A Transfer Learning Model-Based Patient-Specific ECG Lead Synthesis Algorithm

Arjun Malpani[1], Joel Xue[1,2], David Albert[1], Gari Clifford[2,3]

[1]Alivecor Inc., Mountain View, CA, USA
[2]Emory University, Atlanta, GA, USA
[3]Georgia Institute of Technology, Atlanta, GA, USA

Abstract

The 12-lead electrocardiogram (ECG) is a fundamental diagnostic tool in clinical cardiology, yet often, attaching the necessary electrodes can be cumbersome, particularly with wearable devices that are optimized for at-home monitoring and record a reduced-lead set. Machine learning-based approaches have been employed to synthesize the remaining leads required for the 12-lead ECG, yet the generalizable nature of these models may not capture the distinctive morphology of a given patient's ECG. Herein, this study aims to incorporate a transfer learning-based personalized model to improve the accuracy of ECG lead synthesis on a per-patient basis. A generalized deep learning-based synthesis framework is initially developed, processing inputs from four measured ECG leads (I, II, V2, V4) to generate outputs for four synthesized leads (V1, V3, V5, V6). Trained instances of this model are subsequently fine-tuned on five full 12-lead recordings per patient, a notable reduction in the number of ECGs needed for personalization compared to similar work. Significant improvements in the personalized models over the generalized models were observed in patients with both normal and abnormal morphologies, indicating promising implications for the clinical application of personalized reduced-lead ECG lead synthesis.

1. Introduction

Cardiovascular diseases (CVDs) represent the predominant cause of mortality globally; in 2021 they accounted for approximately 20.5 million deaths, one-third of all worldwide [1]. The 12-lead electrocardiogram (ECG) remains the most widely utilized tool in the diagnosis of CVDs, and a particular emphasis has been placed on the early attainment of a 12-lead ECG in improving mortality and health-related outcomes [2]. However, doing so requires more frequent monitoring, which accompanies numerous practicality challenges– primarily including correct attachment of the 10 electrodes to the body. Considering this, advances in the development of wireless personal ECG recording devices have enabled the possibility of circumventing these drawbacks on a large scale. Such devices often record fewer leads, necessitating mathematical or machine learning techniques in synthesizing the remaining leads. This is made possible by the existence of redundancies among the 12 ECG leads, which have been successfully leveraged for ECG lead synthesis [3]. The efficacy of such synthesized leads in serving as a diagnostic tool is well-documented; various studies have demonstrated statistically significant lead reconstruction and equal or near-equal performance on cardiac abnormality classification from a partially synthesized 12-lead ECG [4,5,6].

However, the majority of such methods employ a singular generalized approach to generate the ECG leads; one model with fixed parameters is used for all patients, a process that may fail to capture unique waveform morphologies. We aim to investigate the effectiveness of utilizing a personalized approach to ECG lead synthesis via transfer learning, enabling fine-tuning of the generalized model with patient-specific recordings. This approach has been minimally explored in existing literature. One of the few such works utilized transfer learning to create a personalized long short-term memory-based lead synthesis model for any patient [7]. However, while observing highly significant improvement after the personalization, hundreds of recordings were utilized per patient for fine-tuning, which is highly impractical in practice. As such, our work explores the viability of transfer learning in enhancing ECG lead synthesis using a smaller, practical number of recordings per patient.

Eight leads are utilized for analysis in this study; leads I, II, V2, and V4 serve as measured input leads, and leads V1, V3, V5, and V6 serve as synthesized output leads. Linear mathematical relationships between the unipolar limb leads (aVL, aVR, aVF) and the bipolar limb leads (I, II, III) have been well-reported [8, 9]; given two from this group, the other four can be derived. Deep learning aims to capture more nuanced patterns, hence the choice of leads in this work.

2. Methods

2.1. Data

All ECG data in this study is sourced from the Mayo Clinic. Two primary datasets are utilized to train and evaluate the generalized and personalized ECG lead synthesis models–referenced as group G and group P, respectively. All waveforms undergo preprocessing to derive a median beat consisting of 1.2 seconds of 150 sps downsampled data. Each recording comprises 12 leads from which 8 are selected for analysis: leads I, II, V2, and V4 as input, and leads V1, V3, V5, and V6 as output. Group G–employed for training the initial generalized synthesis model–encompasses 20,000 such recordings from unique patients, 15,000 of which are randomly categorized for training and 5000 for validation. Group P, comprising recordings from 971 patients (40 ± 10 recordings each) independent from those in group G, is employed for evaluating the personalized models on an individual patient basis. Group P is further categorized into morphology-based subgroups, with a particular focus on ischemia, old and acute myocardial infarction (MI), and normal morphologies.

