KG-MultiMS: a Knowledge Graph-Enhanced Multimodal Framework for Multiple Sclerosis

Riccardo Francia*
Università Degli Studi del Piemonte Orientale, Viale Teresa Michel, 11, Alessandria, Italy
Università Campus Bio-Medico di Roma, Via Álvaro del Portillo, 21, Roma, Italy
riccardo.francia@uniupo.it

Abstract

Despite advances in the computational analysis of multiple sclerosis, the ability to accurately predict disease progression remains limited due to unimodal approaches in the context of deep learning. KG-MultiMS proposes as a multimodal deep-learning framework that integrates data from three heterogeneous sources: (1) MRI volumes using the Swin-UNETR architecture to learn both semantic embeddings and segmentation masks of brain lesions; (2) tabular clinical variables using the FT-Tokenizer architecture to learn clinical embeddings; (3) medical knowledge graph based on the Harvard-Oxford atlas using a Heterogeneous Graph Transformer to learn contextual embeddings from patient-specific subgraphs. A cross-attention mechanism will allow for leveraging multimodal contributions by dynamically weighting them based on the anatomical-functional concordance represented in the contextual embeddings.

1 Introduction

Multiple sclerosis (MS) is a chronic autoimmune disease of the central nervous system that affects 2.8 million people worldwide, with a marked prevalence among women (2:1). MS is characterized by a pathological process affecting nerve fibers (demyelination) that causes motor, sensory, and cognitive dysfunction. It is estimated that there is a delay of approximately 1.5 years between the onset of symptoms and diagnosis.

The scientific literature extensively documents the *clinical-radiological paradox*, whereby patient disability is weakly correlated with the amount of lesions observable via magnetic resonance imaging (MRI). While recent insights suggest this disconnect is driven by lesion topography [Hartmann *et al.*, 2023], this paradox remains a challenge in the computational analysis of MS. Current research, however, relies predominantly on the use of unimodal pipelines with a focus on segmentation tasks (68%) compared to tasks predicting the clinical course of the disease (6.45%) [Belwal and Singh, 2025]. My research questions are:

RQ1. Is it possible to improve upon the predictive capabilities of unimodal architectures by integrating MRI and clinical tabular data into a multimodal deep learning architecture?

RQ2. Is it possible to overcome the clinical-radiological paradox by explicitly incorporating medical knowledge in the form of a knowledge graph (KG)?

2 Background and Related Work

State of the art. Recent studies demonstrate that multimodal approaches have the ability to take advantage of latent synergies (which are not apparent with unimodal approaches) between lesion morphology and clinical progression by integrating heterogeneous data [Al-iedani *et al.*, 2024]. The lack of explicit integration of medical domain knowledge in most approaches limits their ability to resolve the clinical-radiological paradox.

Prior work. My first contribution to research was the development of **AutoML-Med**, a framework that facilitates the automated optimization of preprocessing pipelines for tabular data, with a focus on the clinical context. I validated it on Italian MS patients from the WISDOM project [1], publishing the results at IEEE BIBM 2025 [Francia *et al.*, 2025]. The framework was subsequently validated on a clinical study for mantle cell lymphoma[2]. AutoML-Med has allowed me to understand that unimodal approaches, even when highly optimized, struggle to capture the complexity of diseases like MS.

3 Proposed Architecture: KG-MultiMS

To address the challenges discussed, I propose **KG-MultiMS** (Knowledge Graph-Enhanced Multimodal Framework for Multiple Sclerosis): a multimodal approach that integrates images, clinical data, and structured medical knowledge.

(1) Multimodal feature extraction. To process volumetric MRI data derived from sequences (T1, T2, FLAIR), I propose the use of the Swin-UNETR architecture due to its ability to utilize the three-dimensional nature of volumetric acquisitions by extracting semantic embeddings and segmentation masks

*. R. Francia is a Phd student enrolled in the National PhD in Artificial Intelligence for Health and Life Sciences, Università Campus Bio-Medico, Roma.

1. European Union's Horizon Europe Research and Innovation Actions, grant no. 101137154, https://emsp.org/projects/wisdom

