Inverse ECG Reconstruction Using Kolmogorov–Arnold Networks

Yunus Emre Tüysüz[1,2], Gözde Bozdağı Akar[1], Yeşim Serinağaoğlu Doğrusöz[1]

[1]Electrical-Electronics Engineering Department, Middle East Technical University, Ankara, Turkey
[2]ASELSAN A.Ş., Ankara, Turkey

Abstract

Electrocardiographic imaging aims to reconstruct epicardial electrical activity from body-surface potential measurements, but the inverse ECG problem is highly ill-posed and becomes even more challenging under realistic clinical conditions such as missing electrodes and measurement noise. Although recent deep learning methods have shown promising performance, many of them degrade substantially when the input data are corrupted by lead loss or noise. In this study, we investigate the use of a Kolmogorov–Arnold Network (KAN)-based model for robust inverse ECG reconstruction. We evaluate the proposed approach on paced body-surface and epicardial recordings and benchmark it against prominent neural network methods, including a multilayer perceptron and a residual network. Our experiments show that while conventional deep models can achieve very high reconstruction accuracy under clean conditions, their performance deteriorates more rapidly under severe lead dropout. In contrast, the KAN-based model preserves higher temporal and spatial reconstruction fidelity in missing-lead scenarios, indicating improved robustness to incomplete measurements.

1. Introduction

Electrocardiographic imaging (ECGI) aims to reconstruct cardiac electrical activity from body-surface potential (BSP) measurements [1]. This is a severely ill-posed inverse problem, since the torso acts as a spatial low-pass filter and different cardiac source distributions can produce very similar body-surface signals. Classical ECGI methods address this instability through regularization, but they often oversmooth the solution and depend strongly on an accurate forward model. In recent years, deep learning methods have become increasingly popular in inverse ECGI because they can learn the nonlinear BSP-to-epicardial mapping directly from data and reduce the computational burden of iterative inverse solvers [2, 3]. At the same time, their reliability under realistic acquisition degradations remains an open challenge.

Kolmogorov–Arnold Networks (KANs) provide a structured alternative to standard neural regressors by representing multivariate mappings as compositions of learnable one-dimensional nonlinear functions combined through linear mixing [4]. In practical implementations such as FastKAN, these scalar nonlinearities are parameterized via localized basis expansions followed by lightweight linear transformations, enabling efficient and flexible function approximation. This structure is particularly appealing for ECGI, where the underlying mapping is high-dimensional yet governed by smooth and spatially correlated waveform transformations across sensors and epicardial nodes. This motivates studying whether KAN-based models can provide not only strong reconstruction accuracy, but also improved robustness under degraded measurement conditions.

In this work, we compare a FastKAN-based inverse ECG model with two strong baselines, a multilayer perceptron (MLP) and a residual network (ResNet), on the Bordeaux Exp16 vest–sock torso-tank dataset with LV and RV pacing recordings obtained from the EDGAR data repository [5]. Performance is assessed using the pooled temporal correlation coefficient (TCC), spatial correlation coefficient (SCC), and activation-time correlation coefficient (ATCC), and localization error (LE) reported as median (IQR) across beats.

2. Methods

2.1. Inverse ECG formulation

After discretization ECGI forward problem can be written as

$$\mathbf{Y} = \mathbf{A}\mathbf{X} + \mathbf{N}, \tag{1}$$

where $\mathbf{Y} \in \mathbf{R}^{\mathbf{m}\times\mathbf{T}}$ denotes BSP measurements, $\mathbf{X} \in \mathbf{R}^{\mathbf{n}\times\mathbf{T}}$ denotes epicardial potentials, $\mathbf{A}$ is the forward transfer operator, and $\mathbf{N}$ represents measurement noise and modeling error. Instead of explicitly inverting $\mathbf{A}$, we learn a direct nonlinear inverse mapping from BSPs to epicardial potentials.

