Neural Circuit Mapping and Integration Analysis:

A Graph Neural Network Framework for Decoding Stem Cell Integration in Epileptic Circuits

Daniel KALU

Cortex Labs

Abstract

Drug-resistant epilepsy (DRE) affects approximately 30% of the global epilepsy population of 50 million, remaining refractory to all available pharmacological and surgical interventions. Regenerative stem cell therapy — in particular, the transplantation of GABAergic inhibitory interneurons to restore the excitatory-inhibitory (E/I) balance disrupted in mesial temporal lobe epilepsy (MTLE) — represents a paradigm shift from symptomatic management toward genuine circuit repair. However, decoding whether and how transplanted cells form functional synaptic connections within pathologically remodeled host circuits remains the central unresolved challenge in translating this approach to broad clinical use. This paper presents a formal framework for AI-driven neural circuit mapping and integration analysis, centered on Graph Neural Networks (GNNs) as the primary computational architecture. We systematically review the multimodal data substrates required — spanning scalp and intracranial electrophysiology (EEG/iEEG/SEEG), volume electron microscopy, Connectome-seq, and two-photon live imaging — and analyze how their complementary spatiotemporal resolutions can be unified into structured graph representations amenable to GNN learning. We introduce a suite of mathematical formulations covering graph construction from neural data, a dynamic rewiring evolution model, a multimodal integration loss function, and a composite Integration Success Score (ISS), providing a quantitative scaffold for experimental and trial design. We further survey GNN architectures — including graph convolutional, dynamic, attention-based, and heterogeneous variants — and evaluate their suitability for modeling post-transplantation circuit reorganization at both synaptic and network scales. A comprehensive multimodal biomarker framework spanning electrophysiological, imaging, and molecular indicators is presented, along with a structured analysis of translational barriers, including data harmonization, model interpretability, computational scalability, and the preclinical-to-clinical gap. Strategic recommendations covering data acquisition protocols, explainable AI integration, longitudinal biomarker validation, and open-science practices are provided. This framework is intended to serve as both a research roadmap and an evidence base for applying AI-driven stem cell integration analysis to DRE, supporting the case for targeted research investment and collaborative infrastructure development.

Keywords: *drug-resistant epilepsy · stem cell therapy · GABAergic interneurons · graph neural networks · neural circuit mapping · multimodal data integration · synaptic biomarkers · explainable AI · circuit rewiring · functional integration*

1 Introduction

1.1 Clinical Imperative: The Limits of Symptomatic Management

Epilepsy is a chronic neurological disorder characterized by recurrent, unprovoked seizures arising from excessive or hypersynchronous neuronal activity. Its clinical burden extends far beyond seizure frequency: patients face elevated rates of cognitive decline, psychiatric comorbidity, social marginalization, and premature mortality [1]. Despite the availability of more than twenty licensed anti-seizure medications (ASMs), approximately 30% of patients meet criteria for drug-resistant epilepsy (DRE) — defined as failure of two appropriately trialed ASMs — and continue to experience disabling seizures with no adequate pharmacological recourse [2].

Surgical options, including temporal lobectomy and laser interstitial thermal therapy, offer seizure freedom to a subset of structurally eligible candidates but are inherently destructive and irreversible, carrying risks of permanent cognitive deficit, including memory impairment and language disruption [3]. Neuromodulatory approaches such as deep brain stimulation (DBS) and responsive neurostimulation (RNS) provide an additional therapeutic tier yet rarely achieve complete seizure freedom [3]. The common denominator of these interventions is their symptomatic orientation: they suppress seizure expression without addressing the underlying cellular and circuit-level pathology that sustains epileptogenicity.

1.2 The Circuit Repair Hypothesis

Mesial temporal lobe epilepsy (MTLE), the most prevalent and pharmacoresistant adult epilepsy syndrome, is characterized at the cellular level by selective depletion of inhibitory GABAergic interneurons — particularly parvalbumin (PV)- and somatostatin (SOM)-expressing subtypes derived from the medial ganglionic eminence (MGE) — from hippocampal subfields [3]. This loss disrupts the E/I balance of hippocampal circuits, generating the sustained hyperexcitability and pathological hypersynchrony that drives seizure propagation and network remodeling. Dysregulation of axon guidance molecules (e.g., semaphorins, ephrins) and neuroinflammatory cascades further compound circuit disruption, promoting aberrant sprouting and maladaptive plasticity [1].

The circuit repair hypothesis posits that if functionally appropriate inhibitory interneurons can be transplanted into the epileptic focus, engage with host excitatory neurons through genuine inhibitory synapses, and persist within the host tissue, they will restore E/I balance and attenuate seizure generation. This represents a mechanistically targeted intervention — replacing a precisely identified cellular deficit — rather than a pharmacological proxy. Pioneering first-in-human trials, including intrahippocampal transplantation of hPSC-derived GABAergic interneurons (NRTX-1001), have reported promising early safety and efficacy signals, lending clinical credibility to this hypothesis [4].

1.3 The Integration Decoding Problem

The central unresolved challenge in advancing cell therapy for epilepsy from proof-of-concept to reliable clinical practice is the integration decoding problem: given a transplanted cell population within a living, dynamically evolving neural circuit, how do we determine — with specificity, sensitivity, and sufficient temporal resolution — whether and to what degree those cells have established functional inhibitory connections, and how those connections are reshaping network-level seizure dynamics? This problem is

inherently multimodal, multi-scale, and longitudinal. It cannot be resolved by any single experimental modality, and the data it generates — graphs of neuronal connectivity spanning nanometre-scale synaptic ultrastructure to centimeter-scale network topology, evolving over weeks to months — requires computational approaches that can natively represent and learn from complex relational structures.

Graph Neural Networks (GNNs) are precisely such an approach. Their capacity to represent entities as nodes and relationships as edges, and to learn hierarchical relational features through iterative neighborhood aggregation, makes them ideally suited for the analysis of neural circuits — both in their healthy and pathological states, and in their dynamic reorganization following therapeutic intervention [22,26].

