Real-Time Motion Correction in Magnetic Resonance Spectroscopy: AI solution inspired by fundamental science

Alberto Annovi[1], Benedetta Argiento[1,2⋆], Silvia Capuani[3], Matteo Cacioppo[4], Andrea Ciardiello[1,5], Roberto Coccurello[3], Stefano Giagu[1,4], Federico Giove[4,5], Alessandro Lonardo[1], Francesca Lo Cicero[1], Alessandra Maiuro[1,4], Carlo Mancini Terracciano[1], Mario Merola[1,2], Marco Montuori[3], Emilia Nisticò[4], Pierpaolo Perticaroli[1], Biagio Rossi[1], Cristian Rossi[1], Elvira Rossi[1,2], Francesco Simula[4], and Cecilia Voena[1]

1 Istituto di Fisica Nucleare, Italia
2 Università degli Studi di Napoli Federico II, Napoli, Italia
3 Consiglio Nazionale delle Ricerche, ISC, Roma, Italia
4 Sapienza Università di Roma, Roma, Italia
5 Istituto Superiore di Sanità, Roma, Italia
4 Centro Ricerche Enrico Fermi, Roma, Italia
5 Fondazione Santa Lucia IRCCS, Roma, Italia
[
] ⋆ bargient@cern.ch

The 2nd European AI for Fundamental Physics Conference (EuCAIFCon2025)
Cagliari, Sardinia, 16-20 June 2025

Abstract

Magnetic Resonance Spectroscopy (MRS) is a powerful non-invasive tool for metabolic tissue analysis but is often degraded by patient motion, limiting clinical utility. The RECENTRE project (REal-time motion CorrEctioN in magneTic Resonance) presents an AI-driven, real-time motion correction pipeline based on optimized GRU networks, inspired by tagging and fast-trigger algorithms from high-energy physics. Models evaluated on held-out test sets achieve good predictive performance (R^2>0.87) and overall positive framewise displacement (FD) gains. These results demonstrate feasibility for prospective scanner integration; future work will complete in-vivo validation.

Received Date
Accepted Date
Published Date

Contents

1 Introduction

Magnetic Resonance (MR) is a well-established, non-invasive modality for the study of tissue structure and function. Magnetic Resonance Spectroscopy (MRS) [1] extends MR to metabolic and biochemical assessment, with potential clinical applications. However, patient motion remains a critical obstacle: even modest displacements produce spectral distortions, baseline shifts and loss of quantification that reduce reproducibility and limit diagnostic value. The RECENTRE project proposes a real-time motion correction [2] [3] [4] pathway for MRS based on compact deep recurrent networks. The design is inspired by algorithmic techniques from high-energy physics, such as tagging and fast-triggering [5], which enable low-latency decisions on streaming data. Training and evaluation emphasise two complementary goals. First, the model must accurately [6] predict motion-related parameters that can be used for prospective adjustment during the acquisition. Second, the training objective explicitly favours reductions in framewise displacement while preserving spectral fidelity.

2 Neural Network and Framewise Displacement approach

Recurrent neural networks (RNNs) are naturally suited for modeling sequential data, as they capture temporal dependencies across consecutive observations. Among their variants, Gated Recurrent Units (GRUs) provide an effective balance between modeling capacity and computational efficiency, achieving performance comparable to long short-term memory (LSTM) networks with fewer parameters and reduced training time. Within the RECENTRE project, a GRU-based architecture was therefore adopted to predict motion corrections directly from short temporal sequences of MR acquisitions.

Dataset

The data used for training were obtained from the Human Connectome Project (HCP) [7]. A total of 1113 subjects were scanned on Siemens 3T MRI systems with a repetition time (TR) of 720 ms and an echo time (TE) of 33.1 ms. Three acquisition types have been included in this work, differing in the number of frames per sequence: Resting State (1200 frames), Working Memory (316 frames), and Language (405 frames). The motion parameters used as input to the network consist of three translations and three rotations, extracted through rigid realignment during post-processing of the MRI data. As illustrated in Fig. 1, the network input was formed by 7 sequences of two data points each (14 in total), then the model predicts the following 15th point.

