

Improving the safety of atrial fibrillation monitoring systems through human verification

FAUST, Oliver <<http://orcid.org/0000-0002-0352-6716>>, CIACCIO, Edward J., MAJID, Arshad and ACHARYA, U. Rajendra

Available from Sheffield Hallam University Research Archive (SHURA) at:

<https://shura.shu.ac.uk/24755/>

This document is the Accepted Version [AM]

Citation:

FAUST, Oliver, CIACCIO, Edward J., MAJID, Arshad and ACHARYA, U. Rajendra (2019). Improving the safety of atrial fibrillation monitoring systems through human verification. *Safety science*, 118, 881-886. [Article]

Copyright and re-use policy

See <http://shura.shu.ac.uk/information.html>

Improving the safety of atrial fibrillation monitoring systems through human verification

Abstract

In this paper we propose a hybrid decision-making process for medical diagnosis. The hypothesis tested is that a deep learning system can provide real-time monitoring of Atrial Fibrillation (AF), a prevalent heart arrhythmia, and a human cardiologist will then verify the results and reach a diagnosis. The verification step adds the necessary checks and balances to increase the safety of the computer-based diagnostic process.

In order to test hybrid-decision making, we created a prototype AF monitoring service. The service is based on Heart Rate (HR) sensors for signal acquisition as well as Internet of Things (IoT) technology for data communication and storage. These technologies enable transfer of HR data from patient to central cloud server. A deep learning system is used to analyze the data, which is then presented to a cardiologist when a dangerous condition is detected. This human specialist then works to verify the deep learning results based on the HR data and additional knowledge obtained through patient records or by personal interaction with the patient.

A prerequisite for safety in any computer expert system is the clarity of purpose for the decision-making process. Health-care providers are considered customers who register patients with the AF monitoring service. The service delivers real-time diagnostic support by providing timely alarm messages and HR analysis. The safety critical decision then lies with the human practitioner.

1. Introduction

Atrial Fibrillation (AF) is a condition that causes irregular heartbeat, and in many cases a rapid heart-beat. AF is the most common clinically significant cardiac arrhythmia. It can cause problems including tiredness, shortness of breath, and dizziness. These symptoms may reduce quality of life, functional status, and cardiac performance. AF is associated with substantial medical cost, as well as an increased risk of death. It is a potent risk factor for ischemic stroke (Wolf et al., 1991). The absolute impact of AF on the stroke risk depends on comorbid conditions and age (Gage et al., 2001). There are several types of AF, usually categorized as paroxysmal, persistent and permanent. Treatment of AF is possible via medication that controls the heart-beat, as well as more invasive procedures, such as cardioversion, that work to restore the normal heart rhythm (Lafuente-Lafuente et al., 2015). Treatment monitoring provides the necessary feedback to ensure a positive outcome for the patient. Current clinical practice is based on sampling patient health via consultation, and forming an opinion concerning treatment success based on limited data. This is less desirable than continuous monitoring, which would provide an accurate representation on how AF develops over time. The presence of AF can be established via Heart Rate (HR) measurements (Hagiwara et al., 2018). As such, HR is a physiological signal with high information content, because structural features are largely absent in the signal waveform, and a low data-rate. The low data-rate is ideal for online monitoring, with the data stored and assessed in a central location (Yang et al., 2015). The complex structure of the signal, which is a direct consequence of the high information content, makes it difficult for a human practitioner to entirely analyze the waveform in direct detail (Faust and Bairy, 2012). Hence, computational methods could be useful to extract human readable features from HR signals. State-of-the-art diagnostic support tools typically incorporate Artificial Intelligence (AI) structures to interpret the HR signals (Faust et al., 2018a). In addition, computer support is useful to reduce inter- and intra-operator variability, because it does not depend on subjective decision-making. However,

that form of machine-based AF diagnosis is not safe to implement in a clinical setting, because the machine learning system might err when presented with wrong or corrupt data. Furthermore, the computing system will be unaware of the general circumstances of the patient.

