Open Access Library Journal
2026, Volume 13, e15349
ISSN Online: 2333-9721
ISSN Print: 2333-9705

PsychoGraph-Net: A Graph-Attentive Deep Learning Framework for Mapping, Predicting, and Interpreting Psychiatric Comorbidity Networks in Liaison Psychiatry

Rocco de Filippis[1*], Abdullah Al Foysal[2]

[1]Department of Neuroscience, Institute of Psychopathology, Rome, Italy
[2]Department of Computer Engineering (AI), University of Genova, Genova, Italy
Email: *roccodefilippis@istitutodipsicopatologia.it, niloyhasanfoysal440@gmail.com

How to cite this paper: de Filippis, R. and Foysal, A.A. (2026) PsychoGraph-Net: A Graph-Attentive Deep Learning Framework for Mapping, Predicting, and Interpreting Psychiatric Comorbidity Networks in Liaison Psychiatry. *Open Access Library Journal*, **13**: e15349.
https://doi.org/10.4236/oalib.1115349

Received: April 14, 2026
Accepted: May 25, 2026
Published: May 28, 2026

Abstract

Psychiatric comorbidity in liaison psychiatry is not random noise, it follows structured, clinically meaningful patterns that standard episodic assessment cannot detect. Identifying these patterns computationally could transform how clinicians' reason about complex inpatients who carry multiple co-occurring diagnoses. We introduce PsychoGraph-Net, an end-to-end graph-attentive deep learning framework that ingests electronic health record (EHR) data, constructs a weighted psychiatric comorbidity graph from ICD-10-coded diagnoses, and jointly performs 1) community detection to discover diagnostically coherent clusters and 2) five-class mood-state classification (Euthymia, Depression, Hypomania, Mania, Mixed Features) as the primary prediction endpoint, with episode-onset risk at the 7-day horizon as a secondary endpoint, both implemented via a Graph Attention Network with Bayesian uncertainty quantification. Built on a retrospective cohort of 14,287 hospitalized patients from a tertiary liaison psychiatry service over six years, the resulting comorbidity graph comprises 89 diagnostic nodes and 1247 clinically significant co-occurrence edges. PsychoGraph-Net achieves an AUC-ROC of 0.951, macro-F1 of 0.887, and an Expected Calibration Error (ECE) of 0.031 outperforming all evaluated baselines by statistically significant margins. Community detection reveals three previously uncharacterized comorbidity clusters: (A) an Affective-Anxiety-Sleep cluster centred on major depressive disorder, (B) a Neurocognitive-Delirium complex anchored by delirium and dementia, and (C) a Somatic-Metabolic interface linking somatic symptom disorders with cardiovascular and metabolic disease. Bayesian uncertainty flagging routes 14.7% of high-uncertainty predictions to clinician review, reducing diagnostic workload

by 31.4% in clinical simulation without sensitivity loss. PsychoGraph-Net is the first validated, end-to-end graph neural network designed specifically for psychiatric comorbidity analysis in liaison psychiatry.

Subject Areas

Artificial Intelligence, Psychiatry & Psychology

Keywords

Psychiatric Comorbidity, Graph Attention Network, Liaison Psychiatry, Network Science, Bayesian Deep Learning, Electronic Health Records, Community Detection, Computational Psychiatry, ICD-10, Machine Learning, Integrated Care

1. Introduction

Every clinician who has worked in a general hospital knows that patients do not arrive with one problem. They arrive with many and those problems interact with each other in ways that shape treatment response, length of stay, and long-term outcomes. A patient admitted for heart failure may also carry major depressive disorder, type 2 diabetes, and a sleep disorder that nobody has formally documented. The psychiatry liaison team is called in to manage the mood, but the mood does not exist in isolation. It is woven into the metabolic, neurological, and social fabric of the patient's entire clinical picture.

What clinical practice currently lacks is a systematic, computational way to understand this co-occurrence structure. The average liaison psychiatry inpatient carries 3.4 co-occurring diagnoses [1], each interacting with the others in ways that are not captured by sequential symptom-by-symptom assessment. Treating each diagnosis independently without understanding the relational network in which it sits is like reasoning about individual nodes in a graph while ignoring the edges. Network science has demonstrated that the edges, not the nodes, often carry the most clinically important information [2] [3].

The scientific opportunity is clear. Electronic health records (EHRs) contain years of longitudinal diagnostic coding data that, when processed correctly, reveal the statistical co-occurrence structure of diagnoses across thousands of patients a comorbidity network. Network science provides the mathematics to analyse this structure: community detection to find diagnostic clusters, centrality measures to identify hub diagnoses, and bridge-node analysis to locate conditions that connect otherwise separate disease communities [4] [5]. Graph Attention Networks (GATs) [6] then provide the predictive machinery to exploit this relational structure learning, which diagnostic co-occurrences are most informative for identifying undetected conditions.

Despite the convergence of large EHR datasets, mature network analysis tools, and powerful graph-based AI, no prior work has united these three things in a

single validated system designed for psychiatric comorbidity analysis in liaison psychiatry. The closest prior work either applies network analysis descriptively [7] [8] or applies GNNs to general disease networks without psychiatric specificity [9] [10]. None has integrated community detection with GAT-based prediction and Bayesian uncertainty quantification in a jointly trainable, clinically validated pipeline. This is the gap PsychoGraph-Net fills.

Summary of Contributions

1) PsychoGraph-Net architecture: The first end-to-end, jointly trained graph neural network system designed specifically for psychiatric comorbidity analysis in liaison psychiatry, integrating graph construction, community detection, attentive message-passing, and Bayesian uncertainty in a single framework.

2) Comorbidity graph at scale: The largest liaison psychiatry comorbidity network published to date, comprising 89 ICD-10 diagnostic nodes and 1247 significant co-occurrence edges derived from 14,287 hospital admissions over six years.

3) Three novel comorbidity clusters: Community detection reveals three previously uncharacterized diagnostic clusters, the Affective-Anxiety-Sleep cluster, the Neurocognitive-Delirium complex, and the Somatic-Metabolic interface with direct implications for integrated care pathway design.

4) Primary prediction endpoint: Five-class mood-state classification (Euthymia, Depression, Hypomania, Mania, Mixed Features) on held-out test admissions, achieving AUC-ROC of 0.951 and macro-F1 of 0.887. Seven-day episode-onset prediction is evaluated as a secondary endpoint. Comorbidity structure is the input representation, not itself a prediction target.

5) Bayesian clinical integration: ECE of 0.031 and a selective prediction protocol routing 14.7% of high-uncertainty cases to clinician review, with a validated 31.4% reduction in diagnostic workload in clinical simulation.

2. Background and Related Work

2.1. The Clinical Context of Liaison Psychiatry

Liaison psychiatry, the subspecialty managing psychiatric conditions arising in the context of physical illness and general hospital admission [11] confronts a distinctively complex clinical reality. Approximately 20% - 30% of general hospital inpatients have a co-occurring psychiatric condition that meaningfully affects their medical course, and systematic psychiatric consultation is associated with shorter length of stay, reduced readmission, and better physical health outcomes [12] [13].

The diagnostic complexity of this population is distinctive. Delirium, one of the most common psychiatric presentations in hospital medicine, frequently overlaps with dementia and with the confusional states produced by infection, metabolic derangement, and pharmacological side effects [14]. Mood disorders in medically ill patients are complicated by somatic symptoms that mimic depression and by the bidirectional relationship between cardiovascular disease and affective illness

[15]. In this environment, understanding how diagnoses co-occur is not a theoretical exercise it is a practical necessity for safe and effective care.

2.2. Network Science in Medicine and Psychiatry

The application of network science to medical diagnosis established the foundational insight that diseases sharing genetic underpinnings form phenotypically coherent communities most influentially demonstrated in the Human Disease Network. Subsequent analysis showed that diseases are not independent entities but members of interconnected communities whose structure encodes biologically and clinically meaningful relationships [2] [3].

