from the increasing number of trainable parameters. In very recent years, Weakly Supervised Object Location (WSOL) and Weakly Supervised Semantic Segmentation (WSSS) have gained relevance. Weakly supervised techniques do not need explicit pixel-wise labeling to be trained. WSOL [22] is the task of learning how to place a bounding box or a loose ROI around a relevant object in the image, according to the label provided at image level. Many kinds of WSSS [23] are reported in the literature. The most common task in this respect is learning explicit (pixel-wise) segmentation of the object starting from either an image-wise or a box-wise label. Some recent WSSS approaches leverage the Class Activation Maps (CAM) [24]. CAM is a well known XAI approach that generates a map for each class to describe which features were most activated for predicting the class itself. In general, CAM based WSSS use the maps as "seeds" for predicting a first ROI on the object to be segmented. Then the seed ROI is used as a weak mask to train the proper semantic segmentation model [25].

ROI identification through weakly supervised techniques is the best option for the diagnostic purpose because the physician is not hurried in precise pixel level analysis; rather she or he is simply alerted that some interesting thing may be present in a certain zone of the image. Moreover, the model itself does not need to be extremely precise, and this is an advantage in the medical imaging scenario where few training data are at disposal. In view of the previous considerations we decided to design our AI-supported diagnostic scenario as a ROI proposal one. The

proposed workflow is agnostic with respect to the use of either bounding boxes or generic ROI contour because, in any case, the vertexes of the ROI can be properly stored in a DICOM RT Structure Set. Moreover, there is no difference in using IODep to describe a network that learns bounding boxes or a model trained to predict ROIs. Finally, our workflow prescribes that each ROI/bounding box has to be validated manually by the physician; thus we ensure clinical explainability of the whole process.

3. Motivation of the work

Even if it is well known from the literature that radiologists would like(semi-)automated decision support systems provided that they are non-invasive [26], DL models have not entered the routine diagnostic workflow due to lack of standardization. Indeed, the clinical community started addressing this issue from the DICOM data harmonization viewpoint with the aim of introducing a proper use of DICOM metadata for mixing established data sets and true clinical images to enable development, validation, and clinical translation of AI tools in a standardized way [27]. We claim that the main way to standardization is including DL in the DICOM standard. Such an integration would allow for the development of consistent and interoperable diagnostic tools across heterogeneous imaging platforms, making it easier for healthcare providers to integrate AI into their workflow. Second, a standardized framework for training and validating DL models is the only way to design distributed computing infrastructures where federated learning can be enabled. As well debated in the literature [28,29] federated learning is the key for implementing very effective models that are trained on adequate amounts of data thus providing high accuracy for a plethora of diseases, while scientific data sets are almost limited as regards both their size and the clinical goal they are designed for. In turn, very accurate models can be used effectively by XAI techniques to assess the most relevant image features which guided the model to a certain prediction [30].

Recent literature reports several frameworks in the field of medical imaging that can be integrated with the DICOM standard. The work in [31] introduces a vendor-independent platform that offers interfaces aimed at managing digitized slides and medical reports. It provides digital image analysis services compatible with other current standards. The solution incorporates the open-source Dicoogle PACS plugin architecture for interoperability and extensibility, enabling customization of the proposed solution. In [32] the Visilab Viewer is presented. It is a web-based platform compliant with the DICOM standard web architecture to solve the Whole Slide Image (WSI) format visualization from multi-frame DICOM images.

The literature reports also several frameworks that make use of the DICOM standard to achieve a sort of "loose" AI integration into the clinical workflow rather than trying to include them in the standard itself. In recent years, a great effort has been done for DICOM standardization of digital pathology where particularly effective pan-and-zoom characteristics are needed to visualize WSI data in a PACS system along with the connection to external AI applications. A fully DICOM implementation of a platform for digital pathology has been proposed in [33] where both visualization and ML applications can connect using the DICOM PS3.18 RESTful web services (DICOMweb™).³ In this architecture external applications have to implement the DICOMweb client API that connects to the DICOMweb server using the HTTPS protocol. Moreover, the authors selected a suitable set of DICOM metadata to accomplish the clinical requirements needed for diagnosing microscope images. In [34] the open decentralized platform devised by the EMPAIA Consortium is illustrated that enables pathologists to connect transparently using a web browser both to a medical and a computing infrastructure to man-

³ <https://dicom.nema.org/medical/dicom/current/output/html/part18.html>.

age their data while running third-party AI applications on the same data. The platform provides an abstraction on the DICOM infrastructure, and it is intended to be agnostic with respect to both PACS and WSI scanner vendors.

In [35] a roadmap is proposed for integrating AI image analysis algorithms into existing radiology workflows: a case study on brain MRI is presented. In [36] a framework called *Niffler* is presented, that is a Machine Learning (ML) framework retrieving images from the PACS using DICOM network listeners. Niffler extracts and processes metadata from the acquired images at the research clusters, then it executes both ML and real-time analytics pipelines on the radiological images, and their textual metadata. In [37] a ML model for the automatic analysis of medical images in the radiological workflow is integrated into a PACS system by using DICOM services provided by open-source tools. A DL architecture is trained for classifying chest X-ray images and it is reported as a case study. A DICOM Imaging Router is presented in [38] that incorporates CNNs for categorizing unknown DICOM X-ray images into five anatomical groups: “abdominal”, “adult chest”, “pediatric chest”, “spine”, and “others”.

The frameworks outlined above, connect themselves to a DICOM infrastructure to carry out their task. Therefore, their use requires many tricks that radiologists must learn [39]. In essence, the DICOM standard and the AI techniques remain two separate worlds that play their role in synergy when needed, without real integration.

In [40] the PyRaDiSe package has been developed, which goes in the direction of a tight integration. PyRaDiSe is an open-source Python package which is independent of DL frameworks, and addresses the issue of artifacts caused by 2D reconstruction as it provides a framework for developing auto-segmentation solutions that can directly operate on DICOM data. Authors claim that PyRaDiSe helps to bridge the gap between data science and clinical radiotherapy by facilitating the implementation of deep learning segmentation models in clinical research practice. Actually, PyRaDiSe has the same research objective as IODeep, and the authors leverage DICOM RT Structure Sets (RTSS) to allow data conversion from DICOM to other image formats thus enabling easy auto-segmentation routines. It is well known in the DICOM related literature [41] that DICOM RTSS are a suitable place to store AI related information like the labels for models training, and also the framework we devised for IODeep makes use of DICOM RTSS to store ROI contours. Differently from PyRaDiSe, in IODeep there is no need to code the segmentation solution. IODeep provides the physician with predictions about relevant ROIs in a completely transparent way.

