POLITECNICO MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE E DELL'INFORMAZIONE

Executive Summary of the Thesis

A New Feature Extraction Approach for Temporal Lobe Epilepsy Classification: Integrating Conventional Diffusion and Structural Network Metrics in a Multi-Scale Framework

Laurea Magistrale in Mathematical Engineering - Ingegneria Matematica

Author: Elisa Broserà

Advisor: Prof. Marta Gandolla

Co-advisor: Alfonso Mastropietro

Academic year: 2024-2025

1. Introduction

Temporal Lobe Epilepsy (TLE) is the most common drug-resistant form of epilepsy, primarily originating in mesial temporal structures, specifically the Hippocampus. While traditionally regarded as a focal disorder, TLE has been recently redefined as a large-scale *network disease* characterized by whole-brain alterations in structural connectivity [1]. To investigate these systemic changes, Diffusion MRI (dMRI) has emerged as a fundamental tool. By measuring the random motion of water molecules, dMRI allows for the non-invasive mapping of white matter architecture. Specifically, Diffusion Tensor Imaging (DTI) quantifies the direction and strength of microstructural connectivity, while tractography enables the extraction of whole-brain connectivity matrices, which can be rigorously analyzed using graph theory. Existing literature has extensively explored statistically significant structural differences between TLE patients and healthy controls, using both conventional diffusion metrics and network-based indices. Furthermore, several studies have successfully applied machine learning for TLE classification. However, a significant gap remains: the integration of established statistical findings into the design of machine learning pipelines. Rather than feeding a black box model with raw data, we propose a methodology where the feature space is constructed based on the connectivity signatures of TLE.

Using the *Epilepsy Connectome Project (ECP)* dataset, this thesis aims to optimize TLE classification by leveraging significant differences in structural connectivity. Our approach integrates conventional microstructural diffusion metrics with graph-theoretical metrics to quantify brain reorganization both locally and globally. The novelty of this work lies in the introduction of a *multi-domain and multi-scale framework* for feature extraction, which accounts for both anatomical and functional brain networks, such as the Limbic System [2] and the Default Mode Network (DMN) [4], notably the two most affected by TLE. This structured approach to feature extraction allows an informed feature selection process, bridging the gap between advanced structural connectivity analysis and machine learning classification and providing a competitive performance and interpretable model.

2. Materials and Methodology

2.1. Dataset and Preprocessing

The study utilized Diffusion-Weighted Imaging (DWI) data from the Epilepsy Connectome Project (ECP), including 105 TLE patients and 76 Healthy Controls (HC). The patient cohort exhibits clinical heterogeneity, including left, right, and bilateral TLE and subject with and without Medial Temporal Sclerosis (MTS). The diffusion data were reconstructed using *Generalized Q-Sampling Reconstruction (QSDR)* [6]. This method was specifically chosen to overcome the inherent limitations of the traditional Diffusion Tensor Imaging (DTI) model, particularly in handling crossing fibers within the white matter. QSDR reconstructs the Spin Distribution Function (SDF) directly in the MNI space, ensuring precise voxel-to-voxel alignment across subjects. In addition to conventional DTI metrics, such as Fractional Anisotropy (FA), Mean Diffusivity (MD), and Axial/Radial Diffusivity (AD/RD), advanced QSDR indices were also extracted: Quantitative Anisotropy (QA) to assess directional spin density, Restricted Diffusion Index (RDI) to quantify cellular crowding through water restriction and Quantitative Isotropy Reconstruction (QIR) to measure interstitial free fluid between cells. To build regional measures from voxel-wise data, we utilized *DSI Studio* to compute the mean and standard deviation of these scalars within each Region of Interest (ROI) of the AAL2 atlas. This atlas provides a comprehensive parcellation of 120 cortical and subcortical regions.

2.2. Tractography and Network Analysis

To further investigate structural connectivity, we implemented in *DSI Studio* a deterministic tractography algorithm with a seeding density of 1,000,000 streamlines per subject. By mapping these streamlines onto the anatomical parcellations, we constructed both weighted and binary connectivity matrices. Weighted matrices were built from the number of streamlines between two ROIs, while binary matrices simply indicate the existence of a connection. Once the matrices were obtained, the brain was modeled as a mathematical graph using the *NetworkX* library in *Python*. Using graph theory, we computed a comprehensive set of global and nodal metrics, on weighted and binary graphs and subnetworks. Graph measures of segregation, integration and nodal centrality were considered. Since certain graph properties require a minimum topological complexity to be meaningful, specific size thresholds were implemented. When a graph has a limited number of nodes, only the metrics compatible with its dimension were computed.

