

# Virtual Scanner: Leveraging Resilient Generative AI for Radiological Imaging in the Era of Medical Digital Twins

Carolina Adornato<sup>1</sup>, Cecilia Assolito<sup>1</sup>, Ermanno Cordelli<sup>1</sup>, Francesco Di Feola<sup>1,2</sup>, Valerio Guarrasi<sup>1,\*</sup>, Giulio Iannello<sup>1</sup>, Lorenzo Marcoccia<sup>1</sup>, Elena Mulero Ayllon<sup>1</sup>, Rebecca Restivo<sup>1</sup>, Aurora Rofena<sup>1</sup>, Rosa Sicilia<sup>1</sup>, Paolo Soda<sup>1,2</sup>, Matteo Tortora<sup>1</sup> and Lorenzo Tronchin<sup>1,2</sup>

<sup>1</sup>Research Unit of Computer Systems and Bioinformatics, Department of Engineering, Università Campus Bio-Medico di Roma, Via Álvaro del Portillo 21, Rome, 00128, Italy

<sup>2</sup>Department of Diagnostics and Intervention, Radiation Physics, Biomedical Engineering, Umeå University, Umeå, 90187, Sweden

## Abstract

Advancements in generative artificial intelligence (AI) are setting the stage for transformative changes in medical imaging, particularly through the development of the Virtual Scanner. This innovative approach leverages resilient generative AI to synthesize radiological images, addressing critical challenges in the field such as data scarcity, patient exposure to radiation, and the limitations of current imaging technologies. By harnessing the power of Generative Adversarial Networks (GANs) and focusing on the resilience of these algorithms, the Virtual Scanner aims to enhance diagnostic accuracy, improve patient care, and fill gaps in multimodal datasets. Our research explores both unimodal and multimodal techniques, including GAN ensembles, latent augmentation, and advanced texture synthesis, to create robust and adaptable generative models. Through extensive experimentation and analysis, we demonstrate the potential of the Virtual Scanner to revolutionize medical diagnostics by providing a safer, more efficient, and comprehensive imaging solution. The implications of this work extend beyond immediate medical applications, offering insights into the development of AI technologies capable of navigating the complexities of real-world data.

## Keywords

Medical Imaging, Generative Artificial Intelligence, Virtual Scanner, Resilient AI, Multimodal Learning, Radiology

## 1. Introduction

In recent years, the intersection of artificial intelligence (AI) and healthcare has opened up novel possibilities for enhancing diagnostic accuracy, optimizing patient care, and tailoring treatment plans towards precision medicine. One of the most promising developments in this domain is the concept of the Medical Digital Twin, a virtual representation of a patient's health status, enabling personalized medical interventions and predictive healthcare analytics. Central to the utility and effectiveness of Medical Digital Twins is the capability for detailed and accurate radiological imaging, which provides a window into the internal workings of the human body without invasive procedures.

Radiological imaging, encompassing a range of modalities such as X-rays, MRI, and CT scans, plays a pivotal role in the diagnosis, monitoring, and treatment planning for a myriad of health conditions. However, the acquisition of these images often requires patients to undergo multiple scans, exposing them to potential risks associ-

ated with radiation and contrast agents. Furthermore, the reliance on comprehensive multimodal imaging data presents challenges in scenarios where certain modalities are unavailable or unsuitable for some patients, leading to gaps in the data that can hinder diagnostic processes and the development of AI models in healthcare [1, 2, 3].

The advancement of generative AI, particularly through the deployment of Generative Adversarial Networks (GANs), offers a novel solution to these challenges. By enabling the virtual generation of radiological images where real ones are unavailable or undesirable, AI not only mitigates the risks to patients but also bridges the data gaps in multimodal learning applications [4, 5, 6]. We introduce the concept of the Virtual Scanner as a cornerstone of the Medical Digital Twin paradigm, aiming to revolutionize the field of radiology by synthesizing high-fidelity, modality-specific images through the power of AI, thus enhancing patient care and supporting radiologists in delivering more accurate diagnoses.