For each beat b and time index t, the input is defined as

$$\mathbf{x_{b,t}} = [\mathbf{y_{b,t}};\, \tau_\mathbf{t}], \tag{2}$$

where $\mathbf{y_{b,t}}$ is the BSP vector and $\tau_t \in [0, 1]$ represents the normalized temporal position within the beat, providing the model with relative timing information in addition to the BSP measurements. The model predicts the corresponding epicardial potential vector

$$\hat{\mathbf{x}}_{b,t} = F_\theta(\mathbf{x_{b,t}}). \tag{3}$$

2.2. Models

Our main model is FastKAN [6], a Kolmogorov–Arnold Network (KAN) variant. In FastKAN, each scalar input is expanded through Gaussian radial basis functions (RBFs),

$$\phi_g(x) = \exp\left(-\frac{(x-\mu_g)^2}{\sigma_g^2}\right), \qquad g = 1, \ldots, G, \tag{4}$$

and the expanded responses are linearly mixed. A FastKAN layer combines this RBF branch with a residual base branch:

$$\mathbf{h}^{(\ell+1)} = \mathbf{W}_{\text{rbf}}^{(\ell)} \mathbf{\Phi}(\mathbf{h}^{(\ell)}) + \mathbf{W}_{\text{base}}^{(\ell)} \sigma(\mathbf{h}^{(\ell)}), \tag{5}$$

where $\Phi(\cdot)$ denotes the concatenated scalar basis expansion and $\sigma(\cdot)$ is GELU as an activation function. Layer normalization is used before the basis expansion, and residual connections are retained where dimensions match.

We compared FastKAN with two neural baselines: a fully connected MLP and a ResNet. The MLP used three hidden layers with LayerNorm, GELU, and dropout. The ResNet used an input projection, residual convolutional blocks, and a final linear output layer. All three models were matched to a similar parameter budget (3.6–3.8 million). Classical model-based ECGI baselines were not included in the main comparison, since our goal was to evaluate FastKAN against operator-free learning-based methods under a fair and consistent setting. Accordingly, we restricted the comparison to data-driven baselines that, like FastKAN, do not require an explicit forward transfer matrix during reconstruction.

2.3. Beat extraction and preprocessing

In this study, we used the Bordeaux Exp16 vest–sock torso-tank dataset from the EDGAR repository [5]. This dataset provides torso-tank body-surface potential recordings and unipolar ventricular epicardial sock recordings sampled at 2048 Hz, with experiment-specific geometries reconstructed from tank MRI data. Although the dataset also includes sinus-rhythm recordings following LBBB induction by ablation, the present study used only the paced LV and RV recordings. After excluding invalid channels, 124 BSP leads and 103 epicardial nodes were retained. Beat boundaries were identified by cross-correlation with the provided single-beat template. Both BSP and epicardial signals were preprocessed by temporal centering, implemented as subtraction of the per-channel mean over time, and spatial centering, implemented as subtraction of the across-node mean at each frame. Pacing stimulus artifacts were removed by linear interpolation, and the QRS interval was identified from a derivative-energy envelope of the epicardial signal. Only QRS samples were retained for training and evaluation.

2.4. Training protocol

The dataset was divided into disjoint training, validation, and test sets at the beat level, with 13952, 1551, and 12035 data samples, respectively, for a total of 128 beats (70 RV and 58 LV). BSP inputs were z-score normalized, whereas epicardial targets were scaled separately; the normalized temporal position was appended to each frame. All models were trained using AdamW (learning rate 10^{-4}, weight decay 5×10^{-4}) together with cosine-annealing warm restarts, and the checkpoint with the lowest validation loss was retained. To improve robustness, BSP inputs were augmented during training using random channel masking (maximum ratio 0.30) and additive Gaussian noise (standard deviation 0.05). The training objective combined a correlation-based term with a gradient-aware reconstruction term. The latter combined mean-squared error with first- and second-order temporal-difference penalties to preserve waveform morphology and activation-related structure

2.5. Evaluation metrics and robustness experiments

Performance was evaluated on LV and RV pacing beats and summarized as pooled LV+RV median (IQR) values across beats. Temporal correlation coefficient (TCC) was computed as the median nodewise Pearson correlation between predicted and reference epicardial waveforms. Spatial correlation coefficient (SCC) was computed as the median framewise correlation across epicardial nodes. In addition, activation times were estimated from the minimum first temporal derivative within the QRS interval [7]. ATCC was computed as the Pearson correlation between predicted and reference activation-time maps across epicardial nodes, and LE was defined as the Euclidean distance between the predicted and reference earliest-activation sites.

