“Localizing the invisible: Graph Neural Networks for Biomedical Signals.”

Syed Muhammad Hamza Zaidi · Myra Spiliopoulou

AIRA SEMINAR

Knowledge Management & Discovery Lab (KMD)

Otto-von-Guericke-University Magdeburg, Germany

22/01/2026

	AGENDA
1.	Sensor-level anomalies in HealthCare: The problem
2.	A Common Structure Behind These Problems
3.	Solution: Temporal GNNs
4.	Case study 1: Localization of Breathing abnormalities
5.	XAI for GNNs (Explainability in TGNNs)
6.	Future Direction

	Appendix
7.	Case study 2: Artefact Detection in Gerbil Electrophysiology

1. Sensor-Level Anomalies in Healthcare: Problem

- Sensor loosening, broken contacts, movement artifacts.
- In clinical and safety-critical environments, unnoticed sensor faults or anomalies can lead to misdiagnosis or false alarms.
- True physiological issue or a sensor fault?
- Healthcare professionals demand interpretable AI systems, which sensor is showing abnormality and why?
- Black-Box alerts are not trusted and evidence are required.

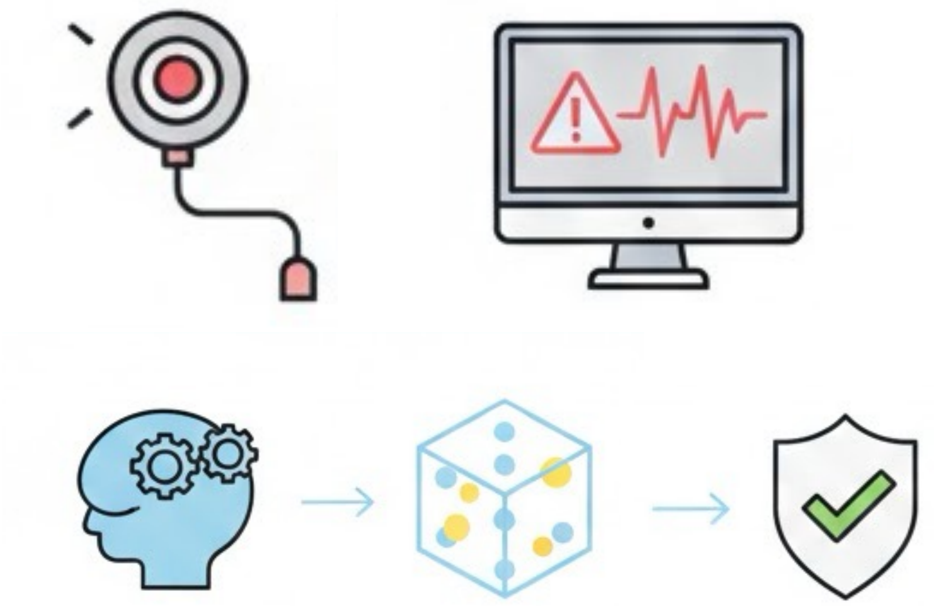

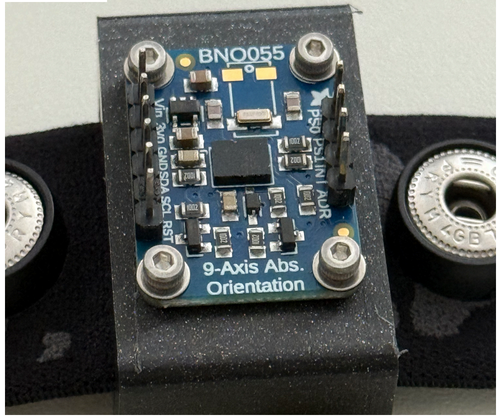

2. A Common Structure Behind These Problems

- Healthcare systems = sensors + time
- Sensors interact with each other
- Data forms a **graph over time**
- Nodes = sensors / channels / regions
- Edges = proximity, correlation

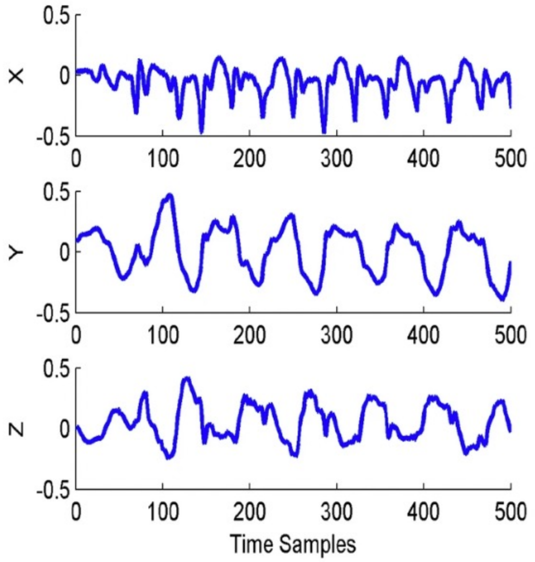

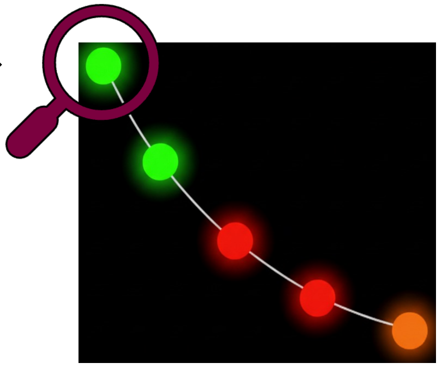

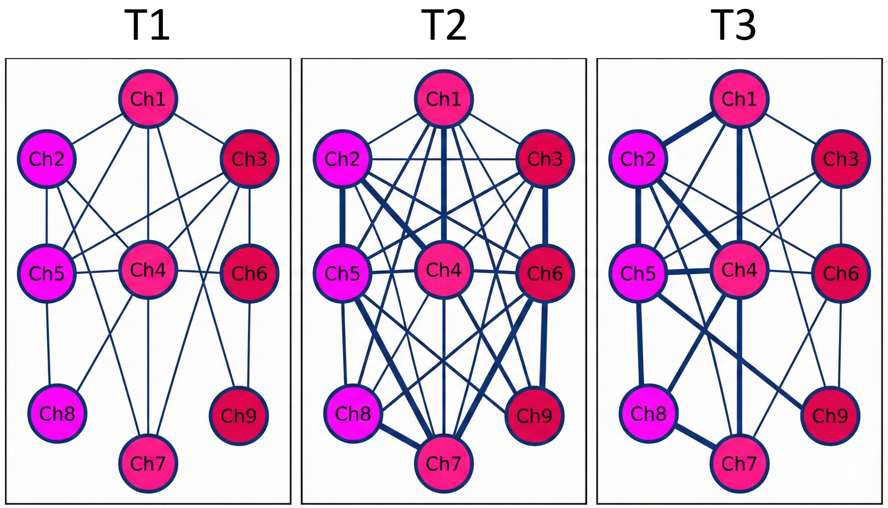

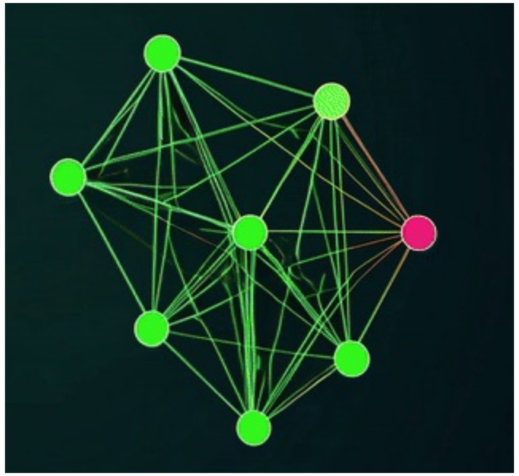

Animation source: SORA AI

3. Solution: Temporal GNN - Combining Space & Time

- TGNN: At each time step, do Graph Convolution → get spatial features
- Feed spatial features into GRU → integrates information over time steps
- After last time step, each sensor node's final state → classify severity (Normal/Mild/Moderate/Severe)
- Also outputs an overall normal/abnormal status for the whole graph (patient)
- Network of sensors + memory → outputs where and how bad is the issue

