AI-Driven Innovations in Healthcare: Bridging Imaging and Genomics for Advanced Disease Insights

Carlo Adornetto[1,**], Pierangela Bruno[1], Francesco Calimeri[1], Edoardo De Rose[1], Gianluigi Greco[1] and Alessandro Quarta[1,2]

[1]Department of Mathematics and Computer Science, University of Calabria, Via Pietro Bucci, Rende, Italy
[2]Department of Computer, Control and Management Engineering "Antonio Ruberti", Sapienza University of Rome, Via Ariosto 25, Rome, Italy

Abstract
The application of Artificial Intelligence (AI) techniques for analyzing medical images and omics data is revolutionizing the healthcare industry by offering profound insights into various diseases. Achieving precise diagnoses and formulating effective treatment plans, however, demands intricate and multimodal analysis of complex, sensitive, and diverse medical datasets. Recent advancements in Machine Learning and Deep Learning have proven to be formidable in identifying and classifying specific diseases. This paper outlines the current projects undertaken by our research group in this innovative domain.

Keywords
Artificial Intelligence, Medical Imaging, Genomics, Deep Learning

1. Introduction

The rapid advancement of technology and increased data availability have positioned Artificial Intelligence (AI) as a cornerstone in healthcare. AI significantly enhances patient care, refines treatment protocols, and accelerates the diagnosis of diverse health conditions. Notably, AI has advanced medical imaging and omics analysis, refining diagnostic accuracy and personalizing treatment strategies.

Deep Learning (DL), a subset of AI, excels in analyzing medical images. Its ability to autonomously identify critical features and yield accurate interpretations has made it essential for analyzing complex visual data in medical imaging modalities such as X-rays, MRI, CT scans, PET, and ultrasound. These capabilities are crucial for diagnosing complex conditions like cancers, and cardiovascular and neurological disorders.

However, the assembly of extensive datasets poses significant challenges. To address this, Continual Learning (CL) has emerged as a solution, enabling models to adapt through ongoing data streams, thus enhancing scalability and application efficiency resulting in more sustainable and less resource-demanding systems [1, 2, 3].

In omics analysis, DL has excelled by exploring the vast arrays of biological molecules, aiding in disease understanding and treatment customization across fields like *genomics*, *transcriptomics*, *proteomics*, and *metabolomics*. Advancements in high-throughput and next-generation sequencing technologies have fueled significant progress in functional genomics, especially in understanding cancer-related genomic factors [4].

Despite the potential, DL models often suffer from a lack of interpretability, a critical challenge in bioinformatics. The rise of Explainable Artificial Intelligence (XAI) aims to enhance model transparency and improve feature selection. Techniques like Shapely Additive exPlanations (SHAP) and Gradient-weighted Class Activation Mapping (Grad-CAM) have become pivotal in demystifying the decisions of Neural Networks (NNs), providing clearer insights into their predictive mechanisms [5, 6].

This paper, following our previous work [7], outlines our recent advancements in medical imaging and omics data analysis, paving the way for an in-depth exploration of AI's evolving role in healthcare. The forthcoming sections discuss medical imaging in Section 2, and omics-scale data analysis in Section 3.1, concluding with a comprehensive overview in Section 5.

2. Medical Imaging and AI

2.1. Vessel segmentation of cine-angiography

The methodology adopted in this study systematically enhances the evaluation of vascular complexity in Peripheral Arterial Occlusive Disease (PAOD) patients through the integration of advanced imaging segmenta-

Ital-IA 2024: 4th National Conference on Artificial Intelligence, organized by CINI, May 29-30, 2024, Naples, Italy
*Corresponding author.
carlo.adornetto@unical.it (C. Adornetto); pierangela.bruno@unical.it (P. Bruno); francesco.calimeri@unical.it (F. Calimeri); edoardo.derose@unical.it (E. D. Rose); gianluigi.greco@unical.it (G. Greco); alessandro.quarta@uniroma1.it (A. Quarta)
0000-0002-9734-1017 (C. Adornetto); 0000-0002-0832-0151 (P. Bruno); 0000-0002-0866-0834 (F. Calimeri); 0000-0002-0032-9434 (E. D. Rose); 0000-0002-5799-6828 (G. Greco); 0000-0001-6319-2466 (A. Quarta)

