Modelling and personalisation techniques for behavioural prediction and emotion recognition

Marta Kwiatkowska

Department of Computer Science, University of Oxford

HSB 2019, Prague, 6th April 2019

A typical scene today

Smartphones, wearables everywhere…

- Rely on the phone ion your pocket for
  - communication
  - shopping
  - navigation…

- Multitude of data being collected
  - map, location, GPS, heart rate, gait, preferences, …

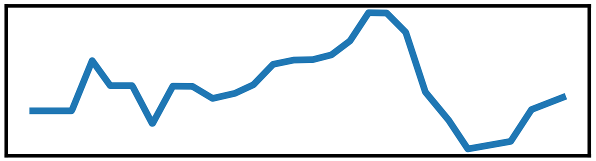

- Can we build accurate models to predict mood & behaviour?
  - emotion, stress, trust, intention, …
- For good uses…

Why predict mood & behaviour?

- Monitoring of affective disorders
  - stress, depression, bipolar, cognitive decline

  and their management/regulation
  - suggest coping strategies
  - send alerts
  - deliver medical intervention

- Also regulation of chronic medical conditions (diabetes, cardiac disorders, etc)

- Longer-term, effective human-robot collaboration
  - assistive robotics and shared control
  - cobotics

AffecTech project

People can experience different types of mental health disorders. From anxiety and panic attacks to bipolar disorder, depression and eating disorders, these problems can affect your thinking, mood, and behavior.

To combat affective health disorders, a coalition of scientists from around the world have launched the AffecTech project. In the next four years, this will look into developing new technologies that will empower people to better understand their emotions and deal with them on a daily basis. With some 4.88 million euros in funding, the European based research project kicked off this month.

- Personalised and adaptive emotion regulation
  - wearable systems for capturing emotion regulation
  - apps for understanding emotions and regulatory processes
  - personalised adaptive emotion regulation
  - automated synthesis of emotion regulation strategies

AffecTech:Personal Technologies for Affective Health ITN. http://www.cs.ox.ac.uk/projects/AFFECTech/

Modelling challenges

- Cyber-physical systems
  - hybrid combination of continuous and discrete dynamics, with stochasticity
  - autonomous control
- Data rich, data enabled models
  - achieved through learning
  - parameter estimation
  - continuous adaptation

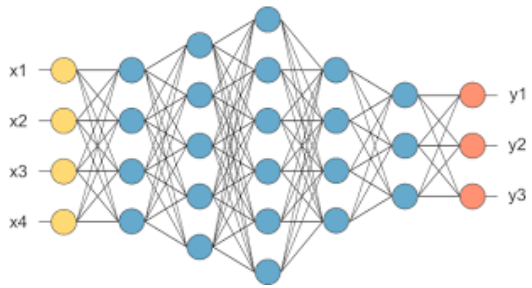

- Personalisation: key enabler of personalised healthcare
  - automation of intervention strategies
  - uniquely adapted to the individual

This lecture…

- Selected recent advances in quantitative modelling
  - focus on physiological signals
- The pacemaker case study
  - real CPS: non-linear hybrid dynamics, stochasticity
  - optimal parameter synthesis
  - personalisation
  - in silico testing
  - and more
- Multiple uses of quantitative models…
  - attacks on biometric security
  - intention prediction
  - emotion recognition
  - and more

Case study: Cardiac pacemaker

- Hybrid model-based framework
  - timed automata model for pacemaker software
  - hybrid heart models in Simulink
- http://www.veriware.org/heart_pm_methods.php
- Properties
  - (basic safety) maintain 60-100 beats per minute
  - (advanced) detailed analysis energy usage, plotted against timing parameters of the pacemaker
  - parameter synthesis: find values for timing delays that optimise energy usage