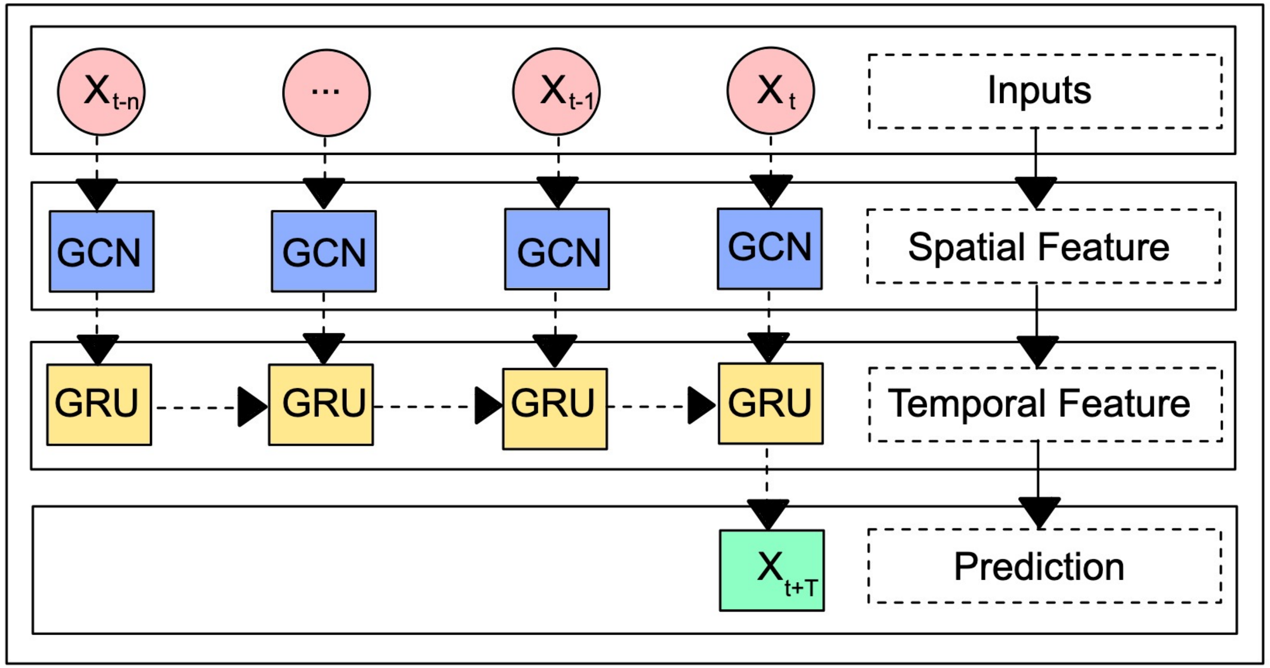

Image source: "Graph Neural Networks for the Localization of Breathing Abnormalities", S.M.H.Zaidi , Myra Spiliopoulou, 10.1109/CBMS65348.2025.00131

3.1 Graph Convolutional Network (GCN) – Capturing Spatial Patterns

- Graph Convolution: Each sensor node shares info with its neighbors
- At a given moment, combine readings from a node + nearby nodes
- Learns spatial dependencies: e.g., how one rib's motion relates to adjacent ribs
- Flags when one area's movement is off compared to its neighbors (injury indicator)

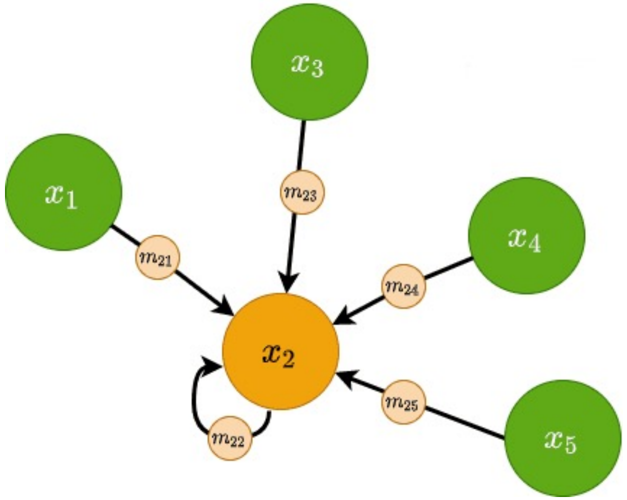

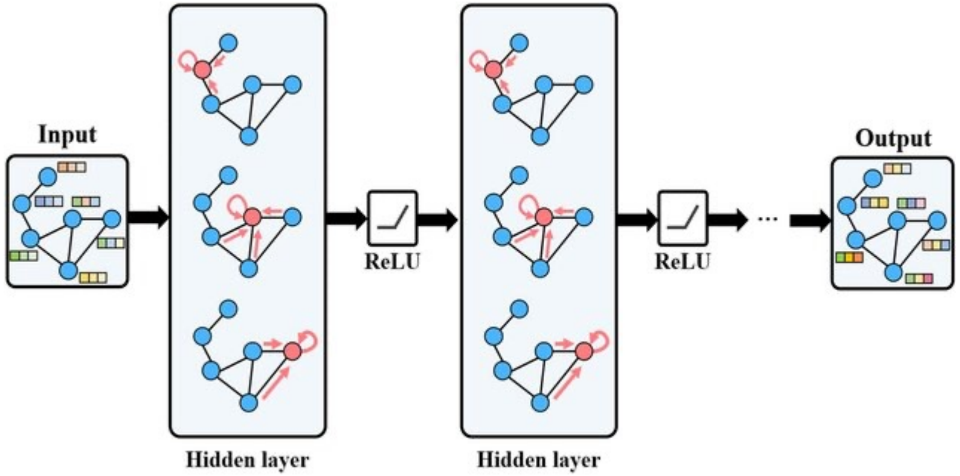

Image sources: 1) Adaloglou (AI Summer), Best Graph Neural Network architectures, 2021. https://theaisummer.com/gnn-architectures
2) Kim et al., Eng. with Computers, 2023. Doi: 10.1007/s00366-023-01811-0

3.2 Gated Recurrent Unit (GRU) – Capturing Temporal Patterns

- Breathing = time-series (inhale/exhale cycles); need memory of past states
- GRU (Recurrent Neural Net): carries forward important info from previous breaths
- Gating mechanism: decides what to remember vs forget (tracks breathing rhythm)
- Models patterns like consistent shallow breaths or changes over time (rate, depth)