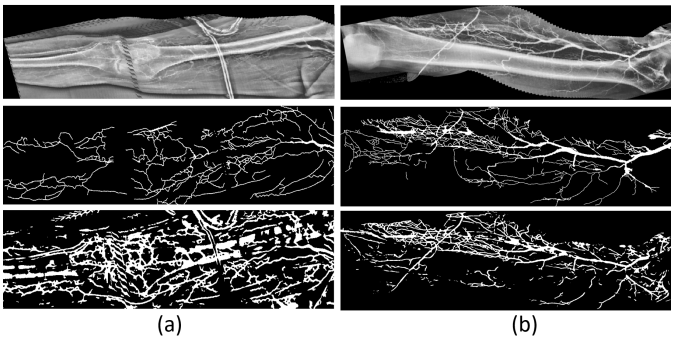

Figure 1: Outcomes from automatic segmentation on two distinct patients are depicted. The figure on the left, Fig. a, represents the patient with the lowest AUC value across the entire study group, whereas Fig. b on the right displays the patient with the highest AUC value. Each figure progresses from top to bottom, beginning with the stitched image (i.e., the grayscale input image), followed by the ground truth segmentation, and concluding with the automatically segmented image.

tion and computational analysis. Utilizing deep learning techniques, this research successfully transforms cine-angiography videos into detailed static images, markedly enhancing the clarity and reliability of vascular assessments. Furthermore, the adoption of fractal dimension as a quantitative metric for vascular complexity introduces a novel, objective criterion to the field. This dual approach not only promises to mitigate the subjectivity inherent in current diagnostic practices but also establishes a robust correlation with conventional clinical evaluations, potentially revolutionizing PAOD management strategies [8]. Incorporating advanced imaging segmentation and computational analysis, our method significantly refines the assessment of vascular complexity in PAOD patients. Figure 1 vividly illustrates the segmented vascular trees from cine-angiography, alongside their corresponding fractal dimension analysis, showcasing the clarity and precision of our deep learning-based approach. The study achieved significant findings, demonstrating that the deep learning-based segmentation method resulted in an Area Under the Curve mean value of 0.77 ± 0.07, with a range from 0.57 to 0.87. This method significantly improved the reliability of visual assessments of vascular complexity, achieving an Inter-Class Correlation Coefficient (ICC) of 0.96 for segmented images, compared to 0.76 for video assessments. Additionally, the Fractal Dimension (FD) analysis correlated well with clinical scores, showing coefficients of 0.85 for manually segmented images and 0.75 for automatically segmented images.

2.2. Segmentation

Semantic segmentation, a process that entails labeling each pixel of an image with a specific class, represents a major leap forward within the realm of medical imaging. This method has been widely adopted for its critical role in identifying tumors, recognizing various organs, and classifying different tissue types, proving to be an invaluable tool in enhancing diagnostic accuracy and patient care [9].

2.2.1. Laryngeal Endoscopic Images

In this work, we present a novel approach using deep learning (DL) for performing semantic segmentation on laryngeal endoscopy images, building upon the foundations laid by previous research [10, 11]. The dataset utilized in this study includes 536 color images manually segmented from in vivo laryngeal examinations, all at a resolution of 512×512 pixels, originating from two separate surgical procedures. These images are categorized into seven distinct groups: *void, vocal folds, other tissue, glottal space, pathology, surgical tool*, and *intubation*. Our model's predictive capabilities were significantly enhanced by leveraging the capabilities of rule-based languages, especially Answer Set Programming (ASP). Incorporating ASP allowed us to navigate the neural network's (NN) decision-making with greater precision, applying penalties for inaccuracies grounded in well-established knowledge. Moreover, rule-based methods were applied to refine our model's output, successfully rectifying minor mistakes, such as single pixels mislabeled, and adjusting misclassified categories that were inconsistent with medical guidelines.