The scarcity of comprehensive radiological images presents significant challenges in medical diagnostics, affecting the efficacy of diagnostic processes and the development of AI tools. This scarcity arises from limited access to advanced imaging technologies, concerns over radiation exposure, and the difficulty of compiling diverse, multimodal datasets. Such challenges hinder the creation of effective AI models for diagnostics, impacting

*Ital-IA 2024: 4th National Conference on Artificial Intelligence, organized by CINI, May 29-30, 2024, Naples, Italy*

\*Corresponding author.

✉ valerio.guarrasi@unicampus.it (V. Guarrasi)

ORCID 0000-0002-1860-7447 (V. Guarrasi)

© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

their accuracy and real-world applicability. The Virtual Scanner, leveraging resilient generative AI algorithms, addresses these issues by synthesizing radiological images to fill dataset gaps and reduce the need for repeated scans. Resilient generative AI refers to the development of models that not only excel in their designated tasks under ideal conditions but also maintain their performance when confronted with data that deviate from the norm, known as “data in the wild”. Such resilience is crucial in ensuring that the AI tools developed for medical imaging are robust against the variations inherent in patient data across different demographics, equipment used, and pathological conditions.

By embedding resilience at the core of our generative AI algorithms, we aim to create a foundation for the Virtual Scanner that is not only technologically advanced but also reliable and effective across the spectrum of medical imaging needs. This approach positions our work not just as a technical achievement but as a meaningful contribution to the field of radiology, where the capacity to handle data in the wild can significantly enhance diagnostic processes and patient care.

## 2. Research Activities

Embarking on the journey to realize the Virtual Scanner, our investigation delves into a series of research activities, as shown in Figure 1, each designed to push the boundaries of what’s possible with generative AI in the field of radiology [7, 8]. These activities are categorized into two main areas: “Resilient Generative AI” and “Virtual Scanner Applications”, enveloping a diverse array of methodologies and applications aimed at enhancing the generation and translation of medical imaging data. By addressing many aspects of generative AI, from improving algorithm resilience to creating virtual modalities, these activities underscore our commitment to advancing diagnostic capabilities and patient care through technological innovation.

### 2.1. Resilient Generative AI

#### 2.1.1. GAN Ensemble

In tackling the complexities of synthetic data generation within medical imaging, our research delves into optimizing generative AI through the use of GAN ensembles. This strategy is born from the necessity to overcome inherent limitations in single-model GAN applications, such as mode collapse and the inadequate representation of real data distributions, a common obstacle in generating high-quality and diverse medical images. The core of our approach lies in creating an ensemble of GANs that jointly optimizes the visual quality and diversity of synthetic images from a set of GANs. We aim to solve a

Pareto multi-objective optimization problem that simultaneously covers the real training set, aiming to generate high-quality GANs and using as few GANs as possible.

We tested out methodology across three distinct medical datasets, employing 22 GANs with differing architectures, loss functions, and regularization techniques. Moreover, we uniformly sampled each model every 20000 training iterations, i.e., resulting in a total search space of 110 models. The experiments showcase that using synthetic datasets generated from such an ensemble improves the performances in classification downstream tasks compared to single GANs and Naive selection approaches, i.e., using all available 110 GANs or randomly selecting a subset.

#### 2.1.2. LatentAugment

Data Augmentation (DA) is a crucial strategy in AI to enhance the volume and diversity of training datasets, thereby mitigating the risk of overfitting and bolstering model generalization to unseen data. Standard DA methods in image recognition tasks transform the images via geometric rigid and non-rigid transformations using image processing primitives, such as translation, rotation, cropping, etc. However, such transformations rely on human experts with prior knowledge of the dataset and fail to generate sufficiently diverse synthetic data. GANs offer a valuable addition to the available augmentation techniques. However, GANs generate high-quality samples rapidly, but they suffer from poor mode coverage, i.e., the variation and variety of the samples that can be generated, limiting their utility for DA purposes in the medical field.

