
12-lead ECG Reconstruction via Koopman Operators

Tomer Golany¹ Daniel Freedman² Saar Minha³ Kira Radinsky¹

Abstract

32% of all global deaths in the world are caused by cardiovascular diseases. Early detection, especially for patients with ischemia or cardiac arrhythmia, is crucial. To reduce the time between symptoms onset and treatment, wearable ECG sensors were developed to allow for the recording of the full 12-lead ECG signal at home. However, if even a single lead is not correctly positioned on the body that lead becomes corrupted, making automatic diagnosis on the basis of the full signal impossible. In this work, we present a methodology to reconstruct missing or noisy leads using the theory of Koopman Operators. Given a dataset consisting of full 12-lead ECGs, we learn a dynamical system describing the evolution of the 12 individual signals together in time. The Koopman theory indicates that there exists a high-dimensional embedding space in which the operator which propagates from one time instant to the next is linear. We therefore learn both the mapping to this embedding space, as well as the corresponding linear operator. Armed with this representation, we are able to impute missing leads by solving a least squares system in the embedding space, which can be achieved efficiently due to the sparse structure of the system. We perform an empirical evaluation using 12-lead ECG signals from thousands of patients, and show that we are able to reconstruct the signals in such way that enables accurate clinical diagnosis.

1. Introduction

Cardiovascular diseases are responsible for about a third of all deaths globally (Roth et al., 2018). The electrocardiogram (ECG) is a noninvasive tool for detecting diseases of the heart, and as such is one of the most common tests performed by cardiologists. The short-duration standard 12-lead ECG is the most commonly used ECG exam in medical facilities (Maron et al., 2014). In this test, ten electrodes are placed on a patient and the overall electrical potential amplitude of the heart is then measured from twelve different angles referred to as “leads”, and is recorded over a period of time (10 seconds in the standard 12-lead ECG exam). This evaluation provides a full diagnosis of heart activity, including arrhythmia, acute coronary syndrome, ventricular dysfunction and cardiac chamber hypertrophy.

However, it is still a challenge to conveniently and robustly track 12-lead ECG in people’s daily lives. To reduce the time between the onset of symptoms and their treatment, wearable ECG sensors such as (Laguna et al., 1990) been developed to allow the recording of a 12-lead ECG at home. Accurate ECG monitoring from those devices is of high importance. For example, Atrial Fibrillation, the most common serious cardiac arrhythmia, affects an estimated 2.7-6.1 million people and increases a person’s risk of a life-changing stroke, heart failure and death. It can occur without symptoms, thus its timely detection could help physicians and their patients get an earlier confirmed diagnosis.

In order to properly rely on these sensors for clinical interpretation, each lead measurement must be well grounded. If even a single lead is not correctly positioned on the body that lead becomes corrupted, making diagnosis on the basis of the full 12-lead ECG signal impossible. To overcome this challenge, the problem of ECG reconstruction has gained considerable attention and several solutions based on machine learning have been proposed (Scherer et al., 1989; Nelwan, 2005; Atoui et al., 2004; Zhou et al., 2019). However, all of these methods assume a fixed set of pre-specified leads to have clean signals. This is a challenge with wearable devices where arbitrary leads can be corrupted. In this work, we introduce a framework which is able to reconstruct 12-lead ECG from any subset of available leads, without training a new model for each subset.

Our framework begins by learning the dynamics of 12-lead

^{*}Equal contribution ¹Technion - Israel Institute of Technology, Haifa, Israel ²Google Research ³Shamir Medical Center, Zerifin, Israel and Sackler School of Medicine, Tel-Aviv University, Israel. Correspondence to: Tomer Golany <tomer.golany@cs.technion.ac.il>, Kira Radinsky <kirar@cs.technion.ac.il>, Daniel Freedman <danielfreedman@google.com>, Saar Minha <minha.saar@gmail.com>.

ECGs using the theory of Koopman Operators (Koopman, 1931; Koopman & Neumann, 1932), a decades old theory which has recently re-emerged as a leading candidate for the systematic linear representation of nonlinear systems (Mezić & Banaszuk, 2004) (Mezić, 2005). The key aspect of Koopman theory which we leverage is its linear structure: the signal of interest can be embedded in a high-dimensional space in which the operator which propagates from one time instant to the next is linear. Learning the dynamical system is therefore equivalent to learning both the mapping to this embedding space, as well as the corresponding linear operator. Due to the linear structure, missing lead reconstruction can be posed as a least squares problem in the embedding space. Minimization leads to an explicit solution in the form of a sparse linear system, which can be solved efficiently. We emphasize again that this method of reconstruction may be applied no matter which subset of leads have been corrupted, giving it a crucial advantage over existing techniques (Zhou et al., 2019). Figure 1 presents an example of two Koopman-reconstructed leads.

We empirically evaluate our reconstruction technique in 3 separate ways: (1) We compute the reconstruction error of our algorithm, and show that it is lower than competitor techniques. (2) We learn classifiers for common classes of abnormalities, and analyze the change in performance of these classifiers as clean signals in the test are replaced with signals in which some of the leads have been reconstructed. We show that classification accuracy remains high when using signals with reconstructed leads; and this remains the case even when a large number of leads have been corrupted. (3) We perform a small clinical experiment, in which clinicians are given examples of ECG signals with missing vs. reconstructed corrupted leads. We demonstrate that our reconstruction improves clinicians’ diagnosis capabilities.

Our contributions in this work are threefold. (1) We present a methodology to learn 12-lead ECG dynamics using Koopman operators, which are represented by deep neural nets. We learn a separable representation of the Koopman embedding functions that can be applied on each ECG lead separately. (2) We introduce a least squares system which is able to impute missing leads efficiently from any partial sub-leads ECG. We share the code for the reproducibility of our results ¹ (3) We empirically show that our method is able to reconstruct any partial-lead ECG signal to a 12-lead ECG without hurting clinical diagnosis. This is demonstrated by empirical experiments showing increased performance of both clinicians and state-of-the-art deep-learning models to identify ECG abnormalities using the reconstructed data.