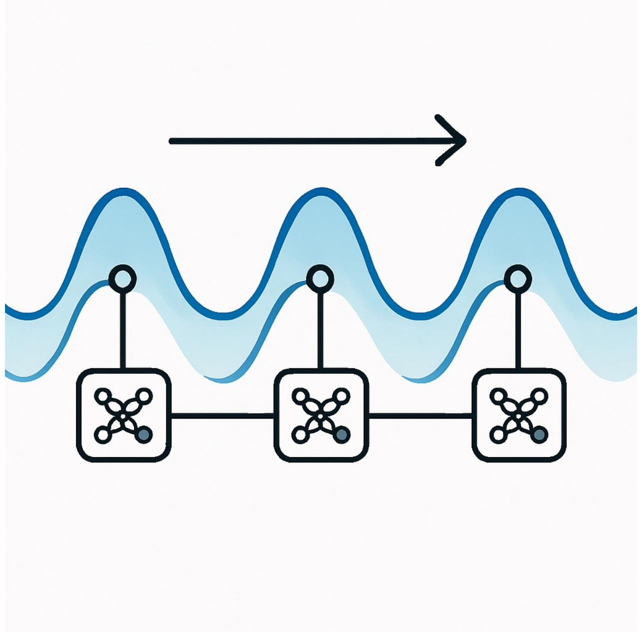

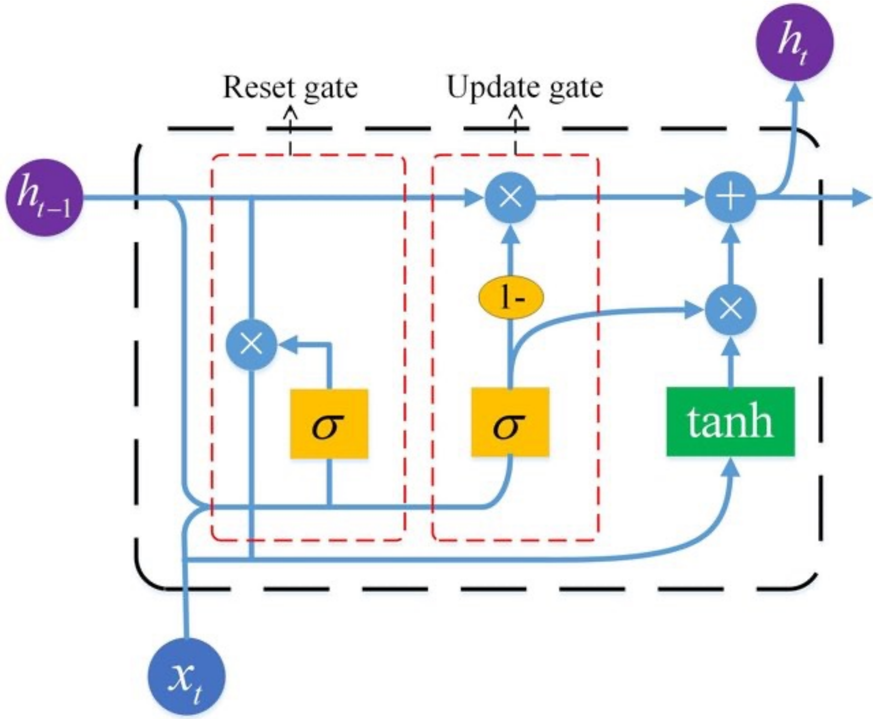

Image source: 1) Deep Learning Bible – Part C: LSTM, GRU, RNN, CNN, WikiDocs. https://wikidocs.net/166316

4. Case study 1: Localization of Breathing abnormalities

- Approximately 545 million people worldwide live with chronic respiratory diseases
- 4.1 million new cases of sternum and rib fractures in 2019
- Early detection & localization of respiratory problems
- Current methods (X-ray, CT) show injury location but are radiation-based and not continuous
- Clinicians ask “In which region is the pain and how severe?
- Particular location have abnormality (sensors are not faulty)

Overcrowded waiting rooms and patients queuing up are therefore a daily occurrence.

Image source: 1) https://www.aerztezeitung.de/Politik/Mir-fehlten-im-OP-die-Instrumente-235114.html
2) Gemini Nano Banana AI

4.1 Challenges

- Complex chest motion: Many regions move in different ways during breathing.
- Traditional sensors: Often single-point (can't tell where an issue is).
- Imaging limits: Great detail, but can't be done continuously (hospital-only, radiation).
- Lack of clinical data recorded from the patients.

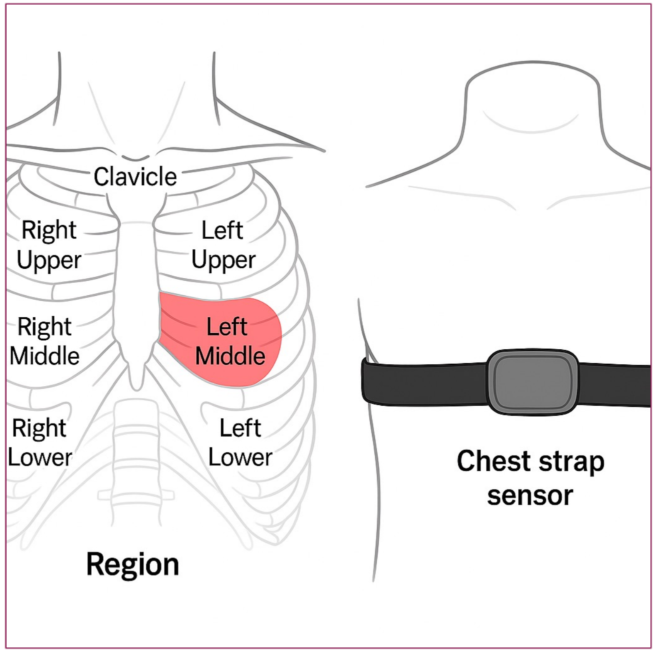

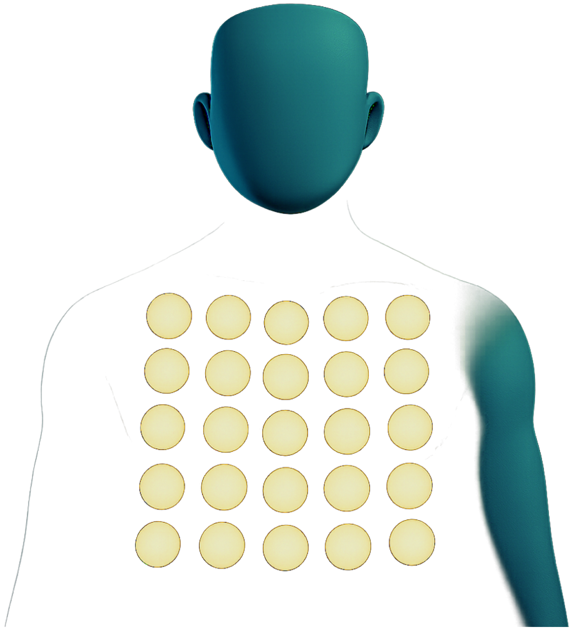

4.2 Synthetic Dataset

- Simulated 5 min abnormal breathing signals at 50Hz (reference dataset = PAMAP2)
- Severity levels: Mild, Moderate, Severe – different impact on signal amplitude & frequency
- Introduced “injuries” at specific sensor locations with neighbours also slightly affected (Influence Factor)
- Shallow breaths, rapid breathing and gaussian noise added to mimic real breathing abnormal motion
- This synthetic setup is ideal for explainability research because we know the true causal structure.

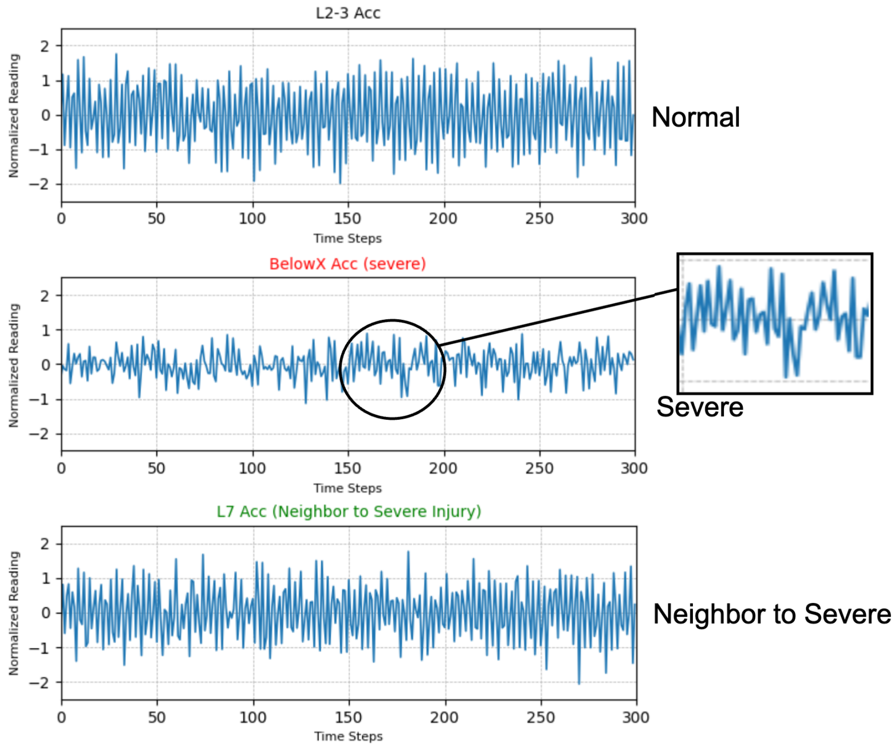

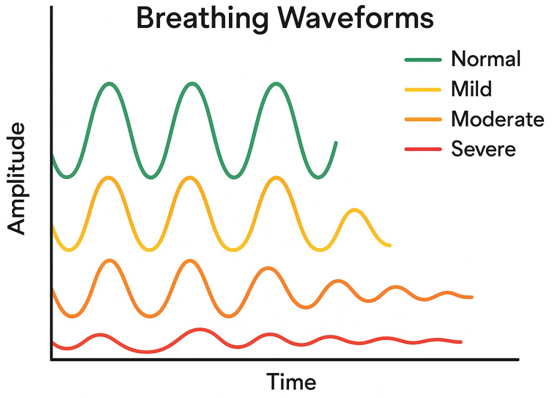

4.3 Why Graphs for Breathing?

- Performance of thorax decreases as the area of affected segment increases
- Capture this performance via chest wearable sensors
- Temporal GNN learns spatial and temporal dependencies
- Goal: Detect abnormal breathing and pinpoint its location simultaneously

