Detecting Malignant Ventricular Arrhythmias in Electrocardiograms by Gaussian Process Classification

Kimmo Suotsalo and Simo Särkkä

This article has been accepted for publication in *Proceedings of 27th IEEE International Workshop on Machine Learning for Signal Processing (MLSP)*, to be held in Tokyo, Japan, on September 25–28, 2017. When citing this article, use the following reference:

> K. Suotsalo and S. Särkkä, "Detecting Malignant Ventricular Arrhythmias in Electrocardiograms by Gaussian Process Classification," in *27th IEEE International Workshop on Machine Learning for Signal Processing (MLSP)*, 2017. Accepted for publication.

2017 IEEE INTERNATIONAL WORKSHOP ON MACHINE LEARNING FOR SIGNAL PROCESSING, SEPT. 25–28, 2017, TOKYO, JAPAN

DETECTING MALIGNANT VENTRICULAR ARRHYTHMIAS IN ELECTROCARDIOGRAMS BY GAUSSIAN PROCESS CLASSIFICATION

Kimmo Suotsalo and Simo Särkkä

Aalto University
Department of Electrical Engineering and Automation
Rakentajanaukio 2, 02150 Espoo, Finland

ABSTRACT

Ventricular tachycardia, ventricular flutter, and ventricular fibrillation are malignant forms of cardiac arrhythmias, whose occurrence may be a life-threatening event. Several methods exist for detecting these arrhythmias in the electrocardiogram. However, the use of Gaussian process classifiers in this context has not been reported in the current literature. In comparison to the popular support vector machines, Gaussian processes have the advantage of being fully probabilistic, they can be re-casted in Bayesian filtering compatible state-space form, and they can be flexibly combined with first-principles physical models. In this paper we use Gaussian process classification to detect malignant ventricular arrhythmias in the electrocardiogram. We describe how Gaussian process classifiers can be used to solve the detection problem, and show that the proposed classifiers achieve a performance that is comparable to that of the state-of-the-art methods henceforth laying down promising foundations for more general electrocardiogram-based arrhythmia detection framework.

Index Terms— Electrocardiography, cardiac arrhythmias, classification algorithms, Gaussian processes.

1. INTRODUCTION

Ventricular arrhythmias (VAs) account for a significant portion of sudden cardiac deaths and heart disorders [1]. Malignant forms of VAs include ventricular tachycardia (VT), ventricular flutter (VFL), and ventricular fibrillation (VF), all of which are potentially life-threatening events. It is crucial to detect the presence of these arrhythmias promptly so that the normal sinus rhythm can be restored by applying a defibrillation shock to the patient [2].

Several studies have addressed the problem of detecting malignant VAs in the electrocardiogram (ECG) [3–5]. Machine learning methods, such as neural networks and support vector machines (SVMs) [6], have also been proposed [7, 8]. A comparison of 11 different techniques is presented in [9], where the SVM is found to yield the best results.

In this article, the aim is to use Gaussian processes (GPs) [10] for malignant VA detection. In ECG analysis, GP classification has previously been proposed for premature ventricular contraction (PVC) detection [11–13]. However, there seems to be no research on the use of GP classification in detecting malignant VAs. This can be seen as a gap in the current research, since GPs provide many advantages: (1) GPs are fully probabilistic models allowing for the use of sophisticated hierarchical Bayesian models and methods [10, 14], (2) in the case of time series and especially in the case of periodic signals, GPs can be re-casted as state-space models allowing for efficient Bayesian filtering solution methods [15–17], and (3) GPs can be flexibly combined with first-principles physical models [18]. State-space Gaussian process methods have been previously applied to cardiac signal processing, for example, in [15, 19, 20].

The authors would like to thank the Academy of Finland for funding.

