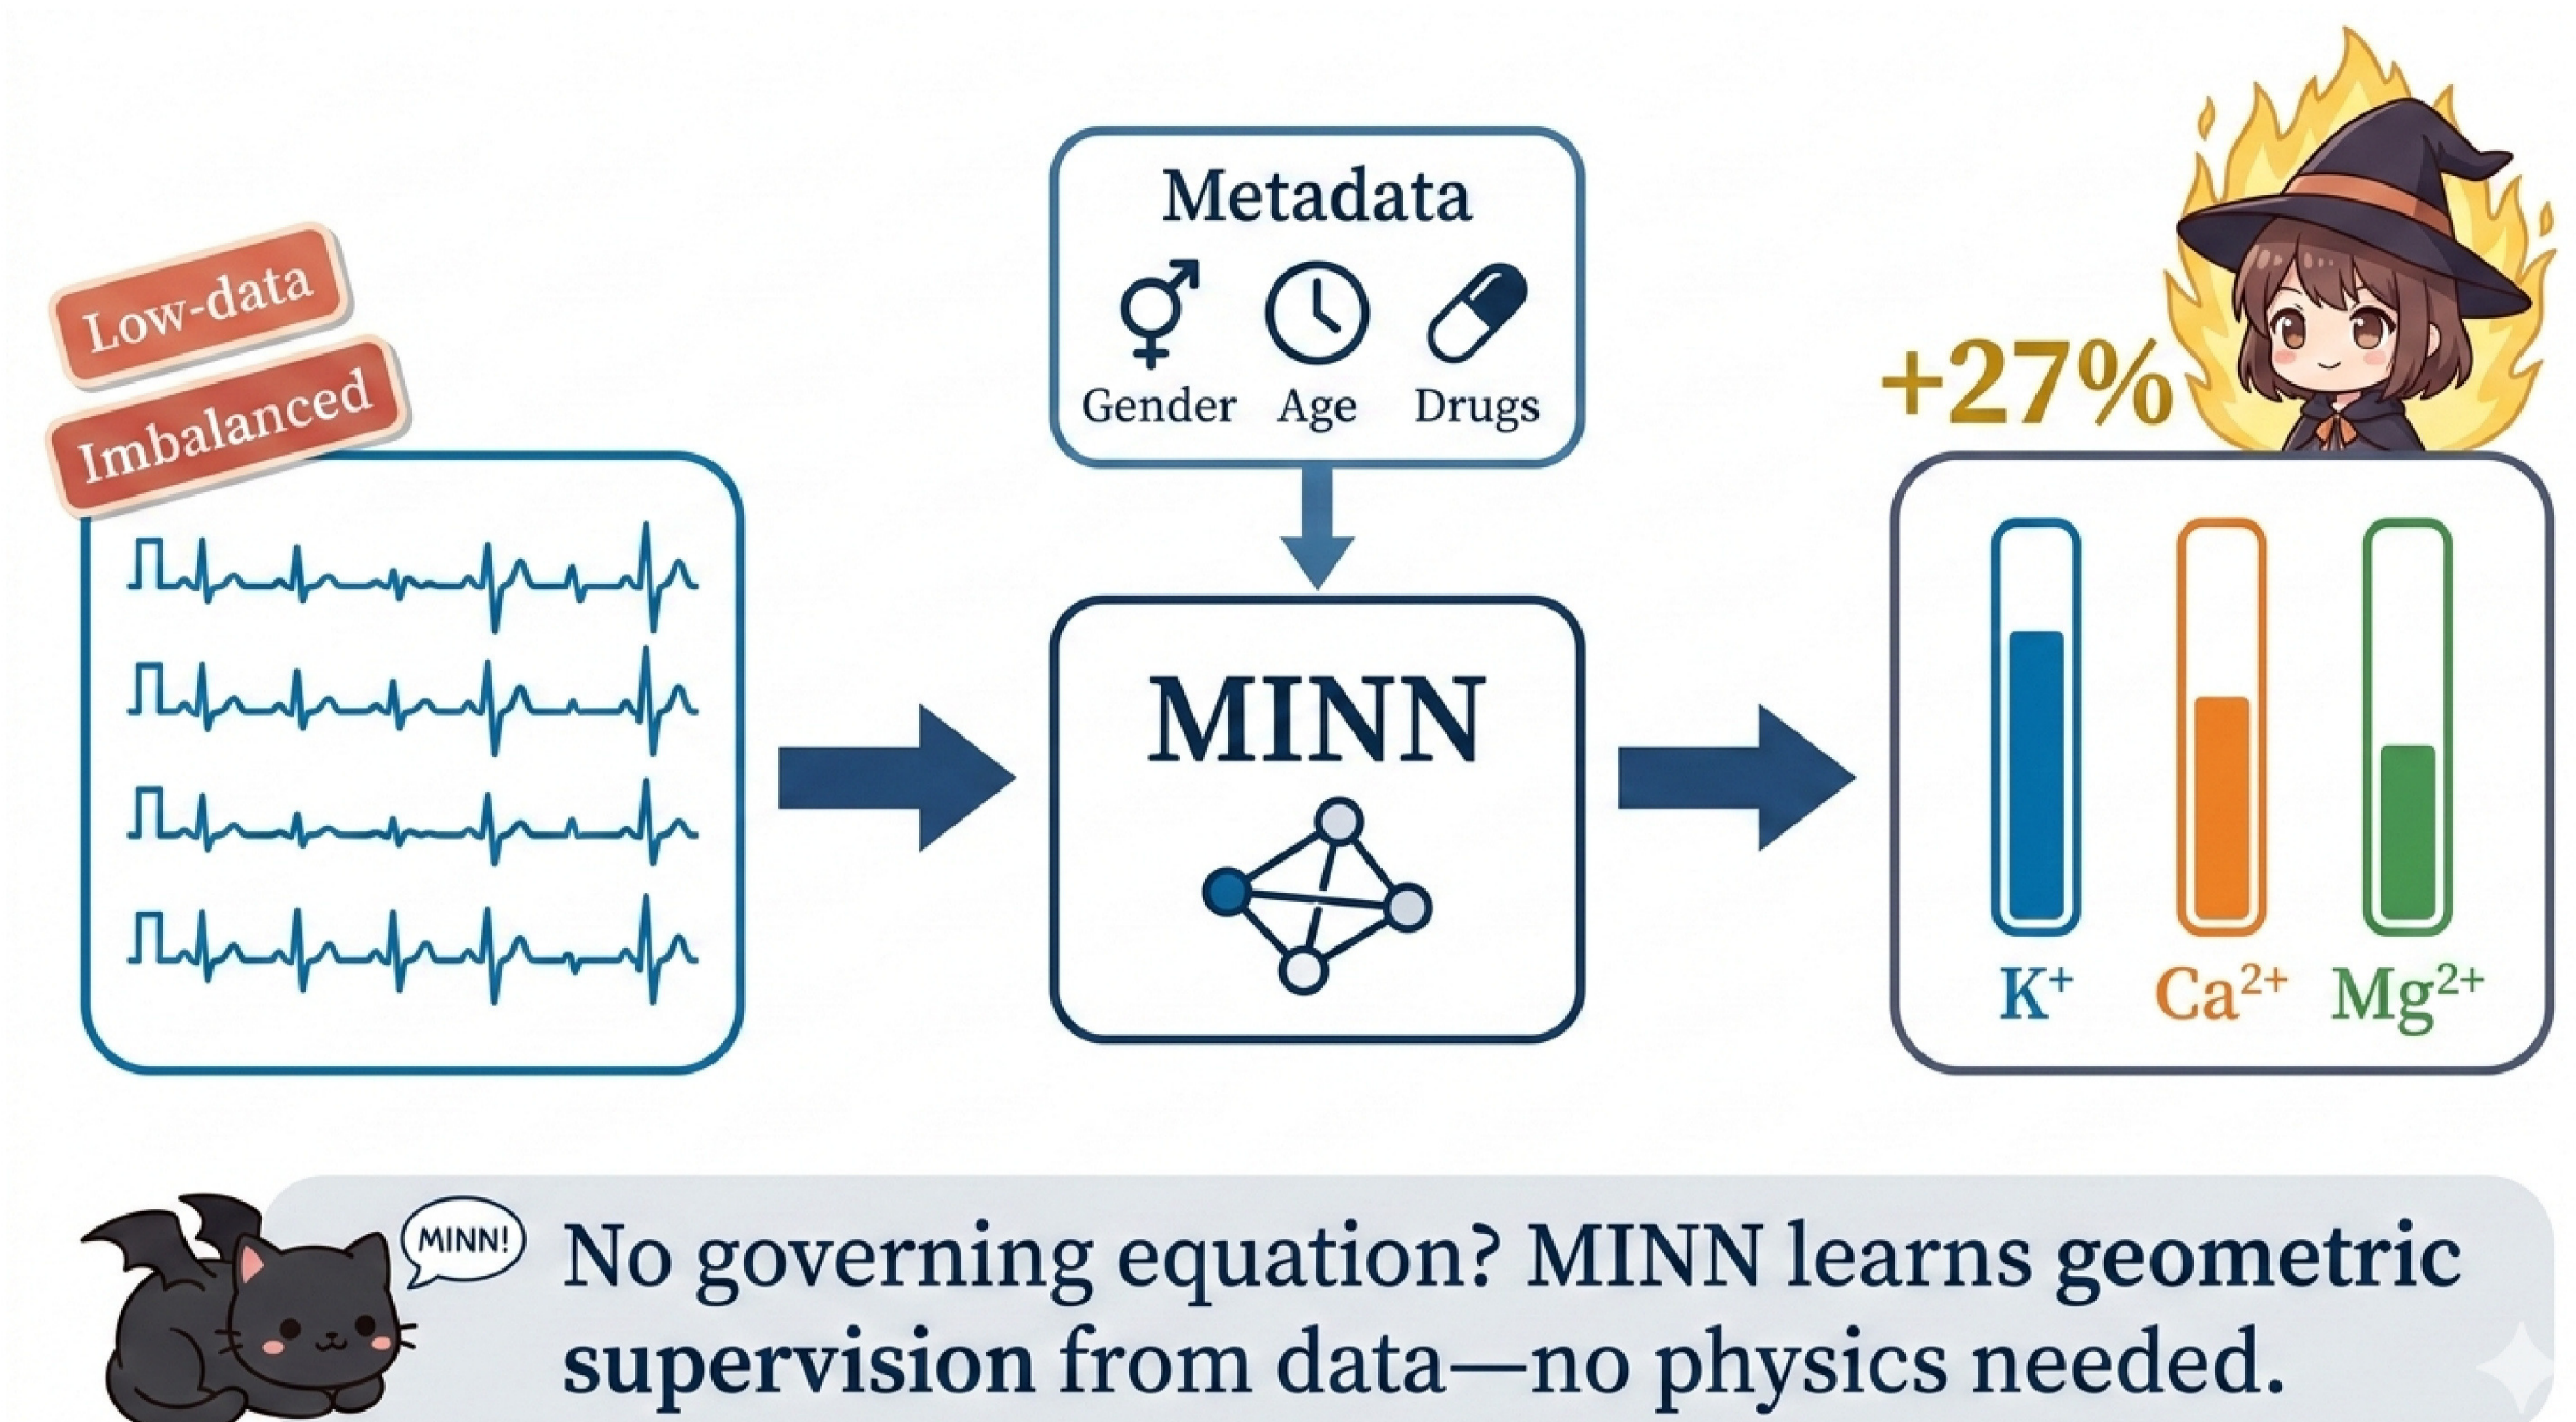




Why It Matters

- Electrolyte imbalances → fatal cases; diagnosis requires blood draws
- ECGs encode electrolyte effects non-invasively, but labeled data is **scarce and imbalanced**
- These electrolyte-driven ECG morphologies live on a **manifold**
- MINN exploits this geometry as a **supervision**—fewer labels needed