We propose LatentAugment [9], a DA strategy that overcomes the low diversity of GANs, opening up for use in DA applications. LatentAugment addresses the three-fold challenge of producing synthetic samples that are not only of high fidelity and quality but also diverse and rapidly generated. LatentAugment modifies the latent vectors of the real training set, moving them towards regions that maximize their diversity and fidelity. We applied LatentAugment to improve the performances of a downstream model performing of MRI-to-CT image translation. The results showed LatentAugment’s superiority over common DA methods and naive GAN-sampling, i.e., creating data sampling from the GAN’s latent space without any control.

#### 2.1.3. Paired vs. Unpaired Image Translation

Image-to-Image translation in medical imaging presents a critical challenge due to the predominance of unpaired datasets, where the direct correspondence between source and target images is not established [10]. While paired methods, e.g., Pix2Pix, rely on direct map-

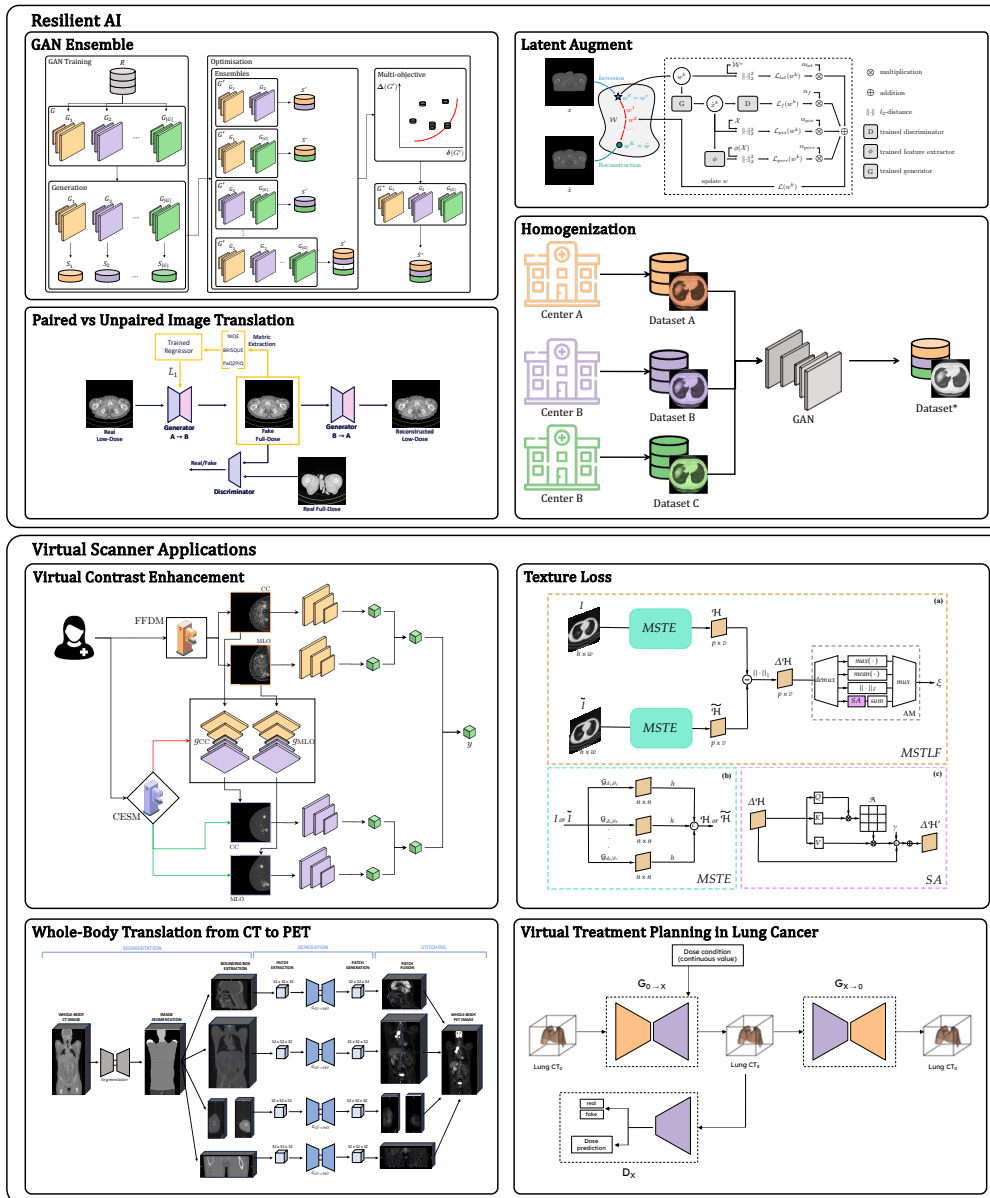


Figure 1: Overview of the presented research activities.

plings between paired images, unpaired techniques, e.g., CycleGAN, offer flexibility by ignoring this requirement. Given the logistical and ethical complexities in acquiring paired medical images, unpaired models have gained significant attention despite the potential compromises in performance attributed to training instability and variable outcomes.

Our work introduces a novel approach to narrow the performance drop between paired and unpaired image

translation methodologies. By integrating a novel paired virtual loss function into the unpaired CycleGAN framework, we enhance the stability and accuracy of unpaired image translation without necessitating direct image pairs. This innovation finds practical application in Low-Dose Computed Tomography (LDCT) denoising, a process aimed at reducing radiation exposure while maintaining image quality. Through this approach, we leverage the abundance of unpaired LDCT and Full-Dose CT

(FDCT) images to validate our model, demonstrating its effectiveness in producing high-quality, denoised images that closely approximate FDCT standards, thereby mitigating health risks associated with radiation without compromising diagnostic integrity.