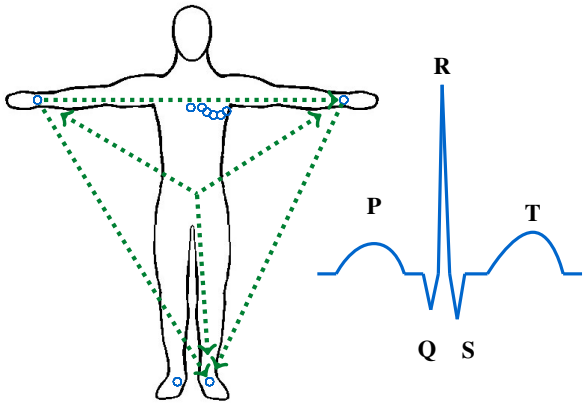

Fig. 1. The principle of electrocardiography. Potential differences between electrodes are measured and presented as one or more signals, each of which corresponds to a specific channel or *lead*. During the normal sinus rhythm, a single heartbeat usually appears in the ECG as a sequence of five waves: the P, Q, R, S, and T wave. The number and location of the electrodes vary but typically follow some standard configuration.

The specific research problem of this paper is to detect VT, VFL, and VF episodes in single-channel ECG recordings that may also contain other types of arrhythmias, abnormal beats, and noise. The novelty of this paper lies in the combination of the clinical problem and the classification method. Previous research [11–13] does consider GP classification in ECG analysis, but the addressed problem as well as the dataset are clearly distinct from ours. In [11–13], the problem is to classify previously detected individual heartbeats into normal or ventricular, whereas our approach classifies 10 second signal segments into those that contain VT, VF, or VFL, and those that do not. It should be noted that such a signal segment may contain several types of beats and that there is no need to detect the individual beats nor their boundaries prior to the malignant VA analysis. Also, the use cases are very different, since infrequent PVCs do not

978-1-5090-6341-3/17/$31.00

typically require treatment while a malignant VA may quickly lead to cardiac arrest.

The contribution of this paper is to investigate the use of GP classification in malignant VA detection and to show that GPs achieve a performance comparable to that of the SVM. From the results we can conclude that GPs provide a promising framework for malignant VA detection and, due to their extensibility, are a better alternative to this task than SVMs.

The organization of the paper is as follows. The remaining of this section provides an introduction to electrocardiography, cardiac arrhythmias, and GP classification. Section 2 describes the materials and methods, and Section 3 presents the experimental results. Section 4 is a conclusion with a short summary and future insights.

1.1. Electrocardiography and cardiac arrhythmias

Electrocardiography is a non-invasive technique for examining the heart. A number of electrodes are placed on the body surface according to one of the standard configurations. The potential differences between the electrodes are measured and converted into digital signals. The principle is illustrated in Fig. 1.

A cardiac arrhythmia is any heart rhythm other than the normal sinus rhythm [21]. Cardiac arrhythmias can be characterized as supraventricular, such as sinus tachycardia, atrial fibrillation, and atrial flutter; or ventricular, such as PVCs, VT, VFL, and VF. It is important to note that while some arrhythmias pose a serious threat of cardiac arrest, others are typically benign and require little or no treatment.

Malignant types of VAs include VT, VFL, and VF. In VT, the heart beats rapidly at a rate of 120 bpm or more, which may result in insufficient blood supply [21]. In VFL, the heart rate is even higher, as much as 300 bpm [22]. Both VT and VFL may turn into VF [21,22], which is a life-threatening condition and requires immediate resuscitation [22].

All the three types of malignant VAs can usually be seen in the ECG as somewhat sinusoidal waveforms. This is in contrast with the normal sinus rhythm, where each heartbeat typically appears as a sequence of the P, Q, R, S, and T wave. Fig. 2 shows three examples of the normal sinus rhythm turning into a malignant VA.

It should be noted that the differences between the normal sinus rhythm and malignant VAs are not always as clear as in Fig. 2. ECG signals vary considerably between individuals, and may contain noise and artifacts. The signals may also encompass other, less critical arrhythmias such as atrial fibrillation or ventricular bigeminy.

1.2. Gaussian process classification

GPs provide an infinite-dimensional generalization of Gaussian probability distributions. Their use in machine learning covers both regression and classification. A comprehensive review of GPs is given in [10]. The purpose of this section is to describe the central ideas of GP classification in the present context.

