

000
001
002
003
004
005
006
007
008
009
010
011
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028
029
030
031
032
033
034
035
036
037
038
039
040
041
042
043
044
045
046
047
048
049
050
051
052
053

GRAPH MIXING ADDITIVE NETWORKS

Anonymous authors

Paper under double-blind review

ABSTRACT

Real-world temporal data often consists of multiple signal types recorded at irregular, asynchronous intervals. For instance, in the medical domain, different types of blood tests can be measured at different times and frequencies, resulting in fragmented and unevenly scattered temporal data. Similar issues of irregular sampling occur in other domains, such as the monitoring of large systems using event log files. Effectively learning from such data requires handling sets of temporally sparse and heterogeneous signals. In this work, we propose Graph Mixing Additive Networks (GMAN), a novel and interpretable-by-design framework for learning directly from sets of graphs that represent such signals. GMAN provides diverse interpretability capabilities, including node-level, graph-level, and subset-level importance, and enables practitioners to trade finer-grained interpretability for greater expressivity when domain priors are available. GMAN achieves state-of-the-art performance in real-world high-stakes tasks, including predicting Crohn’s disease onset and hospital length of stay from routine blood test measurements and detecting fake news. Furthermore, we demonstrate how GMAN’s interpretability properties assist in revealing disease development phase transitions and provide crucial insights in the healthcare domain.

1 INTRODUCTION

Modern clinical data consist of diverse types of signals, often collected at irregular and asynchronous time intervals. For example, a patient’s medical record may include a set of blood tests taken over their lifetime, where each type of test is performed at its own frequency. As a result, the data can be viewed as a set of sparse temporal signals, each with its own fragmented temporal structure. Similar patterns occur in other domains. For instance, the spread of news articles in social media networks often follow asynchronous tree-like patterns of dissemination. Another example is system event logs, which typically include different types of events occurring at varying times and rates.

A common approach for learning from such data is to align the signals to a fixed-size time grid, thereby enforcing a shared timeline. This is typically achieved by trimming or aggregating signals and filling in missing values through interpolation or learned imputation models (Cao et al., 2018; Tashiro et al., 2021; Wu et al., 2022; Du et al., 2023). However, these procedures can lead to substantial information loss and ignore the informative patterns in the irregularity itself, such as the varying time intervals between different measurement types.

In this work, we introduce Graph Mixing Additive Networks (GMAN), a novel framework designed to learn directly from sets of graphs. GMAN can learn over sparse, irregular, temporal heterogeneous signals, without information loss or imputation, by representing them as sets of graphs. For example, in blood test data, each biomarker is modelled as a graph whose nodes correspond to individual measurements. These graphs may differ in structure, size, and feature space, reflecting the diversity of real-world signals.

GMAN provides multiple interpretability capabilities, including node-level, graph-level, and subset-level importance. It also enables practitioners to trade fine-grained interpretability for greater expressivity when domain priors are available. Specifically, signal graphs can be grouped into subsets, shifting interpretability from individual nodes or graphs to the subset level. GMAN builds on Graph Neural Additive Networks (GNAN) (Bechler-Speicher et al., 2024), an interpretable class of GNNs that operate on individual nodes or graphs. **In contrast, GMAN is expressly designed for sets of graphs and introduces a graph-grouping mechanism that aggregates signals within subsets**

054 **while preserving an additive decomposition across subsets.** Rather than applying GNAN directly,
 055 GMAN employs an extended and more flexible variant, which we denote as ExtGNAN. **Within each**
 056 **graph, ExtGNAN generalises GNAN’s univariate feature-shape networks to multivariate feature**
 057 **groups, allowing non-linear dependencies among related features while retaining additive**
 058 **transparency at the group level.** In analogy to signal grouping, ExtGNAN supports grouping
 059 features into subsets, thereby replacing feature-level interpretability with subset-level interpretability,
 060 such as importance scores. We prove that grouping signals or features strictly increases expressivity,
 061 and that GMAN is strictly more expressive than GNAN.

062 We demonstrate the effectiveness of GMAN on real-world high-stakes medical datasets, where it
 063 achieves state-of-the-art (SoTA) performance while also providing valuable clinical and biological
 064 insights through its interpretability. In addition, we show that GMAN attains SoTA performance in
 065 fake news detection, highlighting its flexibility in operating on sets of graphs with arbitrary structure.
 066 This contrasts with existing approaches, which are restricted to sets of path-like signals and offer no
 067 interpretability. Finally, we demonstrate how GMAN’s interpretability reveals phase transitions and
 068 provides crucial insights in healthcare.

069 Our main contributions are as follows:

- 071 1. We introduce GMAN, a novel framework for learning directly from sets of sparse, irregular
 072 temporal signals, without information loss or imputation.
- 073 2. We allow practitioners to integrate domain priors, when available, by grouping features or
 074 signal types into subsets. This shifts interpretability from fine-grained (node- or feature-
 075 level) to subset-level, while strictly improving expressivity. This capability is particularly
 076 valuable in the medical domain, where such priors are common.
- 077 3. We provide a theoretical analysis proving that groupings of features and signals make
 078 GMAN strictly more expressive.
- 079 4. We demonstrate the effectiveness of GMAN on real-world high-stakes medical datasets and
 080 fake news detection, achieving state-of-the-art performance in both domains.
- 081 5. We show that GMAN provides valuable interpretability capabilities, including node-level,
 082 graph-level, and subset-level importance, which yield meaningful clinical and biological
 083 insights.

085 2 RELATED WORK

087 **Graph Neural Additive Networks** Graph Neural Networks (GNNs) (Kipf & Welling, 2016;
 088 Gilmer et al., 2017; Velickovic et al., 2017; Xu et al., 2018) have become the dominant framework for
 089 learning over graph-structured data, enabling flexible representation learning across diverse domains
 090 such as healthcare (Paul et al., 2024; Ochoa & Mustafa, 2022; Peng et al., 2023), chemistry (Reiser
 091 et al., 2022; Jumper et al., 2021) and social networks (Li et al., 2023; Sharma et al., 2024), amongst
 092 others. GNNs leverage both the graph topology and node features to compute learned representations
 093 for individual nodes or for entire graphs. Recently, Bechler-Speicher et al. (2024) introduced Graph
 094 Neural Additive Networks (GNANs) a novel interpretable-by-design graph learning framework
 095 inspired by generalized additive models (GAMs) (Hastie & Tibshirani, 1986; 1987). GNAN applies
 096 univariate neural networks to each feature of the nodes separately, and then linearly combines their
 097 outputs across nodes to produce node-level and graph-level representations. As features are not mixed
 098 non-linearly, GNAN is fully interpretable, and provides feature-level and node-level interpretability
 099 which shows exactly how each feature and each node contribute to the final target variable.

100 **Learning from sparse data** In many real-world settings, data often presents missing values from
 101 irregular sampling and variable feature availability. Recurrent models (Cao et al., 2018) treat missing
 102 data as latent variables, while attention-based methods (Du et al., 2023; Wu et al., 2022; Tipirneni
 103 & Reddy, 2022; Labach et al., 2023) reconstruct them via contextual masking and temporal blocks.
 104 Diffusion models (Tashiro et al., 2021; Alcaraz & Strodthoff, 2022; Senane et al., 2024; Dai et al.,
 105 2024) learn conditional distributions over missing values using score-based processes. Graph-based
 106 approaches use GNNs to model feature dependencies through bipartite graphs, adaptive message
 107 passing, or spatio-temporal attention (You et al., 2020; Cini et al., 2022; Marasca et al., 2022; Ye et al.,
 2021; Chen et al., 2024). These imputation methods often distort dynamics and may not improve

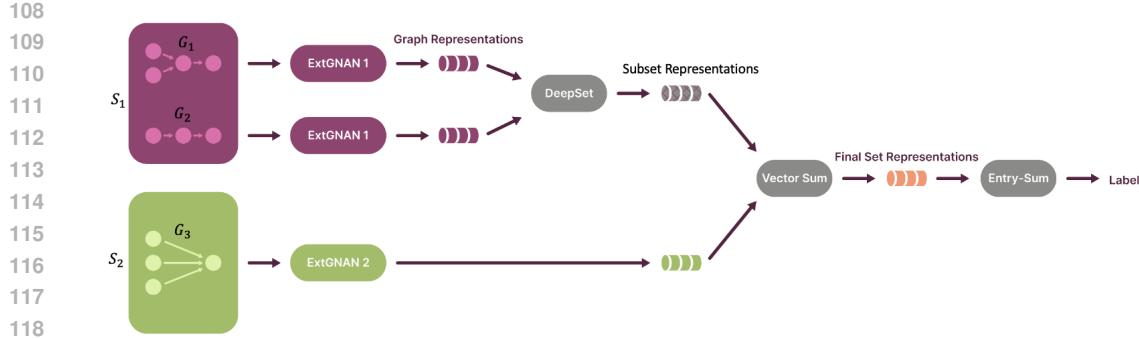


Figure 1: In this example, the input is a set of three graphs, G_1, G_2, G_3 , grouped into two subsets S_1 and S_2 . Within each subset, the same ExtGNAN instance is applied to all graphs to produce their graph-level representations. For subsets containing multiple graphs, a DeepSets module aggregates these graph representations into a single subset representation. For subsets of size one, the subset representation is simply the graph representation itself. The final set representation is then obtained by summing the subset representations, and the final label prediction is produced by summing the entries of this set representation.

prediction (Qian et al., 2025). An alternative is to model sparsity directly. Neural and Latent ODEs (Chen et al., 2018; Rubanova et al., 2019) address irregular gaps via continuous dynamics but are compute-intensive and rely on missingness encodings. Recent models (Zhang et al., 2021) have used graph representations to capture sparsity without imputation. Importantly, none of the aforementioned methods offer built-in multi-grain interpretability in the way that GMAN does.

3 GRAPH MIXING ADDITIVE NETWORKS

In this section, we present GMAN, an interpretable and flexible method for learning over sets of arbitrary graphs.

Preliminaries GMAN acts on a set of m graphs $S = \{G_1, \dots, G_m\}$, which can be directed or undirected. Each node $v \in G_i$, $1 \leq i \leq m$ is associated with a feature vector $x_v \in \mathbb{R}^d$ and a time-stamp t_v . GMAN utilizes the distances between each pair of nodes in each graph, denoted by Δ_{uv} , where

$$\Delta_{uv} = \begin{cases} t_u - t_v, & \text{if there exists a path from } u \text{ to } v \\ 0, & \text{otherwise.} \end{cases}$$

Graphs may be provided directly through an explicit layout, e.g., the news propagation graphs we evaluate GMAN over in section 4. More commonly, however, sparse temporal heterogeneous data does not come with a predefined graph structure. In such cases, we construct directed path-graphs for each signal type based on its measurement time stamps, as demonstrated for the medical dataset in section 4. For instance, in patient blood test records, each graph corresponds to a specific biomarker. Nodes in each graph represent individual measurements of that biomarker, annotated with the observed test value as a feature and the associated time stamp. We mark vectors with bold, and denote the entry c of a vector \mathbf{h} by $[\mathbf{h}]_c$, and the set of entries corresponding to a set of features S by $[\mathbf{h}]_S$.

Signal Grouping To incorporate domain priors, the graphs in S can be partitioned into k disjoint subsets S_1, \dots, S_k with $\bigcup_{i=1}^k S_i = S$. If at least one subset S' contains multiple graphs, the model's expressivity increases, as we prove in Theorem 3.2, at the cost of shifting interpretability from individual nodes and graphs in S' to the subset level. In practice, this means we can attribute importance to S' as a whole, but not to its individual components. This trade-off—enhanced expressivity at reduced granularity of interpretability—is especially valuable in domains such as medicine, where priors often suggest natural groupings of signals and interpretability is only needed at the subset level. This is demonstrated in Section 4.

162 GMAN linearly aggregates representations of the subsets of S to form a final set representation, and
 163 then assigns a single label to S .

