FireGNN: Neuro-Symbolic GNN with Trainable Fuzzy Rules

Prajit Sengupta, and Islem Rekik

BASIRA Lab, Imperial-X and Department of Computing, Imperial College London, London, UK

This work is part of my research at BASIRA lab:
https://basira-lab.com
GitHub code: https://github.com/basiralab/FireGNN

Imperial College London

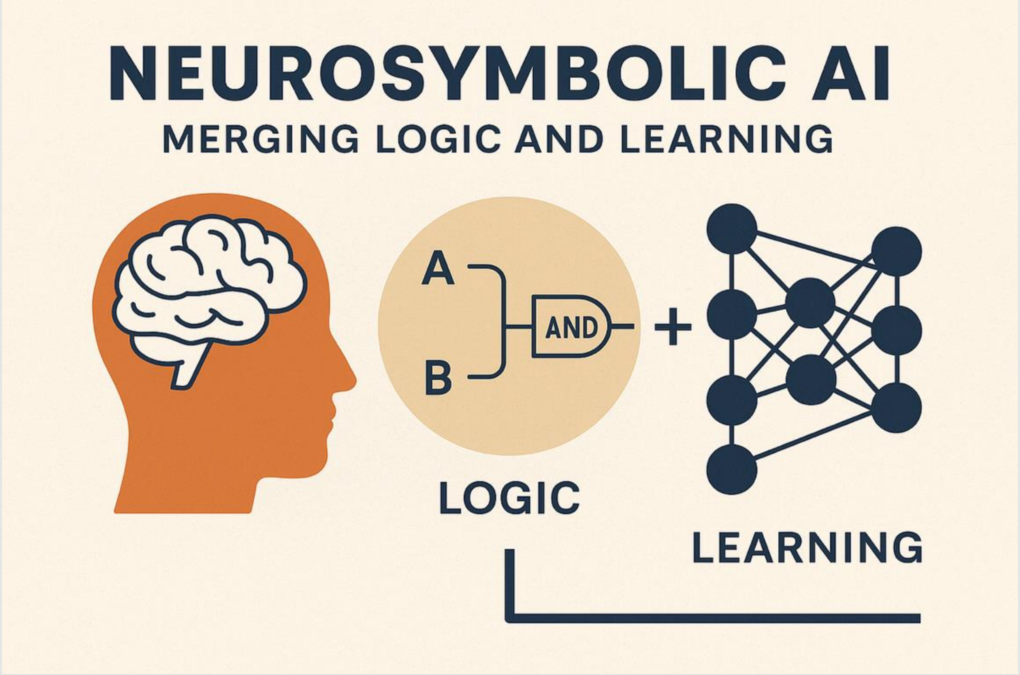

Source: Bing

NEURAL INFORMATION PROCESSING SYSTEMS

The Medical AI Dilemma!!

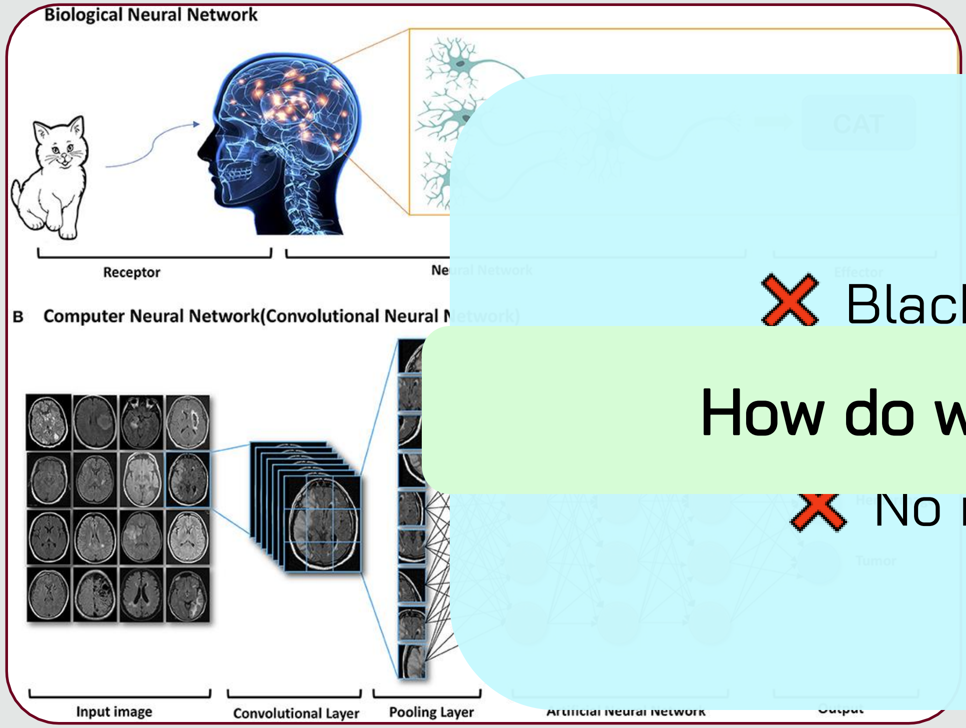

Source: Bing

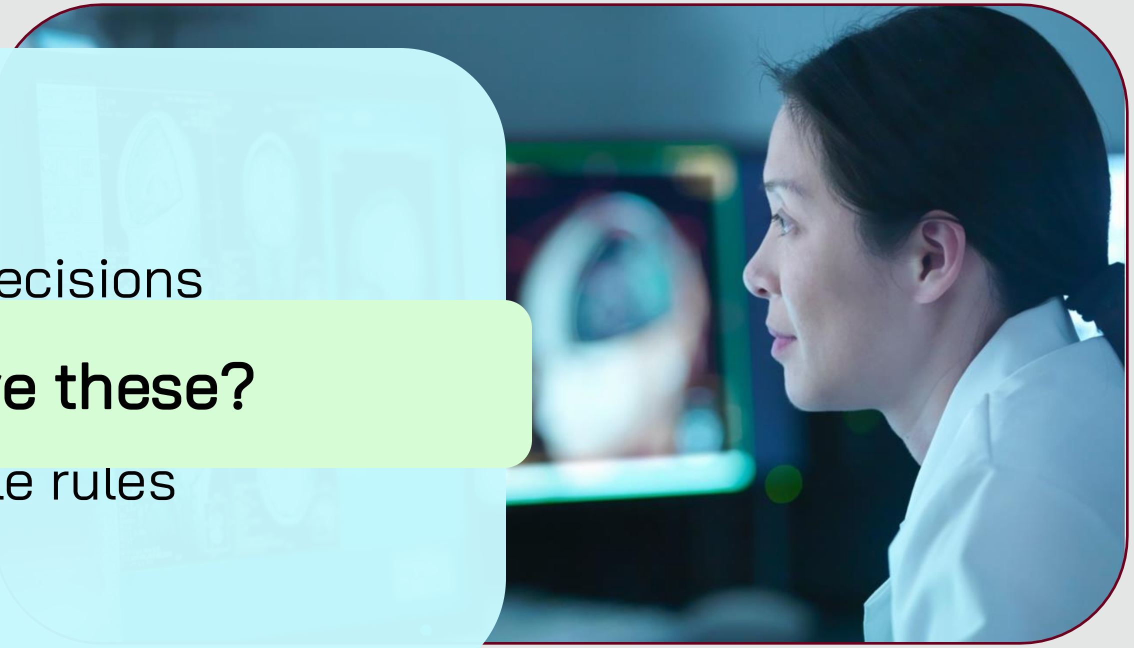

Source: Bing

❌ Black box decisions

How do we solve these?