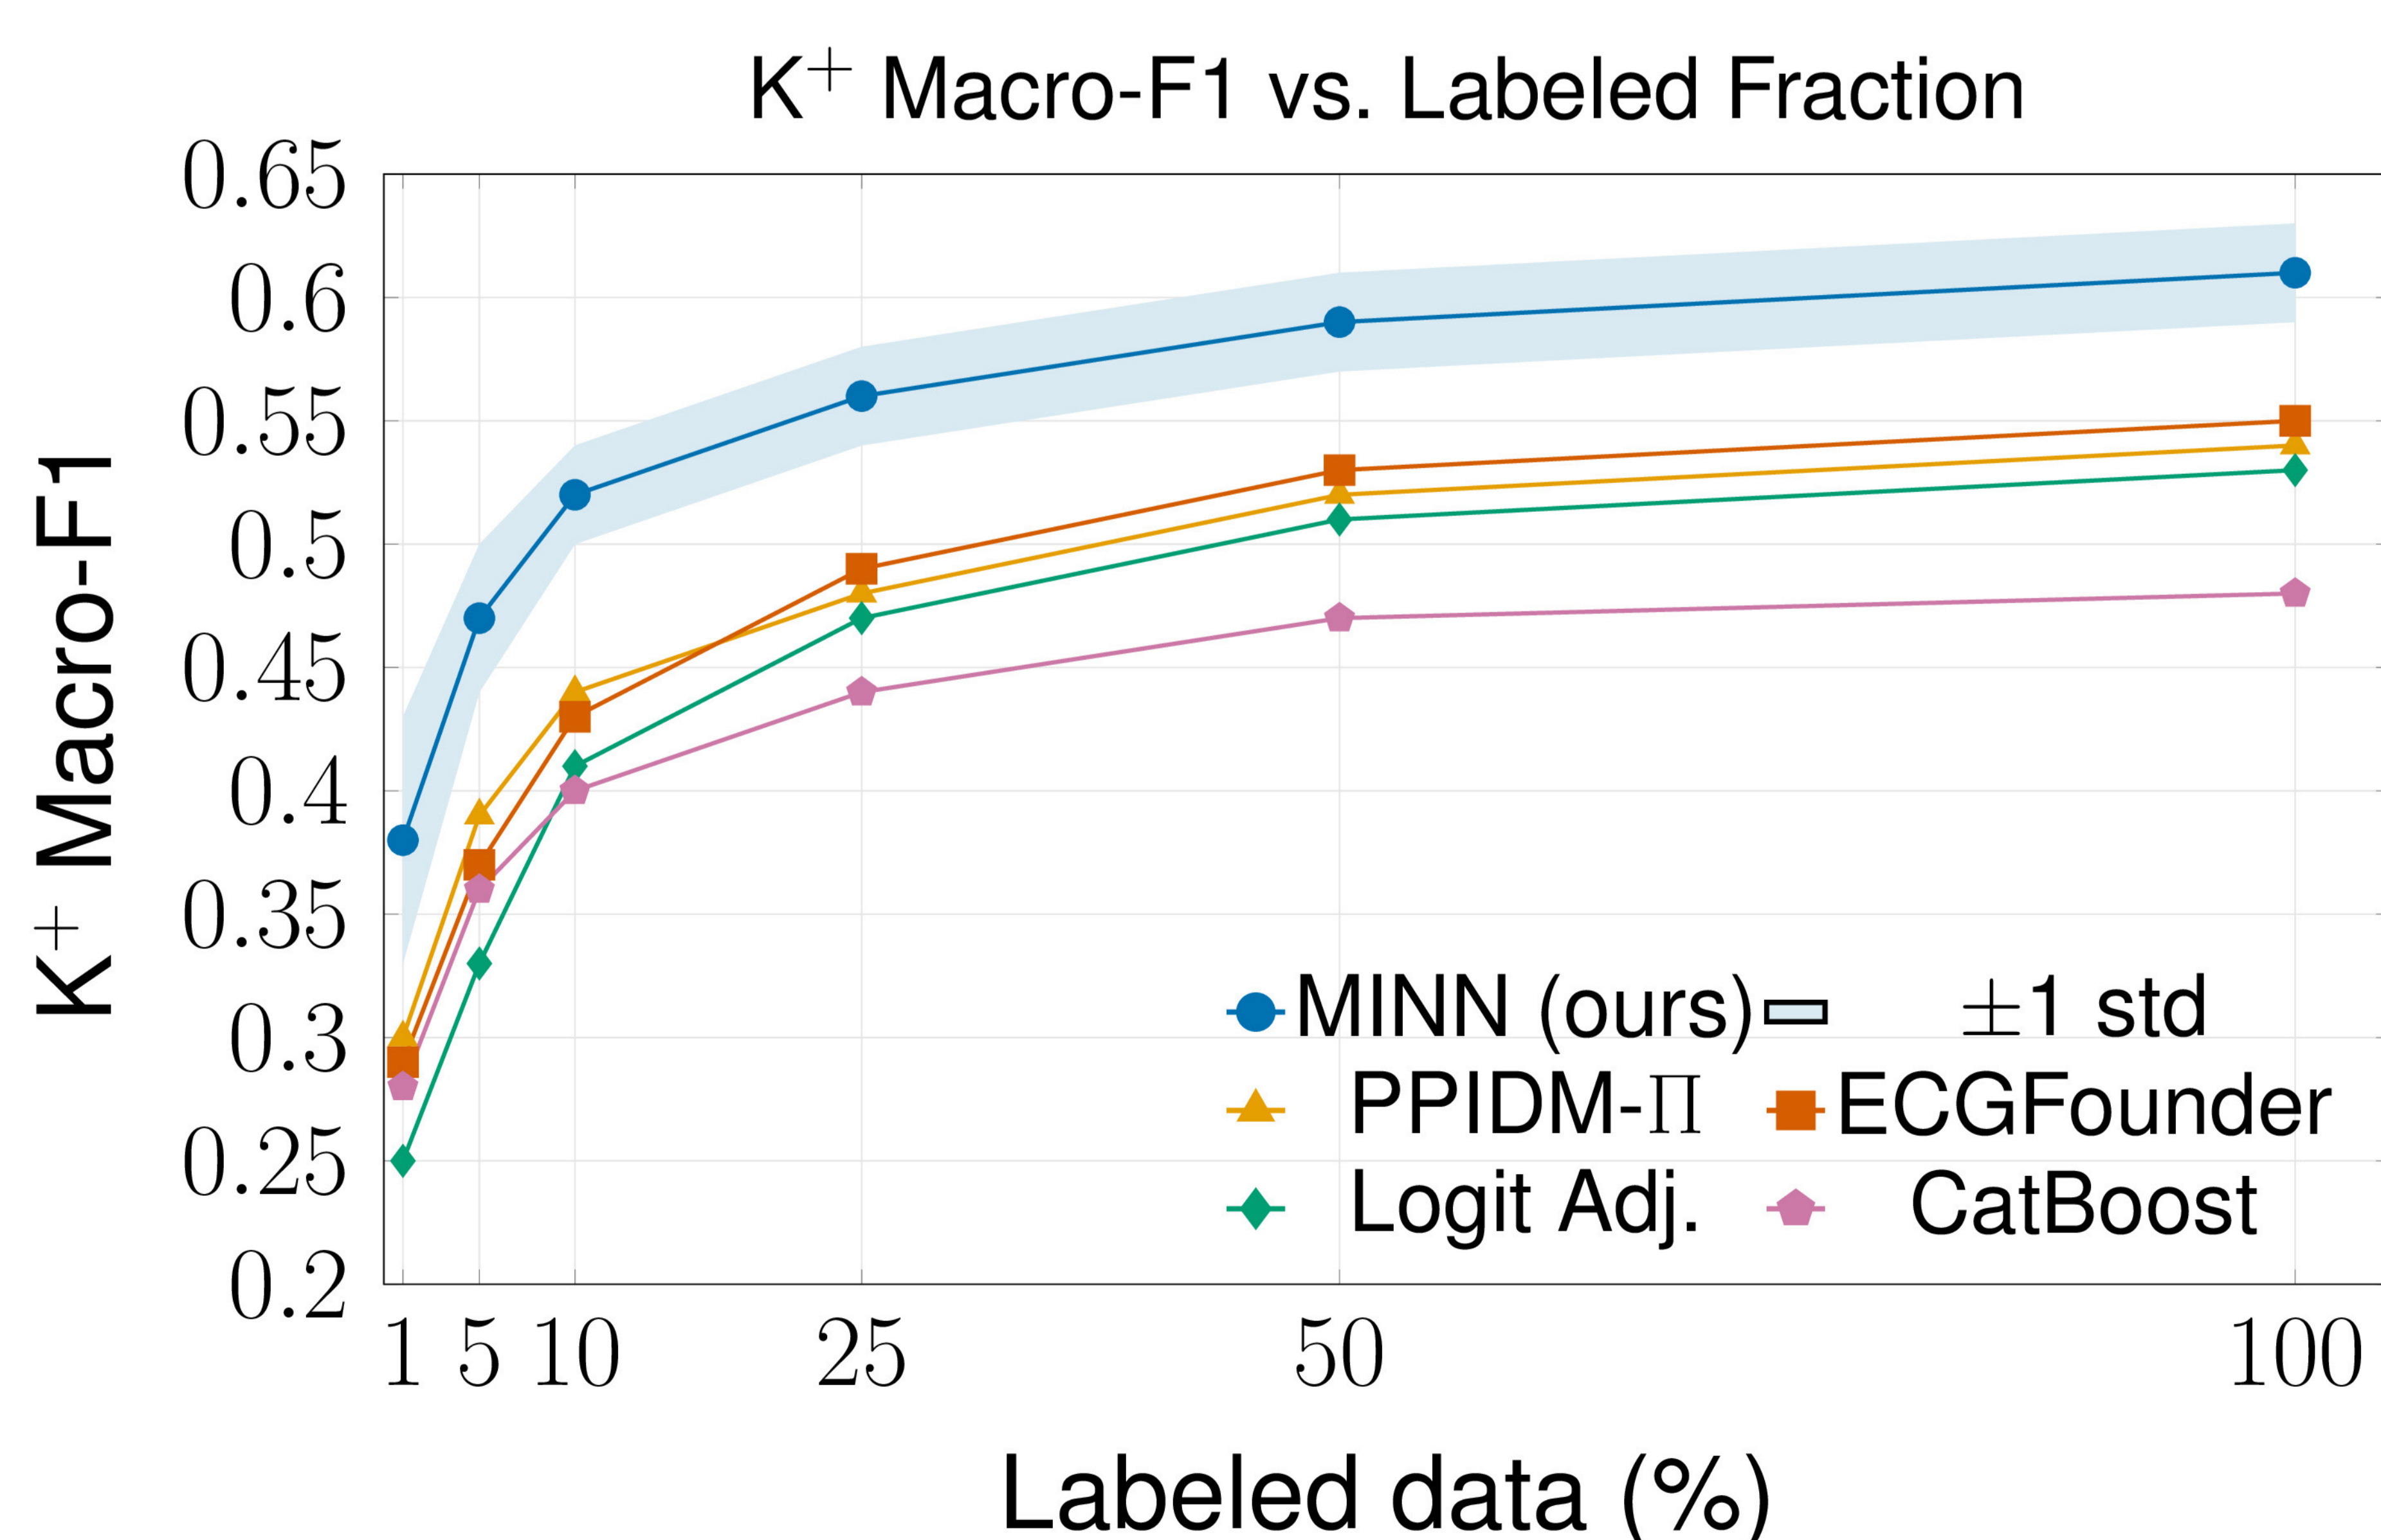
Results

MIMIC-IV (public ICU database). 800K ECGs ± 1 hr lab pairing
~10K matched recordings. 3-class (hypo/normal/hyper), imbalanced up to 1:10. Mean \pm std over 5 seeds.