164 First, GMAN applies a function Φ_i to each subset S_i to obtain a representation of the subset S_i ,
 165 denoted as $\mathbf{h}_i \in \mathbb{R}^d$.

$$167 \quad \mathbf{h}_i = \Phi_i(S_i),$$

168 Then, it produce a representation for the whole set, \mathbf{h}_S by summing the subsets' $\mathbf{h}_S = \sum_{i=1}^k h_i$.
 169 Finally, to produce the label, it sums over the d entries of \mathbf{h}_S . Overall:

$$170 \quad \text{GMAN}(S) = \sum_{c=1}^d \sum_{i=1}^k [\Phi_i(S_i)]_c \quad (1)$$

171 Where $\Phi_i(S_i) = \mathbf{h}_{S_i}$ is a representation of the subset S_i .

172 For subsets of size one, $\Phi_i(S_i)$ applies an Extended GNAN (EXTGNAN), as described in Section 3.1.
 173 For subsets containing multiple graphs, a featuregroupnan is applied to each graph, followed by
 174 a DeepSet aggregation (Zaheer et al., 2018) over the resulting vectors. Importantly, each subset is
 175 assigned its own EXTGNAN, and all graphs within a subset share the same one. A DeepSet first
 176 applies a neural network (NN) $f : \mathbb{R}^d \rightarrow \mathbb{R}^d$ for each vector in the set $\{h_l\}_{G_l \in S_i}$, sums the results,
 177 and then applies another NN $g : \mathbb{R}^d \rightarrow \mathbb{R}^d$.

$$178 \quad g \left(\sum_{i \in S_2} f(h_i) \right)$$

179 Here, g and f are NNs of arbitrary depth and width. A high-level visual overview of GMAN is
 180 presented in Figure 1. We now turn to define EXTGNAN.

181 3.1 EXTGNAN

182 In GNAN, univariate NNs are applied to each feature of each node in isolation, to learn a representation
 183 for a graph. This has the benefit of generating interpretable models as features do not mix non-
 184 linearly. Nonetheless, when interactions between features are crucial for the task, or feature-level
 185 interpretability is not required for all features, it may result in sub-par performance. Therefore,
 186 EXTGNAN extends GNAN by allowing multivariate NNs to operate on groups of features to gain
 187 accuracy at the cost of reducing the feature-level interpretability only for features that are grouped
 188 together, and obtaining interpretability for their subset as a whole instead.

189 Assume that the features are partitioned into K subsets $\{F_l\}_{l=1}^K$. For any subset of features greater
 190 than one, EXTGNAN applies a multivariate NN for all the features in the subset together, instead of
 191 a univariate NN for each one separately. To learn a representation of a graph G , EXTGNAN first
 192 computes representations for the nodes of G as follows.

193 EXTGNAN learns a distance function $\rho(x; \theta) : \mathbb{R} \rightarrow \mathbb{R}$ and a set of feature shape functions
 194 $\{\psi_l\}_{l=1}^K, \psi_l(X; \theta_k) : \mathbb{R}^{|F_l|} \rightarrow \mathbb{R}^{|F_l|}$. Each of these functions is a NN of arbitrary depth. For brevity,
 195 we omit the parameterization θ and θ_k for the remainder of this section.

196 The entries of the representation of node j corresponding to the indices of the features in F_l , denoted
 197 as $[\mathbf{h}_j]_{F_l}$, is computed by summing the contributions of the features in the subset F_l from all nodes in
 198 the graph:

$$199 \quad [\mathbf{h}_j]_{F_l} = \sum_{w \in V} \rho(\Delta(w, j)) \cdot \psi_l([\mathbf{X}_w]_{F_l}),$$

200 where $\Delta(w, j) = t_w - t_j$ and $[\mathbf{X}_w]_{F_l}$ are the features of node w corresponding to the subset F_l .

201 Overall, the full representation of node j can be written as:

$$202 \quad \mathbf{h}_j = ([\mathbf{h}_j]_{F_1}, [\mathbf{h}_j]_{F_2}, \dots, [\mathbf{h}_j]_{F_K}).$$

216 Then EXTGNAN produces a graph representation by summing the node representations,
 217
 218
 219

$$\mathbf{h}_G = \sum_{i \in V} \mathbf{h}_i. \quad (2)$$

220 This concludes the description of EXTGNAN, which computes graph-level representations. We
 221 provide a complexity analysis of GMAN in the appendix, where we also discuss how ρ can be
 222 masked to adapt to any complexity limitation, including a linear one. Next, we describe how GMAN
 223 combines these representations across sets and enables multi-level interpretability.
 224

225 3.2 NODE, GRAPH AND SUBSET IMPORTANCE

226 GMAN retains all interpretability properties of GNAN, including feature-level and node-level
 227 importance. However, it extends beyond GNAN by operating on sets of graphs rather than single
 228 graphs, enabling additional forms of interpretability such as graph-level and subset-level importance.
 229 Because GMAN allows a flexible trade-off between interpretability and expressivity, permitting
 230 non-linear mixing within graph subsets, some adaptations are required to obtain importance scores. A
 231 key property of GMAN’s interpretability is that its importance scores directly reflect the contribution
 232 of each node, graph, or subset to the predicted label, since these terms are combined additively to
 233 produce the final output. Section 4 illustrates how node-level importances yield insightful real-world
 234 insights.

235 We can extract the total contribution of each node j to the prediction by summing the contributions of
 236 the node across all feature subsets. This is only valid when the node belongs to a graph that is not
 237 combined non-linearly with other graphs, i.e., it belongs to a subset of size one.
 238

239 Therefore, the contribution of node j is

$$\text{TotalContribution}(j) = \sum_{l=1}^K [\mathbf{h}_j]_{F_k} = \sum_{w \in V} \rho(\Delta(w, j)) \sum_{l=1}^K \psi_k([\mathbf{x}_w]_l, l \in F_k). \quad (3)$$

240 The contribution of a graph G is then
 241
 242

$$\text{TotalContribution}(G) = \sum_{v \in G} \text{TotalContribution}(v).$$

243 For graphs that are mixed non-linearly, i.e., graphs that belong in subsets of size greater than one, we
 244 provide instead the total contribution of the set to the final prediction
 245
 246

$$\text{TotalContribution}(S) = \sum_{l=1}^K [\mathbf{S}]_{F_k}. \quad (4)$$

247 3.3 EXPRESSIVITY PROPERTIES

248 In this section, we provide a theoretical analysis of the expressiveness of GMAN. Proofs are provided
 249 in the Appendix

250 **Theorem 3.1.** GMAN is strictly more expressive than GNAN.

251 The following theorem shows that a GMAN which is applied to subsets of graphs of size at least two,
 252 is more expressive than a GMAN that is applied to only subsets of size one:

253 **Theorem 3.2.** Let S be a set of graphs $\{G_i\}_{i=1}^m$. Let $S_1 = \{S_i\}_{i=1}^m$ be a partition of S such that
 254 $|S_i| = 1$. Let $S_2 = \{S_i\}_{i=1}^k$ such that there exists k with $|S_k| > 1$. with a subset partition $\{S_i\}_{i=1}^k$.
 255 Then a GMAN trained over S_2 is strictly more expressive than a GMAN trained over S_1 .

256 4 EMPIRICAL EVALUATION

257 In this section, we evaluate GMAN on real-world tasks, and demonstrate its interpretability properties
 258 ¹.

259 ¹Code is provided in Supplementary Materials

270 4.1 MEDICAL PREDICTIONS
271

272 We evaluate GMAN on two high-impact clinical prediction tasks: LoS of intensive care patients
273 and onset of CD. In both settings, each individual is represented as a set of time-stamped biomarker
274 trajectories, where each trajectory forms a directed path graph with nodes corresponding to test
275 results and edges encoding the time elapsed between measurements. This representation preserves the
276 temporal structure of each biomarker independently while enabling joint reasoning across biomarkers
277 during learning.

278 **Data** Next, we describe the two medical datasets used in our evaluation.
279

280 *P12 (ICU Length of Stay):* The PhysioNet2012 (P12) dataset, introduced by Goldberger et al. (2000),
281 contains records from 11,988 intensive-care unit (ICU) patients, following the exclusion of 12
282 samples deemed inappropriate according to the criteria in Horn et al. (2020). For each patient,
283 longitudinal measurements from 36 physiological signals were recorded over the initial 48 hours
284 of ICU admission. Additionally, each patient has a static profile comprising 9 features, including
285 demographic and clinical attributes such as age and gender. The dataset is labelled for a binary
286 classification task: predicting whether or not the total LoS in the ICU exceeded 72 hours. **The dataset**
287 **is highly imbalanced, with ~93% positive samples. To balance the classes we perform batch**
288 **minority class upsampling, following the example in Zhang et al. (2021).**

289 *Crohn’s Disease (CD Onset):* The Danish health registries (DHR) are comprehensive, nationwide
290 databases covering healthcare interactions for over 9.5 million individuals (Pedersen, 2011). A key
291 resource is the Registry of Laboratory Results for Research (RLRR), which has collected laboratory
292 test results from hospitals and general practitioners since 2015 (Arendt et al., 2020). We first identify
293 8,567 individuals in the DHR who are later diagnosed with CD and use them as our patient class. We
294 then construct a control pool by randomly sampling individuals from the DHR and downsampling by
295 age to match the expected frequency of blood tests in the prediagnostic period, reflecting that CD
296 typically manifests in early adulthood (20–30 years). From this age-matched control pool, we sample
297 8,567 controls, yielding two balanced classes. For each person, we extracted temporal trajectories
298 of 17 routinely measured biomarkers, reflecting key physiological processes. The complete list and
299 descriptions of these biomarkers are provided in the Appendix. The task is binary classification:
300 predicting future CD onset from pre-diagnostic medical histories.

301 **Setup** We compare GMAN to 8
302 **baselines from Zhang et al. (2021),**
303 **spanning both sequential and graph-**
304 **based models, including:** *Trans-*
305 *former* (Vaswani et al., 2017), a
306 **vanilla self-attention model applied**
307 **to irregular time series;** *Trans-mean*,
308 **which combines a Transformer with**
309 **mean imputation of missing values;**
310 **GRU-D** (Che et al., 2016), a gated
311 **recurrent model with decay terms**
312 **that encode informative missingness;**
313 **SeFT** (Horn et al., 2020), which treats
314 **each record as a set of timestamped**
315 **feature-value pairs and aggregates**
316 **them with permutation-invariant en-**
317 **coders;** *mTAND* (Shukla & Marlin,
318 **2021), a multi-time attention archi-**
319 **ecture that outperforms a broad range of RNN- and ODE-based models on irregular data;**
320 ***DGM*²** (Wu et al., 2021) and *MTGNN* (Wu et al., 2020), graph-based method originally proposed
321 **for multivariate time-series forecasting; and** *Raindrop* (Zhang et al., 2021), a state-of-the-art
322 **graph model for sparse, irregular EHR time series.** For the P12 we used the splits as in (Zhang
323 et al., 2021). For CD, we randomly split the data into train (80%), validation (10%), and test (10%)
324 sets. We conducted a grid search by training on the training set and evaluating on the validation set.
325 We then selected the best performing model over the validation set and report results over the test set.
326 We define each individual biomarker as a unique signal type and group them according to **coarse**

Table 1: Evaluation of GMAN on two real-world medical tasks. The metric reported is the mean AUPRC with standard deviation, calculated over 3 random seeds.

