A Hierarchical Generative Model of Electrocardiogram Records

Andrew C. Miller *
School of Engineering and Applied Sciences
Harvard University
Cambridge, MA 02138
acm@seas.harvard.edu

Sendhil Mullainathan
Department of Economics
Harvard University
Cambridge, MA 02138
mullain@fas.harvard.edu

Ziad Obermeyer
Brigham and Women's Hospital Harvard Medical School
Boston, MA 02115
zobermeyer@bwh.harvard.edu

Abstract

We develop a probabilistic generative model of electrocardiogram (EKG) tracings. Our model describes multiple sources of variation in EKGs, including patient-specific cardiac cycle morphology and between-cycle variation that leads to quasi-periodicity. We use a deep generative network as a flexible model component to describe variation in beat-specific morphology. We apply our model to a set of 549 EKG records, including over 4,600 unique beats, and show that it is able to discover interpretable information, such as patient similarity and meaningful physiological features (e.g., T wave inversion).

1 Introduction

An electrocardiogram (EKG) is a common non-invasive medical test that measures the electrical activity of a patient's heart by recording the time-varying potential difference between electrodes placed on the surface of the skin. The resulting data is a multivariate time-series that reflects the depolarization and repolarization of the heart muscle that occurs during each heartbeat. These raw waveforms are then inspected by a physician to detect irregular patterns that are evidence of an underlying physiological problem.

In this paper, we build a hierarchical generative model of electrocardiogram signals that disentangles sources of variation. We are motivated to build a generative model of EKGs for multiple reasons. First, we would like our inferences to properly cope with nuisance variation present in EKG signals (e.g. variation in cardiac cycle with breathing and inadvertent movement). Second, generative models can be used for semi-supervised tasks — not all EKG observations are paired with test results or diagnoses. Semi-supervised modeling allows us to leverage a large amount of unlabeled EKG data to improve predictions when training with a smaller labeled dataset (under an appropriate model of censoring or missingness). Third, in medical diagnoses, correlating model features with underlying physiological realities is important for several reasons, including model checking and interpretability. A generative model provides an intuitive mechanism to examine the inner workings of the model — one can always draw a sample from the inferred generative distribution to reveal what the latent features themselves represent in terms of observed data features. Finally, a statistical model can be much more sensitive and robust to difficult-to-detect patterns.

*http://andymiller.github.io

Machine Learning for Health (ML4H) Workshop at NeurIPS 2017.

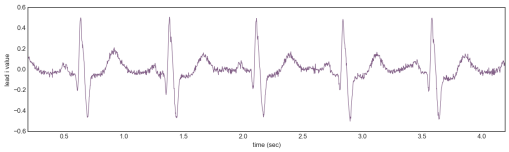

(a) Patient 1

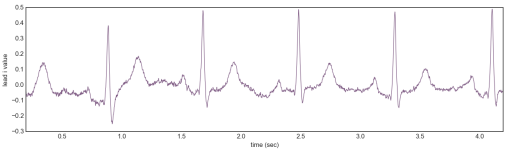

(b) Patient 2

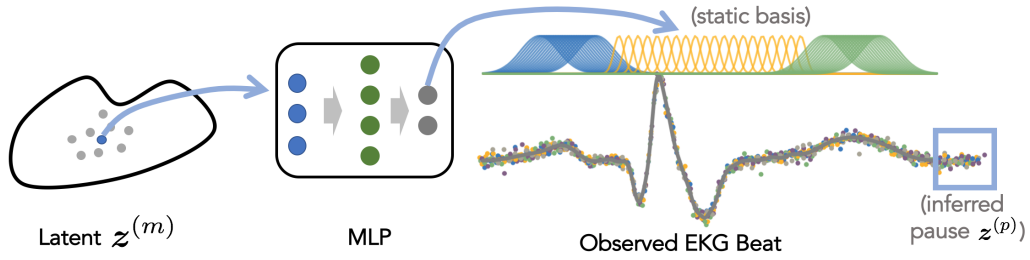

(c) Generative model

Figure 1: Top: Example EKG tracings (single lead) from two patients. Bottom: The generative procedure — each beat is represented in a low-dimensional latent space (left). To generate a beat, this vector is up-sampled through a multi-layer perceptron (center), resulting in a set of coefficients (depicted as grey dots above). These coefficients are used with an over-complete set of fixed radial basis functions (top right) to describe the raw EKG signal, excluding the inferred pause duration.