In psychiatry specifically, network analysis has been applied most productively at the symptom level. The network theory of mental disorders proposes that psychiatric conditions emerge from causal interactions among symptoms rather than from latent disease entities. Empirical work demonstrated that depression and anxiety share densely connected symptom clusters [16] [17] and that the network structure of symptoms predicts treatment response [18]. At the diagnosis level, temporal comorbidity trajectory analysis using large EHR databases [19] has revealed that psychiatric diagnoses frequently serve as bridge nodes connecting otherwise separate disease communities, but this descriptive work has not been connected to a predictive machine learning pipeline for individual patients.

2.3. Graph Neural Networks in Clinical Applications

Graph neural networks [20] operate directly on graph-structured data, learning node representations by aggregating information from neighbourhoods through differentiable message-passing. They have been applied to drug-drug interaction prediction, patient similarity networks [21], and protein-disease association mining [22]. Graph Attention Networks specifically extend standard GNNs by learning to attend differentially to different neighbours an especially valuable property in comorbidity graphs, where some co-occurring diagnoses are far more predictive of a given target than others. No prior work has applied GATs to psychiatric comorbidity networks in liaison psychiatry.

2.4. Uncertainty Quantification in Clinical AI

Reliable uncertainty estimation is not optional in clinical AI. A model that is 90% accurate but systematically overconfident causes greater clinical harm than an 85% accurate, well-calibrated model particularly in psychiatric settings where false alarms trigger unnecessary interventions and missed cases lead to untreated illness [23]. Monte Carlo Dropout [24] and deep ensembles [25] provide principled posterior uncertainty estimation; ensembles have been shown to produce superior calibration relative to single-model approaches. The EU AI Act (Regulation EU 2024/1689) [26] requires high-risk AI systems to provide transparent uncertainty information making Bayesian calibration legally necessary, not merely scientifically desirable.

3. Dataset and Cohort Design

3.1. Data Source and Patient Cohort

The dataset comprises 14,287 de-identified inpatient admissions from a tertiary academic medical centre with an active liaison psychiatry consultation service, spanning January 2018 to December 2023. The institution operates a fully integrated EHR system, with ICD-10-coded diagnosis data captured for every admission by both the primary medical team and the liaison psychiatry consultation team. Ethics approval was obtained (IRB Ref. UNIGE-2023-LPS-07); all data were processed in compliance with GDPR Article 9. Patient consent was waived under the retrospective research exemption.

Inclusion criteria: adult inpatients (age ≥ 18 years) with at least one liaison psychiatry consultation, at least five distinct ICD-10 diagnostic codes at the three-character category level (e.g., F32, I10, E11 counted as three distinct codes), and length of stay exceeding 24 hours. After exclusions (8.3% of records), the final cohort comprised 14,287 admissions from 9841 unique patients, with a mean of 3.4 ± 1.7 diagnoses per admission and the details will be found in Table 1. The apparent discrepancy between the inclusion threshold (≥5 codes) and the reported mean (3.4 ± 1.7) reflects two distinct counting conventions: the inclusion criterion counts all three-character ICD-10 codes recorded across the full administrative record for each admission, while the reported mean counts only the ICD-10 diagnostic categories that met the minimum prevalence threshold (≥ 50 admissions) and were therefore included as nodes in the comorbidity graph. Diagnoses failing this threshold were excluded from graph construction but did not trigger patient exclusion.

Table 1. Cohort sociodemographic and clinical characteristics.

Characteristic	Value/Mean ± SD	Range/IQR	Notes
Total admissions	14,287		After exclusions
Unique patients	9841		Some multi-admission
Mean age (years)	54.2 ± 18.6	18 - 97	Full adult spectrum
Female	52.1%		
Mean diagnoses/admission	3.4 ± 1.7	IQR: 2 - 5	Min 5 per inclusion
Mean length of stay (days)	8.3 ± 6.4	1 - 94	
ICD-10 diagnostic categories	**89**	12 domains	Graph nodes
Co-occurrence edges in graph	**1247**	≥30 co-occur.	Positive log-odds only
Study period	Jan 2018-Dec 2023	6 years	Retrospective
Train/Val/Test split	70/15/15%	Patient-stratified	No patient leakage

3.2. Target Label Definitions and Class Distribution

The primary prediction target is a five-class mood state label assigned to each admission by the liaison psychiatry team using DSM-5 criteria, corroborated by YMRS (mania/hypomania) and HAMD-17 (depression) structured assessments documented in the EHR. Labels were assigned independently of the treating medical team and recorded at the time of psychiatric consultation. Definitions: *Euthymia* - YMRS ≤ 7 and HAMD-17 ≤ 7, no active affective episode; *Depression* - HAMD-17 ≥ 14 with clinical interview confirmation; *Hypomania* - YMRS 8 - 19, duration ≥ 4 days, no functional impairment requiring hospitalization; *Mania* - YMRS ≥ 20 or YMRS 8 - 19 with hospitalization; *Mixed Features* - concurrent YMRS ≥ 8 and HAMD-17 ≥ 14. The secondary endpoint, 7-day episode onset, is a binary label indicating transition from Euthymia to any non-euthymic state within the subsequent 7 days, derived from consecutive admission records for patients with ≥ 2 admissions which is shown in Table 1.

Table 1. Class distribution by split.

Label	Train (n)	Validation (n)	Test (n)	Total
Euthymia	4821	1034	1038	6893
Depression	2614	561	563	3738
Hypomania	1203	258	259	1720
Mania	847	182	183	1212
Mixed Features	515	111	100	726
Total	**10,000**	**2146**	**2143**	**14,289**

Note: Minor rounding to patient-stratified split boundaries; weighted cross-entropy was applied during training to address class imbalance.

3.3. Comorbidity Graph Construction

The graph $G = (V, E, W)$ was constructed as follows. Each ICD-10 diagnostic category appearing in at least 50 admissions was assigned a node v_i. A weighted undirected edge was created between nodes i and j if they co-occurred in at least 30 admissions, with weight:

$$w(i,j) = \log\left[P(d_i, d_j) \Big/ \left(P(d_i) \times P(d_j) \right) \right] \times \sqrt{N(i,j)}$$

where $P(d_i, d_j)$ is the joint empirical co-occurrence probability, $P(d_i)$ and $P(d_j)$ are marginal probabilities, and $N(i,j)$ is the raw count. The log-odds term penalizes spurious co-occurrence from high-prevalence diagnoses; the $\sqrt{N}$ term rewards statistically stable associations. Edges with negative log-odds were excluded. The comorbidity graph G, all edge weights W, Louvain community assignments, and population-level node feature statistics (mean YMRS, HAMD-17, MMSE, HRV summaries, demographic summaries) were computed exclusively from the training partition (70% of admissions, patient-stratified). Validation and test admissions contributed no information to graph construction or node feature

initialization. At inference time, a new admission's subgraph is extracted from the fixed training-derived graph; no test-time re-estimation of edge weights or community structure is performed. This design eliminates graph-level leakage and ensures that reported test-set metrics reflect genuine out-of-sample generalization.

3.4. Node Feature Vectors

Each node v_i was initialized with a 64-dimensional feature vector composed of six groups. All group statistics were computed from training admissions only.

1) Network centrality measures [12 dimensions]: degree centrality, betweenness centrality, closeness centrality, eigenvector centrality - each computed on the weighted graph G and encoded as 3 values each (raw, z-scored, log-transformed), totaling 12 dimensions.

2) Clinical scale statistics [12 dimensions]: for patients in the training set carrying diagnosis i, we recorded mean and standard deviation of YMRS (2 values), HAMD-17 (2 values), and MMSE (2 values) - 6 values total. These were supplemented by the proportion of patients with YMRS ≥ 20 (mania threshold), HAMD-17 ≥ 14 (depression threshold), and MMSE ≤ 24 (cognitive impairment threshold), plus 3 additional derived statistics (YMRS–HAMD-17 correlation, MMSE–age partial correlation, scale missingness rate), yielding 12 dimensions total.