Methods	LoS in ICU	CD onset
Transformer	96.06 ± 0.32	75.60 ± 0.52
Trans-mean	96.44 ± 0.17	75.96 ± 0.92
GRU-D	95.91 ± 2.10	83.36 ± 0.40
SeFT	95.89 ± 0.08	71.22 ± 2.30
mTAND	93.02 ± 1.04	83.17 ± 0.67
DGM ²	97.00 ± 0.40	83.02 ± 0.56
MTGNN	96.20 ± 0.78	75.26 ± 3.04
RAINDROP	96.32 ± 0.13	82.60 ± 0.82
GMAN	97.41 ± 0.38	83.93 ± 0.27

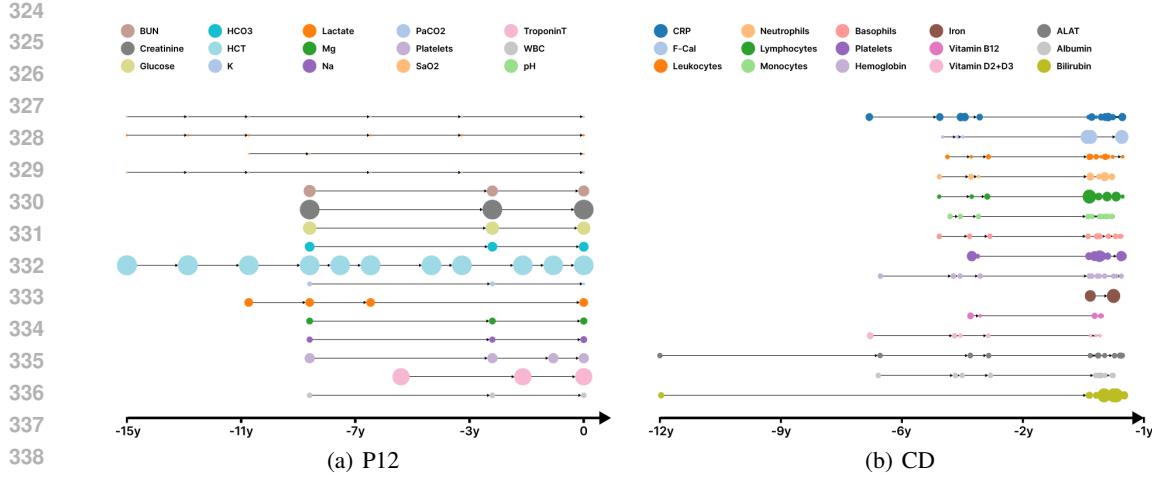


Figure 2: Node-level importances for two individuals from (a) the P12 ICU LoS and (b) CD onset datasets. Node size indicates the exact node (measurement) contribution to the prediction.

physiological categories. Importantly, these clinically-guided groups are based on broad domain knowledge and do not rely on specialized expertise. We also examine data-driven grouping that does not rely on any priors, based on the insights presented in Subsection 4.1.1. Due to space limitations these are differed to the Appendix (Table 4, App. F). We tune the biomarker subsets over no grouping at all (i.e, full interpretability), and 5 additional subset groupings motivated by public common clinical knowledge. To account for class imbalance, and as commonly done for these datasets, we report the average AUPRC score and standard-deviation of the selected configuration with 3 random seeds. Additional details on the datasets, experimental setup, subset groupings, and hyperparameter configurations are provided in the Appendix We also present Calibration and robustness evaluation in Appendix G.

Results The results in Table 1 show that GMAN achieves the highest AUPRC across both medical prediction tasks. On Length of Stay in ICU, it improves upon the best baseline with a relative uplift of about 0.41 points, while on Crohn’s Disease onset, it outperforms the best baseline by 0.57 points. Notably, these improvements are achieved while GMAN also provides interpretability properties, as shown in detail in the next subsection.

4.1.1 CLINICAL INSIGHTS THROUGH INTERPRETABILITY

Beyond predictive performance, a central strength of GMAN is that interpretability is built into the model by design, rather than added post hoc. This allows for fine-grained, temporally resolved explanations that go beyond global feature importance, enabling users to understand how and when specific biomarkers influence the model’s predictions. In high-stakes domains such as healthcare, this level of transparency is crucial. Clinical decision-making often depends not only on the outcome of a prediction but on a clear understanding of the reasoning behind it. Models that can provide such insights are far more likely to be trusted, audited, and integrated into clinical workflows. To highlight the interpretability of GMAN, we conduct attribution analyses on both the CD and P12 datasets, examining how the model assigns importance across time and signals.

Critical phase detection through node-level importance In the GMAN framework, nodes represent individual signals within a biomarker trajectory. As such, highly influential nodes can highlight critical phases where specific signals most strongly impact the prediction. We use Equation (3) to quantify node-level contributions across biomarkers in both the CD and P12 datasets. To comply with privacy and data protection legislation, CD data is anonymised via noise and temporal shifting. Figure 2 displays node-level importance across biomarker trajectories for two randomly selected individuals from each dataset, with node size reflecting the magnitude of each node’s contribution to the model’s prediction. In both clinical tasks, the model’s attributions appear consistent with established biomedical knowledge. For CD prediction, GMAN highlights key inflammatory and

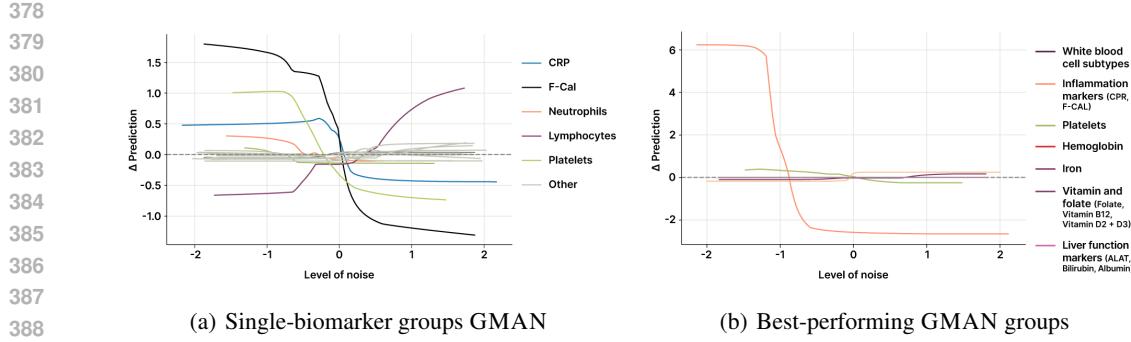


Figure 3: Subset-level contribution curves for Crohn’s Disease prediction. Each curve shows how the GMAN’s output changes as increasing noise is added to the latent representation of a biomarker group. (a) uses individual biomarkers; (b) uses physiologically coherent groups.

immune markers, such as F-Cal, platelets, and lymphocytes, as primary contributors. All of these are known to play central roles in disease onset (Vestergaard et al., 2023). In the P12 example, the model assigns high importance to markers of renal function, liver injury, cardiac stress, and metabolic imbalance, aligning well with clinical predictors of severity in intensive care settings.

Total biomarker contribution In addition to node-level importances, GMAN provides subset-level importance scores, quantifying the contribution of entire biomarker groups to the model’s prediction. This enables flexible, system-level interpretability, allowing users to assess the collective influence of physiologically related biomarkers on the target variable. We conduct a subset-level importance analysis on the CD dataset using two clinically motivated grouping strategies. In the first, we assess individual biomarkers to determine whether the model prioritizes features known to be associated with CD onset, allowing us to verify its alignment with established biomedical knowledge. In the second, we group biomarkers into clinically coherent subsets based on physiological function, such as immune response; inflammation; oxygen transport; and liver function; to examine whether the model captures system-level patterns consistent with disease progression. Full details on the clinical significance of these groupings are provided in the Appendix.

For each grouping strategy, we trained a separate instance of GMAN. To quantify the importance of each biomarker subset, we use Equation (4), measuring how the model’s prediction changes when increasing the noise added to the features of nodes within that subset. To introduce noise in a structured way, we apply PCA (Abdi & Williams, 2010) to the feature vectors of all nodes corresponding to biomarkers in the group, and progressively perturb the input along the 1st principal component. This procedure is performed independently for each subset and allows us to assess the model’s sensitivity to perturbations in biologically meaningful groupings, offering insight into the relative predictive weight of each group. Importantly, GMAN is interpretable by design, which makes perturbation analysis directly reflect its internal computation. Unlike post-hoc attribution, we do not estimate influence indirectly but observe how the additive contribution changes. Adding structured noise (via PCA) and measuring the resulting prediction shift faithfully traces the model’s internal mechanism rather than relying on an external approximation. Thus, perturbation effects align exactly with feature importance, making the experiment both natural and principled. Figure 3 presents the results of the subset-level attribution analysis. In the single-biomarker setting (Figure 3(a)), F-Cal, platelets, and lymphocytes show strong directional effects on model output. F-Cal and platelets are positively associated with CD risk, while lymphocytes have an inverse effect, findings that are both biologically grounded and consistent with prior work (Vestergaard et al., 2023). In the clinically-coherent group setting (Figure 3(b)), the inflammation subset, emerges as the most influential, with a pronounced non-linear effect on predictions, aligning with established diagnostic relevance in CD (Vestergaard et al., 2023). **We note that the above also serves as a quantitative measure of faithfulness, as in GMAN, interpretability is not post-hoc but built into the model architecture by design. Specifically, the contributions of nodes, graphs, or subsets are explicitly and additively used in computing the final prediction. This means that importance scores correspond directly to the actual values that are summed to produce the output label, making them faithful by construction.**



Figure 4: Node importance for fake-news spread graphs, over the GOS dataset. The node size corresponds to its importance learned by tGMAN, according to Equation (3). All graphs with a single node are grouped into one subset. Therefore, the importance is provided on the subset level rather than the node level (green node).

4.2 FAKE-NEWS DETECTION

Data GossipCop (GOD) is a dataset of news articles annotated by professional journalists and fact-checkers, containing both content-based labels and social context information verified through the GossipCop fact-checking platform. It is composed of 5,464 tree-structured graphs based on sharing information, where the news article is the root node and sharing users are subsequent nodes in the cascade, with edges signifying sharing relationships. Each node in these graphs is associated with 4 features: 768-dimensional embeddings generated using a pretrained BERT model on user historical posts, 300-dimensional embeddings from a pretrained word2vec (Mikolov et al., 2013) on the same historical posts, 10-dimensional features extracted from user profiles, and a 310-dimensional “content” feature that combines the 300-dimensional embedding of user comments with the 10-dimensional profile features. Each graph is labelled to indicate whether it originates from a fake news post or not. We decomposed the tree into a set of directed graphs rooted at the origin, each representing a distinct path of news propagation.

Setup While GMAN can be applied to sets of graphs with any structure, the baselines used in the CD and P12 experiments are limited to path-like graphs and are unable to act on more complex graphs as in this dataset. Therefore, we instead evaluate GMAN against 4 GNNs, including GATv2 (Brody et al., 2022), GraphConv (Morris et al., 2021), GraphSAGE (Hamilton et al., 2018) and GCNFN (Monti et al., 2019). We use random splits of train (80%), validation (10%), and test (10%) sets over the data and selected the best performing model over the validation set. We then report the average Accuracy score and std of the selected configuration with three seeds. Since many subgraphs reduce to a single node after decomposition, we group all size-one graphs into a shared subset and combine them non-linearly. For the remaining graphs, whose identities are not uniquely distinguishable, we apply a shared EXTGNAN. In total, we use two distinct EXTGNAN instances for this experiment.

Results and node-level importance Results are provided in Table 2. GMAN outperforms all baselines, with an uplift of 0.57 accuracy points. Figure 4 presents the node importance of a random sample from the datasets, where the size of a node corresponds to its importance score according to Equation (3).

4.2.1 ABLATION STUDY

Finally, we carry out an ablation study on the CD dataset, to isolate the impact of key components on performance. Specifically, we take the best-performing configuration from the training-data grid search and re-train it with individual components ablated to assess their effect on performance. We test: (i) replacing DeepSet with mean pooling to assess the value of learned non-linear aggregation; (ii) replacing the distance function NN (ρ) with a constant 1 value.

This tests the usefulness of structural information; (iii) substituting ExtGNAN with a node-wise MLP to test the importance of graph inductive bias; (iv) replacing ExtGNAN with an identity mapping as a lower bound without feature learning; and (v) using standard GNAN (no multivariate feature

Table 2: Evaluation of GMAN on the Gossipcop (GOS) fake news detection dataset.