2 Modeling Electrocardiograms

Looking at an EKG, a few features of the tracing stand out. First, there are individual heartbeats — discrete periods of active contraction of the heart muscle, that are responsible for pumping blood to the lungs and the rest of the body. These are the bursts of activity in the tracing. Small parts of the signal in this area are carefully scrutinized by physicians for signs of disturbance in the heart's electrical conduction system, or heart attack. Second, there are the periods between heartbeats, when the muscle is resting as the heart fills with blood. Since there is little electrical activity during this period, the electrodes record it as a flat line. The time between beats can vary as a function of idiosyncratic aspects of the patient's conduction system, or with variations in blood flow to the heart driven by the respiratory cycle. Third, patient movement or artifacts in the recording equipment (conductance of the electrodes, etc) can introduce arbitrary changes in the signal.

Our generative model tackles these multiple sources of variation in electrocardiogram data directly. We model: (i) the morphology of an individual beat; (ii) the variation in that morphology from beat to beat; (iii) the variability in the periodicity of the beat; and (iv) nuisance variability in the measurement process (e.g. overall drift due to movement). The output of this model will be a set of features that explain these sources of variability, which can be useful in exploratory and predictive tasks.

Our data are multi-dimensional temporal observations $\boldsymbol{Y} = \boldsymbol{y}_1, \ldots, \boldsymbol{y}_T$ where $\boldsymbol{y}_t \in \mathbb{R}^D$ are sampled on a regular time grid, $\boldsymbol{t}^{(obs)} = t_1, \ldots, t_T$ (given in seconds). Our model separately parameterizes the morphology of the cardiac cycle, its duration, and the length of time between cycles. We express these sources of variation as a hierarchical probabilistic model

$$\boldsymbol{z}^{(m)}, \boldsymbol{z}^{(p)} \sim p(\boldsymbol{z}^{(m)}, \boldsymbol{z}^{(p)}; \boldsymbol{\theta})$$
$$\boldsymbol{y}_t \mid \boldsymbol{z}^{(m)}, \boldsymbol{z}^{(p)} \ldots = f(\boldsymbol{z}^{(m)}, \boldsymbol{z}^{(p)}; \boldsymbol{\theta}) + \epsilon_t \ , \epsilon_t \sim \mathcal{N}(\boldsymbol{m}_t, \sigma^2)$$

where the variables are

- $\boldsymbol{z}^{(m)}$: the morphology of a patient's beat in a low-dimensional (D) latent space.
- $\boldsymbol{z}^{(p)}$: the pause between cardiac cycles and the length of cycle (in seconds).
- $\boldsymbol{\theta}$: global parameters, including the morphology basis parameters, and prior parameters.
- $\boldsymbol{y}_t$: observed voltage, conditionally Gaussian given parameters and link function $f(\cdot)$.

The pause variable $\boldsymbol{z}^{(p)}$ measures the amount of time on each side of the cardiac cycle that can be explained by a constant offset. The morphology variable $\boldsymbol{z}^{(m)}$ explains the shape of the cardiac cycle measured by the EKG.

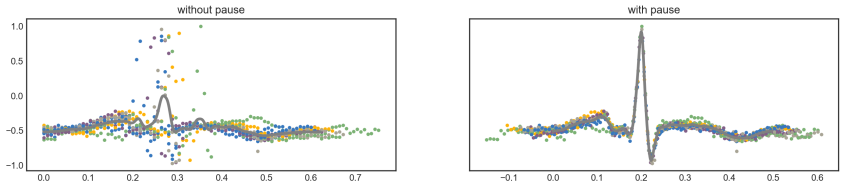

(a) Model fit, with and without beat-specific warping

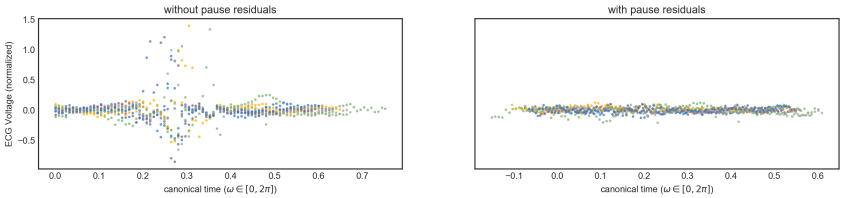

(b) Model residuals, with and without beat-specific warping

Figure 2: Top: comparison of model fit without pauses between beats (left) and with pause latent variables (right). The model average beat is shown in grey. Bottom: comparison of model residuals without (left) and with (right) pause latent variables.