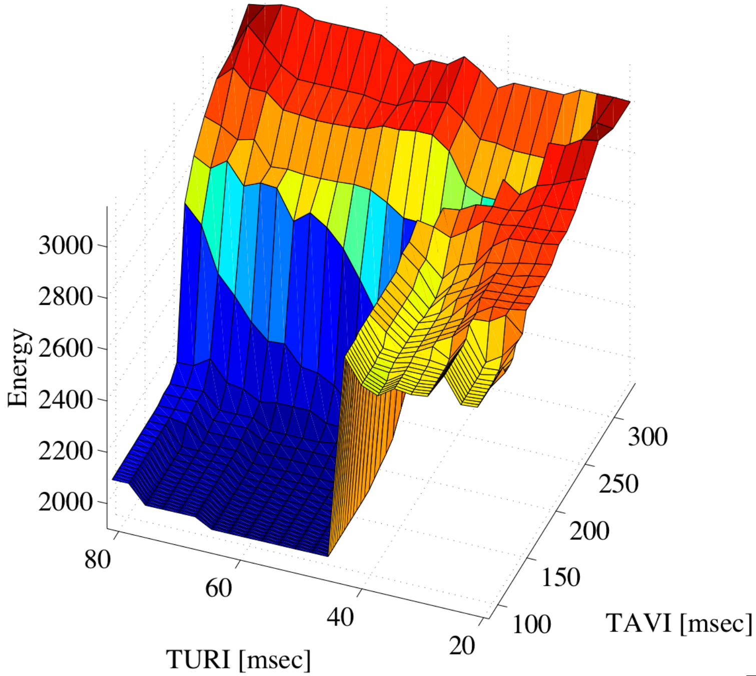

Synthesising robust and optimal parameters for cardiac pacemakers using symbolic and evolutionary computation techniques. Kwiatkowska, Mereacre, Paoletti and Patane, HSB'16

Quantitative verification for pacemakers

- Model the pacemaker and the heart, compose and verify

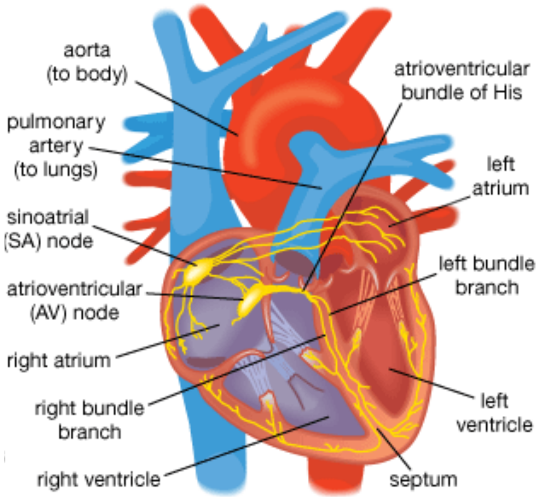

Quantitative verification for pacemakers

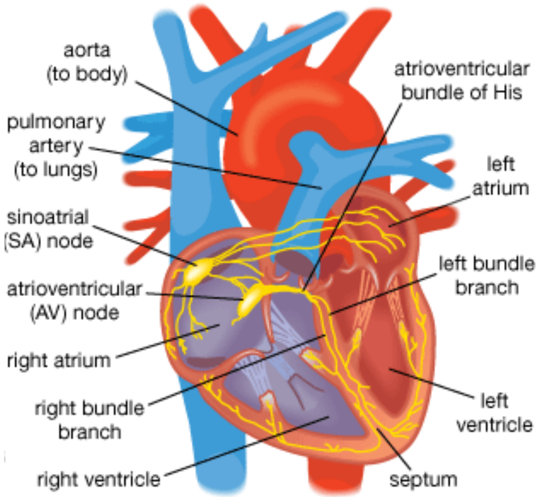

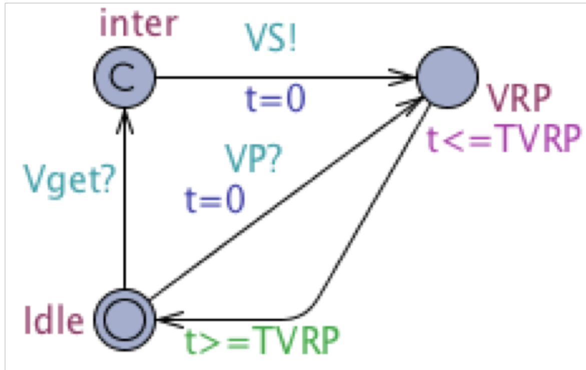


```
module VRP

s_vrp:[0..2] init 0;
t_vrp : clock;

// Invariants for clock t_vrp
  invariant
       (s_vrp = 2 => (t_vrp <= TVRP)) &
       (s_vrp = 1 => (t_vrp <= 0 ))
  endinvariant

[Vget] (s_vrp = 0) -> (s_vrp' = 1) & (t_vrp'=0);
[VP]  (s_vrp = 0) -> (s_vrp' = 2) & (t_vrp' = 0);
```