A binary classification problem can be described as follows. Let $\mathbf{x}_i \in \mathbb{R}^{d\times 1}$ denote an input, that is, a vector representation of an object that belongs to either class $\mathcal{C}_1$ or class $\mathcal{C}_2$. Let y_i be the corresponding output, and let $y_i = +1$ if the object belongs to $\mathcal{C}_1$; otherwise, $y_i = -1$. Given a training set $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$ and a test input $\mathbf{x}_*$, the problem is to predict the test output y_*.

In GP classification [10], the outputs y_i are modeled to depend on the inputs $\mathbf{x}_i$ through the probability distribution $p(y_i = +1 \mid \mathbf{x}_i) = \sigma(f(\mathbf{x}_i))$, where $f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'))$ is a GP with mean function m and covariance (or kernel) function k; these are defined as $m(\mathbf{x}) \triangleq \mathbb{E}\{f(\mathbf{x})\}$ and $k(\mathbf{x}, \mathbf{x}') \triangleq \mathbb{E}\{[f(\mathbf{x}) - m(\mathbf{x})][f(\mathbf{x}') - m(\mathbf{x}')]\}$, respectively.

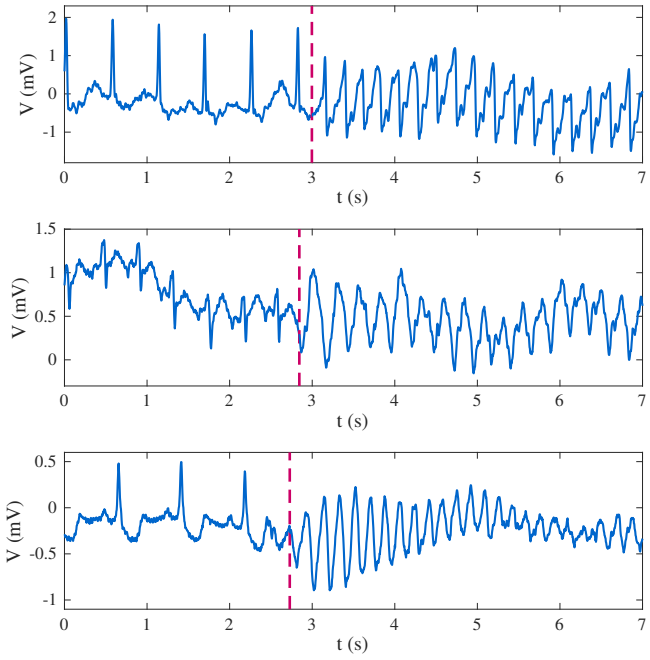

Fig. 2. Examples of VAs in the MIT-BIH Malignant Ventricular Arrhythmia Database: VT (top), VFL (middle), and VF (bottom). The transitions from the normal sinus rhythm to a VA are marked with dashed vertical lines.

One possible choice for σ is the probit function that satisfies

$$\sigma(f(\mathbf{x}_i)) = \int_{-\infty}^{f(\mathbf{x}_i)} \mathcal{N}(x \mid 0, 1)dx. \tag{1}$$

The function f is latent in the sense that its values cannot be observed directly.

Common choices for covariance functions include the linear kernel k_{linear}, the constant kernel k_{const}, the squared exponential kernel k_{sqexp}, and the Matérn kernel $k_{\text{Matérn}}$. These can be defined (see [10, 23]) as

$$k_{\text{linear}}(\mathbf{x}, \mathbf{x}') = \mathbf{x}^\top \boldsymbol{\Sigma}_d \mathbf{x}', \tag{2}$$

$$k_{\text{const}}(\mathbf{x}, \mathbf{x}') = \sigma_{\text{c}}^2, \tag{3}$$

$$k_{\text{sqexp}}(\mathbf{x}, \mathbf{x}') = \sigma_{\text{s}}^2 \exp\left(-\frac{1}{2}r^2\right), \tag{4}$$

$$k_{\text{Matérn}}(\mathbf{x}, \mathbf{x}') = \sigma_{\text{m}}^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\sqrt{2\nu} r\right)^\nu K_\nu \left(\sqrt{2\nu} r\right), \tag{5}$$

where

$$r = \sqrt{\sum_{k=1}^{d} \left(\frac{x_k - x'_k}{\ell_k}\right)^2}, \tag{6}$$

$\boldsymbol{\Sigma}_d = \text{diag}(\sigma_1^2, \ldots, \sigma_d^2)$, $\sigma_{\text{c}}^2, \sigma_{\text{s}}^2, \sigma_{\text{m}}^2, \sigma_1^2, \ldots, \sigma_d^2, \ell_k, \nu$ are constant parameters, and K_ν is a modified Bessel function. Different kernels can also be combined, for example, by summing them together.

