

Deep Neural Networks in Medical Imaging: Privacy Preservation, Image Generation and Applications

Diana Ioana Stoian ^{1,2}, Horia Andrei Leonte ^{1,2}, Anamaria Vizitiu ^{1,2}, Constantin Suciuc ^{1,2} and Lucian Mihai Itu ^{1,2,*}

¹ Advanta, Siemens SRL, 15 Noiembrie Bvd, 500097 Brasov, Romania; diana.stoian@siemens.com (D.I.S.); horia.leonte.ext@siemens.com (H.A.L.); anamaria.vizitiu@siemens.com (A.V.); constantin.suciu@siemens.com (C.S.)

² Automation and Information Technology, Transilvania University of Brasov, Mihai Viteazu nr. 5, 5000174 Brasov, Romania

* Correspondence: lucian.itu@unitbv.ro

1. Introduction

Medical Imaging plays a key role in disease management, starting from baseline risk assessment, diagnosis, staging, therapy planning, therapy delivery, and follow-up. Each type of disease has led to the development of more advanced imaging methods and modalities to help clinicians address the specific challenges in analyzing the underlying disease mechanisms. Imaging data is one of the most important sources of evidence for clinical analysis and medical intervention as it accounts for about 90% of all healthcare data. Researchers have been actively pursuing the development of advanced image analysis algorithms, some of which are routinely used in clinical practice. These developments were driven by the need for a comprehensive quantification of structure and function across several imaging modalities such as Computed Tomography (CT), X-ray Radiography, Magnetic Resonance Imaging (MRI), Ultrasound, Nuclear Medicine Imaging, and Digital Pathology [1].

In the context of the availability of unprecedented data storage capacity and computational power, Deep learning has become the state-of-the-art machine learning technique, providing unprecedented performance at learning patterns in medical images and great promise for helping physicians during clinical decision-making processes. Previously reported deep learning-related studies cover various types of problems, e.g., classification, detection, segmentation, for different types of structures, e.g., landmarks, lesions, organs, in diverse anatomical application areas [2].

The aim of this special issue is to present and highlight novel methods, architectures, techniques, and applications of deep learning in medical imaging. Papers both from theoretical and practical aspects were welcome, including ongoing research projects, experimental results, and recent developments related to, but not limited to, the following topics: image reconstruction; image enhancement; segmentation; registration; computer aided detection; landmark detection; image or view recognition; automated report generation; multi-task learning; transfer learning; generative learning; self-supervised learning; semi-supervised learning; weakly supervised learning; unsupervised learning; federated learning; privacy preserving learning; explainability and interpretability; robustness and out-of-distribution detection; and uncertainty quantification.

2. The Papers

In this Special Issue, we published a total of 14 papers that span across four interesting topics as outlined below.

2.1. Privacy-Preserving Learning

Deep Learning heavily relies on existing and forthcoming patient data to yield precise and dependable outcomes within the realm of healthcare applications. Despite the



Citation: Stoian, D.I.; Leonte, H.A.; Vizitiu, A.; Suciuc, C.; Itu, L.M. Deep Neural Networks in Medical Imaging: Privacy Preservation, Image Generation and Applications. *Appl. Sci.* **2023**, *13*, 11668. <https://doi.org/10.3390/app132111668>

Received: 13 October 2023

Accepted: 23 October 2023

Published: 25 October 2023



Copyright: © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

copiousness of biomedical data, its dissemination and retrieval are hindered by ethical limitations, particularly concerning safeguarded health-related information pertaining to patients. Consequently, the actualization of medical AI systems encounters challenges, as the requisite data for their development and training are ensnared within the confines of hospital security protocols. To engender resilient algorithms, the databases employed for training, validation, and testing must encompass the complete spectrum of pathological deviations and permutations. Additionally, it is imperative to leverage the entirety of accessible information to formulate a more tailored solution. In instances where training datasets lack heterogeneity, there is a propensity for algorithms to exhibit partiality or inclination towards specific patient profiles [2].

In the field of privacy-preserving learning three papers are presented.