In summary, our approach has shown substantial effectiveness, attaining an average Intersection over Union (IoU) score above 0.7, a figure significantly improved by subsequent post-processing strategies.

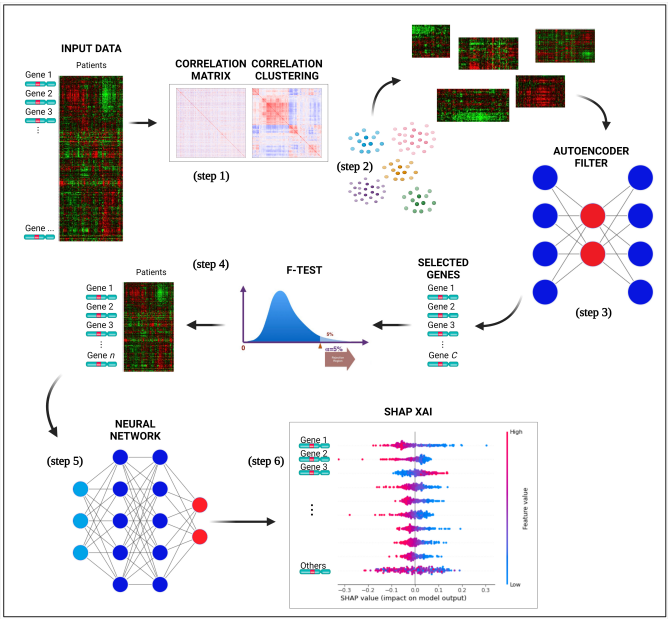

Figure 2: The proposed algorithm for selecting a subset of genes relevant to classify CLL patients. The input data is used to compute the genes pairwise correlation matrix (step 1), and the correlation matrix is clustered (step 2) to group similarly correlated genes. The clusters are then mapped to the original input data and transposed. AEs are trained for each cluster to select the most representative gene, reducing dimensionality (step 3). The genes are ranked with F-test, selecting a subset with the highest F-value (step 4). A neural network is trained with a selected set of genes to perform binary classification of the CLL patients (step 5). The best NNs architecture is determined through model selection, and the SHAP XAI method explains each gene's importance in the predictions (step 6).

3. Engineered Data Encoding for Medical Advancements

In this section, we delve into the innovative intersection of feature engineering and medicine, focusing on manipulating latent spaces to enable new AI-based solutions. We explore a series of our works in which we exploit suitably defined latent spaces to design new gene selection algorithms and Generative AI approaches. In the following, we will discuss a new algorithm for gene selection and its application to Chronic Lymphocytic Leukemia (CLL), and two new generative AI approaches used for automatic report generation and inverse design of materials and molecules. Our works not only showcase the potential of latent spaces in enhancing precision and efficiency in medical research but also highlight their role in fostering the development of novel therapeutic strategies, marking a significant stride toward the future of personalized medicine.

3.1. AI for Omics Data Analysis

Functional genomics data, particularly GEP datasets, are crucial in medical science for diagnosis, prevention, and tailored treatments, yet their analysis is complex due to three main reasons: (1) *course of dimensionality*: a genomics dataset typically consists of a very large number of genes (features) for a small number of patients (samples); (2) *imbalanced classes*: there is often a significant difference between the number of instances in each group of interest; (3) *Noise* sequencing data are typically collected from multiple sources, different laboratories, and sequencing tools resulting in noisy datasets difficult to analyze.