Right now, the information architecture of the DICOM standard does not foresee any type of inclusion of AI, neither as a proper IOD nor in terms of defining a suitable Information Object Module (IOM) to be included in a more general structure. Currently, DICOM is moving towards the integration of AI applications, and a Work Group has been created purposely that is the “WG-23: Artificial Intelligence/Application Hosting”.⁴ The main activity of the WG-23 has been oriented in defining mechanisms for discovering heterogeneous AI services that can expose a suitable manifest for declaring the DICOM services provided to the imaging platform. In the present work we adopt an approach that relies on a direct extension of the DICOM information architecture. This architectural choice derives from the previous experience by some of the authors in developing a framework for adaptive configuration of the PACS viewers’ GUIs based either on the content (i.e. reason for study, modality and body part) of the images to be displayed or on explicit preferences issued by the radiologist. Configuration information is stored in a dedicated IOM which extends the DICOMDIR IOD [42]. We think that the design of new components of the DICOM information architecture, as in the case of IODeep, makes the extension to the standard simpler than defining the interfaces to interact with an external software ecosystem. Moreover, the architectural solution devised




Fig. 1. E-R diagram representing the connection between IODeep and the Dicom Model of the Real World.

by the WG-23 could be prone to problems in a training and/or fine-tuning scenario, when a huge amount of data would have to be moved across the interface to feed the DL model. We developed both a “monolithic” PACS client that implements all the workflow related to the use of IODeep and the service architecture that makes use of IODeep according to the indications of the WG-23. In the next section we detail our implementation, and we compare the two solutions.

4. IODeep architecture

In Fig. 1 we report the E-R diagram describing how IODeep is connected to the DICOM Model of the Real World.⁵ As it can be noticed, the information contained in IODeep is used to instantiate a suitable DNN to predict ROIs on different images. In turn, references to such images are used to generate and store ROI data in the PACS infrastructure.

The overall IODeep structure is reported in Table 1. The operations related to the selection of the proper neural model for the image under investigation, and making the actual prediction, rely on the information contained in four IOMs: “DNN”, “Image Pixel”, “General Study”, and “General Series”. The “DNN” module is a brand new one, which has been defined purposely to store the core info about the architecture, and the weights of the network. Both the “Image Pixel” and the “General Series” module provide information to the network selection algorithm, but some tags in the “Image Pixel” module contain also information about image data formatting to give the DNN input tensor the correct shape. Finally, the “General Study” module is mandatory because the *StudyInstanceUID* has to be filled with a proper value to create the Service Object Pair (SOP) Class aimed at storing/retrieving IODeep instances from the PACS server.

As required by the DICOM WG-23, IODeep is a non-Patient IOD, and it does not contain any patient information. Nevertheless, all the mandatory tags belonging to the “Patient” module (not reported in Table 1) along with the other modules composing IODeep, were included in its structure but they were left empty, as stated by the standard. The *StudyInstanceUID* in the “General Study” module and the *SeriesInstanceUID* in the “General Series” module are a mandatory tag, but they are not useful in the proposed architecture, and we decided to set them to the same value as the *DnnUID* tag.

The “DNN” module contains the information about the network. Apart from the obvious meaning of the *DnnUID*, and *DnnName* tags, both *DnnArchitecture* and *DnnWeights* contain textual information about the layers and the weights of the network respectively. No particular formatting has been indicated in our tag specification, and that’s why we used the Unlimited Text (UT) Value Representation. One can imagine to fill these tags with either a structured text like JSON and XML or a simple URI pointing to a binary file containing network information according to either a Tensorflow or a PyTorch file format. The second option is viable because in our application scenario, the DNN back-end is provided through a computing infrastructure in the same network

⁴ <https://www.dicomstandard.org/activity/wgs/wg-23>.

⁵ https://dicom.nema.org/medical/dicom/current/output/html/part03.html#chapter_7.

Table 1

The general *IODeep* structure reporting both the modules and the tags required for selecting and instantiating a DNN architecture.

| Name | TAG | VR | Values |
|----------------------------------|--------------|----|--|
| DNN module | | | |
| <i>DnnArchitecture</i> | (0017, 00XX) | UT | DNN architecture |
| <i>DnnWeights</i> | (0017, 00X1) | UT | DNN weight |
| <i>DnnName</i> | (0017, 00X2) | PN | e.g. “Brain Tumor segmentation Unet” |
| <i>DnnUID</i> | (0017, 00X3) | UI | UID* |
| Image Pixel module | | | |
| <i>PhotometricInterpretation</i> | (0028, 0004) | CS | RGB, MONOCHROME1... (DICOM standard Section C.7.6.3.1.2) |
| <i>SamplesPerPixel</i> | (0028, 0002) | US | 3 or 1 |
| <i>PatientOrientation</i> | (0020, 0020) | CS | ["P", "F"]; ["L", "P"]; ["L", "F"] ... (DICOM standard Section C.7.6.1.1.1) |
| <i>PlanarConfiguration</i> | (0028, 0006) | US | 0 or 1 |
| General Study module | | | |
| <i>StudyInstanceUID</i> | (0020, 000D) | UI | UID* |
| General Series module | | | |
| <i>SeriesInstanceUID</i> | (0020, 000E) | UI | UID* |
| <i>Modality</i> | (0008, 0060) | CS | CT, MR, PT ... (DICOM standard Section C.7.3.1.1.1) |
| <i>BodyPartExamined</i> | (0018, 0015) | CS | BREAST, ABDOMEN, CHEST ... (DICOM Part 16: Content Mapping Resource) |

*These tags share the same Unique Identifier that is the one defined for the *DnnUID* tag

where the PACS is. One can imagine either a dedicated server or better an instance of the DNN back-end running on a medical workstation because it will be used only in inference mode thus the computational effort is very limited. In both cases, no security concerns arise because the DNN back-end does not need to connect to the Internet, and the network instanced through *IODeep* predicts the ROIs starting from an image that is already open on the PACS viewer. This ensures that no privacy issues arise because the radiologist is authorized to see patient’s data, and the network uses only the values stored in the “Frame Of Reference” and “Image Pixel” modules of the slice.