Quantitative verification for pacemakers

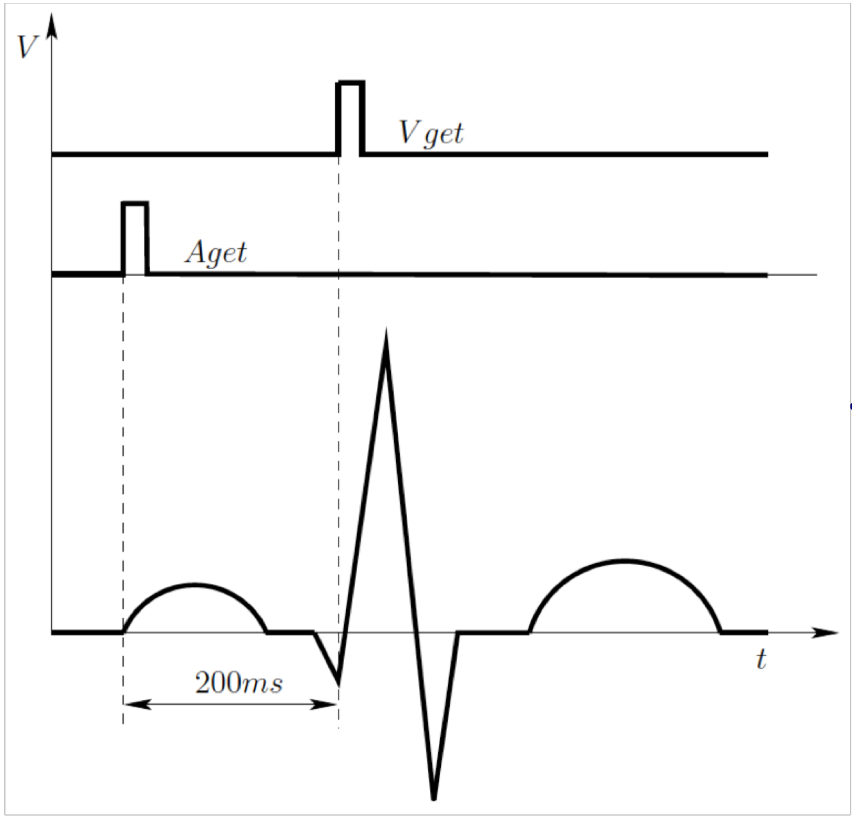

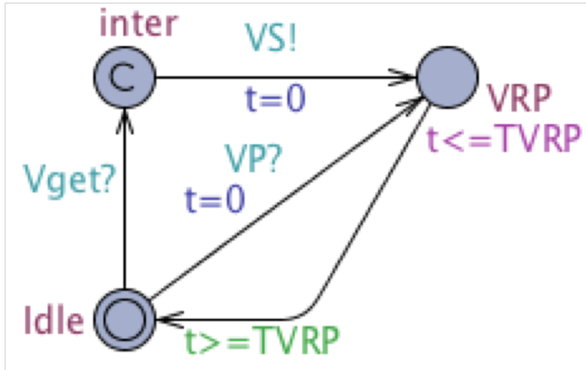


```
module VRP

s_vrp:[0..2] init 0;
t_vrp : clock;

// Invariants for clock t_vrp
  invariant
      (s_vrp = 2 => (t_vrp <= TVRP)) &
      (s_vrp = 1 => (t_vrp <= 0 ))
  endinvariant

[Vget] (s_vrp = 0) -> (s_vrp' = 1) & (t_vrp'=0);
[VP]  (s_vrp = 0) -> (s_vrp' = 2) & (t_vrp' = 0);
```