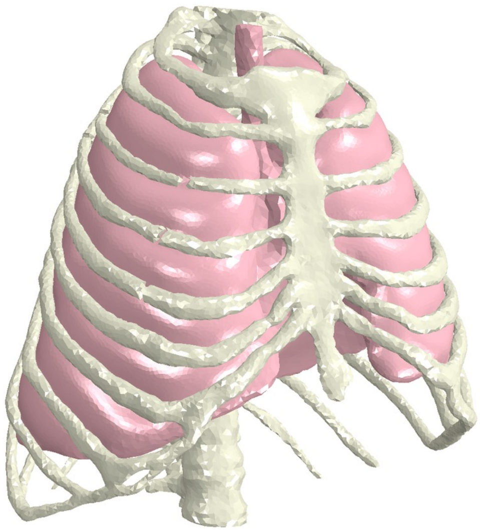

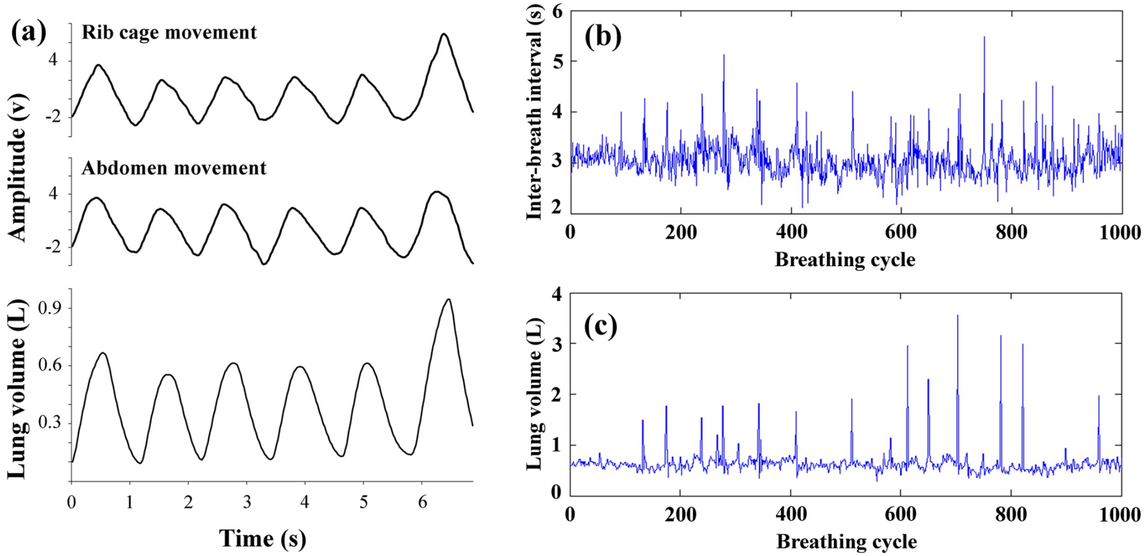

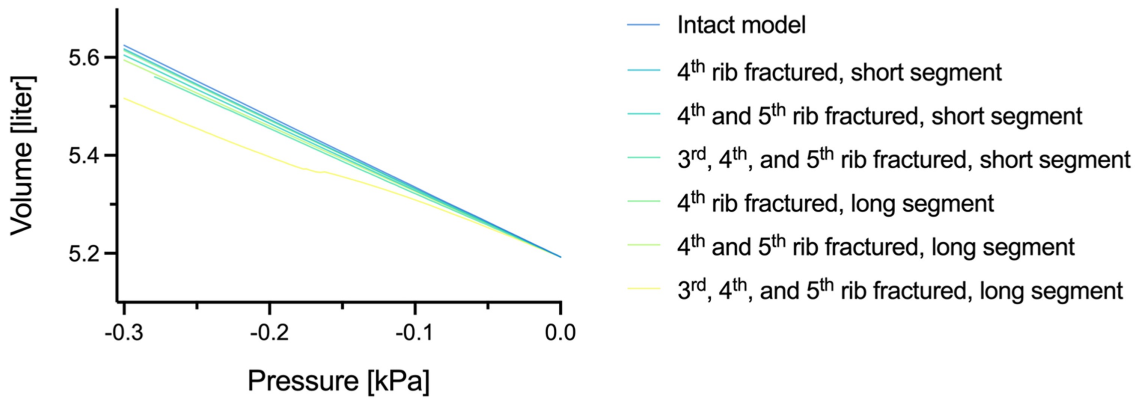

Image sources: 1) Raoufy et al., (2016). Classification of Asthma Based on Nonlinear Analysis of Breathing Pattern. Doi: 10.1371/journal.pone.0147976
2) Zierke et al., (2025). Biomechanics of flail chest injuries: tidal volume and respiratory work changes in multiple segmental rib fractures. Doi: 10.1007/s00068-024-02754-x

4.4 Wearable Sensor Network Modelled as a Graph

- 14 IMUs sensors covering the thorax region
- Each sensor = a node in a graph (captures motion in one area)
- Edges: capture spatial proximity and signal correlation (exists only if both close & correlated)
- Graph = map of the chest – Each sensor’s time-series as each node embedding

Spatial Proximity: Connect sensors within a distance threshold τs

Temporal Correlation:
1) Pearson correlation on time series
2) Keep edges with correlation ≥ τc

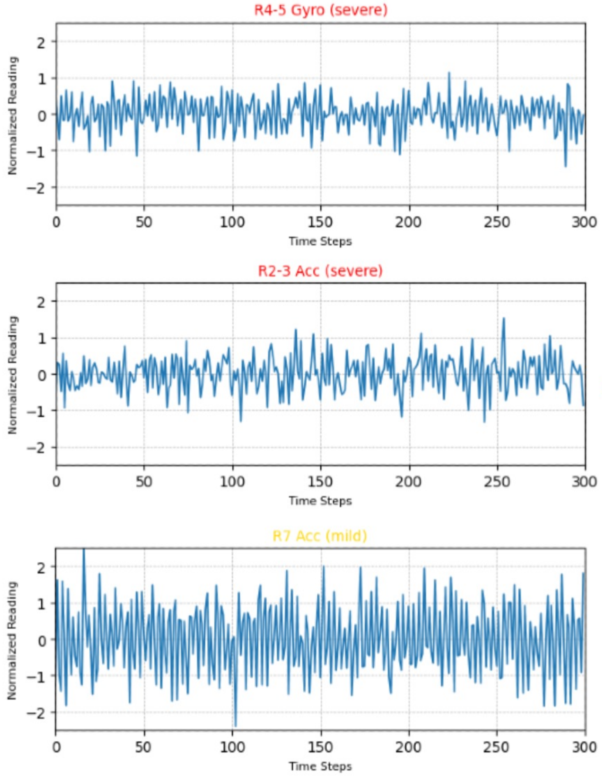

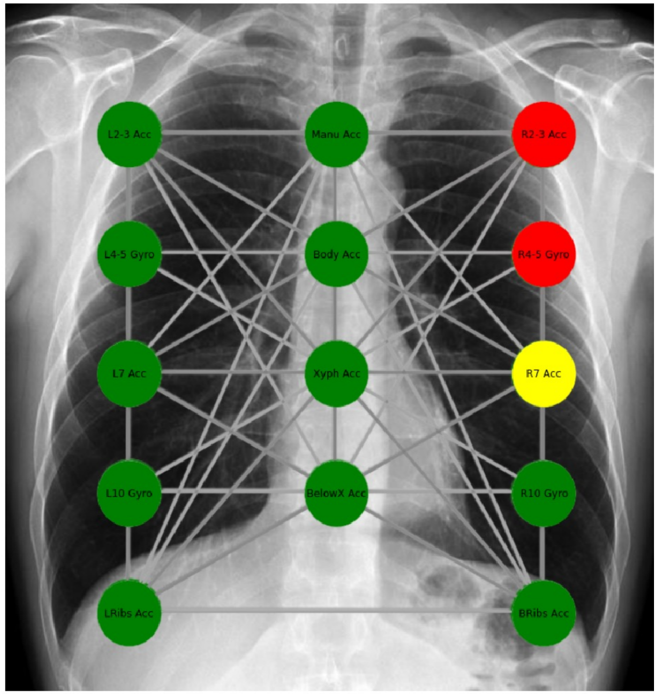

Image source: “Graph Neural Networks for the Localization of Breathing Abnormalities”, S.M.H.Zaidi , Myra Spiliopoulou, [10.1109/CBMS65348.2025.00131](https://doi.org/10.1109/CBMS65348.2025.00131)