In our implementation the *DnnArchitecture* tag contains a JSON document whose structure has been developed purposely by some of the authors along with the related parser, which allows to instantiate the same abstract model seamlessly in TensorFlow or PyTorch. The parser executable code is freely available in a suitable repository⁶ linked by the project’s one. The repository contains also the command line specifications to invoke the parser, the complete documentation for writing user-defined JSON descriptions of a network architecture, and the JSON code for building the example U-net [43] used in our implementation. The JSON code is showed also in Fig. 2 along with the U-net architecture. As it can be seen, the key `input_shape` contains the information about the shape of the input tensor, which will be used to check if reshaping is needed for the slice to be processed. The key `architecture` shows the type and sequential order of the layers that make up the neural network, while the keys `convolution`, `kernel`, and so on report the parameters to be set to implement the convolution properly for each layer. In the example a simple sequential model is reported, but our parser deals with all the layer types both in TensorFlow and PyTorch. As an example, our format supports the key `skip_connection` to build U-net models that are typical architectures for medical imaging. A detailed description of the implemented JSON keys, and the executable of the CLI version of the parser are available in the project repository.

In the proposed implementation, we used a PyTorch back-end so the tag *DnnWeights* contains the URI pointing to a file with `.pth` extension that is the binary format used by Torch, storing the weights as a serialized Python object. In TensorFlow we adopted the `tf save` format to allow the implementation of custom layers. We choose to save separately the weights and the architecture to allow simple creation of the network by the user, who does not need to have strong coding skills. When running the parser, either a `torch` or a `tensorflow` environment is set according to the value passed in the `-backend` command line parameter. The environment is used just to instantiate a specific model that will be saved using the path specified in the `-name` command line parameter. In this way, a network developer has no need to instantiate a specific model class inherited from either `torch.nn.Module` or `tensorflow.keras.Model`. The `dcm` file containing the *IODeep* used to instantiate our U-net example network is available in the project repository.

4.1. The ROI prediction workflow

Fig. 3 reports a simple UML sequence diagram showing the workflow of the ROI prediction scenario. We developed a lightweight PACS client to show the effectiveness of this workflow.

All the *IODeep* instances corresponding to networks that can be served by the DNN back-end are within the PACS server, and they can be made available to the physician who wants to perform automatic prediction of the ROIs in a particular slice. To test our implementation, we trained a couple of example Unets on the “Brain Tumor Classification (MRI)” data set⁷ and the “UW-Madison GI Tract Image Segmentation” data set⁸ respectively.

After the study request, and the visualization of the series to be diagnosed, the physician can click on a suitable “AI ROI” button at the viewer GUI while she or he is viewing whatever slice in the series. This

⁶ <https://github.com/CHILab1/DNN-parser>.

⁷ <https://www.kaggle.com/datasets/sartajbhuvaji/brain-tumor-classification-mri>.

⁸ <https://www.kaggle.com/competitions/uw-madison-gi-tract-image-segmentation/overview>.

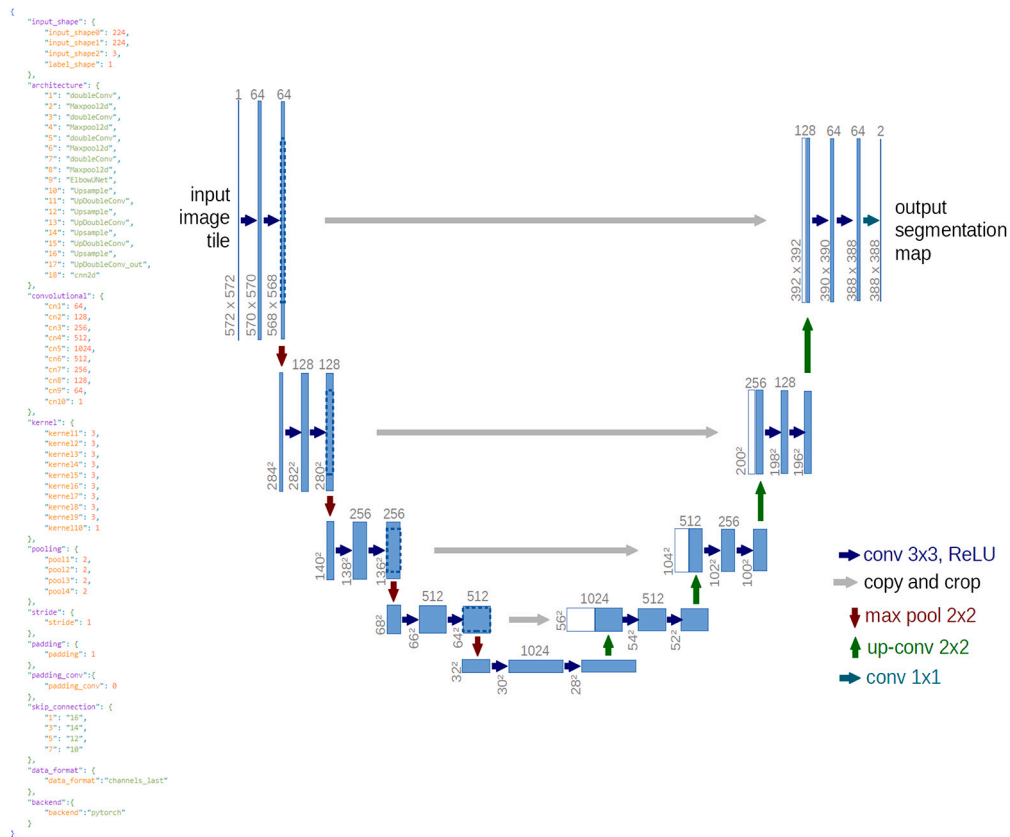


Fig. 2. Representation of the U-net neural network used in our example, and its JSON description.

event will start the actual ROI prediction process. We implemented a thick PACS client that performs both the network selection and the true ROI prediction by instantiating the network on the DNN back-end. In this scenario, the client reads the tags stored in the “General Series” and “Image Pixel” modules inside the Image IOD of the current slice, and it retrieves all the IODeep instances from the server to start network selection.

Specifically, the network selection algorithm will use the *Modality* tag (0008, 0060) which describes the type of data acquisition, the *BodyPartExamined* tag (0018, 0015) which describes the anatomic region of the study using Code Strings (CS) that refer to the standardization reported in the DICOM Part 16: Content Mapping Resource, and the *SamplePerPixel* tag (0028, 0002) which describes the number of channels in the data. In case the *BodyPartExamined* tag is empty, the content of the *StudyDescription* tag (0008, 1030) is used. The matching algorithm for selecting the neural network is shown below in Algorithm 1.

In case of successful match, the SOP UID of the IODeep is returned by the algorithm. If no IODeep instance satisfies requirements for network prediction, a warning message will be returned blocking the work pipeline. The returned SOP UID will be used to query the PACS server for obtaining both the network architecture and the URI pointing to the weights file. In turn, the JSON document in the *DnnArchitecture* tag will be parsed to create and save the network instance using the DNN back-end, which is selected according to the value of the *backend* key. The parser itself is invoked by the PACS client using a system call. After the input tensor shape check, actual prediction will be made. The entire ROI prediction workflow is reported in the Algorithm 2.