2.3. Statistical Analysis

Differences between TLE and HC groups were evaluated using the non-parametric Mann-Whitney U test, as data are not normally distributed. To account for multiple comparisons, p-values were corrected using the False Discovery Rate (FDR). Beyond strict statistical significance ($p<0.05$), identifying statistical trends was considered equally important to characterize TLE.

2.4. Classification Pipeline

To classify TLE patients vs. healthy controls, a supervised machine learning pipeline was developed in *Python*. Its major steps include: feature extraction, feature selection, model training and evaluation, feature importance analysis. Missing values ($<5\%$) were handled via median imputation, while variables exceeding this threshold were excluded to maintain the integrity of the 181-subject dataset. The high-dimensional feature space was filtered through a two-step selection process. First, a *Mann-Whitney U-test* was performed on each feature to retain those with a p-value below a specific threshold α. A subsequent *Spearman correlation* filter was applied; for feature pairs with a correlation coefficient greater than a threshold β, the variable exhibiting the weaker correlation with the target label was discarded. After a *RobustScaler* normalization to mitigate the influence of outliers, a Support Vector Machine (SVM) was trained and optimized via *GridSearchCV*. To prevent data leakage, feature selection and parameter tuning were strictly performed on the training set (85% of the data) using 5-fold cross-validation. Balanced Accuracy was chosen as the primary evaluation metric to account for class imbalance, alongside other standard performance metrics.

3. Multi-Scale Framework

The core methodological innovation of this work lies in the extraction of features across four *hierarchical levels* of cerebral organization, ranging from local tissue properties to large-scale anatomical and functional subsystems. Inspired by the nature of TLE as a complex network disease, this multi-scale framework integrates diverse findings across several brain parcellations and not only anatomical but also functional subsystems.
The analysis begins at Level 1, where structural connectivity is modeled as a global graph with nodes corresponding to the 120 Regions of Interest (ROIs) of the AAL2 atlas. To mitigate the noise inherent in individual ROI analysis, Level 2 groups these regions into ten anatomical macro-regions based on neuroanatomical literature (Figure 1). Connectivity at this level is assessed intra-regionally through anatomical subgraphs; subsequently, super-graphs, using macro-regions as nodes, are built to quantify the interactions between broader anatomical regions.

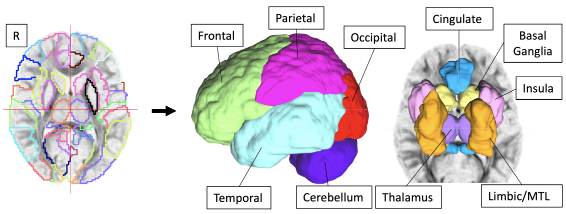

Figure 1: From AAL2 ROI-based parcellation (Level 1) to macro-anatomical regions (Level 2).

Level 3 introduces a functional dimension focusing exclusively on the Default Mode Network (DMN) and the Limbic System (Figure 2). The super-nodes of these functional systems are defined aggregating AAL2 ROIs. Finally, Level 4 adopts an edge-based approach, directly analyzing the macro-edges that connect the anatomical and functional systems previously defined.
For every level of the hierarchy, we compute a dual set of features: the mean and standard deviation of conventional diffusion metrics and a comprehensive set of network-based properties. This structured organization ensures that the subsequent machine learning model is trained on a strategically designed feature space.

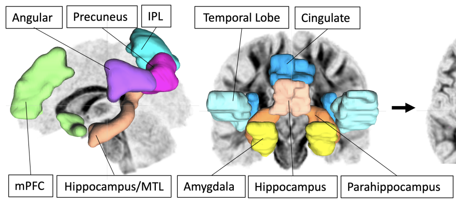

Figure 2: From functional networks, specifically DMN and Limbic Sistem (Level 3) to macro-edge connectivity (Level 4).

The complete architecture of this approach is presented as a simplified illustrative sketch in Figure 3.