The first paper presents a novel approach for protecting sensitive data in the healthcare setting, as nowadays the trend is to send it outside the facility for it to be processed, trained on etc. [3]. The researchers designed the solution to make it robust against both human perception, as well as against different software attacks such as AI reconstruction attempts. Therefore, by using the proposed pipeline, the data will be obfuscated before leaving the healthcare facility, and the external processing (such as training of AI models) can be performed with a satisfactory privacy-accuracy trade-off, i.e., without a significant drop in accuracy. The three main objectives of the paper are to hide the content from any person viewing the images, to make it difficult for an AI to reconstruct the image, as well as to facilitate AI model training on such data. Regarding the technical aspects, a Variational Autoencoder (VAE) is used, trained on around 30,000 images from the Medical MNIST dataset. The VAE has two different output channels, one based on the mean of the normal distribution, which offers more privacy, but limits the performance in further ML applications, and a second one which is based on the standard deviation of the normal distribution, which trades privacy for training performance.

The second paper's focus lies on maintaining and improving the training and prediction accuracy of AI based solutions in heart disease diagnosis, while overcoming data privacy issues [4]. The main mechanism is federated learning, which aims to keep the data on a single device while training via a collaborative system of a shared model. The optimizer framework—the Modified Artificial Bee Colony (M-ABC) has been chosen due to its flexibility and user-friendliness, has less parameters than other algorithms and has a fast convergence rate. These two methods work together, as the federated matched averaging (FedMA) is constructing a privacy-aware framework for a global cloud model, and the M-ABC framework serves as the feature selector on the client's side. The pipeline was trained on the heart disease dataset of UCI Cleveland, with 303 records and 76 attributes.

In the third paper, the same approach is used, but instead of employing the M-ABC optimization, a hybrid M-ABC with Support Vector Machine is used [5]. The SVM is less prone to overfitting, works well in high dimensional spaces, and has good handling of non-linear data. Moreover, it is a suitable candidate for classifying multiple classes. The M-ABC acts as the feature extractor and the SVM as the classification algorithm. The dataset combines over eleven common features, such as blood pressure, cholesterol serum, sex, age etc. from the datasets of Cleveland, Stalog, Hungary, Long Beach, and Switzerland. As far as results are concerned, the proposed solution is both more efficient and more accurate than the previous ones.

2.2. Image Generation

Within the domain of medical imaging, generative models pursue two primary trajectories: (i) transformation from noise to image and (ii) transition between images. The former encompasses methodologies focused on artificially augmenting the dataset, often referred to as augmentation, by training a deep learning-grounded model to transmute a noise vector into authentic-looking images. To illustrate, when confronted with an extensive array of breast imaging data, such as mammograms, a generative model endeavors to fabricate novel images resembling constituents of the imaging dataset. In the context of

image-to-image generation, generative models undertake the conversion of a given input image into an alternative representation. A diverse array of challenges within the healthcare domain can be effectively addressed through image-to-image generative models. These encompass tasks spanning the enhancement of image fidelity via denoising, amplification of resolution, image inpainting, amalgamation of multi-modal images, along with image reconstruction and alignment.

In the field of reconstruction and image generation four papers are presented. The first one suggests employing an end-to-end deep learning network for the correction of metal artifacts in CT images [6]. The input of the network is represented by metal-affected NMAR sinograms, and the outputs are artifact-free reconstructed images. The architecture consists of three parts: sinogram refinement used to filter the sinogram, back projection used to reconstruct the image into the image domain, and image refinement used to further refine the reconstruction. All parts are trained simultaneously and furthermore, the network performs the complete CT image reconstruction, and does not require a predefined back projection operator or the exact X-ray beam geometry.

A second paper in this area proposes a method to increase the field of view of intraoperative images obtained from Computer Tomographs [7]. This method is used as a prior step to the registration of two volumes: thin intraoperative volume and preoperative volume. The method consists in extrapolating the thin volume by generating additional slices from the existing ones using a GAN architecture. By enhancing the context information required for the matching process, the results appear to be comparable to those obtained after aligning two high-resolution images having the same field of view.

The third paper presents a transfer learning enhanced GAN technique for image reconstruction using under-sampled MR data [8]. The model was tested on an open-source knee dataset, and a private brain dataset with two different acceleration factors: 2 and 4. Both datasets were divided into training and test sets. The training sets were used for finetuning the model after transfer learning. The results indicate that the proposed model outperforms the other reconstruction techniques for both acceleration factors, suggesting that, by using transfer learning, the variation in image contrast, acceleration factor and anatomy between training and test dataset is smaller. Moreover, the distribution of the reconstructed images, produced by transfer learning is more similar to the distribution of the completely sampled image.