¹Link anonymized




Figure 1. Examples of reconstructed leads. Blue: real signal. Red: reconstructed signal by Koopman framework.

2. Related Work

12-Lead ECG Reconstruction The first attempt to reconstruct 12-lead ECG from a subset of leads was introduced by (Frank, 1956). Later, classical machine learning methods were proposed using simple linear regression techniques (Scherer et al., 1989; Nelwan, 2005). An early method which used neural networks for the purposes of lead reconstruction is presented in (Atoui et al., 2004). More recent methods based on CNNs (Zhou et al., 2019) and LSTMs (Zhang & Frick, 2019) have successfully reconstructed 9-lead ECG from the 3-lead ECG. All prior works assume that specific indices of leads are recorded cleanly, and attempt to reconstruct the remaining leads. For example, it might be assumed that leads V1 and V2 are clean, and the remaining 10 leads require reconstruction. However, each 12-lead ECG recording coming from a wearable device might have a different set of leads that are cleanly recorded. Therefore, a model which expects a specific subset of leads might fail to reconstruct the full 12-lead ECG from such devices. By contrast, our framework is able to reconstruct 12-lead ECG from any subset of available leads, without training a new model for each subset.

Learning ECG Dynamics of a Single Lead Formulating the dynamics as a system of differential equations often admits compact and efficient representations for many natural systems (Brunton et al., 2016). This holds true in the case of single-lead ECG signals, one-dimensional signals of voltage values representing the electrical activity of the heart through time. The ECG signal is a periodic signal of cardiac muscle depolarization followed by repolarization, with each period corresponding to a single heartbeat. An ECG heartbeat follows a prototypical pattern of a P wave, followed by a QRS complex, and finally a T wave. To capture this pattern, (McSharry et al., 2003) proposed a physics-based model of ECG dynamics consisting of a system of three coupled ordinary differential equations (ODE), parameterized by specific heart rate statistics, such as the frequency-domain characteristics of the heart rate variability (Malik & Camm, 1990). While this model is able to generate synthetic ECG signals with somewhat realistic PQRST morphology as well as prescribed heart rate dynamics, it has

limited expressiveness. A more recent work (Golany et al., 2020) introduced a GAN-based setup enriched with additional knowledge from this physics-based ECG model, and showed that using the synthetically generated ECG heartbeats from the GAN significantly improved ECG heartbeat classification. Others (Golany et al., 2021) attempted to learn a new set of ODEs from data rather than relying on predefined set of ODEs to represent the dynamics of a single ECG heartbeat. This prior work that learns data-driven ECG Dynamics attempts to capture the dynamics of a single ECG heartbeat within a single lead. By contrast, we focus on the dynamics of an entire ECG signal, consisting of multiple heartbeats, with all 12 leads. The data and the corresponding modelling problem are concomitantly more complex.

Koopman Theory The original Koopman theory was introduced nearly one hundred years ago (Koopman, 1931; Koopman & Neumann, 1932). Renewed interest in Koopman analysis has been driven by a combination of theoretical advances (Mezić & Banaszuk, 2004) (Mezić, 2005) (Budišić et al., 2012) (Mezić, 2013), improved numerical methods such as dynamic mode decomposition (Schmid, 2010) (Rowley et al., 2009), and an increasing abundance of data. Recently, (Lusch et al., 2018) utilized the power of deep learning for flexible and general representations of the Koopman framework, while enforcing a network structure that promotes parsimony and interpretability of the resulting models. Although it was applied on small scale toy problems, such as pendulum motion prediction (Erichson et al., 2019; Pan & Duraisamy, 2020), to the best of our knowledge, it was yet to be applied in a large-scale machine learning application.

3. Koopman-Based ECG Reconstruction

3.1. Koopman Theory of Dynamical Systems

Throughout this paper, we will consider discrete-time dynamical systems of the form

$$x_{t+1} = F(x_t) \quad (1)$$

where $x_t \in \mathcal{X} \equiv \mathbb{R}^L$ is the state of the dynamical system and F represents the nonlinear transformation (the dynamics) which maps the state of the system to its future state. Note that this formulation subsumes discretizations of ordinary differential equations (ODEs). That is, suppose that the underlying continuous signal is given by $x(\tau)$ for $\tau \in [0, \bar{\tau}]$; and the dynamics is described by the ODE $dx/d\tau = f(x)$. Then the signal may be discretized as $x_t = x(t\Delta)$ for $t = 0, \dots, T$ with $\Delta = \bar{\tau}/T$, and the dynamics approximated as $x_{t+1} \approx x_t + f(x_t)\Delta \equiv F(x_t)$. The approximation becomes increasingly exact as Δ gets smaller.

(Koopman, 1931) offers a different and useful viewpoint for examining dynamical systems. In particular, rather than consider the state space x , Koopman considers the space of possible measurements on x . A measurement on x is defined as a scalar-valued function on the state space \mathcal{X} , that is

$$y : \mathcal{X} \rightarrow \mathbb{R} \quad (2)$$

The space of all measurements is denoted as \mathcal{Y} , which is an infinite-dimensional space. For a dynamical system of the form in Equation (1) given by dynamics F , we define the corresponding Koopman operator which maps from measurements to measurements, $\mathcal{K} : \mathcal{Y} \rightarrow \mathcal{Y}$ by

$$\mathcal{K}y \equiv y \circ F \quad (3)$$

where \circ indicates function composition. (Note that $y \circ F$ is indeed a measurement, as it maps \mathcal{X} to \mathbb{R} .) In this case, the dynamical system of Equation (1) can be rewritten as

$$\begin{aligned} y(x_{t+1}) &= y(F(x_t)) = y \circ F(x_t) \\ &= (\mathcal{K}y)(x_t) \end{aligned} \quad (4)$$

Thus, if a measurement y evolves forward with the operator \mathcal{K} , then it satisfies the ‘‘pullback’’ property given in Equation (4). However, what makes the formulation most interesting is the fact that the Koopman operator \mathcal{K} is *linear*. This fact is easily shown:

$$\begin{aligned} \mathcal{K}(\alpha_1 y_1 + \alpha_2 y_2) &= (\alpha_1 y_1 + \alpha_2 y_2) \circ F \\ &= \alpha_1 y_1 \circ F + \alpha_2 y_2 \circ F \\ &= \alpha_1 \mathcal{K}y_1 + \alpha_2 \mathcal{K}y_2 \end{aligned}$$

The linearity of the Koopman operator is crucial to the development of our method, as we shall see in Section 3.3.