Before the actual creation of the network instance in the DNN back-end, the shape of the current slice is checked to assess if it is compatible with the network input tensor. The tags *SamplePerPixel* (0028, 0002), *Rows* (0028, 0010), and *Columns* (0028, 0011) are used to check if the dimensions of the image are compatible with the neural network input

or if reshape is required. Moreover, The tag *PhotometricInterpretation* (0028, 0004), is used to verify the interpretation of the given pixels (e.g. MONOCHROME2, MONOCHROME1, RGB, and so on). Finally, the tag *PixelRepresentation* (0028, 0103) describes the pixel representation (signed or unsigned) and will be the last tag checked for the final formatting of the image to be tested. This process is summarized by the `check_tensor_shape()` call in the trace of Algorithm 2.

Once the shape of the image has been checked w.r.t. the expected network input tensor, the actual model is created in the DNN back-end starting from the object parsed from the JSON document in the *DnnArchitecture* tag, and the corresponding weights are loaded. The actual ROI prediction is a very fast process, which will output a list containing all the poly-lines that have been predicted.

Each ROI will be displayed by the viewer, giving the physician the option of either validating or rejecting the region. Fig. 4 shows how the ROI can be validated in our simple viewer that was implemented using the Python library PyQt. The viewer interfaces with the PACS server, and returns to the user a list of the available studies. Once a study has been selected, a list of that study’s series is retrieved in order to select the slice to be viewed. The display will show two information blocks in the upper section. The left block displays patient’s information: *PatientName* (0010, 0010), *PatientID*(0010, 0020), *PatientBirthdate* (0010, 0030), *PatientSex* (0010, 0040). The second block displays clinical information: *AccessionNumber*(0008, 0050), *InstitutionName*(0008, 0080), *ReferringPhysicianName* (0008, 0090), *StudyDate* (0008, 0020), *StudyDescription* (0008, 1030), *StudyID* (0020, 0010), *StudyInstanceUID* (0020, 000D), and *StudyTime* (0008, 0030). The central body of the viewer will display the image. In this part of the layout there are the “autoplay” and the navigation buttons apart from the “AI ROI” button, already mentioned. The interaction between the viewer and the DICOM standard

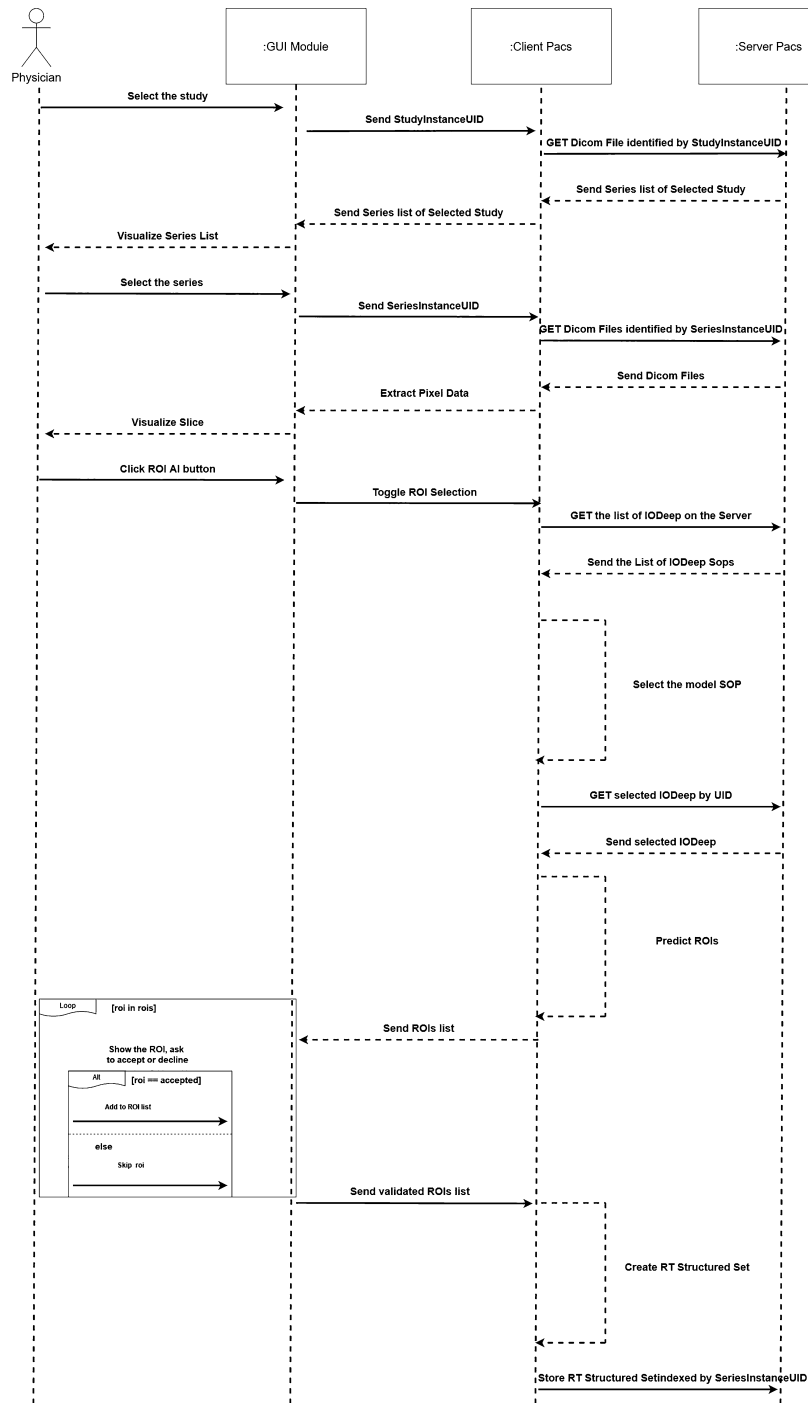


Fig. 3. Sequence diagram of the ROI prediction scenario.

was developed using the Pydicom library, while the server was an instance of the Orthanc DICOM server.⁹

The ROIs validated by the physician will be stored within a DICOM RT Structure Set, a data structure used to store all information that is collected through virtual instruments/workstations. In this data structure, all the patient information contained in the source IOD will be kept. Moreover, both the “Frame of Reference” and the “ROI Contour” module will be added to instance the *FrameReferenceUID* (0020, 0052) and the *ROIContourSequence* (3006, 0039) tag respectively, which contain the slice reference and the actual ROI data. Finally, the validation

information will be stored through the “Approval” module that will contain the tags:

- *ApprovalStatus* (300E, 0002)
- *ReviewDate* (300E, 0004)
- *ReviewTime* (300E, 0005)
- *ReviewerName* (300E, 0008)

The RT Structure Set will be uploaded to the PACS server to be reused should the physician wish to retrieve ROI information. The DICOM file containing the RT Structure Set generated from the ROI displayed in Fig. 4 is available in the project repository. As already pointed out in section 2, even if we used an example U-net trained at predicting

⁹ <https://www.orthanc-server.com/>.