Methods	Accuracy
GATv2	96.10 ± 0.3
GraphConv	96.77 ± 0.1
GraphSAGE	94.45 ± 1.5
GCNFN	96.52 ± 0.2
GMAN	97.34 ± 0.2

486 groups) to evaluate the benefit of grouped feature processing. We report the performance degradation
 487 (AUPRC difference) in Table 3. The ablation results demonstrate that the core components of GMAN
 488 are critical to its effectiveness, as their removal consistently leads to notable performance degradation.
 489

490 5 CONCLUSION 491

492 We introduced Graph Mixing Additive Net-
 493 works (GMAN), a framework for learning
 494 from sets of graphs that represent irregular
 495 and asynchronous temporal signals. GMAN
 496 was designed to handle real-world scenar-
 497 ios where multiple signal types are collected
 498 at uneven intervals, such as medical records
 499 with heterogeneous blood tests or event logs
 500 in complex systems. The framework com-
 501 bined strong predictive performance with
 502 built-in interpretability, offering importance
 503 scores at the node, graph, and subset levels. GMAN allows practitioners to integrate domain priors
 504 when available, trading fine-grained interpretability for greater expressivity. Across experiments on
 505 real-world high-stakes tasks, GMAN achieved state-of-the-art results. Beyond predictive accuracy,
 506 GMAN’s interpretability capabilities proved particularly valuable in domains like healthcare, where
 507 uncovering phase transitions and providing actionable insights is critical.
 508
 509
 510
 511
 512
 513
 514
 515
 516
 517
 518
 519
 520
 521
 522
 523
 524
 525
 526
 527
 528
 529
 530
 531
 532
 533
 534
 535
 536
 537
 538
 539

Table 3: Ablation study of GMAN components.

Ablation	AUPRC drop
(i) DeepSet→mean pooling	- $19.98\% \pm .28\%$
(ii) $\rho \rightarrow 1$	- $12.39\% \pm 1.39\%$
(iii) ExtGNAN→MLP	- $15.00\% \pm 2.09\%$
(iv) ExtGNAN→Identity	- $17.70\% \pm 0.15\%$
(v) ExtGNAN→GNAN	- $4.38\% \pm 2.85\%$

540 REPRODUCIBILITY STATEMENT
541542 The full code is provided in the Supplementary Material. We will release our code upon acceptance,
543 including all training and evaluation scripts. The README file in the code provides all necessary
544 details to ensure our results and evaluations can be reproduced faithfully.
545546 REFERENCES
547548 Hervé Abdi and Lynne J Williams. Principal component analysis. *Wiley interdisciplinary reviews: 549 computational statistics*, 2(4):433–459, 2010.551 Juan Miguel Lopez Alcaraz and Nils Strodtthoff. Diffusion-based time series imputation and forecast-
552 ing with structured state space models. *arXiv preprint arXiv:2208.09399*, 2022.553 Johan Frederik Håkonsen Arendt, Anette Tarp Hansen, Søren Andreas Ladefoged, Henrik Toft
554 Sørensen, Lars Pedersen, and Kasper Adelborg. Existing data sources in clinical epidemiology:
555 laboratory information system databases in denmark. *Clinical epidemiology*, pp. 469–475, 2020.
556557 Maya Bechler-Speicher, Amir Globerson, and Ran Gilad-Bachrach. The intelligible and effective
558 graph neural additive network. *Advances in Neural Information Processing Systems*, 37:90552–
559 90578, 2024.560 Shaked Brody, Uri Alon, and Eran Yahav. How attentive are graph attention networks?, 2022. URL
561 <https://arxiv.org/abs/2105.14491>.563 Peter Buneman. A note on the metric properties of trees. *J. Combin. Theory Ser. B*, 17(1):48–50,
564 1974.566 Wei Cao, Dong Wang, Jian Li, Hao Zhou, Lei Li, and Yitan Li. Brits: Bidirectional recurrent
567 imputation for time series. *Advances in neural information processing systems*, 31, 2018.569 Zhengping Che, Sanjay Purushotham, Kyunghyun Cho, David Sontag, and Yan Liu. Recurrent neural
570 networks for multivariate time series with missing values, 2016. URL <https://arxiv.org/abs/1606.01865>.572 Qianyi Chen, Jiannong Cao, Yu Yang, Wanyu Lin, Sumei Wang, and Youwu Wang. Multistage
573 graph convolutional network with spatial attention for multivariate time series imputation. *IEEE
574 Transactions on Neural Networks and Learning Systems*, 2024.576 Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary
577 differential equations. *Advances in neural information processing systems*, 31, 2018.578 Andrea Cini, Ivan Marisca, and Cesare Alippi. Filling the g_ap_s: Multivariate time series imputation
579 by graph neural networks, 2022. URL <https://arxiv.org/abs/2108.00298>.581 Zongyu Dai, Emily Getzen, and Qi Long. Sadi: Similarity-aware diffusion model-based imputation
582 for incomplete temporal ehr data. In *International Conference on Artificial Intelligence and
583 Statistics*, pp. 4195–4203. PMLR, 2024.585 Wenjie Du, David Côté, and Yan Liu. Saits: Self-attention-based imputation for time series. *Expert
586 Systems with Applications*, 219:119619, 2023.587 Manik Garg, Marcin Karpinski, Dorota Matelska, Lawrence Middleton, Oliver S Burren, Fengyuan
588 Hu, Eleanor Wheeler, Katherine R Smith, Margarete A Fabre, Jonathan Mitchell, et al. Disease
589 prediction with multi-omics and biomarkers empowers case-control genetic discoveries in the uk
590 biobank. *Nature Genetics*, 56(9):1821–1831, 2024.592 Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural
593 message passing for quantum chemistry. In *International conference on machine learning*, pp.
1263–1272. PMLR, 2017.

594 Ary L Goldberger, Luis AN Amaral, Leon Glass, Jeffrey M Hausdorff, Plamen Ch Ivanov, Roger G
595 Mark, Joseph E Mietus, George B Moody, Chung-Kang Peng, and H Eugene Stanley. Physiobank,
596 physiotoolkit, and physionet: components of a new research resource for complex physiologic
597 signals. *Circulation*, 101(23):e215–e220, 2000.

598 William L. Hamilton, Rex Ying, and Jure Leskovec. Inductive representation learning on large graphs,
599 2018. URL <https://arxiv.org/abs/1706.02216>.

600 Trevor Hastie and Robert Tibshirani. Generalized additive models. *Statistical science*, 1(3):297–310,
601 1986.

602 Trevor Hastie and Robert Tibshirani. Generalized additive models: some applications. *Journal of the
603 American Statistical Association*, 82(398):371–386, 1987.

604 Max Horn, Michael Moor, Christian Bock, Bastian Rieck, and Karsten Borgwardt. Set functions for
605 time series, 2020. URL <https://arxiv.org/abs/1909.12064>.

606 John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger,
607 Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate
608 protein structure prediction with alphafold. *Nature*, 596(7873):583–589, 2021.

609 Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks.
610 *arXiv preprint arXiv:1609.02907*, 2016.

611 Alex Labach, Aslesha Pokhrel, Xiao Shi Huang, Saba Zuberi, Seung Eun Yi, Maksims Volkovs, Tomi
612 Poutanen, and Rahul G Krishnan. Duett: dual event time transformer for electronic health records.
613 In *Machine Learning for Healthcare Conference*, pp. 403–422. PMLR, 2023.

614 Xiao Li, Li Sun, Mengjie Ling, and Yan Peng. A survey of graph neural network based recommenda-
615 tion in social networks. *Neurocomputing*, 549:126441, 2023.

616 Ivan Marasca, Andrea Cini, and Cesare Alippi. Learning to reconstruct missing data from spatiotem-
617 poral graphs with sparse observations. *Advances in neural information processing systems*, 35:
618 32069–32082, 2022.

619 Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representa-
620 tions in vector space. *arXiv preprint arXiv:1301.3781*, 2013.

621 Federico Monti, Fabrizio Frasca, Davide Eynard, Damon Mannion, and Michael M. Bronstein. Fake
622 news detection on social media using geometric deep learning, 2019. URL <https://arxiv.org/abs/1902.06673>.

623 Christopher Morris, Martin Ritzert, Matthias Fey, William L. Hamilton, Jan Eric Lenssen, Gaurav
624 Rattan, and Martin Grohe. Weisfeiler and leman go neural: Higher-order graph neural networks,
625 2021.

626 Juan G Diaz Ochoa and Faizan E Mustafa. Graph neural network modelling as a potentially effective
627 method for predicting and analyzing procedures based on patients’ diagnoses. *Artificial Intelligence
628 in Medicine*, 131:102359, 2022.

629 Showmick Guha Paul, Arpa Saha, Md Zahid Hasan, Sheak Rashed Haider Noori, and Ahmed
630 Moustafa. A systematic review of graph neural network in healthcare-based applications: Recent
631 advances, trends, and future directions. *IEEE Access*, 12:15145–15170, 2024.

632 Carsten Bøcker Pedersen. The danish civil registration system. *Scandinavian journal of public health*,
633 39(7_suppl):22–25, 2011.

634 Wei Peng, Tielin Chen, Hancheng Liu, Wei Dai, Ning Yu, and Wei Lan. Improving drug response
635 prediction based on two-space graph convolution. *Computers in Biology and Medicine*, 158:
636 106859, 2023.

637 Linglong Qian, Yiyuan Yang, Wenjie Du, Jun Wang, Richard Dobsoni, and Zina Ibrahim. Beyond
638 random missingness: Clinically rethinking for healthcare time series imputation, 2025. URL
639 <https://arxiv.org/abs/2405.17508>.

648 Patrick Reiser, Marlen Neubert, André Eberhard, Luca Torresi, Chen Zhou, Chen Shao, Houssam
 649 Metni, Clint van Hoesel, Henrik Schopmans, Timo Sommer, et al. Graph neural networks for
 650 materials science and chemistry. *Communications Materials*, 3(1):93, 2022.

651

652 Yulia Rubanova, Ricky TQ Chen, and David K Duvenaud. Latent ordinary differential equations for
 653 irregularly-sampled time series. *Advances in neural information processing systems*, 32, 2019.

654 Zineb Senane, Lele Cao, Valentin Leonhard Buchner, Yusuke Tashiro, Lei You, Paweł Andrzej
 655 Herman, Mats Nordahl, Ruibo Tu, and Vilhelm Von Ehrenheim. Self-supervised learning of
 656 time series representation via diffusion process and imputation-interpolation-forecasting mask. In
 657 *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp.
 658 2560–2571, 2024.

659 Kartik Sharma, Yeon-Chang Lee, Sivagami Nambi, Aditya Salian, Shlok Shah, Sang-Wook Kim,
 660 and Srijan Kumar. A survey of graph neural networks for social recommender systems. *ACM
 661 Computing Surveys*, 56(10):1–34, 2024.

662

663 Satya Narayan Shukla and Benjamin M. Marlin. Multi-time attention networks for irregularly
 664 sampled time series, 2021. URL <https://arxiv.org/abs/2101.10318>.

665 Yusuke Tashiro, Jiaming Song, Yang Song, and Stefano Ermon. Csdi: Conditional score-based diffu-
 666 sion models for probabilistic time series imputation. *Advances in neural information processing
 667 systems*, 34:24804–24816, 2021.

668

669 Sindhu Tipirneni and Chandan K Reddy. Self-supervised transformer for sparse and irregularly
 670 sampled multivariate clinical time-series. *ACM Transactions on Knowledge Discovery from Data
 (TKDD)*, 16(6):1–17, 2022.

671

672 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz
 673 Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing
 674 systems*, 30, 2017.

675 Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, Yoshua Bengio,
 676 et al. Graph attention networks. *stat*, 1050(20):10–48550, 2017.

677

678 Marie Vibeke Vestergaard, Kristine H Allin, Gry J Poulsen, James C Lee, and Tine Jess. Char-
 679 acterizing the pre-clinical phase of inflammatory bowel disease. *Cell Reports Medicine*, 4(11),
 680 2023.

681 Jakob Woerner, Thomas Westbrook, Seokho Jeong, Manu Shivakumar, Allison R Greenplate,
 682 Sokratis A Apostolidis, Seunggeun Lee, Yonghyun Nam, and Dokyoon Kim. Plasma protein-based
 683 and polygenic risk scores serve complementary roles in predicting inflammatory bowel disease. In
 684 *Pacific Symposium on Biocomputing. Pacific Symposium on Biocomputing*, volume 30, pp. 522,
 685 2025.