Generative model of beat morphology The cardiac cycle exhibits difficult-to-prescribe variation from patient to patient and beat to beat. To address this, we represent a beat's shape with a low-dimensional latent variable that is passed through a set of non-linear basis functions (i.e. a deep neural network), parameterized by $\boldsymbol{\theta}$. The output of the deep generative model is a set of regression coefficients, applied to a fixed, temporally separated and over-complete basis. These two components model the de-noised EKG tracing; the output basis is held fixed to maintain a degree of interpretability. For inference, we use a variational autoencoder-style inference network within a variational inference framework to maximize a lower bound to the marginal likelihood of the data as a function of $\boldsymbol{\theta}$ and recognition network parameters $\boldsymbol{\phi}$ [Kingma and Welling, 2014]. The generative procedure is illustrated in Figure 1, with further detail in Appendix A.

Related Work McSharry et al. [2003] describe a generative model of EKG records defined ordinary differential equations. This model similarly includes a periodic basis, and instantiates an angular velocity to model the quasi-periodicity of the signal. However, inference for datasets of EKG records is not discussed. Oster et al. [2015] describe a switching Kalman filter approach to EKG modeling, using discrete latent states to cluster similar beat types, in contrast to our continuous latent-space description of EKG beat morphology. This approach uses discrete clusters, while our approach is more similar to a non-linear factor model. Chia and Syed [2014] align EKG beats using dynamic time warping, while our approach directly models the biologically plausible sources of temporal variation, the cycle length and pause duration.

3 Empirical Evaluation

We fit our model to the PhysioNet PTB Diagnostic EKG database [Bousseljot et al., 1995, Kreiseler and Bousseliot, 1995]. This dataset contains 549 EKG records from 290 subjects with over 4,600 individual beats. We first look at consequence of inferring beat-specific pauses, and see that it offers a sort of alignment that enables coherent modeling of the morphology. We then examine the latent space inferred by the beat morphology model.

Pause model checks Inferring pause parameters provides an effective way to align the salient features (e.g. P wave, QRS complex, T wave) in a way that can be coherently modeled. Without the pause latent variable, the temporal variability exhibited by the patient washes out features of the beat. With it, the temporal variability is appropriately separated from morphological features, which are modeled with a basis function regression model. In Figure 2 we see that residuals around the QRS complex decrease significantly when we incorporate a model of the pause between beats.

Nearest Neighbor Evaluation We examine the nearest neighbors (in different EKG records) in the latent morphology space, $\boldsymbol{z}^{(mor)}$. In Figure 3, we show a source beat and depict the five nearest neighbors in the latent morphology space. We see that the similarity in distance for this example corresponds in part to similarity in T-wave direction, while the rest of the beat remains unchanged.

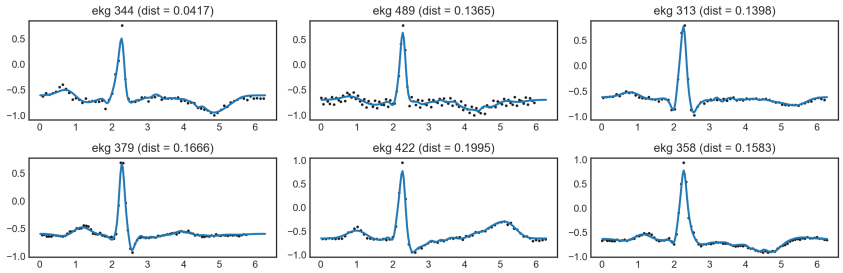

Figure 3: Nearest record examples. The source record is in the upper left (EKG 344). The following five are example beats from the nearest neighbors in latent $\boldsymbol{z}^{(m)}$ space. The solid line is the generative model reconstruction.

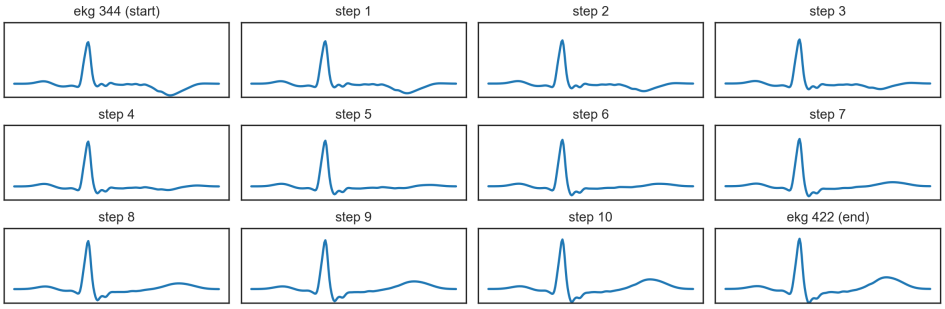

(a) Interpolation in the latent space. We start at the embedding of EKG 344 (with an inverted T wave), and linearly interpolate to EKG 422 (without an inverted T wave).