❌ No readable rules

"Dr. Sarah has a 95% accurate AI system, but when it says 'malignant tumor,' she asks 'Why?' The AI stays silent..."

Neuro-Symbolic AI

Bridging Neural Intuition with Logical Reasoning

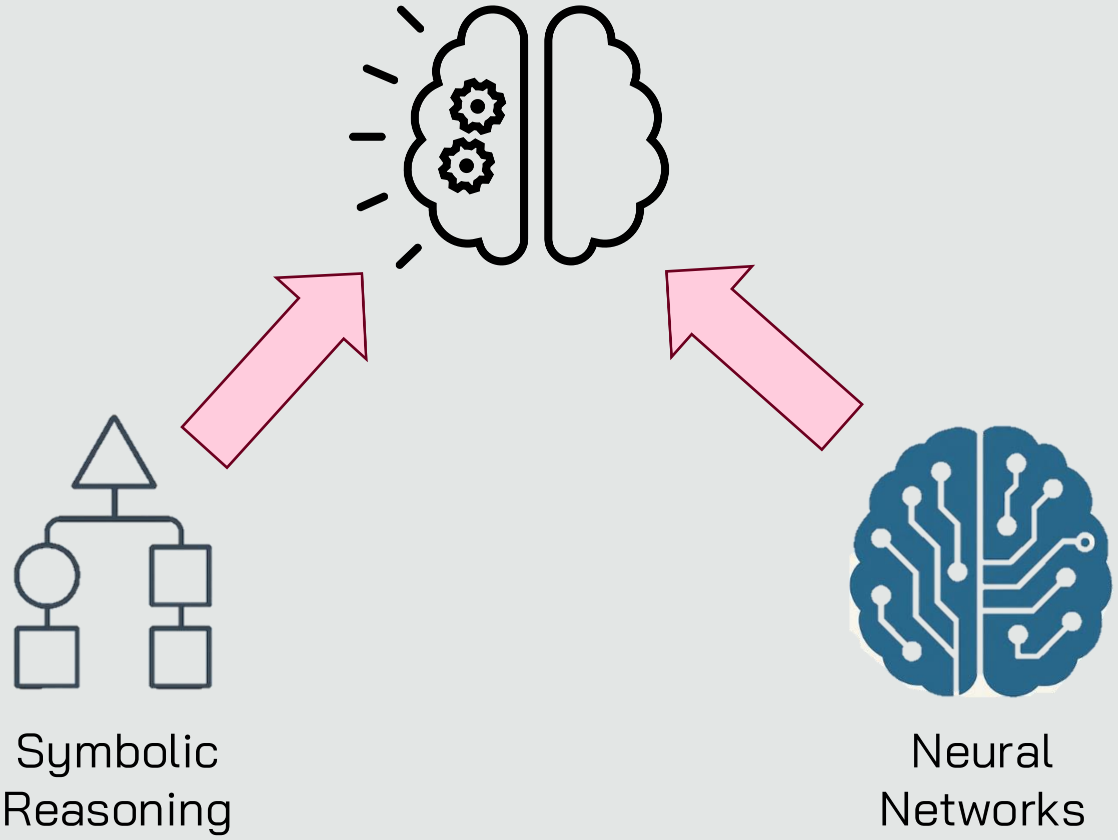

Examples of Neuro-Symbolic AI

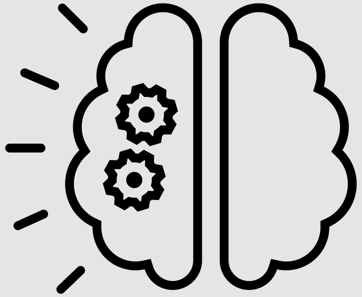

Logic Rule: $\text{IF Cat}(x) \Rightarrow \text{Mammal}(x)$

Visual: $\text{IF Red}(x) \wedge \text{Cube}(x) \wedge \text{LeftOf}(x, y) \wedge \text{Sphere}(y) \Rightarrow \text{True}$

Relational: $\text{IF Parent}(x, y) \wedge \text{Parent}(y, z) \Rightarrow \text{Grandparent}(x, z)$

RL Constraint: $\text{IF At}(agent, \text{RedZone}) \Rightarrow \text{Penalty}$

Probabilistic Fact: $\text{Pr}(\text{Rain}) = f_\theta(\text{WeatherImage})$

A More Flexible Logic: What are Fuzzy Rules?

Instead of thinking in black-and-white (a rule is either TRUE or FALSE),
fuzzy logic allows for degrees of truth

$$\text{IF degree} \geq \theta_1 \text{ THEN high\_connectivity}$$

$$\text{IF degree} \geq 10 \text{ THEN high\_connectivity } \text{“}\theta = 0.5\ true\text{”}$$

$$\text{IF degree} \geq 15 \text{ THEN high\_connectivity } \text{“}\theta = 0.6\ true\text{”}$$

*“A fuzzy rule doesn't just turn on or off; it activates with a certain **strength**.”*

BASIRA

Problem with Today's Neuro-Symbolic AI?

