

Improving the transition modelling in hidden Markov models for ECG segmentation

Benoît Frénay, Gaël de Lannoy* and Michel Verleysen

Université catholique de Louvain - Machine Learning Group
EPL/ELEC/DICE - Place du Levant, 3 1348 Louvain-la-Neuve, Belgium

Abstract. The segmentation of ECG signal is a useful tool for the diagnosis of cardiac diseases. However, the state-of-the-art methods use hidden Markov models which do not adequately model the transitions between successive waves. This paper uses two methods which attempt to overcome this limitation: a HMM state scission scheme which prevents ingoing and outgoing transitions in the middle of the waves and a bayesian network where the transitions are emission-dependent. Experiments show that both methods improve the results on pathological ECG signals.

1 Introduction

Physicians often use the electrocardiogram (ECG) signal to detect anomalies in the heart dynamics of their patients. The ECG signal is a record of the heart electrical activity which consists of successive beats made up of waves and baselines. The delays between key points of these waves are of particular interest [1] for the physicians to diagnose cardiac diseases such as the long QT syndrome or the torsades de pointes. Therefore, an important problem called ECG segmentation consists of delineating the ECG signal.

Most of the state of the art techniques [2, 3] rely on the use of the wavelet transform [4, 5, 6] and hidden Markov models (HMMs, see [7]). Rather good results are obtained on both normal and pathological ECG signals, but the fact that transitions between successive waves are inherently difficult to model with HMMs suggests that more complex models could lead to even better results. A first solution proposed in [8] is to use hidden semi-Markov models (HSMs, see [9]) which include a modelling of the wave duration. However, HSMs still fail to model adequately the transitions for abnormal signals when the wave lengths are subject to important variations. Therefore, we propose to use two methods which attempt to lift the limitations of both HMMs and HSMs: (i) an HMM state scission scheme which prevents ingoing and outgoing transitions in the middle of the waves and (ii) a bayesian network where the transitions are emission-dependent. We show experimentaly that both methods improve the segmentation quality on pathological signals.

2 State of the art in ECG analysis

FIG. 1 shows two annotated beats extracted from a real ECG signal. The segmentation has been carried out by an expert; we can see that the beat consists

*Gaël de Lannoy is funded by a Belgian F.R.I.A. grant.

of three waves (P wave, QRS complex and T wave) separated by three baselines (B1, B2 and B3). In practice, algorithms do not work directly on this signal : it has been shown easier [2] to use a high-dimensional representation obtained by differentiating the signal or applying a wavelet transform [4, 5, 6].

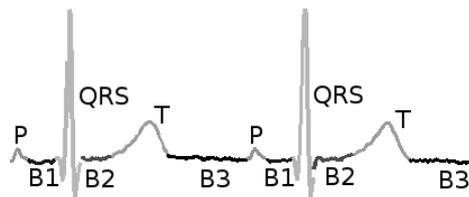


Figure 1: Two annotated beats extracted from a real ECG signal.

In order to model an ECG signal, we can use HMMs in which the states S are the waves and baselines, and the emissions O are the voltage levels of the data. The HMMs make several assumptions about the data: (i) each voltage measure is an emission O_t whose probability only depends on the state S_t (or wave) we are in and (ii) the transition probabilities $P(S_{t+1}|S_t)$ between states only depend on the starting and ending states. The HMM architecture used in this paper is shown on FIG. 2; the arrows represent transitions between states and the small squares are the emission models.

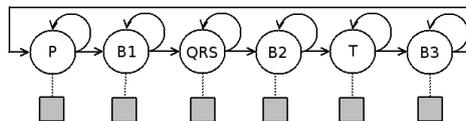


Figure 2: HMM architecture for ECG segmentation.

The computation of the optimal state sequence for a given HMM and a given emission sequence can be easily and efficiently carried out using the Viterbi algorithm [10]. In fact, the main computational costs come from the estimation of the emission probabilities. In the ECG segmentation case, it occurs in two steps. Firstly, the ECG signal is mapped to a high-dimensional representation using e.g. a wavelet transform which computes the convolution of the signal and a wavelet at different scales. Secondly, a PDF estimator is built in the new space from the training data.

A major limitation of HMMs is their poor modelling of transitions. Indeed, in state $S_t = a$, the transition probability $P(S_{t+1} = b|S_t = a)$ to switch to state $S_{t+1} = b$ does neither depend on the current value of the ECG signal nor on the time already spent in a . Therefore, even if transitions outgoing from a given wave W only occur on the downward slope of W , the Viterbi algorithm could leave W on its rising slope if it better fits the data.

A first solution to get rid of these limitations is to include an explicit modelling of the state duration by using HSMMs. In HSMMs, the transition probabilities depend on the starting and ending states and on the duration spent in the starting state. [8] has shown that using HSMMs increases the segmentation accuracy, but it is still inadequate. For example, if a very short wave occurs, HSMMs will tend to extend its downward slope over the next baseline.

In the two following sections, we will see how we can improve HMMs to better model transitions and enhance the segmentation of pathological ECG signals.