3.2. Learning a Koopman Representation for ECG Dynamics

In this section, we adapt the Koopman framework to learn the dynamics of 12-lead ECG signals. We begin by describing two necessary modifications to the Koopman theory, after which we show how to learn the dynamical system.

We begin with some notation. The number of ECG leads is denoted as $L = 12$. The standard 12-lead electrocardiogram is a representation of the heart’s electrical activity recorded from electrodes on the body surface, sampled at a fixed frequency. The ℓ^{th} lead sampled at time t is denoted by x_t^ℓ ; all L leads taken together at time t are denoted $x_t = [x_t^1, \dots, x_t^L] \in \mathbb{R}^L$, taken to be a column vector.

Finite-Dimensional Approximation The first modification we must make to the standard Koopman theory concerns dimensionality. The space \mathcal{Y} of measurements is infinite-dimensional and the Koopman operator \mathcal{K} is likewise an infinite-dimensional operator. For computational

purposes, we approximate the entire Koopman framework by mapping in into a finite-dimensional setting. In particular, suppose that

$$\Gamma : \mathbb{R}^D \rightarrow \mathcal{Y} \quad (5)$$

maps a finite-dimensional space to the space of measurements. (For concreteness, the reader may imagine mapping the coefficients of a basis expansion to the function y itself, though we will not use this representation.) In this case, we will approximate the Koopman operator \mathcal{K} by

$$\mathcal{K} = \Gamma K \Gamma^{-1} \quad (6)$$

where K is a $D \times D$ matrix.

In this case, we can rewrite the dynamical system in Equation (4) as

$$\begin{aligned} y(x_{t+1}) &= \Gamma K \Gamma^{-1} y(x_t) \\ \Rightarrow \Gamma^{-1} y(x_{t+1}) &= K \Gamma^{-1} y(x_t) \end{aligned} \quad (7)$$

Now, letting $\Phi = \Gamma^{-1} y$ so that $\Phi : \mathbb{R}^L \rightarrow \mathbb{R}^D$, we have that

$$x_{t+1} = \Phi^{-1}(K \Phi(x_t)) \quad (8)$$

This modification is standard, and follows the practice of prior works, e.g. (Lusch et al., 2018). We refer to Φ as the *Koopman embedding*.

Note that we use K rather than \mathcal{K} to emphasize this move to a finite-dimensional framework, but we abuse notation slightly by continuing to use the symbol y to represent its finite-dimensional version, i.e.

$$y = \Phi(x) \quad (9)$$

In this case Equation (8) may be rewritten as

$$\begin{aligned} y_t &= \Phi(x_t) \\ y_{t+1} &= K y_t \\ x_{t+1} &= \Phi^{-1}(y_{t+1}) \end{aligned} \quad (10)$$

which illustrates the fact that in the embedding space, the dynamics are linear.

Separable Koopman Embedding We make a second modification to the standard Koopman theory, which is necessary for our reconstruction algorithm. We assume that the Koopman embedding is *separable*: that is, each lead has its own separate embedding. More specifically, we map the ℓ^{th} lead x_t^ℓ to its corresponding embedding y_t^ℓ as follows:

$$y_t^\ell = \phi_\ell(x_t^\ell) \quad (11)$$

where $\phi_\ell : \mathbb{R} \rightarrow \mathbb{R}^{D/L}$. The overall Koopman embedding Φ is then derived by concatenating the per-lead embeddings:

$$y_t = [y_t^1, \dots, y_t^L] \in \mathbb{R}^D \quad (12)$$

so that

$$\Phi(x_t) = [\phi_1(x_t^1), \dots, \phi_L(x_t^L)] \quad (13)$$

The importance of separability to the reconstruction algorithm will become clear in Section 3.3. We note that separability is not guaranteed by the Koopman theory; nevertheless, there is nothing which prevents us from imposing it as a constraint during our learning procedure. In spite of this lack of theoretical guarantees, we show empirically in Section 5 that separability does not impair the learning of an accurate dynamical system. In this context, we also note that due to the separable structure, all of the coupling between the leads is encapsulated by the matrix K .

Learning the Dynamical System Given the above Koopman framework, learning the dynamical system entails learning two things: the Koopman embedding Φ , and the Koopman operator K . A variety of methods have been proposed for learning the Koopman framework based on neural networks (Wehmeyer & Noé, 2018; Mardt et al., 2018; Takeishi et al., 2017; Yeung et al., 2019). We choose to follow the technique of (Lusch et al., 2018) and outline this method briefly.

A multilayer perceptron (MLP) specifies the Koopman embedding Φ ; in our case, we impose the separable structure on the embedding, so that the network’s structure is tantamount to L separate MLPs $\{\phi_\ell\}_{\ell=1}^L$. An additional MLP is learned to represent the inverse transformation Φ^{-1} , which is again tantamount to learning L separate MLPs $\{\phi_\ell^{-1}\}_{\ell=1}^L$. The Koopman operator is simply a $D \times D$ matrix K . To learn the networks Φ and Φ^{-1} and matrix K , three separate losses are used:

- (1) Reconstruction: $\|x_t - \Phi^{-1}(\Phi(x_t))\|$
 - (2) Linear Dynamics: $\|\Phi(x_{t+m}) - K^m \Phi(x_t)\|$, $m \geq 1$
 - (3) State Prediction: $\|x_{t+m} - \Phi^{-1}(K^m \Phi(x_t))\|$, $m \geq 1$
- Further details, including values of m to use, are described in (Lusch et al., 2018).