To address these problems, we propose an AF monitoring service that incorporates a hybrid approach to medical diagnosis. A deep learning system will be used to monitor patient heart-beat in real-time. The machine decision is reached without feature engineering, which reduces design time errors that would impact negatively on the learning model. When a serious patient condition is detected, the service framework will be directed to broadcast alarm messages to the appropriate medical practitioners. In response to the alarm messages, a human specialist will review the evidence contained in the HR trace, and fuse this information with further knowledge and experience concerning the patient, in order to reach a diagnosis. As such, the validation process requires substantially less time than the real-time signal analysis. However, validation and subsequent diagnosis requires higher-order cognitive abilities. Hence, the proposed hybrid decision-making process will utilize both the objectiveness and diligence of deep learning systems along with human ability, to fuse a multitude of information sources and reach a diagnosis that could be implemented in a clinical setting.

To support our hypothesis that the hybrid decision-making process can improve the accuracy and safety of AF monitoring, we have structured the remainder of the paper in the following way. The Methods section introduces the technologies used to implement the AF monitoring service. Having that practical background enables us to discuss the safety aspects of AF monitoring. The Discussion section is a main focus of this paper, because in it we consider checks and balances required to improve the safety of an automated decision-making system. In the Conclusion, the main findings of the paper are stated and the merits of the proposed hybrid decision-making system are highlighted.

2. Methods

The safety of an implemented clinical system can be associated with the question: How does the system react in case of failure? Therefore, we introduce the proposed system architecture before discussing the safety aspects of the paradigm. Figure 1 shows an overview block diagram of the proposed AF monitoring system. With that system setup, we follow Internet of Things (IoT) design principles, wherein the data flows from point of measurement, i.e. the patient, to a central location for storage and processing (cloud).




Figure 1: Setup of the AF monitoring service

The main impetus behind the system setup is that HR data travels from patient to a central cloud server. The data acquisition for the implemented AF monitoring service is captured with commercial HR sensors, which communicate data to a smart-phone via a low power Bluetooth link. IoT technology was employed to transfer data from the phone to a central cloud server for storage. Having the data in




Figure 2: Overview block diagram for the deep learning inference

a central location enabled us to readily deploy a deep learning system for real-time analysis. The deep learning system is used to determine the **AF** probability for a block of 100 **HR** samples. If the deep learning system detects a dangerous condition, the medical practitioner receives an alarm message. This prompts the human expert to review the machine decision and reach a diagnosis.

A safety relevant aspect for measurement, communication and storage processes is data integrity (Faust et al., 2012). As such, this can be established by using adequate engineering principals (Bahr, 2018). Engineering best practices can be useful to reduce data corruption to a minimum (Faust et al., 2011). However, these best practice considerations are insufficient for deep learning systems, because formal correctness of the decision functionality cannot be established during the design procedure. In our architecture, the input vector to the deep learning system was assigned 100 dimensions (block of heart rate samples), wherein each element is encoded by a 2 byte integer. Therefore, the input vector would project a state space of $2^{16 \times 100} = 4.4 \times 10^{481}$. The time to computationally check all these states would be enormous. Hence, a formal proof of the system as specified is not currently implementable. However, in the subsequent sections of this paper, we propose a system that utilizes human verification to improve the system safety for clinical use.

2.1. Deep learning

Deep learning is used to estimate the probability that a **HR** signal segment was measured during **AF**. To realize this functionality, the **HR** signal is partitioned into 100 beat segments, and these segments are input to the Long Short-Term Memory (**LSTM**) based deep learning system. As such **LSTM** is suitable for **HR** processing, because it is optimized for processing sample sequences (Hochreiter and Schmidhuber, 1997). The block diagram, shown in Figure 2, documents the functionality. Based on the data, the Deep Learning (**DL**) system generates a value between 0 and 1. The value indicates the estimated **AF** probability for the input segment. A value of 0 indicates that the deep learning system is certain that the analyzed signal segment does not present an **AF** rhythm. Conversely, if the outcome is 1 then the deep learning system is certain that the signal segment shows an **AF** rhythm.