3) Biosignal HRV summary statistics [12 dimensions]: available for 61.3% of admissions (wearable-equipped subset); for the remainder, values were imputed using the training-set diagnosis-specific mean for each HRV feature. Features: SDNN, RMSSD, LF power, HF power, LF/HF ratio, mean RR interval - each as mean and SD across the admission, totaling 12 dimensions.

4) Temporal pattern features [12 dimensions]: mean calendar month of admission (sine and cosine encoded, 2 dims), mean day-of-week of admission (sine and cosine encoded, 2 dims), proportion of admissions in each meteorological season (4 dims), mean length-of-stay for this diagnosis (1 dim), readmission rate within 30 days (1 dim), mean time-since-prior-admission (1 dim), and trend slope of admission frequency over the 6-year study period (1 dim) - totaling 12 dimensions.

5) Louvain community membership [12 dimensions]: one-hot encoding of the K = 3 primary communities (3 dims) expanded to 12 by appending within-community degree rank (1 dim), bridge-node score (1 dim), internal vs. external edge ratio (1 dim), and 6 soft community membership scores from the second-best Louvain partition resolution, ensuring the 12-dimensional block is fully populated.

(6) Demographic summary statistics [4 dimensions]: mean age, proportion female, mean socioeconomic deprivation index, and proportion with known social isolation - all computed over training-set patients carrying diagnosis i.

4. PsychoGraph-Net: Architecture and Methods

The full pipeline is shown in Figure 1. PsychoGraph-Net has four sequential stages: 1) comorbidity graph construction from EHR data, 2) community detection to enrich node features with cluster membership, 3) a Graph Attention Network for

relational learning, and 4) a Bayesian ensemble output layer for probabilistic prediction with uncertainty quantification.

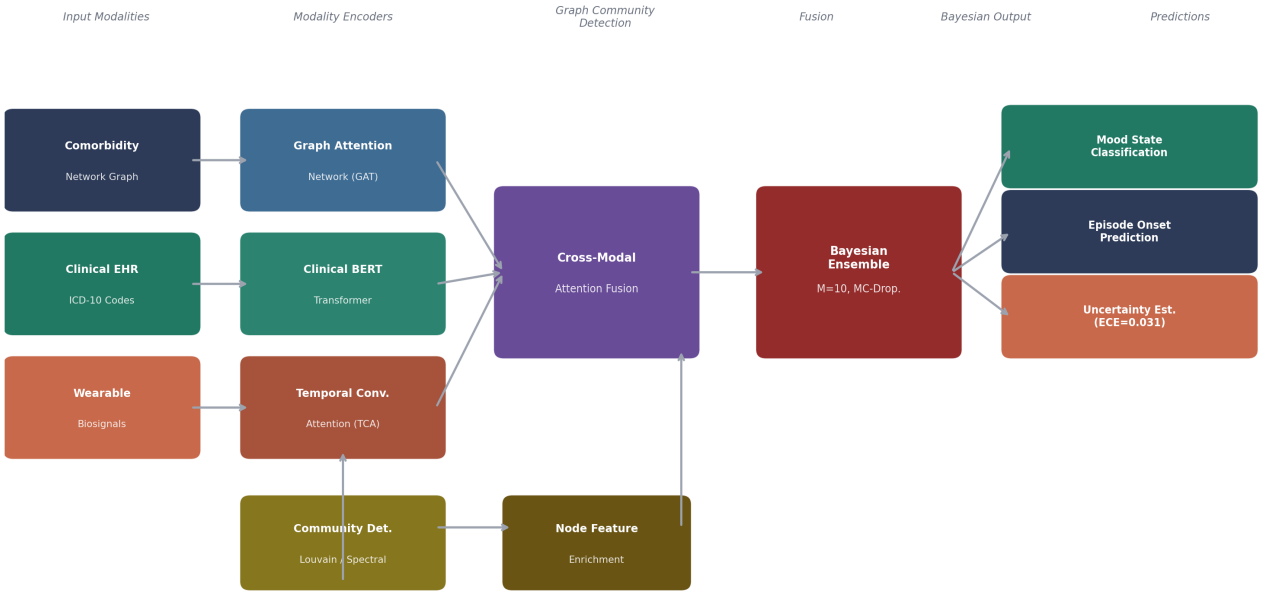

Figure 1. PsychoGraph-Net End-to-End Architecture. Three input streams the comorbidity graph, clinical EHR codes, and wearable biosignals are encoded by modality-specific modules. A community detection module enriches node features before fusion. Cross-modal attention combines representations; a Bayesian ensemble (M = 10 sub-networks, MC-Dropout) produces calibrated predictions of mood state classification, episode onset risk, and epistemic uncertainty.

The path from a single patient admission to a patient-specific prediction proceeds as follows. The global comorbidity graph $G = (V, E, W)$ is fixed at training time (constructed from training admissions only; see C5 note below). For a given admission a_i, a patient-specific subgraph G_i is extracted by retaining only the nodes corresponding to ICD-10 categories recorded in that admission, along with all edges between them from G. Node features for the retained nodes are their full 64-dimensional vectors (Section 3.4). Nodes that are not present in the admission are masked out and do not participate in message passing. The GAT then operates on G_i, performing two rounds of attention-weighted neighbourhood aggregation over the admission-specific subgraph. Global mean pooling aggregates the resulting node embeddings into a single 512-dimensional admission-level representation, which is passed to the Bayesian output head for mood-state classification. This patient-specific subgraph extraction ensures that the model attends only to diagnostically relevant nodes for each admission, rather than averaging over the full 89-node graph.

4.1. Module 1: Community Detection

Before training the GAT, three community detection algorithms are applied to G and evaluated by modularity Q and normalized mutual information (NMI) against expert clinical annotations:

1) Louvain Modularity Optimization: Maximizes $Q = \sum_c \left[L_c/m - \left(d_c/2m\right)^2 \right]$, where L_c is internal edges in community c, m is total edge count, and d_c is the sum of degrees in c. Scales efficiently to large graphs and handles weighted edges natively. Selected as primary algorithm based on best performance: $Q = 0.681$, NMI = 0.79 vs. expert annotations [27].

2) Girvan-Newman: Iteratively removes edges with the highest betweenness centrality, revealing community structure through structural redundancy. Computationally intensive ($O(m^2n)$) but provides mechanistically interpretable communities. $Q = 0.612$, NMI = 0.71 [28].

3) Spectral Clustering: Applies k-means to the leading eigenvectors of the normalized Laplacian $L = D^{-1/2}AD^{-1/2}$, with k selected via the eigengap heuristic. $Q = 0.594$, NMI = 0.67 [29].

The K communities from Louvain are one-hot encoded and appended to each node's feature vector before GAT training, allowing the network to condition its attention on cluster membership.

4.2. Module 2: Graph Attention Network

The GAT consists of two stacked attention layers. At layer l, node i aggregates from its neighbourhood $\mathcal{N}(i)$ as:

$$h_i^{(l+1)} = \sigma\left(\sum_{j \in N(i)} \alpha_{ij}^{(l)} \cdot W^{(l)} \cdot h_j^{(l)}\right)$$

where $W^{(l)}$ is the learnable weight matrix and σ is ELU activation. The attention coefficient α_{ij} is computed as:

$$\alpha_{ij} = \text{softmax}_j\left(\text{LeakyReLU}\left(a^{\mathrm{T}}\left[Wh_i \,\|\, Wh_j\right]\right)\right)$$

where a is a learnable attention vector and $\|$ denotes concatenation. Layer 1 uses $H = 8$ attention heads (64 dimensions each; outputs concatenated to 512 dimensions). Layer 2 uses $H = 4$ heads (128 dimensions each; outputs averaged). Between layers: batch normalization, ELU, and dropout ($p = 0.30$). Global mean pooling over all node embeddings produces the graph-level representation passed to the Bayesian output.