Two robustness experiments were performed at test time: a lead-dropout experiment, in which a fraction of BSP channels was randomly removed, and an additive Gaussian noise experiment, in which zero-mean noise was injected into the BSP signals. For each corruption level, pooled TCC, SCC, ATCC and LE values were computed, and representative qualitative waveform panels were gen-

Table 1. Robustness comparison (median(IQR) over test beats). **Bold**: best per condition. Dropout: fraction of BSP leads channels zeroed. Noise: σ as a fraction of per-lead standard deviation.

		Lead Dropout				Additive Noise			
Level/σ	Method	TCC	SCC	ATCC	LE (mm)	TCC	SCC	ATCC	LE (mm)
0.00	FastKAN	0.995(0.002)	0.992(0.004)	0.993(0.002)	**7.8(9.9)**	0.995(0.002)	0.992(0.004)	0.993(0.002)	**7.8(9.9)**
	MLP	0.995(0.004)	0.994(0.003)	0.993(0.002)	9.9(9.9)	0.995(0.004)	0.994(0.003)	0.993(0.002)	9.9(9.9)
	ResNet	**0.997(0.002)**	**0.996(0.003)**	**0.995(0.002)**	9.9(14.3)	**0.997(0.002)**	**0.996(0.003)**	**0.995(0.002)**	9.9(14.3)
0.10	FastKAN	**0.993(0.003)**	**0.990(0.004)**	**0.993(0.007)**	**9.2(13.4)**	0.992(0.003)	0.989(0.005)	0.945(0.030)	21.9(49.5)
	MLP	0.985(0.008)	0.984(0.009)	0.991(0.004)	9.9(14.2)	0.992(0.004)	0.992(0.004)	0.958(0.046)	18.8(55.3)
	ResNet	0.988(0.007)	0.988(0.007)	0.990(0.023)	9.9(32.7)	**0.995(0.002)**	**0.995(0.003)**	**0.962(0.033)**	**18.0(55.7)**
0.25	FastKAN	**0.989(0.004)**	**0.985(0.005)**	**0.989(0.020)**	**9.9(13.3)**	0.977(0.007)	0.969(0.007)	0.904(0.048)	**28.7(55.7)**
	MLP	0.958(0.018)	0.955(0.018)	0.977(0.033)	13.4(34.0)	0.976(0.007)	0.974(0.013)	0.919(0.062)	28.8(55.7)
	ResNet	0.959(0.022)	0.958(0.025)	0.962(0.036)	21.4(54.3)	**0.980(0.005)**	**0.978(0.009)**	**0.926(0.051)**	29.2(48.5)
0.55	FastKAN	**0.923(0.030)**	**0.896(0.039)**	**0.902(0.056)**	**19.6(34.6)**	0.878(0.024)	0.797(0.057)	0.882(0.124)	37.3(55.7)
	MLP	0.829(0.046)	0.786(0.068)	0.872(0.067)	36.7(55.7)	0.865(0.024)	0.798(0.073)	0.742(0.103)	42.9(36.9)
	ResNet	0.790(0.054)	0.737(0.072)	0.800(0.184)	52.9(47.9)	**0.892(0.022)**	**0.825(0.091)**	**0.912(0.058)**	**34.9(43.5)**

erated to visualize model behavior under degraded measurement conditions.

3. Results

Table 1 summarizes pooled LV+RV reconstruction metrics under lead-dropout and additive-noise conditions. Under clean inputs, all three models achieved near-perfect reconstruction, with ResNet performing best overall and FastKAN and MLP following closely.

Under lead dropout, the left half of Table 1 reveals a clear separation among models as the dropout fraction increases. At 10% lead removal, FastKAN already leads in all three metrics, whereas MLP and ResNet show the first signs of degradation. The gap widens substantially at 25% dropout. Under severe dropout (55%), the differences become pronounced. FastKAN maintains TCC = 0.923 and SCC = 0.896. In contrast, MLP falls to TCC = 0.829 and SCC = 0.786, and ResNet degrades even further to TCC = 0.790 and SCC = 0.737. The ATCC metric follows a similar trend: FastKAN preserves 0.902, compared to 0.872 for MLP and 0.800 for ResNet. LE follows the same pattern under lead dropout, with FastKAN achieving the lowest LE at all dropout levels.