#### 2.1.4. Homogenization

In the domain of lung CT imaging, the heterogeneity of images stemming from varied scanners and reconstruction kernels poses a significant challenge. This variability can severely impact the performance of automated analysis tools, notably in tasks relying on deep learning models such as 3D Convolutional Neural Networks (CNNs), which are crucial for predicting patient outcomes like overall survival rates. To address this challenge, our work introduces an innovative approach based on StarGAN, a state-of-the-art image-to-image translation generative model, for the homogenization of lung CT images.

Our objective is to transform disparate lung CT images, regardless of their originating scanner types or reconstruction kernels, into a standardized format that retains critical diagnostic features while presenting a uniform appearance. By employing StarGAN, we leverage its capacity for multi-domain image translation to achieve the goal of not only enhancing the quality of the dataset but also to significantly improve the performance of downstream tasks. This approach paves the way for more generalized and robust AI tools in medical diagnostics, ultimately contributing to better patient care and outcomes.

## 2.2. Virtual Scanner Applications

### 2.2.1. Virtual Contrast Enhancement (VCE)

In the evolving landscape of medical imaging, Contrast Enhanced Spectral Mammography (CESM) represents a significant advancement, offering detailed insights for breast cancer diagnosis by utilizing a dual-energy technique that integrates both low and high-energy images. This method, however, necessitates the administration of an iodinated contrast medium and subjects patients to higher radiation doses than standard mammography, raising concerns about potential side effects and increased radiation exposure.

Addressing these critical limitations, our work introduces a novel approach to VCE in CESM using deep generative models [11]. By eliminating the need for contrast mediums and aiming to reduce radiation doses, this research not only mitigates the associated risks but also preserves the diagnostic benefits of CESM. Our methodology employs GAN, e.g., Pix2Pix or CycleGAN, to generate synthetic recombined images from solely low-energy images.