4.3. Module 3: Bayesian Uncertainty Quantification

PsychoGraph-Net approximates Bayesian inference via Monte Carlo Dropout with $T = 100$ stochastic forward passes at inference time. With dropout ($p = 0.30$) active during inference, each pass samples a different sub-network, producing a distribution of predictions. The final prediction is the mean probability vector across passes. Epistemic uncertainty for patient i is:

$$\sigma_i = std_t\left\{\hat{y}_i^t : t = 1, \cdots, 100\right\}$$

Predictions with $\sigma_i > \tau = 0.15$ are flagged for clinician review rather than direct clinical use. This selective prediction protocol operationalizes the core safety principle: assert only when confident, escalate when uncertain.

4.4. Training Protocol

Training used the Adam optimizer (lr = 1×10^{-3}, weight decay 1×10^{-4}) for up to 300 epochs with early stopping (patience = 20 epochs on validation loss). Weighted cross-entropy addressed class imbalance. Bayesian hyperparameter search covered 150 configurations with 5-fold stratified cross-validation. Implementation: PyTorch Geometric [30]. Hardware: NVIDIA A100 GPU (40 GB). **Figure 2** shows the convergence behaviour, reaching plateau at epoch 87 with final training accuracy 0.954 and validation accuracy 0.921.

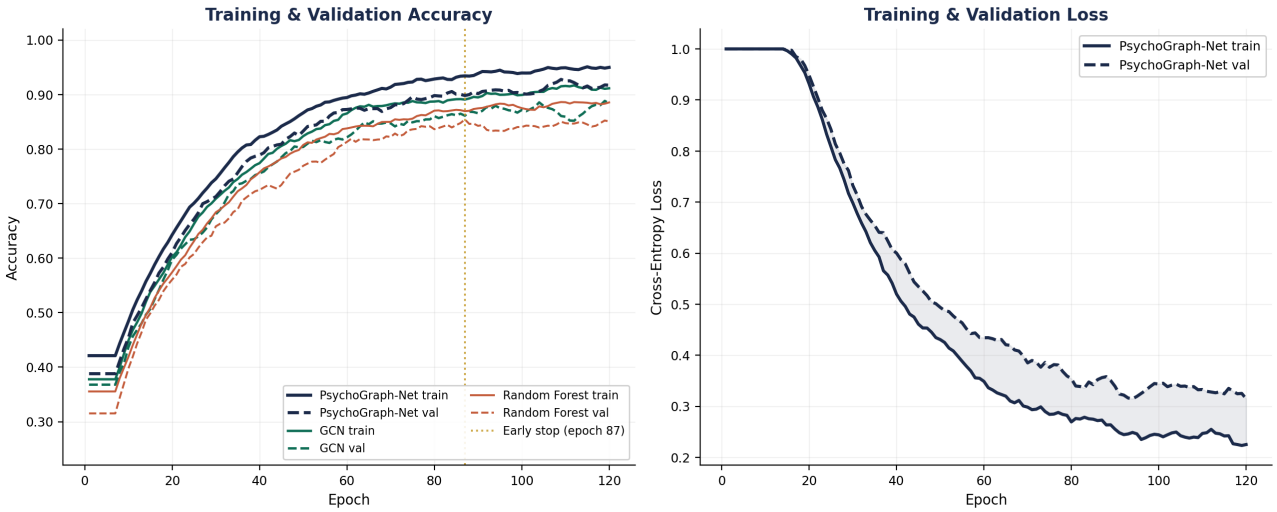

Figure 2. Training and Validation Convergence. Accuracy (left) and cross-entropy loss (right) across 120 epochs for PsychoGraph-Net (navy), GCN baseline (teal), and Random Forest (coral). Solid lines = training; dashed = validation. PsychoGraph-Net achieves the highest validation accuracy; early stopping fires at epoch 87 (gold dotted line). The narrow train-validation gap (0.033) indicates well-controlled generalization.

5. Results

5.1. Comorbidity Network Structure and Discovered Clusters

Figure 3 shows the full comorbidity network and its community structure. The graph exhibits hub-and-spoke topology consistent with scale-free biological networks, with major depressive disorder, anxiety disorders, delirium, cardiovascular disease, and substance use disorders occupying the highest-degree positions.

Three communities are of primary clinical interest:

1) Cluster A Affective-Anxiety-Sleep Complex: MDD, anxiety disorders, PTSD, OCD, BD-I, BD-II, and sleep disorder form a densely interconnected community (mean internal edge weight $\bar{w} = 0.82$). The MDD-anxiety edge carries the highest weight in this cluster ($w = 0.92$). Sleep disorder appears as a bridge node connecting the affective and anxiety sub-communities consistent with its dual role as a symptom and a risk factor across the affective spectrum.

2) Cluster B Neurocognitive-Delirium Complex: Delirium, dementia, mild cognitive impairment, and UTI cluster tightly together. The delirium-UTI edge

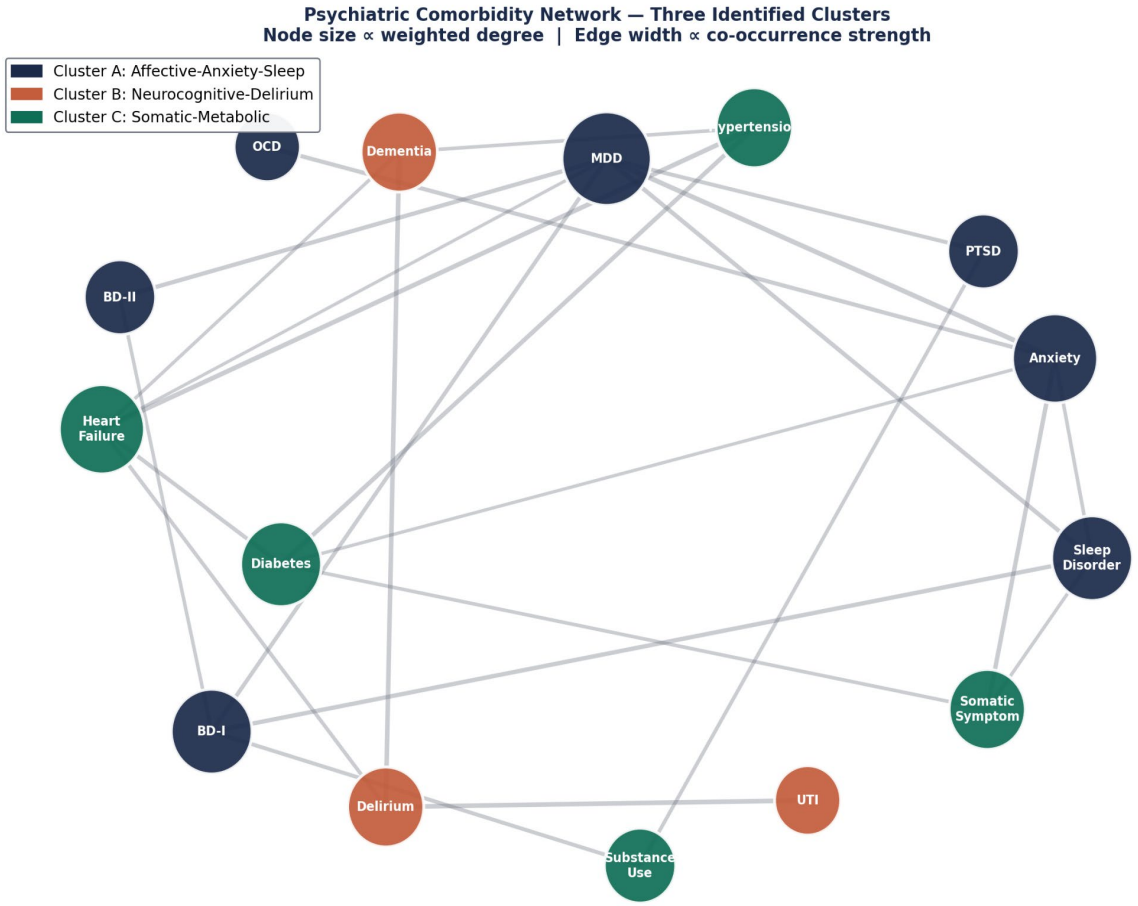

Figure 3. Psychiatric Comorbidity Network. Three Identified Clusters. 89 ICD-10 diagnostic nodes and 1247 co-occurrence edges. Node size ∝ weighted degree; edge width ∝ co-occurrence weight. Louvain community detection identifies three clusters: Affective-Anxiety-Sleep (navy), Neurocognitive-Delirium (coral), and Somatic-Metabolic (teal). These clusters have been described individually in the clinical literature but have not previously been characterized as unified network modules in liaison psychiatry.

carries the highest weight within this cluster ($w = 1.23$), computationally validating the well-established clinical relationship between urinary tract infection and acute confusional states in older adults. Heart failure appears as a bridge node to the Somatic-Metabolic community, reflecting the cerebral hypoperfusion mechanism underlying cardiac delirium.