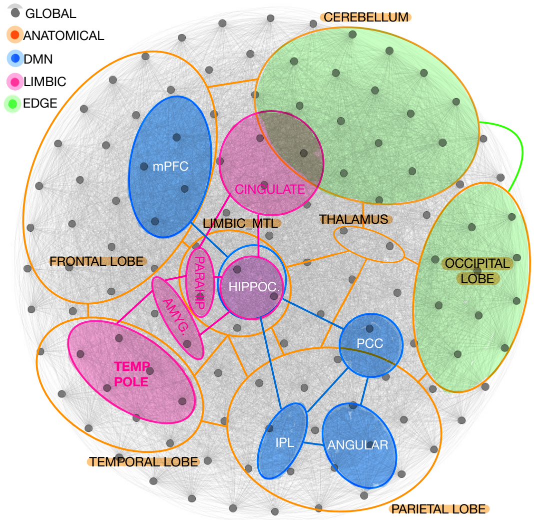

Figure 3: Illustrative sketch of the multi-scale framework. Levels are shown in different colors.

4. Results

4.1. Significant Differences: Conventional Diffusion Metrics

Conventional diffusion analysis revealed widespread microstructural degradation. Significant reductions, after FDR corrections, in Fractional Anisotropy (FA) were identified within the Temporal, Cingulate, and Frontal Lobes of TLE patient, consistent with the findings in [2]. Beyond regional means, the evaluation of structural heterogeneity (variance) revealed significant increases in RDI and QIR within the Hippocampus and Thalamus. This increased heterogeneity likely reflects a complex pathological environment where cell necrosis coexists with the compensatory formation of novel local re-entrant connections.

4.2. Significant Differences: Network-based Metrics

The network-based investigation across the four hierarchical levels showed a significant topological reorganization in TLE patients. At Level 1, patients exhibited a significant (uncorrected) increase in the number of communities, indicating a *fragmentation* of the brain network into isolated modules. This was accompanied by patterns of lower global efficiency and average strength alongside an increased clustering coefficient, depicting a pattern of network *regularization* consistent with the literature. This suggests a pathological shift characterized by a reduction in long-range connections and super-hubs, leading to localized hypersynchronization and inefficient global information transfer.

In Level 2, although FDR-corrected results were limited, a pattern of diffuse degradation was observed in the Frontal Lobe, Occipital Lobe, and Cerebellum. Conversely, intra-regional connectivity within the Limbic and Temporal areas appeared relatively preserved, likely due to neuroplastic reorganization masking focal damage. Furthermore, nodal isolation was observed within the anatomical super-graph, where regions no longer act as effective bridges for information exchange.

Regarding the Default Mode Network (DMN) at Level 3, TLE patients demonstrated a significant decrease in average strength (p_{FDR}=0.019) and density (p_{FDR}=0.049), coupled with increased shortest path lengths (p_{FDR}=0.049), following DeSalvo et al. [4]. A pivotal finding was the *de-hubbing* of the Medial Prefrontal Cortex and Posterior Cingulate/Precuneus, which were downgraded from central super-hubs to peripheral nodes. Regarding the Limbic system, a significant reduction in the standard deviation of node strength was observed. This decrease in heterogeneity further supports the *regularization* of the seizure-prone network.

Finally, the Level 4 edge-based analysis identified a significant reduction in the structural connection between the Hippocampus/MTL and the mPFC (p_{FDR}=0.037). This specific breakdown between encoding centers and executive control hubs provides a quantitative structural basis for the memory loss and attention deficits frequently observed in TLE patients. Collectively, these hierarchical findings confirm that TLE involves a complex *multi-scale reorganization*, spanning over the domain of diffusion metrics and network-based ones.

4.3. Classification

The final stage of this work involved the classification of TLE patients versus healthy controls using the previously defined hierarchical features. We evaluated three distinct scenarios: using exclusively *conventional diffusion metrics*, using only *network-based metrics* and incorporating *all features* within the multi-scale framework.