Framewise displacement gain

Framewise displacement (FD) was employed as a subject-specific index of motion, providing a scalar measure of head movement at each time point.
FD was computed both from the ground-truth (Eq. 1) and the predicted motion parameters

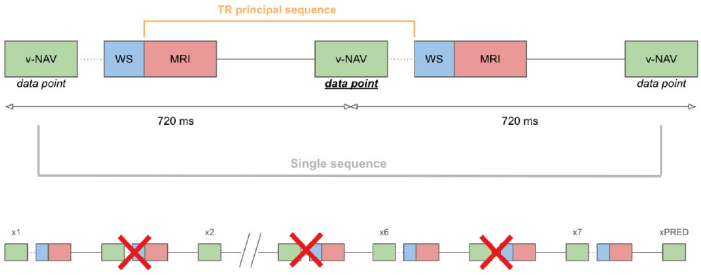

Figure 1: Each input sample consists of 7 sequences, each covering two consecutive acquisition points (for a total of 14 time points). The model is trained to predict the subsequent 15th point.

(Eq. 2).

$$FD_{\text{total}} = \sum_{i=1}^{3} \left| T_{i,t} - T_{i,t+1} \right| + \frac{50\pi}{180} \sum_{i=1}^{3} \left| R_{i,t} - R_{i,t+1} \right| \tag{1}$$

$$FD_{\text{predicted}} = \sum_{i=1}^{3} \left| \hat{T}_{i,t+1} - T_{i,t+1} \right| + \frac{50\pi}{180} \sum_{i=1}^{3} \left| \hat{R}_{i,t+1} - R_{i,t+1} \right| \tag{2}$$

where T_i denote the translational motion parameters (in mm) and R_i the rotational motion parameters (in rad), with the index $i = 1, 2, 3$ corresponding to the three spatial directions. The quantities $\hat{T}_i$ and $\hat{R}_i$ denote the corresponding predicted parameters. The FD gain was then defined as:

$$FD_{\text{gain}} = \frac{FD_{\text{total}} - FD_{\text{predicted}}}{FD_{\text{total}}}. \tag{3}$$

A positive FD_{gain} indicates an improvement, i.e. reduced motion relative to the original sequence.

Neural network model and training objective

The adopted model is composed of a GRU layer with hidden size 128, followed by normalization, non-linear activations and two fully connected layers. The total number of trainable parameters is approximately 270k. The network outputs the predicted motion parameters together with their associated uncertainty estimates. Training was performed with the Adam optimizer and early stopping based on validation performance. The loss function combined a probabilistic negative log-likelihood (NLL) term with a term (βFD_{gain}) that promote positive FD gains and effective motion reduction on motion reduction:

$$\mathcal{L} = \text{NLL} - \beta\, FD_{\text{gain}} \qquad \beta = 0.1 \tag{4}$$

This design enforces both predictive accuracy and effective reduction of head motion. The corresponding training and validation loss curves for the three acquisition types are reported in Fig. 2.

3 Results

Results were obtained on the held-out test set across the three acquisition types (Resting State, Working Memory, and Language), as described in Section 2.
Training curves (Fig. 2) demonstrate stable convergence of the loss across all acquisition types. No significant overfitting was observed.

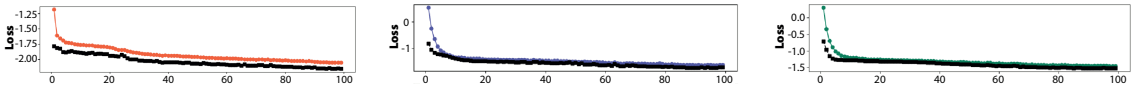

Figure 2: Training and validation loss curves for the three acquisition types: on the left Resting state set, in the middle Working memory set and on the right Language set.

Prediction accuracy was evaluated by comparing the predicted motion parameters with the corresponding ground-truth values, scatter plots show good agreement, as illustrated in Fig. 3. Quantitative performance was assessed using the coefficient of determination (R^2), computed between predicted and ground-truth parameters on the test sets. Across all tasks, R^2 values consistently exceeded 0.87, indicating reliable predictive performance.

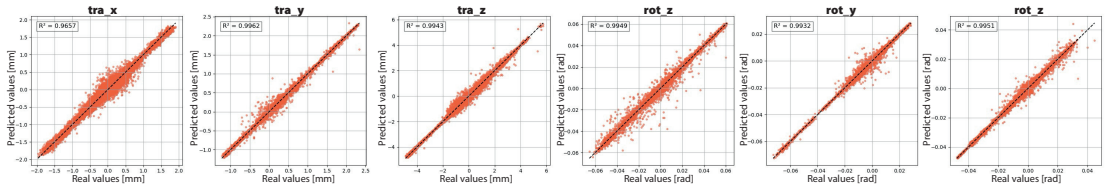

Figure 3: Predicted vs. true motion parameters on the test set for the Resting State acquisition. The figure reports the six estimated motion parameters, namely the three rotational displacements (in rad) and the three rotational angles (in mm).