We proposed a new algorithm for genomic-scale analysis, based on DL and XAI, whose aim is threefold: first, select the most meaningful genes for a regression/classification problem; second, provide a more accurate prediction model; third, quantify and evaluate the feature's contribution to the predictions through XAI [12]. The proposed algorithm is based on two main ideas: (1) recognize similarly correlated features using clustered correlation matrix and then filter the redundant information for each group by using Autoencoders (AEs). In contrast with previous works, where AEs are used for dimensionality reduction [13], we implemented a mechanism to still work at the level of the original features. We hence provide a more treatable dataset in terms of dimensionality,

without affecting interpretability; (2) we train NNs and we iteratively select the most meaningful features using a new ad-hoc defined XAI score. We eventually use the set of selected features (from all the iterations) to train and explain a final model.
We used a preliminary version of this algorithm (depicted in Figure 2) for the GEP analysis of CLL patients. In our work [14] we introduced the DeepSHAP Autoencoder Filter for Genes Selection (DSAF-GS), a deep learning and explainable AI-based method for gene selection in genomics-scale data analysis. Through the SHAP explainable AI techniques, we identified key genes influencing CLL prognosis with high accuracy. Our findings pave the way for more targeted bio-molecular research in CLL, suggesting novel paths for investigating disease mechanisms and therapy timing.

3.2. Building and Exploring Meaningful Latent Spaces for Generative AI in Medicine

Automatic Medical Report Generation via Latent Space Conditioning and Transformers
In this work, we explore the integration of artificial intelligence within healthcare, focusing on automatic medical report generation. We introduce the VAE-GPT architecture, combining Variational Autoencoder (VAE) and Generative Pre-trained Transformer (GPT) for generating medical reports from images. The VAE learns a latent representation of images, capturing underlying patterns, while the GPT uses this representation to generate coherent text. For the purpose the VAE is jointly trained with a pre-trained text generator (GPT) and a tags predictor such that images belonging to the same context (e.g. diseases) are placed in the same region of the latent space. Furthermore, we propose a novel metric, Medical Embeddings Attention Distance (MEAD), to measure the semantic similarity between generated and reference reports. Our experiments demonstrate state-of-the-art performance in creating informative medical reports, highlighting the potential of AI in enhancing diagnostic processes [15].

GIDnets: Generative Neural Networks for Solving Inverse Design Problems via Latent Space Exploration
In fields such as Engineering, Molecular Biology, and Physics, the design of technological tools and device structures is progressively supported by Inverse Design methods, providing suggestions on crucial architectural choices based on the properties that these tools and devices should exhibit. The *inverse* design problem aims at designing proper devices according to a set of desired properties and it is typically an ill-posed problem suffering from the non-uniqueness of the solution where, moreover, very different devices can share identical properties. Furthermore, the design spaces are likely high-dimensional and subjected to feasibility constraints.
Most of the state-of-the-art DL methods for inverse design share the idea of looking for the design solution by directly working at the level of the design space; indeed, they have been mainly conceived to deal with applications where such a space is a low-dimensional space. By departing from these approaches, a few works in the literature have already advocated the benefits of mapping the input space into a continuous latent space. This perspective influenced our work which proposes a neural network architecture, named GIDnet (Generative Inverse Design Network), where the suitable solutions are additionally constrained to the only feasible region of the latent design space, and an exploration algorithm is used to end up with more accurate solutions [16]. A thorough experimental activity over several state-of-the-art benchmark datasets evidenced the superior performance of GIDnet for inverse design problems.
In a promising future scenario, our approach can be built using GNNs to generate specific social networks, molecules, and topological representations starting from the prior desired properties. Our generative approach, indeed, demonstrated breakthrough performances in such scenarios where the design space is large, discrete, and constrained, taking into account such feasibility constraints during the design process itself.

4. Other Research Activities

This research group has also engaged in a variety of studies including the impact of a Nutrition Education Program combined with physical activity on the Mediterranean Diet adherence and inflammatory biomarkers in adolescents, showing significant improvements [17]. Additionally, they have examined the dynamics of opinion diffusion within social networks, identifying effective strategies based on centrality measures for influencing opinion adoption [18]. Furthermore, [19] have proposed a neuro-symbolic AI approach for the compliance verification of electrical control panels in Industry 4.0, utilizing a combination of deep learning and Answer Set Programming to detect anomalies with limited data. In [20] developed a Graph Neural Network model to assess lateral spreading displacement in New Zealand, aiming to enhance earthquake impact predictions. In [21] is presented a statistical framework to learn more effectively from algorithm validation challenges, specifically for medical image analysis in laparoscopic videos, identifying underexposure and motion as significant sources of errors. [22] introduced a deep learning framework using heatmaps for disease classification based on gene expression data,

demonstrating its effectiveness in tumor classification. In [23] detailed a method for reducing and visualizing data for automatic diagnosis using gene expression and clinical data, achieving high recall rates in diagnoses. Lastly, we also developed a system to improve the interpretability of automatic diagnosis by analyzing the internal decision-making processes of neural networks [24].