3 SME decomposition

In ECG signals, transitions between states only occur at the beginning and at the end of waves and baselines. Therefore, a simple idea is to split each state s corresponding to a wave or baseline W into the three new states s_s , s_m and s_e which correspond to the starting, middle and ending parts of W , respectively. For each observed beat, s_s (s_e) corresponds to the w first (last) points of W whereas s_m correspond to the in-between points, where w is called the *window width*. FIG. 3 shows the resulting HMM architecture.

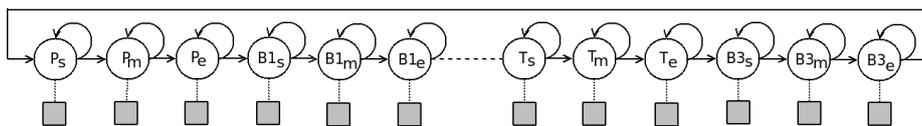


Figure 3: HMM architecture for ECG segmentation with SME decomposition.

The advantage of using a HMM with this start-middle-end (SME) decomposition is that transitions can no longer occur anywhere. Indeed, the only way to begin (end) W is now to enter s_s (leave s_e), what is only possible if the current emission is probable enough w.r.t. the emission model in s_s (s_e). Moreover, the middle of the wave corresponds to state s_m which can be only reached from s_s and leaved toward s_e . Notice that previous works (see e.g. [5]) used the Baum-Welch algorithm to discover several substates for each wave; however this last solution is unsupervised and does not aim to better model the transitions between states.

FIG. 4(a) and FIG. 4(b) show the emission probabilities of a T wave and of a B3 baseline obtained after applying a wavelet transform with a Coiflet wavelet and two dyadic scales. We can see that the probability distributions for s_s and s_e are distinct and localised compared to the distribution for s_m . Therefore, we can expect that a HMM using the SME decomposition will easily discriminate these three states. Leaving the middle part of a wave without passing through the end part will then be hardly achievable.

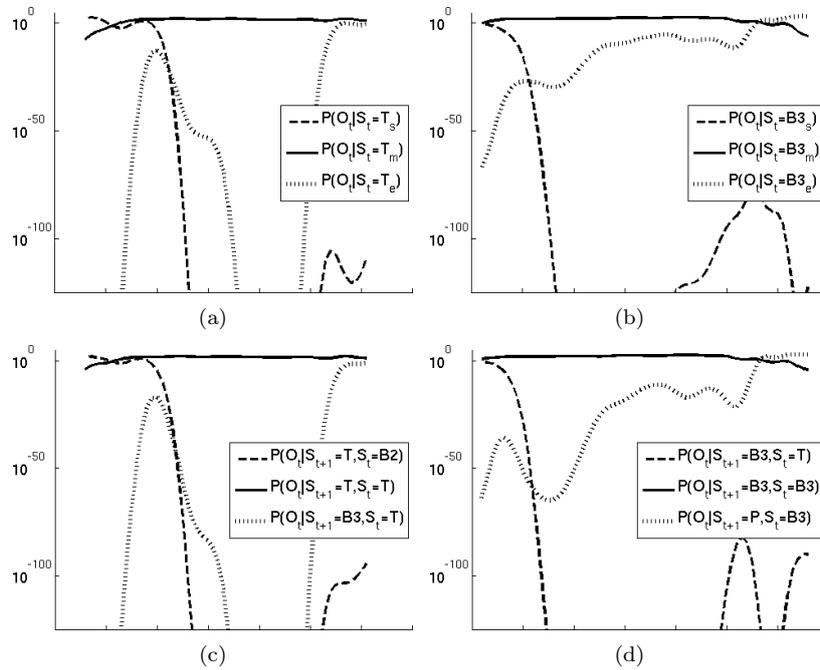


Figure 4: Distributions of the emission probabilities for a T wave and a B3 baseline obtained with (a,b) the SME decomposition and (c,d) the EHMM model.

4 Emission-dependent transitions HMMs

Another possibility is to use a more powerful Bayesian network. Indeed, a HMM is a specific case of more general Bayesian network [11] which allows modelling probabilistic dependencies in data. In the model considered in this section, the transitions depend on the starting and ending states, but also on the current emission, as illustrated by FIG. 5 which shows the statistical dependencies between states and emission values. In fact, this emission-dependent transitions HMM (EHMM) can be seen as a generalisation of the SME decomposition described above: the probability to enter (leave) a given wave will increase as the emission becomes closer from a statistical point of view to the emissions corresponding to incoming (outgoing) transitions observed in the training data.

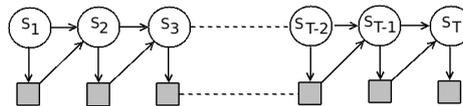


Figure 5: Statistical dependencies in an EHMM.

On one hand, one drawback of EHMMs is that leaving the state corresponding to a particular wave in its middle part is just very unlikely, not impossible. But on the other hand, the training of EHMMs does not require to choose a window width w ; so that EHMMs should be able to provide more accurate descriptions of the ECG signal dynamics.