$$\text{IF degree} \geq 10 \text{ THEN high\_connectivity}$$

Fixed Rule Hand Crafted

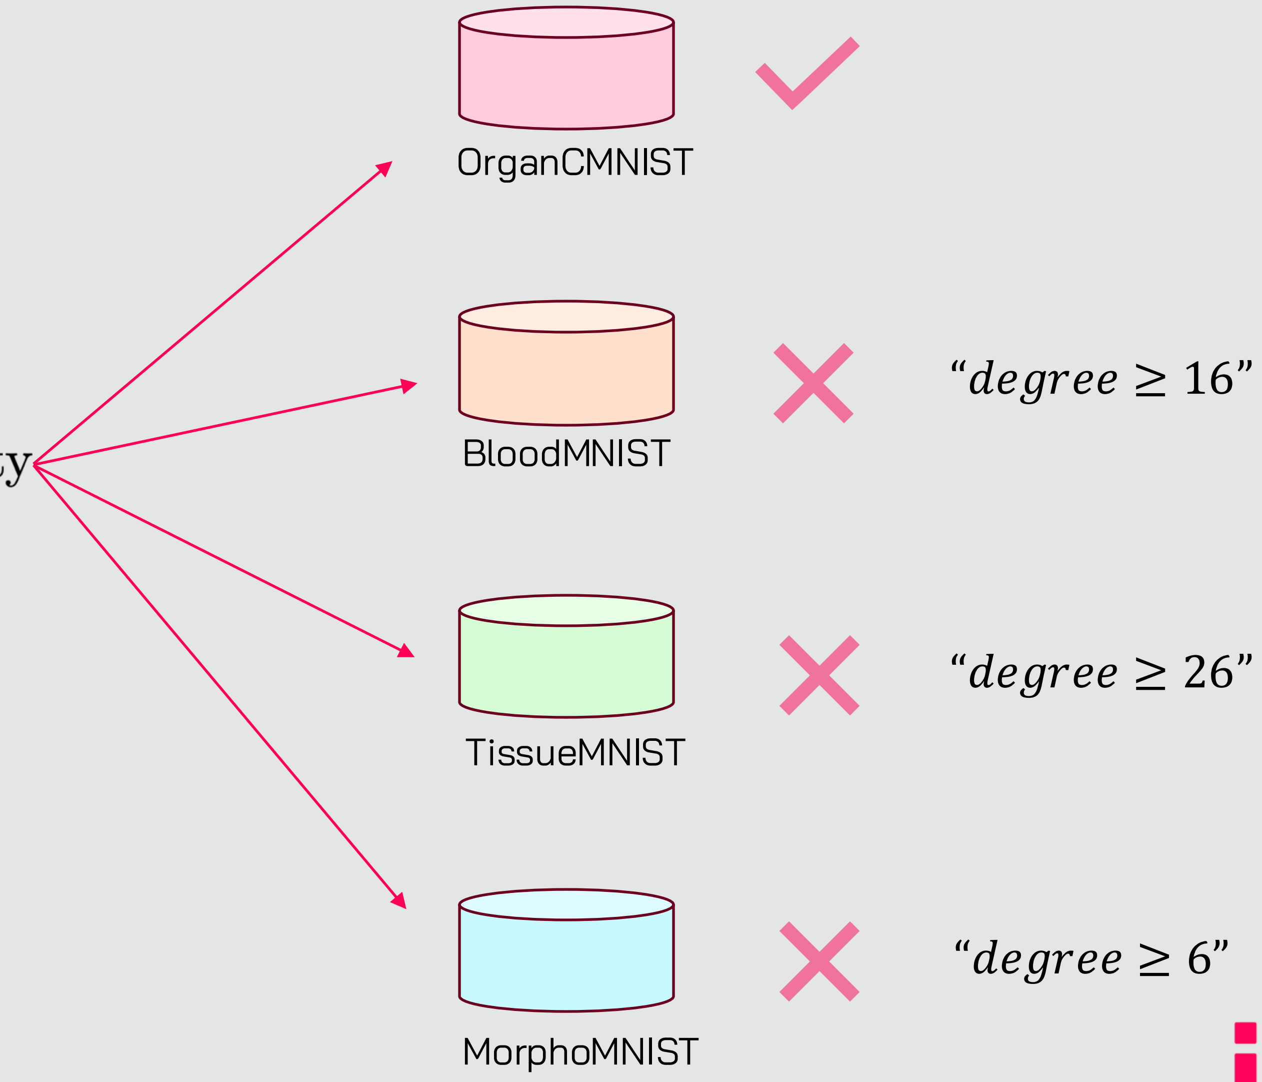

Diverse Datasets

FireGNN Architecture Overview

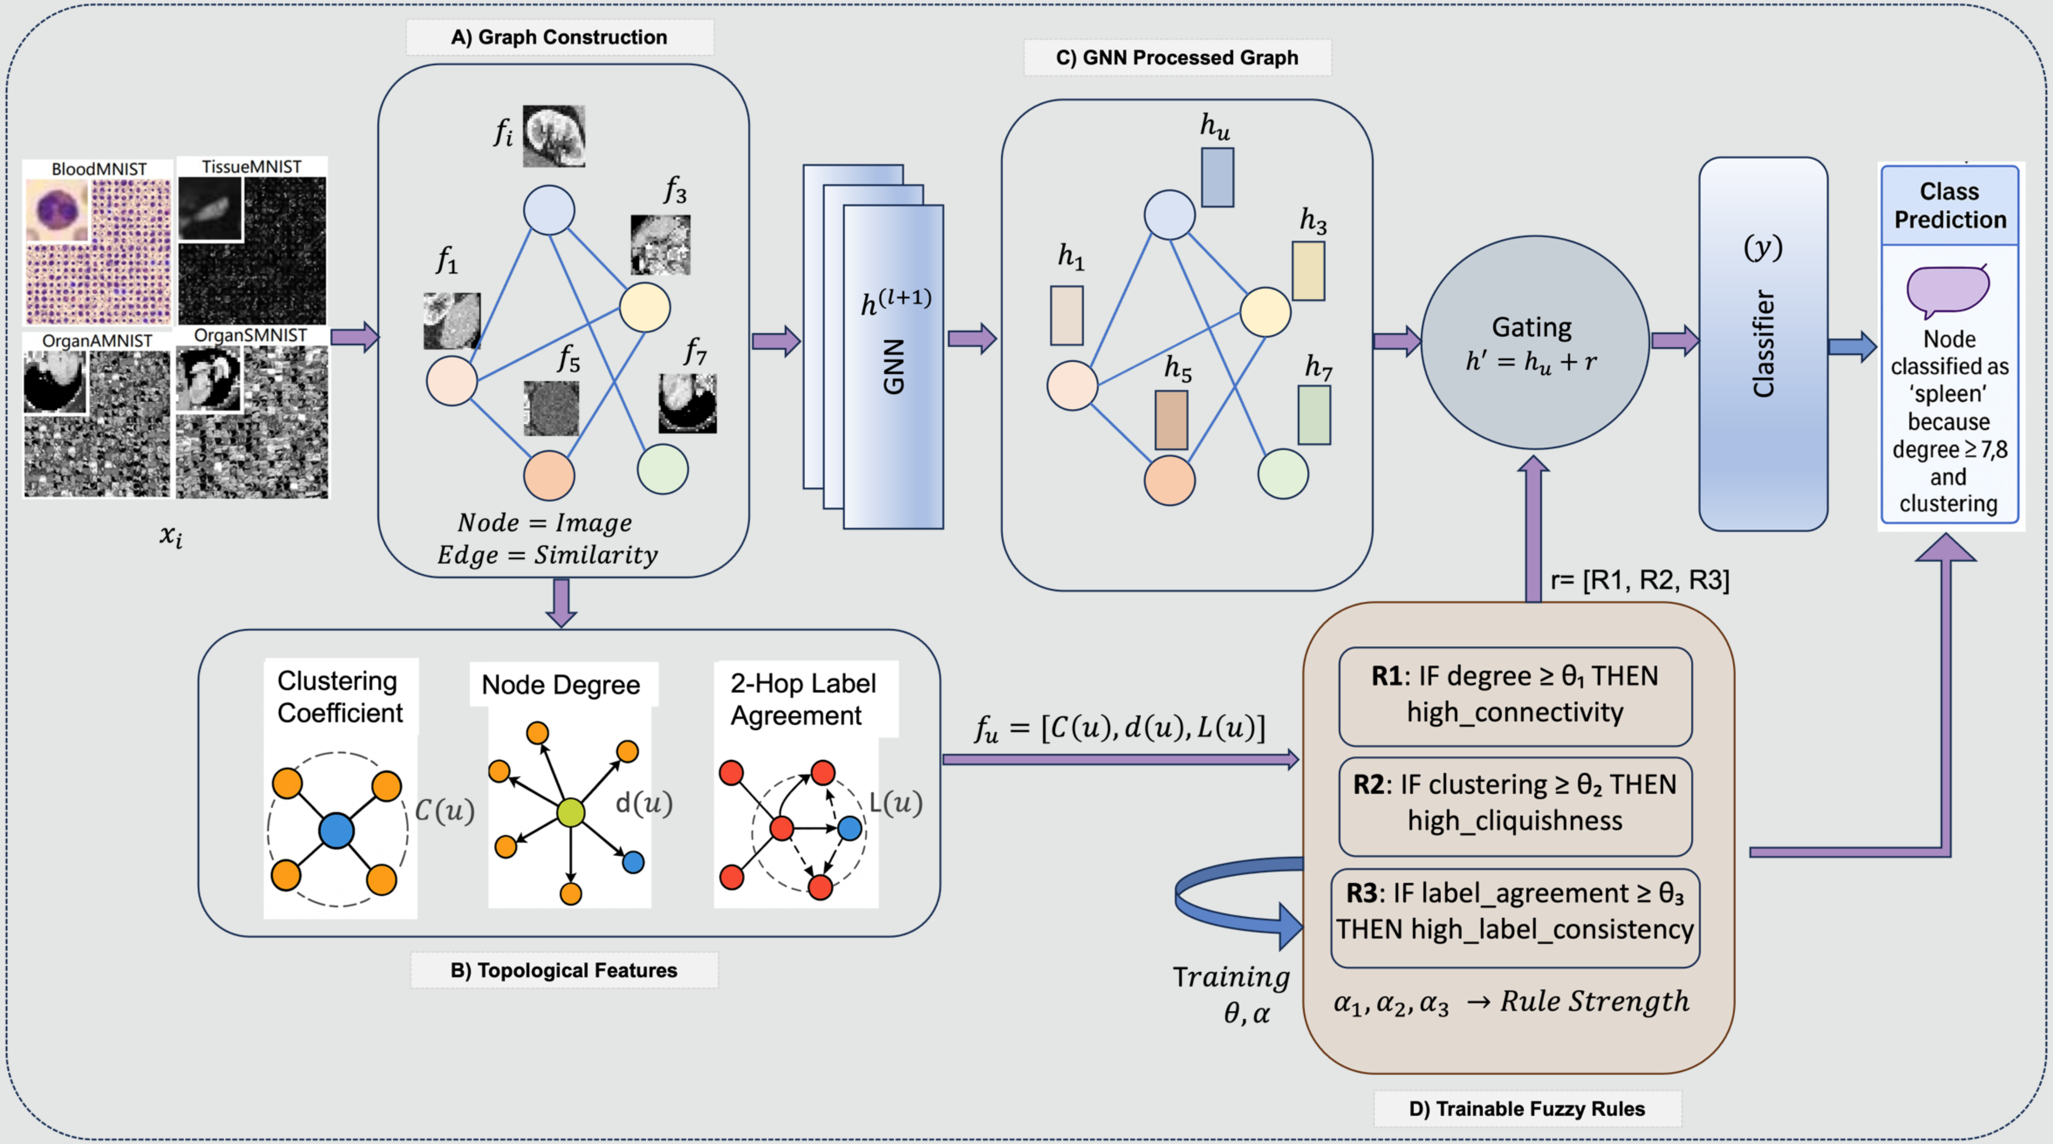

Step A & B - Building the Graph & Features

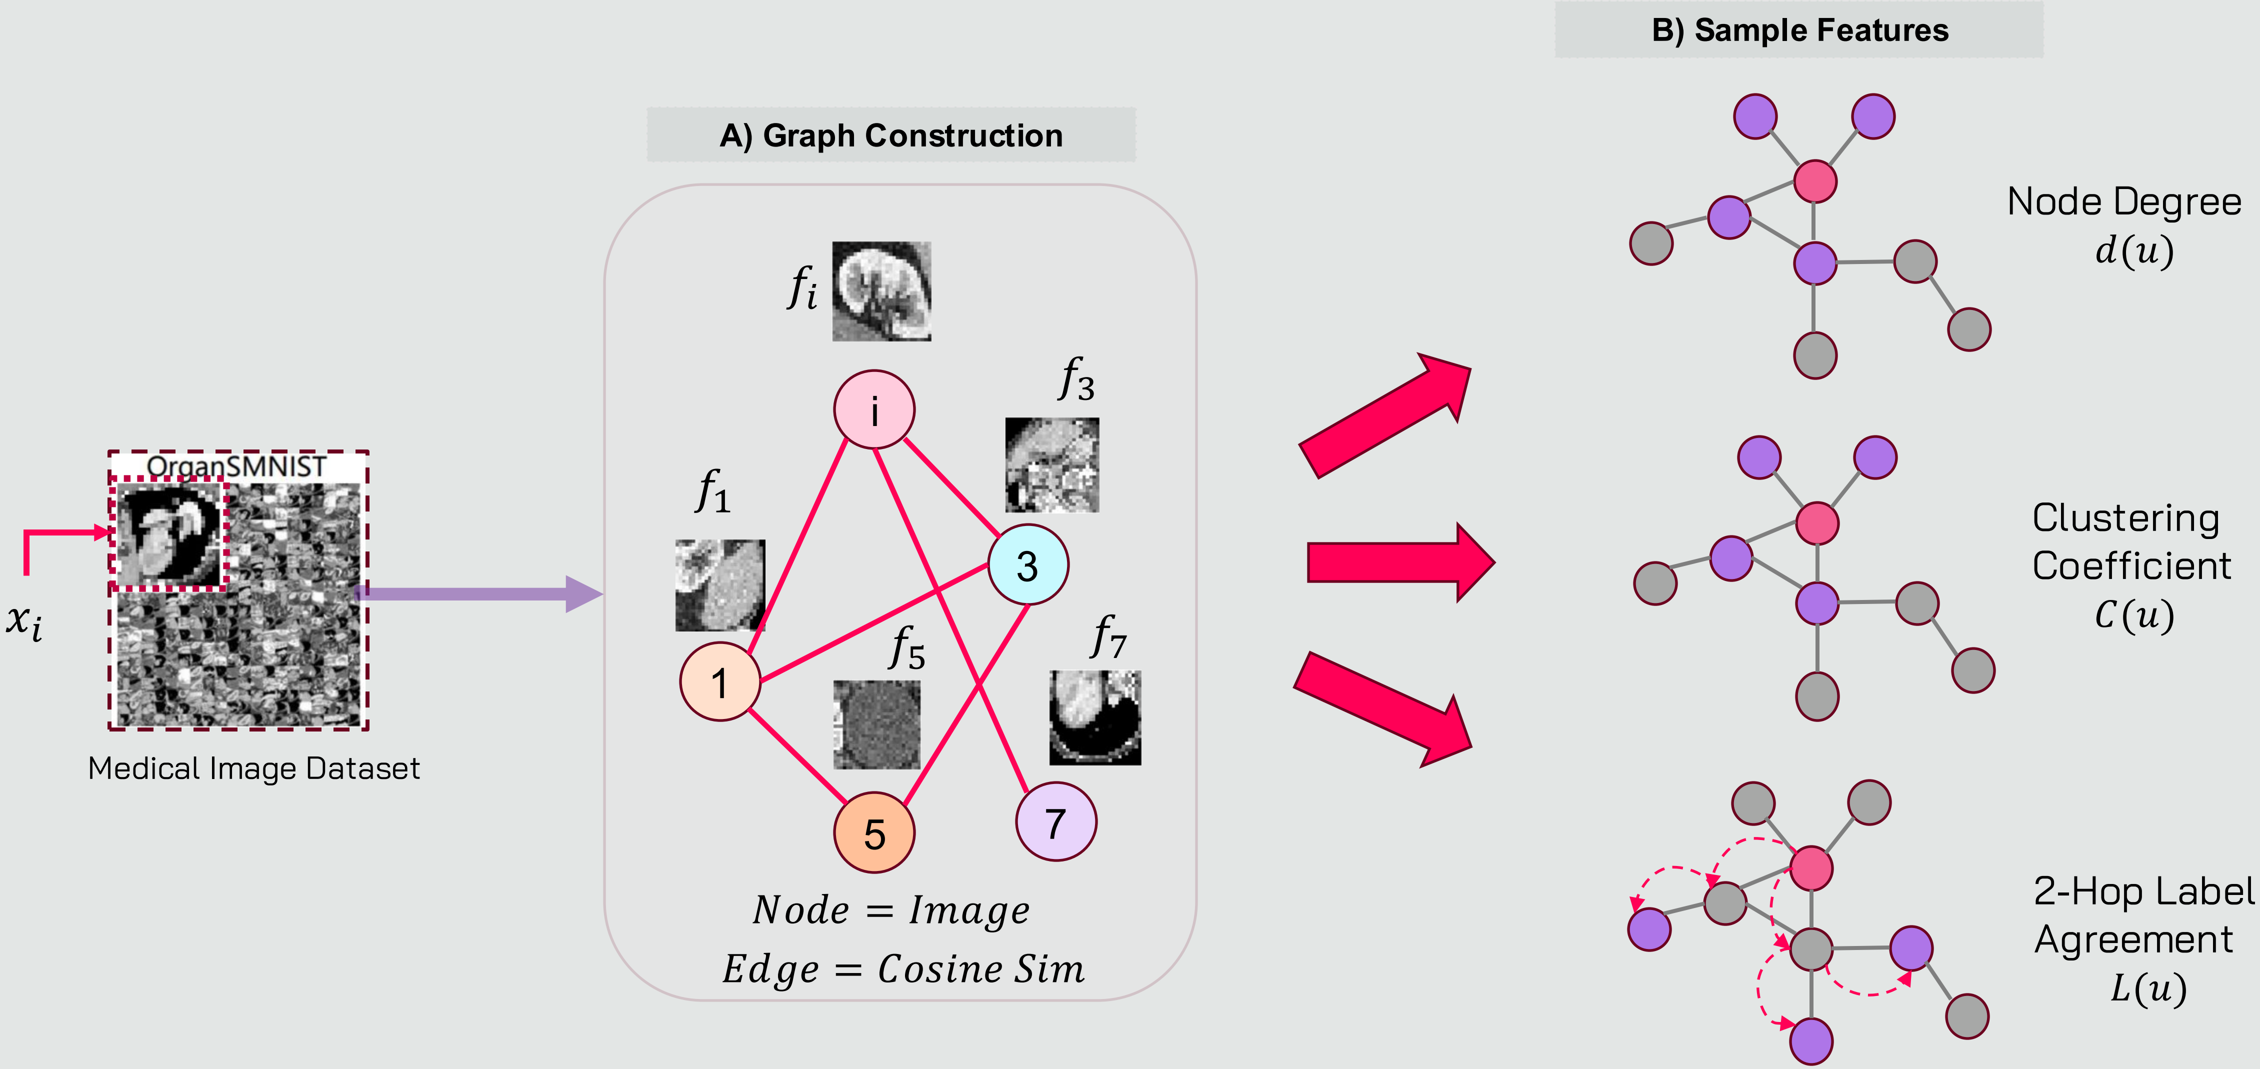

Rules

- To build its explanations, FireGNN grounds its logic in three fundamental, human-understandable properties of a node's structure within the graph.

Decision Boundary

Rule Activation Sharpness

$$\text{IF degree IS high\_connectivity } (\theta_1 =?,\ \alpha_1 =?) \text{ THEN } \ldots$$

- These rules form the symbolic building blocks for its reasoning.

But these rules are fixed??

$$\text{Rule 1:}\quad \text{IF degree} \geq \theta_1 \text{ THEN high\_connectivity}$$

$$\text{Rule 2:}\quad \text{IF clustering} \geq \theta_2 \text{ THEN high\_cliquishness}$$

$$\text{Rule 3:}\quad \text{IF label\_agreement} \geq \theta_3 \text{ THEN high\_label\_consistency}$$

What if we make then trainable??

What if we make the rules trainable?