1.4 Scope and Contributions

This paper makes the following contributions to the field:

- A structured technical framework for applying GNNs to the problem of post-transplantation neural circuit mapping, grounded in the biological specifics of GABAergic interneuron integration in MTLE.
- A formal mathematical treatment of graph construction, dynamic rewiring modeling, multimodal loss formulation, and composite integration scoring.
- A systematic review and comparison of multimodal data acquisition techniques and GNN architectures relevant to this application domain.
- A comprehensive multimodal biomarker taxonomy for assessing stem cell integration success, spanning electrophysiological, imaging, and molecular dimensions.
- A candid analysis of translational barriers and a prioritized strategic research agenda to guide future investment and collaboration.

---

2 Background and Related Work

2.1 Neural Circuit Pathology in MTLE

The pathological substrate of MTLE is well-characterized at multiple scales. Macroscopically, hippocampal sclerosis — marked by pyramidal cell loss and reactive gliosis — is the most common structural finding in surgically resected tissue. At the circuit level, loss of inhibitory interneurons generates a net excitatory bias, while pathological axonal sprouting of dentate granule cell mossy fibers creates recurrent excitatory loops that amplify and sustain seizure discharges [1]. Neuroinflammatory processes, including microglial activation and upregulation of complement and phagocytic pathway genes (notably TREM2), further remodel synaptic connectivity in ways that can be both pro- and anti-epileptic depending on context and timing [11].

These circuit-level abnormalities are not static: epilepsy is a progressive condition in which repeated seizures induce further network remodeling through activity-dependent plasticity mechanisms. Any AI framework for monitoring post-transplantation circuit dynamics must therefore account for this baseline trajectory of pathological evolution, distinguishing graft-induced network changes from the disease's natural history.

2.2 Stem Cell Therapy: Mechanisms Relevant to Circuit Mapping

Stem cell types under investigation for epilepsy include embryonic stem cells (ESCs), iPSCs, mesenchymal stem cells (MSCs), neural stem cells (NSCs), and adipose-derived regenerative cells (ADRCs) [2]. For the

circuit repair strategy most directly amenable to GNN-based integration analysis, the relevant population is MGE-derived GABAergic interneuron progenitors, which have been shown in preclinical studies to: differentiate into PV- and SOM-positive interneuron subtypes; maintain mature electrophysiological properties, including fast-spiking and regular-spiking firing patterns; receive both excitatory and inhibitory synaptic input from host neurons; and form inhibitory synapses on host pyramidal cells [6].

A critical recent finding from multimodal synapse analysis is that transplanted neurons can exhibit signs of immature or incomplete integration even months post-transplantation — including spines without synapses and shaft synapses that are not inhibitory — and that host TREM2-mediated microglial pruning significantly restricts the brain-wide input connectome of transplanted neurons [11]. This underscores that integration is not a binary event but a graded, time-dependent, microenvironment-sensitive process that must be tracked longitudinally and at multiple biological scales simultaneously. MSC transplantation operates through a distinct, primarily paracrine mechanism — anti-inflammatory signaling, trophic factor secretion (BDNF, GDNF, IGF-1), and promotion of endogenous neurogenesis — rather than direct circuit integration, and presents a different biomarker profile [2].

2.3 Graph Neural Networks in Neuroscience: Prior Work

GNN applications in neuroscience have expanded substantially over the past five years. Core use cases include brain state classification from functional MRI connectivity matrices, multimodal data integration combining structural and functional connectivity, disease biomarker discovery in conditions including Alzheimer's disease and schizophrenia, and epilepsy surgery planning [22,26]. Of particular relevance, topology-aware GNNs that integrate intracranial EEG features, electrode topology, and structural MRI have been shown to outperform traditional neural networks in automated epilepsy surgery planning, including identifying alternative therapeutic targets beyond the primary ictal onset zone [13].

Dynamic GNN architectures that incorporate temporal edges to model time-variant network activity have demonstrated the ability to capture brain developmental trajectories and disease progression dynamics [22]. This temporal modeling capacity is directly applicable to the post-transplantation integration timeline, where circuit reorganization unfolds over weeks to months as grafted cells mature and form synapses. The present paper extends this literature by providing the first systematic framework specifically oriented toward GNN-based decoding of stem cell integration in the epileptic brain.

---

3 Methodology

3.1 Literature Review Protocol

A systematic narrative review was conducted across PubMed, Scopus, IEEE Xplore, bioRxiv, and arXiv preprint servers, with searches executed through May 2025. Search terms combined MeSH headings and free-text variants of: (epilepsy OR seizure) AND ("stem cell" OR "interneuron transplant" OR "GABAergic") AND ("graph neural network" OR "circuit mapping" OR "synaptic integration" OR "neural circuit") AND ("electrophysiology" OR "EEG" OR "microscopy" OR "connectome"). Supplementary targeted searches were performed for specific named techniques (Connectome-seq, volume EM, iEEG/SEEG multimodal profiling), GNN frameworks (PyTorch Geometric, DGL, NVIDIA GNN), and open datasets (OpenNeuro, EBRAINS,

NWB). Empirical studies, methodological papers, and authoritative reviews were included; purely theoretical contributions were selectively included where they provided essential conceptual scaffolding.

3.2 Framework Development

The mathematical framework presented in Section 4 was developed de novo by the authors to provide a unifying quantitative language for the GNN-based integration analysis pipeline. Formulations were derived from first principles of graph theory, neural network learning theory, and established biophysical models of neuronal dynamics (including the FitzHugh-Nagumo model as a theoretical substrate for seizure transition modeling [31]). All equations are proposed as analytical constructs for future experimental validation and trial design rather than empirically fitted models.

3.3 Scope Limitations

This paper does not report primary experimental data or conduct a formal meta-analysis. The framework is developed primarily with reference to MTLE and GABAergic interneuron transplantation; its applicability to other epilepsy syndromes or stem cell types (e.g., MSCs operating via paracrine mechanisms) may require adaptation. Clinical recommendations are not made. Readers should consult ILAE consensus guidelines for clinical management decisions.

