MMIA 2026

Generative AI Meets Super-Resolution: A Unified Pipeline for Cancer Imaging

Pedro Sousa [1,2,*] Tiago Gonçalves [1,2]
Luís Teixeira [1,2] Tania Pereira [1,2] Hélder P. Oliveira [1,3]

[1] Institute for Systems and Computer Engineering, Technology and Science (INESC TEC) Porto, Portugal
[2] Faculty of Engineering, University of Porto (FEUP), Portugal
[3] Faculty of Sciences, University of Porto (FCUP), Portugal
* pedro.fernandes.sousa@inesctec.pt

Introduction

- Cancer remains a leading cause of mortality worldwide, breast and lung cancer being among the most prevalent types. Early and accurate diagnosis of these ailments, be it using deep learning AI models or traditional methods, still depends heavily on the acquisition of high-quality medical image acquisition, which can be both very costly and time-consuming. In other words, clinically-used datasets often suffer from limited quality and scarcity of annotated samples.
- Innovations in generative AI methods have proven to be successful at overcoming such limitations. As such, this paper presents a **two-stage pipeline for medical imaging generation by combining lower-resolution synthesis algorithms with super-resolution methods** in order to guarantee both quality and availability.
- This method places a great emphasis on the combined effort of structural variability generation and fine-detail reconstruction as a way to produce realistic, yet high-quality medical imaging volumes.
- This approach has been validated on breast MRI and lung CT scan data, demonstrating potential for expanding existing datasets and supporting reliable diagnosis workflows.

Dataset

For each of the two clinical imaging contexts, two datasets were explored: Breast MRI data was sourced from the **Duke-Breast-Cancer-MRI dataset** [4] and a private aggregation of 89 volumes in partnership with the *Champalimaud Foundation* ; Lung CT data was acquired from the **LIDC-IDRI** [1] and **RIDER** [2] datasets. All samples were pre-processed in the same manner, consisting of: min-max normalisation, HU unit adjustment (for CT data) and sample resizing from 512^3 to 64^3 and 128^3.

Methodology

- Regarding the generative task, the chosen architecture to produce 64^3 and 128^3-sized samples was a Wasserstein Generative Adversarial Network (GAN) with Gradient Penalty (**WGAN-GP**) [3].
- For generative task experiments, breast MRI data was kept as a control group for the measuring of adversarial performance, while lung CT datasets were tested within two different training regimes: the first experiment used a generator-discriminator update per iteration ratio of 1-to-1 and the second one of 5-to-1. This approach allowed for the evaluation of the models' sensitivity to overfitting and mode collapse.
- As for the super-resolution task, Real Enhanced Super-Resolution GAN (**Real-ESRGAN**) [6] was the model used for $4x$ resolution enhancement (from 64^3 to 256^3 and from 128^3 to 512^3). Bicubic interpolation served as the baseline comparison for this task.
- Super-resolution experiments were performed using each imaging modality separately and also a combination of the two. The resulting model from the latter was then also tested for each modality, as a way to measure the architecture's generalisability. Final experiments also included fine-tuning and the introduction of heavy blurring to LR samples. Further experimental details are also included in [5].
- Performance evaluation comprised not only quantitative metrics, but also a Visual Turing Test (VTT) with radiologists so as to assess perceptual realism and bridge algorithmic quality with clinical applicability.