The problem of predicting y_* is solved by assessing the conditional probability $\pi(\mathbf{x}_*) \triangleq p(y_* = +1 \mid \mathcal{D}, \mathbf{x}_*)$. This probability

can then be used to predict the output by assigning $y_* = +1$ if and only if $\pi(\mathbf{x}_*) \geq p_{\text{thr}}$. The threshold p_{thr} can be chosen between 0 and 1 to reflect the desired level of confidence.

The conditional probability $\pi(\mathbf{x}_*)$ can be written as

$$p(y_* = +1 \mid \mathcal{D}, \mathbf{x}_*) = \int \sigma(f_*) p(f_* \mid \mathcal{D}, \mathbf{x}_*) df_* \tag{7}$$

where $f_* \triangleq f(\mathbf{x}_*)$. Denoting $\mathbf{X} = [\mathbf{x}_1 \cdots \mathbf{x}_N]$, $\mathbf{y} = [y_1 \cdots y_N]^\top$, and $\mathbf{f} = [f_1 \cdots f_N]^\top$, the second term in the integrand satisfies

$$p(f_* \mid \mathcal{D}, \mathbf{x}_*) = \int p(f_* \mid \mathbf{X}, \mathbf{x}_*, \mathbf{f}) p(\mathbf{f} \mid \mathbf{X}, \mathbf{y}) d\mathbf{f}; \tag{8}$$

this is the distribution of the latent variable f_* given the training data and the test input. By the Bayes' rule,

$$p(\mathbf{f} \mid \mathbf{X}, \mathbf{y}) \propto p(\mathbf{y} \mid \mathbf{f}) p(\mathbf{f} \mid \mathbf{X}), \tag{9}$$

where the conditioning on $\mathbf{X}$ has been dropped as unnecessary in the likelihood $p(\mathbf{y} \mid \mathbf{f})$.

The integrals in Eqn. (7) and Eqn. (8) or the normalization constant of the above distribution cannot usually be evaluated in closed form. Instead, one has to use numerical approximations, such as the Laplace method or the expectation propagation algorithm [10, 23]. In this article we have used the Laplace method as implemented in the GPstuff toolbox [23].

In practice, we also have unknown hyperparameters $\boldsymbol{\theta}$ in the covariance function. In that case we need to consider the joint posterior distribution of the GP and its parameters, which is given by

$$p(\mathbf{f}, \boldsymbol{\theta} \mid \mathbf{X}, \mathbf{y}) \propto p(\mathbf{y} \mid \mathbf{f}) p(\mathbf{f} \mid \mathbf{X}, \boldsymbol{\theta}) p(\boldsymbol{\theta}), \tag{10}$$

and need to perform an additional integral over the hyperparameters:

$$p(\mathbf{f} \mid \mathbf{X}, \mathbf{y}) = \int p(\mathbf{f}, \boldsymbol{\theta} \mid \mathbf{X}, \mathbf{y}) d\boldsymbol{\theta}. \tag{11}$$

This integral cannot be computed in closed form either and in this paper we have chosen to use the central composite design (CCD) method from the GPstuff toolbox [23] for its numerical computation.

In the context of malignant VA detection, the binary classification framework can be applied as follows. The training signals are concatenated and divided into N segments. The inputs $\mathbf{x}_i$ are the segments' vector representations, and the elements of $\mathbf{x}_i$ are features computed for each segment. The output $y_i = +1$ if the ith segment contains VT, VFL, or VF; otherwise, $y_i = -1$.