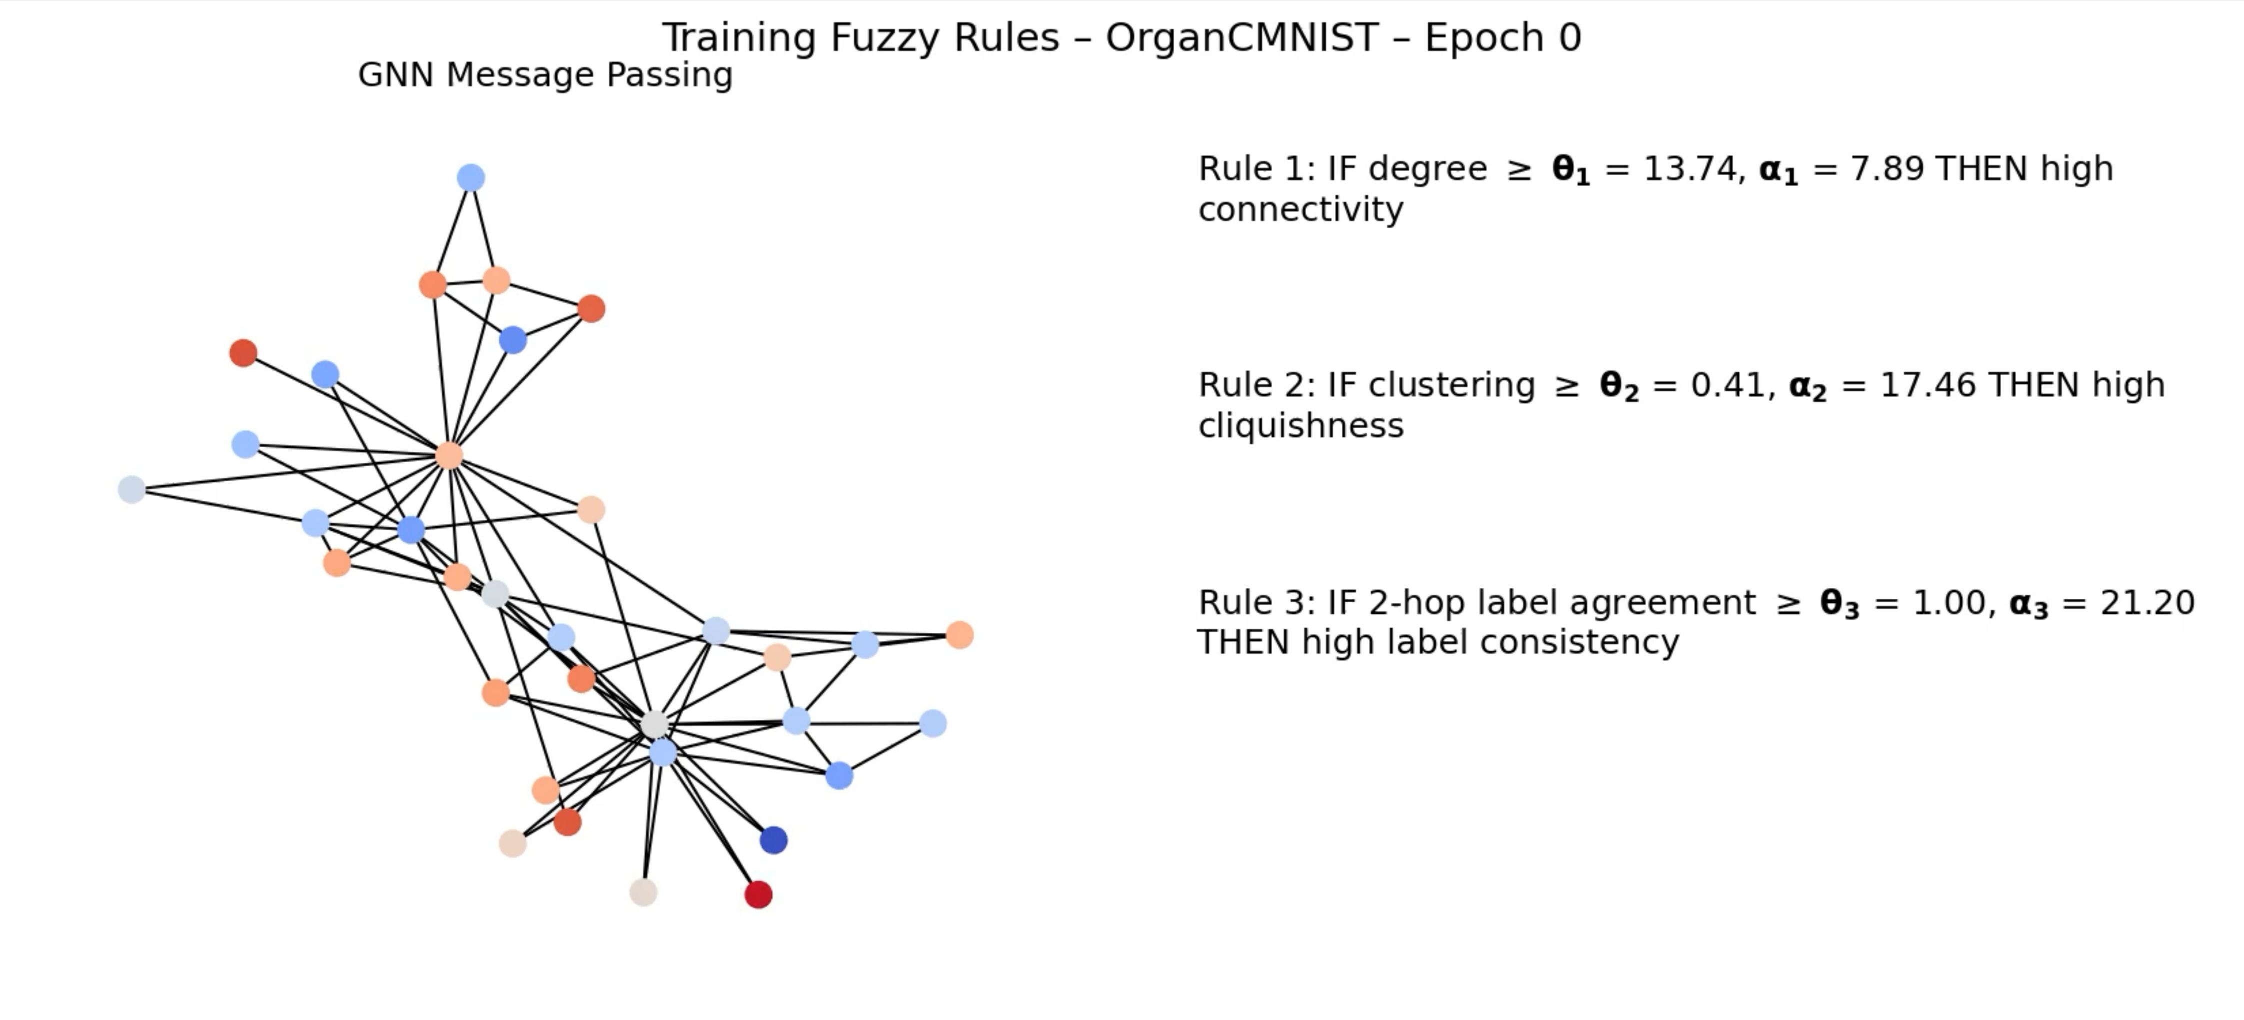

Step C & D - Learning Patterns and Logic

- **Fact Vector:** $$f_u = \left[d(u),\, C(u),\, L(u)\right]$$

Rule Activation Sharpness

Decision Boundary

- **Activation of Rule i:** $$r_i(u) = \sigma(\alpha_i(f_u[i] - \theta_i)) \in [0,1]$$

For instance, $r_1(u) \approx 1$ suggests *"node has high degree"*

- $r_i(u)$ is ***close to 1*** when the feature $f_u[u]$ significantly exceeds its threshold.