Fig. 4. The viewer interface for ROI validation. (a) Predicted ROIs are displayed and outlined in red. (b) Validated ROIs are outlined in green.

Algorithm 1 Network selection algorithm.

Require: *SliceTagList*: list(*Modality*, *SamplePerPixel*, *BodyPartExamined*, *StudyDescription*) ▷ The list of relevant tags from the current slice

Require: *IODeepList* ▷ The list of all the IODeep instances retrieved from the PACS server

```

1: if SliceTagList.BodyPartExamined == ∅ then
2:   SliceTagList.remove(BodyPartExamined)
3: else
4:   SliceTagList.remove(StudyDescription)
5: end if
6: for net in IODeepList do
7:   IODeepTagList ← list(net.Modality, net.SamplePerPixel,
                        net.BodyPartExamined)
8:   control ← 0
9:   for slice_tag, net_tag in (SliceTagList, IODeepTagList) do
10:    if slice_tag == net_tag and not SliceTagList.isLast(slice_tag) then
11:      control ← control + 1
12:    end if
13:    if SliceTagList.isLast(slice_tag) and (slice_tag == net_tag or
      slice_tag.substring(net_tag) == True) then
14:      control ← control + 1 ▷ When using StudyDescription a pattern
      search is used in place of strict equality
      on the last tag
15:    end if
16:  end for
17:  if control == 3 then
18:    net_sop_uid ← net.DnnUID ▷ the SOP UID of the selected IODeep
    instance
19:    return net_sop_uid
20:  else
21:    net_sop_uid ← None
22:  end if
23: end for
24: return net_sop_uid

```

ROIs, IODeep can be used equally to instantiate both networks that deal with ROIs and models trained at devising bounding boxes. In any case the DICOM RT Structure Set will store the vertex list used to render the poly-line contour in the PACS viewer.

4.2. Comparison with the DICOM WG-23 proposal

People in the DICOM community are working actively towards the integration of AI algorithms in the standard, and a suitable Working Group has been founded in this respect: the WG-23 “Artificial Intelligence and Application Hosting”. The current proposal of the WG-23 has

Algorithm 2 ROI prediction algorithm.

Require: *IPModule* ▷ The ImagePixel module of the current slice

Require: *FRModule* ▷ The FrameOfReference module of the current slice

Require: *IODeep* ▷ The selected DNN

Require: *Backend* ▷ The reference to the DNN back-end

```

1: rois ← list()
2: network ← parser(IODeep.DnnArchitecture)
3: check_tensor_shape(IPModule, network.input_shape) ▷ tensor shape
   analysis
4: model ← Backend.createNetwork(network) ▷ instance of the network in
   the DNN back-end
5: model.load(IODeep.DnnWeights) ▷ loading method of the DNN back-end
6: preds ← model.predict(IPModule.pixelData) ▷ prediction method of the
   DNN back-end
7: for roi in preds do:
8:   slice_id ← FRModule.FrameOfReferenceUID
9:   rois.add(dict("sliceID" : slice_id, "polyline" : roi) ▷ These value will
   populate the RT Structure Set containing the ROIs
10: end for
11: return rois

```

been submitted in the form of *work item*¹⁰ that is the standard way for working groups to publish their results.

The core of the WG-23 work item is to pursue AI integration using a mechanism of (micro-)service discovery according to the model proposed by the Open Application Model (OAM).¹¹ In this model, AI applications are deployed as services running on heterogeneous computing infrastructures that can be either on the same network of the PACS or in an external cloud. Each application owns a “Manifest” in the form of a YAML file, which encapsulates the definition of a workload and the information needed to run it. OAM defines also an “application operator” that is the infrastructure that actually deploys the application, and defines more information from the operational point of view to allow the so called “Platform” (the PACS infrastructure using the services, in our case) to integrate them explicitly in its workflow. In the vision of the WG-23, applications can either register to the Platform or they can be detected through a service discovery mechanism.

Our proposal moves from a monolithic implementation perspective where information is fully specified according to the DICOM Information Architecture, and only a TCP/IP socket connection is needed to reach the DNN back-end that runs on the same network of the PACS. Apart from the privacy and security concerns already mentioned, we adopted this solution for efficiency reasons. The services ecosystem proposed by the WG-23 is a scalable solution with very low implementation

¹⁰ <https://www.dicomstandard.org/workitems>.

¹¹ <https://oam.dev/>.

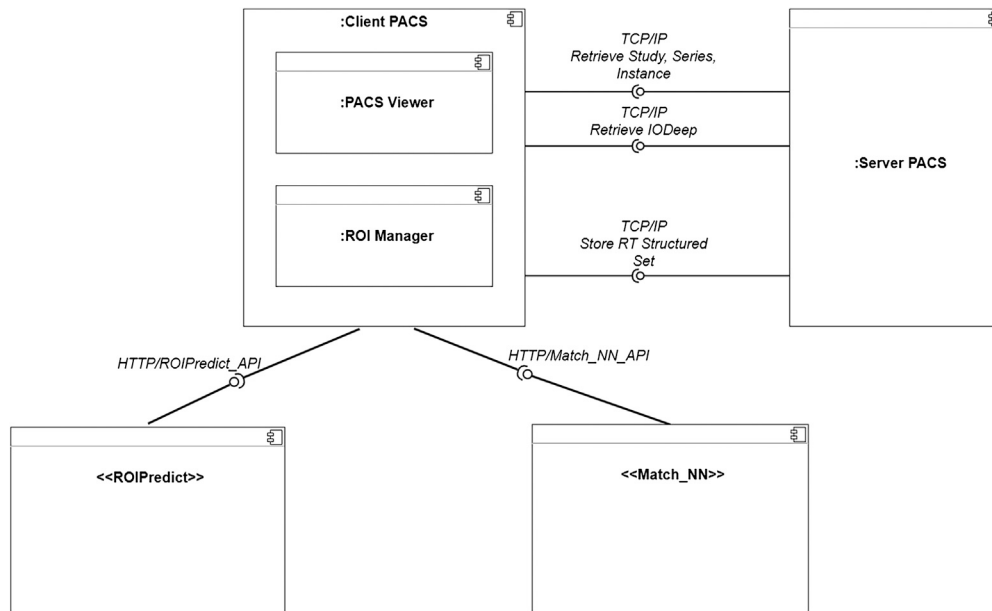


Fig. 5. Component diagram representing the service implementation of the proposed framework.