Results & Discussion

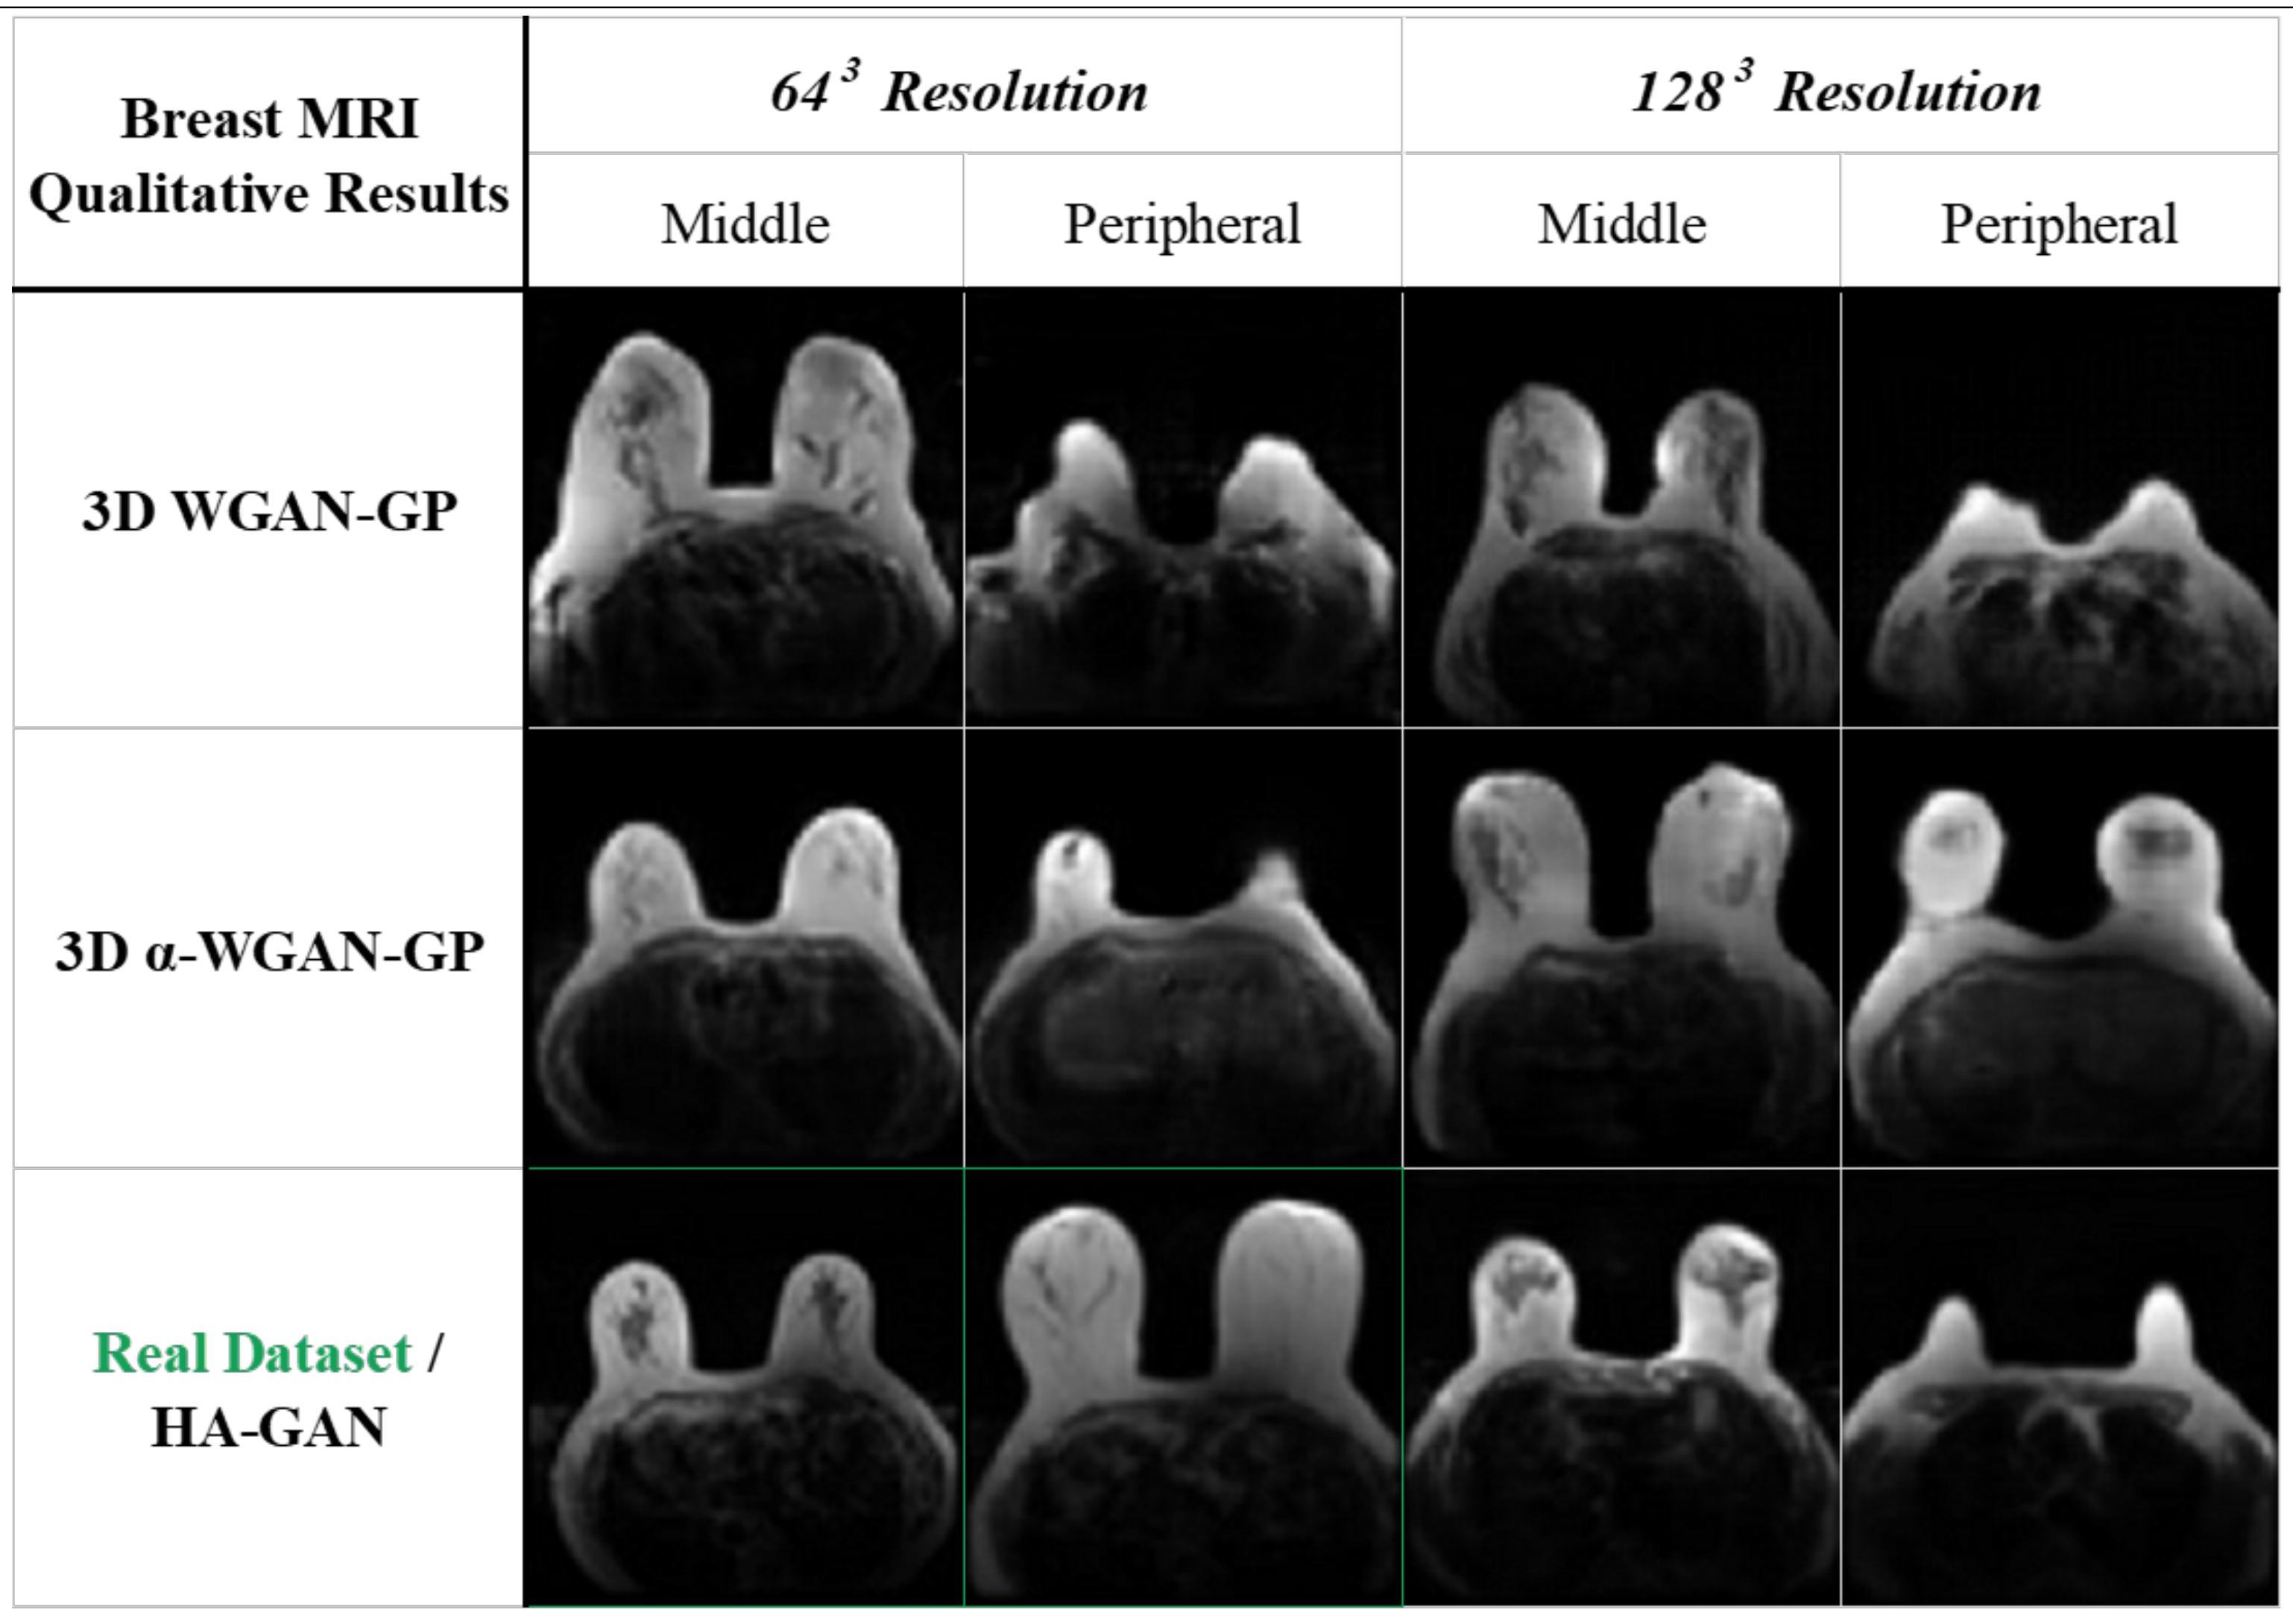

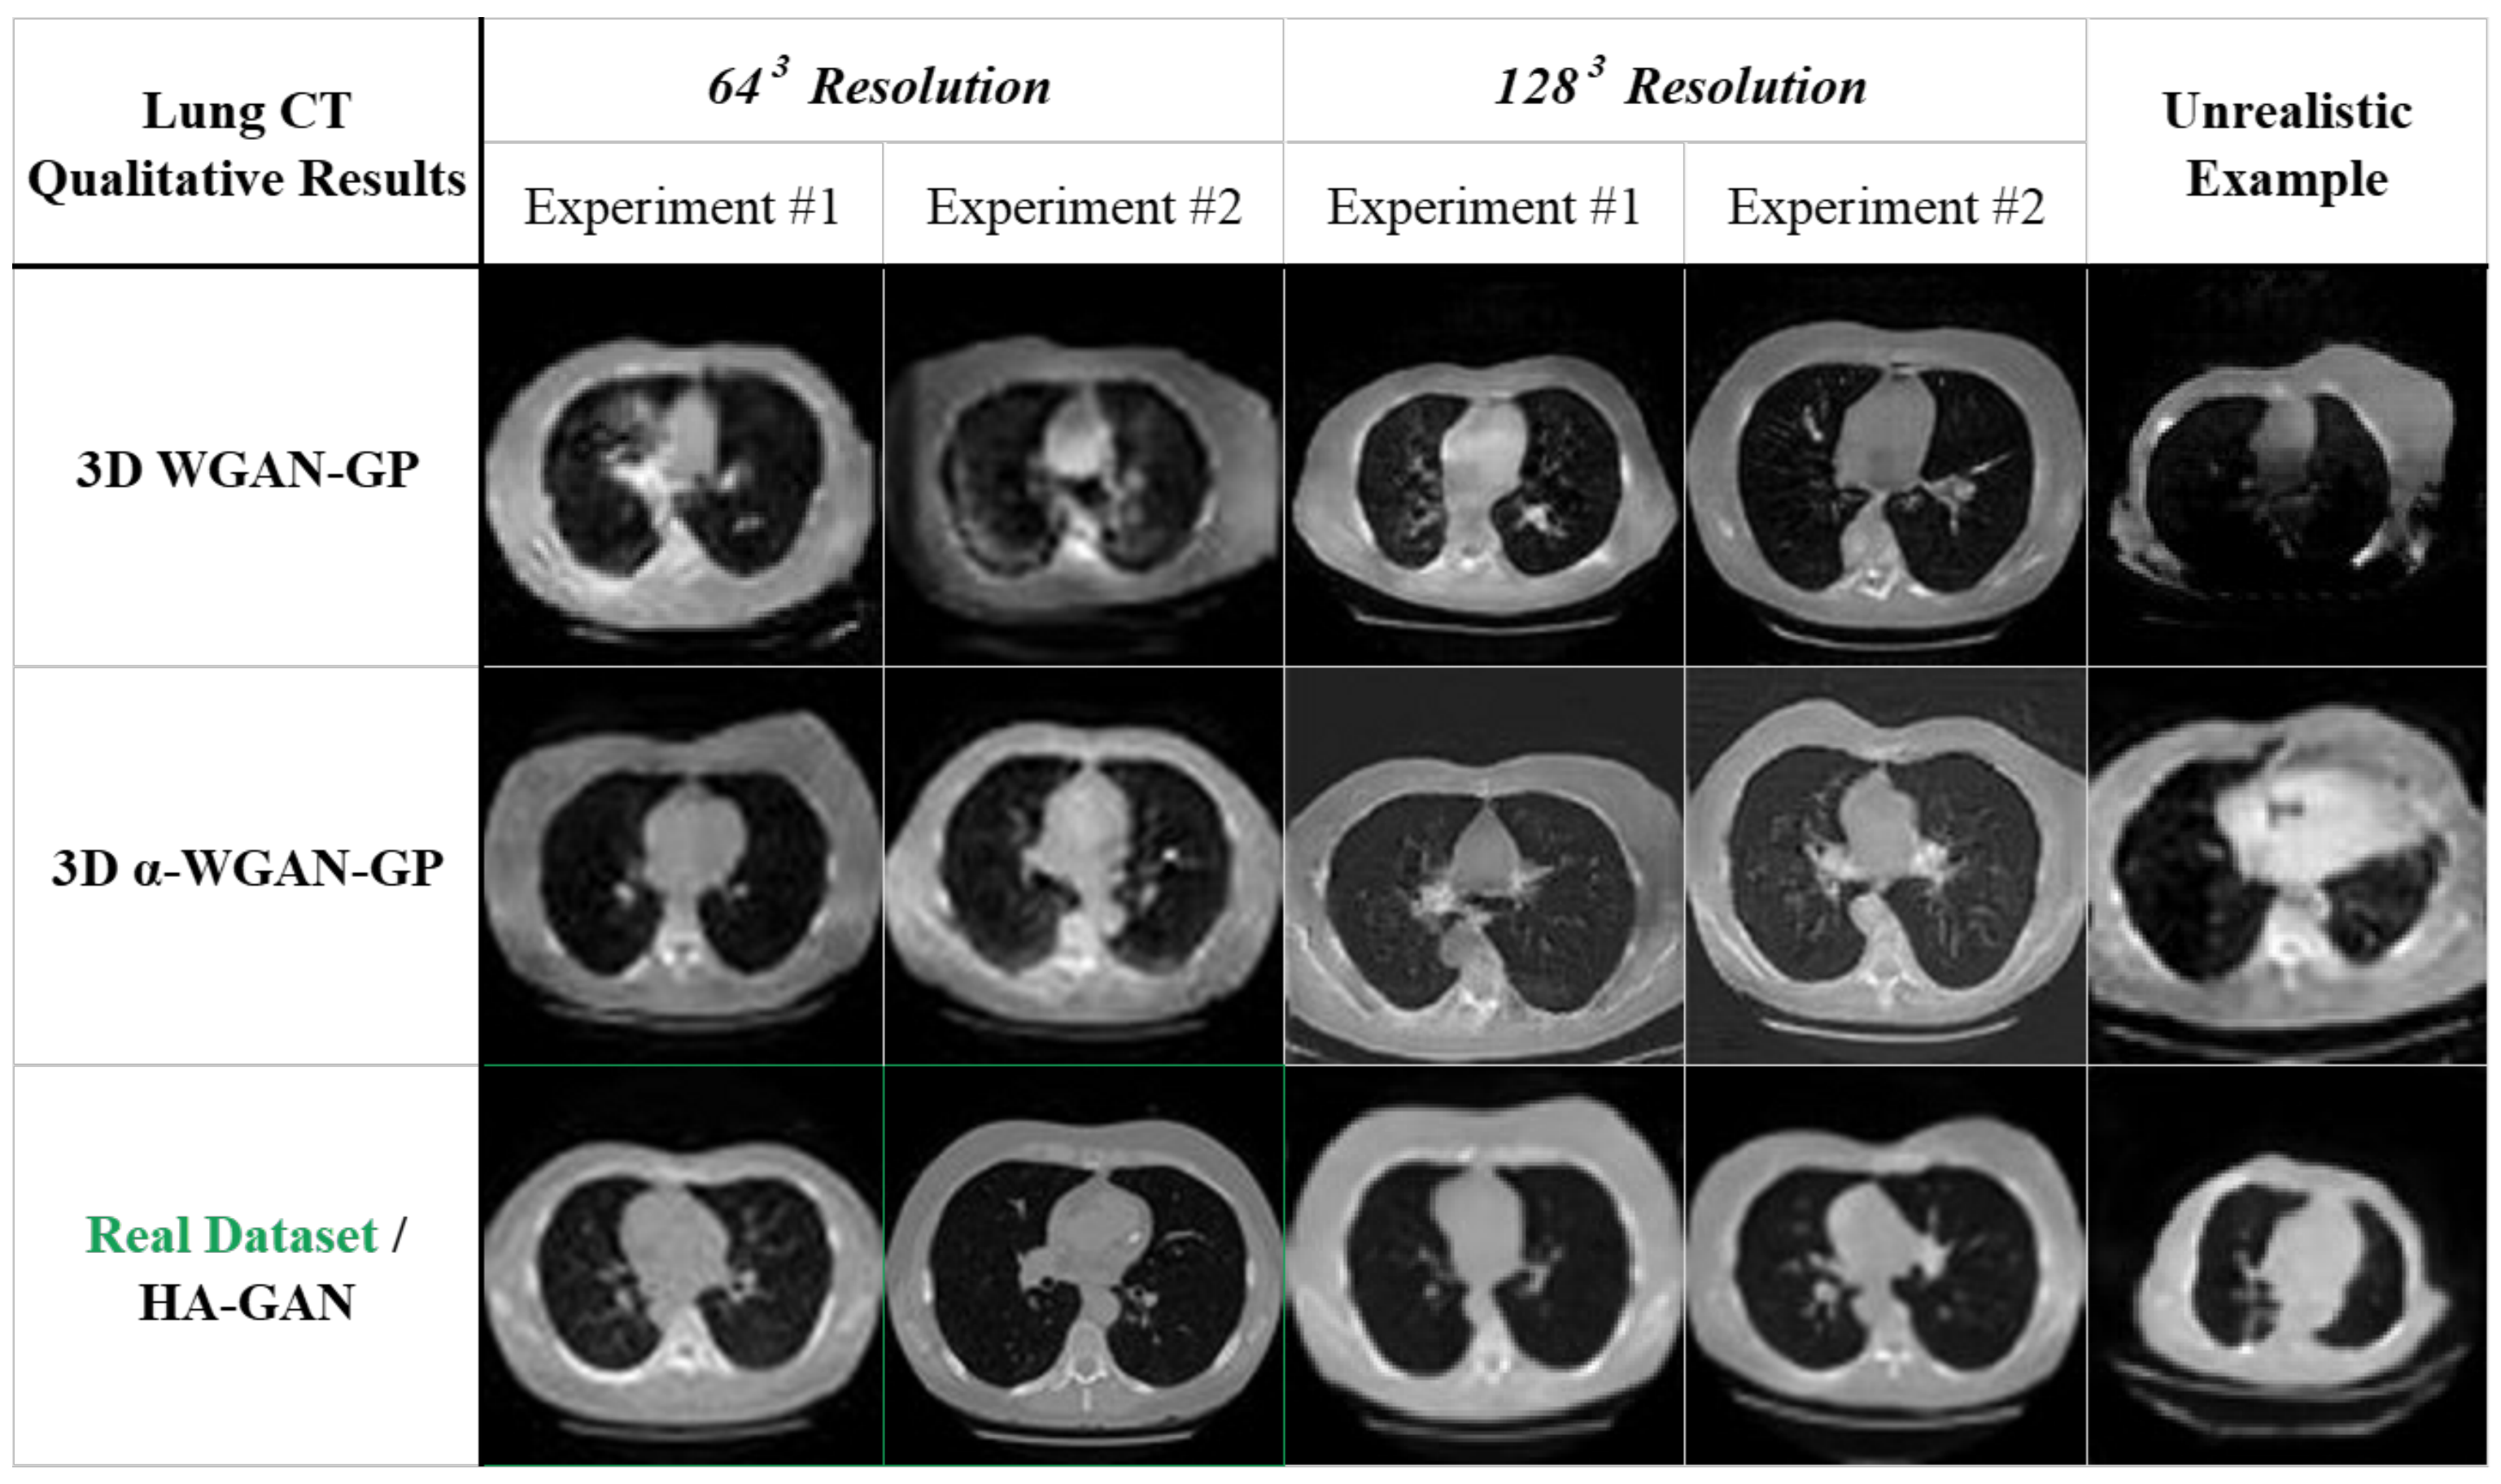

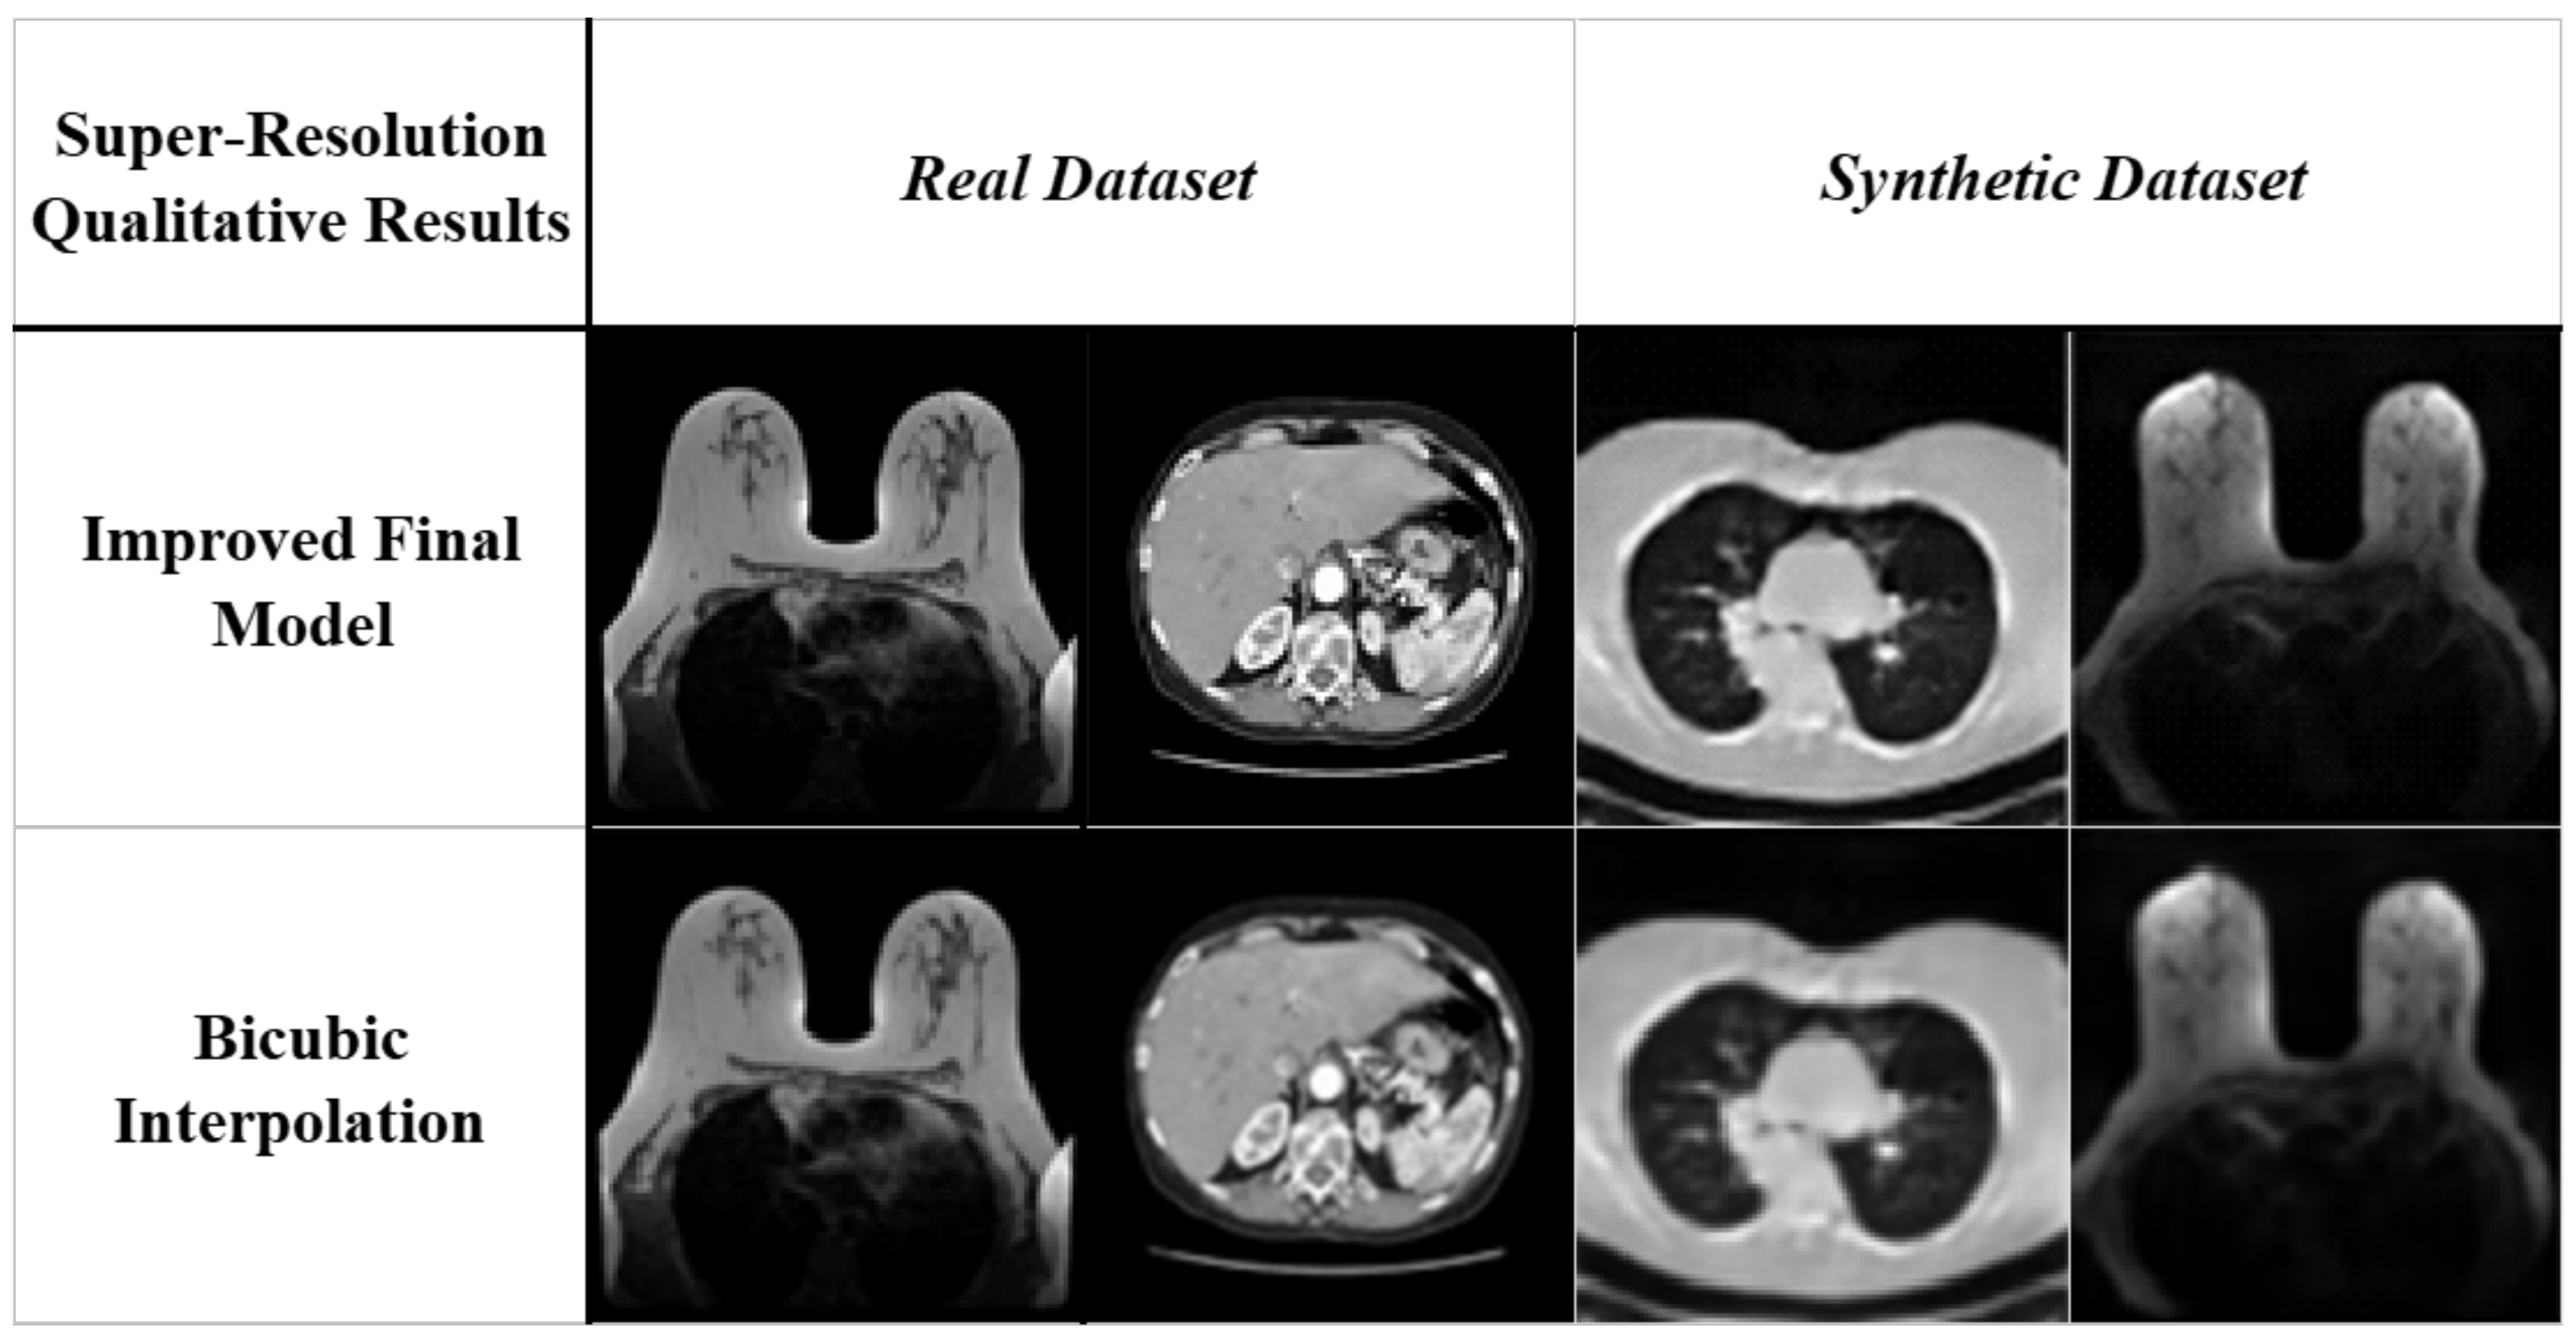

Acknowledgments

The authors would like to acknowledge the support provided by the Centre for Telecommunications and Multimedia of INESC TEC. This work