Model-based framework

- We advocate a model-based framework
  - models are networks of communicating hybrid I/O automata, realised in Matlab Simulink
    - discrete mode switching and continuous flows: electrical conduction system
    - quantitative: energy usage and battery models
    - patient-specific parameterisation
  - framework supports plug-and-play composition of
    - **heart models** (timed/hybrid automata, some stochasticity)
    - **pacemaker models** (timed automata)

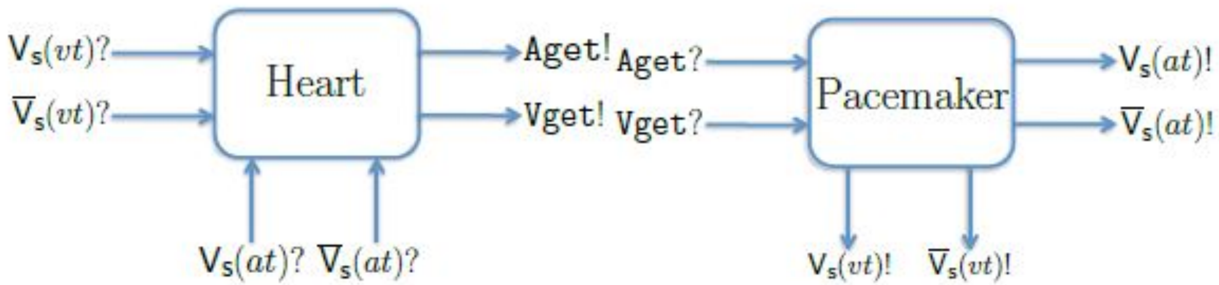

(a) Heart model (b) Pacemaker model

Quantitative Verification of Implantable Cardiac Pacemakers over Hybrid Heart Models. Chen *et al*, *Information and Computation*, 2014

Cardiac cell heart model

- Based on model of electrical conduction [Grosu et al]
  - abstracted as a network of cardiac cells that conduct voltage

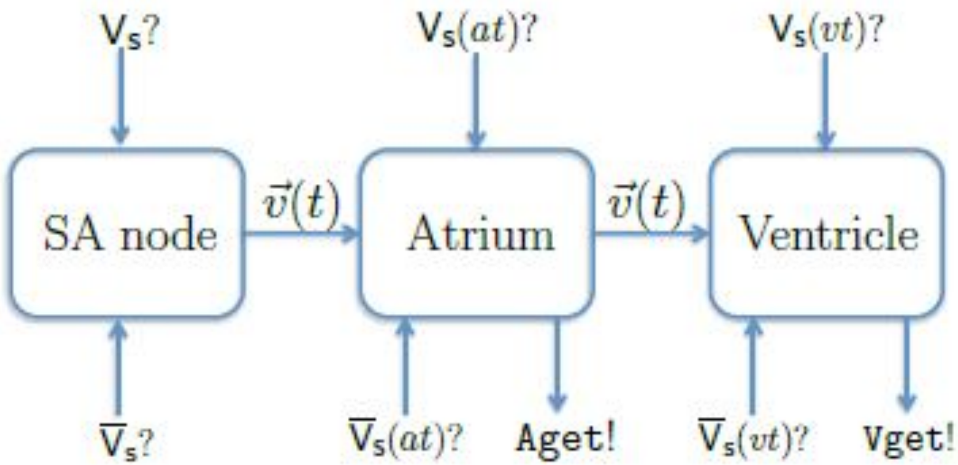

  - cells connected by pathways, modelled using Simulink delay and gain components
  - SA node is the natural pacemaker

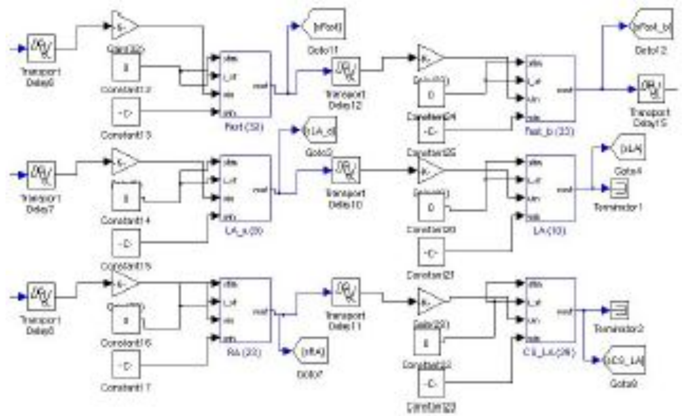

Cardiac cell heart model: single cell

- Single ventricular cell [Grosu et al]
  - four modes: resting and final repolarisation (q_0), stimulated (q_1), upstroke (q_2) and plateau and early repolarisation (q_3)