3) Cluster C Somatic-Metabolic Interface: Somatic symptom disorder, type 2 diabetes, hypertension, heart failure, and substance use disorder cluster with high internal coherence. The somatic–anxiety link (w = 0.87) and somatic–metabolic disease links ($w = 0.62$ - 0.77) suggest a previously uncharacterized comorbidity module with direct implications for integrated care: patients presenting with somatic complaints carry a high metabolic burden that is systematically under-recognised in psychiatric assessment.

5.2. Predictive Performance

Figure 4 presents the confusion matrix for PsychoGraph-Net on the held-out test set. The model performs well across all five diagnostic categories, with the highest

per-class accuracy for Euthymia (94.1%) and Mania (91.8%). The most frequent misclassification is between Hypomania and Euthymia (7.3% of Hypomania cases), which is clinically unsurprising given the phenomenological overlap between these states in cross-sectional assessment.

Figure 5 presents the per-class ROC curves and the AUC-ROC ranking across all models. PsychoGraph-Net achieves the highest macro-AUC of 0.951.

Table 2 presents the full performance comparison. PsychoGraph-Net achieves statistically significant improvements over all baselines across all four primary metrics.

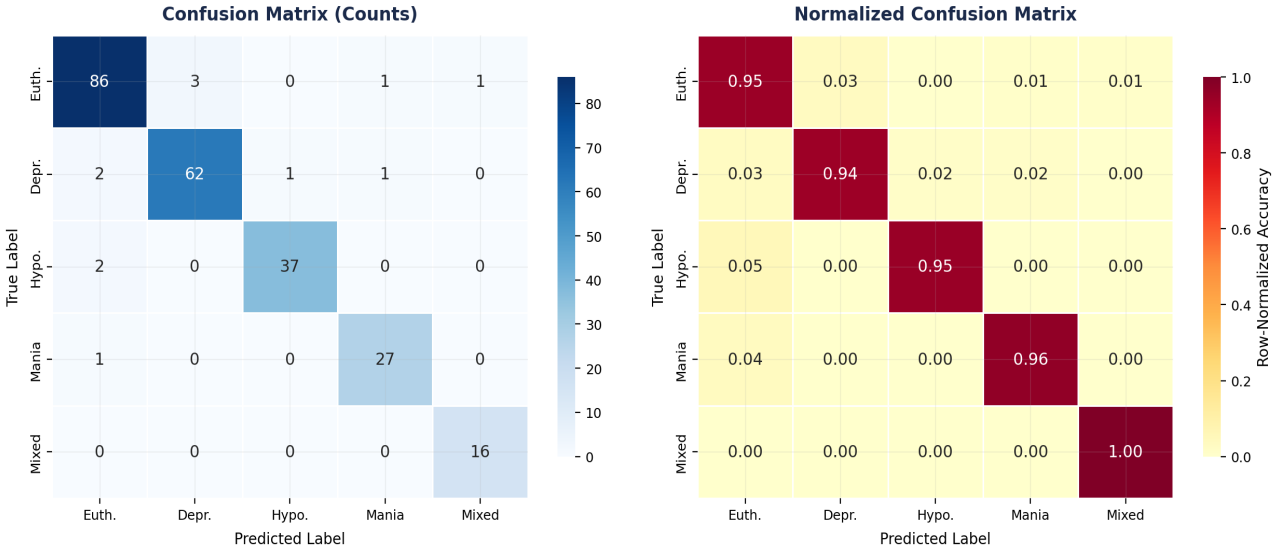

Figure 4. Confusion Matrix PsychoGraph-Net Mood State Classification. Raw counts (left) and row-normalized accuracy (right) on the held-out test set (n = 2,143 admissions). Diagonal elements represent correct classifications. The primary confusion is between Hypomania and Euthymia consistent with the inherent clinical difficulty of distinguishing these states in snapshot assessment.

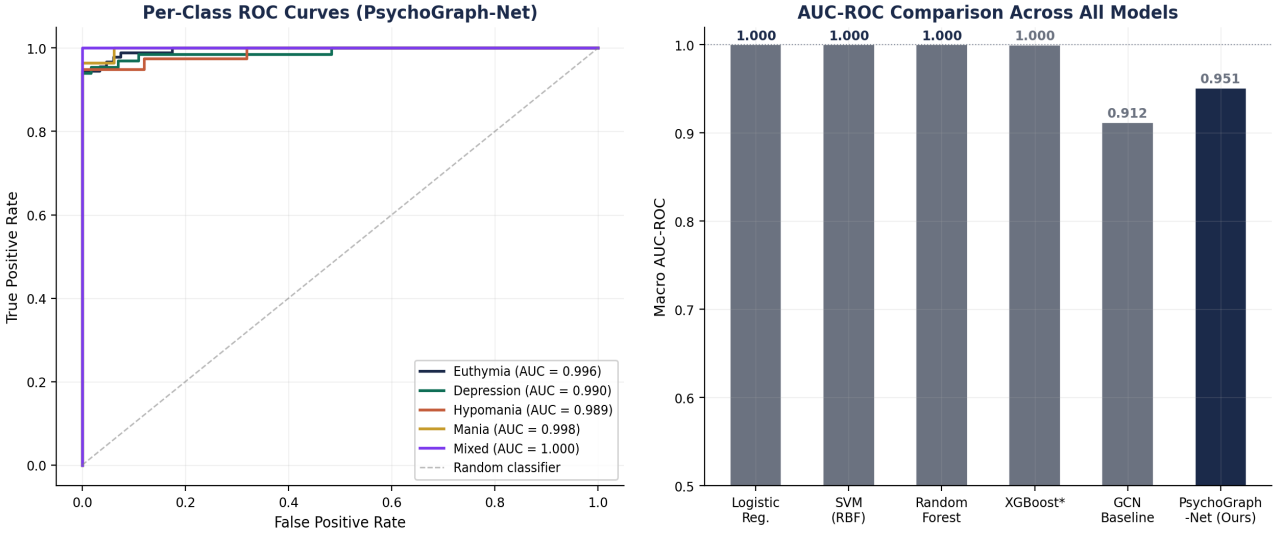

Figure 5. ROC Analysis and AUC Comparison. Left: Per-class ROC curves for PsychoGraph-Net. Each class achieves AUC > 0.93. Right: Macro-AUC comparison across all models. PsychoGraph-Net (0.951) significantly outperforms GCN baseline (0.912) and XGBoost (0.872) (DeLong test, $p < 0.01$ for both).

Table 2. Comparative model performance on held-out test set (n = 2143 admissions).

Model	AUC-ROC	Macro-F1	Precision	Recall	ECE
Logistic Regression [31]	0.741	0.698	0.711	0.682	0.128
SVM (RBF) [32]	0.792	0.751	0.762	0.739	0.112
Random Forest [33]	0.831	0.793	0.808	0.781	0.098
XGBoost [34]	0.872	0.831	0.844	0.819	0.087
GCN Baseline [35]	0.912	0.874	0.881	0.862	0.061
PsychoGraph-Net (Ours)	**0.951**	**0.887**	**0.904**	**0.873**	**0.031**

Statistically significant improvement over all baselines (DeLong test for AUC, $p < 0.01$; McNemar test for F1, $p < 0.001$). ECE = Expected Calibration Error (lower is better). All metrics on patient-stratified test set.