5. Conclusion

This work advances the application of Artificial Intelligence (AI) and Deep Learning (DL) in medical diagnostics and genomics, demonstrating their transformative potential for enhancing diagnostic accuracy and enabling personalized medicine. By employing advanced imaging segmentation, computational analysis, and introducing fractal dimension as a novel metric for vascular complexity, we offer innovative solutions to the challenges in medical imaging and omics data analysis. Our findings highlight the effectiveness of these methods in improving the reliability of medical assessments and the interpretability of complex data through Explainable Artificial Intelligence (XAI) techniques. The integration of AI in healthcare, as illustrated by our research, promises to refine diagnostic processes, optimize treatment plans, and contribute significantly to the future of personalized patient care.

Acknowledgements

This work has been partially funded by PON "Ricerca e Innovazione" 2014-2020, CUP: H25F21001230004, and has been carried out while Alessandro Quarta was enrolled in the Italian National Doctorate on Artificial Intelligence run by Sapienza University of Rome with University of Calabria.

References

[1] C. S. Lee, A. Y. Lee, Clinical applications of continual learning machine learning, The Lancet Digital Health 2 (2020) e279–e281.

[2] A. Quarta, P. Bruno, F. Calimeri, Continual learning for medical image classification, in: F. Calimeri, M. Dragoni, F. Stella (Eds.), 1st AIxIA Workshop on Artificial Intelligence For Healthcare (HC@AIxIA 2022) co-located with the 21st International Conference of the Italian Association for Artificial Intelligence (AIxIA 2022), Udine, Italy, November 30 2022, volume 3307 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2022, pp. 67–76. URL: https://ceur-ws.org/Vol-3307/paper7.pdf.

[3] E. De Rose, Continual learning: an approach via feature maps extrapolation, in: DC@AIxIA23 Doctoral Consortium of AIxIA 2023 conference, CEUR Proceedings, volume 3537, Italian Association for Artificial Intelligence, 2023, p. To be assigned.

[4] E. Alhenawi, R. Al-Sayyed, A. Hudaib, S. Mirjalili, Feature selection methods on gene expression microarray data for cancer classification: A systematic review, Computers in Biology and Medicine 140 (2022) 105051.

[5] S. M. Lundberg, S.-I. Lee, A unified approach to interpreting model predictions, Advances in neural information processing systems 30 (2017).

[6] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, D. Batra, Grad-cam: Visual explanations from deep networks via gradient-based localization, in: Proceedings of the IEEE international conference on computer vision, 2017, pp. 618–626.

[7] C. Adornetto, P. Bruno, F. Calimeri, E. DE ROSE, G. Greco, et al., Artificial intelligence in medicine: From imaging to omics, in: CEUR WORKSHOP PROCEEDINGS, volume 3486, 2023, pp. 140–145.

[8] P. Bruno, M. F. Spadea, S. Scaramuzzino, S. De Rosa, C. Indolfi, G. Gargiulo, G. Giugliano, G. Esposito, F. Calimeri, P. Zaffino, Assessing vascular complexity of paod patients by deep learning-based segmentation and fractal dimension, Neural Computing and Applications 34 (2022) 22015–22022.

[9] T. Dhamija, A. Gupta, S. Gupta, R. Katarya, G. Singh, Semantic segmentation in medical images through transfused convolution and transformer networks, Applied Intelligence 53 (2023) 1132–1148.

[10] M.-H. Laves, J. Bicker, L. A. Kahrs, T. Ortmaier, A dataset of laryngeal endoscopic images with comparative study on convolution neural network-based semantic segmentation, International journal of computer assisted radiology and surgery 14 (2019) 483–492.