By contrast, the right half of Table 1 presents a different pattern under additive noise. Here the ranking is reversed, with ResNet consistently achieving the best metrics across all noise levels. At $\sigma = 0.10$, all models remain above TCC = 0.99, but ResNet leads with 0.995. At $\sigma = 0.25$, ResNet maintains TCC = 0.980 and SCC = 0.978, slightly ahead of both FastKAN and MLP. Under severe noise ($\sigma = 0.55$), ResNet retains TCC = 0.892 and ATCC = 0.912, outperforming FastKAN (TCC = 0.878, ATCC = 0.882) and MLP (TCC = 0.865, ATCC = 0.742). For additive noise, LE only partially matches the correlation-based ranking: ResNet yields the lowest LE at $\sigma = 0.10$ and $\sigma = 0.55$, whereas FastKAN is marginally best in the clean case and at $\sigma = 0.25$.

Figure 1 illustrates both robustness experiments on a representative RV beat. In the top row, which shows increasing lead dropout, all three models match the reference waveform closely at low dropout levels, whereas at higher dropout FastKAN better preserves the QRS morphology and timing than MLP and ResNet. In the bottom row, which shows increasing additive Gaussian noise, all three models retain the overall waveform shape, but ResNet produces the smoothest reconstruction under severe noise and remains closest to the reference signal.

4. Discussion and Conclusion

The results reveal a complementary robustness profile between FastKAN and ResNet. FastKAN is consistently more robust under lead-loss conditions, whereas ResNet performs best under additive-noise corruption. The LE results broadly support the same pattern, favoring FastKAN for missing-lead scenarios and ResNet for stronger noise levels. A possible explanation lies in how the architectures process corrupted inputs. In FastKAN, each BSP channel is expanded independently through learnable radial basis functions before linear mixing. This may help preserve useful information when some channels are missing. In contrast, ResNet appears to handle additive noise more effectively, likely due to its residual structure, which promotes smoother and more stable reconstructions. Although this interpretation remains tentative, it is consistent with the observed trends. These findings are relevant in practice, where electrode detachment or poor skin contact may lead to incomplete BSP measurements [8].

Several limitations should be noted. First, the evaluation was conducted on a single torso-tank experiment, and although the training, validation, and test sets were dis-