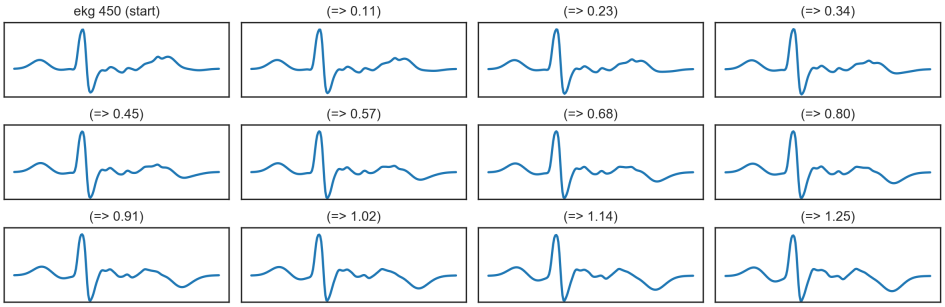

(b) Generated beats along the T wave inversion direction determined by EKG 422 to EKG 344 (above), starting from very different EKG 450. We see that this direction does correspond to T wave inversion, leaving other features relatively unchanged.

Interpolation To illustrate the generative capacity of our model, we visualize the latent path between two beats. In Figure 4a, we start out at the $z^{(m)}$ value of an observed beat, and linearly interpolate to the $z^{(m)}$ value of another observed beat. Each transition beat is generated from our model — we see that following this direction corresponds to inverting the T wave. To further explore this concept, we take a beat with a different morphology (and standard T wave) and follow the direction "T wave inversion" direction found in the previous example. Similar to neural word embeddings [Mikolov et al., 2013], we find that following this direction inverts the T wave while leaving other features of the morphology relatively unchanged, shown in Figure 4b.

Identifying Patients A subset of patients have multiple EKG records in the PTB dataset. We expect a good representation of EKG beats to cluster together the same patients across different records. To test this, we compare the average distance of $\boldsymbol{z}^{(m)}$ between the same patient across different EKG records ($d^{(same)}$) to the average distance of $\boldsymbol{z}^{(m)}$ to some random EKG record ($d^{(diff)}$). We compare this difference to the same value measured using PCA, and find that our model measures the same patient to be significantly closer. More experiment details are in Appendix B.

4 Discussion

We developed a latent variable model for large datasets of electrocardiogram records using a flexible deep neural network component and an interpretable pause and beat duration component. In preliminary model exploration, we show that the latent morphology space encodes information about patient similarity and physiological features that correlate to biological processes. We plan to further criticize our model and hope to predict patient outcomes that are less easily observed.

References

R Bousseljot, D Kreiseler, and A Schnabel. Nutzung der ekg-signaldatenbank cardiodat der ptb über das internet. *Biomedizinische Technik/Biomedical Engineering*, 40(s1):317–318, 1995.

Chih-Chun Chia and Zeeshan Syed. Scalable noise mining in long-term electrocardiographic time-series to predict death following heart attacks. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 125–134. ACM, 2014.

Diederik P Kingma and Max Welling. Auto-encoding variational Bayes. In *Proceedings of the International Conference on Learning Representations*, 2014.

D Kreiseler and R Bousseliot. Automatisierte ekg-auswertung mit hilfe der ekg-signaldatenbank cardiodat der ptb. *Biomedizinische Technik/Biomedical Engineering*, 40(s1):319–320, 1995.

Patrick E McSharry, Gari D Clifford, Lionel Tarassenko, and Leonard A Smith. A dynamical model for generating synthetic electrocardiogram signals. *IEEE transactions on biomedical engineering*, 50(3):289–294, 2003.

Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119, 2013.

Julien Oster, Joachim Behar, Omid Sayadi, Shamim Nemati, Alistair EW Johnson, and Gari D Clifford. Semisupervised ecg ventricular beat classification with novelty detection based on switching kalman filters. *IEEE Transactions on Biomedical Engineering*, 62(9):2125–2134, 2015.

A Model and Inference

Our inference procedure is two step — we first maximize the data likelihood with respect to latent variables $\boldsymbol{z}^{(p)}$, which finds the per-beat pause and EKG cardiac cycle length for each heartbeat. We do this by using the closed-form posterior mean solution for the regression coefficients with respect to the fixed basis. We can think of this as a sort of alignment procedure — the part of the waveform corresponding to the P-wave, QRS complex, and T-waves can now be jointly modeled.

