POLITECNICO MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

EXECUTIVE SUMMARY OF THE THESIS

HTA for the Adoption of AI-Based MRI Reconstruction Software: Enhancement and Image Quality Quantification

TESI MAGISTRALE IN BIOMEDICAL ENGINEERING – INGEGNERIA BIOMEDICA

AUTHOR: SIMONE MUSUMECI

ADVISOR: PROF. GIUSEPPE BASELLI

ACADEMIC YEAR: 2024-2025

1. Introduction and Objective

Deep learning–based MRI reconstruction software has shown significant potential in enhancing image quality while reducing acquisition times. By training neural networks on large high-quality datasets, these systems enable improved signal-to-noise ratio and diagnostic performance when integrated with established reconstruction techniques such as parallel imaging and compressed sensing. Conventional MRI requires extensive k-space sampling to ensure adequate image quality, resulting in prolonged acquisition times. Deep learning–based reconstruction aims to mitigate this limitation by enabling accelerated acquisitions without compromising diagnostic reliability.

This thesis presents a multidomain local Health Technology Assessment (HTA) conducted to support the decision-making process regarding the adoption of the technology at Ospedale di Erba. The assessment was developed according to an integrated framework evaluating clinical effectiveness, safety, economic impact, organizational implications, and regulatory positioning.

The analysis was based on the real annual MRI volumes of the healthcare facility and on data derived from a one-month pivotal trial conducted within the hospital. The objective was to address the policy question of whether the adoption of the technology provides a tangible clinical advantage in terms of image quality and diagnostic performance, as well as a sustainable economic benefit for the institution.

2. Technologies and current clinical gold standard

Currently, accelerated MRI relies primarily on two established strategies: parallel imaging, which leverages the spatial sensitivity profiles of multiple receiver coils, and compressed sensing, which exploits the inherent sparsity of MR images to enable reconstruction from undersampled k-space data. While effective, both approaches face limitations related to signal-to-noise ratio degradation, reconstruction complexity, and the need to balance acceleration with diagnostic reliability.

Deep learning–based reconstruction has emerged as an advanced solution to overcome these constraints. SwiftMR is based on a U-Net–derived architecture trained on pairs of fully sampled k-space data subjected to multiple degradation strategies. These include noise addition, uniform and random undersampling, k-max reduction, elliptical sampling, and partial Fourier acquisition. By learning to recover high-quality images from degraded inputs, the model aims to enable accelerated acquisitions while preserving structural details and diagnostic information.

The architecture is designed to be vendor-agnostic and operates as a standardized reconstruction module. Contextual training strategies allow denoising and sharpness optimization tailored to clinical imaging requirements. The undersampling patterns used during training are aligned with those adopted in conventional parallel imaging and compressed sensing pipelines, enabling recovery of spatial resolution along encoding directions.

From an integration perspective, the system is fully interoperable with existing MRI hardware and hospital IT infrastructure, including DICOM-compliant PACS systems. Reconstruction is performed on a secure cloud platform, ensuring scalability while maintaining compatibility with established clinical workflows.

Validation testing demonstrated effective denoising performance and recovery capability consistent with the training objectives. In addition to technical validation, this work includes a clinical assessment of diagnostic confidence to evaluate the real-world impact of the technology in the hospital setting.

3. Clinical effectiveness and Safety

Clinical validation of the software was analyzed through peer-reviewed evidence comparing deep learning–based reconstruction with standard-of-care MRI techniques. Across studies, deep learning reconstruction demonstrated improved image quality and increased diagnostic confidence when compared with conventional reconstruction methods, either at equal acquisition time or under accelerated protocols.

In neuroimaging, (Oh et al. 2025) reported increased sensitivity in the detection of small brain metastases (<5 mm) on T1-weighted black-blood imaging using deep learning reconstruction. However, potential false positives were associated with unsuppressed vascular signals mimicking metastatic lesions. In vascular imaging, (Kim et al. 2024) showed that deep learning reconstruction significantly enhanced 4D TWIST MR angiography, narrowing the quality gap with time-of-flight imaging while reducing acquisition time and increasing field of view, particularly relevant in emergency settings such as acute ischemic stroke. Residual venous contamination related to dynamic acquisition was occasionally amplified by the reconstruction process.