An extensive quantitative and qualitative analysis underpins our research, including evaluations by professional radiologists on a novel CESM dataset comprising 1,138 images. This dataset has been made publicly available to foster ongoing research and development in the field. Among the models tested, CycleGAN emerged as the most effective, showcasing its ability to produce high-quality synthetic recombined images that closely mimic those obtained with traditional contrast-enhanced techniques.

### 2.2.2. Virtual Treatment Planning in Lung Cancer

Monitoring the progression and response to therapy is fundamental in lung cancer treatment. Traditional approaches rely on a series of CT scans taken before and during treatment to evaluate the efficacy of the interventions. In our previous work [12], we developed an ODE-based Digital Twin by using patient-specific CT scans to train a deep reinforcement learning controller, which can adapt to different tissue aggressiveness and outperform the current radiotherapy clinical practice of uniform dose delivery. However, this methodology often exposes patients to additional radiation and can be logistically challenging. Our innovative research introduces a novel application of AI in virtual treatment planning, leveraging conditioned CycleGANs to simulate the potential progression of lung cancer treatment based on varying doses. By conditioning the CycleGAN on specific treatment doses, our model can generate virtual CT scans that predict how the patient's anatomy and the tumor itself might respond to different levels of treatment. This approach allows for the creation of a virtual time series of CT scans without the need for repeated radiation exposure. The ability to accurately forecast the treatment's progression through these synthetic scans offers a significant advantage in personalizing treatment plans, enabling more precise adjustments to therapy regimens based on predicted outcomes. By reducing the reliance on multiple physical CT scans and minimizing patient exposure to radiation, we pave the way for a more patient-centric approach to cancer treatment monitoring. Additionally, the predictive insights gained from this technology could significantly enhance decision-making processes in treatment planning, potentially improving patient outcomes in lung cancer care.

### 2.2.3. Whole-Body Translation from CT to PET

The integration of CT and PET scans is essential in oncological diagnostics, combining the structural clarity of CT with the metabolic insights of PET imaging. While CT scans provide detailed anatomical structure, PET scans offer a window into the metabolic activity within the body, making the combined PET/CT an invaluable tool in the

diagnosis, staging, and management of cancer patients. Despite their clinical significance, the dual-modality approach of PET/CT scanning is not without drawbacks, e.g., additional radiation exposure and higher costs compared to CT-only scans. These limitations restrict the widespread availability of PET/CT imaging in numerous medical centers globally, underscoring the need for alternative methods that can replicate the integrative insights of PET/CT imaging while mitigating its drawbacks.

Recognizing the challenges inherent in translating CT images to PET, especially given the variability in translation effectiveness across different anatomical regions, our methodology introduces a district-specific approach. Drawing from the current literature, which suggests the potential for improved accuracy through organ-specific networks, we propose a novel strategy that segments whole-body images into four major anatomical districts. Each district is then processed through independently trained GANs to generate district-specific PET images. The final step involves stitching these district-specific PET images together to reconstruct a comprehensive whole-body PET scan.

Employing two GAN architectures, Pix2Pix and CycleGAN, our approach facilitates a comparative analysis to evaluate the effectiveness and precision of the image translation process. Through standard evaluation metrics, we quantify the quality of the generated images, highlighting the advantages of our district-specific translation methodology over traditional approaches that rely on a single GAN trained on entire whole-body images. This innovative strategy not only promises to reduce the time, cost, and radiation exposure associated with PET/CT imaging but also offers a tailored approach that accounts for the unique characteristics of different anatomical regions.

#### 2.2.4. Texture Loss

In the quest to enhance the quality of medical images through denoising, the application of GANs emerges as a promising task. Yet, a critical challenge lies in the GAN-based algorithms' capacity to accurately capture and replicate the intricate textural details inherent in medical images. This task's complexity is significantly amplified by the diverse and complex relationships that define image textures, making conventional denoising approaches inadequate for preserving or restoring fine-grained textural fidelity.