Method	K ⁺		Ca ²⁺	
	Macro-F1	BAcc	Macro-F1	BAcc
<i>Tabular</i>				
TabPFN	0.47 \pm .02	0.54 \pm .02	0.43 \pm .01	0.50 \pm .02
CatBoost	0.48 \pm .02	0.55 \pm .02	0.44 \pm .01	0.51 \pm .02
<i>ECG Foundation Models</i>				
ECGFounder	0.55 \pm .01	0.62 \pm .01	0.44 \pm .03	0.51 \pm .02
ECG-FM	0.52 \pm .03	0.58 \pm .02	0.42 \pm .03	0.49 \pm .03
von Bachmann	0.50 \pm .02	0.57 \pm .02	0.40 \pm .02	0.47 \pm .02
<i>Manifold & Imbalanced Learning</i>				
Logit Adj.	0.53 \pm .02	0.60 \pm .02	0.43 \pm .02	0.50 \pm .02
Manifold Mixup	0.51 \pm .03	0.58 \pm .02	0.41 \pm .03	0.48 \pm .02
Poisson Learn.	0.49 \pm .04	0.56 \pm .04	0.40 \pm .03	0.46 \pm .03
<i>Physics-Informed</i>				
PPIDM-II	0.54 \pm .03	0.61 \pm .02	0.43 \pm .03	0.50 \pm .02
MINN (Ours)	0.61\pm.02	0.67\pm.02	0.50\pm.02	0.56\pm.02