2.2. Generalized synthesis model

The initial stage of this study involves the development of a generalized lead synthesis model, which should serve as both a robust baseline for generating a complete ECG for any patient and a deep learning framework that is effectively adaptable for subsequent patient-specific personalization. Given that ECGs are one-dimensional time series by nature, a one-dimensional convolutional neural network (CNN) is chosen as a suitable architecture for the generalized model. Additionally, the structure of a CNN enables increased flexibility in controlling the number of parameters that remain trainable for the transfer-learning phase of this work. Differing deep-learning approaches reported for lead syntheses, such as the multi-layer perceptron or long short-term memory network [6,7], have rigid, typically output size-dependent final layers, which often remain unfrozen for transfer learning. CNNs allow for independent adjustment of the number of filters and kernels at any given point in the model regardless of output size, increasing the degree of control and adaptability during transfer learning.

The generalized model is loosely constructed in an encoder-decoder format. Multi-head attention was implemented in the initial encoder phase of the model but was not conducive to performance improvements, perhaps due to the smaller input shape of median-beat data. A simple grid search was utilized to determine an optimal set of model hyperparameters. Various forms of regularization were utilized: L2 regularization with a decay rate of 0.00025 and a single batch normalization layer after the first decoder block.

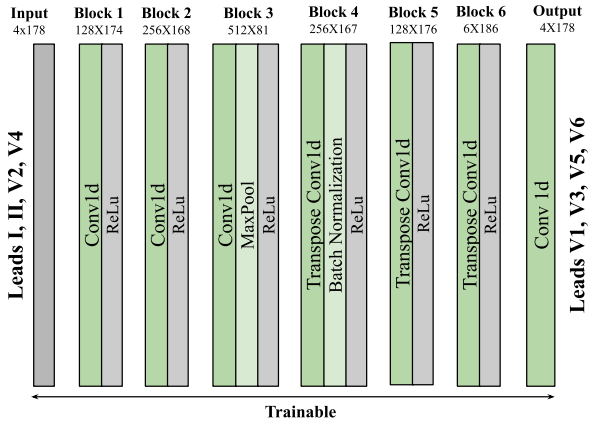

Figure 1. Generalized deep learning model architecture

2.3. Personalized synthesis model

Transfer learning is the subsequent phase to the development of the generalized model; new, fine-tuned personalized models are created for each patient within the Group P dataset. A key consideration of this phase is the number of recordings per patient necessary for the personalized model to observe improved performance over the generalized model. In this work, for any given patient, groups of ten recordings arranged chronologically are isolated. The first five recordings per group–representing the number of recordings a patient would need for personalization–are utilized for fine-tuning the associated personalized model, with the second five recordings being delineated for evaluation (618 ± 570 days between fine-tuning and evaluation recordings). To obtain these groups, for each patient in group P, a sliding window of size ten and stride one is applied to their recordings. This results in the obtainment of 21-40 groups of ten recordings per patient, with the intent of capturing potential morphological changes that a patient may experience over time. This process is further repeated exclusively for patients with severe morphological abnormalities–namely ischemia and previous/acute myocardial infarction (MI). For acute MI, the 5 training recordings represent normal waveforms and the 5 evaluation recordings indicate acute MI, to simulate real-world situations.

For each of these groups, a new pre-trained version of the generalized model with identical parameters is subsequently initialized. All parameters in this model are frozen, excluding the final convolutional layer (220 trainable parameters) to be used for transfer learning-based personalization. It is well-known that the number of trainable parameters within machine learning and deep learning contexts must be calibrated with training data

size to prevent overfitting–this is where the flexible CNN structure becomes particularly useful. A grid search was utilized to determine the number of kernels and channels within the final layer. No form of regularization yielded improved synthesis, and 60 epochs at a learning rate of 0.002 were found to produce optimal results while taking <2 seconds per group for training.

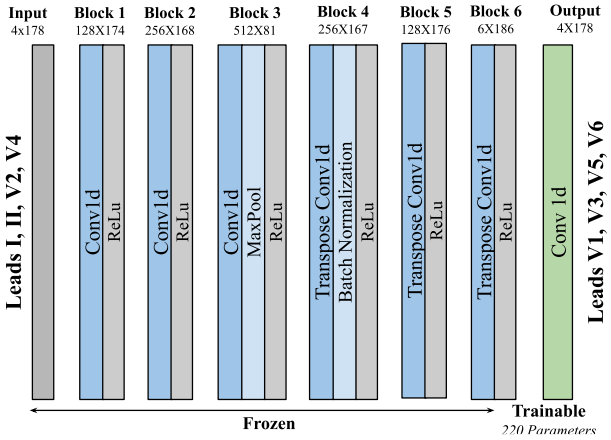

Figure 2. Personalized deep learning model architecture

3. Results and analysis

Patient-specific models are evaluated relative to the performance of the generalized model for each group of five validation recordings. Figure 3 compares the correlation (averaged over the five validation ECGs) of the synthesized leads and the ground truth between the generalized and personalized model for all leads and groups of recordings among the 971 group P patients. A notable improvement is observed–particularly within the lower-performance ranges. Moreover, the performance improvement remains relatively consistent across all four leads, further exemplified by the evaluation metrics presented per morphology-based subgroup in Table 1.