- **Firing strength vector**: $r(u) = [r_1(u), r_2(u), r_3(u)] \in [0,1]^3$

Node Degree | Clustering Coefficient | 2-Hop Label Agreement

Analysis of θ_1 Values During Training

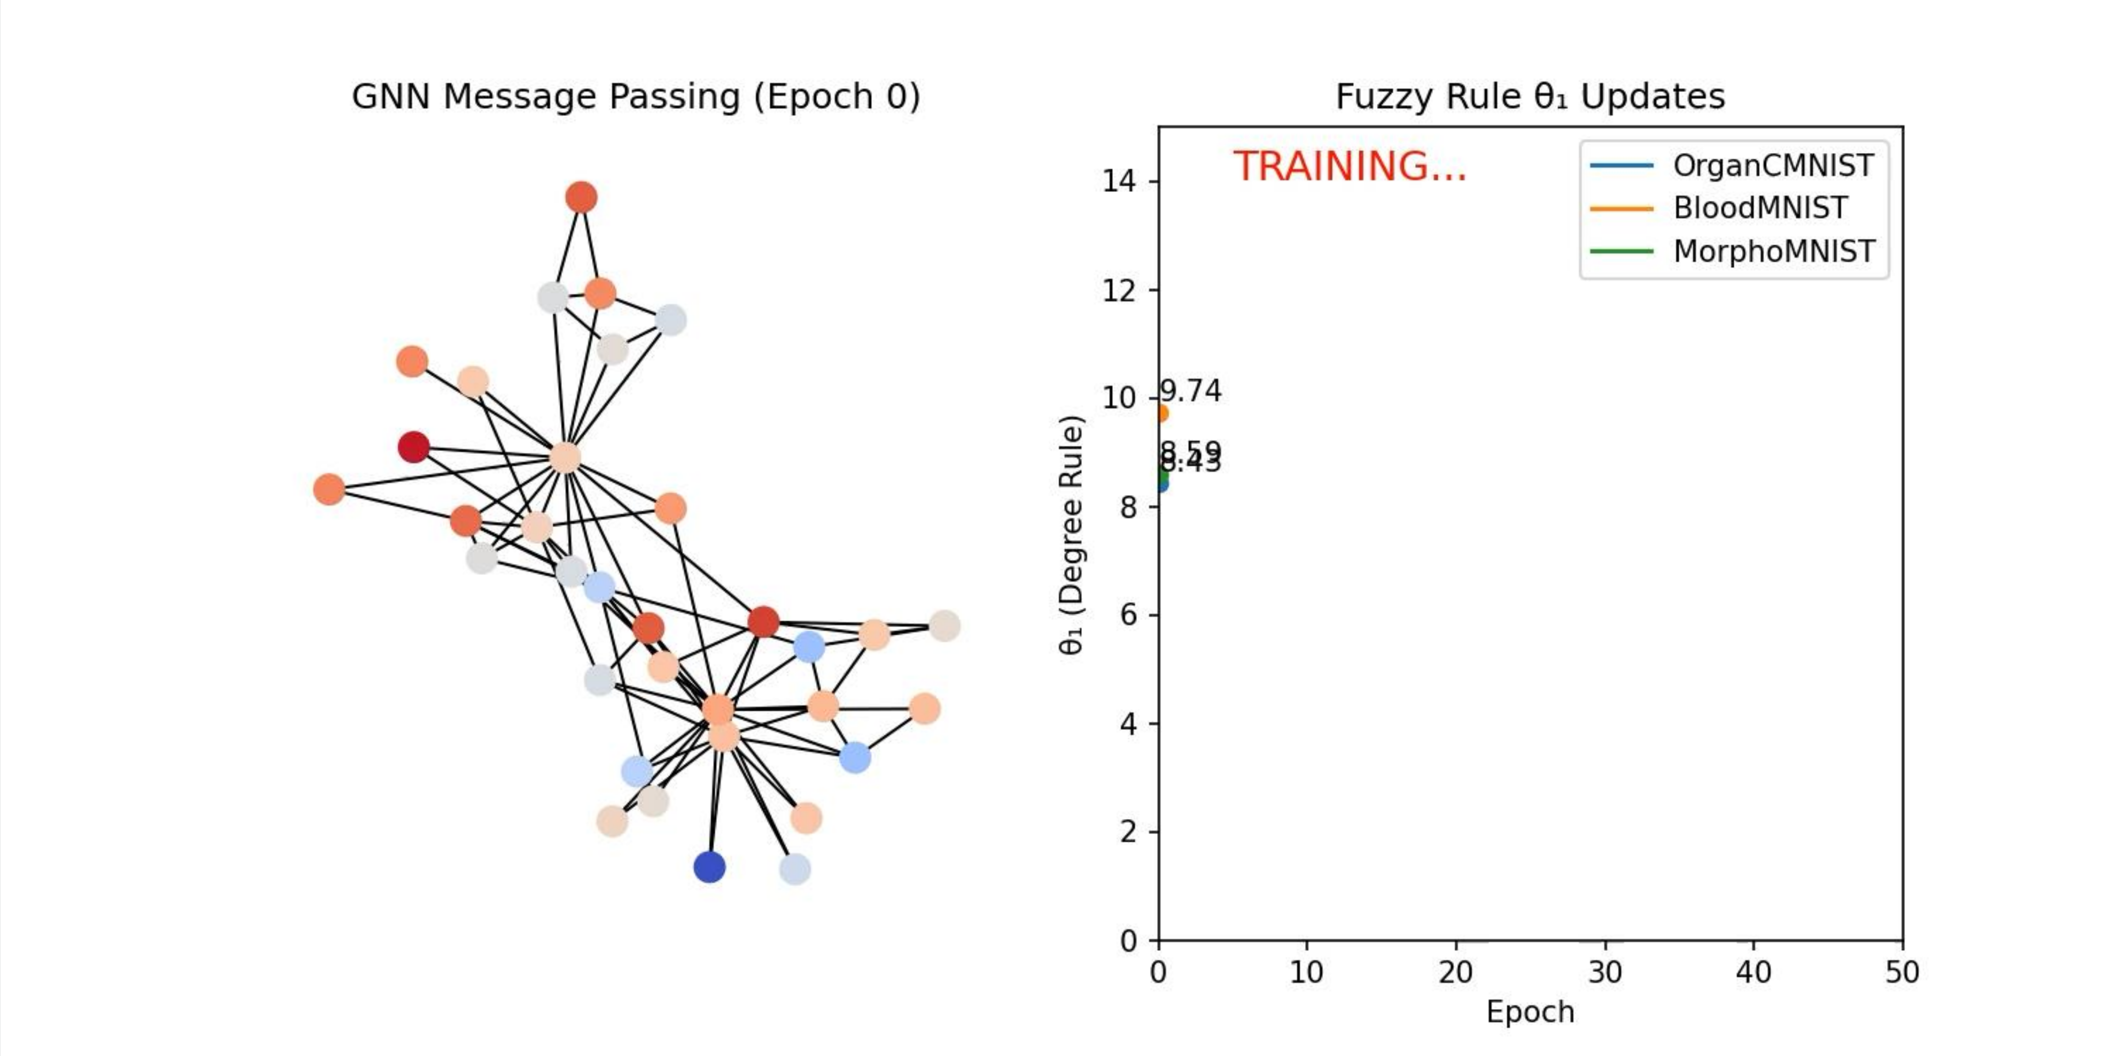

Fusing Logic and Patterns for the Final Prediction

- FireGNN intelligently combines the GNN's learned embedding ($\boldsymbol{h_u}$) with the fuzzy rules' firing strengths (**r**)
- A **gating mechanism** learns to weigh how much influence the symbolic rules should have on the final prediction for each node.

1. Project the rule vector into the GNN embedding space:

$$e_u = W_r\, r(u) + b_r, \quad e_u \in \mathbb{R}^d$$

2. Then, concatenate $[h_u \parallel e_u] \in R^{2d}$ and compute a gate vector:

$$g_u = \sigma(W_g\,[\,h_u \parallel e_u\,] + b_g) \in [0,1]^d$$

3. The updated node embedding is given by:

$$h'_u = g_u \odot h_u + (1 - g_u) \odot e_u$$

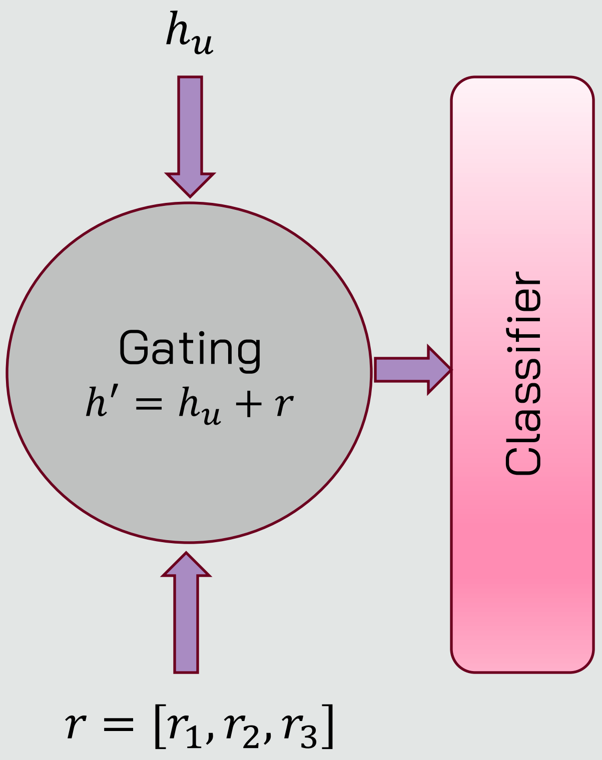

Intuitively, if rule-based information is highly relevant, g_u increases the influence of e_u; otherwise, it favors the original embedding h_u

FireGNN Explainability

We can now look at the calculated firing strengths $(r_1(u), r_2(u), r_3(u))$ for that specific prediction.

If $r_1(u)$ was high (e.g., 0.60) and $r_3(u)$ was high (e.g., 0.56), we can generate a human-readable explanation: *"The node was classified as 'spleen'* because the 'high_connectivity' and 'high_label_consistency' rules were strongly activated". This explanation is an intrinsic part of the model's computation, not a post-hoc guess.

BASIRA

Benchmarking Against Auxiliary Tasks

- To see if we could improve performance simply by forcing the GNN to learn about topology, without using explicit symbolic rules.

The two benchmark tasks:

1. **Homophily Prediction:** To teach the GNN about local label consistency
2. **Similarity Entropy Prediction:** To teach the GNN about the certainty of connections in a neighborhood.

Finding: Auxiliary tasks provide only **modest, incremental improvements** and sometimes don't improve performance at all.