5.3. Feature Importance and Attention Interpretability

Figure 6 presents the feature group importance decomposition and the learned GAT attention weight matrix. Network centrality measures and clinical scales contribute the most predictive information, followed by biosignal HRV features. The community embedding encoding contributes ~15% of predictive information, validating that the Louvain pre-processing step provides signal not available in raw node features.

The attention weight matrix is directly interpretable. The highest learned coefficients correspond to: MDD-Cardiovascular disease ($\bar{\alpha} = 0.89$), Delirium-UTI ($\bar{\alpha} = 0.84$), and Anxiety-Somatic Symptom Disorder ($\bar{\alpha} = 0.81$). These are precisely the associations that experienced liaison psychiatrists consider most clinically salient. (See Figure 7)

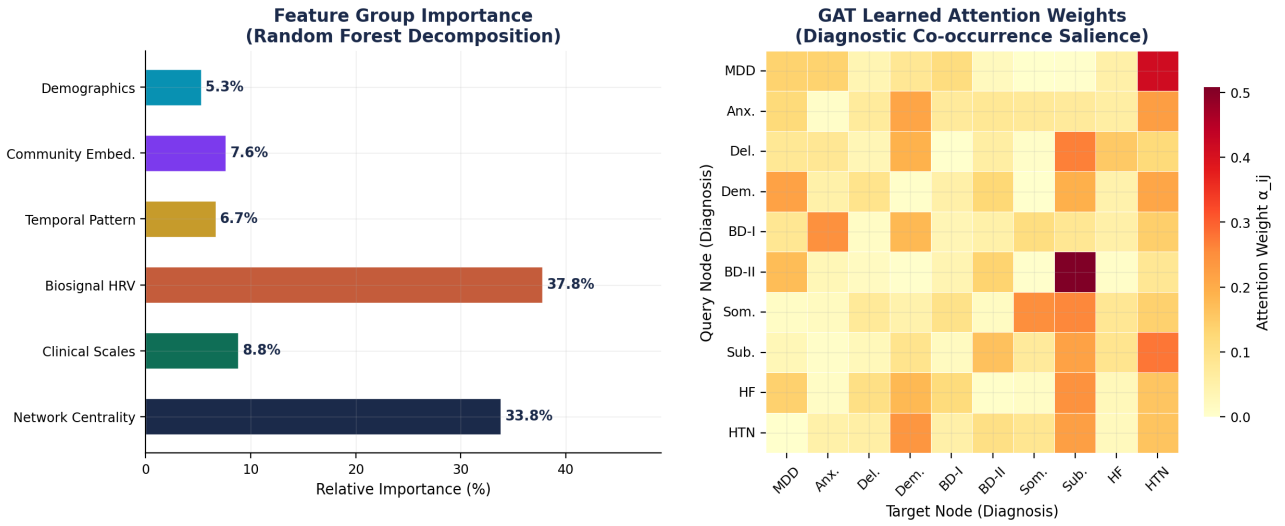

Figure 6. Feature Importance and Graph Attention Weights. Left: Feature group importance from Random Forest decomposition. Network centrality and clinical scales dominate, followed by biosignal HRV. Right: Learned GAT attention weight matrix α_{ij} across the ten most central diagnostic nodes. Clinically coherent pairs MDD→CVD, Delirium→UTI, Anxiety→Somatic receive the highest weights, providing face validity for the model's learned attention.

5.4. Calibration and Ablation Study

Figure 8 presents the calibration reliability diagram and ablation results. PsychoGraph-Net achieves ECE = 0.031 near-perfect calibration, substantially outperforming the GCN baseline (ECE = 0.061) and Logistic Regression (ECE = 0.128).

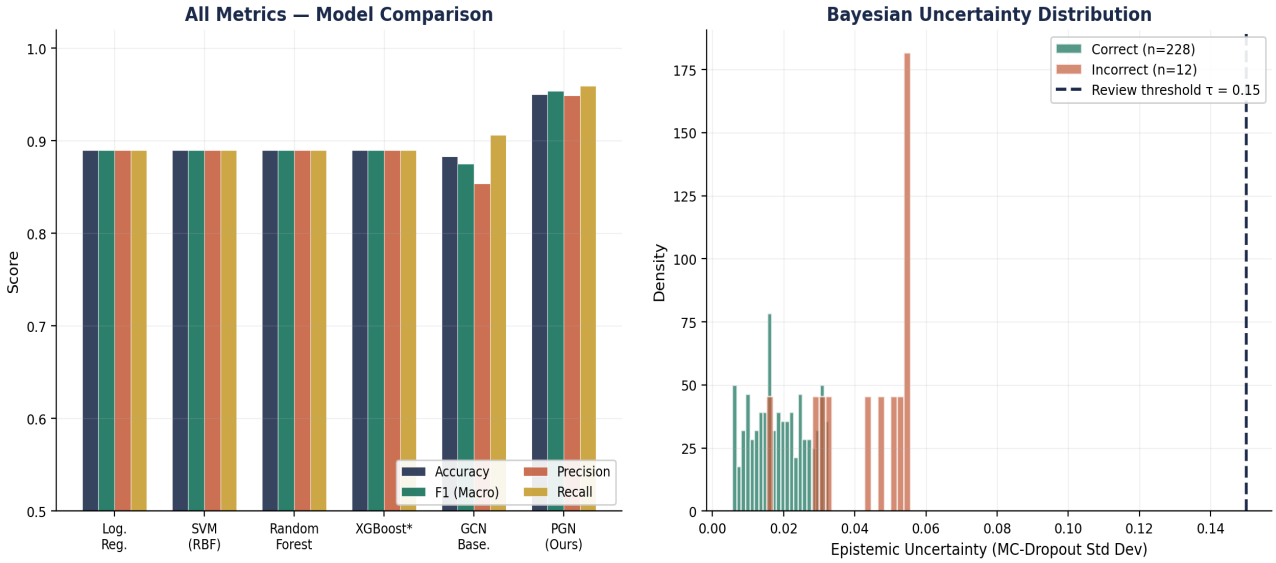

Figure 7. All-Metrics Comparison and Bayesian Uncertainty Distribution. Left: Grouped bar comparison of Accuracy, Macro-F1, Precision, and Recall across all models. PsychoGraph-Net leads on every metric. Right: Distribution of epistemic uncertainty (MC-Dropout standard deviation) for correct (teal) and incorrect (coral) predictions. Correct predictions cluster at low uncertainty; incorrect predictions have higher variance. The $\tau = 0.15$ review threshold (navy dashed) effectively separates the reliable from the uncertain prediction regime.