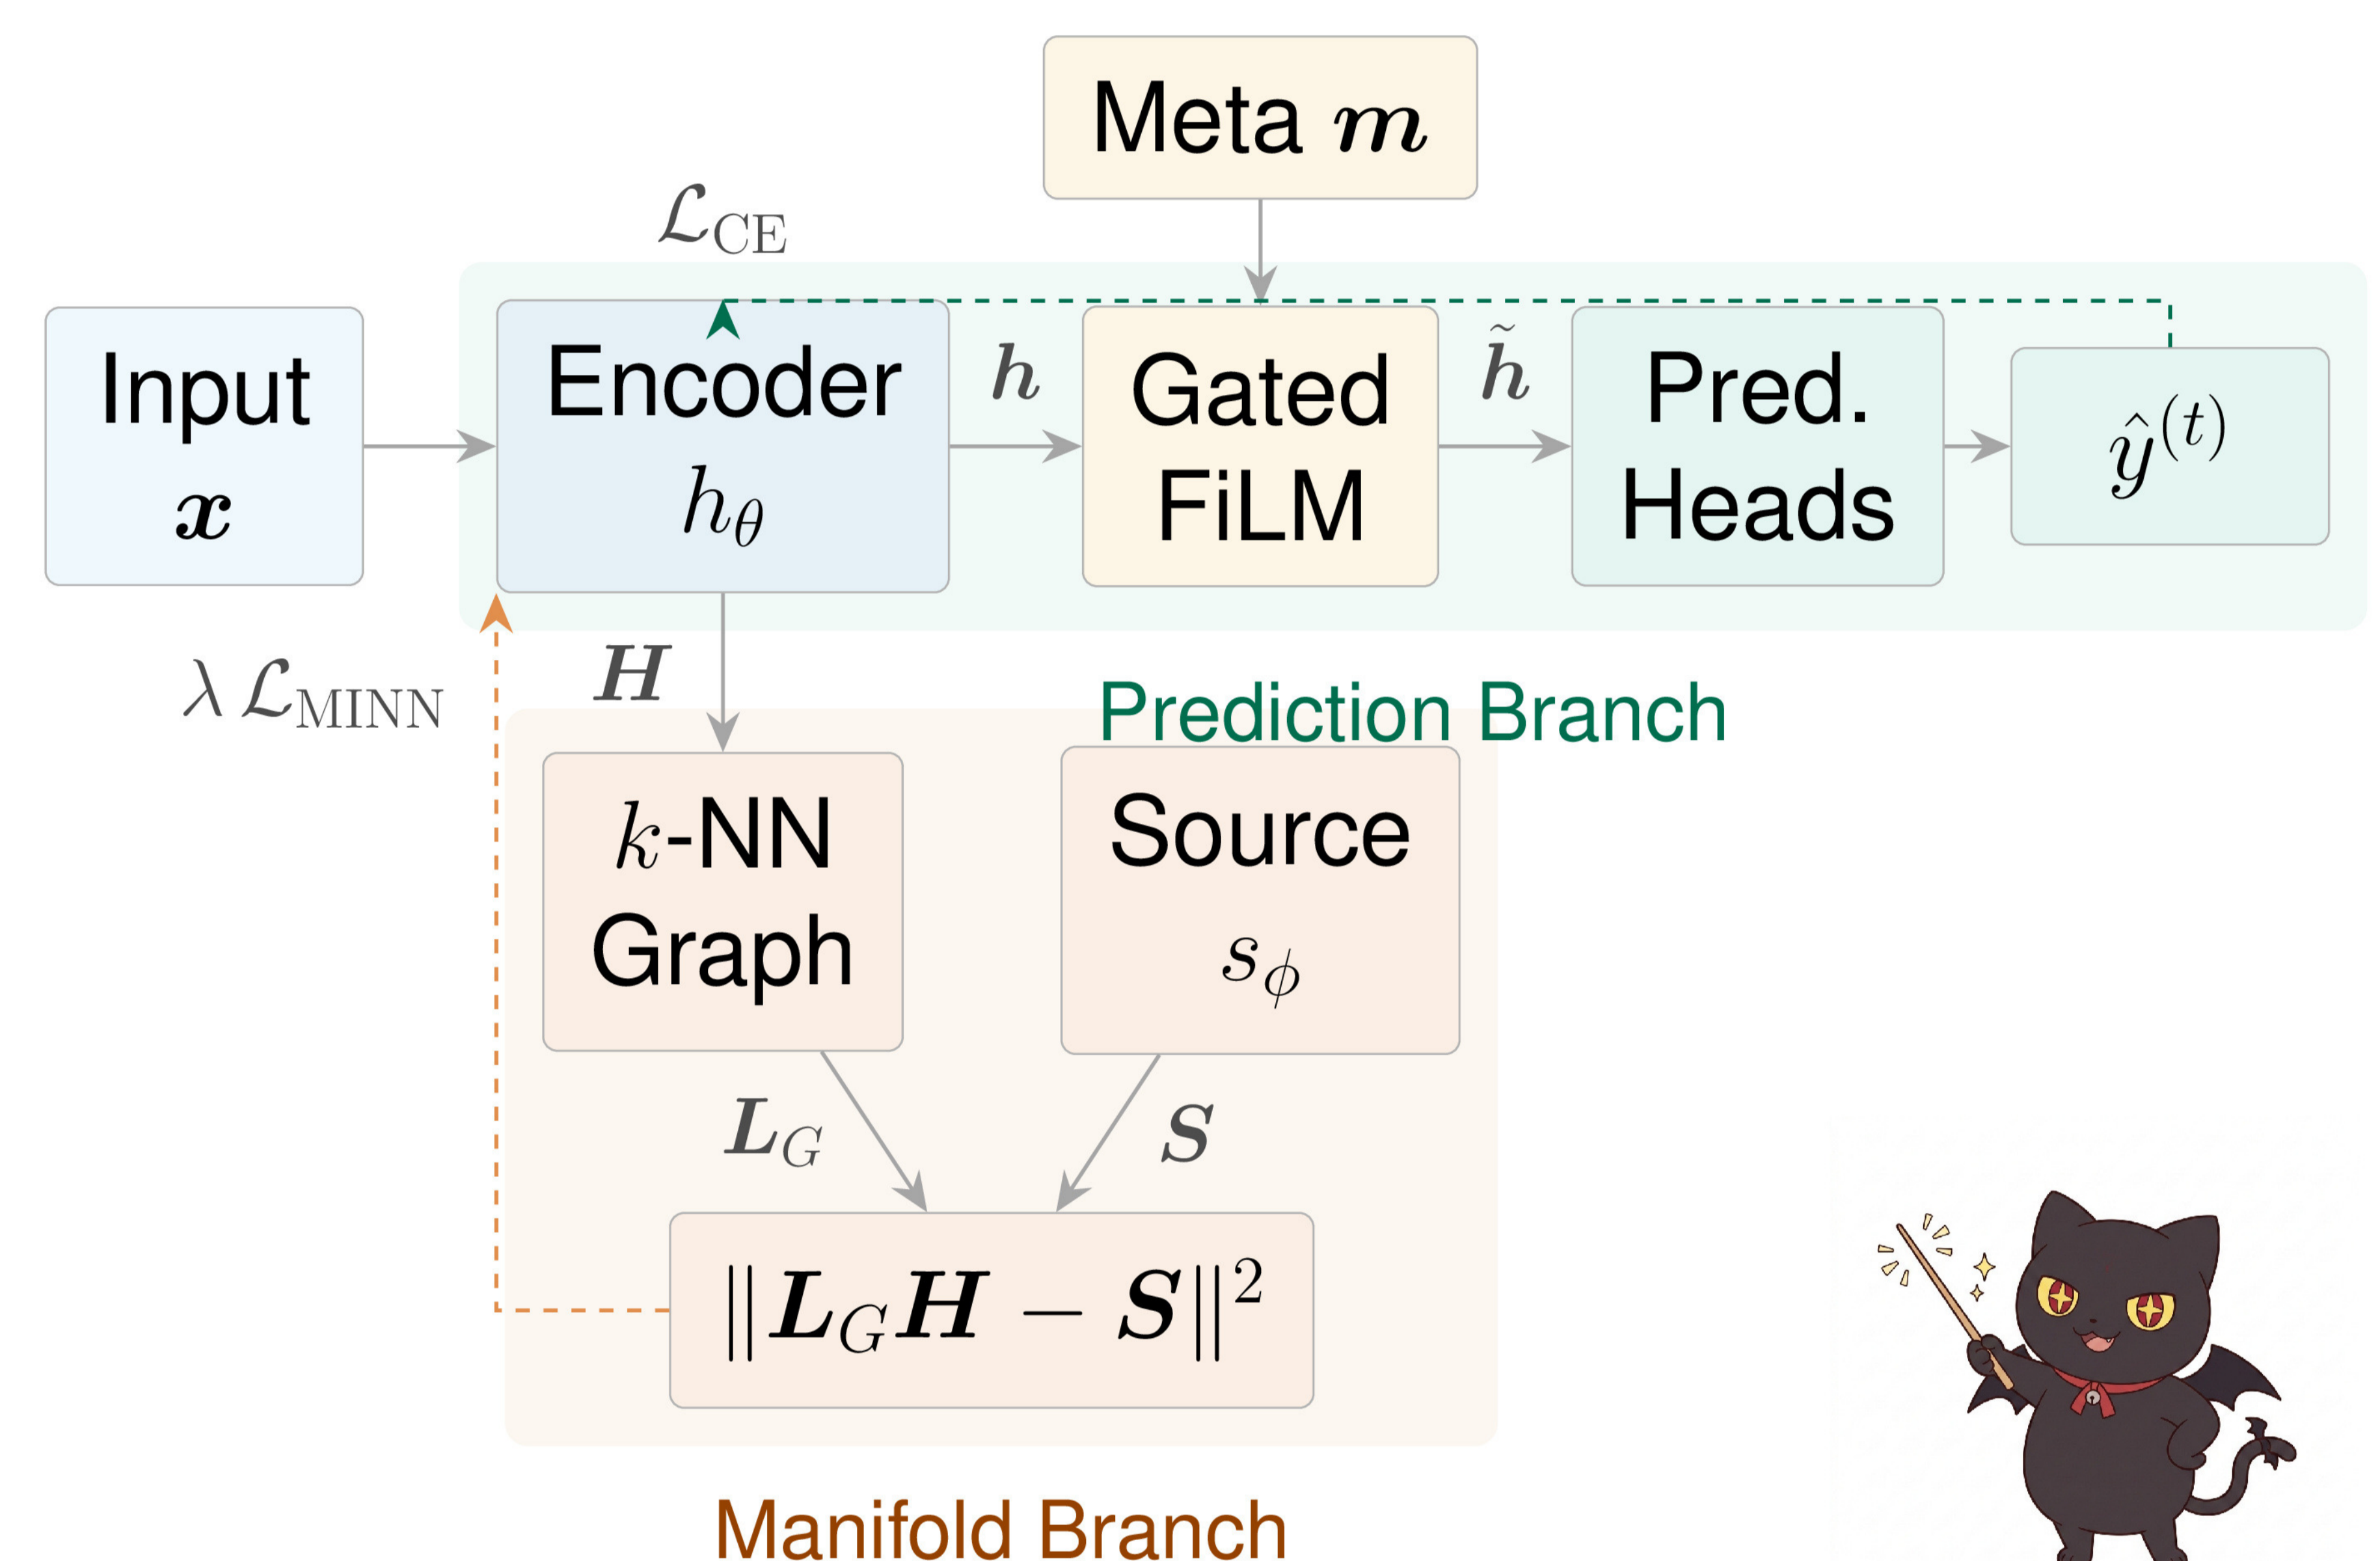
Best **Macro-F1** on all three electrolytes (Mg²⁺ omitted; 0.40 vs. 0.37). **+0.06** K⁺ F1 over next-best ECGFounder—geometric > pretraining scale.

