Effect of Token Granularity on Self-Supervised ECG Representations for Myocardial Infarction Detection

Background: Self-supervised learning (SSL) has emerged as a promising approach for learning ECG representations from large unlabeled datasets. However, the impact of temporal representation granularity on downstream clinical tasks remains poorly understood. In particular, it is unclear whether increasing the number of latent tokens extracted from a fixed ECG window improves representation quality.

Methods: We trained a self-supervised encoder using a SimSiam architecture with a bidirectional gated recurrent unit (BiGRU) backbone on multiple ECG databases, including PTB-XL, LUDB, and QTDB. The encoder processes 4-second 12-lead ECG segments sampled at 125 Hz. To study representation granularity, the number of latent tokens extracted from each segment was varied (1, 10, 20, 50, 100). Token representations were aggregated into a fixed-length embedding using either average or maximum pooling. After SSL training, the encoder was frozen and a multilayer perceptron classifier was trained to detect MI using PTB-XL. Model performance was evaluated at both window and patient levels, with external validation on the PTB Diagnostic Database (PTBDB).

Results: Increasing token granularity did not improve MI detection performance. Models using a small number of tokens achieved comparable or better results than those using higher token counts. In addition, maximum pooling consistently outperformed average pooling, suggesting that preserving the strongest local features is beneficial for the MI detection task. The best configuration achieved a patient-level AUC of approximately 0.93 on the external PTBDB test set.

Conclusions: In this setting, compact token representations combined with max aggregation provided the most effective embeddings for MI detection. Future work will investigate whether larger encoder capacities or alternative architectures can better leverage higher token granularity.

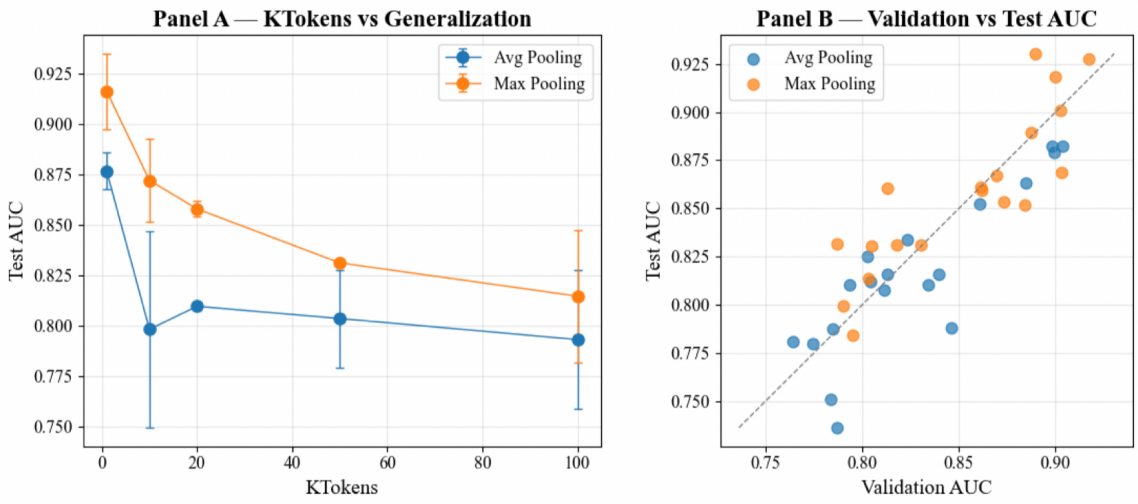

Figure. Effect of token granularity and pooling strategy on MI detection. (A) Patient-level test AUC on the external PTBDB dataset as a function of the number of latent tokens. (B) Validation versus external test AUC for all experimental configurations. Each point corresponds to a combination of token count, pooling method, and random seed.