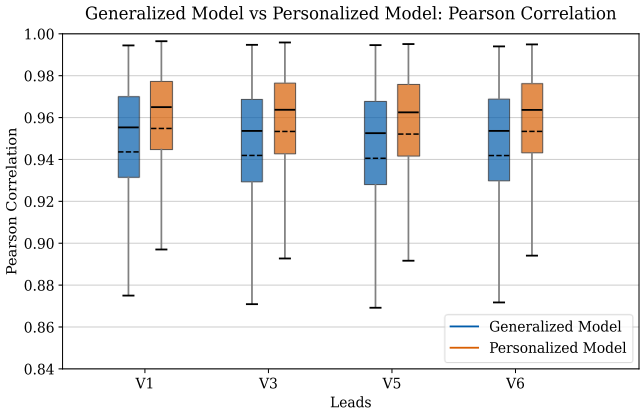

Figure 3. Generalized vs personalized model Pearson correlation comparison

3.1. Morphology subgroup analysis

		Ischemia				Previous MI			
		V1	V3	V5	V6	V1	V3	V5	V6
G	*r*	.939	.938	.936	.936	.936	.933	.932	.934
	rmse	.062	.061	.061	.062	.053	.053	.054	.053
P	*r*	.958	.955	.944	.956	.952	.950	.949	.951
	rmse	.052	.051	.053	.052	.045	.046	.046	.046
		Normal				**Acute MI**			
		V1	V3	V5	V6	V1	V3	V5	V6
G	*r*	.973	.972	.970	.968	.925	.925	.924	.925
	rmse	.033	.034	.035	.035	.060	.060	.061	0.60
P	*r*	.978	.977	.975	.975	.925	.924	.922	.923
	rmse	.029	.030	.031	.030	.061	.061	.062	.062

Table 1. Morphology-specific correlation (r) and root-mean-squared-error (rmse) comparisons between global model (G) and personalized models (P)

Patient-specific performance is further evaluated with respect to morphology-categorized subgroups, as seen in Table 1. Figure 4 plots the mean correlation improvement across all leads per recording group as a function of their original correlation. Lead-by-lead visualization is not vital; as aforementioned, there is very little variability between performance improvements per lead. Notably, performance improvement via transfer learning becomes increasingly pronounced as the original performance declines. Furthermore, patients with abnormal ECG morphologies–specifically ischemia and previous MI–demonstrate more significant improvement than those with normal morphologies. The lack of consistent improvement observed in patients experiencing acute MI may be attributed to its more localized nature and the fact that training only occurs on normal recordings before the event.

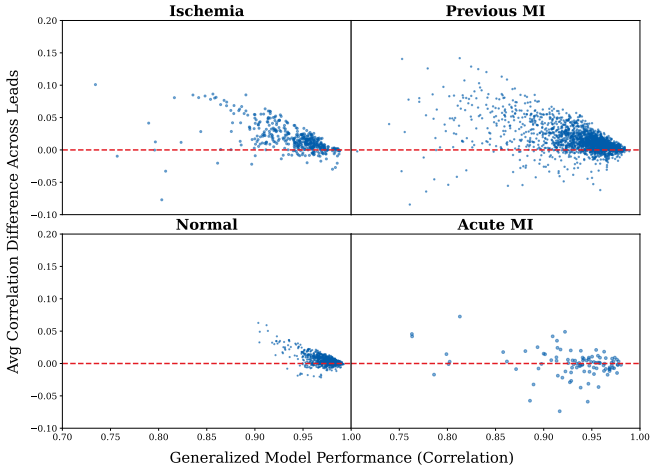

Figure 4. Morphology-specific personalized correlation improvement as a function of original generalized model performance