---

4 Mathematical Framework

We propose a formal mathematical scaffold for GNN-based neural circuit integration analysis, encompassing graph construction, dynamic rewiring modeling, multimodal learning, and composite integration scoring.

4.1 Graph Construction from Multimodal Neural Data

Let G(t) = (V, E(t), X_V, X_E(t)) denote the neural circuit graph at time t post-transplantation, where V is the set of nodes (representing individual neurons or electrode-defined brain regions), E(t) is the set of directed edges (synaptic or functional connections), X_V ∈ R^(|V| × d_v) is the matrix of node feature vectors (encoding cell type, molecular identity, firing properties, or spatial coordinates), and X_E(t) ∈ R^(|E| × d_e) is the matrix of edge feature vectors (encoding connection weight, synapse type — excitatory/inhibitory — ultrastructural morphology, or functional coherence). Node features are populated from multimodal data: electrophysiological recordings (firing rate, action potential shape, oscillatory phase) from patch-clamp or iEEG; molecular identity (GABA expression, PV/SOM subtype markers) from immunostaining or Connectome-seq; and morphological properties (soma size, spine density) from two-photon or volume EM imaging. Edge features are derived from functional connectivity metrics (coherence, Granger causality, phase-locking values) and structural connectivity measures (synapse density from volume EM).

```
G(t) = (V, E(t), X_V, X_E(t))  where  X_V from {EP,          (Eq. 1)
mol, morph},  X_E(t) from {FC, SC}
```

4.2 Dynamic Circuit Rewiring Model

Post-transplantation circuit rewiring is modeled as a temporally evolving graph process. Let A(t) ∈ {0,1}^(|V|×|V|) denote the adjacency matrix at time t. The temporal evolution of connectivity is governed by:

$$A(t+\Delta t) = A(t) + \Delta A_{graft}(t) + \Delta A_{plast}(t) - \Delta A_{prune}(t) \quad \text{(Eq. 2)}$$

where ΔA_graft(t) represents new edges formed by grafted interneurons establishing inhibitory synapses on host excitatory neurons (informed by volume EM and Connectome-seq data), ΔA_plast(t) captures activity-dependent synaptic potentiation and depression in existing host-host connections (informed by iEEG functional connectivity changes), and ΔA_prune(t) represents synapse elimination — including pathological over-pruning mediated by TREM2-activated microglia [11]. Dynamic GNNs operating on the sequence {G(t_0), G(t_1), ..., G(t_T)} learn to predict A(t_{T+1}), enabling forecasting of circuit-state trajectories and the identification of rewiring patterns predictive of therapeutic success.

4.3 Multimodal GNN Learning Objective

The GNN is trained to simultaneously solve multiple biologically motivated prediction tasks. The composite loss function L is:

$$L = \lambda_1 \cdot L_{seizure} + \lambda_2 \cdot L_{connect} + \lambda_3 \cdot L_{subtype} + \lambda_4 \cdot L_{inflamn} \quad \text{(Eq. 3)}$$

where L_seizure is the seizure prediction loss (binary cross-entropy on seizure/non-seizure EEG windows), L_connect is the link prediction loss (predicting which new inhibitory synapses will form based on molecular compatibility of node pairs), L_subtype is the node classification loss (predicting PV vs. SOM vs. host cell identity from node features), and L_inflamn is a regression loss on TREM2 expression level as a predictor of pruning-mediated integration restriction. The weighting coefficients λ_1...λ_4 ≥ 0 (with Σλ_i = 1) are tuned via cross-validation on held-out patient data. Multi-task learning in this formulation encourages the GNN to learn shared representations across modalities that generalize to biological scales and time points.

4.4 Integration Success Score (ISS)

To operationalize the qualitative concept of 'successful integration' as a single trackable endpoint for clinical trial design, we define the Integration Success Score (ISS) as a weighted composite of four normalized indicator domains:

$$ISS(t) = \alpha \cdot \Phi_{EP}(t) + \beta \cdot \Phi_{struct}(t) + \gamma \cdot \Phi_{mol}(t) + \delta \cdot \Phi_{net}(t) \quad \text{(Eq. 4)}$$

where Φ_EP(t) ∈ [0,1] is the electrophysiological domain score (normalised reduction in IED rate, seizure frequency, and restoration of E/I balance metrics), Φ_struct(t) ∈ [0,1] is the structural domain score (normalised inhibitory synapse density formed by grafted cells per unit volume), Φ_mol(t) ∈ [0,1] is the molecular domain score (normalised GABAergic marker expression — PV, SOM, GABA — in grafted cell population), and Φ_net(t) ∈ [0,1] is the network topology domain score (GNN-derived change in graph

connectivity metrics including clustering coefficient, path length, and synchronisability relative to pre-transplantation baseline). The weights α, β, γ, δ (summing to 1) are to be determined empirically from clinical trial data. An ISS threshold (e.g., ISS ≥ 0.6 at 6 months) may serve as a composite primary endpoint in future adaptive trial designs, replacing or supplementing seizure frequency reduction alone.

4.5 GNN Message Passing Formalization

At each layer l of the GNN, the hidden representation h_v^(l) of node v is updated by aggregating messages from its neighborhood N(v):

```
h_v^(l+1) = UPDATE( h_v^(l), AGGREGATE({ MSG(h_v^(l),          (Eq. 5)
h_u^(l), e_vu) : u ∈ N(v) }) )
```

where MSG(·) is a learnable message function (e.g., edge-conditioned convolution incorporating edge feature e_vu such as synapse type or functional coherence weight), AGGREGATE(·) is a permutation-invariant pooling function (sum, mean, or attention-weighted), and UPDATE(·) is a GRU-based state update for dynamic GNNs or an MLP for static variants. For heterogeneous GNNs processing multiple node types (transplanted interneurons, host pyramidal cells, host interneurons, electrode nodes), separate MSG functions are learned for each edge type, enabling type-specific relational reasoning essential for correctly modeling inhibitory versus excitatory circuit contributions.