We note that in practice, we have found that learning a single per-lead embedding ϕ which is the same for all leads is sufficient, i.e. $\phi_\ell = \phi$ for all ℓ . However, this is not necessary for the reconstruction algorithm described next, so we leave the derivation there in the general setting.

3.3. Reconstruction of Missing Leads

We now turn to our main goal: the reconstruction of corrupted 12-lead ECG signals. As we have already outlined, the corruption may be due to either missing leads or noisy values, a frequent scenario when measuring ECG from wearable sensors such as Holter monitoring (DiMarco & Philbrick, 1990) and ECG patches (Steinhubl et al., 2018). The reconstruction will rely on the Koopman-based dynamical system we have learned.

Setup The set of missing leads is denoted $\mathcal{M} \subset \{1, \dots, L\}$; our goal is therefore to reconstruct $\{x_t^\ell\}_{t=0}^T$ for each missing lead $\ell \in \mathcal{M}$. The set of available leads is just the complement of the set of missing leads $\mathcal{A} = \{1, \dots, L\} - \mathcal{M}$, which has corresponding indicator vector

$$a_\ell = \begin{cases} 1 & \ell \in \mathcal{A} \\ 0 & \ell \notin \mathcal{A} \end{cases} \quad (14)$$

Step 1: Mapping Available Leads to their Koopman Embeddings We begin by mapping the available leads to their corresponding Koopman embeddings. The available leads are given by $\{\bar{x}_t^\ell\}_{t=0}^T$ for each $\ell \in \mathcal{A}$; we therefore let

$$\bar{y}_t^\ell = \begin{cases} \phi_\ell(\bar{x}_t^\ell) & \ell \in \mathcal{A} \\ 0 & \ell \notin \mathcal{A} \end{cases} \quad (15)$$

Missing values have been filled in with zeros for convenience, so that the overall Koopman embeddings have the correct size, i.e. $\bar{y}_t \in \mathbb{R}^D$; however, the missing entries can take on any values, as they will not be used.

Step 2: Reconstructing the Missing Leads in Embedding Space Given the available leads' Koopman embeddings, we can now solve for the missing leads by leveraging the fact that the Koopman operator is linear. For convenience, we let

$$A = \text{diag}(a \otimes \mathbf{1}_{D/L}) \quad (16)$$

where \otimes is the Kronecker product. In this case, we can formulate our leads reconstruction problem as one of solving the following optimization problem:

$$\begin{aligned} \min_{y_0, \dots, y_T} L(y_0, \dots, y_T) = \\ \frac{1}{2} \sum_{t=0}^{T-1} \|y_{t+1} - Ky_t\|^2 + \frac{\lambda}{2} \sum_{t=0}^T (y_t - \bar{y}_t)^T A (y_t - \bar{y}_t) \end{aligned} \quad (17)$$

The first term ensures that the dynamical system holds at each time instant; crucially, due to the linearity of the Koopman formulation of the dynamics, this can be formulated nicely as a convex quadratic term. The second term is a data fidelity term for the available leads only: the matrix A picks out only the available leads. $\lambda > 0$ is the weighting factor between the two terms, where a larger λ ensures greater consistency to the given leads. In the limit as $\lambda \rightarrow \infty$, we have a hard constraint.

Due to the fact that L is convex, we can solve for the globally optimal values of y . Furthermore, L is quadratic, giving us an explicit solution. Specifically, let

$$C = K^T K + I + \lambda A; \quad (18)$$

then the solution is given by

$$\begin{aligned} -K^T y_{t+1} + (C - I)y_t &= \lambda A \bar{y}_t & t = 1 \\ -K^T y_{t+1} + C y_t - K y_{t-1} &= \lambda A \bar{y}_t & t \in [2, T-1] \\ (C - K^T K)y_t - K y_{t-1} &= \lambda A \bar{y}_t & t = T \end{aligned} \quad (19)$$

The above is a system of linear equations, and furthermore is quite sparse. As a result, the solution can be achieved efficiently using standard methods. In this work we leverage least squares method (Levenberg, 1944) to solve these equations.

Step 3: Mapping the Missing Leads Back to Signal Space Finally, given the optimal values y_t^ℓ from the solution to Equation (19), we can map back to signal space. This is achieved by applying the inverse of the separable Koopman embedding function:

$$x_t^\ell = \phi_\ell^{-1}(y_t^\ell) \quad \text{for } \ell \in \mathcal{M} \quad (20)$$

This yields the final reconstruction of the missing ECG leads. We note in passing that it is also possible to compute a reconstruction of the ECG signals for the available leads $\ell \in \mathcal{A}$; if the data fidelity weight $\lambda \rightarrow \infty$, it is straightforward to show that these will precisely replicate the data, i.e. $x_t^\ell = \bar{x}_t^\ell$ for all t and $\ell \in \mathcal{A}$.

Comparison with Seq2Seq We draw the reader's attention to a key distinction between our method and the commonly used seq2seq-style techniques for signal reconstruction applied for ECG reconstruction (Zhou et al., 2019). The seq2seq techniques require learning a separate model from each different subset of available leads; by contrast, the methodology presented learns a single model, which can be easily applied with equal ease to any subset. More specifically, in the seq2seq setting, learning to map from lead 1 to lead 2 is different from lead 1 to lead 3, or leads 1 and 7 to the rest. In our formulation, they may all be reconstructed directly from the ECG signal's master dynamical system.

4. Experimental Framework

4.1. ECG Dataset

The Georgia 12-lead ECG dataset, referred to as G12EC, was introduced in the 12-lead ECG Physionet Challenge 2020 (Alday et al., 2020) and is considered one of the largest public 12-lead ECG datasets. It represents a large population from the southeastern United States and contains 10,344 12-lead ECGs (male: 5,551, female: 4,793). Each ECG signal is 10 seconds in length with a sampling frequency of 500 Hz, yielding a total of 5,000 time samples per signal.