effort on the PACS side, but it is tailored on an inference scenario, and it is not well suited to perform training or fine-tuning. In general, Radiology wards are focused on particular diseases due to their location in the territory, and the presence of precise medical specializations. As a consequence, fine tuning a trained network by means of the peculiar data generated in a particular PACS infrastructure, appears a reasonable extension of using DNNs in Medical Imaging. Making a data set available at the DNN back-end for training a model is a bandwidth consuming task. A service implementation is not well suited for training and/or fine tuning because, in general, a service infrastructure is intended to be hosted in a remote cloud with respect to the applications that require a service. In the model training scenario, this leads to both high latency and privacy issues. On the other hand, our solution requires only a data transfer on the internal network of the ward, and this is not so resource consuming. IODeep can be used explicitly also in this context: a PACS client devoted to back-end administration could select a data set made by several slice series, and start a network selection procedure, that is the same we described before, followed by the training in the back-end.

Nevertheless, we rearranged our implementation according to the indications of the WG-23 work item. The software architecture of this solution is reported in Fig. 5. In this scenario, Algorithm 1 has been redesigned to provide a service discovery REST API (Application Programming Interface) and it is available at a known endpoint on the service infrastructure. Such an URI is the only thing the PACS platform has to know, and IODeep stores the information used to call the API at the service discovery endpoint. This interaction provides the platform with the endpoint of the selected DNN service. The PACS platform calls the DNN service API at the service endpoint passing the slice, and the information for checking the tensor shape stored in the selected IODeep. Finally, the selected DNN service runs the Algorithm 2. This service architecture does not need to store network information in the IODeep architecture as we can assume that all the networks are already instantiated in the DNN back-end. In our implementation the PACS server, the client, and each REST service run in separate Docker containers that communicate through using a virtual LAN defined through the YAML configuration file. The actual containers can be deployed on either a single machine or a cluster. We choose explicitly this kind of infrastructure because we strongly believe that effective extension of the service architecture to the training/fine-tuning scenario can be achieved only in a LAN context where the concerns related to the connection bandwidth are negligible.

5. Conclusion

In this paper, we introduced IODeep, a new DICOM Information Object Definition aimed at integrating all the information needed to allow a pre-trained DNN model to be selected w.r.t. the features of the slice under investigation, and to be instantiated in a DNN back-end running in the same network of the PACS server to perform inference about ROIs in that slice. We presented the detailed information architecture of IODeep along with a precise workflow for selecting the correct DNN, instancing the abstract model description in the back-end, making predictions, and storing the ROIs validated by the physician in a RT Structure Set for further reuse. Finally, we implemented a suitable PACS client to show the effectiveness of our proposal.

We strongly believe that the key for using Artificial Intelligence in the everyday diagnostic activity for medical images is related to standardization. Our approach provides a complete answer to the first statement in Section 1, and it can be reviewed as a technological starting point for answering the other two statements.

DNNs are useful if they can be fine-tuned on the real data sets available at the Radiology wards, and our solution allows a simple extension to deal with this scenario. In fact, the same information architecture in IODeep can be used to select a model to be retrained provided that the bandwidth concerns to move data from the PACS to the DNN back-end are easy to solve. In this respect, a solution that runs entirely in the same LAN is strongly preferable to a services ecosystem running in cloud.

We also showed that IODeep can be used effectively to implement a service based solution for AI integration in DICOM, as the one devised by the WG-23. In fact, our network selection and ROI prediction algorithms can be implemented as a service discovery procedure, and DNN service respectively to pursue this goal.

Many DNNs for Medical Imaging work on multi-modal data: this is the case of registered CT-PET volumes, or the very recent generative AIs that integrate both textual and visual information. The current implementation of IODeep does not support multi-modal input data, but this is not a strong limitation for two reasons. First, the general two-stage approach based on both DNN selection and ROI prediction does not depend on the particular input data. Second, the structure of IODeep in Table 1 can be easily extended to utilize multi-modal and/or textual data. In this respect, one can devise a “Study IODeep” that is the proper IOD containing the tags related to both the architecture and the

weights of the DNN. Moreover, different “Series IODeep” modules will contain the information related to each kind of input tensor, which can be either a visual or a textual one. Each “Series IODeep” module will reference the *DnnUID* of the “Study IODeep” to allow proper retrieval of all the information needed to select the multi-modal network, and to check the shape of all the input tensors. We are now working actively on this issue.

Finally, we want to point out that IODeep is also a technological key for explainability of DNN models, and their acceptance in the radiological community. First of all, we claim the need to support the diagnostic process just with a “ROI suggestion” task. A model suggesting a loose ROI does not make a precise pixel level segmentation. In fact, validating an exact segmentation represents a huge cognitive load for the doctor. Moreover, precise ROI prediction can be achieved only with extensive training on very huge amounts of data. Such a goal can be pursued using Federated Learning (FL) that is a distributed learning paradigm where the model is trained while “moving” it across the data repositories where the data set is split, but data are always kept where they reside without moving them at all. The technology proposed in this paper enables a distributed scenario where IODeeps are shared between the PACS servers of various public and private hospital institutions without violating any rule related to either data security or patients’ privacy, while allowing FL through the data stored within the individual PACS servers.

CRedit authorship contribution statement

Salvatore Contino: Writing – original draft, Software, Methodology, Investigation, Data curation. **Luca Cruciata:** Writing – original draft, Software, Methodology, Investigation, Formal analysis, Data curation. **Orazio Gambino:** Writing – original draft, Conceptualization. **Roberto Pirrone:** Writing – review & editing, Supervision, Project administration, Formal analysis, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

This work has been partially supported by project:

1. “SAMOTHRACE” (Sicilian MicronanoTech Research And Innovation Center), cup project B73C22000810001, project code ECS_00000022.
2. PON “Ricerca e Innovazione” 2014-2020, Asse IV “Istruzione e ricerca per il recupero” con riferimento all’Azione IV.4 “Dottorati e contratti di ricerca su tematiche dell’innovazione” e all’Azione IV.5 “Dottorati su tematiche green”, DM 1061/2021.