---

5 Multimodal Data Acquisition for Graph Construction

5.1 Electrophysiology: Functional Dynamics from Population to Synapse

Electrophysiological methods provide the functional backbone of the circuit graph. Scalp EEG captures population-level electrical dynamics with millisecond temporal resolution and serves as the primary clinical seizure-monitoring signal, but its low spatial resolution (regional at best) limits its utility for localized integration mapping [1]. EEG source localization (ESL) algorithms improve cortical source estimation by inverting scalp signals, partially mitigating this limitation.

Intracranial EEG (iEEG), and in particular stereoelectroencephalography (SEEG), provides direct measurement of local field potentials and single-unit spiking from discrete hippocampal subfields with both high temporal and spatial resolution [12]. SEEG is the modality of choice for seizure onset zone localization in presurgical evaluation and has the important additional capability — relevant to this project — of enabling simultaneous RNA, DNA, and protein sampling from electrode sites via the MoPEDE framework, directly linking molecular profiles to recorded neurophysiology at the same spatial location [12]. At the single-cell level, patch-clamp electrophysiology provides the gold standard for characterizing intrinsic membrane properties, synaptic currents, and firing patterns of individual grafted interneurons, confirming functional maturation [6].

5.2 High-Resolution Microscopy: Structural and Molecular Circuit Mapping

Volume electron microscopy (volume EM) provides ultrastructural 3D reconstruction of individual synapses at nanometre resolution, enabling classification of synapse type (excitatory versus inhibitory based on postsynaptic density morphology), measurement of synapse density and spatial distribution, and identification of immature or malformed synaptic junctions — all critical for evaluating the quality of graft-derived inhibitory connections [11,15]. While technically demanding and limited in tissue volume, volume EM provides the ground truth structural data against which other biomarkers must be validated.

Connectome-seq represents a breakthrough capability for this project: by combining engineered synaptic proteins with RNA barcoding and parallel single-nucleus and single-synaptosome sequencing, it maps neuronal connectivity at single-synapse resolution while simultaneously capturing the molecular identities of pre- and post-synaptic neurons across large brain volumes [16]. This directly addresses the critical need to determine not only whether grafted cells form synapses, but also whether those synapses are molecularly appropriate — i.e., inhibitory contacts on host excitatory pyramidal cells rather than nonfunctional or aberrant connections.

Two-photon microscopy enables live imaging of thousands of neurons simultaneously, tracking morphological dynamics (dendritic spine turnover, axonal growth), calcium transients reflecting neuronal activity, and population-level coding properties over time [17]. When combined with spatial light modulators (SLMs) or fiber photometry for simultaneous optogenetic stimulation and recording, two-photon imaging provides a uniquely powerful window into the temporal evolution of grafted cell integration in awake, behaving animals.

5.3 Multimodal Integration: Technical Challenges

The fundamental challenge of multimodal integration lies in reconciling data collected at vastly different spatiotemporal scales: volume EM operates at nanometre spatial resolution on fixed tissue (no temporal dimension), two-photon microscopy at micrometer resolution with second-to-minute temporal sampling, iEEG at millimeter resolution with millisecond temporal resolution, and fMRI at millimeter resolution but with second-scale temporal resolution constrained by hemodynamic response delay [21]. Converting these heterogeneous modalities into a unified graph representation requires modality-specific preprocessing pipelines, spatial co-registration (aligning microscopy volumes with electrode coordinates and MRI anatomy), temporal alignment (mapping electrophysiological dynamics to structural snapshots), and feature engineering that extracts biologically meaningful graph node and edge attributes from each modality.

Community data standards — Brain Imaging Data Structure (BIDS) for neuroimaging and Neurodata Without Borders (NWB) for neurophysiology [37] — provide essential scaffolding for harmonized data management and should be adopted as the organizational framework for all datasets generated in this project. For public data resources, the OpenNeuro Imaging Database for Epilepsy and Surgery (IDEAS) [42] and EBRAINS [44] offer relevant neuroimaging and multimodal epilepsy datasets, though dedicated datasets for neural stem cell transplantation remain scarce and primarily cover hematopoietic rather than neural applications [45].

Table 1. Summary of key data acquisition techniques, their spatiotemporal properties, and relevance to GNN-based circuit integration analysis.

Technique	Spatial Res.	Temporal Res.	Information Type	Strengths	Limitations	Relevance
Scalp EEG	Low (regional)	High (ms)	Population electrical activity	Wide clinical use; high temporal resolution; low cost	Low spatial resolution; noise-prone	Macro-level functional dynamics; epileptiform activity screening
iEEG / SEEG	High (local circuits)	High (ms)	Local field potentials; single-unit spiking	Direct neural measurement; seizure-onset zone localization	Invasive; limited coverage	Essential for the functional mapping of transplanted cell integration
Patch-clamp	Single-cell	Very high (µs)	Membrane potential; synaptic currents	Gold standard for single-cell physiology	Low throughput; technically demanding	Confirms functional maturation and synapse formation of grafted cells
Two-photon microscopy	Cellular / spines	Seconds-minutes (live)	Morphology; calcium transients	Live imaging; deep penetration; large FOV	Restricted to superficial layers	Tracks spine dynamics and activity of transplanted neurons longitudinally
Volume EM	Ultrastructural (synapse)	Static	3D synapse density, shape, type	Gold standard for ultrastructure; excitatory/inhibitory discrimination	Destructive; limited volume; labour-intensive	Ground truth for synapse formation and maturation of graft
Connectome-seq	Single-synapse + molecular	Static	Synaptic wiring + molecular identity	High-throughput; unbiased; captures molecular profiles	Requires engineered proteins; computation-intensive	Maps the precise inhibitory wiring and molecular identity of integrated cells

6 Graph Neural Network Architectures for Circuit Rewiring Analysis

6.1 Graph Convolutional Networks for Static Circuit Snapshots

Graph Convolutional Networks (GCNs) perform neighborhood feature aggregation via spectral or spatial convolution, learning representations of nodes (neurons/regions) that reflect both their intrinsic properties and their relational context within the circuit graph. Applied to static circuit snapshots derived from volume EM synapse maps or fixed-timepoint iEEG connectivity matrices, GCNs can classify circuit states (pre-transplant, early integration, mature integration), identify pathological versus normalized connectivity patterns, and detect aberrant graft connectivity (e.g., excitatory rather than inhibitory synapse formation) from node and

edge feature patterns. GCN implementations are available in PyTorch Geometric (PyG) and Deep Graph Library (DGL), both of which are GPU-accelerated via NVIDIA AI frameworks [28,29].