4.5 Results – Classification Performance

- Normal class: near-perfect precision/recall (breathing normal regions easily identified)
- Severe class: high precision (~0.86) and recall (~0.94) – model rarely misses a severe issue
- Moderate class: well classified (Precision= 0.77, recall= 0.86)
- Mild class: high recall (0.82) but lower precision (0.51) – some mild cases flagged as moderate (Confusion distinguishing mild vs moderate severity)

CONFUSION MATRIX

Actual \ Predicted	Normal	Mild	Moderate	Severe
Normal	734	30	3	0
Mild	20	32	6	0
Moderate	15	1	37	5
Severe	0	0	2	31

CLASSIFICATION REPORT

Class	Precision	Recall	F1-score
Normal	1.00	0.99	0.99
Mild	0.51	0.82	0.63
Moderate	0.77	0.86	0.81
Severe	0.86	0.94	0.90

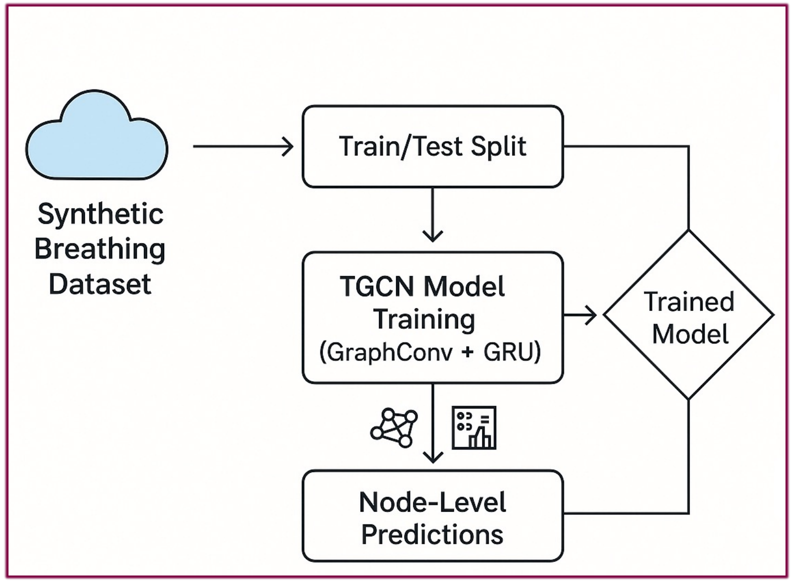

Table source: "Graph Neural Networks for the Localization of Breathing Abnormalities", S.M.H.Zaidi , Myra Spiliopoulou, 10.1109/CBMS65348.2025.00131

4.6 Results – Localizing an Abnormality (Example)

- Model pinpoints where an abnormality is located on the chest (matches ground-truth location)
- Predicted an abnormality on upper right side (near sensors “R2-3 Acc” & “R4-5 Gyro”)
- Visualization: normal = green, mild = yellow, moderate = orange, severe = red

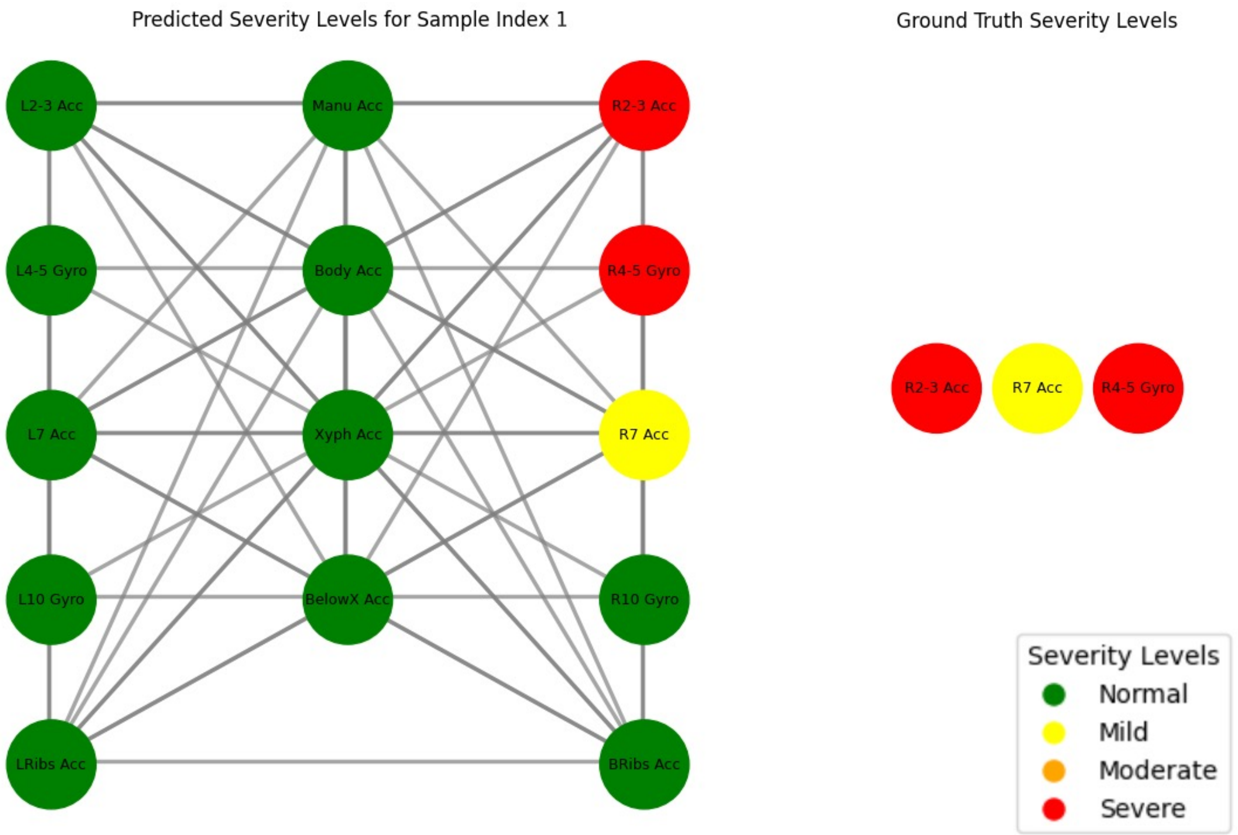

Image source: “Graph Neural Networks for the Localization of Breathing Abnormalities”, S.M.H.Zaidi , Myra Spiliopoulou, [10.1109/CBMS65348.2025.00131](https://doi.org/10.1109/CBMS65348.2025.00131)

5. Why XAI is Essential Here

- Why upper-right rib region?
- Black-box localization is insufficient in medical decision support.
- Yellow Sensor got some influence from the neighbouring sensors?
- Need to justify localization.
- Monitoring the changes in trend of each sensor with its neighbourhood.