The effectiveness of the network in reducing apparent motion was quantified using the FD gain metric (see Eq. (3)). Positive FD gains were consistently observed across all acquisition (Fig. 4), demonstrating improved motion estimates. Results were comparable among Resting State, Working Memory, and Language tasks.

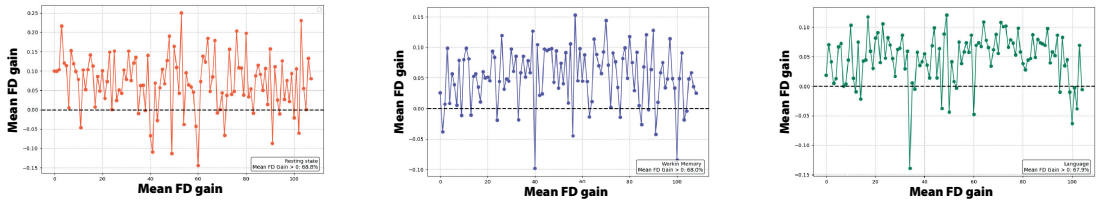

Figure 4: Mean FD gain per patient (average over all predictions). Negative FD gains indicate that the network sometimes slightly overestimates motion.

4 Conclusions

The GRU-based predictor achieves high predictive accuracy ($R^2 > 0.87$) across evaluated dimensions, capturing both large and subtle motion patterns. FD-gain analysis shows mostly positive values, reflecting a net reduction of estimated motion for the majority of patients, with only occasional slight overestimation. The GRU network will be integrated into the Siemens syngo MR MAGNETOM workflow, where the Siemens Image Calculation Environment (ICE) will use the predicted roto-translation parameters for motion-corrected image reconstruction. The model will run on the Siemens Framework for Image Reconstruction Environment (FIRE) with deployment on the Siemens MARS workstation, leveraging on-board GPUs to accelerate prediction. Validation will be performed on in-vivo data acquired with the Siemens 3T Prisma scanner at IRCCS Santa Lucia in Rome.

Acknowledgements

The RECENTRE Project (Prot. P202294JHK) was supported by the PRIN PNRR 2022 program which enabled this research and its presentation.

References

[1] M. G. Saleh, R. A. E. Edden, L. Chang and T. Ernst, *Motion correction in magnetic resonance spectroscopy*, Magn. Reson. Med. **84**, 2312 (2020), doi:10.1002/mrm.28287.

[2] O. C. Andronesi, P. K. Bhattacharyya, W. Bogner and et al., *Motion correction methods for mrs: experts' consensus recommendations*, NMR Biomed. **34**, e4364 (2020), doi:10.1002/nbm.4364.

[3] B. Keating, W. Deng, J. C. Roddey, N. White, A. Dale, V. A. Stenger and T. Ernst, *Prospective motion correction for single-voxel 1h mr spectroscopy*, Magn. Reson. Med. **64**, 672 (2010), doi:10.1002/mrm.22448.

[4] A. Marsman, E. T. Lind, M. Petersen, V. Andersen and V. O. Boer, *Prospective frequency and motion correction for edited 1h magnetic resonance spectroscopy*, Neuroimage **233**, 117922 (2021), doi:10.1016/j.neuroimage.2021.117922.

[5] S. Giagu, *Fast and resource-efficient deep neural network on fpga for the phase-ii level-0 muon barrel trigger of the atlas experiment*, In *EPJ Web Conf.*, vol. 245, p. 01021, doi:10.1051/epjconf/202024501021 (2020).

[6] D. Tamada, *Review: Noise and artifact reduction for mri using deep learning*, arXiv:2002.12889 [physics], Available at: https://arxiv.org/abs/2002.12889 (2020).

[7] W.-M. H. C. P. Consortium, *WU-Minn HCP 1200 Subjects Release: Reference Manual*, Updated April 2018; available at https://humanconnectome.org/storage/app/media/documentation/s1200/HCP_S1200_Release_Reference_Manual.pdf (2018).