6.2 Dynamic GNNs for Longitudinal Rewiring Trajectories

Dynamic GNNs extend the static GCN paradigm by incorporating temporal edges, enabling the graph structure and feature matrices to evolve over time. Applied to longitudinal post-transplantation datasets — in which circuit connectivity is measured at multiple post-operative timepoints via repeated iEEG recording sessions, serial two-photon imaging, and longitudinal MRI — dynamic GNNs can learn the characteristic trajectory of successful versus unsuccessful integration. Temporal message passing using GRU-based update functions (as formalized in Eq. 5) enables the model to maintain a memory of prior circuit states, capturing the cumulative effect of graft-induced synaptic additions and TREM2-mediated pruning events on network topology [22].

6.3 Attention-based GNNs for Interpretable Biomarker Discovery

Attention mechanisms within GNN architectures assign learnable weights to individual neighboring nodes or edges during message aggregation, allowing the model to differentially emphasize the most informative connections. In the context of transplantation analysis, attention weights can highlight which specific host-graft synaptic connections are most predictive of seizure suppression, or which molecular features of connected neurons (e.g., TREM2 expression level in adjacent microglia) most strongly modulate integration quality. These attention weights provide directly interpretable biological insights — a critical requirement for clinical adoption — and can be visualized as circuit-level saliency maps overlaid on anatomical images for communication with clinical collaborators [27].

6.4 Heterogeneous GNNs for Multimodal Fusion

Heterogeneous GNNs handle graphs containing multiple node types (e.g., transplanted GABAergic interneurons, host pyramidal cells, host interneurons, microglial nodes, iEEG electrode nodes) and multiple edge types (inhibitory synapses, excitatory synapses, functional coherence edges, proximity edges in physical space). Each edge type is assigned a separate learnable message function, enabling type-specific relational reasoning. This architecture is the natural computational representation of the biological reality of the transplanted brain, in which grafted cells with a specific molecular identity must form specific types of connections with specific host cell populations. Heterogeneous GNN implementations are supported in PyG via the HeteroData and HeteroConv modules [29].

Table 2. GNN architecture comparison: mechanisms, prior neuroscience applications, and relevance to post-transplantation circuit analysis.

GNN Architecture	Mechanism	Neuroscience Applications	Relevance to This Project
Graph Convolutional Networks (GCNs)	Aggregates neighbor node features via spectral or spatial convolution to learn relational representations	Brain state classification; disease diagnosis (AD, schizophrenia); functional connectivity analysis	Models structural/functional relationships in epileptic circuits; integrates electrophysiology and microscopy features as node attributes
Dynamic GNNs	Incorporates temporal edges, allowing graph	Brain development modeling; seizure	Critical for modeling post-transplantation circuit rewiring

GNN Architecture	Mechanism	Neuroscience Applications	Relevance to This Project
	topology to evolve, capturing time-variant network activity	detection; disease progression tracking	as grafted cells mature and integrate over weeks to months
Attention-based GNNs	Assigns learnable weights to neighbors, enabling the model to focus on the most informative connections	Biomarker identification; pathology-relevant hub detection; improved classification accuracy	Highlights which connections or cell types are most predictive of successful integration; supports interpretability
Heterogeneous GNNs	Handles graphs with multiple node/edge types representing diverse biological entities	Multimodal data integration (structural, functional, molecular, clinical)	Combines electrophysiology nodes/edges with microscopy-derived structural and molecular nodes into a single unified network model

7 Multimodal Biomarker Framework for Integration Assessment

7.1 Electrophysiological Biomarkers

At the clinical level, the primary efficacy endpoint is a reduction in disabling seizure frequency, with a secondary endpoint of a reduction in the interictal epileptiform discharge (IED) rate on continuous EEG monitoring [5]. However, seizure frequency is an imprecise, high-variance outcome measure that requires months of observation to achieve sufficient statistical power. Finer-grained electrophysiological biomarkers — directly measurable earlier in the post-transplantation period — include: reduction in high-frequency oscillation (HFO) rate on iEEG (a marker of local hyperexcitability), restoration of normal theta/gamma coupling in hippocampal recordings (a marker of cognitive network function), and the emergence of fast-spiking and regular-spiking firing patterns in single-unit recordings attributed to grafted interneuron subtypes [6,33]. Changes in network-level iEEG connectivity metrics — coherence, phase-locking value, Granger causality — quantifiable via GNN analysis of time-series graph evolution (Eq. 2) — provide macro-scale evidence of circuit normalization.

7.2 Imaging and Structural Biomarkers

Post-transplantation MRI provides longitudinal non-invasive tracking of graft survival (volume of engrafted tissue), migration (distribution of transplanted cells relative to injection site), and host tissue response (T2 signal changes reflecting edema or gliosis) [4,35]. Diffusion tensor imaging (DTI) enables assessment of changes in white matter integrity that reflect circuit remodeling at the tract level. Volume EM and two-photon microscopy provide the structural ground truth: inhibitory synapse density per unit volume of transplanted cell territory, spine maturity index in grafted cells, and three-dimensional reconstruction of the inhibitory synaptic architecture formed at the graft-host interface [11,15].