The main objective of the fourth paper is OCT image enhancement through denoising and deblurring of the image on a single step process [9]. The applied method is an unsupervised learning technique with unpaired images and disentangled representation, combined with a GAN architecture. The framework consists in encoders (used to extract relevant features from the raw images: image content, image noise, blur features, and blur-noise features), generators (used to generate from the extracted features blurred, noisy, blurred-noisy and clean images) and discriminators (used to discriminate between generated and real images for each feature). The obtained model was compared with state-of-the-art methods for OCT image enhancement, which were outperformed. Also, a quantitative comparison with state-of-the-art methods indicates that the proposed enhancer performs better than all the other methods, with the best processing speed when the computations were run on a GPU.

2.3. Applications—Cardiovascular Diseases

Cardiovascular ailment (CVD) poses a substantial peril to human well-being and stands as the primary global fatality determinant [10]. The incidence of both mortality and morbidity linked to CVD exhibits an escalating trajectory, particularly within burgeoning territories. This malady precipitates considerable financial ramifications, approximated at 351.2 billion USD in the United States, thereby engendering persistent compromise to the quality of life [11]. Within the European Union, the annual expenditure has been assessed at 210 billion euros, apportioned amongst direct healthcare outlays (53%), diminished productivity (26%), and informal caregiving for individuals afflicted with CVD (21%) [12].

In the field of cardiovascular disease applications four papers are presented. The first paper in this field proposes two methods for binary classifying the risk of CAD based on the CAC ($CAC > 400$ represents high risk of CAD, while $CAC < 400$ low risk of CAD) in diabetic patients [13]: the first method consists in employing a state-of-the-art CNN architecture for CAD risk assessment, based on the retina images and the second method consists in employing classical machine learning classifiers on the clinical data (age and presence of diabetic retinopathy). The DL algorithm considered therein is a VGG16 architecture trained on ImageNet and finetuned on the available retina images. By using the proposed methods, two protocols were established that target two specific applications. The statistics (accuracy, precision, recall, F1 score, confusion matrix) were computed to evaluate each method and the protocols. Results show acceptable accuracies when evaluating the methods independently, while when combining the methods either the precision or recall improve depending on the protocol (the protocol that is created based on the particular needs of each application).

The second paper focuses on obtaining a smaller processing time when using a semi-automated approach for the task of segmenting coronary artery lumen by pre-selecting vessel locations likely to require manual inspection and editing [14]. The pre-selection step is formulated as an Out-of-Distribution (OoD) detection problem with the task of detecting mismatched pairs of CCTA lumen images and their corresponding lumen segmentations. Two Normalizing Flows architectures are employed and assessed: a Glow-like baseline, and a NF architecture which uses a novel coupling layer which exhibits an inductive bias favoring the exploitation of semantical features instead of local pixel correlations. The models were assessed on both synthetic mask perturbations and expert annotations. On synthetic perturbations, the results indicate a better performance for the proposed model, when compared with the baseline model. The proposed model also outperforms the baseline, having a sensitivity for detecting faulty annotations close to inter-expert agreement.

The main objective in the third paper is to evaluate the feasibility of using neural networks in predicting invasively measured FFR from the radius of the coronary lumen that is extracted along the centerline of the coronary artery from OCT images [15]. Three different approaches were used for solving this task: a regression, a classification and an FSL (few shot learning) approach, where the task was formulated also as a classification problem. For each approach different types of architectures were considered: ANN, CNN and RNN. The evaluation step is performed on ensembles for each architecture type: each proposed architecture is trained 20 times, with different random seeds, and the final prediction is performed by the mean value (for regression)/probabilities (for classification) of all 20 models. The FSL CNN based ensemble shows the best diagnostic performance, while being the most robust approach by having the smallest standard deviation and uncertainty. Moreover, compared with baseline approaches based on MLD or %DS FSL CNN reaches improved results. Also, the authors demonstrated that the dataset size has a significant impact on the accuracy: a linear increase in performance was observed as a function of dataset size.

The aim of the last paper of this section is to classify fundus images into five classes of diabetic retinopathy by using neural networks with transfer learning and data augmentation [16]. Three architectures were employed for this task: VGG19, ResNet50, DenseNet169. All models were first finetuned on a public dataset (APTOS). Since the public dataset was imbalanced the models were enhanced by further finetuning on the augmented public dataset (APTOS augmented). The resulting models were tested on a blinded test dataset. Results indicate that ResNet50 performs better than all the other models on all classes.