The deep learning model was established with data from PhysioNet’s Atrial Fibrillation Database (**AFDB**). The data from 20 patients was used to train and test the model. Ten-fold cross validation yielded an accuracy of 98.51% (Faust et al., 2018b). This value was established by thresholding the estimated **AF** probability, i.e. all result values below 0.5 were classed as non-**AF**, and all result values of 0.5 and above were classed as **AF**. From a safety perspective, that accuracy is the upper limit of the system performance. However, from a medical perspective, the thresholding is not satisfactory, because it obscures information related to the uncertainty of the deep learning system. The reading cardiologist must be able to trace the analytic results, which implies that it must be possible to observe both the **HR** waveform and the estimated **AF** probability. We address this problem with the cardiologist support software program, which is described in the next section.

2.2. Cardiologist support

The cardiologist needs to reach a safe, reliable and accurate diagnosis, in order to achieve the best outcome for a patient. To do so, the specialist requires the relevant information at the right time, rather than just all the available information. In this case, relevant information means an indication of the suspicious **HR** traces and reliable features from these traces. The modified Heart Rate Variability Analysis Software (**HRVAS**) program addresses this need. Figure 3 shows a screenshot which depicts the major elements of the modified **HRVAS** program (Ramshur, 2010). The **HRVAS** program was modified such that it is possible to fetch, display and process data from the cloud server. The proposed

use case scenario, for this cardiologist support tool, initiates by the user selecting the patient from the drop-down list at top of the Graphical User Interface (GUI). Pressing ‘Fetch Data’ loads the estimated AF probability from the cloud server. That data is displayed in the form of a two-dimensional graph, presented in the upper left corner of the GUI. The graph provides an overview of the estimated AF probability over time. It is envisioned that this graph will provide concise information concerning disease progression. For further analysis, the cardiologist can then select a region of interest on the first graph. This action will cause the tool to download the HR segments. The user selects a range of deep learning results; each deep learning result is associated with a block of 100 RR intervals. The tool downloads the RR interval blocks for which the results were selected. Furthermore, the tool calculates a wide range of features for all of the downloaded RR intervals. Ramshur described all of the available features in his original HRVAS documentation report (Ramshur, 2010). The graph in the lower left corner of the GUI displays the downloaded RR intervals. These RR intervals are color-coded in accordance with their associated estimated AF probability. The color bar, shown below the graph, indicates the mapping between color and estimated AF probability.

The GUI of the modified HRVAS program displays the estimated AF probability first. This behavior was a deliberate design decision, because the estimated AF probability provides an overview of patient health – with respect to AF. The operator is required to select a region of interest before an augmented version of the HR trace is displayed, which corresponds to the selected region. Hence, the cardiologist sees the original evidence, augmented by the deep learning results, and the features which provide disease relevant information. This workflow paradigm therefore guides the cardiologist from the overview toward the detailed signals and feature analysis. Even at the signal level, the program thereby provides support by color-coding the signal plot in terms of the estimated AF probability. That functionality allows the cardiologist to be in control of the diagnosis process, none of the evidence is hidden while at the same time a wide range of algorithmic support is provided.

3. Discussion

Health care applications are often difficult to develop, because there is a human element which must be taken into account. With the hybrid decision-making process, we have proposed to synergize with, rather than replace, human experts. That cooperation combines the speed and diligence of machine classification with the overview and instinct of the human practitioners. This will maximize the detection of important clinical events, while at the same time minimizing the possibility of misdiagnosis. Thus, an increase in systemic efficacy is attained.

From a medical perspective, the proposed system is based on the HR being a satisfactory predictor of human health. Extracting disease-relevant information from the HR signals is an active area of research. Hence, new knowledge becomes available over time. There is an opportunity to integrate this new knowledge into the decision-making process. The proposed hybrid decision-making paradigm would be ideal for continuously retraining a deep learning network. Having a human practitioner validating the deep learning result can be used to provide relevant training data. In the case when the human practitioner agrees with the deep learning result, the validated HR sequence can be added to the labeled data used for training. Even more valuable would be the case where the human practitioner rejects the deep learning result, meaning that the machine classification was thought to be incorrect. The incorrectly identified HR traces are then assigned a higher weight for subsequent retraining of the network, in order to correct the misconceptions of the network.