2. MATERIALS AND METHODS

2.1. Datasets

The proposed method was tested on the MIT-BIH Malignant Ventricular Arrhythmia Database (VFDB) [24, 25]. This database consists of 22 ECG records, each of which comprises two separate channels. The length of each record is 35 minutes. The records have been annotated for rhythm changes, that is, the beginnings of each rhythm. Individual beats have not been annotated.

The records contain episodes of different kinds of arrhythmia, such as atrial fibrillation, ventricular bigeminy, and sinus bradycardia. Importantly, each record contains one or more episodes of VT, VFL, or VF. There are also episodes of asystole, noise, paced rhythm, and other events that cause the signal to deviate from the normal sinus rhythm.

2.2. Feature extraction

The preprocessing of the ECG recordings followed the procedure described in [9]. More specifically, the signals were filtered with a cascade of a high-pass, a low-pass, and a notch filter, with cut-off frequencies of 1 Hz, 30 Hz, and 59/61 Hz, respectively. The purpose of the filters was to remove baseline wander, noise, and power line interference.

As a part of the preprocessing, the signals were split into 10 second segments, and the true class of each segment was determined using the annotations in the database. A segment was assigned to the positive class (with VA) if it contained at least 1 second of ventricular arrhythmia; otherwise, the segment was assigned to the negative class.

As in [9], the preprocessing step also comprised the rejection of signal segments that contained asystole or noise. The decision was made in a similar fashion as in assigning the segments to positive and negative classes: a segment was rejected if it contained at least 5 seconds of asystole or at least 5 seconds of noise.

The second step in the experiments was to compute the features. The features were the same as those used in [9], with one additional feature from [26]. This amounts to a total of 15 features.

The 14 features presented in [9] are: *Complexity*, *Leakage*, *FSMN*, *A1*, *A2*, *A3*, *TimeDelay*, *Count1*, *Count2*, *Count3*, *CovarBin*, *FreqBin*, *AreaBin*, and *Kurtosis*. The reader is referred to [9] for the definitions of these features.

The 15th feature, *spectral purity index (SPI)*, was originally developed for EEG signal analysis [27], and is discussed further in [28] and [26]. The idea of SPI is to quantify the bandwidth of the signal by a single number in the interval $[0, 1]$: for narrowband signals, SPI is close to unity; for wideband signals, SPI is close to zero. SPI is defined as

$$\text{SPI}(n) = \frac{m_2^2(n)}{m_0(n) m_4(n)}, \tag{12}$$

where m_0, m_2, and m_4 are the zeroth, second, and fourth spectral moment, respectively. The spectral moments can be estimated directly in the time domain using the procedure described in [28].

After the features were computed for each signal segment, they were normalized to the interval [0, 1] by applying a linear transformation.

2.3. Classifiers

The experiments were made using the MATLAB software. The GP classifiers were implemented with the GPstuff toolbox [23]. A total of eight classifiers were included in the experiments. Three of the classifiers were SVMs, and the remaining five were GP classifiers.

The three SVMs used a linear kernel, a polynomial kernel, and a squared exponential kernel, respectively. Henceforth, the corresponding classifiers are denoted as $\text{SVM}_{\text{linear}}$, SVM_{poly}, and $\text{SVM}_{\text{sqexp}}$. The hyperparameters – box constraint, kernel scale, and polynomial order (for SVM_{poly}) – were found by optimization.

The GP classifiers utilized the following kernels: linear, squared exponential, linear + constant + squared exponential, linear + constant + Matérn with $\nu = 3/2$, and linear + constant + Matérn with $\nu = 5/2$. From now on, these classifiers are denoted as $\text{GP}_{\text{linear}}$, GP_{sqexp}, GP_{lcs}, GP_{lcm32}, and GP_{lcm52}, respectively. The probit function was used as the squashing function σ. The numerical approximations were carried out using the Laplace method. The parameters were integrated out numerically using the CCD method.

2.4. Cross-validation

The ECG records in the VFDB database were split into separate training and validation sets. The division was similar to the one used in [9]: the records were first sorted by their names and then placed alternately in the training and validation sets. This way, both sets came to comprise a total of 11 records.