is supported by the European Commission funded PHASE IV AI project (Grant Agreement No. 101095384) under the European Union's Horizon Europe research and innovation programme and co-financed by Component 5 - Capitalization and Business Innovation, integrated in the Resilience Dimension of the Recovery and Resilience Plan within the scope of the Recovery and Resilience Mechanism (MRR) of the European Union (EU), framed in the Next Generation EU, for the period 2021 - 2026, within project HfPT, with reference 41.

References

[1] Samuel G. Armato III et al. Data From LIDC-IDRI, 2015.

[2] Samuel G. Armato III, Cr Meyer, Mf McNitt-Gray, G McLennan, Ap Reeves, By Croft, Lp Clarke, and The RIDER Research Group. The reference image database to evaluate response to therapy in lung cancer (RIDER) project: A resource for the development of change-analysis software. 84(4):448–456.

[3] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron Courville. Improved training of wasserstein GANs.

[4] Ashirbani Saha, Michael R. Harowicz, Lars J. Grimm, Connie E. Kim, Sujata V. Ghate, Ruth Walsh, and Maciej A. Mazurowski. A machine learning approach to radiogenomics of breast cancer: A study of 922 subjects and 529 DCE-MRI features. 119(4):508–516.

[5] Pedro Sousa, Diogo Campas, João Andrade, Pedro Pereira, Tiago Gonçalves, Luís F. Teixeira, Tania Pereira, and Hélder P. Oliveira. Enhancing medical image analysis: A pipeline combining synthetic image generation and super-resolution. In Nuno Gonçalves, Hélder P. Oliveira, and Joan Andreu Sánchez, editors, *Pattern Recognition and Image Analysis*, pages 122–133. Springer Nature Switzerland.

[6] Xintao Wang, Liangbin Xie, Chao Dong, and Ying Shan. Real-ESRGAN: Training real-world blind super-resolution with pure synthetic data.