Computational Efficiency

Method	OrganCMNIST		OrganAMNIST		OrganSMNIST		TissueMNIST		BloodMNIST		MoprhoMNIST	
	Time	Memory	Time	Memory	Time	Memory	Time	Memory	Time	Memory	Time	Memory
GCN	0.48	758.95	1.15	1234.64	0.49	1024.92	2.76	2472.47	0.35	557.36	4.20	2648.61
GCN + Aux	1.90	944.25	1.78	1439.06	1.25	889.30	4.76	1781.39	0.35	718.44	4.71	1717.73
GCN + FR	**3.12**	519.28	**12.50**	1197.50	3.47	592.38	**17.96**	1035.08	**1.92**	547.02	4.86	1628.05
GAT	0.48	1895.61	1.49	1698.78	0.61	1486.59	**5.95**	1950.41	0.47	1576.44	3.79	2130.48
GAT+Aux	2.33	1486.59	2.89	1516.97	2.31	1425.03	OOM	OOM	1.55	1191.38	OOM	OOM
GAT + FR	**44.50**	827.34	**19.75**	864.91	**16.13**	803.56	OOM	OOM	**3.85**	1196.73	OOM	OOM
GIN	0.32	913.77	1.77	1293.22	2.26	801.11	3.22	1255.64	0.56	649.13	**5.49**	1526.40
GIN+ Aux	0.68	1004.95	1.98	1258.81	**4.36**	1022.61	3.87	1368.44	1.36	893.95	4.88	1480.20
GIN + FR	1.10	1221.91	3.31	1384.21	3.96	729.97	4.96	734.70	1.89	656.72	**16.61**	1660.84

Notes: Top two highest epoch times of each dataset are bolded.

The fuzzy rule module is **memory-efficient**; GCN+FR consistently uses less peak memory than the baseline GCN

Adding Fuzzy Rules (FR) **increases epoch time** due to the extra logic computations, with GAT+FR showing the highest overhead.

Performance

3-4% ↑ across datasets

OOM on large graphs

Sensitivity ↑ 3-5%

F1 Score ↑ 2-3%

	OrganCMNIST				OrganAMNIST			
Method	ACC	F1	Sensitivity	ROC-AUC	ACC	F1	Sensitivity	ROC-AUC
GCN	88.20±0.61	86.03±0.89	86.13±0.06	99.09±0.08	91.85±0.30	91.19±0.33	91.18±0.03	99.51±0.03
GCN+ Aux	88.41±0.44	86.48±0.58	86.43±0.04	98.84±0.10	93.11±0.24	92.60±0.26	92.50±0.02	99.28±0.05
GCN+ FR	91.41±0.61	89.71±0.58	89.74±0.06	99.54±0.05	94.32±0.18	93.88±0.18	93.81±0.02	99.77±0.01
GAT	90.31±0.28	88.47±0.34	88.51±0.03	99.38±0.05	93.69±0.36	93.26±0.28	93.36±0.04	99.69±0.02
GAT+ Aux	90.88±0.49	89.07±0.59	88.93±0.62	99.45±0.08	93.70±0.46	93.36±0.44	93.33±0.04	99.54±0.02
GAT+ FR	91.66±0.48	90.02±0.52	90.12±0.05	99.56±0.05	94.52±0.31	94.08±0.30	94.13±0.03	99.78±0.02
GIN	87.96±0.59	85.61±0.75	85.56±0.75	98.86±0.09	91.54±0.71	90.45±0.73	90.50±0.69	99.41±0.08
GIN+ Aux	88.53±0.44	86.37±0.48	86.37±0.48	98.65±0.13	92.18±0.33	91.16±0.39	91.13±0.39	99.26±0.05
GIN+ FR	89.12±1.18	86.81±1.51	86.82±1.56	99.38±0.16	92.48±1.82	91.32±2.69	91.27±2.71	99.59±0.33

	OrganSMNIST				TissueMNIST			
Method	ACC	F1	Sensitivity	ROC-AUC	ACC	F1	Sensitivity	ROC-AUC
GCN	78.62±0.82	73.74±0.99	73.85±0.08	97.80±0.11	50.90±0.32	32.61±0.79	32.51±0.07	81.98±0.31
GCN+ Aux	79.19±0.74	74.21±0.84	74.28±0.07	97.56±0.14	52.70±0.22	35.67±0.36	35.60±0.04	82.99±0.19
GCN+ FR	85.05±0.43	80.56±0.61	80.74±0.04	98.95±0.05	65.73±0.88	51.63±1.22	51.59±0.09	93.56±0.36
GAT	81.80±0.68	77.22±0.73	77.16±0.07	98.39±0.09	51.53±0.35	33.10±0.56	33.06±0.07	83.11±0.12
GAT+Aux	81.69±0.68	77.33±0.73	77.06±0.07	98.28±0.09	OOM	OOM	OOM	OOM
GAT+FR	84.82±0.52	80.23±0.80	80.43±0.05	98.93±0.07	OOM	OOM	OOM	OOM
GIN	77.23±0.62	71.65±0.65	71.77±0.76	97.36±0.10	50.51±1.09	30.31±3.70	32.50±2.05	81.72±0.85
GIN+Aux	79.26±1.02	73.73±1.97	74.29±1.23	97.44±0.11	51.31±1.42	31.12±3.54	33.20±1.70	82.42±0.27
GIN+ FR	81.46±3.14	76.69±3.85	76.81±3.84	98.51±0.52	64.07±2.44	48.59±3.93	48.59±3.96	92.84±1.25

	BloodMNIST				MorphoMNIST			
Method	ACC	F1	Sensitivity	ROC-AUC	ACC	F1	Sensitivity	ROC-AUC
GCN	80.49±0.66	77.46±0.96	77.15±0.09	96.56±0.18	90.89±0.05	90.82±0.02	90.73±0.02	98.88±0.03
GCN+ Aux	80.17±0.66	77.07±0.84	76.96±0.10	96.05±0.13	92.84±0.10	92.83±0.10	92.78±0.03	99.00±0.04
GCN+ FR	88.31±0.37	86.36±0.45	86.22±0.04	99.15±0.05	94.76±1.02	94.76±1.03	94.58±0.34	99.63±0.25
GAT	82.02±0.31	79.38±0.40	79.18±0.04	97.45±0.08	91.50±0.15	91.48±0.16	91.33±0.05	98.73±0.04
GAT+ Aux	81.87±2.08	79.70±2.21	79.73±0.23	97.13±0.46	OOM	OOM	OOM	OOM
GAT+ FR	87.79±2.08	85.86±2.21	85.80±0.23	99.01±0.46	OOM	OOM	OOM	OOM
GIN	80.30±0.60	77.21±0.91	76.44±0.95	96.58±0.21	91.60±0.19	91.59±0.18	91.60±0.19	98.75±0.05
GIN+ Aux	80.47±0.20	77.13±1.31	76.83±0.38	96.30±0.42	92.30±0.24	92.11±0.32	92.50±0.43	99.25±0.12
GIN+ FR	84.83±3.04	82.41±3.60	82.39±3.81	98.36±0.94	93.72±1.55	93.73±1.54	93.72±1.55	99.55±0.22

Conclusion of FireGNN

Strengths of our model:

Data-Adaptive Reasoning

Learns fuzzy thresholds & sharpness directly from graph structure.

Topology-Aware Learning

Captures structural cues like degree, clustering, and label consistency.

Potential Research Directions:

Expanding Rule Vocabulary

Incorporate global and community-level graph metrics.

Dynamic Graph Extension:

Apply FireGNN to evolving biomedical networks.

"The requisites of knowledge: a quick mind, zeal for learning, humility, foreign land, a professor's inspiration, and a life of long span."
Juwaini of Nishapur (d.1085)

BASIRA

Brain **A**nd **SI**gnal **R**esearch & **A**nalysis Lab

https://basira-lab.com/

Imperial College London

Subscribe
Like
Share
&
\cite{}