The training set with 11 records was used to optimize the classifiers and to compare their performance. A leave-one-out cross-validation was performed between the records, giving rise to a total of 11 validation folds. The 10 training records in each fold were used to optimize the classifiers.

To compare the classifiers' performances, the following measures were computed for each validation fold: accuracy (Acc), sensitivity (Se), specificity (Sp), receiver operating characteristic (ROC) curve, and area under ROC curve (AUC). These measures were then averaged over the 11 folds, giving estimates of the corresponding means. For Acc, Se, and Sp, the standard deviation over the 11 folds was also computed.

2.5. Classification of records not used in training

Based on the cross-validation results, $\text{SVM}_{\text{sqexp}}$ and GP_{lcm32} were selected to demonstrate the performance of SVMs and GPs classifiers on the 11 records not used in training. These two classifiers were trained once more on the whole training set of 11 records, and then used to classify the signal segments in each of the 11 records in the validation set.

3. RESULTS

The cross-validation results are listed in Table 1. Columns 2 to 5 present the average Acc, Se, Sp, and AUC, respectively. The ROC curves are shown in Fig. 3.

Table 1. Cross-validation results

Classifier	Acc	Se	Sp	AUC
$\text{SVM}_{\text{linear}}$	0.88	0.83	0.88	0.95
SVM_{poly}	0.90	0.92	0.88	0.95
$\text{SVM}_{\text{sqexp}}$	0.90	**0.93**	0.87	0.96
$\text{GP}_{\text{linear}}$	0.88	0.84	0.88	0.95
GP_{sqexp}	0.90	0.89	0.87	0.96
GP_{lcs}	0.90	0.89	0.87	0.96
GP_{lcm32}	**0.91**	0.89	0.88	0.96
GP_{lcm52}	0.90	0.89	0.87	0.95

The average standard deviations of Acc, Se, and Sp are 0.06, 0.04, and 0.09, respectively.

The cross-validation results show no clear difference between the performances of the SVMs and the GP classifiers. While the sensitivity of $\text{SVM}_{\text{sqexp}}$ is somewhat higher than the sensitivities of the GP classifiers, GP_{lcm32} yields the highest accuracy. Nonetheless, these differences are statistically insignificant, as implied by the standard deviations reported in Table 1.

The classification results for $\text{SVM}_{\text{sqexp}}$ and GP_{lcm32} on the validation set are listed in Table 2. These results suggest that GP classifiers can obtain a VA detection accuracy of ca. 0.96 on signals that have not been used in training. This is comparable with the results listed in [9]; there, however, all methods but the proposed one were tested on the training set. Using a separate validation set provides a more realistic estimate of the classifiers' performance.

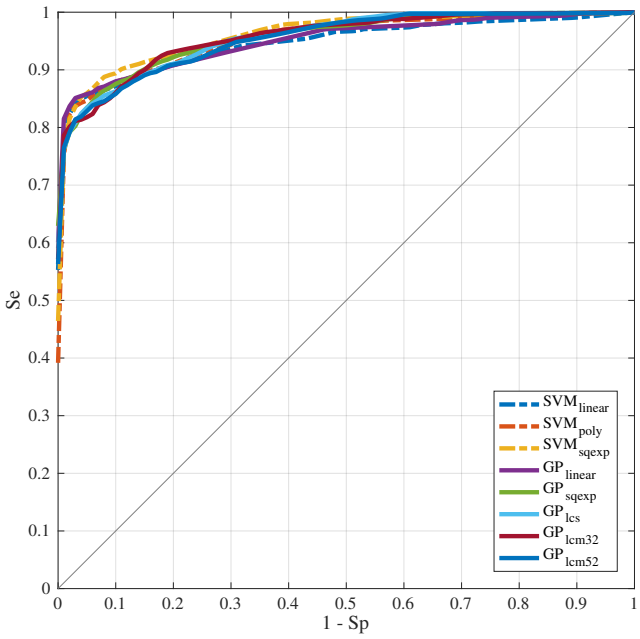

Fig. 3. ROC curves for the different classifiers, obtained by cross-validation.