Each 12-lead ECG exam is annotated with 27 diagnoses. These 27 classes represent relatively common diagnoses which are of clinical interest, with the potential to be recognizable from ECG recordings. Note that the classes are

not mutually exclusive: each 12-lead ECG exam may hold multiple diagnoses. In our experiments we focus on the following six common types of diagnosis: AF - Atrial fibrillation; TAb - T wave abnormal; QAb - Q wave abnormal; VPB - Ventricular premature beats; LAD - Left axis deviation, and SA - Sinus arrhythmia. Our dataset is divided as follows: the train set contains 8,233 ECG signals, while the test set contains the remaining 2,059 signals.

4.2. Baselines

We compared our reconstruction model with the state-of-the-art (SOTA) model for 12-lead reconstruction. (Zhou et al., 2019) proposed a seq2seq approach using a CNN-based model for reconstruction of short 12-lead ECG segments from a 3-lead ECGs. We extend this approach and build a model for each n available leads. That is, given n leads the model reconstructs the 12-lead ECG. Note that the model receives *any* n leads and reconstructs the missing k leads.

4.3. Experimental Setup

We train the baselines and our model (Section 3) on the training set of G12EC. To mimic a partial 12-lead ECG reading (as often occurs in a home setting when using a wearable), we remove $k \in \{1, 4, 8, 11\}$ random leads from each 12-lead ECG recording in the test-set. Each test instance represents a random subsample of $12 - k$ leads. For example, for $k = 4$ we might remove leads 1, 2, 3, and 4 from one recording, leads 4, 7, 9 and 11 from another. The resulting test-set contains ECG signals of shape $\mathbb{R}^{T \times n}$, where $n = 12 - k$ is the number of leads left in each signal. On the resulting test-set we apply the baselines and solve the system of linear equations described in Section 3 to reconstruct the missing leads.

We perform experiments showing the performance of reconstruction via two types of experiments:

1. Reconstruction Error: Measuring the **distance** between the reconstructed lead and the corresponding ground truth lead (Section 5.1).

2. Classification Accuracy: Measuring **clinical diagnosis** based on the reconstructed leads.

- We perform a small clinical experiment with clinicians (Section 5.3). They received 52 12-lead ECG reading from the test (where k leads are reconstructed) and are asked to make a diagnosis. This diagnosis is compared to the ground truth diagnosis.
- To perform a larger experiment, we leverage the state-of-the-art machine-learning model for 12-lead ECG classification (Attia et al., 2019; Ribeiro et al., 2020) and measure its performance on reconstructed leads (Section 5.2). The model is trained on G12EC training

set, and we report its performance over the test set, where each test set contains reconstructed leads. We compare the classifier diagnosis with the ground-truth diagnosis. We next describe the architecture of the machine-learning model (Section 4.4).

4.4. Classification Network Details

Recently (Ribeiro et al., 2020) and (Attia et al., 2019) showed superior results for classification of ECG abnormalities from 12-lead ECG signals. They trained a Residual Neural Network (He et al., 2016) based architecture. We follow this practice and use in our experiments a Residual Neural Network model. The input to the model is a 10 seconds 12-lead ECG signal sampled at 500Hz. That is, input of shape $\mathbb{R}^{5000 \times 12}$, where the first dimension represents the temporal dimension and the second dimension represents the spatial dimension. The network consists of a convolution layer, followed by a max pooling layer, followed by six residual blocks. Each residual block consists of 3 convolution layers, and between each convolution layer, Batch-normalization and Relu activation are performed. A skip connection is applied between the input of the block to the output of the third convolution layer. The output of the last residual block is fed into a global average pooling layer, followed by a dense layer. Since multiple abnormalities may occur in the same 12-lead ECG signal (classes are not mutually exclusive), the last activation function we use is a Sigmoid function which gives a separate probability score for each predicted abnormal class. The first convolution layer has 16 filters of size 7×7 . The residual blocks start with 16 filters and are increased to 32 filters in the last block. The size of the kernel in the residual blocks starts in 5×5 , and decreases to 3×3 . In all the residual blocks, except the first one, the first convolution layer down-samples the input temporal dimension by a stride of 2. The neural network weights were initialized as in (He et al., 2016), and the bias was initialized with zeros. The network was trained by feeding 12-lead ECG batches of size 128 from the training data. The binary cross-entropy loss was minimized using Adam Optimizer with initial learning rate 0.0001. The training ran for 100 epochs, with the final model being the one with the best accuracy on the validation set.

5. Experimental Results

5.1. Leads Reconstruction Performance

We first present the results of Koopman-based ECG reconstruction. We measure the distance of the reconstructed 12-lead ECG signal \hat{x}_t^ℓ to the ground truth signal x_t^ℓ . We report our results by the Mean Absolute Deviation (MAD) error function:

$$\text{MAD} = \frac{1}{|\mathcal{M}|T} \sum_{\ell \in \mathcal{M}} \sum_t |\hat{x}_t^\ell - x_t^\ell| \quad (21)$$

Table 1. Evaluation of the SOTA ECG Classifier (Section 4) on reconstructed 12-lead ECG testset. Results are shown for different number of reconstructed leads both for Koopman-reconstruction and baseline-reconstruction.