2.4. Applications—Other

In the last category we included three other applications based on medical imaging, related to cancer, gastrointestinal disorders, and respectively medical report generation. In the first paper, the authors try to optimize the radiological workload, by using the knowledge graph method, a novel method that enhances search engines in general proposed by

Google in 2012 [17]. Firstly, there is an initial knowledge association between disease labels, that are defined as nodes. This is done in two steps, with the help of CheXpert tagger, that classifies the reports into 14 different categories, and the SentencePiece tagger/de-tagger tool, from which the nouns with top k occurrences are selected as additional disease categories. Based on this, a graph convolutional neural network is used to aggregate information between nodes, creating prior knowledge. This is done by generating a hybrid image-text feature, with features extracted from X-ray images with the help of a CNN, and text features extracted from the associated clinical reports using transformers. The transformers represent a better option compared to the classic RNN approach, as radiology reports tend to consist of longer sentences. This hybrid pair is sent through the graph convolutional network, and the node features are split into two branches: a linear classifier for disease classification and a generator head for the report itself. The result is fine-tuned by re-running it through the text classifier. The results are evaluated by the quality of the Natural Language Generation and as well as the clinical efficacy.

The second paper uses the DDSM dataset (Digital Database for Mammography Screening), with a total of 2620 films, of which 695 are normal and 1925 are abnormal, to improve the detection of Breast Cancer, a pathology encountered worldwide, which puts at risk many lives [18]. The employed model, unlike the classic CNN architectures such as VGG, ResNet, MobileNet, etc. is represented by the novel CoroNet: based on the Xception CNN architecture, which consists of a 71-layer deep CNN architecture, pretrained on an ImageNet dataset. The key efficiency improvement of this architecture is the depth wise separable convolution layers with residual connections, which enable a decrease in the number of operations. Separable convolution replaces the classical $n \times n \times k$ convolution with a $1 \times 1 \times k$ point-wise convolution, followed by a channel-wise $n \times n$ spatial convolution, and the residual connections represent “skip connections” which enable the flow of gradients without the need for non-linear functions of activation. This mitigates the disappearing gradient issue. As per the results, this solution outperforms alternative networks.

Finally, in the third paper, a CNN backbone, ResNet-50 pre-trained on ImageNet, is used to extract features from static images, and a GCN (Graph Convolutional Network) is employed for classifying the relationship between labels [19]. An LSTM architecture was used for the temporal association between subsequent frames in the gastroscopy. Those form the proposed GL-Net architecture, which combines the label extraction feature and temporal correlation and dependencies, for real-time predictions in gastroscopy videos. The dataset consists of 49 videos and after video processing and 5 Hz sampling, 23,471 training images and 5798 test images are obtained, with multi-label annotations.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Sharma, P.; Suehling, M.; Flohr, T.; Comaniciu, D. Artificial intelligence in diagnostic imaging: Status quo, challenges, and future opportunities. *J. Thorac. Imaging* **2020**, *35*, S11–S16. [\[CrossRef\]](#) [\[PubMed\]](#)
2. Puiu, A.; Vizitiu, A.; Nita, C.; Itu, L.; Sharma, P.; Comaniciu, D. Privacy-preserving and explainable AI for cardiovascular imaging. *Stud. Inform. Control.* **2021**, *30*, 21–32. [\[CrossRef\]](#)
3. Popescu, A.; Taca, I.; Vizitiu, A.; Nita, C.; Suci, C.; Itu, L.; Scafa-Udriste, A. Obfuscation Algorithm for Privacy-Preserving Deep Learning-Based Medical Image Analysis. *Appl. Sci.* **2022**, *12*, 3997. [\[CrossRef\]](#)
4. Yaqoob, M.; Nazir, M.; Yousafzai, A.; Khan, M.; Shaikh, A.; Algarni, A.; Elmannai, H. Modified Artificial Bee Colony Based Feature Optimized Federated Learning for Heart Disease Diagnosis in Healthcare. *Appl. Sci.* **2022**, *12*, 12080. [\[CrossRef\]](#)
5. Yaqoob, M.; Nazir, M.; Khan, M.; Qureshi, S.; Al-Rasheed, A. Hybrid Classifier-Based Federated Learning in Health Service Providers for Cardiovascular Disease Prediction. *Appl. Sci.* **2023**, *13*, 1911. [\[CrossRef\]](#)
6. Bauer, D.; Ulrich, C.; Russ, T.; Golla, A.; Schad, L.; Zöllner, F. End-to-End Deep Learning CT Image Reconstruction for Metal Artifact Reduction. *Appl. Sci.* **2022**, *12*, 404. [\[CrossRef\]](#)
7. Puiu, A.; Reaungamornrat, S.; Pfeiffer, T.; Itu, L.; Suci, C.; Ghesu, F.; Mansi, T. Generative Adversarial CT Volume Extrapolation for Robust Small-to-Large Field of View Registration. *Appl. Sci.* **2022**, *12*, 2944. [\[CrossRef\]](#)
8. Yaqub, M.; Jinchao, F.; Ahmed, S.; Arshid, K.; Bilal, M.; Akhter, M.; Zia, M. GAN-TL: Generative Adversarial Networks with Transfer Learning for MRI Reconstruction. *Appl. Sci.* **2022**, *12*, 8841. [\[CrossRef\]](#)