Table 2. Classification results on the validation set

Record	$\text{SVM}_{\text{sqexp}}$			GP_{lcm32}		
	Acc	Se	Sp	Acc	Se	Sp
419	0.79	0.71	0.82	**0.85**	**0.82**	**0.86**
421	0.92	**0.72**	1.00	0.92	0.69	1.00
423	1.00	0.99	1.00	1.00	0.99	1.00
425	0.94	1.00	0.94	**0.98**	1.00	**0.97**
427	0.96	0.93	1.00	0.96	**0.94**	1.00
429	0.98	0.92	**0.99**	0.98	**1.00**	0.98
602	**0.92**	**0.67**	0.99	0.90	0.58	0.99
607	0.95	0.83	0.96	**0.99**	**0.89**	**0.99**
610	0.97	0.50	0.97	**0.98**	0.50	**0.98**
612	0.98	0.94	0.99	**0.99**	**0.95**	0.99
615	0.97	0.19	1.00	**0.98**	**0.38**	1.00
Average	0.94	0.76	0.97	**0.96**	**0.79**	**0.98**

4. CONCLUSION

In this paper we have proposed to use GP classification to detect malignant VAs in the ECG. We have described how GP classifiers can be used to solve the detection problem. Moreover, we have shown that GP classifiers achieve a performance that is comparable to that of the SVM. One of the major advantages of GPs is the possibil-

ity to reformulate them as Bayesian filtering and smoothing problems [15, 16], which leads to huge computational benefits and the possibility for online algorithms and hybrid first-principles models. In the future our aim is to explore this path further.

5. REFERENCES

[1] R. M. John, U. B. Tedrow, B. A. Koplan, C. M. Albert, L. M. Epstein, M. O. Sweeney, A. L. Miller, G. F. Michaud, and W. G. Stevenson, "Ventricular arrhythmias and sudden cardiac death," *The Lancet*, vol. 380, no. 9852, pp. 1520 – 1529, 2012.

[2] M. S. Link, D. L. Atkins, R. S. Passman, H. R. Halperin, R. A. Samson, R. D. White, M. T. Cudnik, M. D. Berg, P. J. Kudenchuk, and R. E. Kerber, "Part 6: Electrical therapies: Automated external defibrillators, defibrillation, cardioversion, and pacing. 2010 American Heart Association guidelines for cardiopulmonary resuscitation and emergency cardiovascular care.," *Circulation*, vol. 122, no. 18 suppl 3, pp. S706–S719, 2010.

[3] N. V. Thakor, Y.-S. Zhu, and K.-Y. Pan, "Ventricular tachycardia and fibrillation detection by a sequential hypothesis testing algorithm," *IEEE Transactions on Biomedical Engineering*, vol. 37, no. 9, pp. 837–843, 1990.

[4] A. Amann, R. Tratnig, and K. Unterkofler, "Detecting ventricular fibrillation by time-delay methods," *IEEE Transactions on Biomedical Engineering*, vol. 54, no. 1, pp. 174–177, 2007.

[5] I. Jekova, "Shock advisory tool: Detection of life-threatening cardiac arrhythmias and shock success prediction by means of a common parameter set," *Biomedical Signal Processing and Control*, vol. 2, no. 1, pp. 25 – 33, 2007.

[6] C. Cortes and V. Vapnik, "Support-vector networks," *Machine Learning*, vol. 20, no. 3, pp. 273–297, 1995.

[7] K. Minami, H. Nakajima, and T. Toyoshima, "Real-time discrimination of ventricular tachyarrhythmia with Fourier-transform neural network," *IEEE Transactions on Biomedical Engineering*, vol. 46, no. 2, pp. 179–185, 1999.

[8] F. Alonso-Atienza, E. Morgado, L. Fernández-Martínez, A. García-Alberola, and J. L. Rojo-Álvarez, "Detection of life-threatening arrhythmias using feature selection and support vector machines," *IEEE Transactions on Biomedical Engineering*, vol. 61, no. 3, pp. 832–840, 2014.

[9] Q. Li, C. Rajagopalan, and G. D. Clifford, "Ventricular fibrillation and tachycardia classification using a machine learning approach," *IEEE Transactions on Biomedical Engineering*, vol. 61, no. 6, pp. 1607–1613, 2014.