ABNORMAL CLASS	KOOPMAN BASED RECONSTRUCTION								BASELINE (ZHOU ET AL., 2019)					
	RECALL (SENSITIVITY)				SPECIFICITY				RECALL (SENSITIVITY)			SPECIFICITY		
	12-LEAD	11-LEAD	8-LEAD	4-LEAD	12-LEAD	11-LEAD	8-LEAD	4-LEAD	11-LEAD	8-LEAD	4-LEAD	11-LEAD	8-LEAD	4-LEAD
AF	0.91	0.91	0.90	0.90	0.85	0.85	0.72	0.80	0.75	0.76	0.79	0.65	0.65	0.62
TAB	0.85	0.85	0.83	0.81	0.77	0.77	0.70	0.70	0.60	0.61	0.56	0.60	0.55	0.52
QAB	0.85	0.87	0.82	0.78	0.70	0.70	0.66	0.62	0.83	0.73	0.57	0.40	0.52	0.47
VPB	0.77	0.76	0.79	0.77	0.56	0.59	0.58	0.67	0.89	0.85	0.81	0.20	0.30	0.37
SA	0.66	0.68	0.68	0.62	0.56	0.50	0.55	0.56	0.50	0.64	0.46	0.47	0.57	0.40
LAD	0.94	0.95	0.88	0.81	0.87	0.90	0.80	0.70	0.62	0.61	0.55	0.50	0.55	0.47

where \mathcal{M} is the set of missing leads.

Table 2 shows the reconstruction results as a function of the number of missing leads. We note, that as expected as the number of missing leads in the corrupted signal increases, the reconstruction error increases for both the baseline the Koopman-based reconstruction. While our Koopman-based method is better in all cases than the baseline, it is considerably better when there are 10 missing leads, i.e. when most of the information is absent.

MISSING LEADS	MAD KOOPMAN	MAD BASELINE
1	0.130	0.134
4	0.135	0.137
8	0.138	0.139
10	0.142	0.196

Table 2. Mean Absolute Deviation (MAD) error between the reconstructed ECG leads and the ground truth. In Bold are statistically significant results. Lower numbers indicate better reconstruction.

5.2. ECG Classification using Reconstructed Leads

In this section, we compare the performance of the SOTA ECG classifier when applied on 12-lead ECGs where some of the leads are reconstructed. We experiment on several number of reconstructed leads (k).

Comparison to SOTA ECG Reconstruction Figures 2(a)-(f) show the ROC curves of each of the six classified diagnoses (Sec. 4.1). For each diagnosis we compared the results of the 12-lead ECG classifier evaluated on a different reconstructed test-set. The purple curve, the blue curve and the red curve, corresponds to a corrupted test reconstructed via our methods using Koopman operators (Sec. 3.3), with valid 11-leads, 8-leads, and 4-leads respectively. The green, pink, and brown curves in each subfigure corresponds to a corrupted test reconstructed by the CNN-based methods of (Zhou et al., 2019), with a valid 11-leads, 8-leads, and 4-leads respectively. Sensitivity and Specificity metrics are also reported in Table 1. Our reconstruction method outperforms the state-of-the-art method with respect to the ROC evaluation metric for each number of corrupted leads and precision-recall points in a statistically significant manner (t-test with p-value < 0.05). We observe that for all type of diagnosis, our method is better than the CNN-based recon-

struction. This emphasizes the ability of our method to learn to reconstruct any subset of ECG leads to 12-lead ECG.

Comparison to Complete 12-Lead ECG We notice that when comparing to the gold standard – classification using 12-Lead ECG with no missing leads – we see a very small loss in performance. This indicates that ECG classifiers can be considered for automated classification of ECGs from devices with smaller amount of leads than 12 leads, reconstructed using our method and yet reaching similar performance of full 12-lead devices.

5.3. Clinician’s Diagnosis Performance using Reconstructed Leads

We perform a small clinical experiment. We choose to focus on the T wave abnormality (TAB), as abnormalities of this form are associated with several life-threatening diseases. The electrocardiographic T wave represents ventricular repolarization and are usually hard to identify without the V1 and L leads. We randomly selected 52 ECGs from the test set where 38% had an abnormal T wave. We mimic a situation where the V1 and L leads are corrupted. For each example, we showed the cardiologist the 10 non-corrupted leads and asked to make a diagnosis of whether the patient exhibits TAB. We then showed the additional 2 leads (the V1 and L leads) which were reconstructed using our Koopman framework and asked the cardiologist to make the diagnosis again. Table 3 summarizes the results. Our methodology enabled the cardiologist to identify all of the patients with TAB abnormalities. Notice that without the reconstructed leads, only by observing the non-corrupted leads, the cardiologist identified only 60% of the patients with TAB. We observe a loss in precision (though marginal compared to the recall improvement) and points to the fact that additional

	Recall	Precision	F_1
Cardiologist using 10 leads	0.6	0.75	0.67
Cardiologist using 10 leads + Koopman-reconstructed 2 leads	1.0	0.63	0.77

Table 3. Clinical experiment results for the TAB abnormality. Each line presents the the diagnosis accuracy of the clinician. The first represents the performance results given no reconstructed leads whereas the second with reconstructed leads.