7.3 Molecular Biomarkers

Expression of PV and SOM in grafted cells — assessed by immunostaining, single-cell RNA sequencing, or Connectome-seq — directly confirms differentiation into the targeted inhibitory interneuron subtypes [6]. Spatial transcriptomics applied to post-operative tissue sections, or RNA profiling via MoPEDE at SEEG electrode sites, enables characterization of the host inflammatory microenvironment at the transplant site, including quantification of TREM2, complement genes (C1q, C3), and pro- versus anti-inflammatory microglial gene expression programs [11,12]. TREM2 upregulation, as noted earlier, predicts excessive synaptic pruning and restricted integration; its early detection and potential pharmacological modulation may represent a targetable co-intervention to optimize graft integration.

7.4 Composite Biomarker Strategy

The ISS formulation (Eq. 4) operationalizes the principle that no single biomarker is sufficient: robust assessment of integration requires convergent evidence from electrophysiological, structural, and molecular domains. AI-driven multimodal fusion — the GNN learning framework in Eq. 3 — provides the computational mechanism to combine these signals into a unified predictive model. The most powerful clinical biomarker signatures will likely emerge from GNN-identified combinations of modality-specific features that individually lack discriminative power but together reliably separate successful from unsuccessful integration trajectories.

Table 3. Comprehensive multimodal biomarker taxonomy for stem cell integration assessment, spanning electrophysiological, imaging, and molecular domains.

Category	Specific Biomarker	Assessment Method(s)	Significance
Electrophysiological	Seizure frequency reduction; reduction in epileptiform / HFO activity	EEG, iEEG/SEEG monitoring	Primary clinical efficacy outcome: reflects macro-level circuit normalization
Electrophysiological	Restored excitatory-inhibitory (E/I) balance; fast-spiking / regular-spiking firing of grafted cells	Patch-clamp; iEEG; computational modelling	Confirms functional rebalancing by transplanted inhibitory interneurons at single-cell resolution
Electrophysiological	Altered network connectivity patterns (oscillatory, coherence metrics)	EEG, iEEG, fMRI analyzed by GNNs	Reflects large-scale circuit rewiring and functional normalization post-transplantation
Imaging	Transplanted cell survival and migration; precise deposit placement	MRI; histological staining; stereotactic imaging	Prerequisites for integration: confirm cells reach and remain in the target zone
Imaging	Reduction in inflammation/gliosis; reorganization of network structure	MRI (T2, contrast), PET, DTI, spatial transcriptomics	The host inflammatory environment is a major determinant of integration quality; reduction indicates a permissive microenvironment
Imaging	Synaptic density and ultrastructure	Volume EM; two-photon microscopy	Direct structural evidence of new synapse formation and maturation

Category	Specific Biomarker	Assessment Method(s)	Significance
Molecular	Expression of PV- and SOM-positive GABAergic markers in grafted cells	Immunostaining; scRNA-seq; Connectome-seq	Confirms differentiation into the desired inhibitory subtypes essential for E/I rebalancing
Molecular	Molecular profiles of connected host neurons; TREM2 inflammatory signature	Connectome-seq; spatial transcriptomics; MoPEDE from SEEG	Reveals circuit-level molecular determinants of integration quality; TREM2 upregulation predicts restricted connectivity

8 Translational Challenges

8.1 Data Integration and Harmonization

Neuroscience data are characteristically heterogeneous, noisy, and institutionally siloed. Variability in EEG electrode configurations, MRI scanner field strengths and pulse sequences, microscopy imaging parameters, and tissue processing protocols introduces systematic biases that degrade AI model generalization across sites and studies [20]. The multiscale nature of the integration problem — spanning nanometre synaptic features to centimeter-scale network topology — further amplifies this challenge, as no standardized pipeline currently exists to align data across these scales into a unified graph representation. Robust data harmonization pipelines, adoption of community standards (BIDS, NWB), and federated learning approaches that enable model training across distributed datasets without centralizing sensitive clinical data are essential enabling infrastructure investments.

8.2 Interpretability and Clinical Trust

The black-box character of deep learning models is a fundamental barrier to clinical adoption in high-stakes medical decision-making [20]. Clinicians and regulatory agencies require not just accurate predictions but transparent, biologically grounded explanations of model behavior. Explainable AI (XAI) techniques — attention mechanism visualization, gradient-based saliency maps, concept-based explanations — must be embedded natively within GNN architectures rather than applied post hoc. Model explanations should be validated against established biological mechanisms (e.g., that attention is appropriately concentrated on GABAergic rather than glutamatergic connection types when predicting E/I balance restoration) to confirm biological plausibility and build clinical trust [20,27].

8.3 Computational Scalability

Processing neural graphs with thousands to millions of nodes (individual synapses in volume EM reconstructions, individual neurons in whole-hippocampus models) and dynamic edges evolving across longitudinal timepoints imposes substantial computational demands. GNNs are known to suffer from over-smoothing at depth (excessive aggregation homogenizing node representations) and scalability limitations as the graph size grows [27]. GPU cluster infrastructure or cloud computing resources are required, alongside optimized GNN frameworks (NVIDIA AI-accelerated GNN libraries, sparse matrix

operations in PyG/DGL) and algorithmic strategies such as graph sampling, mini-batch training on subgraphs, and hierarchical graph pooling to manage computational cost [28].

8.4 Preclinical-to-Clinical Translation

The translational gap between animal model findings and human clinical outcomes is a persistent challenge across neuroscience [21]. Rodent models of MTLE provide tractable experimental systems for multimodal data collection and GNN training but differ from the human epileptic brain in circuit architecture, temporal dynamics of graft integration, and inflammatory microenvironment. iPSC-derived brain organoid systems offer an intermediate platform — genetically human, experimentally controllable — for generating training data for AI models before clinical deployment [2]. Biomarkers identified and validated in preclinical models must be prospectively evaluated in human participants through longitudinal clinical studies with standardized multimodal assessment protocols and long-term follow-up sufficient to capture the full integration timeline.