FIG. 4(c) and FIG. 4(d) show the emission probabilities of a T wave and of a B3 baseline obtained after applying a wavelet transform with a Coiflet wavelet and two dyadic scales. We can see that the probability distributions for different transitions are distinct and localised. Therefore, we can expect that an EHMM will easily discriminate between ingoing transition, self transitions and outgoing transitions, so that it will hardly leave a wave from its middle part.

5 Experiments

In order to check if the SME decomposition and the EHMMs obtain better segmentations than the state-of-the-art HMMs and HSMMs, we have applied these four methods on three kinds of pathological signals from the QT database [12]:

- 3 records from the MIT-BIH arrhythmia database;
- 3 records from the MIT-BIH ST change database;
- 3 records from the MIT-BIH supraventricular arrhythmia database.

Each signal is filtered using a 3Hz-30Hz band-pass filter and preprocessed using a wavelet transform with a Coiflet wavelet and 7 dyadic scales. For each method, the emission probability density functions are estimated using Gaussian mixture models (GMMs). The number of components of the GMMs was chosen using 10-fold cross-validation. The window width for the SME decomposition is $w = 1$. Notice that the first and last annotations were imposed.

TAB. 1 and TAB. 2 show the results of the four methods. The quality of each segmentation is assessed using (i) the percentage of correct classifications (or accuracy), (ii) the ratio of the number of detected and real waves/baselines (or single beat percentage) and (iii) the average error in ms between the bounds of the detected and those of real waves (or consistency). Notice that the super-numerary waves are not used to compute the consistencies. As mentioned in [8], HSMMs outperform classical HMMs. However, they are themselves slightly outperformed by the SME decomposition and EHMMs in terms of accuracies and consistencies. The best single beat percentages are achieved with HSMMs.

6 Conclusion

This paper shows that the adequacy of the transition modelling is an important issue for ECG segmentation using HMMs. Moreover, it proposes to use two methods called SME decomposition and EHMMs which solve this problem and obtain better segmentations than the state-of-the-art solution based on HSMMs. We plan to perform more tests on a larger number of ECG signals.

Method	Accuracy	Single beat percentage	Consistency
HMM	94.83%	96.34%	9.70ms
HSMM	95.03%	99.18%	9.53ms
HMM+SME	95.23%	99.09%	9.18ms
EHMM	95.12%	95.82%	9.30ms

Table 1: Means of the accuracies, single beat percentages and consistencies on 9 ECG signals for the HMM, HSMM, SME decomposition and EHMM methods.

Method	Accuracy	Single beat percentage	Consistency
HSMM	0.20%	2.83%	0.17ms
HMM+SME	0.40%	2.74%	0.53ms
EHMM	0.29%	-0.52%	0.41ms

Table 2: Means of the improvements in accuracies, single beat percentages and consistencies on 9 ECG signals for the HSMM, SME decomposition and EHMM methods w.r.t the results obtained with the HMM method.

References

- [1] M. Laakso, A. Aberg, J. Savola, P. J. Pentikäinen, and K. Pyörälä. Diseases and drugs causing prolongation of the QT interval. *American journal of Cardiology*, 59(8):862–865, 1987.
- [2] G. D. Clifford, F. Azuaje, and P. E. McSharry. *Advanced Methods and Tools for ECG Data Analysis*. Artech House Publishers, 2006.
- [3] A. Koski. Modelling ECG signals with hidden markov models. *Artificial intelligence in medicine*, 8(5):453–471, 1996.
- [4] P. D. Addison. Wavelet transform and the ECG: A review. *Physiological Measurement*, 26(5):155–199, 2005.
- [5] R. V. Andreão and J. Boudy. Combining wavelet transform and hidden markov models for ECG segmentation. *Journal on Applied Signal Processing*, 2007(1):1–9, 2007.
- [6] S. Mallat. *A Wavelet Tour Of Signal Processing (Wavelet Analysis And Its Applications)*. IEEE Press, 1999.
- [7] L. R. Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. pages 267–296, 1990.
- [8] N. P. Hughes, L. Tarassenko, and S. J. Roberts. Markov models for automated ECG interval analysis. In *Advances in Neural Information Processing Systems 16*, pages 611–618. MIT Press, 2004.
- [9] H. Bourlard and C.J. Wellekens. Connected speech recognition by phonemic semi-markov chains for state occupancy modelling. In *Signal Processing III: Theories and Applications*, pages 511–514. Elsevier Science Publishers, 1986.
- [10] G. D. Forney. The Viterbi algorithm. *Proceedings of the IEEE*, 61(2):268–278, 1973.
- [11] C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2007.
- [12] A. L. Goldberger, M. Luis, A. N. Amaral, L. Glass, J. M. Hausdorff, P. Ch. Ivanov, R. G. Mark, J. E. Mietus, G. B. Moody, C k Peng, and P. H. E. Stanley. Physiobank, physioooolkit, and physionet components of a new research resource for complex physiologic signals. *Circulation*, 101(22):215–220, 2000.