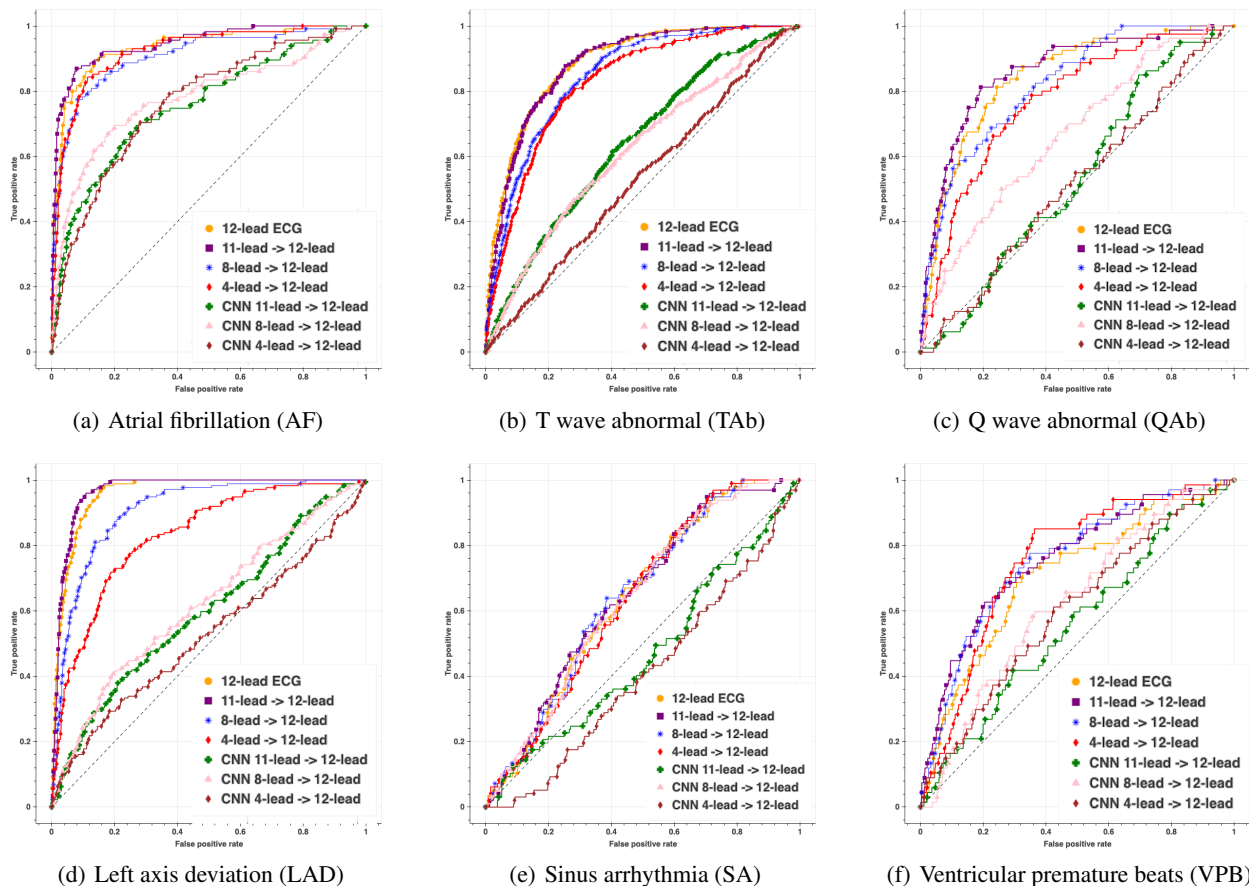


Figure 2. ROC curves of the 6 diagnosis classes evaluated on the test-set. The orange curve at each subfigure corresponds to the results on the complete 12-lead test-set. The other curves correspond to a corrupted 12-lead test which was reconstructed either by our approach via Koopman operators (Section 3.3) or by the baseline (Zhou et al., 2019).

training on using computer-generated ECGs is needed and should be further explored. Overall, the F_1 score with the reconstruction is considerably higher than without.

6. Conclusions

To reduce the time between cardiac symptoms onset and treatment, wearable ECG sensors were developed to allow for the recording of the full 12-lead ECG signal at home. To rely on such sensors for clinical interpretation, each lead measurement must be well grounded. However, it is enough for one lead not to be well-positioned on the body for the entire lead signal to be corrupt. This has prevented the wider usage of those sensors from home. In this work, we presented a methodology to reconstruct missing or noisy leads using the theory of Koopman Operators. To the best of our knowledge, this is one of the first applications of this theory for a large-scale machine-learning real-life application. We learn the dynamical system describing the evolution of the 12 individual signals together in time. Koopman theory

allows us a linear structure: the signal of interest can be embedded in a high-dimensional space in which the operator which propagates from one time instant to the next is linear. Learning the dynamical system is therefore equivalent to learning both the mapping to this embedding space, as well as the corresponding linear operator and then solving a least squares system in the embedding space. An additional key benefit of this system is its ability to reconstruct any number of corrupted leads without the need to retrain a machine learning model. We empirically show that our reconstruction error is rather small and that classifiers trained on 12-leads ECGs perform well in the presence of reconstructed leads. A small-scale clinical experiment shows the value of presenting the reconstructed leads to a clinician during diagnosis. The results are staggering – the recall of a severe abnormality rises from 60% to 100% with a tolerable number of false positives. For future work, we plan to expand the clinical trial and to better understand how to best present the reconstructed leads to humans for better benefit of diagnosis.