From an engineering perspective, standard safety assessment methods are not applicable for DL systems. Any such standard methods would attempt to establish, beyond a reasonable doubt, that the software app is safe, i.e. the information it provides will not harm humans under any circumstances (Hoare, 1978; Brookes et al., 1984; Abrial and Abrial, 2005; Lamport, 1999). However, proving this statement would involve the exhaustive checking of all possible circumstances. For design realization, all circumstances would mean the state space of the deep learning system (Faust et al., 2006). For the LSTM based DL algorithm, the state space is determined by the input vector. As outlined in Section 2.1, in a practical implementation that input vector can encode 4.4×10^{481} states. It is not practical

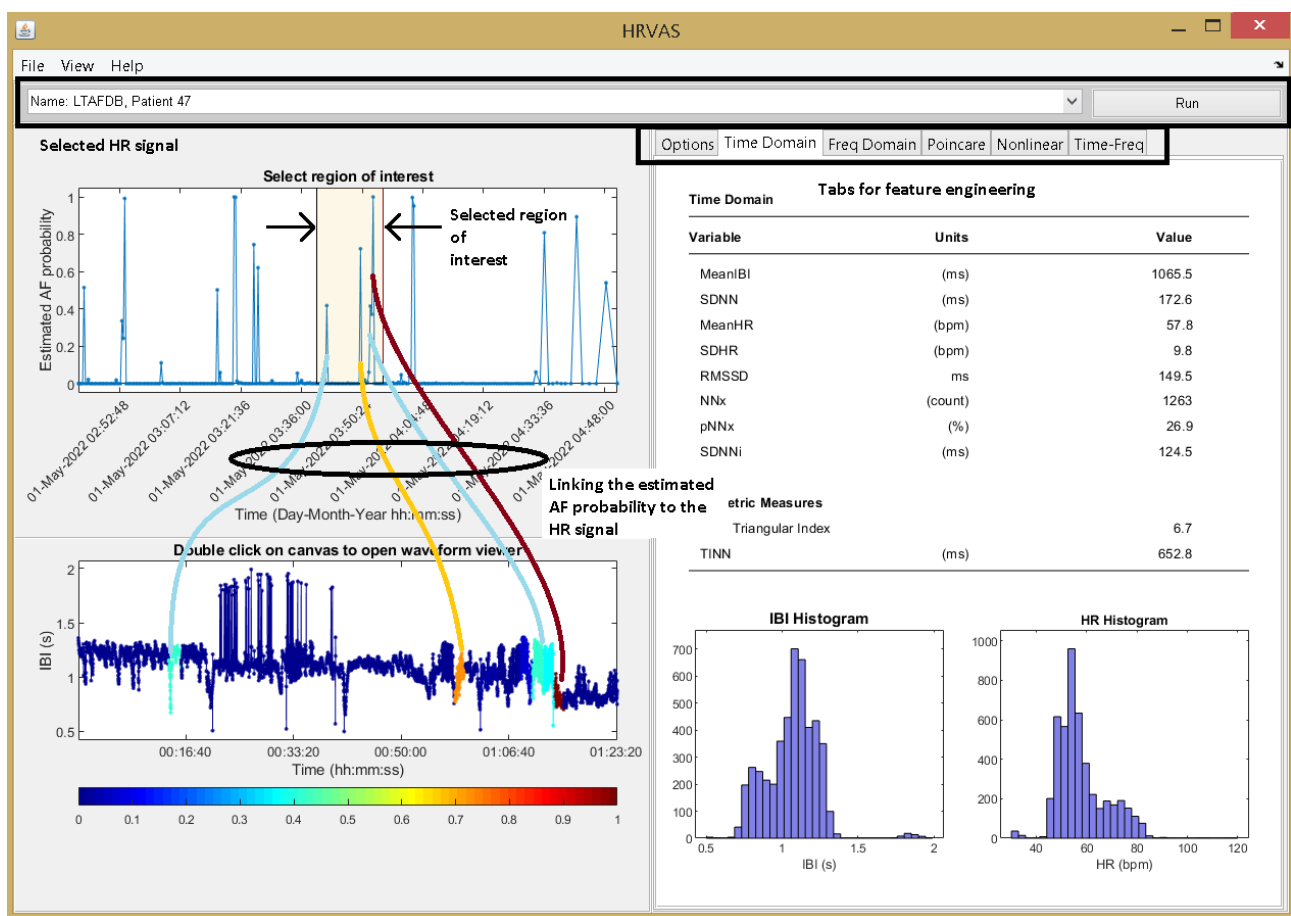


Figure 3: Cardiologist support with the modified HRVAS program (Ramshur, 2010).

to explore such a state space (Anders and Grevesse, 1989). Furthermore, testing the DL algorithm with a subset of the state space will not establish the safety of the system. The coverage of such tests is minute, i.e. the test set will always be small when compared to the state space. Even choosing a salient test set is not an option, because our understanding of the DL model, i.e. the weighting pattern, is limited. Thus, it is not possible to know or to accurately estimate where the corner cases will be. The possible corner cases include a range in the state space wherein a small change in the input vector will cause a dramatic change in the estimated AF probability. With a subset of the state space, we can only establish the system functionality. As an alternative, we have used ten-fold cross validation to estimate the decision-making quality of the system. The RR interval blocks from the known training data constitute discrete points in the state space. Testing these points are used to inform the operator as to whether or not the system is sufficient for the purpose. However, it does not establish system safety.

System safety for deep learning systems must originate in checks and balances procedures (Lindsay, 1992). Having the data at a central location is a prerequisite for traceability, where an operator traces back a sequence of events to find the root cause of a safety critical fault. As such, traceability is an important systemic safety aspect. Misdiagnosis becomes evident if there is a negative outcome for the patient. From a system safety