---

9 Strategic Recommendations

9.1 Data Acquisition and Infrastructure

1. Design experimental protocols enabling simultaneous electrophysiology and microscopy (e.g., two-photon calcium imaging with concurrent patch-clamp or iEEG), minimizing spatiotemporal misalignment between modalities [17].
2. Adopt NWB and BIDS standards from study inception to ensure data portability, reproducibility, and compatibility with community repositories (OpenNeuro, EBRAINS, BRAIN Initiative archives) [37,42,44].
3. Invest in patient-derived iPSC organoid platforms as high-throughput, controlled data sources for initial GNN model training, enabling systematic exploration of parameter space before scarce clinical data is used [2].
4. Develop and publish open-source preprocessing pipelines for harmonizing multimodal neural circuit data, contributing to the community infrastructure needed for multi-site model generalization.

9.2 GNN Architecture and Training

5. Prioritize dynamic heterogeneous GNN architectures that natively accommodate multi-type nodes and temporally evolving edges, reflecting the biological reality of post-transplantation circuit dynamics [22,26].
6. Embed attention mechanisms and XAI modules as first-class components of GNN design, enabling biological validation of model interpretability against known circuit mechanisms [20].
7. Explore hybrid architectures combining data-driven GNNs with biophysically grounded computational models (e.g., conductance-based neuron models, FitzHugh-Nagumo seizure dynamics) to incorporate domain knowledge and improve sample efficiency [31].
8. Benchmark GNN models against established baselines (linear connectivity metrics, classical ML classifiers on connectivity matrices) to rigorously demonstrate added value.

9.3 Biomarker Validation and Clinical Integration

9. Establish a longitudinal biomarker validation pipeline that tracks ISS components at standardized time points (1, 3, 6, and 12 months post-transplantation) and correlates early integration indicators with long-term seizure and cognitive outcomes [33].
10. Prioritize non-invasive biomarker development (scalp EEG with advanced source modeling, structural and functional MRI) to enable broad clinical applicability, while using invasive methods (iEEG, patch-clamp, volume EM) for gold-standard validation in research contexts [35].
11. Engage clinical, regulatory, and patient stakeholders early in biomarker development to ensure that ISS components and thresholds are clinically meaningful, practically measurable, and aligned with regulatory endpoints [35].

9.4 Open Science and Collaboration

12. Contribute trained GNN models, anonymized datasets, and analysis code to public repositories (GitHub, BRAIN Initiative Distributed Archives) to promote reproducibility and accelerate community progress.
13. Build interdisciplinary consortia integrating neuroscience, epileptology, stem cell biology, AI/ML engineering, clinical trials, and health economics expertise — the complexity of this problem domain exceeds any single disciplinary boundary.
14. Proactively engage with ILAE working groups and regulatory bodies (FDA, EMA) to co-develop the evidentiary standards that AI-derived biomarkers must meet for clinical and regulatory acceptance.

---

10 Discussion

10.1 GNNs as the Natural Computational Language of Circuit Integration

The case for GNNs as the appropriate computational architecture for this problem rests on structural alignment: neural circuits are graphs, and GNNs are the class of models designed to learn from graphs. The alternative — flattening connectivity matrices into vectors for conventional ML or CNN processing — discards the relational structure that is biologically and mechanistically informative. The emergence of topology-aware GNNs that outperform conventional methods in epilepsy surgery planning [13] provides a proof of concept that graph-native representations carry clinically relevant information not captured by traditional approaches.

The integration decoding problem adds temporal and multimodal dimensions that standard static GCNs cannot address. Dynamic GNNs operating on longitudinal graph sequences, heterogeneous GNNs processing multi-type biological entities, and attention mechanisms surfacing the most integration-relevant connections represent the architectural evolution required to match the biological complexity of the problem. The mathematical framework presented in Section 4 is intended to make these requirements precise and tractable for experimental design and model development.

10.2 The TREM2 Finding as a Prioritization Signal

The discovery that TREM2-mediated microglial pruning significantly restricts the brain-wide input connectome of transplanted neurons — and that integration quality is substantially improved in a TREM2-deficient microenvironment [11] — is a finding of immediate strategic importance for this framework. It identifies a specific, measurable molecular variable (TREM2 expression at the transplant site)

that modulates the graph connectivity outcome (ΔA_prune in Eq. 2) and suggests a targetable co-intervention (modulation of the TREM2 pathway) to optimize integration. The GNN framework proposed here — specifically the L_inflamn loss component in Eq. 3 — is designed to incorporate this variable, enabling the model to learn the relationship between inflammatory microenvironment features and integration trajectory outcomes. This exemplifies the broader principle that the GNN integration framework should be updated iteratively as new biological discoveries identify additional mechanistically relevant variables.

10.3 Limitations

This paper presents a theoretical and programmatic framework; it does not report experimental validation of the proposed GNN architectures or mathematical formulations. The ISS weights (α, β, γ, δ in Eq. 4) and the GNN loss coefficients (λ_1...λ_4 in Eq. 3) require empirical determination from clinical trial data that does not yet exist at a sufficient scale. The framework is developed with primary reference to MTLE and GABAergic interneuron transplantation; extension to other epilepsy syndromes, stem cell types, or transplantation paradigms will require biological context-specific adaptation. Computational feasibility for whole-hippocampus-volume EM graphs remains an open engineering challenge.

---

11 Conclusions

This paper has presented a comprehensive framework for AI-driven neural circuit mapping and integration analysis in the context of stem cell therapy for drug-resistant epilepsy. The central argument is that the integration decoding problem — determining whether and how transplanted GABAergic interneurons form functional inhibitory synapses within the pathologically remodeled MTLE circuit — is fundamentally a graph learning problem, and that Graph Neural Networks operating on multimodal, longitudinally acquired circuit data provide the appropriate computational architecture for its solution.

The mathematical framework introduced here — covering dynamic graph construction (Eq. 1-2), multimodal learning objectives (Eq. 3), composite integration scoring (Eq. 4), and heterogeneous message passing (Eq. 5) — provides a quantitative scaffold that operationalizes biological hypotheses as testable computational constructs, enabling rigorous experimental design, model benchmarking, and clinical trial endpoint development.