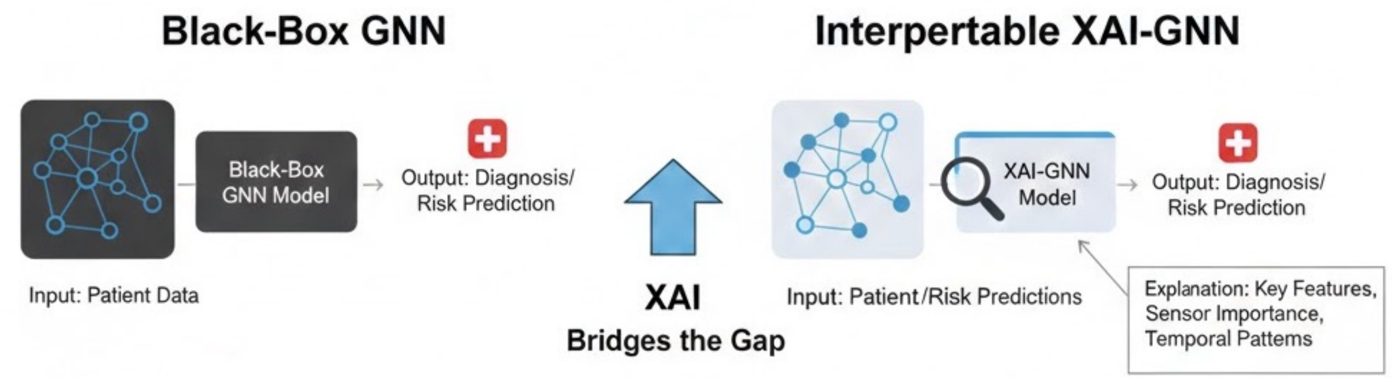

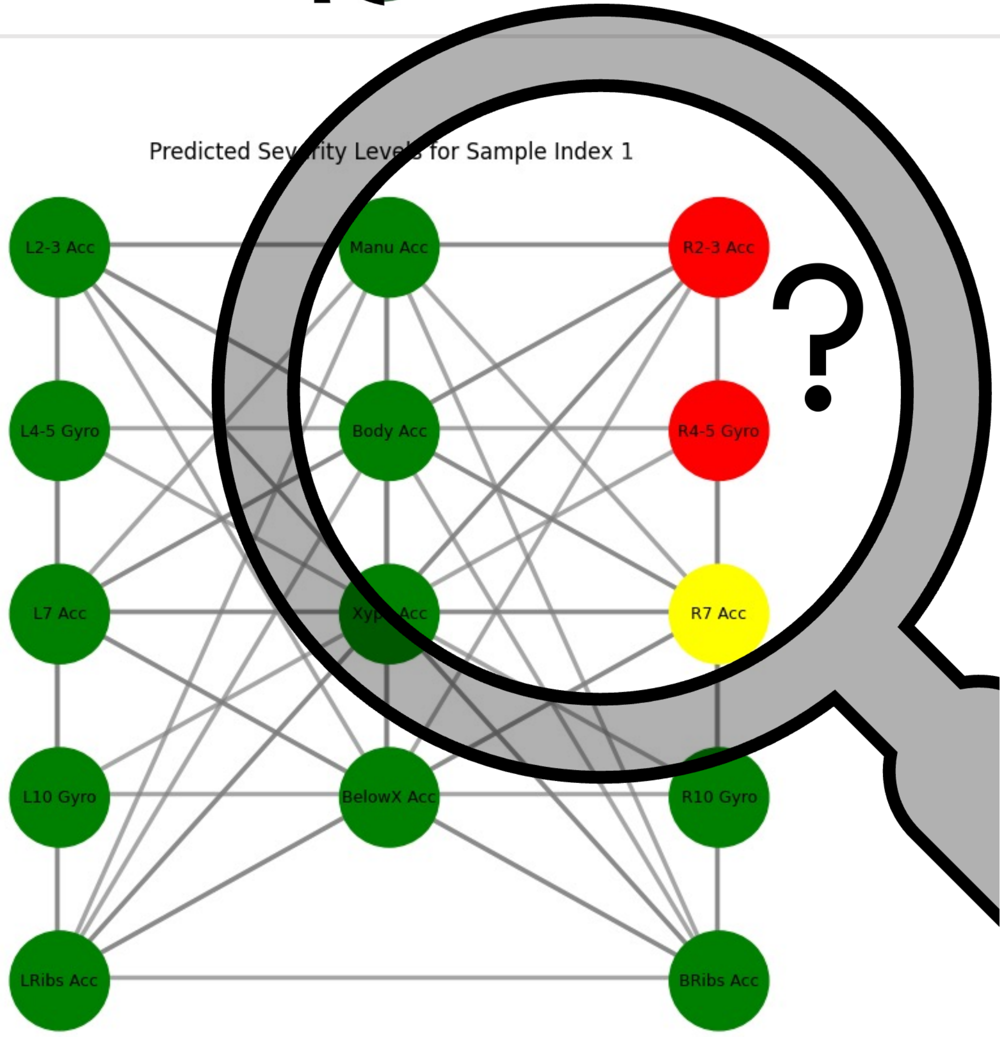

5.1 XAI for GNNs: State of the Art

GNN explainability:

- GNNExplainer
- GNN-Monte Carlo Tree Search (MCTS)
- PGExplainer
- SubGraphX

What they explain:

- Node Importance
- Edge Importance
- Subgraph explanations

❌ Limitation:

- Mostly static & Snapshot-based
- Ignore time dimension
- Do not explain memory

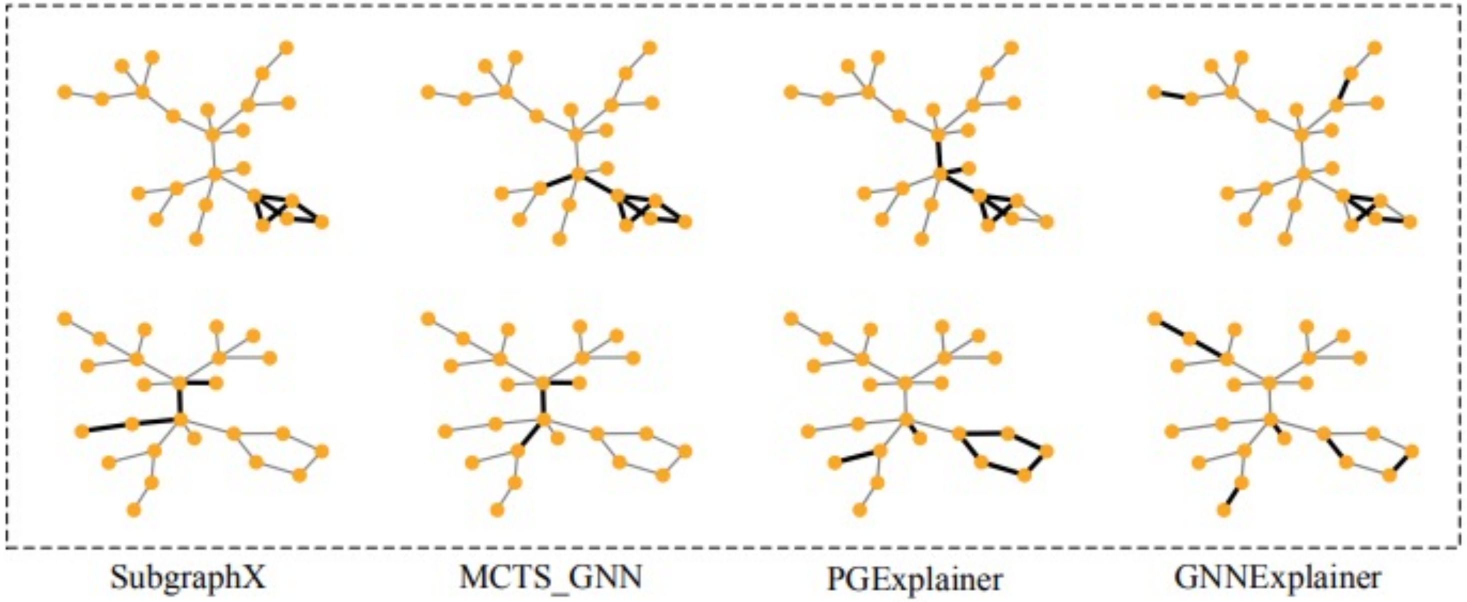

Method	MCTS*	MCTS†	SubgraphX	GNNExplainer	PGExplainer
TIME	>10 hours	865.4 ± 1.6s	77.8 ± 3.8s	16.2 ± 0.2s	0.02s (Training 362s)
FIDELITY	N/A	0.53	0.55	0.19	0.18

Image source: Yuan, H., Yu, H., Wang, J., Li, K., & Ji, S. (2021). *On Explainability of Graph Neural Networks via Subgraph Explorations*.

5.2 XAI for TGNNs – How to Explain Abnormailites

- **When and Where**: Need explanations that highlight not just which sensor (node) was important, but at what time step (or range) it became important. (I.e. identifying the exact breath cycle when the anomaly was detected)
- **Dynamic relationships:** The graph structure itself can change with time (i.e. correlations between sensors might shift during abnormal events)
- **Local vs Global**: need a way to communicate a sequence of events (i.e. “Sensor A’s drop followed by Sensor B’s drop caused the alert”).
- **Computational load**: Generating explanations for complex models like TGNN are computationally expensive and slow.