In epilepsy imaging, (Suh et al. 2024) demonstrated increased diagnostic confidence in detecting temporal and hippocampal lesions. Nonetheless, isolated cases of incoherent signal intensity between adjacent slices were reported, potentially reducing specificity and increasing false positive interpretation.

Regarding accelerated acquisitions, (Lee et al. 2023) for knee MRI and (Yoo et al. 2023) for spine MRI reported scan time reductions of 41% and 32.3%, respectively, with non-inferior diagnostic confidence compared to standard protocols. In both musculoskeletal applications, image sharpness and denoising improvements were observed. However, enhanced sharpness occasionally increased false positive findings, and in some cases amplified residual parallel imaging artifacts.

Importantly, most reported safety-related effects did not reach statistical significance in clinical validation studies. Nevertheless, the correlation between observed imaging artifacts and known deep learning failure modes highlights the importance of integrating technological understanding with clinical evidence when assessing safety.B

ased on literature review and technological analysis of the SwiftMR architecture, five domains of AI-related risk were identified and evaluated within the HTA framework (Table 3.1), with risk levels classified as low, medium, or high depending on potential clinical impact and detectability.

Domains	Definition	Assessment
Domain Shifts	Performance degradation in different acquisition setting	Low-Moderate
Hallucinations	Plausible but not-existing image features	Moderate (task-dependent)
Texture Alteration and oversmoothing	Suppression of fine structural details	Low-Moderate
Silent Failure	Plausible but inaccurate reconstruction	Moderate
Adversarial Perturbation Instability	Sensitivity of small perturbation in undersampled k-space	Low

Table 1: Assessment of the safety domains for SwiftMR.

Overall, the HTA indicates that the integration of deep learning–based reconstruction is supported by consistent evidence of improved image quality and non-inferior diagnostic performance in accelerated acquisition scenarios. When interpreted within a structured safety framework that accounts for known algorithmic limitations and ensures appropriate clinical oversight, the identified risks remain controllable and proportionate to the expected benefits.

From an HTA perspective, the combined assessment of clinical effectiveness and safety domains supports the adoption of the technology in medium-to-high volume clinical settings, provided that governance mechanisms and continuous performance monitoring are implemented.

4. Economic and Organizational Impact

The economic evaluation was conducted using real-world MRI activity data from Ospedale di Erba for the year 2025. A baseline acceleration scenario of 30% was defined as a conservative estimate, reflecting routine clinical and organizational constraints. This assumption was supported by results from a one-month pivotal trial, which demonstrated an average acceleration potential of approximately 40%.

To assess the robustness of the investment under varying productivity conditions, a volume sensitivity analysis was performed across acceleration scenarios ranging from 20% to 40%. The model incorporated the acquisition cost of the software (€71,260), an assumption of 240 working days per year, and the average weighted reimbursement per MRI examination derived from national tariff nomenclature (Sistema Sanitario Nazionale), adjusted according to the hospital's actual examination mix.

Across all simulated scenarios, including the most conservative 20% acceleration case, the investment demonstrated economic sustainability. In the baseline scenario, projected additional annual revenue reached approximately €285,840, with a payback period of three months and a positive return on investment (ROI).

Scenario	ROI	Additional revenue	Payback period
Base case (30%)	401.12%	€285840	3 months
20%	267.42%	€190560	5 months
40%	534.83%	€381120	2 months

Table 2: Volume sensitivity analysis of SwiftMR.

5. Regulatory and Governance Considerations

From a regulatory perspective, the AI Act introduces specific requirements for artificial intelligence systems classified as high-risk, including those embedded in medical devices. AI-based MRI reconstruction software therefore falls within the intersection of the AI Act and the Medical Device Regulation (MDR), requiring full compliance with both regulatory frameworks. In addition, the use of a cloud-based reconstruction platform necessitates adherence to the General Data Protection Regulation (GDPR) to ensure secure data transmission, storage, and processing.