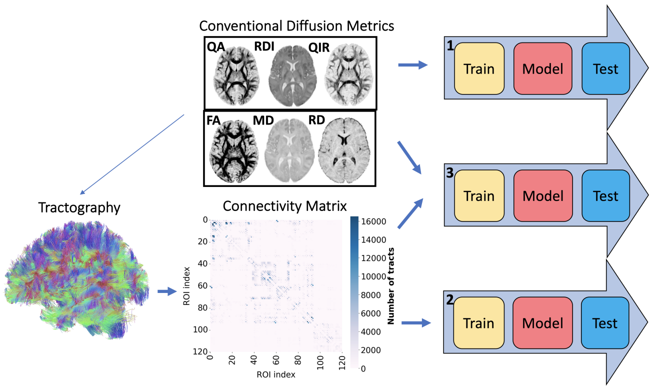

Figure 4: Overview of the workflow: from feature extraction to model classification.

Regarding *conventional diffusion metrics*, the optimal classification pipeline utilized a Mann-Whitney α-threshold of 0.05 and a correlation β-threshold of 0.55 to reduce the high-dimensional feature space to 19 variables. An optimized SVM (kernel=rbf, C=10, γ=0.001) was employed, with 2 features subsequently removed following permutation importance analysis, resulting in a final 17-feature model. The most predictive variables included diffusion indices from the right Orbitofrontal Cortex and the right Inferior Temporal Lobe. Furthermore, the model selected the standard deviation of diffusion metrics within the Hippocampus, Thalamus, Cingulate and Frontal Lobe, alongside features of the functional level such as the mean Fractional Anisotropy (FA) across the Default Mode Network (DMN). Performance metrics for this model are provided in Table 1.

Moving on to *network-based metrics*, the optimized pipeline yielded a model with performance metrics ranging between 62.3% and 70.1 in cross-validation. Specifically, the model

achieved a Balanced Accuracy of 60.4% and a ROC-AUC of 66.7% on the test set. While these network-based metrics contain a predictive signal confirming topological brain reorganization, it is less specific and more subtle. The performance dropped substantially compared to microstructural models.
Lastly, we combined *all features* across both domains and optimized thresholds at 0.05 for Mann-Whitney and 0.6 for correlation, selecting 43 variables from the initial 1500. The optimized SVM parameters were kernel=rbf, C=1, and γ=0.01. Following permutation importance, 9 irrelevant features were dropped, resulting in a refined 34-feature model. Both models performances are in Table 1.

Table 1: Classification performance comparison: CV and Test set results.

Metric	**Set**	**Diffusion Metrics**	**All Features (43 features)**	**All Features (34 features)**
Bal. Acc.	CV	$0.748 \pm .03$	$0.725 \pm .02$	$0.753 \pm .05$
	Test	0.708	0.760	0.729
F1 Score	CV	$0.792 \pm .04$	$0.766 \pm .03$	$0.790 \pm .03$
	Test	0.750	0.833	0.800
Precision	CV	$0.789 \pm .04$	$0.782 \pm .05$	$0.805 \pm .06$
	Test	0.750	0.750	0.736
Recall	CV	$0.807 \pm .10$	$0.762 \pm .09$	$0.786 \pm .07$
	Test	0.750	0.937	0.875
ROC-AUC	CV	$0.791 \pm .04$	$0.743 \pm .03$	$0.844 \pm .05$
	Test	0.791	0.755	0.760
Accuracy	CV	$0.757 \pm .03$	$0.732 \pm .02$	$0.758 \pm .04$
	Test	0.714	0.785	0.750
Specificity	CV	$0.689 \pm .11$	$0.689 \pm .13$	$0.721 \pm .14$
	Test	0.667	0.583	0.583

The most predictive features in this scenario include the diffusion metrics identified in the first scenario, alongside several key network-based features: the Small-worldness of the Occipital network and the Betweenness Centrality of the Parietal weighted network. Slightly less significant contributors include the Closeness Centrality in the Limbic area, the macro-edge between the Frontal and Parietal regions, and the global number of communities in the weighted network. The comprehensive distribution of the selected features, which span all levels and domains of the framework, is in Figure 5.

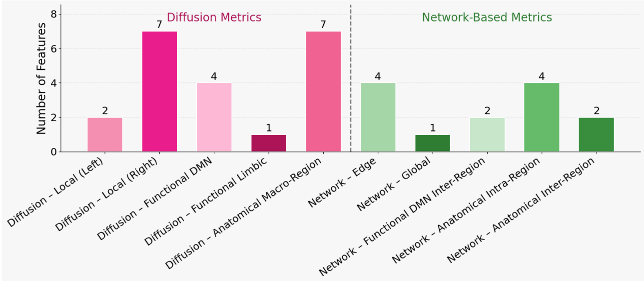

Figure 5: Distribution of selected features across the hierarchical levels and metric domains in the model integrating all features.