2. S. Ferrero et al., Stud Health Technol Inform, 2026. DOI:10.3233/SHTI260115

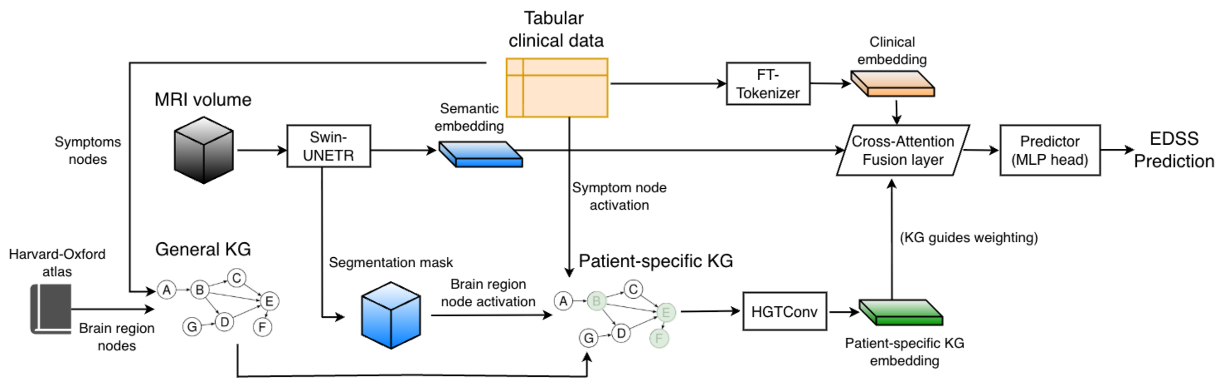

Figure 1: KG-MultiMS architecture overview.

of brain lesions. For the processing of tabular clinical data, I propose the use of the FT-Tokenizer architecture to learn semantic embeddings from clinical variables.

(2) Medical knowledge graph. Building a KG based on the Harvard-Oxford atlas[3] and clinical manifestations recorded in the dataset allows me to encode causal relationships between the location of brain lesions and the symptoms reported by the patient. For each subject, a specific subgraph will be generated by activating the nodes corresponding to the lesioned brain regions via a mapping of the segmentation masks and the nodes corresponding to clinical symptoms using tabular data. The specific subgraphs will be processed by a Heterogeneous Graph Transformer to learn patient-specific contextual embeddings.

(3) Knowledge-guided fusion. To fuse the semantic embeddings extracted from MRI volumes, the clinical embeddings extracted from tabular data, and the contextual embeddings extracted from the patient-specific subgraph, I propose the use of a Cross-Attention-guided Fusion Layer where the contextual embeddings will guide the contribution of imaging and tabular modalities. The resulting fused representation is passed to an MLP prediction head that produces the final EDSS estimate. This mechanism aims to address **RQ2** by explicitly encoding why certain lesion locations influence specific functional domains, linking radiological findings to clinical outcomes.

4 Experimental Design and Validation

I will conduct the model training and validation phase using a publicly available MS brain MRI dataset (Muslim, 2022[4]). Specifically, validation will occur in two phases, each dedicated to addressing the research questions: For **RQ1**, I will compare KG-MultiMS with the respective unimodal baseline counterparts, namely Swin-UNETR, FT-Tokenizer, and the KG encoder, each paired with an independent MLP prediction head. For **RQ2**, I will conduct ablation studies to isolate the KG component, comparing the performance of models with and without cross-attention conditioning. Performance will be evaluated using accuracy, F1 score, and macro-average AUC, with statistical significance assessed via bootstrap confidence intervals and the McNemar test.

3. https://nilearn.github.io/dev/modules/description/harvard_oxford.html

4. https://data.mendeley.com/datasets/8bctsm8jz7/1

5 Expected Impact and Future Prospects

The use of KG-MultiMS in a clinical setting could provide personalized prognostic predictions for high-risk patients. Early diagnosis and optimized treatment can improve patients' quality of life and resource allocation. Concluding model validation by the second year of the PhD program, future prospects include incorporating temporal modeling via longitudinal data to develop KG-MultiMS into a clinical tool for direct practice.

References

[Al-iedani *et al.*, 2024] Oun Al-iedani, Stasson Lea, A. Alshehri, Vicki E. Maltby, Bente Saugbjerg, Saadallah Ramadan, Rodney Lea, and Jeannette Lechner-Scott. Multimodal neuroimaging signatures predict cognitive decline in multiple sclerosis: A 5-year longitudinal study. *Multiple Sclerosis and Related Disorders*, 81:105379, January 2024.

[Belwal and Singh, 2025] Priyanka Belwal and Surendra Singh. Deep learning techniques to detect and analysis of multiple sclerosis through mri: A systematic literature review. *Computers in Biology and Medicine*, 185:109530, February 2025.

[Francia *et al.*, 2025] Riccardo Francia, Giorgio Leonardi, Stefania Montani, Marzio Pennisi, Manuel Striani, Maurizio Leone, and Sandra D'Alfonso. Automl-med: A tool for optimizing pipeline generation in medical ml. In *2025 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 6989–6996, 2025.

[Hartmann *et al.*, 2023] Ana Hartmann, Fabio Noro, Paulo Roberto Valle Bahia, Fabricia Lima Fontes-Dantas, Rodrigo Ferrone Andreiuolo, Fernanda Cristina Rueda Lopes, Valeria Coelho Santa Rita Pereira, Renan Amaral Coutinho, Amanda Dutra de Araujo, Edson Marchiori, and Soniza Vieira Alves-Leon. The clinical-radiological paradox in multiple sclerosis: myth or truth? *Arquivos de Neuro-Psiquiatria*, 81(01):055–061, January 2023.