686

687 Haixu Wu, Tengge Hu, Yong Liu, Hang Zhou, Jianmin Wang, and Mingsheng Long. Timesnet:
 688 Temporal 2d-variation modeling for general time series analysis. *arXiv preprint arXiv:2210.02186*,
 689 2022.

690

691 Yinjun Wu, Jingchao Ni, Wei Cheng, Bo Zong, Dongjin Song, Zhengzhang Chen, Yanchi Liu,
 692 Xuchao Zhang, Haifeng Chen, and Susan Davidson. Dynamic gaussian mixture based deep
 693 generative model for robust forecasting on sparse multivariate time series, 2021. URL <https://arxiv.org/abs/2103.02164>.

694

695 Zonghan Wu, Shirui Pan, Guodong Long, Jing Jiang, Xiaojun Chang, and Chengqi Zhang. Connecting
 696 the dots: Multivariate time series forecasting with graph neural networks, 2020. URL <https://arxiv.org/abs/2005.11650>.

697

698 Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural
 699 networks? *arXiv preprint arXiv:1810.00826*, 2018.

700

701 Yongchao Ye, Shiyao Zhang, and James JQ Yu. Spatial-temporal traffic data imputation via graph
 702 attention convolutional network. In *International Conference on artificial neural networks*, pp.
 241–252. Springer, 2021.

702 Jiaxuan You, Xiaobai Ma, Yi Ding, Mykel J Kochenderfer, and Jure Leskovec. Handling missing
 703 data with graph representation learning. *Advances in Neural Information Processing Systems*, 33:
 704 19075–19087, 2020.

705 Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Ruslan Salakhutdinov, and
 706 Alexander Smola. Deep sets, 2018. URL <https://arxiv.org/abs/1703.06114>.

708 Xiang Zhang, Marko Zeman, Theodoros Tsiligkaridis, and Marinka Zitnik. Graph-guided network
 709 for irregularly sampled multivariate time series. *arXiv preprint arXiv:2110.05357*, 2021.

712 A APPENDIX

714 B THEORETICAL FRAMEWORK

716 B.1 PROOF OF THEOREM 3.1

718 We will prove that GMAN is strictly more expressive than GNAN. To prove this, we use a ground
 719 truth function that is a feature-level XOR. Let a single-node graph be endowed with binary features
 720 $x = (x_1, x_2) \in \{0, 1\}^2$ and define the target $f_{\oplus}(x) = x_1 \oplus x_2$.

721 First we will show that GNAN cannot express f_{\oplus} . A GNAN scores the graph by $\hat{y} = \sigma(\phi_1(x_1) +$
 722 $\phi_2(x_2))$, where each ϕ_i is univariate. Put $a = \phi_1(0)$, $b = \phi_1(1)$, $c = \phi_2(0)$, $d = \phi_2(1)$. To match
 723 the XOR truth-table there must exist a threshold τ such that

$$725 \quad a + c < \tau, \quad b + c > \tau, \quad a + d > \tau, \quad b + d < \tau.$$

726 Summing the first and last inequalities yields $a + b + c + d < 2\tau$, while the middle pair gives
 727 $a + b + c + d > 2\tau$ —a contradiction. Thus no GNAN realises f_{\oplus} .

729 Now we will show that GMAN can express f_{\oplus} . Place the two features in the same subset $F =$
 730 $\{x_1, x_2\}$ and choose the subset-network

$$731 \quad \psi_F(x_1, x_2) = x_1 + x_2 - 2x_1x_2.$$

733 For the four binary inputs this mapping returns $(0, 1, 1, 0)$, exactly f_{\oplus} . Hence GMAN represents a
 734 function unattainable by GNAN, proving that GMAN is strictly more expressive.

736 B.2 PROOF OF THEOREM 3.2

737 Let S be a set of graphs $\{G_i\}_{j=1}^m$. Let $S_1 = \{S_i\}_{i=1}^m$ be a partition of S such that $|S_i| = 1$. Let
 738 $S_2 = \{S_i\}_{i=1}^k$ such that there exists k with $|S_k| > 1$. with a subset partition $\{S_i\}_{i=1}^k$. We will prove
 739 that a GMAN trained over S_2 is strictly more expressive than a GMAN trained over S_1 .

741 To prove this, we use a ground truth function that is a set-level XOR. Let every graph G_i carry a
 742 single binary feature $x_i \in \{0, 1\}$ and let the ExtGNAN encoder return this feature unchanged, i.e.
 743 $h(G_i) = x_i$. Denote a set containing two graphs by $S = \{G_1, G_2\}$ and define the permutation-
 744 invariant target

$$745 \quad f_{\oplus}(S) = x_1 \oplus x_2.$$

746 **Singleton partition (S_1).** If each graph is placed in its own subset, GMAN aggregates *additively*:
 747 the model output is

$$749 \quad \hat{y} = \phi(x_1) + \phi(x_2),$$

750 because the final GMAN stage simply sums subset scores :contentReference
 751 [oai:icite:0]index=0:contentReference[oai:icite:1]index=1. Write $a = \phi(0)$ and $b = \phi(1)$.
 752 To realise f_{\oplus} via a threshold τ we would need

$$753 \quad a + a < \tau, \quad b + a > \tau, \quad a + b > \tau, \quad b + b < \tau.$$

755 Adding the first and last inequalities yields $a + b < \tau$, while the middle pair gives $a + b > \tau$ —a
 contradiction. Hence GMAN_{S_1} cannot represent f_{\oplus} .

756 **Paired partition (S_2).** Group the two graphs together and use a DeepSet $\Phi(S_2) = g(\sum_{i=1}^2 f(x_i))$
 757 with $f(x) = x$ and $g(s) = s(2 - s)$. Then
 758

$$759 \quad 760 \quad 761 \quad g(x_1 + x_2) = \begin{cases} 0 & (x_1, x_2) = (0, 0) \text{ or } (1, 1), \\ 1 & (x_1, x_2) = (0, 1) \text{ or } (1, 0), \end{cases}$$

762 exactly f_{\oplus} . The final GMAN sum over feature channels leaves this value unchanged, so GMAN_{S_2}
 763 realises f_{\oplus} .
 764

765 **Strict separation.** Because f_{\oplus} is representable by GMAN_{S_2} but not by GMAN_{S_1} , the former is
 766 strictly more expressive.
 767

768 B.3 FOUR-POINT CONDITION AND RECOVERABILITY

770 We now turn to structural identifiability. We prove that when input graphs are connected, acyclic,
 771 and positively weighted (i.e., trees), the pairwise distance matrix learned by GMAN encodes the
 772 full structure of the graph, up to isomorphism. This provides theoretical justification for the model's
 773 ability to reason over temporal structure without needing explicit graph supervision.

774 The following theorem shows that if the graph satisfies the four-point condition (Buneman, 1974),
 775 GMAN can reconstruct the original graph from the transformed distance matrix that is fed to GMAN
 776 as the graph input:

777 **Theorem B.1.** *Let G be a graph represented by an adjacency matrix A , and D be the transformed
 778 distance-matrix for GMAN. Then if D satisfies the four-point condition, GMAN can learn ρ such
 779 that $\rho(D) = A$.*
 780

781 *Proof.* Let $G = (V, E, w)$ be a positively weighted path graph, i.e.
 782

$$783 \quad V = \{v_1, \dots, v_n\}, \quad E = \{\{v_i, v_{i+1}\} \mid i = 1, \dots, n-1\}, \quad w(\{v_i, v_{i+1}\}) > 0.$$

785 Define the pair-wise distance matrix $D \in \mathbb{R}^{n \times n}$ by
 786

$$787 \quad D_{uv} = \sum_{e \in P_G(u, v)} w(e),$$

789 where $P_G(u, v)$ is the unique u - v path in G . Then:
 790

- 792 (a) Tree-metric property. D satisfies the four-point condition of Buneman (Buneman, 1974);
 793 hence (V, D) is a *tree metric*.
 794
- 795 (b) Uniqueness (*no information loss*). By Buneman's theorem the tree that realises D is unique
 796 up to isomorphism. For a path graph the only automorphism is the reversal $(v_1, \dots, v_n) \mapsto$
 797 (v_n, \dots, v_1) , so D determines G completely except for left-right orientation.
 798
- 799 (c) Efficient reconstruction. G can be reconstructed from D in $O(n^2)$ time:
 - 800 i. Choose an endpoint $s = \arg \max_{v \in V} \max_{u \in V} D_{vu}$.
 801 ii. Order the vertices $v_1 = s, v_2, \dots, v_n$ so that $D_{sv_1} < D_{sv_2} < \dots < D_{sv_n}$.
 802 iii. Set edge weights $w(\{v_i, v_{i+1}\}) = D_{sv_{i+1}} - D_{sv_i}$ for $i = 1, \dots, n-1$.
 803

804 \square
 805

806 C COMPUTATIONAL COMPLEXITY AND EFFICIENT IMPLEMENTATION

808 In this section, we provide a big-O analysis of the time complexity of GMAN. We also provide an
 809 efficient approach to implement GMAN.
 810

Computational Complexity The computational complexity of GMAN is as follows:

- **Scales linearly with the number of graphs m :** Each graph is processed independently or in small subsets, so total cost is $\mathcal{O}(m)$ assuming fixed per-graph cost.
- **Scales quadratically with the number of nodes per graph n :** Due to the dense aggregation over all node pairs in ExtGNAN, the per-graph cost is $\mathcal{O}(n^2)$.
- **Overall complexity** for a set of graphs is:

$$\mathcal{O}(m \cdot K \cdot n^2 \cdot d_\psi)$$

where K is the number of feature groups and d_ψ is the cost of evaluating the multivariate neural networks.

Adapting the complexity We note that while the per-graph complexity of ExtGNAN is $O(n^2)$ in the most general case, the distance function Δ can be masked to enforce any desired sparsity pattern between nodes, and therefore adapt to any complexity limitation, including a linear one. For example, one can define the Δ to only be the distance between adjacent nodes in a trajectory, and then it will be both linear in memory and time.

Efficient Implementations In the main paper, we present GMAN with the objective of maximal clarity, e.g., by presenting vector entry-wise operations. Nonetheless, the operations of GMAN can be done in an optimized fashion for GPU, through tensor operations.

D DATASET DETAILS

D.1 PHYSIONET P12

We provide the full list of the 36 physiological signals and 3 static patient features used in our experiments.

1. Alkaline phosphatase (ALP): A liver- and bone-derived enzyme; elevations suggest cholestasis, bone disease, or hepatic injury.
2. Alanine transaminase (ALT): Hepatocellular enzyme; increased values mark acute or chronic liver cell damage.
3. Aspartate transaminase (AST): Enzyme in liver, heart, and muscle; rises indicate hepatocellular or muscular injury.
4. Albumin: Major plasma protein maintaining oncotic pressure and transport; low levels reflect inflammation, malnutrition, or liver dysfunction.
5. Blood urea nitrogen (BUN): End-product of protein catabolism cleared by the kidneys; elevation signals renal impairment or high catabolic state.
6. Bilirubin: Hemoglobin breakdown product processed by the liver; accumulation indicates hepatobiliary disease or hemolysis.
7. Cholesterol: Circulating lipid essential for membranes and hormones; dysregulation is linked to cardiovascular risk.
8. Creatinine: Waste from muscle metabolism filtered by the kidneys; higher levels imply reduced glomerular filtration.