CE marking remains mandatory under MDR, while conformity with the AI Act introduces additional obligations related to risk management, transparency, human oversight, and post-market monitoring. The AI Act entered into force in August 2024, with specific provisions for high-risk AI systems becoming applicable from 2 August 2026. For AI systems integrated into medical devices, an extended transitional period applies until 2 August 2027.Evidence collected during the pivotal trial at Ospedale di Erba, together with findings from the literature, confirms that deep learning reconstruction may amplify patient-related motion artifacts or residual artifacts associated with parallel imaging and compressed sensing when acceleration is pushed to extreme levels. These findings reinforce the regulatory requirement for structured human oversight and continuous post-market surveillance, as explicitly mandated by the AI Act.

6. Conclusions

Overall, the multidomain HTA—integrating technological analysis, clinical validation, safety assessment, economic modeling, and regulatory evaluation—addresses the research question by demonstrating a net benefit associated with the adoption of deep learning–based MRI reconstruction in medium-to-high volume clinical settings.Clinical evidence supports improved image quality and maintained or enhanced diagnostic confidence, while identified safety risks remain proportionate and manageable under appropriate clinical supervision and governance frameworks. Economic modeling, based on real-world hospital data, demonstrates financial sustainability and operational efficiency gains, particularly in routine MRI environments with substantial examination volumes.Within a compliant regulatory framework and supported by structured human oversight, the adoption of the technology is therefore justified from a service-quality, patient-care, and institutional sustainability perspective.

7. Bibliography

References

B. K. Kim, S. H. You, B. Kim and J. H. Shin, “Deep Learning-Based High-Resolution Magnetic Resonance Angiography (MRA) Generation Model for 4D Time-Resolved Angiography with Interleaved Stochastic Trajectories (TWIST) MRA in Fast Stroke Imaging,” *Diagnostics*, vol. 14, no. 11, 2024.

J. Lee, M. Jung, J. Park, S. Kim, Y. Im, N. Lee, H. T. Song and Y. H. Lee, “Highly Accelerated Knee Magnetic Resonance Imaging Using Deep Neural Network (DNN)–Based Reconstruction: Prospective, Multi-Reader, Multi-Vendor Study,” *Scientific Reports*, vol. 13, no. 1, 2023.

G. Oh, S. Paik, S. W. Jo, H. J. Choi, R. E. Yoo and S. H. Choi, “Deep Learning-Based Image Enhancement for Improved Black Blood Imaging in Brain Metastasis,” *European Radiology*, 2025.

P. S. Suh, J. E. Park, Y. H. Roh, S. Kim, M. Jung, Y. S. Koo, S. A. Lee, Y. Choi and H. S. Kim, “Improving Diagnostic Performance of MRI for Temporal Lobe Epilepsy With Deep Learning-Based Image Reconstruction in Patients With Suspected Focal Epilepsy,” *Korean Journal of Radiology*, vol. 25, no. 4, pp. 374–383, 2024.

H. Yoo, R. E. Yoo, S. H. Choi, I. Hwang, J. Y. Lee, J. Y. Seo, S. Y. Koh, K. S. Choi, K. M. Kang and T. J. Yun, “Deep Learning–Based Reconstruction for Acceleration of Lumbar Spine MRI: A Prospective Comparison with Standard MRI,” *European Radiology*, vol. 33, no. 12, pp. 8656–8668, 2023.

8. Acknowledgements

Desidero ringraziare il Professor Giuseppe Baselli per la guida scientifica e il costante supporto durante lo sviluppo della tesi. Un ringraziamento va inoltre all'ingegnere Andrea Parisi per l'opportunità di tirocinio presso l'Ospedale di Erba e per il supporto fornito nella definizione e nello sviluppo del progetto. Infine, ringrazio Giacomo Rallo per il prezioso contributo tecnico e la disponibilità dimostrata lungo tutto il lavoro.