References

- [1] S. Waite, S. Kolla, J. Jeudy, A. Legasto, S.L. Macknik, S. Martinez-Conde, E.A. Krupinski, D.L. Reede, Tired in the reading room: the influence of fatigue in radiology, *J. Am. Coll. Radiol.* 14 (2) (2017) 191–197, <https://doi.org/10.1016/j.jacr.2016.10.009>, <https://www.sciencedirect.com/science/article/pii/S1546144016310766>.
- [2] N. Savage, How AI is improving cancer diagnostics, *Nature* 579 (7800) (2020) S14–S16, <https://doi.org/10.1038/d41586-020-00847-2>.
- [3] L. Rundo, R. Pirrone, S. Vitabile, E. Sala, O. Gambino, Recent advances of HCI in decision-making tasks for optimized clinical workflows and precision medicine, *J. Biomed. Inform.* 108 (2020) 103479, <https://doi.org/10.1016/j.jbi.2020.103479>.
- [4] A. Ukharov, I. Shlivko, I. Klemenova, O. Garanina, K. Uskova, A. Mironycheva, Y. Stepanova, Skin cancer risk self-assessment using AI as a mass screening tool, *Inform. Med. Unlocked* 38 (2023) 101223, <https://doi.org/10.1016/j.imu.2023.101223>, <https://www.sciencedirect.com/science/article/pii/S2352914823000655>.
- [5] Z. Zhou, M.M.R. Siddiquee, N. Tajbakhsh, J. Liang, U-Net++: a nested U-Net architecture for medical image segmentation, in: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 839–848.
- [6] J. Schlemper, O. Oktay, M. Schaap, M. Heinrich, B. Kainz, B. Glocker, Attention U-Net: learning where to look for the pancreas, *Med. Image Anal.* 52 (2019) 201–211.
- [7] J. Fu, J. Liu, H. Tian, Y. Li, Y. Bao, Z. Fang, Dual attention network for scene segmentation, in: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 3146–3154.
- [8] Q. Dou, H. Chen, L. Yu, J. Qin, P.-A. Heng, 3D anisotropic hybrid network: translating convolutional neural networks for 3D medical image segmentation, *IEEE Trans. Med. Imaging* 39 (5) (2020) 1462–1473.
- [9] F. Milletari, N. Navab, S.-A. Ahmadi, V-Net family: convolutional networks for volumetric medical image segmentation, *IEEE Trans. Med. Imaging* 40 (5) (2021) 1256–1268.
- [10] P.F. Christ, F. Ettliger, F. Grün, M.E.A. Elshaera, J. Lipkova, S. Schlecht, F. Ahmaddy, S. Tatavarty, M. Bickel, P. Bilic, M. Rempfler, F. Hofmann, M.D. Anastasi, S.-A. Ahmadi, G. Kaissis, J. Holch, W. Sommer, R. Braren, V. Heinemann, B. Menze, Automatic liver and tumor segmentation of CT and MRI volumes using cascaded fully convolutional neural networks, [arXiv:1702.05970](https://arxiv.org/abs/1702.05970), 2017.
- [11] Z. Akkus, A. Galimzianova, A. Hoogi, D.L. Rubin, B.J. Erickson, Deep learning for brain MRI segmentation: state of the art and future directions, *J. Digit. Imag.* 30 (4) (2017) 449–459, <https://doi.org/10.1007/s10278-017-9983-4>.
- [12] Y. Qu, X. Li, Z. Yan, L. Zhao, L. Zhang, C. Liu, S. Xie, K. Li, D. Metaxas, W. Wu, Y. Hao, K. Dai, S. Zhang, X. Tao, S. Ai, Surgical planning of pelvic tumor using multi-view CNN with relation-context representation learning, *Med. Image Anal.* 69 (2021) 101954, <https://doi.org/10.1016/j.media.2020.101954>.
- [13] Y. Wang, Y. Zhou, W. Shen, S. Park, E.K. Fishman, A.L. Yuille, Abdominal multi-organ segmentation with organ-attention networks and statistical fusion, *Med. Image Anal.* 55 (2019) 88–102, <https://doi.org/10.1016/j.media.2019.04.005>.
- [14] Y. Yan, D. Zhang, Multi-scale U-like network with attention mechanism for automatic pancreas segmentation, *PLoS ONE* 16 (5) (2021) e0252287, <https://doi.org/10.1371/journal.pone.0252287>.
- [15] H.R. Roth, L. Lu, A. Farag, H.-C. Shin, J. Liu, E.B. Turkbey, R.M. Summers, Deeporgan: multi-level deep convolutional networks for automated pancreas segmentation, in: *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part I 18*, Springer, 2015, pp. 556–564.
- [16] Q. Zhou, Z. Huang, M. Ding, X. Zhang, Medical image classification using lightweight CNN with spiking cortical model based attention module, *IEEE J. Biomed. Health Inform.* 27 (4) (2023) 1991–2002, <https://doi.org/10.1109/JBHI.2023.3241439>.
- [17] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, L. Kaiser, I. Polosukhin, Attention is all you need, <https://proceedings.neurips.cc/paper-files/paper/2017/file/3f5ee243547dee91fdb053c1c4a845aa-Paper.pdf>, 2017.
- [18] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, N. Houlsby, An image is worth 16x16 words: transformers for image recognition at scale, [arXiv:2010.11929](https://arxiv.org/abs/2010.11929) [cs], Jun 2021, <http://arxiv.org/abs/2010.11929>.
- [19] N.S. Punn, S. Agarwal, Modality specific U-Net variants for biomedical image segmentation: a survey, *Artif. Intell. Rev.* 55 (7) (2022) 5845–5889, <https://doi.org/10.1007/s10462-022-10152-1>.
- [20] J. Zhang, Y. Zhang, Y. Jin, J. Xu, X. Xu, MDU-Net: multi-scale densely connected U-Net for biomedical image segmentation, *Health Inf. Syst. Syst.* 11 (1) (2023) 13, <https://doi.org/10.1007/s13755-022-00204-9>.
- [21] H. Song, Y. Wang, S. Zeng, X. Guo, Z. Li, OAU-Net: outlined attention U-net for biomedical image segmentation, *Biomed. Signal Process. Control* 79 (2023) 104038, <https://doi.org/10.1016/j.bspc.2022.104038>.
- [22] D. Zhang, J. Han, G. Cheng, M. Yang, Weakly supervised object localization and detection: a survey, *CoRR*, [arXiv:2104.07918](https://arxiv.org/abs/2104.07918) [abs], 2021, <https://arxiv.org/abs/2104.07918>.
- [23] W. Shen, Z. Peng, X. Wang, H. Wang, J. Cen, D. Jiang, L. Xie, X. Yang, Q. Tian, A survey on label-efficient deep image segmentation: bridging the gap between weak supervision and dense prediction, [arXiv:2207.01223](https://arxiv.org/abs/2207.01223), 2023.
- [24] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, A. Torralba, Learning deep features for discriminative localization, in: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE Computer Society, Los Alamitos, CA, USA, 2016, pp. 2921–2929, <https://doi.ieeecomputersociety.org/10.1109/CVPR.2016.319>.
- [25] J. Lin, G. Han, X. Xu, C. Liang, T.-T. Wong, C.L.P. Chen, Z. Liu, C. Han, Broad-CAM: outcome-agnostic class activation mapping for small-scale weakly supervised applications, [arXiv:2309.03509](https://arxiv.org/abs/2309.03509) [cs], Sep. 2023, <http://arxiv.org/abs/2309.03509>.
- [26] W. Jorritsma, F. Cnossen, P. van Ooijen, Adaptive support for user interface customization: a study in radiology, *Int. J. Hum.-Comput. Stud.* 77 (2015) 1–9, <https://doi.org/10.1016/j.ijhcs.2014.12.008>.
- [27] A. Fedorov, W.J. Longabaugh, D. Pot, D.A. Clunie, S.D. Pieper, D.L. Gibbs, C. Bridge, M.D. Hermann, A. Homeyer, R. Lewis, et al., National cancer institute imaging data commons: toward transparency, reproducibility, and scalability in imaging artificial intelligence, *Radiographics* 43 (12) (2023) e230180.
- [28] E. Darzidehkalani, M. Ghasemi-Rad, P. van Ooijen, Federated learning in medical imaging: part I: toward multicenter health care ecosystems, *J. Am. Coll. Radiol.* 19 (8) (2022) 969–974.
- [29] H. Guan, P.-T. Yap, A. Bozoki, M. Liu, Federated learning for medical image analysis: a survey, [arXiv:2306.05980](https://arxiv.org/abs/2306.05980), 2023.