[10] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*, The MIT Press, 2006.

[11] G. Skolidis, R. H. Clayton, and G. Sanguinetti, "Automatic classification of arrhythmic beats using Gaussian processes," in *2008 Computers in Cardiology*, 2008, vol. 35, pp. 921–924.

[12] Y. Bazi, H. Hichri, N. Alajlan, and N. Ammour, "Premature ventricular contraction arrhythmia detection and classification with Gaussian process and S transform," in *2013 Fifth International Conference on Computational Intelligence, Communication Systems and Networks*, 2013, pp. 36–41.

[13] N. Alajlan, Y. Bazi, F. Melgani, S. Malek, and M. A. Bencherif, "Detection of premature ventricular contraction arrhythmias in electrocardiogram signals with kernel methods," *Signal, Image and Video Processing*, vol. 8, no. 5, pp. 931–942, 2014.

[14] A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin, *Bayesian Data Analysis*, Chapman and Hall/CRC, 3rd edition, 2013.

[15] S. Särkkä, A. Solin, and J. Hartikainen, "Spatiotemporal learning via infinite-dimensional Bayesian filtering and smoothing," *IEEE Signal Processing Magazine*, vol. 30, no. 4, pp. 51–61, 2013.

[16] S. Särkkä, *Bayesian filtering and smoothing*, Cambridge University Press, 2013.

[17] A. Svensson, A. Solin, S. Särkkä, and T. B. Schön, "Computationally efficient Bayesian learning of Gaussian process state space models," *JMLR Workshop and Conference Proceedings Volume 51 (AISTATS 2016)*, 2016.

[18] M. A. Álvarez, D. Luengo, and N. D. Lawrence, "Linear latent force models using Gaussian processes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 11, pp. 2693–2705, 2013.

[19] S. Särkkä, A. Solin, A. Nummenmaa, A. Vehtari, T. Auranen, S. Vanni, and F.-H. Lin, "Dynamic retrospective filtering of physiological noise in BOLD fMRI: DRIFTER," *NeuroImage*, vol. 60, no. 2, pp. 1517–1527, 2012.

[20] A. Solin and S. Särkkä, "Infinite-dimensional Bayesian filtering for detection of quasiperiodic phenomena in spatiotemporal data," *Physical Review E*, vol. 88, no. 5, pp. 052909, 2013.

[21] W. F. Boron and E. L. Boulpaep, *Medical Physiology: A Cellular and Molecular Approach*, Elsevier Saunders, 2nd edition, 2012.

[22] A. Bayés de Luna, *Clinical Arrhythmology*, John Wiley & Sons, first edition, 2011.

[23] J. Vanhatalo, J. Riihimäki, J. Hartikainen, P. Jylänki, V. Tolvanen, and A. Vehtari, "GPstuff: Bayesian modeling with Gaussian processes," *Journal of Machine Learning Research*, vol. 14, pp. 1175–1179, 2013.

[24] S. D. Greenwald, "The development and analysis of a ventricular fibrillation detector," M.S. thesis, Massachusetts Institute of Technology, 1986.

[25] A. L. Goldberger, L. A. N. Amaral, L. Glass, J. M. Hausdorff, P. Ch. Ivanov, R. G. Mark, J. E. Mietus, G. B. Moody, C.-K. Peng, and H. E. Stanley, "PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals," *Circulation*, vol. 101, no. 23, pp. e215–e220, 2000.

[26] S. Fallet, S. Yazdani, and J.-M. Vesin, "False arrhythmia alarms reduction in the intensive care unit: a multimodal approach," *Physiological Measurement*, vol. 37, no. 8, pp. 1217–1232, 2016.

[27] I. I. Goncharova and J. S. Barlow, "Changes in EEG mean frequency and spectral purity during spontaneous alpha blocking," *Electroencephalography and Clinical Neurophysiology*, vol. 76, no. 3, pp. 197–204, 1990.

[28] L. Sörnmo and P. Laguna, *Bioelectrical Signal Processing in Cardiac and Neurological Applications*, Elsevier Academic Press, 2005.