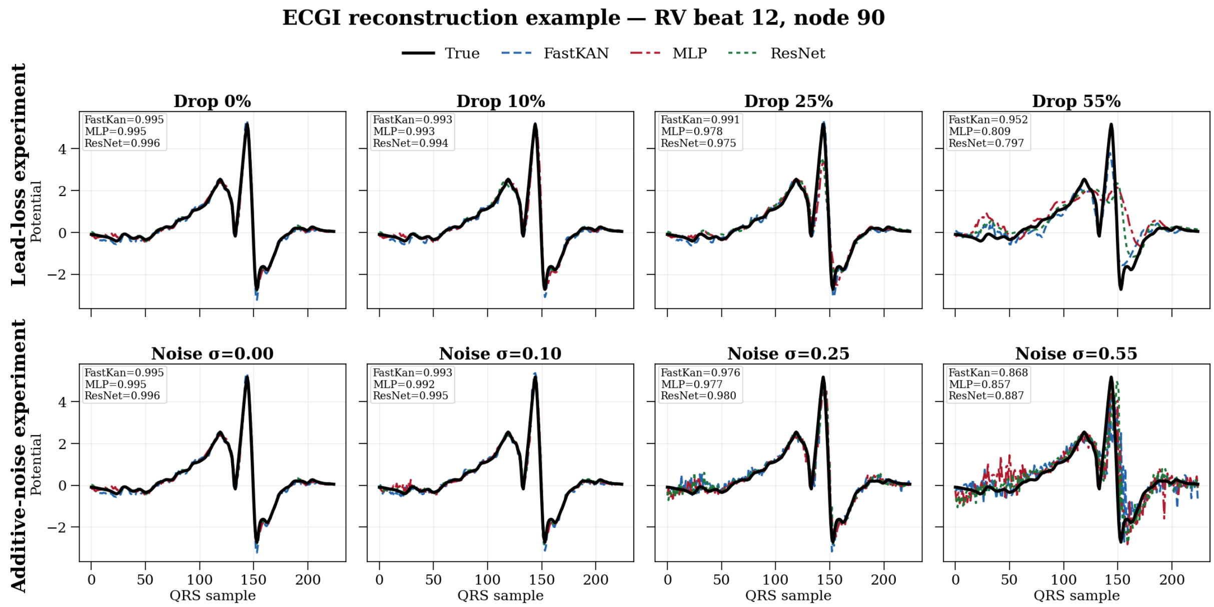

Figure 1. Epicardial potential reconstruction for a beat under increasing lead dropout (top) and additive Gaussian noise (bottom). FastKAN, MLP, and ResNet predictions are compared with the reference waveform across four corruption levels.

joint, they were all derived from recordings from the same experiment; therefore, both anatomical and electrophysiological variation remain limited. Second, the lead-loss experiment assumed random, spatially uniform channel failure, whereas structured electrode loss patterns may produce different behavior.

In conclusion, KAN-based architectures offer a clear advantage under missing-lead conditions while remaining competitive under clean and noisy settings. This robustness stems from the KAN architecture's functional decomposition, where each input dimension is first transformed independently through learnable one-dimensional functions before being combined, thereby reducing reliance on strongly coupled feature interactions and enabling more graceful degradation when inputs are missing. As a result, the model degrades more gracefully under electrode loss compared to standard neural regressors. FastKAN therefore appears promising for robust ECGI, particularly when electrode completeness is a concern. Future work will extend this study to more realistic ECGI scenarios, including structured electrode loss patterns, broader anatomical variability, and clinical data.

References

[1] Ramesh M Gulrajani. The forward and inverse problems of electrocardiography. *IEEE Engineering in Medicine and Biology Magazine*, 17(5):84–101, 1998. doi:10.1109/51.715491.

[2] Ke-Wei Chen, Laura Bear, and Che-Wei Lin. Solving inverse electrocardiographic mapping using machine learning and deep learning frameworks. *Sensors*, 22(6), 2022. doi:10.3390/s22062331.

[3] Ramiro Valdes Jara and Adam Meyers. Geometry-free conditional diffusion modeling for solving the inverse electrocardiography problem, 2026. arXiv:2601.18615.

[4] Ziming Liu, Yixuan Wang, Sachin Vaidya, Fabian Ruehle, James Halverson, Marin Soljačić, Thomas Y Hou, and Max Tegmark. Kan: Kolmogorov-arnold networks, 2024. arXiv:2404.19756.

[5] Kedar Aras, Wilson Good, Jess Tate, Brett Burton, Dana Brooks, Jaume Coll-Font, Olaf Doessel, Walther HW Schulze, Danila Potyagaylo, Linwei Wang, Peter Dam, and Rob MacLeod. Experimental data and geometric analysis repository-edgar. *Journal of Electrocardiology*, 48(6):975–981, 2015. doi:10.1016/j.jelectrocard.2015.08.008.

[6] Ziyao Li. Kolmogorov-arnold networks are radial basis function networks, 2024. arXiv:2405.06721.

[7] Douglas P. Zipes and José Jalife. *Cardiac Electrophysiology: From Cell to Bedside*. Saunders/Elsevier, Philadelphia, PA, USA, 2009.

[8] Y S Dogrusoz, L R Bear, J Bergquist, R Dubois, W Good, R S MacLeod, A Rababah, and J Stoks. Effects of interpolation on the inverse problem of electrocardiography. In *2019 Computing in Cardiology (CinC)*, pages Page 1–Page 4, 2019. doi:10.22489/CinC.2019.100.

Address for correspondence:

Yunus Emre Tüysüz

Electrical and Electronics Engineering Department

Middle East Technical University, Ankara, Turkey

emre.tuysuz@metu.edu.tr