perspective, this is helpful information which can be used to improve the system. For the proposed setup, the cardiologist is tasked with reaching a diagnosis; the DL system combined with the modified HRVAS program are the support tools. These tools are also valuable for tracing back in case a misdiagnosis occurred. The cardiologist can then retrieve the program and inspect the relevant data, to reproduce the situation which led to misdiagnosis. In a tracing scenario, returning to the same situation will train the cardiologist. It might even be necessary to retrain the DL, such that relevant knowledge can be extracted from the misdiagnosis event.

The proposed hybrid decision-making system significantly reduces latency that could lead to negative patient outcomes. The data is streamed from patient to cloud server, and the deep learning system is capable of real-time monitoring. Our testing showed that a modern Central Process Unit (CPU) is capable of calculating the inference result for ca. 1 minute of HR data in 30 seconds. Diagnosis by the clinical expert is a slower procedure. However, compared with state-of-the-art Holter monitoring, where the diagnostic result is reached after days, the proposed method reduces latency between measurement and diagnosis significantly.

State of the art diagnostic approaches for AF include stress test, Holter monitor, clinical examination, Electrocardiogram (ECG), and echocardiogram. Of these methods, ECG is most often used to confirm a diagnosis, because it monitors the electrical activity of the human heart. Therefore, research work on Computer-Aided Diagnosis (CAD) for AF focuses on extracting diagnostically relevant information from ECG and HR signals. Table 1 provides an overview of the performance measures for ECG based studies. Only Acharya et al. (Acharya et al., 2017) used a deep learning approach to establish the diagnosis support, all the other studies used traditional machine learning algorithms. The performance of machine learning tends to deteriorate for large data sets, therefore it is difficult to establish practical decision support in a clinical setting. Furthermore, ECG has a significantly higher data rate when compared to HR. ECG signals are captured with around 256 samples a second whereas HR generates only 1 sample a second. As a consequence, ECG is more difficult, and indeed more costly, to communicate, store and process when compared to HR. That makes HR the better choice for long term monitoring.