References

- Alday, E. A. P., Gu, A., Shah, A. J., Robichaux, C., Wong, A.-K. I., Liu, C., Liu, F., Rad, A. B., Elola, A., Seyedi, S., et al. Classification of 12-lead ecgs: the physionet/computing in cardiology challenge 2020. *Physiological measurement*, 41(12):124003, 2020.
- Atoui, H., Fayn, J., and Rubel, P. A neural network approach for patient-specific 12-lead ecg synthesis in patient monitoring environments. In *Computers in Cardiology, 2004*, pp. 161–164. IEEE, 2004.
- Attia, Z. I., Kapa, S., Lopez-Jimenez, F., McKie, P. M., Ladewig, D. J., Satam, G., Pellikka, P. A., Enriquez-Sarano, M., Noseworthy, P. A., Munger, T. M., et al. Screening for cardiac contractile dysfunction using an artificial intelligence-enabled electrocardiogram. *Nature medicine*, 25(1):70–74, 2019.
- Brunton, S. L., Proctor, J. L., and Kutz, J. N. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the national academy of sciences*, 113(15):3932–3937, 2016.
- Budišić, M., Mohr, R., and Mezić, I. Applied koopmanism. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 22(4):047510, 2012.
- DiMarco, J. P. and Philbrick, J. T. Use of ambulatory electrocardiographic (holter) monitoring. *Annals of internal medicine*, 113(1):53–68, 1990.
- Erichson, N. B., Muehlebach, M., and Mahoney, M. W. Physics-informed autoencoders for lyapunov-stable fluid flow prediction. *arXiv preprint arXiv:1905.10866*, 2019.
- Frank, E. An accurate, clinically practical system for spatial vectorcardiography. *circulation*, 13(5):737–749, 1956.
- Golany, T., Freedman, D., and Radinsky, K. SimGANs: Simulator-based generative adversarial networks for ECG synthesis to improve deep ECG classification. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2020.
- Golany, T., Freedman, D., and Radinsky, K. ECG ODE-GAN: Learning ordinary differential equations of ECG dynamics via generative adversarial learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2021.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Koopman, B. and Neumann, J. v. Dynamical systems of continuous spectra. *Proceedings of the National Academy of Sciences of the United States of America*, 18(3):255, 1932.
- Koopman, B. O. Hamiltonian systems and transformation in hilbert space. *Proceedings of the national academy of sciences of the united states of America*, 17(5):315, 1931.
- Laguna, P., Thakor, N. V., Caminal, P., Jane, R., Yoon, H.-R., De Luna, A. B., Marti, V., and Guindo, J. New algorithm for qt interval analysis in 24-hour holter ecg: performance and applications. *Medical and Biological Engineering and Computing*, 28(1):67–73, 1990.
- Levenberg, K. A method for the solution of certain nonlinear problems in least squares. *Quarterly of applied mathematics*, 2(2):164–168, 1944.
- Lusch, B., Kutz, J. N., and Brunton, S. L. Deep learning for universal linear embeddings of nonlinear dynamics. *Nature communications*, 9(1):1–10, 2018.
- Malik, M. and Camm, A. J. Heart rate variability. *Clinical cardiology*, 13(8):570–576, 1990.
- Mardt, A., Pasquali, L., Wu, H., and Noé, F. Vampnets for deep learning of molecular kinetics. *Nature communications*, 9(1):1–11, 2018.
- Maron, B. J., Friedman, R. A., Kligfield, P., Levine, B. D., Viskin, S., Chaitman, B. R., Okin, P. M., Saul, J. P., Salberg, L., Van Hare, G. F., et al. Assessment of the 12-lead ecg as a screening test for detection of cardiovascular disease in healthy general populations of young people (12–25 years of age) a scientific statement from the american heart association and the american college of cardiology. *Circulation*, 130(15):1303–1334, 2014.
- McSharry, P. E., Clifford, G. D., Tarassenko, L., and Smith, L. A. A dynamical model for generating synthetic electrocardiogram signals. *IEEE transactions on biomedical engineering*, 50(3):289–294, 2003.
- Mezić, I. Spectral properties of dynamical systems, model reduction and decompositions. *Nonlinear Dynamics*, 41(1):309–325, 2005.
- Mezić, I. Analysis of fluid flows via spectral properties of the koopman operator. *Annual Review of Fluid Mechanics*, 45:357–378, 2013.
- Mezić, I. and Banaszuk, A. Comparison of systems with complex behavior. *Physica D: Nonlinear Phenomena*, 197(1-2):101–133, 2004.
- Nelwan, S. *Evaluation of 12-Lead Electrocardiogram Reconstruction Methods for Patient Monitoring*. PhD thesis, Erasmus University Rotterdam, 2005.

- Pan, S. and Duraisamy, K. Physics-informed probabilistic learning of linear embeddings of nonlinear dynamics with guaranteed stability. *SIAM Journal on Applied Dynamical Systems*, 19(1):480–509, 2020.
- Ribeiro, A. H., Ribeiro, M. H., Paixão, G. M., Oliveira, D. M., Gomes, P. R., Canazart, J. A., Ferreira, M. P., Andersson, C. R., Macfarlane, P. W., Meira Jr, W., et al. Automatic diagnosis of the 12-lead ecg using a deep neural network. *Nature communications*, 11(1):1–9, 2020.
- Roth, G. A., Abate, D., Abate, K. H., Abay, S. M., Abbafati, C., Abbasi, N., Abbastabar, H., Abd-Allah, F., Abdela, J., Abdelalim, A., et al. Global, regional, and national age-sex-specific mortality for 282 causes of death in 195 countries and territories, 1980–2017: a systematic analysis for the global burden of disease study 2017. *The Lancet*, 392(10159):1736–1788, 2018.
- Rowley, C. W., MEZI?, I., Bagheri, S., Schlatter, P., Henningson, D., et al. Spectral analysis of nonlinear flows. *Journal of fluid mechanics*, 641(1):115–127, 2009.
- Scherer, J. A., Jenkins, J. M., and Nicklas, J. M. Synthesis of the 12-lead electrocardiogram from a 3-lead subset using patient-specific transformation vectors: an algorithmic approach to computerized signal synthesis. *Journal of electrocardiology*, 22, 1989.
- Schmid, P. J. Dynamic mode decomposition of numerical and experimental data. *Journal of fluid mechanics*, 656: 5–28, 2010.
- Steinhubl, S. R., Waalen, J., Edwards, A. M., Ariniello, L. M., Mehta, R. R., Ebner, G. S., Carter, C., Baca-Motes, K., Felicione, E., Sarich, T., et al. Effect of a home-based wearable continuous ecg monitoring patch on detection of undiagnosed atrial fibrillation: the mstops randomized clinical trial. *Jama*, 320(2):146–155, 2018.
- Takeishi, N., Kawahara, Y., and Yairi, T. Learning koopman invariant subspaces for dynamic mode decomposition. *arXiv preprint arXiv:1710.04340*, 2017.
- Wehmeyer, C. and Noé, F. Time-lagged autoencoders: Deep learning of slow collective variables for molecular kinetics. *The Journal of chemical physics*, 148(24):241703, 2018.
- Yeung, E., Kundu, S., and Hodas, N. Learning deep neural network representations for koopman operators of nonlinear dynamical systems. In *2019 American Control Conference (ACC)*, pp. 4832–4839. IEEE, 2019.
- Zhang, Q. and Frick, K. All-ecg: A least-number of leads ecg monitor for standard 12-lead ecg tracking during motion. In *2019 IEEE Healthcare Innovations and Point of Care Technologies (HI-POCT)*, pp. 103–106. IEEE, 2019.
- Zhou, W., Xing, Y., Liu, N., Movahedipour, M., Zhou, X.-g., et al. 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*, 20(3):405–413, 2019.