Low-Data Performance



27% relative improvement at 1% labeled data.

How It Works



Prediction branch (green): an encoder maps the ECG signal x to a latent embedding h , which is modulated by clinical metadata m (age, sex, medications) via gated FiLM to produce \tilde{h} , then classified by per-electrolyte heads.

Manifold branch (red): a k -NN graph is built in latent space, yielding graph Laplacian L_G . A learned source network s_ϕ produces a target for each sample. The **manifold residual** penalizes the mismatch:

$$\mathcal{L}_{\text{MINN}} = \frac{1}{N} \sum_{i=1}^N \|(L_G \mathbf{H})_i - s_\phi(\mathbf{x}_i, \mathbf{m}_i)\|^2$$

This enforces a discrete Poisson equation on the manifold—providing geometric supervision at every data point, not just labeled ones.

Total objective: $\mathcal{L}_{\text{total}} = \sum_t \mathcal{L}_{\text{CE}}^{(t)} + \lambda \mathcal{L}_{\text{MINN}}$

Conclusions

- Best Macro-F1 on all three electrolytes; advantage grows as labels decrease
- Poisson residual + learned source > classical manifold smoothing (+0.06 K⁺ F1)
- No known physics needed at all—generalizes to any label-scarce domain with manifold structure

Acknowledgments Prof. Shixin Xu (mentor); DKU; my family.