Table 2 provides an overview of current research work on HR based AF detection. The HR is captured by measuring the beat to beat interval of the human heart. Faust et al. were the only authors who proposed to use a deep learning algorithm to discriminate between AF and normal HR signals. None of the papers in this limited review discusses the safety aspects of reaching a medical diagnosis. The papers concentrate on the performance figures of novel and innovative methods to make sense of HR signals. With our work we pick up where these other studies left off and show how deep learning based decision making can be used as part of a hybrid diagnosis process.

Table 1: Research work on automated ECG based AF detection

Author	Method		Decision support quality in %		
	Feature engineering	Classifier	Specificity	Sensitivity	Accuracy
Sufi and Khalil (2011)	Correlation-based feature subset, expectation maximization	Rule-based	—	—	97.00
Martis et al. (2013)	Independent component analysis	Gaussian mixture model	99.33%	99.32%	99.33%
Prasad et al. (2013)	Independent component analysis	K-nearest neighbour	98.75%	98.75%	97.65%
Martis et al. (2014a)	Wavelet coupled with independent component analysis	K-nearest neighbour	100%	99.61%	99.45%
Martis et al. (2014b)	Higher order spectra coupled with independent component analysis	K-nearest neighbour	—	100%	99.50%
Acharya et al. (2016)	Entropies, signal energy, Fractal dimension, Largest Lupnov event	Decision tree	84.10	99.30	96.30
Desai et al. (2016)	Recurrence quantification analysis	Rotation forest	—	—	98.37
Acharya et al. (2017)	—	11-layer Convolutional neural network	93.13	98.09	92.50
Kumar et al. (2018)	Wavelets and entropy measures	Decision tree	97.60%	95.80%	96.84%

Table 2: Research work on automated HR based AF detection

Author	Method		Decision support quality in %		
	Feature engineering	Classifier	Specificity	Sensitivity	Accuracy
Mohebbi and Ghassemian (2011)	Recurrence quantification analysis	Support vector machine	100.0	97.00	-
Kennedy et al. (2016)	Sample entropy, Statistical HR features	Random forest	98.30	92.80	—
Faust et al. (2018b). Used as example in this study.	—	Long short-term memory recurrent neural network	—	—	98.51

3.1. Future work

In many circumstances, **HR** signals are robust to noise, because the the electrical activity of the heart beat, known as R wave peak, typically coincides with the highest amplitude in the **ECG**. However, heart disease may act to alter the electrical activity of the heart sufficiently so that the R peak is no longer the greatest amplitude point in the signal, and beat detection is no longer straightforward (Martínez et al., 2004). Detecting the heartbeat then becomes a matter of interpretation, which can influence Heart Rate Variability (**HRV**). The performance of the deep learning model can also degrade if a different beat detection method is used during training. More studies are needed to confirm that the deep learning model is robust in this scenario.

3.2. Limitations

Herein we have discussed at AF patient safety from the perspective of checks and balances in an automated computerized system. We used two decision processes, with the final decision resting solely on the cardiologist. Systemic problems caused by design-time errors have not yet been addressed. Data loss through the software framework would corrupt the data stream. In this case the automated decision-making process would reach a wrong conclusion. A formal and model driven design is needed to eliminate the problem. Furthermore, the decision-making process is susceptible to design-time errors (Song et al., 2012).

4. Conclusion

DL systems model human decision-making. Therefore, it is not surprising that such systems suffer from similar safety problems as human decision-making. It is necessary to increase the safety of the system to better mimic human decision-making. We introduced checks and balances to better ensure that the correct route is followed, designed to make an **AF** diagnosis safer. The main idea was to combine human and machine decision-making. The machine classifier analysed **HR** signals and published the results as estimated **AF** probabilities in real-time. A cardiologist could then look at the results and investigate specific regions of the **HR** signals.

The combination of machine and human decision-making is symbiotic. Machine decision-making excels at robot tasks, such as the real-time analysis of **HR** signals. However, such algorithms can fail to appropriately address unforeseen circumstances; human decision-making is required to master them. Furthermore, a human cardiologist has personal contact with the patient. The knowledge of this human interaction can help to improve the safety of the diagnosis.