Our research introduces a novel loss function tailored to address these limitations by exploiting the multi-scale textural properties captured by the Gray-Level Co-occurrence Matrix (GLCM) [13]. The GLCM, traditionally utilized in image processing to quantify texture, is redefined in our work as a differentiable module compatible with the gradient-based optimization of the GAN train-

ing processes. By integrating a multi-scale, differentiable GLCM into the loss function, we facilitate a deeper understanding and recognition of complex textural information during the image generation phase.

Furthermore, the incorporation of a self-attention layer represents a pivotal innovation in our methodology, enabling the dynamic synthesis of texture information across various scales. This approach not only enhances the denoising capabilities of GANs but also ensures the preservation of essential textural details, thereby improving the diagnostic utility of the generated images.

Extensive experimental validation of our approach within the field of low-dose CT denoising, aimed at improving noisy CT scans while minimizing radiation exposure, underscores the efficacy of our proposed solution. Utilizing three publicly available datasets, including both simulated and real-world scenarios, our methodology demonstrates a notable improvement over traditional loss functions across a variety of GAN architectures.

#### 2.2.5. Report Generation

Our work focuses on Automatic Medical Reporting (AMR) that, as fostered by escalating digitization of healthcare data and the mounting stress national healthcare systems, aims to produce diagnostic reports from biomedical data. The efforts are currently directed towards chest radiographs assessing solutions based on encoder-decoder and transformer-based models. Alongside all this, Quantum Artificial Intelligence represents a novel field whose theoretical superiority in data representation capabilities and processing speeds makes it the main technology we are forwarding our efforts to, with its numerous methodologies for healthcare, even if it still presents hardware immaturity, scalability issues, and substantial financial costs. Because its application in AMR is still uncharted, we aim to develop an architecture that merges traditional encoder-decoder concepts with quantum computing, to transcribe features obtained using classical binary computation into quantum states, which are then entangled with quantum representations of the shifted predictions for computational and accuracy benefits.

### 3. Future Directions and Conclusion

As we stand on the verge of a new era in medical imaging, propelled by the advancements in generative AI, the journey of our exploration is ongoing. The groundwork laid by the Virtual Scanner and the development of resilient generative AI algorithms opens a myriad of pathways for future research and application. In the quest for further innovation, it is essential to delve deeper into the integra-

tion of AI with emerging imaging technologies, aiming to enhance the precision, efficiency, and accessibility of diagnostic tools. Future research will focus on refining the algorithms for even greater resilience [14], enabling them to adapt more seamlessly to the vast diversity of medical imaging data. Additionally, exploring the potential for AI-driven predictive analytics in patient treatment plans presents a promising frontier, where the insights garnered from virtual scans could inform more personalized and effective treatment strategies. Moreover, the ethical considerations and data privacy concerns associated with deploying AI in healthcare require ongoing attention. Ensuring the security of patient data and the unbiased application of AI tools remains paramount as we advance.

In conclusion, the exploration into generative AI and the Virtual Scanner represents a significant leap toward revolutionizing medical imaging. As we move forward, the presented research activities and the technologies developed will undoubtedly pave the way for a future where diagnostics are more accurate, treatments are more personalized, and patient care is enhanced at every level.

## Acknowledgments

Aurora Rofena and Lorenzo Marcoccia are Ph.D. students enrolled in the National Ph.D. in Artificial Intelligence, course on Health and life sciences, organized by Università Campus Bio-Medico di Roma. We acknowledge financial support from: i) PNRR MUR project PE0000013-FAIR; ii) PRIN 2022 MUR 20228MZFAA-AIDA (CUP C53D23003620008); iii) PRIN PNRR 2022 MUR P2022P3CXJ-PICTURE (CUP C53D23009280001); iv) FCS MISE (CUP B89J23000580005). This work was also partially supported by the following companies: Teleconsys S.p.A..