864 9. Invasive diastolic arterial blood pressure (DiasABP): Pressure during ventricular relaxation;
865 low readings may reflect vasodilation or hypovolemia.
866

867 10. Fraction of inspired oxygen (FiO₂): Proportion of oxygen delivered; values above ambient
868 air denote supplemental therapy.
869

870 11. Glasgow Coma Score (GCS): Composite neurologic score for eye, verbal, and motor
871 responses; scores ≤ 8 indicate severe impairment.
872

873 12. Glucose: Principal blood sugar; hypo- or hyper-glycemia can cause neurologic compromise
874 and metabolic instability.
875

876 13. Serum bicarbonate (HCO₃): Key extracellular buffer; low levels signal metabolic acidosis,
877 high levels metabolic alkalosis or compensation.
878

879 14. Hematocrit (HCT): Percentage of blood volume occupied by red cells; reduced values
880 denote anemia, elevated values hemoconcentration.
881

882 15. Heart rate (HR): Beats per minute reflecting cardiac demand; tachycardia indicates stress or
883 shock, bradycardia conduction disorders.
884

885 16. Serum potassium (K): Crucial intracellular cation; deviations predispose to dangerous
886 arrhythmias.
887

888 17. Lactate: By-product of anaerobic metabolism; elevation marks tissue hypoxia and shock
889 severity.
890

891 18. Invasive mean arterial blood pressure (MAP): Time-weighted average arterial pressure; low
892 values threaten organ perfusion.
893

894 19. Mechanical ventilation flag (MechVent): Binary indicator of ventilatory support; presence
895 denotes respiratory failure or peri-operative care.
896

897 20. Serum magnesium (Mg): Cofactor for numerous enzymatic reactions; abnormalities
898 contribute to arrhythmias and neuromuscular instability.
899

900 21. Non-invasive diastolic arterial blood pressure (NIDiasABP): Cuff-derived diastolic pressure;
901 trends mirror vascular tone without an arterial line.
902

903 22. Non-invasive mean arterial blood pressure (NIMAP): Cuff-based mean pressure; used when
904 invasive monitoring is unavailable.
905

906 23. Non-invasive systolic arterial blood pressure (NISysABP): Cuff-derived systolic pressure;
907 elevations suggest hypertension or pain response.
908

909 24. Serum sodium (Na): Principal extracellular cation governing osmolality; dysnatremias
910 cause neurologic symptoms and fluid shifts.
911

912 25. Partial pressure of arterial carbon dioxide (PaCO₂): Indicator of ventilatory status;
913 hypercapnia implies hypoventilation, hypocapnia hyperventilation.
914

915 26. Partial pressure of arterial oxygen (PaO₂): Measure of oxygenation efficiency; low values
916 denote hypoxemia.
917

918 27. Arterial pH: Measure of hydrogen-ion concentration; deviations from normal reflect
 919 systemic acid-base disorders.
 920

921 28. Platelet count (Platelets): Thrombocyte concentration essential for hemostasis; low counts
 922 increase bleeding risk, high counts thrombosis risk.
 923

924 29. Respiration rate (RespRate): Breaths per minute; tachypnea signals metabolic acidosis or
 925 hypoxia, bradypnea central depression.
 926

927 30. Hemoglobin oxygen saturation (SaO₂): Percentage of hemoglobin bound to oxygen; values
 928 below normal indicate significant hypoxemia.
 929

930 31. Invasive systolic arterial blood pressure (SysABP): Peak pressure during ventricular
 931 ejection; extremes compromise end-organ perfusion.
 932

933 32. Body temperature: Core temperature; fever suggests infection, hypothermia exposure or
 934 metabolic dysfunction.
 935

936 33. Troponin I: Cardiac-specific regulatory protein; elevation confirms myocardial injury.
 937

938 34. Troponin T: Isoform of cardiac troponin complex; rise parallels Troponin I in detecting
 939 myocardial necrosis.
 940

941 35. Urine: Hourly urine volume as a gauge of renal perfusion; oliguria signals kidney
 942 hypoperfusion or failure.
 943

944 36. White blood cell count (WBC): Reflects immune activity; leukocytosis suggests infection or
 945 stress, leukopenia marrow suppression or severe sepsis.
 946

947

948 **Static patient features:** Age; Gender; *ICUType* – categorical code for the admitting intensive care
 949 unit (1 = Coronary Care, 2 = Cardiac Surgery Recovery, 3 = Medical ICU, 4 = Surgical ICU),
 950 capturing differences in case mix and treatment environment.
 951

952 D.2 CROHN'S DISEASE PREDICTION

953 We detail the full list of the 17 biomarkers extracted from the Danish health registries.
 954

955 1. C-reactive protein (CRP): A protein produced by the liver in response to inflammation.
 956 Elevated CRP indicates active inflammation, often associated with inflammatory diseases
 957 like CD.
 958

959 2. Faecal Calprotectin (F-Cal): A protein released from neutrophils into the intestinal lumen,
 960 detectable in stool samples. Elevated levels indicate gastrointestinal inflammation and are
 961 commonly used to detect and monitor inflammatory bowel disease.
 962

963 3. Leukocytes (White Blood Cells): Cells that are central to the body's immune response.
 964 Elevated leukocyte counts typically suggest infection or inflammation, including flare-ups
 965 in CD.
 966

967 4. Neutrophils: A type of leukocyte involved, among other things, in fighting bacterial
 968 infections. High neutrophil counts often indicate acute inflammation or infection, including
 969 intestinal inflammation in CD.
 970

972 5. Lymphocytes: A group of white blood cells that form the core of the adaptive immune
 973 system, including T cells, B cells, and natural killer (NK) cells. They are responsible for
 974 antigen-specific immune responses. Abnormal levels can signal immune dysregulation,
 975 often implicated in autoimmune and chronic inflammatory diseases such as CD.
 976

977 6. Monocytes: A type of white blood cell that circulates in the blood and differentiates
 978 into macrophages or dendritic cells upon entering tissues. These cells are essential for
 979 phagocytosis, antigen presentation, and regulation of inflammation. Elevated levels may
 980 reflect immune activation or tissue damage.
 981

982 7. Eosinophils: Immune cells involved primarily in allergic reactions and parasitic infections.
 983 Elevated eosinophil counts might reflect allergic responses or gastrointestinal inflammation.
 984

985 8. Basophils: The least common type of leukocyte, involved in allergic and inflammatory
 986 responses. Their elevation is uncommon but may accompany certain inflammatory or
 987 allergic conditions.
 988

989 9. Platelets: Cell fragments critical for blood clotting and also involved in inflammatory
 990 responses. High platelet counts (thrombocytosis) are commonly seen during active
 991 inflammation in conditions like CD.
 992

993 10. Hemoglobin (Hb): The protein in red blood cells responsible for oxygen transport. Low
 994 hemoglobin (anemia) is frequently observed in chronic inflammatory conditions such as CD
 995 due to blood loss or nutrient deficiencies.
 996

997 11. Iron: An essential mineral for red blood cell production. Low iron levels often indicate
 998 chronic blood loss or malabsorption, both common in CD due to intestinal inflammation.
 999

1000 12. Folate (Vitamin B9): A vitamin necessary for red blood cell production and DNA synthesis.
 1001 Deficiency may result from impaired absorption in inflamed intestinal tissue.
 1002

1003 13. Vitamin B12 (Cobalamin): Required for red blood cell production and neurological function.
 1004 Deficiencies are common in CD, especially when the ileum is affected.
 1005

1006 14. Vitamin D2+D3 (Ergocalciferol + Cholecalciferol): Vitamins essential for bone health and
 1007 immune regulation. Low levels are often seen in CD due to malabsorption and systemic
 1008 inflammation.
 1009

1010 15. ALAT (Alanine Aminotransferase): An enzyme indicating liver function. Elevated levels
 1011 may reflect liver inflammation, medication effects, or co-occurring autoimmune liver disease.
 1012

1013 16. Albumin: A protein produced by the liver that helps maintain blood volume and transport
 1014 nutrients. Low albumin can reflect chronic inflammation, malnutrition, or protein loss in CD.
 1015

1016 17. Bilirubin: A compound produced from red blood cell breakdown. It is filtered by the liver
 1017 and excreted into the intestine via bile. Elevated levels may indicate liver dysfunction, bile
 1018 duct obstruction, or hemolytic anemia.
 1019

1020 **D.2.1 CLINICAL CONTEXT AND RELATED WORK FOR PREDICTING CD ONSET**

1021 Research on predicting the onset of CD has explored a range of approaches, including the use
 1022 of routinely measured blood-based biomarkers and more complex biological data derived from
 1023 multi-omics technologies.

1026 Several studies have assessed the predictive potential of standard clinical blood tests. For example,
 1027 (Vestergaard et al., 2023) analyzed six routine biomarkers from 1,186 Danish patients eventually
 1028 diagnosed with CD, achieving moderate predictive performance (AUROC of 0.74) approximately
 1029 six months before clinical diagnosis. Larger-scale analyses, such as those using UK Biobank
 1030 data, combined multiple standard biomarkers and basic demographic information, reporting similar
 1031 predictive performances (AUROCs typically between 0.70–0.75). These analyses generally utilized
 1032 methods like logistic regression, random forests, or gradient-boosted trees, favored for structured
 1033 clinical datasets.

1034 Other studies have integrated advanced biochemical data, known as multi-omics, including large-scale
 1035 protein measurements (proteomics), metabolites (metabolomics), or genomic markers. (Garg et al.,
 1036 2024) for instance, combined 67 blood biomarkers with approximately 2,900 plasma proteins from
 1037 the UK Biobank, achieving an AUROC of 0.786. (Woerner et al., 2025) combined genetic risk scores
 1038 with extensive proteomic data, achieving an AUROC of 0.76 for CD prediction up to five years prior
 1039 to diagnosis. Similar multi-omics approaches employing microbiome profiling, immune signaling
 1040 molecules (cytokines), or lipid molecules typically achieve AUROCs between 0.75 and 0.80 but
 1041 often involve significant cost, specialized laboratory analyses, and reduced consistency across diverse
 1042 patient cohorts.

1043 Overall, routine blood tests provide meaningful predictive signals for CD onset, while integrating com-
 1044 plex biochemical measurements can improve predictive accuracy, albeit at greater cost, complexity,
 1045 and variability across clinical populations.

1046

1047 E BIOMARKER SUBSET GROUPINGS

1049 E.1 P12

1051 In the PhysioNet P12 task, we grouped the 36 physiological signals into one multivariate subset and
 1052 29 singleton subsets. Domain knowledge showed that only the respiratory and gas-exchange variables
 1053 shared sufficiently strong, coherent dynamics to benefit from joint modeling. All other signals were
 1054 physiologically diverse, so they were left as singletons to retain their unique predictive information.

1055 1. **Arterial blood gas profile**

1056 *[pH, PaCO₂, PaO₂, SaO₂, HCO₃, FiO₂]*

1057 This group captures systemic acid-base status (pH, HCO₃), carbon dioxide clearance
 1058 (PaCO₂), oxygenation (PaO₂, SaO₂), and inspired oxygen fraction (FiO₂). Together they
 1059 form the canonical arterial blood gas panel, enabling the model to detect respiratory de-
 1060 rangements such as hypoxemia, hypercapnia, or metabolic compensation.

1061 2. **Complete blood count**

1062 *[WBC, HCT, Platelets]*

1063 This cluster summarizes hematologic composition by measuring leukocyte-mediated im-
 1064 mune response (WBC), oxygen-carrying capacity (HCT), and clotting potential (Platelets).
 1065 Joint modeling supports recognition of systemic inflammation, anemia, and coagulopathy.

1066 3. **Comprehensive metabolic panel**

1067 *[Glucose, Na, K, Mg, BUN, Creatinine]*

1068 These biomarkers represent key substrates and electrolytes (Glucose, Na, K, Mg) and renal
 1069 waste products (BUN, Creatinine). Grouping them provides a unified view of metabolic
 1070 balance, electrolyte homeostasis, and kidney function.

1071 4. **Liver function tests**

1072 *[ALT, AST, ALP, Albumin, Bilirubin]*

1073 These biomarkers assess hepatocellular injury (ALT, AST), cholestasis (ALP, Bilirubin),
 1074 and hepatic synthetic function (Albumin). Their combined interpretation reflects multiple
 1075 dimensions of liver health.