We then fix this alignment, and fit the deep generative network that produces regression coefficients for the same fixed output basis. For this, we use a multi-layer perceptron with two hidden layers, each with 50 units

$$\boldsymbol{\beta} = \texttt{MLP}(\boldsymbol{z}^{(m)}; \boldsymbol{\theta}) \,. \tag{1}$$

Our model of beat morphology is essentially a deep generative regression model. Given generative parameters $\boldsymbol{\theta}$, and a fixed observation basis, $B_1(\cdot), \ldots, B_K(\cdot)$, where $B_k(\cdot) : [0, 2\pi] \mapsto \mathbb{R}$ are von Mise-like radial basis functions

$$B_k(\omega; \mu_k, \kappa_k) = \exp\left(\kappa_k \text{cos}(\omega - \mu_k) - \kappa_k\right) \,. \tag{2}$$

The data are then generated using $\boldsymbol{\beta}$ and the static basis $B_1, \ldots, B_K$ at the points where the EKG tracings are observed. For instance, if we observe samples for a single beat at time points $t_1 \ldots, t_T$, and we have inferred the start time and duration of the cardiac cycle, $z^{(p)} = (t^{(start)}, t^{(dur)})$, then the observation can be split into three segments — the pause before the cycle, the cycle itself, and the pause observed after the cycle. For observations within the cycle, we simply transform them to "canonical time" to align with the von-Mise basis

$$\tau_i = (t_i - t^{(start)})/t^{(dur)} \cdot 2\pi \,. \tag{3}$$

If there are $T^{(cycle)}$ cardiac cycle samples, then the cycle portion of each beat thus has a corresponding "design matrix" of size $T^{(cycle)} \times K$

$$\boldsymbol{X}_{i,k} = B_k(\tau_i; \mu_k, \kappa_k) \,. \tag{4}$$

Conditioned on this design matrix, the observed data are normal with a small error term

$$\boldsymbol{y}_i = \mathcal{N}(\boldsymbol{\beta}^\intercal \boldsymbol{X}_i, \sigma^2) \,. \tag{5}$$

The inferential task is to estimate the posterior distribution over $\boldsymbol{z}^{(m)}$ given observations. For this, we use an inference network, which is another multi-layer perceptron (that mirrors the generative network)

$$\mu_z, \sigma_z = \texttt{MLP}(\boldsymbol{\beta}^{(ols)}; \boldsymbol{\phi}) \tag{6}$$

which outputs an estimate of the posterior mean and variance for the latent morphology parameter. The variational objective is now defined with respect to $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$

$$\mathcal{L}(\boldsymbol{\theta}, \boldsymbol{\phi}) = \mathbb{E}_{\boldsymbol{z} \sim q(\cdot; \boldsymbol{\phi}, \boldsymbol{\beta}^{(ols)})} \left[\ln p(\boldsymbol{Y}|\boldsymbol{z}) p(\boldsymbol{z}) - \ln q(\boldsymbol{z}; \boldsymbol{\phi}, \boldsymbol{\beta}^{(ols)}) \right] \tag{7}$$

The "data" we use for the inference network is not the observation vector $\boldsymbol{Y}$, but the posterior mean solution for $\boldsymbol{\beta}$ given the fixed basis, $B_1, \ldots, B_K$. For our experiments, we fix a basis of size $K = 60$, spatially spread over the interval $[0, 2\pi]$.

B Additional Experiments

We measure the average distance of $\boldsymbol{z}^{(m)}$ between the same patient across different EKG records ($d^{(same)}$) to the average distance of $\boldsymbol{z}^{(m)}$ to some random EKG record ($d^{(diff)}$). We compare this difference to the same value measured using PCA on the least squares regression coefficients for the fixed final layer, with the same number of latent dimensions. We find that our model measures the same patient to be significantly closer on average.

For latent dimension $D = 10$, we measure the two averages as

$$\mathbb{E}[d^{(same)} - d^{(diff)}] = -.037 \in [-0.049, -0.025] \qquad \text{PCA} \tag{8}$$

$$\mathbb{E}[d^{(same)} - d^{(diff)}] = -.100 \in [-0.118, -0.083] \qquad \text{VAE} \tag{9}$$

indicating that the VAE significantly improves over PCA. We draw the conclusion that a strictly linear model of EKG beats (given our fixed observation basis) may not be expressive enough to carry patient-specific information, compared to a non-linear latent factor model.