5. Model Comparison and Discussion

Regarding the baseline model using only diffusion metrics, the findings are highly encouraging, achieving a mean Balanced Accuracy of 74.8% and a mean ROC-AUC of 79.1% in cross-validation. Although comparing performance with the literature is difficult and often of limited utility given the differences in methods, classification tasks and datasets, our model outperforms the study by Lee et al. [5] in terms of ROC-AUC (AUC = 60.4%) while maintaining a high Recall (80.8%) to address the low sensitivity (64.7%) seen in Cantor-Rivera et al. [3]. The focus on sensitivity is essential in clinical diagnostics, where missing a patient is often far worse than a false positive. These results indicate that the integration of both anatomical and functional levels, along with local microstructural features, may have improved the model's discriminative power.
Moving to the All Features models, the 43-feature scenario showed a reduced standard deviation compared to the diffusion baseline, indicating that network-based features provide more consistent predictions across diverse patient profiles and making it a reliable tool for standardized assessments. However, the reduced 34-feature model represents the performance peak of this framework, achieving a Balanced Accuracy of 75.3%, a ROC-AUC of 84.4%, and an F1-score of 79.0%, with consistent performance on the test set. Although it introduces higher variance, it exhibits lower bias and greater sensitivity to epilepsy-specific patterns. This model should be favored when the primary goal is to maximize class separation between TLE and controls. While results are not easily compara-

ble, a significant strength of our model compared to many works is the achievement of high recall without excessively compromising specificity.
In summary, these findings suggest, first, that integrating conventional diffusion metrics and network-based ones is optimal in terms of performance and interpretability; and second, that higher hierarchical levels, such as anatomical and functional supergraphs, capture information that is not intrinsically incorporated in lower-level metrics. This multi-level approach allows for a clearer understanding of the pathological signal.

6. Conclusions

This thesis provides a comprehensive characterization of structural connectivity alterations in Temporal Lobe Epilepsy (TLE) through a novel four-level hierarchical framework. The statistical analysis confirmed that TLE is a large-scale network pathology characterized by *systemic reorganization*, where long-range connections are damaged in favor of local re-entrant ones. The hierarchical analysis revealed significant network fragmentation, patterns of isolation across macro-regions and widespread intra-connectivity degradation. Crucially, the Default Mode Network (DMN) and the Limbic System were found to be directly impacted by TLE.
From a classification perspective, the results confirmed that integrating conventional diffusion metrics with network-based indices is optimal, with the *multi-scale framework* providing a different approach for feature extraction. Despite potential limitations, such as the exclusion of TLE lesion laterality, this work successfully bridges the gap between advanced structural connectivity analysis and machine learning classification, paving the way for future developments such as multi-modal approaches and the exploration of state-of-the-art Graph Neural Networks on larger datasets.

References

[1] B. C. Bernhardt, L. Bonilha, and D. W. Gross. Network analysis for a network disorder: the emerging role of graph theory in the study of epilepsy. *Epilepsy Behav.*, 50:162–170, 2015.

[2] L. Bonilha et al. Medial temporal lobe epilepsy is associated with neuronal fibre loss and paradoxical increase in structural connectivity of limbic structures. *J. Neurol. Neurosurg. Psychiatry*, 83:903–909, 2012.

[3] D. Cantor-Rivera et al. Detection of temporal lobe epilepsy using support vector machines in multi-parametric quantitative MR imaging. *Comput. Med. Imaging Graph.*, 46:158–168, 2015.

[4] M. N. DeSalvo et al. Altered structural connectome in temporal lobe epilepsy. *Radiology*, 270:842–848, 2014.

[5] D. A. Lee, H.-J. Lee, and K. M. Park. Identification of focal epilepsy by diffusion tensor imaging using machine learning. *Scientific Reports*, 10(1):1–9, 2020.

[6] F.-C. Yeh and W.-Y. I. Tseng. NTU-90: A high-resolution white matter atlas with Q-space diffeomorphic reconstruction. *NeuroImage*, 58:91–99, 2011.