1076 5. **Lipid and cardiac markers**

1077 *[Cholesterol, TroponinI, TroponinT, HR]*

1078 This group integrates lipid metabolism (Cholesterol), cardiac injury markers (Troponin I,
 1079 Troponin T), and heart rate (HR). Together, they provide insight into cardiovascular stress,
 myocardial injury, and metabolic risk.

1080 6. **Blood pressure profiles**
 1081 *[SysABP, DiasABP, MAP, NISysABP, NIDiasABP, NIMAP]*
 1082 These variables capture invasive and non-invasive arterial blood pressures, reflecting sys-
 1083 temic hemodynamics. Grouping them enables the model to learn coherent pressure dynamics
 1084 rather than treating each measurement in isolation.

1085 7. **Ventilation mechanics**
 1086 *[RespRate, MechVent]*
 1087 This group reflects mechanical and physiological components of ventilation. Their joint
 1088 dynamics provide context for interpreting respiratory compensation and ventilatory support.

1089 8. **Tissue perfusion**
 1090 *[Lactate, Urine]*
 1091 Elevated lactate indicates anaerobic metabolism, while urine output tracks renal perfusion.
 1092 Together they provide complementary signals of global tissue perfusion and shock severity.

1093 9. **Global status indicators**
 1094 *[GCS, Temp]*
 1095 These variables capture overall neurologic responsiveness (GCS) and systemic temperature
 1096 regulation (Temp), providing global context on patient stability and severity of illness.

1097 E.2 CD

1100 In the Crohn’s Disease prediction task, we grouped the 17 selected biomarkers into 7 subsets based
 1101 on shared physiological function, clinical relevance, and correlated patterns observed in exploratory
 1102 analyses. This configuration produced the most robust and interpretable results, balancing domain
 1103 knowledge with empirical performance. The grouping is as follows:

1104 1. **White blood cell subtypes**
 1105 *[Leukocytes, Neutrophils, Lymphocytes, Monocytes, Eosinophils, Basophils]*
 1106 These biomarkers all represent components of the immune system’s cellular response. Grouping
 1107 them enables the model to learn shared immune activation patterns, which are known to be
 1108 dysregulated in inflammatory bowel diseases like CD. Combining them in a multivariate subset
 1109 captures both their relative proportions and total counts, which are clinically relevant for
 1110 distinguishing inflammation subtypes.

1111 2. **Inflammation markers**
 1112 *[CRP, Faecal Calprotectin]*
 1113 These are key indicators of systemic and intestinal inflammation, respectively. CRP reflects
 1114 acute-phase liver response, while F-Cal is specific to intestinal neutrophilic activity. Though
 1115 mechanistically distinct, both are strongly correlated with inflammatory disease activity and
 1116 complement each other in modeling CD-specific inflammation signatures.

1117 3. **Platelets**
 1118 *[Platelets]*
 1119 Thrombocytosis (elevated platelet count) is a well-established marker of chronic inflammation.
 1120 As platelet behavior is relatively independent from other hematological and nutritional markers,
 1121 we model it as its own trajectory.

1122 4. **Hemoglobin**
 1123 *[Hemoglobin]*
 1124 Hemoglobin concentration is a direct measure of anemia, which is prevalent in CD patients due
 1125 to chronic blood loss and inflammation-induced iron sequestration. Its temporal dynamics often
 1126 diverge from those of other blood components, warranting a separate representation.

1127 5. **Iron status**
 1128 *[Iron]*
 1129 Iron metabolism is tightly linked to both hemoglobin levels and systemic inflammation but
 1130 shows distinct dynamics. Modeling it separately allows the model to learn delayed or decoupled
 1131 effects (e.g., iron deficiency preceding hemoglobin drop).

1132 6. **Vitamin and folate status**
 1133 *[Folate, Vitamin B12, Vitamin D2+D3]*
 1134 These nutrients are absorbed in different regions of the gastrointestinal tract (e.g., B12 in the

1134 ileum, folate in the jejunum), and their deficiency profiles can be informative of CD location and
 1135 severity. Grouping them allows the model to detect joint patterns of malabsorption and systemic
 1136 nutrient depletion.

1137 **7. Liver function markers**

1138 *[ALAT, Bilirubin, Albumin]*

1139 These biomarkers reflect hepatic function and protein synthesis. Abnormal liver enzymes and
 1140 hypoalbuminemia are frequently observed in CD due to medication effects, chronic inflammation,
 1141 or comorbid autoimmune liver disease. Combining them supports learning of systemic
 1142 inflammatory effects beyond the gut.

1143 This grouping reflects known biological relationships, enhances the interpretability of the model’s
 1144 subset-level attributions, and improves performance compared to unstructured or purely univariate
 1145 representations. It enables GMAN to exploit interactions among related features while maintaining a
 1146 modular structure that aligns with clinical reasoning.

1147 **F EXPERIMENTAL SETUP, HYPERPARAMETER CHOICES AND GROUPING
 1148 CONFIGURATIONS**

1149 This section outlines key implementation choices and model settings used in our experiments, includ-
 1150 ing the manually tuned biomarker grouping configurations that served as an important hyperparameter
 1151 for performance and interpretability.

1152 **F.1 HYPER-PARAMETERS**

1153 We trained all models with 100 epochs using the Adam optimizer with weight decay 1e-5. We used a
 1154 ReduceLROnPlateau scheduler with a max learning rate in the 1e-2, 1e-4 range, min learning rate in
 1155 the 1e-7, 1e-8 range, factor in the 0.2-0.9 range, and patience=100

1156 We trained all models with batch size of range {16, 32}, dropout rate in {0.1, 0.2}, number of layers
 1157 in the {3, 4, 5} range, hidden channels in the {32, 64} range.

1158 Random seeds were fixed for reproducibility, and results are reported across three independent runs.
 1159 All models were trained on a single NVIDIA Tesla V100-PCIE-16GB GPU.

1160 **F.2 GROUPING CONFIGURATIONS FOR CLINICAL TASKS**

1161 In both clinical tasks, the configuration of input feature subsets (i.e., how we grouped input biomarkers
 1162 into multivariate trajectories) was treated as a manually tuned hyperparameter. These groupings
 1163 determine how GMAN combines individual graph representations prior to final prediction, and they
 1164 affect both the expressivity and interpretability of the model.

1165 **In-Hospital Mortality (P12).** We compared GMAN’s performance under the following grouping
 1166 strategies:

- 1167 • **No grouping:** Each biomarker is in its own size one subset.
- 1168 • **Respiratory** - one group of the biomarkers: *[FiO₂, PaO₂, PaCO₂, SaO₂, RespRate, pH, MechVent]* and the rest are singletons.
- 1169 • **Metabolic Electrolytes** - one group of the biomarkers: *[Na, K, Mg, HCO₃, Lactate, Glucose]* and the rest are singletons.
- 1170 • **Liver Panel** - one group of the biomarkers: *[ALT, AST, ALP, Albumin, Bilirubin, Cholesterol]* and the rest are singletons.
- 1171 • **Pathway-Based Grouping:** Biomarkers are organised based on their molecular or mech-
 1172 anistic roles, grouping them by their function in metabolism, homeostasis, or cellular
 1173 composition: Energy metabolism [Glucose, Cholesterol, Lactate]; nitrogen waste clearance
 1174 [BUN, Creatinine, Urine]; protein synthesis and enzymes [Albumin, ALT, AST]; liver
 1175 function and cholestasis [ALP, Bilirubin]; acid-base balance [pH, HCO₃]; gas transport
 1176 [PaO₂, PaCO₂, SaO₂, FiO₂]; mineral homeostasis [Na, K, Mg]; hematologic composition

1188 [HCT, Platelets, WBC]; cardiovascular dynamics [SysABP, DiasABP, MAP, NISysABP,
 1189 NIDiasABP, NIMAP, HR]; respiratory mechanics [RespRate, MechVent]; and cardiac injury
 1190 [TroponinI, TroponinT]. Global status indicators are grouped separately as [Temp, GCS].
 1191

- 1192 • **Organ-System Grouping** Biomarkers are organized around organ systems and clinical
 1193 monitoring domains (respiratory support, cardiovascular dynamics etc: oxygenation support
 1194 [SaO₂, PaO₂, FiO₂, MechVent]; ventilation and acid–base balance [PaCO₂, RespRate,
 1195 pH, HCO₃]; cardiovascular dynamics [HR, MAP, SysABP, DiasABP, NIMAP, NISysABP,
 1196 NIDiasABP]; perfusion and renal function [Urine, Lactate, Creatinine, BUN]; hepatobil-
 1197 iary function [Bilirubin, ALT, AST, ALP, Albumin]; inflammation and coagulation [WBC,
 1198 Platelets]; electrolyte and oxygen-carrying capacity [Na, K, Mg, HCT]; metabolic reserve
 1199 [Glucose, Cholesterol]; myocardial injury [TroponinI, TroponinT]; and global status indica-
 1200 tors [Temp, GCS].

1201 **Crohn’s Disease Onset.** We evaluated several grouping configurations:

- 1202 • **No grouping:** Each biomarker is in its own size one subset.
- 1203 • **Biologically driven grouping** (see Appendix E.2): Biomarkers are grouped into 7 clinically
 1204 coherent subsets (e.g., inflammation markers, immune cell subtypes, liver function).
- 1205 • **Diagnostic Panel Grouping:** A clinically motivated grouping that mirrors standard blood
 1206 test panels used in routine diagnostics.
- 1207 • **Data-driven grouping:** Groupings are derived from clustering biomarkers based on the
 1208 complementary signal in their attribution curves (see Section 4.1.1).
- 1209 • **Merged coarse groupings:** Broad categories such as inflammation, haematology, and
 1210 micronutrients.
- 1211 • **Minimal Pairwise Interaction :** Emphasizes minimal yet informative combinations that
 1212 capture key axes of immune, inflammatory, and metabolic variational proximity.

1213 We report the AUPRC of different biomarker groupings for predicting the onset of CD in Table 4

1217 G CALIBRATION AND ROBUSTNESS ANALYSES

1219 Here we report additional evidence on the reliability (calibration) and robustness of GMAN on the
 1220 Crohn’s Disease (CD) task.

1223 G.1 CALIBRATION

1224 We compute Expected Calibration Error (ECE) for GMAN and all baselines on the CD task. ECE
 1225 is reported as mean \pm std over three data splits. Lower values indicate better agreement between
 1226 predicted probabilities and observed outcome frequencies. As shown in Table 5, GMAN attains the
 1227 lowest ECE among the compared methods.

1228 To complement ECE, we provide a Q–Q calibration plot for GMAN in Fig. 5. The plot compares
 1229 binned predicted probabilities against empirical outcome frequencies. Close alignment with the
 1230 diagonal indicates good calibration. The near-diagonal trend is consistent with GMAN’s low ECE.

1232 G.2 NOISE ROBUSTNESS UNDER CONTROLLED DISTRIBUTION SHIFT

1234 We evaluate robustness under controlled distribution shifts by training all models on clean data and
 1235 injecting noise only at test time. For each noise level, we re-compute AUROC and AUPRC and report
 1236 the relative performance drop δ (percentage change) from the clean-test baseline.