- [30] B.H. Van der Velden, H.J. Kuijff, K.G. Gilhuijs, M.A. Viergever, Explainable artificial intelligence (XAI) in deep learning-based medical image analysis, *Med. Image Anal.* 79 (2022) 102470.
- [31] R. Jesus, L. Bastião Silva, V. Sousa, L. Carvalho, D. Garcia Gonzalez, J. Carias, C. Costa, Personalizable AI platform for universal access to research and diagnosis in digital pathology, *Comput. Methods Programs Biomed.* 242 (2023) 107787, <https://doi.org/10.1016/j.cmpb.2023.107787>.
- [32] N. Lajara, J.L. Espinosa-Aranda, O. Deniz, G. Bueno, Optimum web viewer application for DICOM whole slide image visualization in anatomical pathology, *Comput. Methods Programs Biomed.* 179 (2019) 104983, <https://doi.org/10.1016/j.cmpb.2019.104983>.
- [33] M.D. Herrmann, D.A. Clunie, A. Fedorov, S.W. Doyle, S. Pieper, V. Klepeis, L.P. Le, G.L. Mutter, D.S. Milstone, T.J. Schultz, R. Kikinis, G.K. Kotecha, D.H. Hwang, K.P. Andriole, A.J. lafrate, J.A. Brink, G.W. Boland, K.J. Dreyer, M. Michalski, J.A. Golden, D.N. Louis, J.K. Lennerz, Implementing the DICOM standard for digital pathology, *J. Pathol. Inform.* 9 (1) (2018) 37, https://doi.org/10.4103/jpi.jpi.42_18, <https://www.sciencedirect.com/science/article/pii/S2153353922003509>.
- [34] C. Jansen, B. Lindequist, K. Strohmenger, D. Romberg, T. Küster, N. Weiss, M. Franz, L.O. Schwen, T. Evans, A. Homeyer, N. Zerbe, The vendor-agnostic EMPAIA platform for integrating AI applications into digital pathology infrastructures, *Future Gener. Comput. Syst.* 140 (C) (2023) 209–224, <https://doi.org/10.1016/j.future.2022.10.025>.
- [35] E. Dikici, M. Bigelow, L.M. Prevedello, R.D. White, B.S. Erdal, Integrating AI into radiology workflow: levels of research, production, and feedback maturity, *J. Med. Imag.* 7 (01) (2020) 1, <https://doi.org/10.1117/1.jmi.7.1.016502>.
- [36] P. Kathiravelu, P. Sharma, A. Sharma, I. Banerjee, H. Trivedi, S. Purkayastha, P. Sinha, A. Cadrin-Chenevert, N. Safdar, J.W. Gichoya, A DICOM framework for machine learning and processing pipelines against real-time radiology images, *J. Digit. Imag.* 34 (4) (2021) 1005–1013, <https://doi.org/10.1007/s10278-021-00491-w>.
- [37] K. Osorno-Castillo, R.D. Fonnegra, G.M. Díaz, Integration of machine learning models in PACS systems to support diagnostic in radiology services, in: J.C. Figueroa-García, F.S. Garay-Rairán, G.J. Hernández-Pérez, Y. Díaz-Gutierrez (Eds.), *Applied Computer Sciences in Engineering*, Springer International Publishing, Cham, 2020, pp. 233–244.
- [38] H.H. Pham, D.V. Do, H.Q. Nguyen, DICOM imaging router: an open deep learning framework for classification of body parts from DICOM X-ray scans, arXiv:2108.06490, 2021.
- [39] M.P. Recht, M. Dewey, K. Dreyer, C. Langlotz, W. Niessen, B. Prainsack, J.J. Smith, Integrating artificial intelligence into the clinical practice of radiology: challenges and recommendations, *Eur. Radiol.* 30 (6) (2020) 3576–3584, <https://doi.org/10.1007/s00330-020-06672-5>.
- [40] E. Rüfenacht, A. Kamath, Y. Suter, R. Poel, E. Ermiş, S. Scheib, M. Reyes, PyRaDiSe: a Python package for DICOM-RT-based auto-segmentation pipeline construction and DICOM-RT data conversion, *Comput. Methods Programs Biomed.* 231 (2023) 107374, <https://doi.org/10.1016/j.cmpb.2023.107374>.
- [41] D.A. Clunie, DICOM Structured Reporting, PixelMed Publishing, 2000.
- [42] O. Gambino, L. Rundo, V. Cannella, R. Pirrone, S. Vitabile, A framework for data-driven adaptive GUI generation based on DICOM, *J. Biomed. Inform.* 88 (2018) 37–52, <https://doi.org/10.1016/j.jbi.2018.10.009>.
- [43] O. Ronneberger, P. Fischer, T. Brox, U-Net: convolutional networks for biomedical image segmentation, in: N. Navab, J. Hornegger, W.M. Wells, A.F. Frangi (Eds.), *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, in: *Lecture Notes in Computer Science*, Springer International Publishing, Cham, 2015, pp. 234–241.