[11] P. Bruno, F. Calimeri, C. Marte, M. Manna, Combining deep learning and asp-based models for the semantic segmentation of medical images, in: Rules and Reasoning: 5th International Joint Conference, RuleML+ RR 2021, Leuven, Belgium, September 13–15, 2021, Proceedings 5, Springer, 2021, pp. 95–110.

[12] C. Adornetto, G. Greco, A new deep learning and xai-based algorithm for features selection in genomics, 2023. arXiv:2303.16914.

[13] P. Danaee, R. Ghaeini, D. A. Hendrix, A deep learning approach for cancer detection and relevant gene identification, in: Pacific symposium on biocomputing 2017, World Scientific, 2017, pp. 219–229.

[14] F. Morabito, C. Adornetto, P. Monti, A. Amaro, F. Reggiani, M. Colombo, Y. Rodriguez-Aldana, G. Tripepi, G. D'Arrigo, C. Vener, et al., Genes selection using deep learning and explainable artifi-

cial intelligence for chronic lymphocytic leukemia predicting the need and time to therapy, Frontiers in Oncology 13 (2023).

[15] C. Adornetto, A. Guzzo, A. Vasile, Automatic medical report generation via latent space conditioning and transformers, in: 2023 IEEE Intl Conf on Dependable, Autonomic and Secure Computing, Intl Conf on Pervasive Intelligence and Computing, Intl Conf on Cloud and Big Data Computing, Intl Conf on Cyber Science and Technology Congress (DASC/PiCom/CBDCom/CyberSciTech), IEEE, 2023, pp. 0428–0435.

[16] C. Adornetto, G. Greco, Gidnets: generative neural networks for solving inverse design problems via latent space exploration, in: Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, 2023, pp. 3404–3413.

[17] C. Morelli, E. Avolio, A. Galluccio, G. Caparello, E. Manes, S. Ferraro, A. Caruso, D. De Rose, I. Barone, C. Adornetto, et al., Nutrition education program and physical activity improve the adherence to the mediterranean diet: impact on inflammatory biomarker levels in healthy adolescents from the dimenu longitudinal study, Frontiers in Nutrition 8 (2021) 685247.

[18] C. Adornetto, V. Fionda, G. Greco, On the effectiveness of compact strategies for opinion diffusion in social environments, in: ECAI 2023, IOS Press, 2023, pp. 11–18.

[19] V. Barbara, M. Guarascio, N. Leone, G. Manco, A. Quarta, F. Ricca, E. Ritacco, Neuro-symbolic ai for compliance checking of electrical control panels, Theory and Practice of Logic Programming 23 (2023) 748–764.

[20] M. Durante, G. Terremoto, C. Adornetto, G. Greco, E. Rathje, A new graph neural network (gnn) based model for the evaluation of lateral spreading displacement in new zealand, in: Proceedings of the 8th International Conference on Earthquake Geotechnical Engineering (VIII ICEGE), Osaka, Japan, 2024, pp. 7–10.

[21] T. Roß, P. Bruno, A. Reinke, M. Wiesenfarth, L. Koeppel, P. M. Full, B. Pekdemir, P. Godau, D. Trofimova, F. Isensee, et al., Beyond rankings: Learning (more) from algorithm validation, Medical image analysis 86 (2023) 102765.

[22] P. Bruno, F. Calimeri, Using heatmaps for deep learning based disease classification, in: 2019 IEEE conference on computational intelligence in bioinformatics and computational biology (CIBCB), IEEE, 2019, pp. 1–7.

[23] P. Bruno, F. Calimeri, A. S. Kitanidis, E. De Momi, Data reduction and data visualization for automatic diagnosis using gene expression and clinical data, Artificial Intelligence in Medicine 107 (2020) 101884.

[24] P. Bruno, F. Calimeri, A. S. Kitanidis, E. De Momi, Understanding automatic diagnosis and classification processes with data visualization, in: 2020 IEEE international conference on human-machine systems (ICHMS), IEEE, 2020, pp. 1–6.