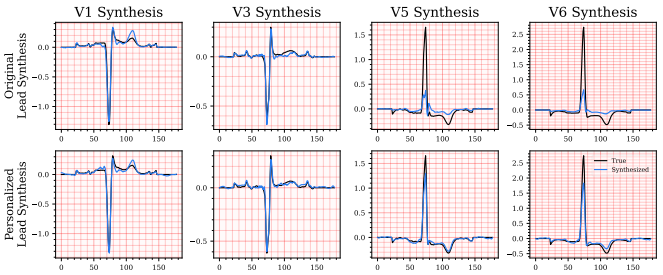

Figure 5. Waveform example of patient with ischemia

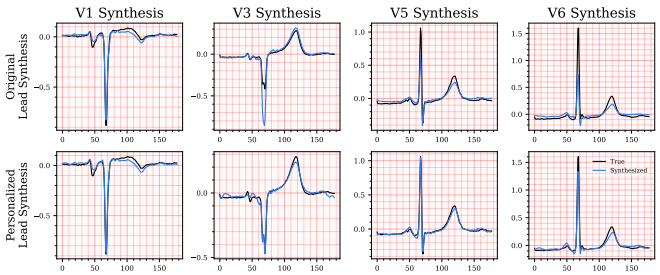

Figure 6. Waveform example of patient with previous MI

3.2. Qualitative waveform analysis

Figure 5 is an example visualization comparing generalized and personalized lead synthesis for an ischemia patient. A particular improvement can be observed in the R-peak amplitude of leads V5 and V6, along with a more accurate inverted T-wave replication. Similar amplitude improvements to a lesser extent can be observed in the synthesis of the V1 and V3 leads in the personalized models.

Figure 6 visualizes the waveforms for a patient with previous MI. The most significant improvement after personalization can be seen in the Q wave of lead V3, while amplitude adjustments are also evident in leads V5 and V6, a pattern also commonly observed in other patients.

4. Discussion and conclusion

In this work, we demonstrate the potential of transfer learning in enhancing ECG lead synthesis from a reduced-lead set with minimal fine-tuning data. Most notably, after personalization, we observe an increasing magnitude of improvement as the original performance from the generalized model decreases–particularly in patients with abnormal morphologies, where generalized models are less accurate on average. In the less common cases where performance decreases after transfer learning, the original lead synthesis is typically adequate to negate this performance drop. This is a promising indication that transfer learning can effectively address the most prominent limitations with current generalized models. With leveraging only five full 12-lead recordings for patient-specific adaptation, on average 618 days before the ECGs used for evaluation were recorded, we believe these results to be practical for clinical application.

Future work involves a more comprehensive investigation of lead synthesis during acute MI events; while some instances of patients with acute MI experienced an improvement after personalization, this was not a consistent finding in our work. Additionally, a comprehensive analysis into the optimal number of recordings for fine-tuning may prove insightful, further solidifying our findings and contributing to the advancement of ECG lead analysis.

References

[1] Summary of Global Burden of Disease Study Methods. Journal of the American College of Cardiology. 2022 Dec;80(25):2372–425.

[2] Seong Yong Yoon, Kim T, Kang EJ, Heo S, Chang H, Seo Y, et al. Feasibility of patch-type wireless 12-lead electrocardiogram in laypersons. Scientific Reports. 2023 Mar 10;13(1). doi: 10.1038/s41598-023-31309-0

[3] Jain U, Butchy A, Leasure M, Covalesky V, McCormick D, Mintz G. Redundancy and Novelty Between ECG Leads Based on Linear Correlation. 2023 Jan 1. doi: 10.5220/0011815700003414

[4] Mason F, Pandey AC, Gadaleta M, Topol EJ, Muse ED, Quer G. AI-enhanced reconstruction of the 12-lead electrocardiogram via 3-leads with accurate clinical assessment. npj Digital Medicine [Internet]. 2024 Aug 1 7(1).

[5] Wang L, Zhou W, Xing Y, Liu N, Movahedipour M, Zhou X. A novel method based on convolutional neural networks for deriving standard 12-lead ECG from serial 3-lead ECG. Frontiers of Information Technology & Electronic Engineering. 2019 Mar;20(3):405–13.

[6]Ato Kapfo, Datta S, Samarendra Dandapat, Prabin Kumar Bora. Artificial Neural Network Based Synthesis of 12-Lead ECG Signal from Three Predictor Leads. Lecture notes in electrical engineering. 2022 Jan 1;625–34.

[7] Shyam Kumar P, Ramasamy M, Kallur KR, Rai P, Varadan VK. Personalized LSTM Models for ECG Lead Transformations Led to Fewer Diagnostic Errors Than Generalized Models: Deriving 12-Lead ECG from Lead II, V2, and V6. Sensors. 2023 Jan 26;23(3):1389.

[8] Madias JE. On recording the unipolar ECG limb leads via the Wilson's vs the Goldberger's terminals: aVR, aVL, and aVF revisited. Indian pacing and electrophysiology journal [Internet]. 2008;8(4):292–7.

[9] 1.Sohn J, Yang S, Lee J, Ku Y, Kim HC. Reconstruction of 12-Lead Electrocardiogram from a Three-Lead Patch-Type Device Using a LSTM Network. Sensors. 2020 Jun 9;20(11):3278.

Address for correspondence:
Arjun Malpani
Arjun.s.malpani@gmail.com