The combination of human and machine-decision advances medical diagnosis and treatment monitoring. Humans and machines work together, and the strength of their combined decision-making ability will improve the safety, functionality and reliability of the diagnosis.

Acknowledgements

The work was partially supported by the NIHR Sheffield Biomedical Research Centre (BRC). The views expressed are those of the authors and not necessarily those of the NHS, the NIHR or the Department of Health and Social Care (DHSC).

Acronyms

AF	Atrial Fibrillation
AFDB	Atrial Fibrillation Database
AI	Artificial Intelligence
CAD	Computer-Aided Diagnosis
CPU	Central Process Unit
DL	Deep Learning

ECG	Electrocardiogram
GUI	Graphical User Interface
HR	Heart Rate
HRV	Heart Rate Variability
HRVAS	Heart Rate Variability Analysis Software
IoT	Internet of Things
LSTM	Long Short-Term Memory

References

- Abrial, J.R., Abrial, J.R., 2005. The B-book: assigning programs to meanings. Cambridge University Press.
- Acharya, U.R., Fujita, H., Adam, M., Lih, O.S., Hong, T.J., Sudarshan, V.K., Koh, J.E., 2016. Automated characterization of arrhythmias using nonlinear features from tachycardia ecg beats, in: 2016 IEEE International Conference on Systems, Man, and Cybernetics (SMC), IEEE. pp. 000533–000538.
- Acharya, U.R., Fujita, H., Lih, O.S., Hagiwara, Y., Tan, J.H., Adam, M., 2017. Automated detection of arrhythmias using different intervals of tachycardia ecg segments with convolutional neural network. *Information sciences* 405, 81–90.
- Anders, E., Grevesse, N., 1989. Abundances of the elements: Meteoritic and solar. *Geochimica et Cosmochimica acta* 53, 197–214.
- Bahr, N.J., 2018. System safety engineering and risk assessment: a practical approach. CRC press.
- Brookes, S.D., Hoare, C.A., Roscoe, A.W., 1984. A theory of communicating sequential processes. *Journal of the ACM (JACM)* 31, 560–599.
- Desai, U., Martis, R.J., Acharya, U.R., Nayak, C.G., Seshikala, G., SHETTY K, R., 2016. Diagnosis of multiclass tachycardia beats using recurrence quantification analysis and ensemble classifiers. *Journal of Mechanics in Medicine and Biology* 16, 1640005.
- Faust, O., Acharya, R., Spath, B.H., Min, L.C., 2011. Systems engineering principles for the design of biomedical signal processing systems. *Computer methods and programs in biomedicine* 102, 267–276.
- Faust, O., Acharya, U.R., Tamura, T., 2012. Formal design methods for reliable computer-aided diagnosis: a review. *IEEE reviews in biomedical engineering* 5, 15–28.
- Faust, O., Bairy, M.G., 2012. Nonlinear analysis of physiological signals: a review. *Journal of Mechanics in Medicine and Biology* 12, 1240015.
- Faust, O., Hagiwara, Y., Hong, T.J., Lih, O.S., Acharya, U.R., 2018a. Deep learning for health-care applications based on physiological signals: A review. *Computer methods and programs in biomedicine* .
- Faust, O., Shenfield, A., Kareem, M., San, T.R., Fujita, H., Acharya, U.R., 2018b. Automated detection of atrial fibrillation using long short-term memory network with rr interval signals. *Computers in biology and medicine* 102, 327–335.
- Faust, O., Spath, B.H., Allen, A.R., 2006. A study of percolation phenomena in process networks., in: CPA, pp. 109–121.