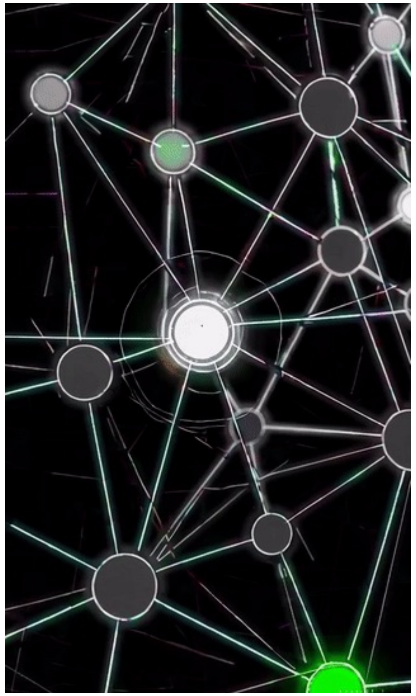

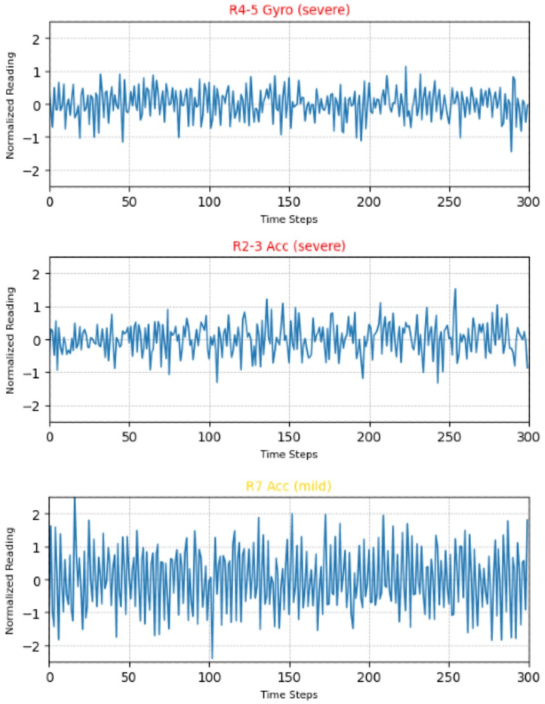

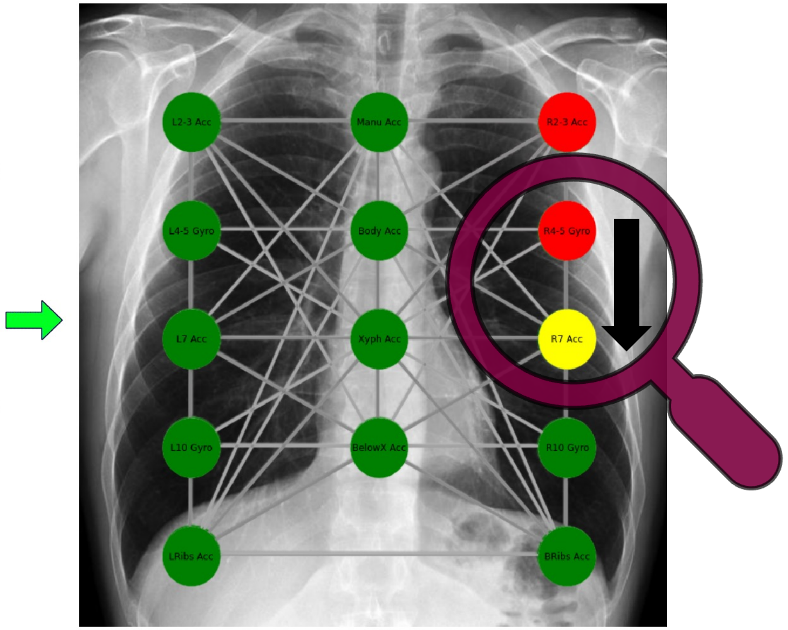

6. Future Direction and Ideas

1. **Temporal counterfactuals**
   - What change in the breathing signal would have prevented this alert? (Signal Deviation)
   - Spatio-temporal saliency maps: highlighting which sensor at which time contributed most to the anomaly (a heatmap over the sensor network × time)

2. **Clinically aligned explanations**
   - Translating sensor-based findings into clinical terms (i.e. reduced movement in upper-right ribcage during exhalation)
   - The ultimate aim is a robust, explainable system that the medical community can trust.

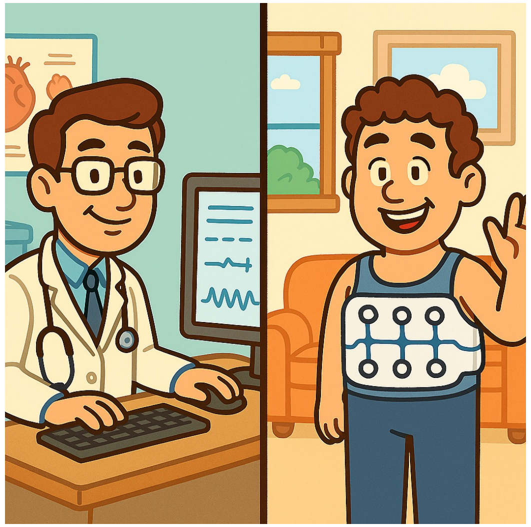

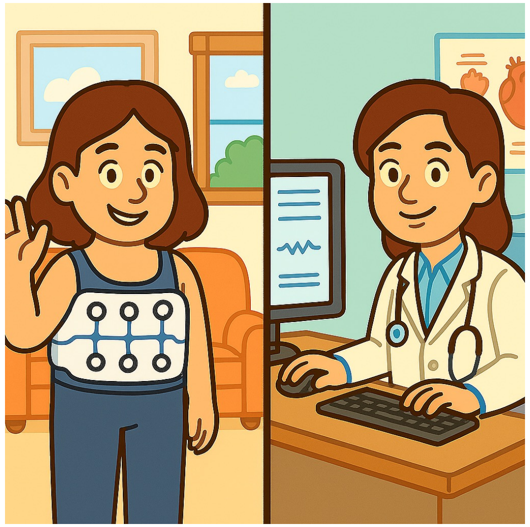

- 9 iEEG channels
- Each channel = graph node
- Goal: detect faulty / anomalous channels
- TGNN highlights inconsistent channels
- Some channels are faulty.

Experiment setup and data type:
Species: Freely moving, healthy Mongolian gerbils
Behavioral state: Passive listening (no task demands)
Stimuli: Pure auditory tones (controlled frequencies, amplitudes and durations)
Data: Extracellular electrophysiological recordings from 2 brain regions : ***primary auditory cortex(A1)*** and ***striatum(STR).***
Continuous voltage traces containing neural events such as spikes, local field potentials, and noise/artifacts from several sources.
*High sampling rate(30 kHz). *Time stamped video recordings are also available.