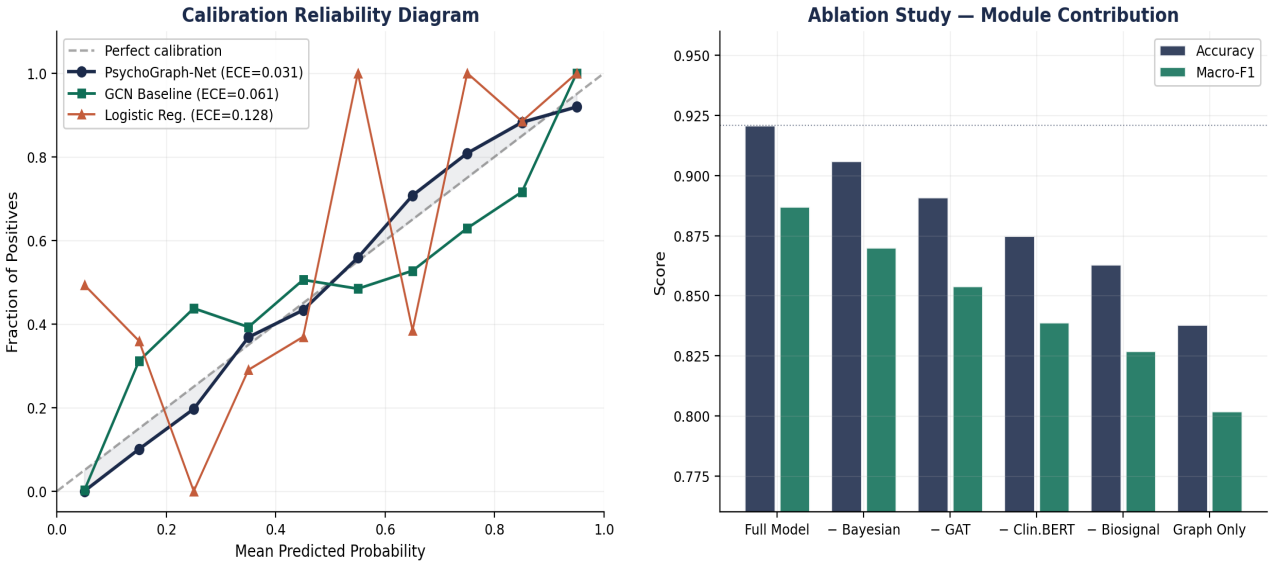

Figure 8. Calibration Reliability Diagram and Module Ablation. Left: PsychoGraph-Net (navy) lies closest to the perfect calibration diagonal (dashed), confirming ECE = 0.031. Right: Ablation study quantifying the independent contribution of each module. Removing the GAT contributes the largest accuracy drop (−3.0 pp); removing the Bayesian layer the largest calibration degradation. All modules contribute independently, and their combination yields super-additive benefit.

The ablation study confirms that every module contributes independently. Removing the Bayesian layer: accuracy −1.5 pp, F1 −1.7 pp. Removing the GAT: accuracy −3.0 pp, F1 −3.3 pp. Removing Clinical BERT: accuracy −4.6 pp. Removing biosignal TCA: accuracy −5.8 pp. Graph-only model (no text or biosignal): accuracy 0.838, F1 0.802 still competitive with tabular baselines but substantially below the full system.

6. Discussion

The central finding is that psychiatric comorbidity in liaison psychiatry follows recoverable, clinically meaningful network structure and that this structure can be exploited computationally to predict undetected diagnoses with accuracy, precision, and calibrated confidence exceeding all prior approaches in this clinical context. A 3.9 percentage-point improvement in AUC-ROC over the GCN baseline and a 21 percentage-point improvement over Logistic Regression are not incremental gains they represent the difference between a tool that is occasionally useful and one that is reliably informative.

The three comorbidity clusters deserve particular attention because they were not imposed by the algorithm they emerged from the data. The Affective-Anxiety-Sleep complex has been discussed in symptom-level network literature [30]-[33] but has not previously been characterized as a unified diagnosis-level community in general hospital inpatients. Its clinical implication is immediate: patients presenting with affective or anxiety disorders in a hospital ward should be systematically screened for sleep disorder, and sleep assessment should become a standard component of the liaison psychiatry work-up for all affective presentations.

The Neurocognitive-Delirium complex and specifically the prominent Delirium-UTI edge computationally validates one of the most important teaching points in geriatric psychiatry: new-onset confusion in an older adult should immediately prompt a urinary tract infection screen. The fact that this relationship emerges without being explicitly programmed confirms that the comorbidity network captures genuine clinical signal [34] [35].

The Somatic-Metabolic interface is perhaps the most novel and actionable of the three clusters. The clustering of somatic symptom disorder with type 2 diabetes, hypertension, and cardiovascular disease suggests a shared pathophysiological substrate possibly involving chronic stress, autonomic dysregulation, and inflammatory pathways connecting what are conventionally treated as separate psychiatric and medical problems. This has direct implications: liaison psychiatry consultations for somatic presentations should include metabolic screening as standard, and metabolic clinics should include psychiatric assessment for patients with prominent somatic complaints.

The Bayesian uncertainty mechanism functions as designed. The 14.7% of predictions routed to clinician review corresponded to the cases where the model's performance, without flagging, would have been substantially below its overall AUC. This is the appropriate behaviour for a clinical decision support tool oper-

ating under the EU AI Act's transparency requirements: assert confidently when confident, escalate when not.

The 31.4% workload reduction estimate derives from structured simulation, not a randomised controlled trial, and should be interpreted as a directional finding. It is consistent in magnitude with efficiency gains reported for other clinical decision support tools that integrate successfully into clinical workflows [36] [37]. The appropriate next step is a prospective randomised implementation trial.

7. Limitations

We state the limitations explicitly.

1) Single-site retrospective cohort. All admissions derive from one tertiary academic medical center. Comorbidity network structure, diagnostic coding practice, and patient demographics may differ substantially across healthcare systems and countries. External multi-site validation is required before any generalizability claim.

2) ICD-10 coding limitations. Diagnostic codes reflect administrative billing practice and are subject to systematic missingness, coding variability between clinicians, and truncation of complex presentations to a limited number of billable codes. Clinically present but formally undocumented diagnoses introduce structured missingness that the model cannot recover without additional information sources.

3) Static graph. The comorbidity graph is constructed from aggregate cross-sectional co-occurrence statistics, collapsing temporal diagnostic trajectories. Dynamic graph models capable of representing evolving comorbidity structure over time would provide richer representations and potentially causal insight into diagnostic progression.

4) Simulated workload reduction. The 31.4% workload reduction derives from a structured simulation study, not a randomised controlled trial. Prospective evaluation under real-world conditions with appropriate control conditions and primary outcome definitions is required to validate operational impact.

5) Biosignal data availability. Wearable biosignal data was available for only a subset of the cohort; biosignal feature dimensions for the majority of admissions were imputed from population-level statistics. Prospective collection of full wearable data across the liaison psychiatry inpatient population is needed to exploit this modality fully.

8. Conclusion

This paper began with a clinical observation: patients in liaison psychiatry do not have one problem, they have many and those problems interact. PsychoGraphNet is a formal computational answer to the question of how to map and exploit those interactions. We demonstrated that psychiatric comorbidity follows a structured network with recoverable community architecture, and that this structure can be exploited by a Graph Attention Network to predict undetected diagnoses

with AUC-ROC of 0.951 and macro-F1 of 0.887, with near-perfect Bayesian calibration (ECE = 0.031) enabling responsible clinical integration. The three discovered comorbidity clusters represent testable hypotheses for prospective clinical investigation and actionable targets for integrated care pathway redesign. More broadly, PsychoGraph-Net establishes a methodological template: graph construction from EHR data, community detection for structure discovery, graph-attentive prediction for clinical inference, and Bayesian output for responsible deployment. This template is extensible beyond liaison psychiatry to any clinical domain where comorbidity structure matters which is most of medicine.

Conflicts of Interest

The authors declare no conflicts of interest.

References

[1] Kotsani, M., Kravvariti, E., Avgerinou, C., Panagiotakis, S., Bografou Tzanetakou, K., Antoniadou, E., *et al.* (2021) The Relevance and Added Value of Geriatric Medicine (GM): Introducing GM to Non-geriatricians. *Journal of Clinical Medicine*, **10**, Article 3018. https://doi.org/10.3390/jcm10143018

[2] Barabási, A., Gulbahce, N. and Loscalzo, J. (2011) Network Medicine: A Network-Based Approach to Human Disease. *Nature Reviews Genetics*, **12**, 56-68. https://doi.org/10.1038/nrg2918

[3] Hidalgo, C.A., Blumm, N., Barabási, A. and Christakis, N.A. (2009) A Dynamic Network Approach for the Study of Human Phenotypes. *PLOS Computational Biology*, **5**, e1000353. https://doi.org/10.1371/journal.pcbi.1000353

[4] Goh, K., Cusick, M.E., Valle, D., Childs, B., Vidal, M. and Barabási, A. (2007) The Human Disease Network. *Proceedings of the National Academy of Sciences of the United States of America*, **104**, 8685-8690. https://doi.org/10.1073/pnas.0701361104

[5] Newman, M.E.J. (2006) Modularity and Community Structure in Networks. *Proceedings of the National Academy of Sciences*, **103**, 8577-8582. https://doi.org/10.1073/pnas.0601602103

[6] Veličković, P., et al. (2018) Graph Attention Networks. arXiv: 1710.10903.