- Gage, B.F., Waterman, A.D., Shannon, W., Boechler, M., Rich, M.W., Radford, M.J., 2001. Validation of clinical classification schemes for predicting stroke: results from the national registry of atrial fibrillation. *Jama* 285, 2864–2870.
- Hagiwara, Y., Fujita, H., Oh, S.L., Tan, J.H., San Tan, R., Ciaccio, E.J., Acharya, U.R., 2018. Computer-aided diagnosis of atrial fibrillation based on ecg signals: a review. *Information Sciences* 467, 99–114.
- Hoare, C.A.R., 1978. Communicating sequential processes. *Communications of the ACM* 21, 666–677.
- Hochreiter, S., Schmidhuber, J., 1997. Long short-term memory. *Neural computation* 9, 1735–1780.
- Kennedy, A., Finlay, D.D., Guldenring, D., Bond, R.R., Moran, K., McLaughlin, J., 2016. Automated detection of atrial fibrillation using rr intervals and multivariate-based classification. *Journal of electrocardiology* 49, 871–876.
- Kumar, M., Pachori, R.B., Acharya, U.R., 2018. Automated diagnosis of atrial fibrillation ecg signals using entropy features extracted from flexible analytic wavelet transform. *Biocybernetics and Biomedical Engineering* 38, 564–573.
- Lafuente-Lafuente, C., Valembois, L., Bergmann, J.F., Belmin, J., 2015. Antiarrhythmics for maintaining sinus rhythm after cardioversion of atrial fibrillation. *Cochrane Database of Systematic Reviews* .
- Lamport, L., 1999. Specifying concurrent systems with tla^+ . *NATO ASI SERIES F COMPUTER AND SYSTEMS SCIENCES* 173, 183–250.
- Lindsay, F., 1992. Successful health and safety management. the contribution of management audit. *Safety science* 15, 387–402.
- Martínez, J.P., Almeida, R., Olmos, S., Rocha, A.P., Laguna, P., 2004. A wavelet-based ecg delineator: evaluation on standard databases. *IEEE Transactions on biomedical engineering* 51, 570–581.
- Martis, R.J., Acharya, U.R., Adeli, H., Prasad, H., Tan, J.H., Chua, K.C., Too, C.L., Yeo, S.W.J., Tong, L., 2014a. Computer aided diagnosis of atrial arrhythmia using dimensionality reduction methods on transform domain representation. *Biomedical signal processing and control* 13, 295–305.
- Martis, R.J., Acharya, U.R., Adeli, H., Prasad, H., Tan, J.H., Chua, K.C., Too, C.L., Yeo, S.W.J., Tong, L., 2014b. Computer aided diagnosis of atrial arrhythmia using dimensionality reduction methods on transform domain representation. *Biomedical signal processing and control* 13, 295–305.
- Martis, R.J., Acharya, U.R., Prasad, H., Chua, C.K., Lim, C.M., 2013. Automated detection of atrial fibrillation using bayesian paradigm. *Knowledge-Based Systems* 54, 269–275.
- Mohebbi, M., Ghassemian, H., 2011. Prediction of paroxysmal atrial fibrillation using recurrence plot-based features of the rr-interval signal. *Physiological measurement* 32, 1147.
- Prasad, H., Martis, R.J., Acharya, U.R., Min, L.C., Suri, J.S., 2013. Application of higher order spectra for accurate delineation of atrial arrhythmia, in: 2013 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), IEEE. pp. 57–60.
- Ramshur, J.T., 2010. DESIGN, EVALUATION, AND APPLICAION OF HEART RATE VARIABILITY ANALYSIS SOFTWARE (HRVAS). Master’s thesis. The University of Memphis.

- Song, Z., Ji, Z., Ma, J.G., Spath, B., Acharya, U.R., Faust, O., 2012. A systematic approach to embedded biomedical decision making. *Computer methods and programs in biomedicine* 108, 656–664.
- Sufi, F., Khalil, I., 2011. Diagnosis of cardiovascular abnormalities from compressed ecg: a data mining-based approach. *IEEE transactions on information technology in biomedicine* 15, 33–39.
- Wolf, P.A., Abbott, R.D., Kannel, W.B., 1991. Atrial fibrillation as an independent risk factor for stroke: the framingham study. *Stroke* 22, 983–988.
- Yang, J.J., Li, J.Q., Niu, Y., 2015. A hybrid solution for privacy preserving medical data sharing in the cloud environment. *Future Generation Computer Systems* 43, 74–86.