## References

- [1] G. Fiscon, F. Salvatore, V. Guarrasi, A. R. Garbuglia, P. Paci, Assessing the impact of data-driven limitations on tracing and forecasting the outbreak dynamics of covid-19, *Computers in biology and medicine* 135 (2021) 104657.
- [2] V. Guarrasi, P. Soda, Multi-objective optimization determines when, which and how to fuse deep networks: An application to predict covid-19 outcomes, *Computers in Biology and Medicine* 154 (2023) 106625.
- [3] V. Guarrasi, L. Tronchin, D. Albano, E. Faiella, D. Fazzini, D. Santucci, P. Soda, Multimodal explainability via latent shift applied to covid-19 stratification, *arXiv preprint arXiv:2212.14084* (2022).
- [4] V. Guarrasi, N. C. D'Amico, R. Sicilia, E. Cordelli, P. Soda, A multi-expert system to detect covid-19 cases in x-ray images, in: *2021 IEEE 34th International Symposium on Computer-Based Medical Systems (CBMS)*, IEEE, 2021, pp. 395–400.
- [5] V. Guarrasi, N. C. D'Amico, R. Sicilia, E. Cordelli, P. Soda, Pareto optimization of deep networks for covid-19 diagnosis from chest x-rays, *Pattern Recognition* 121 (2022) 108242.
- [6] V. Guarrasi, P. Soda, Optimized fusion of cnns to diagnose pulmonary diseases on chest x-rays, in: *International Conference on Image Analysis and Processing*, Springer, 2022, pp. 197–209.
- [7] V. Guarrasi, L. Tronchin, C. M. Caruso, A. Rofena, G. Manni, F. Aksu, D. Paolo, G. Iannello, R. Sicilia, E. Cordelli, et al., Building an ai-enabled metaverse for intelligent healthcare: opportunities and challenges, in: *Ital-IA 2023, Italia Intelligenza Artificiale Thematic Workshops*, co-located with the 3rd CINI National Lab AIIS Conference on Artificial Intelligence (Ital IA 2023), Pisa, Italy, May 29-30, 2023, *CEUR-WS*, 2023, pp. 134–139.
- [8] E. Cordelli, V. Guarrasi, G. Iannello, F. Ruffini, R. Sicilia, P. Soda, L. Tronchin, Making ai trustworthy in multimodal and healthcare scenarios, *Proceedings of the Ital-IA (2023)*.
- [9] L. Tronchin, M. H. Vu, P. Soda, T. Löfstedt, Latent augmentation: Data augmentation via guided manipulation of gan's latent space, *arXiv preprint arXiv:2307.11375* (2023).
- [10] F. Di Feola, L. Tronchin, P. Soda, A comparative study between paired and unpaired Image Quality Assessment in Low-Dose CT Denoising, in: *2023 IEEE 36th International Symposium on Computer-Based Medical Systems (CBMS)*, IEEE, 2023, pp. 471–476.
- [11] A. Rofena, V. Guarrasi, M. Sarli, C. L. Piccolo, M. Sammarra, B. B. Zobel, P. Soda, A deep learning approach for virtual contrast enhancement in contrast enhanced spectral mammography, *arXiv preprint arXiv:2308.00471* (2023).
- [12] M. Tortora, E. Cordelli, R. Sicilia, M. Miele, P. Matteucci, G. Iannello, S. Ramella, P. Soda, Deep reinforcement learning for fractionated radiotherapy in non-small cell lung carcinoma, *Artificial Intelligence in Medicine* 119 (2021) 102137.
- [13] F. D. Feola, L. Tronchin, V. Guarrasi, P. Soda, Multi-Scale Texture Loss for CT denoising with GANs, 2024. *arXiv:2403.16640*.
- [14] C. M. Caruso, V. Guarrasi, S. Ramella, P. Soda, A deep learning approach for overall survival analysis with missing values, *arXiv preprint arXiv:2307.11465* (2023).