The multimodal biomarker taxonomy, spanning electrophysiological, structural, and molecular domains and unified through the Integration Success Score, provides a concrete framework for translating GNN model outputs into clinically meaningful, regulatory-ready endpoints. The translational barriers identified — data harmonization, model interpretability, computational scalability, and the preclinical-clinical gap — are substantial but tractable given appropriate infrastructure investment and interdisciplinary collaboration.

Drug-resistant epilepsy represents one of the clearest unmet needs in neurology, and stem cell-based circuit repair represents one of the most mechanistically principled therapeutic strategies available. The convergence of this biological opportunity with the maturing capabilities of GNN-based circuit analysis creates a timely and compelling case for targeted investment in AI-driven integration decoding as a foundational pillar of next-generation epilepsy care.

Conflict of Interest

The authors declare no conflicts of interest.

Data Availability

This paper presents a theoretical framework and a narrative review. No primary datasets were generated. All referenced data are available in the cited primary sources. Relevant open datasets are accessible via OpenNeuro (IDEAS: https://openneuro.org/datasets/ds005602), EBRAINS (https://www.ebrains.eu), and the BRAIN Initiative Distributed Archives (NWB standard).

References

[1] Neural circuit mechanisms of epilepsy: Maintenance of homeostasis at the cellular, synaptic, and neurotransmitter levels. ResearchGate, 2025. https://www.researchgate.net/publication/391494829

[2] Stem Cell Therapy in Treating Epilepsy. Frontiers in Neuroscience 2022;16:934507. https://doi.org/10.3389/fnins.2022.934507

[3] Stem Cells in Epilepsy: Innovations and Challenges in Neurosurgical Treatment. NeurologyLive, 2025. https://www.neurologylive.com

[4] First-of-its-kind stem-cell transplant treatment targets drug-resistant focal epilepsy. UChicago Medicine, 2025. https://www.uchicagomedicine.org

[5] Stem cell therapy in patients with epilepsy: A systematic review. ResearchGate 2020. https://www.researchgate.net/publication/347383404

[6] Cell therapy for neurological disorders: Progress towards an MGE progenitor-based treatment. Front Neurosci 2023;17:1177678. https://doi.org/10.3389/fnins.2023.1177678

[7] Interneuron Transplantation as a Treatment for Epilepsy. PMC / PubMed Central. PMC4665034.

[8] Mesenchymal stem cell therapy for focal epilepsy: A systematic review of preclinical models and clinical studies. ResearchGate 2022. https://www.researchgate.net/publication/360110799

[9] Duke is conducting first-in-human studies of cell transplantation for epilepsy. Duke Neurosurgery News, 2025.

[10] Human stem cell-derived GABAergic interneuron development reveals early emergence of subtype diversity. bioRxiv 2024. https://doi.org/10.1101/2024.12.03.626662

[11] Multimodal synapse analysis reveals restricted integration of transplanted neurons remodeled by TREM2. bioRxiv 2025. https://doi.org/10.1101/2025.01.31.635250

[12] High-resolution multimodal profiling of human epileptic brain activity via explanted depth electrodes. PMC 2025. PMC11721296.

[13] Leveraging interictal multimodal features and GNNs for automated planning of epilepsy surgery. Brain Communications 2025;7(3):fcaf140. https://doi.org/10.1093/braincomms/fcaf140

[14] Unraveling the Neural Circuits: Techniques, Opportunities, and Challenges in Epilepsy Research. PMC 2024. PMC10914928.

[15] Volume electron microscopy reveals unique laminar synaptic characteristics in the human entorhinal cortex. eLife 2025. https://doi.org/10.7554/eLife.96144

[16] Connectome-seq: High-throughput Mapping of Neuronal Connectivity. bioRxiv 2025. https://doi.org/10.1101/2025.02.13.638129

[17] Simultaneous multi-plane imaging of neural circuits. PMC 2016. PMC4724224.

[20] Artificial Intelligence and Neuroscience: Transformative Synergies in Brain Research and Clinical Applications. ResearchGate / MDPI J Clin Med 2025;14(2):550.

[21] Artificial Intelligence and Neuroscience: Transformative Synergies. MDPI J Clin Med 2025;14(2):550. https://doi.org/10.3390/jcm14020550

[22] Graph Neural Networks for Brain Network Analysis. IECE Technical Report 2024.

[26] Graph Neural Networks in Brain Connectivity Studies: Methods, Challenges, and Future Directions. MDPI Brain Sci 2025;15(1):17. https://doi.org/10.3390/brainsci15010017

[27] Graph Neural Networks in Brain Connectivity Studies. PMC 2025. PMC11763835.

[28] Graph Neural Network (GNN) Frameworks. NVIDIA Developer Documentation, 2025. https://developer.nvidia.com/gnn-frameworks

[29] PyTorch Geometric Documentation. Read the Docs, 2025. https://pytorch-geometric.readthedocs.io

[31] Non-vectorial Integration of Intersectional Short-Pulse Stimulation for Seizure Control. bioRxiv 2025. https://doi.org/10.1101/2025.01.02.631064

[33] Electrophysiological Biomarkers in Genetic Epilepsies. PMC 2021. PMC8609056.

[35] Neuroimaging Biomarkers of Experimental Epileptogenesis and Refractory Epilepsy. PMC 2019. PMC6337422.

[36] A Biomarker for Predicting Responsiveness to Stem Cell Therapy. PubMed 2020. PMID 32402283.

[37] NIH BRAIN Initiative: Data Science and Neurodata Without Borders. NIH RePORTER 2024. Project 10900694.

[42] The Imaging Database for Epilepsy And Surgery (IDEAS). OpenNeuro 2024. https://openneuro.org/datasets/ds005602

[43] Open Access Series of Imaging Studies (OASIS). Washington University, 2024. https://sites.wustl.edu/oasisbrains

[44] Find neuroscience data, models, and tools. EBRAINS, 2025. https://www.ebrains.eu

[45] Publicly Available Datasets — CIBMTR, 2025. https://cibmtr.org/CIBMTR/Resources/Publicly-Available-Datasets