

LETTER

Sequential Gating Is a Restricted Form of Multimodal Fusion in TARGET-AI

Chin Lin , Ph.D.,^{1,2} and Kai-Chieh Chen , Ph.D.^{1,2}

Received: February 8, 2026; Accepted: March 16, 2026; Published: August 27, 2026

To the Editor

Oikonomou et al. report TARGET-AI, a clinically motivated “gate then classify” deployment strategy that uses longitudinal electronic health record (EHR) information to decide who should receive artificial intelligence-enabled electrocardiography (AI-ECG) screening for structural heart disease, and only then applies an ECG model to those selected.¹ They show improved screening efficiency, including fewer false positives, compared with untargeted AI-ECG screening.¹ The results are compelling, and the deployment framing is also excellent. But one of the results (“sequential gating beats multimodal fusion”) should be interpreted with a key qualifier: It is reliable only relative to the particular fusion family evaluated. This is not a criticism of TARGET-AI; it is a clarification of what the evidence can and cannot justify.

In theory, we can define two hypothesis classes: (1) unrestricted joint class $C_{\text{all}} = \delta(E, X)$, where δ is any measurable function; and (2) sequential gating class $C_{\text{seq}} = \delta(E, X) = \mathbb{1}[g(E) \geq t] \cdot \mathbb{1}[h(X) \geq s]$, that is, “pass EHR gate, then pass ECG threshold”, where g and h are any measurable functions, E is EHR, X is ECG, and t and s are thresholds. With $F1(\delta)$ as the F1 score in a specific function using the same EHR-ECG joint database as the input, we can declare:

$$\max_{\delta \in C_{\text{all}}} F1(\delta) \geq \max_{\delta \in C_{\text{seq}}} F1(\delta)$$

Since every sequential classifier is a valid function of (E, X) and $C_{\text{seq}} \subseteq C_{\text{all}}$, maximizing F1 over a set cannot yield a smaller value than maximizing over its subset. In TARGET-AI, fusion was implemented as a concatenation of frozen EHR and ECG embeddings followed by logistic regression. The comparison in TARGET-AI is therefore best read as evidence of a restricted (linear) fusion baseline under specific training constraints.

This can be seen by summarizing each modality into a single score. Linear fusion corresponds to a single threshold on a weighted sum of the EHR and ECG scores (a straight decision boundary). Sequential gating requires both scores to exceed their thresholds, producing an L-shaped boundary (Fig. 1). When cases occupy a “high-high” corner and controls are dispersed elsewhere, the L-shaped rule can reduce false positives even though both strategies use the same two scores; conversely, when the joint score space is close to linearly separable, gating offers little benefit.

Accordingly, the TARGET-AI results should be framed as evidence about a particular fusion family on specific data rather than as a universal statement that sequential strategies

The author affiliations are listed at the end of the article.

Chin Lin can be contacted at xup6fup0629@gmail.com or at Room 7202-1, No. 161, Min-Chun E. Rd., Sec. 6, Neihu, Taipei 114, Taiwan, ROC.




Figure 1. A Geometric View of Multimodal Fusion in EHR-ECG Screening.

Each point represents an individual in the plane of an EHR-derived score (horizontal axis) and an ECG-derived score (vertical axis). Brown points indicate cases and teal points controls. Linear fusion yields a straight decision boundary (left panels); sequential gating yields an L-shaped boundary requiring both scores to exceed thresholds (right panels). Panel A illustrates a corner-like separation favoring gating; Panel B illustrates near-linear separability, where gating may offer limited advantage. EHR denotes electronic health record.

dominate joint modeling. More broadly, this study revealed that future studies should apply at least one nonlinear fusion baseline. For example, tree-based models (“gate then classify” is mathematically a very small decision tree model) should be considered to fuse multimodal information when end-to-end training is not performed. If a study allows full-parameter fine-tuning of $g(E)$ and $h(X)$, the same linear fusion rule could become much more competitive because the underlying nonlinear models (e.g., transformers) may learn better representations to reshape the score space toward linear separability. Linear fusion should not be seen as a baseline, but rather as proof that the proposed foundation model has a strong representational ability.




Figure 2. An Example to Illustrate the Importance of Prospective Accuracy Studies.

Only data from individuals who have both an echocardiogram and an ECG simultaneously can be used to develop and validate the model, but these individuals will not be the target population in the future.

Such reporting would better separate gains attributable to the choice of decision boundary from gains attributable to representation learning and would sharpen guidance for prospective deployment.

Beyond retrospective benchmarking, we view TARGET-AI’s simplicity as a good tool for screening strategy optimization in small-scale prospective studies. Experience from the Columbia group’s work on ValveNet underscores the point: Calibration that appears favorable in retrospective cohorts² may degrade when an AI-ECG model is tested prospectively.³ As conceptually illustrated in [Figure 2](#), the population used for retrospective model development is fundamentally different from the population that stands to benefit from AI-ECG. Even if prospective enrollment is restricted to AI-ECG-positive individuals, an EHR model can be reapplied within that subset to enrich the yield, because gating EHR then ECG or ECG then EHR is mathematically equivalent. This will help us find the rules and apply them to subsequent large-scale EHR-gated AI-ECG trials. As AI-ECG broadens the clinical reach of the conventional ECG,⁴ integrating EHR content and ECG will be essential to improve positive predictive value. Looking ahead, we anticipate that future clinical trials will increasingly use EHR-derived contextual information as enrollment criteria, rather than relying solely on conventional variables such as age, sex, and isolated medical history.

Disclosures

Author disclosures are available at ai.nejm.org.

Author Affiliations

¹Medical Technology Education Center, School of Medicine, College of Medicine, National Defense Medical University, Taipei, Taiwan, ROC

²Military Digital Medical Center, Tri-Service General Hospital, National Defense Medical University, Taipei, Taiwan, ROC

References

1. Oikonomou EK, Batinica B, Dhingra LS, Aminorroaya A, Coppi A, Khera R. TARGET-AI: a foundational approach for the targeted deployment of artificial intelligence electrocardiography in the electronic health record. *NEJM AI* 2026;3(2). DOI: [10.1056/AIoa2500588](https://doi.org/10.1056/AIoa2500588).
2. Elias P, Poterucha TJ, Rajaram V, et al. Deep learning electrocardiographic analysis for detection of left-sided valvular heart disease. *J Am Coll Cardiol* 2022;80:613-626. DOI: [10.1016/j.jacc.2022.05.029](https://doi.org/10.1016/j.jacc.2022.05.029).
3. Poterucha TJ, Jing L, Ricart RP, et al. Detecting structural heart disease from electrocardiograms using AI. *Nature* 2025;644:221-230. DOI: [10.1038/s41586-025-09227-0](https://doi.org/10.1038/s41586-025-09227-0).
4. Lin CS, Liu WT, Chen YH, Lin SH, Lin C. Artificial intelligence-enabled electrocardiography from scientific research to clinical application. *EMBO Mol Med* 2026;18:22-40. DOI: [10.1038/s44321-025-00351-y](https://doi.org/10.1038/s44321-025-00351-y).