1238 **Additive value noise.** For each biomarker trajectory graph, we perturb the primary biomarker
 1239 feature (feature index 0 corresponding to the observed biomarker value) at every node by adding
 1240 zero-mean Gaussian noise with fixed standard deviation σ_{value} , independent of the feature’s magnitude.
 1241 If v is the original biomarker value, we sample $\epsilon \sim \mathcal{N}(0, 1)$ and use $v' = v + \epsilon \cdot \sigma_{\text{value}}$. This models
 1242 absolute measurement variability with a uniform noise scale. The results are shown in Table 6

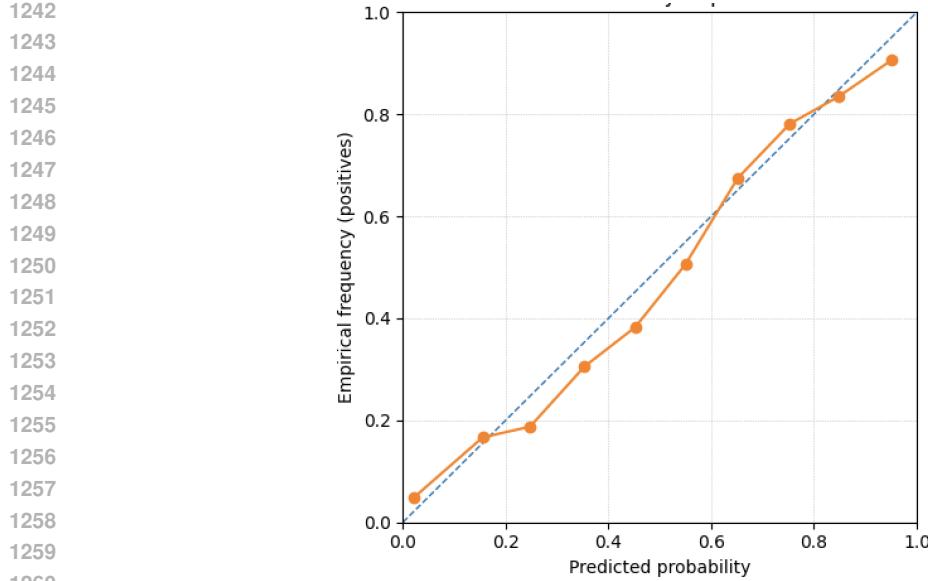


Figure 5: Q–Q calibration plot (reliability diagram) for GMAN on the CD task.

Multiplicative value noise. Using the same feature, we instead add zero-mean Gaussian noise whose scale is proportional to the absolute feature value. If v is the original value, we set the noise scale to $|v| \cdot \sigma_{\text{value}}$ and sample $v' = v + \epsilon \cdot |v| \cdot \sigma_{\text{value}} \approx v \cdot (1 + \epsilon \cdot \sigma_{\text{value}})$. This models relative measurement/units variability, where larger measurements incur larger absolute perturbations. The results are shown in Table 7

Temporal noise. We inject zero-mean Gaussian noise into the temporal distance matrix at test time, with standard deviation σ_{time} (in days). This perturbs pairwise time gaps between visits while keeping biomarker values and graph structure unchanged. The results are shown in Table 8

H LLM USAGE

We relied on Large Language Models solely for grammar and spelling checks, without using them to generate or modify the scientific content.

1296

1297

1298

1299

1300

1301

1302

1303

1304

1305

1306

1307

1308

1309

1310

1311

1312

1313

1314

1315

1316

1317

1318

1319

1320

1321

1322

1323

1324

1325

1326

1327

1328

1329

1330

1331

1332

1333

1334

1335

1336

1337

1338

1339

1340

1341

1342

1343

1344

1345

1346

1347

1348

1349

Table 4: Performance using different biomarker subsets for CD prediction.

Grouping Strategy	Biomarker Groups	AUPRC
Flat grouping	Each biomarker is modeled independently: CRP, F-Cal, Leukocytes, Neutrophils, Lymphocytes, Monocytes, Eosinophils, Basophils, Platelets, Hemoglobin, Iron, Folate, Vitamin B12, Vitamin D2+D3, ALAT, Albumin, Bilirubin.	79.66 ± 0.96
Biologically driven grouping	<ul style="list-style-type: none"> White blood cell subtypes: Leukocytes, Neutrophils, Lymphocytes, Monocytes, Eosinophils, Basophils Inflammatory markers: CRP, F-Cal Platelets Haemoglobin Iron Vitamin and folate status: Folate, Vitamin B12, Vitamin D2 + D3 Liver function markers: ALAT, Albumin, Bilirubin 	83.93 ± 0.27
Diagnostic Panel Grouping	<ul style="list-style-type: none"> Inflammatory markers: CRP, F-Cal WBC count (main): Leukocytes, Neutrophils, Lymphocytes WBC rare subtypes: Monocytes, Eosinophils, Basophils Platelets + Hemoglobin: Platelets, Hemoglobin Nutrient panel: Iron, Folate, Vitamin B12, Vitamin D2+D3 Liver function test panel: ALAT, Albumin, Bilirubin 	79.06 ± 3.4
Data-driven grouping	<ul style="list-style-type: none"> Acute Inflammation Markers: CRP, F-Cal, Platelets Immune–Iron Axis: Lymphocytes, Iron Neutrophils Monocytes Eosinophils Basophils Hemoglobin Folate Vitamin B12 Vitamin D ALAT Bilirubin Albumin Total Leukocytes 	79.04 ± 7.8
Merged coarse groupings	<ul style="list-style-type: none"> F-Cal (local inflammation) Systemic immune/inflammation: CRP, Leukocytes, Neutrophils, Lymphocytes, Monocytes Allergy-linked eosinophils/basophils Hematological status: Platelets, Hemoglobin Nutritional status: Iron, Folate, Vitamin B12, Vitamin D2+D3 Hepatic status: ALAT, Albumin, Bilirubin 	75.21 ± 2.3
Minimal Pairwise Interaction	<ul style="list-style-type: none"> CRP + F-Cal Leukocytes + Neutrophils Lymphocytes + Monocytes Eosinophils + Basophils Platelets Hemoglobin Iron + Folate Vitamin B12 + Vitamin D2+D3 ALAT + Albumin Bilirubin 	81.50 ± 1.3

1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360

Model	ECE (mean \pm std)
GMAN	0.028 \pm 0.004
Raindrop	0.040 \pm 0.070
DGM2	0.035 \pm 0.001
mTAND	2.16 \pm 0.54
Transformer	3.32 \pm 1.08
Trans-mean	3.30 \pm 1.03
SEFT	2.40 \pm 0.71

Table 5: Expected Calibration Error (ECE) on the CD task, reported as mean \pm std over three splits.

1361
1362
1363
1364

σ_v^+	GMAN		Raindrop		mTAND	
	Δ AUROC	Δ AUPRC	Δ AUROC	Δ AUPRC	Δ AUROC	Δ AUPRC
0.0	0.00 \pm 0.00%	0.00 \pm 0.00%	0.00 \pm 0.00%	0.00 \pm 0.00%	0.00 \pm 0.00%	0.00 \pm 0.00%
0.1	-0.07 \pm 0.11%	-0.07 \pm 0.08%	-26.20 \pm 1.97%	-28.72 \pm 2.54%	-0.82 \pm 0.81%	-0.71 \pm 0.64%
0.2	-0.10 \pm 0.15%	-0.02 \pm 0.52%	-32.82 \pm 0.96%	-36.70 \pm 1.49%	-2.98 \pm 1.23%	-2.68 \pm 0.98%
0.3	-0.18 \pm 0.34%	-0.35 \pm 0.47%	-34.74 \pm 1.36%	-38.69 \pm 1.78%	-5.58 \pm 1.71%	-5.16 \pm 1.06%
0.5	-0.51 \pm 0.35%	-0.37 \pm 1.11%	-36.35 \pm 1.54%	-40.60 \pm 1.79%	-10.97 \pm 1.32%	-10.84 \pm 0.73%
0.8	-1.32 \pm 0.97%	-1.20 \pm 1.85%	-37.25 \pm 1.59%	-41.41 \pm 1.68%	-17.13 \pm 1.76%	-17.97 \pm 2.58%
1.5	-1.87 \pm 2.48%	-1.31 \pm 3.96%	-37.73 \pm 1.34%	-42.27 \pm 1.33%	-26.67 \pm 1.44%	-29.26 \pm 2.05%
3.0	-5.46 \pm 5.10%	-4.81 \pm 7.06%	-37.70 \pm 1.52%	-42.39 \pm 1.49%	-34.00 \pm 1.34%	-36.41 \pm 2.36%
5.5	-7.35 \pm 5.45%	-7.48 \pm 7.05%	-38.07 \pm 1.45%	-42.34 \pm 1.73%	-38.56 \pm 1.21%	-41.31 \pm 1.58%
7.0	-9.39 \pm 5.22%	-9.82 \pm 7.54%	-37.59 \pm 1.60%	-42.37 \pm 1.72%	-38.44 \pm 1.67%	-41.40 \pm 1.63%

Table 6: Δ AUROC and Δ AUPRC under additive value noise (σ_v^+).

1375
1376
1377
1378

σ_v^*	GMAN		Raindrop		mTAND	
	Δ AUROC	Δ AUPRC	Δ AUROC	Δ AUPRC	Δ AUROC	Δ AUPRC
0.0	0.00 \pm 0.00%	0.00 \pm 0.00%				
0.1	-0.04 \pm 0.02%	0.00 \pm 0.04%	0.04 \pm 1.11%	0.02 \pm 0.65%	-0.35 \pm 0.99%	-0.35 \pm 0.63%
0.2	0.09 \pm 0.24%	0.26 \pm 0.41%	-0.21 \pm 1.15%	-0.27 \pm 0.64%	-1.36 \pm 0.75%	-1.36 \pm 0.29%
0.3	-0.15 \pm 0.48%	-0.02 \pm 0.54%	-0.40 \pm 1.20%	-0.40 \pm 0.60%	-2.98 \pm 1.21%	-3.18 \pm 1.14%
0.5	-0.57 \pm 0.56%	-0.39 \pm 0.89%	-1.28 \pm 1.31%	-1.26 \pm 0.66%	-7.61 \pm 1.40%	-7.79 \pm 1.08%
0.8	-0.78 \pm 0.66%	-0.68 \pm 1.70%	-2.45 \pm 0.99%	-2.44 \pm 0.47%	-13.79 \pm 1.30%	-13.25 \pm 1.11%
1.5	-1.82 \pm 1.98%	-1.71 \pm 3.41%	-3.19 \pm 1.57%	-3.16 \pm 1.18%	-24.22 \pm 1.96%	-22.58 \pm 1.45%
3.0	-4.63 \pm 3.22%	-4.89 \pm 5.29%	-4.56 \pm 1.18%	-4.60 \pm 0.51%	-31.53 \pm 1.61%	-29.08 \pm 1.56%
5.5	-4.91 \pm 4.85%	-5.31 \pm 5.82%	-6.41 \pm 1.51%	-6.20 \pm 0.84%	-35.16 \pm 1.09%	-32.07 \pm 1.10%
7.0	-7.45 \pm 5.28%	-8.15 \pm 6.59%	-7.23 \pm 1.80%	-7.37 \pm 0.65%	-36.74 \pm 1.82%	-33.55 \pm 2.34%

Table 7: Δ AUROC and Δ AUPRC under multiplicative value noise (σ_v^*).

1390
1391
1392
1393

σ_t (days)	GMAN		Raindrop		mTAND	
	Δ AUROC	Δ AUPRC	Δ AUROC	Δ AUPRC	Δ AUROC	Δ AUPRC
0.0	0.00 \pm 0.00%					
10.0	0.02 \pm 0.17%	0.03 \pm 0.16%	-0.54 \pm 1.17%	-0.72 \pm 0.44%	-2.17 \pm 0.77%	-2.16 \pm 0.65%
30.0	-0.04 \pm 0.05%	0.03 \pm 0.04%	-0.82 \pm 1.05%	-1.44 \pm 0.53%	-1.94 \pm 1.06%	-1.89 \pm 0.80%
90.0	-0.22 \pm 0.13%	-0.39 \pm 0.35%	-1.58 \pm 0.87%	-2.20 \pm 0.15%	-1.99 \pm 0.71%	-2.02 \pm 0.56%
150.0	-0.22 \pm 0.09%	-0.32 \pm 0.08%	-1.89 \pm 0.86%	-2.75 \pm 0.63%	-2.24 \pm 1.13%	-2.28 \pm 0.71%
300.0	-0.45 \pm 0.42%	-0.87 \pm 0.78%	-2.41 \pm 1.21%	-3.26 \pm 0.47%	-1.85 \pm 0.99%	-1.81 \pm 0.74%
500.0	-0.74 \pm 0.04%	-1.12 \pm 0.29%	-2.63 \pm 1.24%	-3.73 \pm 0.90%	-2.14 \pm 0.97%	-2.11 \pm 0.70%

Table 8: Relative change (Δ) in AUROC and AUPRC under temporal noise.

1402
1403