9. Li, S.; Azam, M.; Gunalan, A.; Mattos, L. One-Step Enhancer: Deblurring and Denoising of OCT Images. *Appl. Sci.* **2022**, *12*, 10092. [[CrossRef](#)]
10. Thomas, H.; Diamond, J.; Vieco, A.; Chaudhuri, S.; Shinnar, E.; Cromer, S.; Perel, P.; Mensah, G.A.; Narula, J.; Johnson, C.O.; et al. Global Atlas of Cardiovascular Disease 2000–2016: The Path to Prevention and Control. *Glob. Heart* **2018**, *13*, 143–163. [[CrossRef](#)] [[PubMed](#)]
11. Virani, S.S.; Alonso, A.; Benjamin, E.J.; Bittencourt, M.S.; Callaway, C.W.; Carson, A.P.; Chamberlain, A.M.; Chang, A.R.; Cheng, S.; Delling, F.N.; et al. Heart Disease and Stroke Statistics—2020 Update: A Report from the American Heart Association. *Circulation* **2020**, *141*, e139–e596. [[CrossRef](#)] [[PubMed](#)]
12. Timmis, A.; Townsend, N.; Gale, C.; Grobbee, R.; Maniadakis, N.; Flather, M.; Wilkins, E.; Wright, L.; Vos, R.; Bax, J.J.; et al. European Society of Cardiology: Cardiovascular Disease Statistics 2017. *Eur. Heart J.* **2018**, *39*, 508–579. [[CrossRef](#)] [[PubMed](#)]
13. Barriada, R.; Simó-Servat, O.; Planas, A.; Hernández, C.; Simó, R.; Masip, D. Deep Learning of Retinal Imaging: A Useful Tool for Coronary Artery Calcium Score Prediction in Diabetic Patients. *Appl. Sci.* **2022**, *12*, 1401. [[CrossRef](#)]
14. Ciuşdel, C.; Itu, L.; Cimen, S.; Wels, M.; Schwemmer, C.; Fortner, P.; Seitz, S.; Andre, F.; Buß, S.; Sharma, P.; et al. Normalizing Flows for Out-of-Distribution Detection: Application to Coronary Artery Segmentation. *Appl. Sci.* **2022**, *12*, 3839. [[CrossRef](#)]
15. Hatfaludi, C.; Tache, I.; Ciuşdel, C.; Puiu, A.; Stoian, D.; Itu, L.; Calmac, L.; Popa-Fotea, N.; Bataila, V.; Scafa-Udrisite, A. Towards a Deep-Learning Approach for Prediction of Fractional Flow Reserve from Optical Coherence Tomography. *Appl. Sci.* **2022**, *12*, 6964. [[CrossRef](#)]
16. Mungloo-Dilmohamud, Z.; Heenaye-Mamode Khan, M.; Jhumka, K.; Beedassy, B.; Mungloo, N.; Peña-Reyes, C. Balancing Data through Data Augmentation Improves the Generality of Transfer Learning for Diabetic Retinopathy Classification. *Appl. Sci.* **2022**, *12*, 5363. [[CrossRef](#)]
17. Yu, T.; Hu, H.; Zhang, X.; Lei, H.; Liu, J.; Hu, W.; Duan, H.; Si, J. Real-Time Multi-Label Upper Gastrointestinal Anatomy Recognition from Gastroscope Videos. *Appl. Sci.* **2022**, *12*, 3306. [[CrossRef](#)]
18. Mobark, N.; Hamad, S.; Rida, S. CoroNet: Deep Neural Network-Based End-to-End Training for Breast Cancer Diagnosis. *Appl. Sci.* **2022**, *12*, 7080. [[CrossRef](#)]
19. Zhang, D.; Ren, A.; Liang, J.; Liu, Q.; Wang, H.; Ma, Y. Improving Medical X-ray Report Generation by Using Knowledge Graph. *Appl. Sci.* **2022**, *12*, 11111. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.