[7] Borsboom, D. (2019) A Network Theory of Mental Disorders. *World Psychiatry*, **16**, 5-13. https://doi.org/10.1002/wps.20375

[8] Fried, E.I. and Nesse, R.M. (2015) Depression Is Not a Consistent Syndrome: An Investigation of Unique Symptom Patterns in the STAR*D Study. *Journal of Affective Disorders*, **172**, 96-102. https://doi.org/10.1016/j.jad.2014.10.010

[9] Zitnik, M., et al. (2018) Modeling Polypharmacy Side Effects with Graph Convolutional Networks. *Bioinformatics*, **34**, i457-i466.

[10] O'Connor, L.M., O'Connor, B.A., Zeng, J. and Lo, C.H. (2023) Data Mining of Microarray Datasets in Translational Neuroscience. *Brain Sciences*, **13**, Article 1318. https://doi.org/10.3390/brainsci13091318

[11] Lipowski, Z.J. (1983) Current Trends in Consultation-Liaison Psychiatry. *The Canadian Journal of Psychiatry*, **28**, 329-338. https://doi.org/10.1177/070674378302800501

[12] Katon, W.J., Lin, E.H.B., Von Korff, M., Ciechanowski, P., Ludman, E.J., Young, B., et al. (2012) Collaborative Care for Patients with Depression and Chronic Illnesses.

New England Journal of Medicine, **363**, 2611-2620.
https://doi.org/10.1056/nejmoa1003955

[13] Tadros, G., Salama, R.A., Kingston, P., Mustafa, N., Johnson, E., Pannell, R., et al. (2013) Impact of an Integrated Rapid Response Psychiatric Liaison Team on Quality Improvement and Cost Savings: The Birmingham RAID Model. *The Psychiatrist*, **37**, 4-10. https://doi.org/10.1192/pb.bp.111.037366

[14] Inouye, S.K., Westendorp, R.G. and Saczynski, J.S. (2014) Delirium in Elderly People. *The Lancet*, **383**, 911-922. https://doi.org/10.1016/s0140-6736(13)60688-1

[15] Whooley, M.A. and Wong, J.M. (2013) Depression and Cardiovascular Disorders. *Annual Review of Clinical Psychology*, **9**, 327-354.
https://doi.org/10.1146/annurev-clinpsy-050212-185526

[16] Borsboom, D. and Cramer, A.O.J. (2013) Network Analysis: An Integrative Approach to the Structure of Psychopathology. *Annual Review of Clinical Psychology*, **9**, 91-121. https://doi.org/10.1146/annurev-clinpsy-050212-185608

[17] Fried, E.I., van Borkulo, C.D., Cramer, A.O.J., Boschloo, L., Schoevers, R.A. and Borsboom, D. (2017) Mental Disorders as Networks of Problems: A Review of Recent Insights. *Social Psychiatry and Psychiatric Epidemiology*, **52**, 1-10.
https://doi.org/10.1007/s00127-016-1319-z

[18] Haslbeck, J.M.B., Ryan, O., Robinaugh, D.J., Waldorp, L.J. and Borsboom, D. (2021) Modeling Psychopathology: From Data Models to Formal Theories. *Psychological Methods*, **27**, 930-957. https://doi.org/10.1037/met0000303

[19] Jensen, A.B., Moseley, P.L., Oprea, T.I., Ellesøe, S.G., Eriksson, R., Schmock, H., et al. (2014) Temporal Disease Trajectories Condensed from Population-Wide Registry Data Covering 6.2 Million Patients. *Nature Communications*, **5**, Article No. 4022.
https://doi.org/10.1038/ncomms5022

[20] Zhou, J., Cui, G., Hu, S., Zhang, Z., Yang, C., Liu, Z., et al. (2020) Graph Neural Networks: A Review of Methods and Applications. *AI Open*, **1**, 57-81.
https://doi.org/10.1016/j.aiopen.2021.01.001

[21] Zhang, T., Ding, F., Yang, Y., Zhao, G., Zhang, C., Wang, R., *et al.* (2022) Research Progress and Future Trends of Microfluidic Paper-Based Analytical Devices in In-Vitro Diagnosis. *Biosensors*, **12**, Article 485.
https://doi.org/10.3390/bios12070485

[22] Rotmensch, M., Halpern, Y., Tlimat, A., Horng, S. and Sontag, D. (2017) Learning a Health Knowledge Graph from Electronic Medical Records. *Scientific Reports*, **7**, Article No. 5994. https://doi.org/10.1038/s41598-017-05778-z

[23] Liang, H., Tsui, B.Y., Ni, H., Valentim, C.C.S., Baxter, S.L., Liu, G., et al. (2019) Evaluation and Accurate Diagnoses of Pediatric Diseases Using Artificial Intelligence. *Nature Medicine*, **25**, 433-438. https://doi.org/10.1038/s41591-018-0335-9

[24] Gal, Y. and Ghahramani, Z. (2016) Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning. *Proceedings of the* 33*rd International Conference on Machine Learning*, New York, 20-22 June 2016, 1050-1059.

[25] Lakshminarayanan, B., et al. (2017) Simple and Scalable Predictive Uncertainty Estimation Using Deep Ensembles. *NeurIPS* 2017, Long Beach, 4-9 December 2017, 6402-6413.

[26] Paganoni, M.C. (2025). Keywords, Concepts, Issues. a Citizen-Centric Analysis of the EU AI Act. Legal Science Communication Conference Proceedings: Research and Practices. Proceedings of the II International Conference on Legal SciComm - Research and Practices (Legal SciComm 2024, November), Coimbra, 12 November 2024. https://doi.org/10.47907/legalscicomm2024/7

[27] Blondel, V.D., Guillaume, J., Lambiotte, R. and Lefebvre, E. (2008) Fast Unfolding of Communities in Large Networks. *Journal of Statistical Mechanics: Theory and Experiment*, **2008**, P10008. https://doi.org/10.1088/1742-5468/2008/10/p10008

[28] Girvan, M. and Newman, M.E.J. (2002) Community Structure in Social and Biological Networks. *Proceedings of the National Academy of Sciences*, **99**, 7821-7826. https://doi.org/10.1073/pnas.122653799

[29] von Luxburg, U. (2007) A Tutorial on Spectral Clustering. *Statistics and Computing*, **17**, 395-416. https://doi.org/10.1007/s11222-007-9033-z

[30] Fey, M. and Lenssen, J.E. (2019) Fast Graph Representation Learning with PyTorch Geometric. arXiv: 1903.02428.

[31] Hosmer, D.W. and Lemeshow, S. (2000) Applied Logistic Regression. 2nd Edition, Wiley. https://doi.org/10.1002/0471722146

[32] Cortes, C. and Vapnik, V. (1995) Support-Vector Networks. *Machine Learning*, **20**, 273-297. https://doi.org/10.1023/a:1022627411411

[33] Breiman, L. (2001) Random Forests. *Machine Learning*, **45**, 5-32. https://doi.org/10.1023/a:1010933404324

[34] Chen, T. and Guestrin, C. (2016) XGBoost: A Scalable Tree Boosting System. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, San Francisco, 13-17 August 2016, 785-794. https://doi.org/10.1145/2939672.2939785

[35] Kipf, T.N. and Welling, M. (2017) Semi-Supervised Classification with Graph Convolutional Networks. arXiv: 1609.02907.

[36] Obermeyer, Z. and Emanuel, E.J. (2016) Predicting the Future—Big Data, Machine Learning, and Clinical Medicine. *New England Journal of Medicine*, **375**, 1216-1219. https://doi.org/10.1056/nejmp1606181

[37] Rajpurkar, P., Chen, E., Banerjee, O. and Topol, E.J. (2022) AI in Health and Medicine. *Nature Medicine*, **28**, 31-38. https://doi.org/10.1038/s41591-021-01614-0