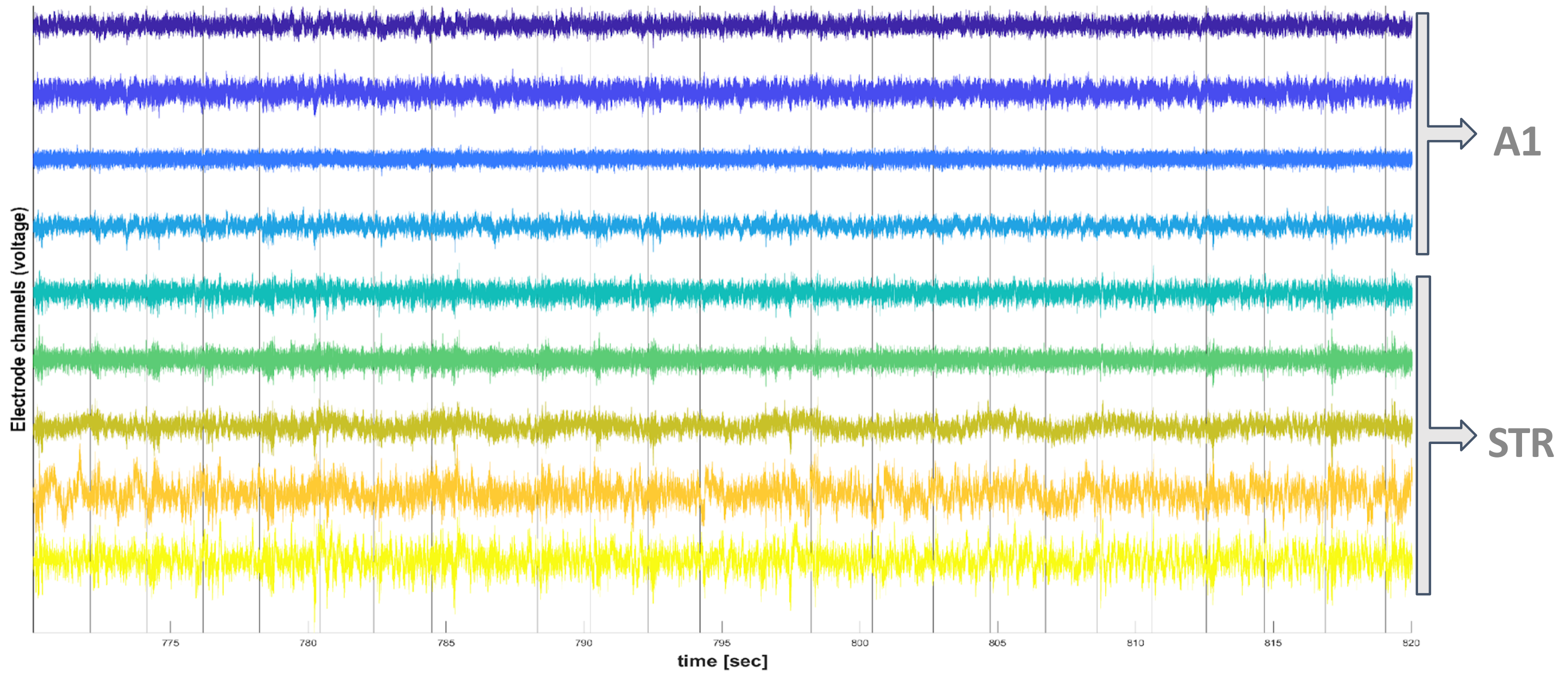

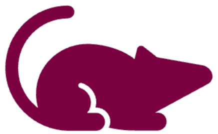

7.1 Another Case Study: Gerbil Electrophysiology

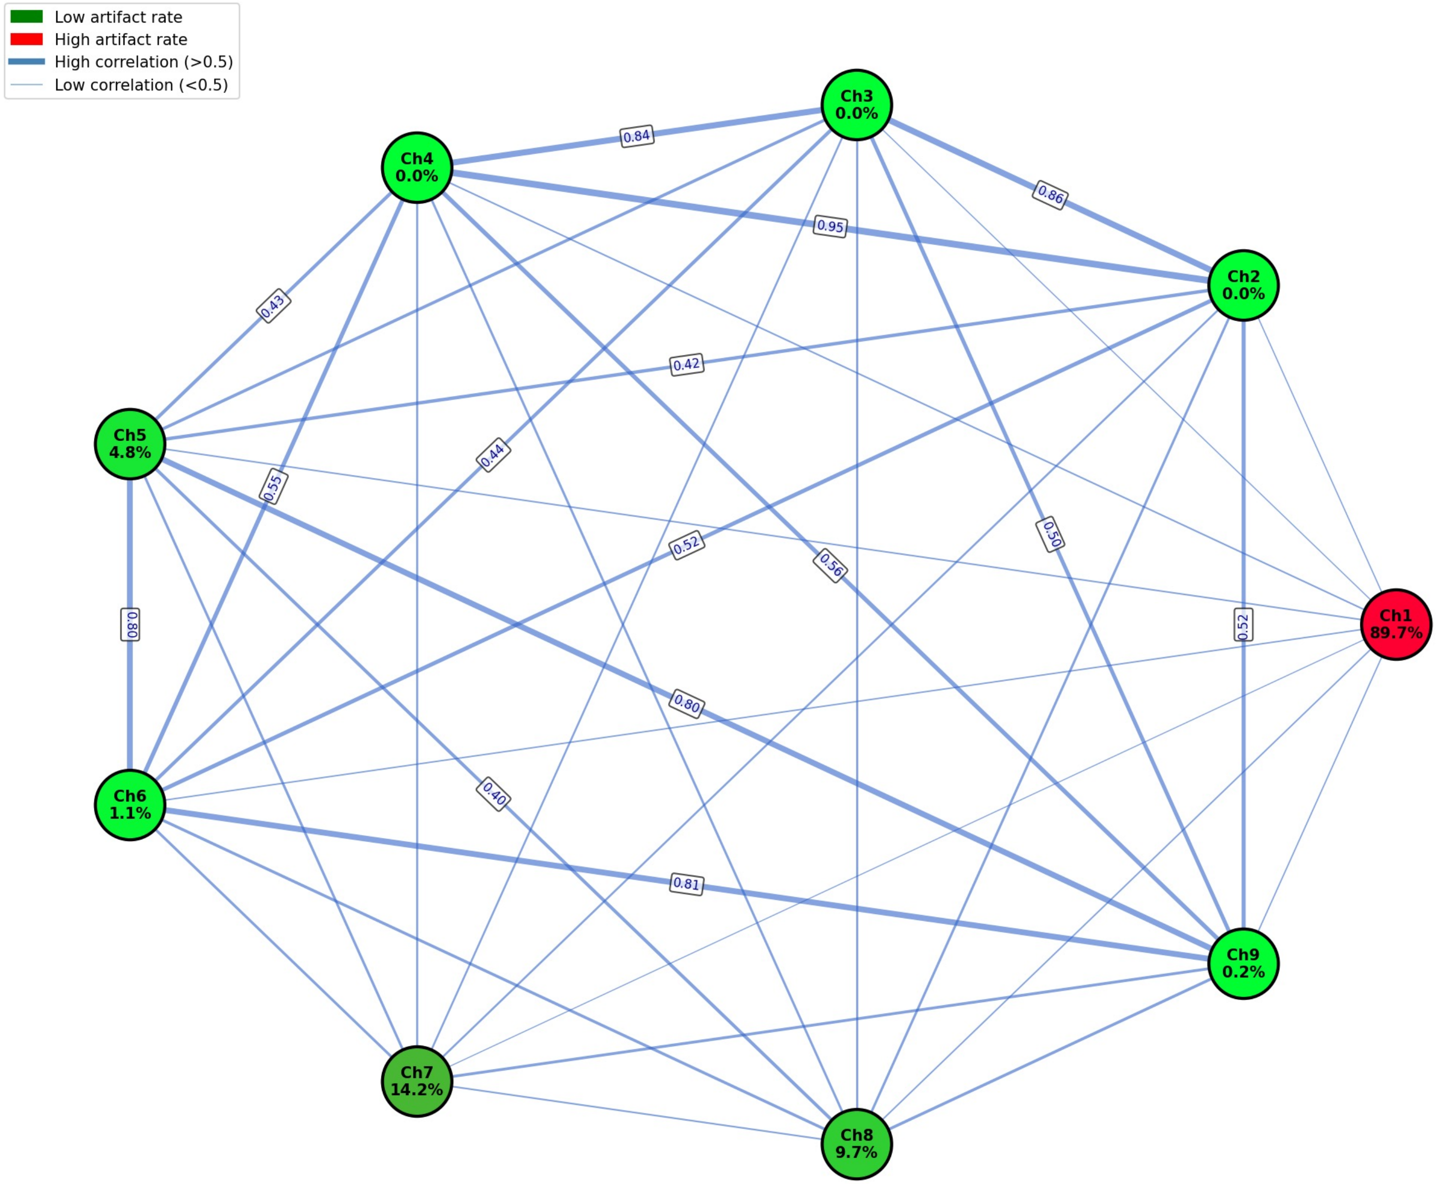

Thank you for your attention!

TACTIC

Towards Co-Evolution in Human-Technology Interfaces

SACHSEN-ANHALT

Finanziert von der Europäischen Union

OTTO VON GUERICKE UNIVERSITÄT MAGDEBURG

FACULTY OF COMPUTER SCIENCE

Knowledge Management & Discovery Lab