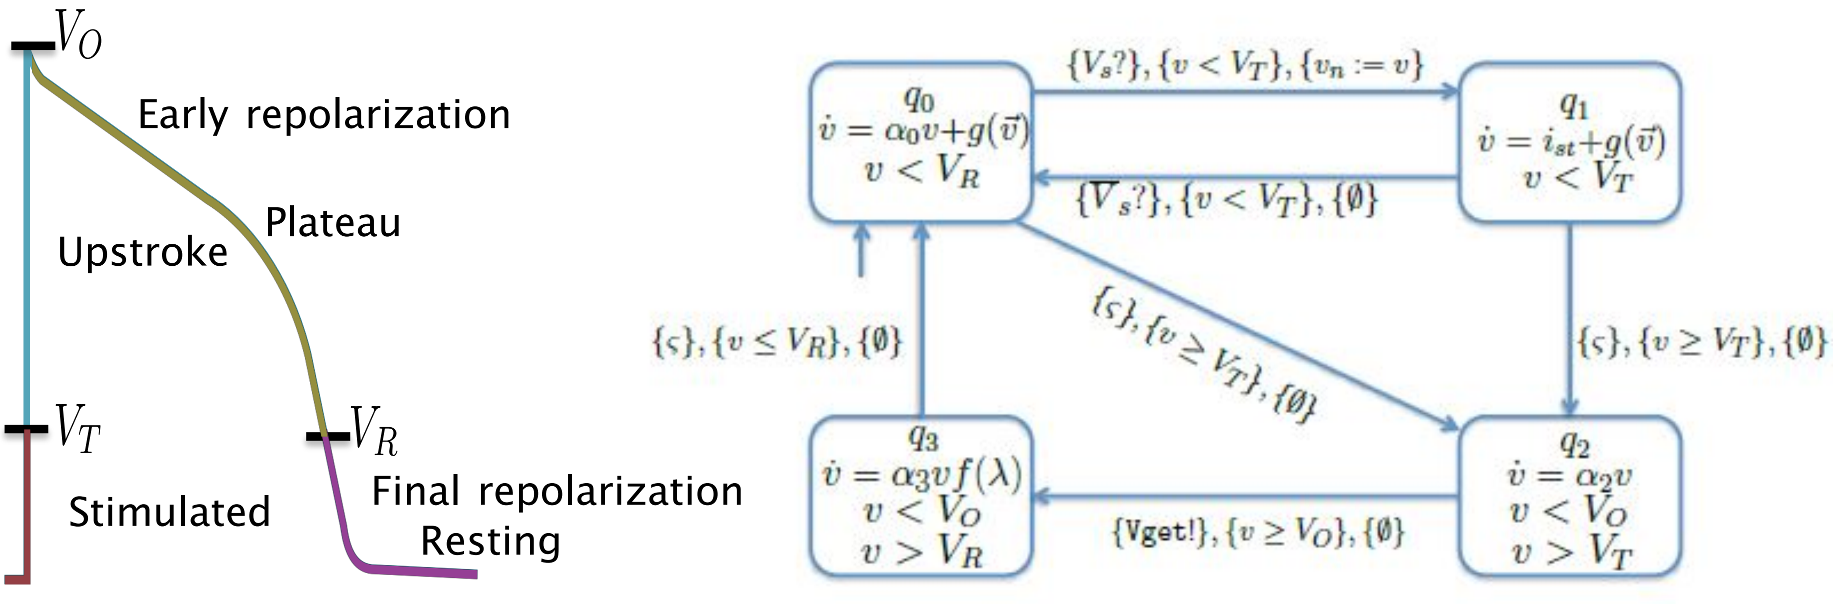

  - variables: v – membrane voltage, i_{st} – stimulus current
  - constants: V_R – repolarisation voltage, V_T – threshold, V_O – overshoot voltage

Property specification: Counting MTL

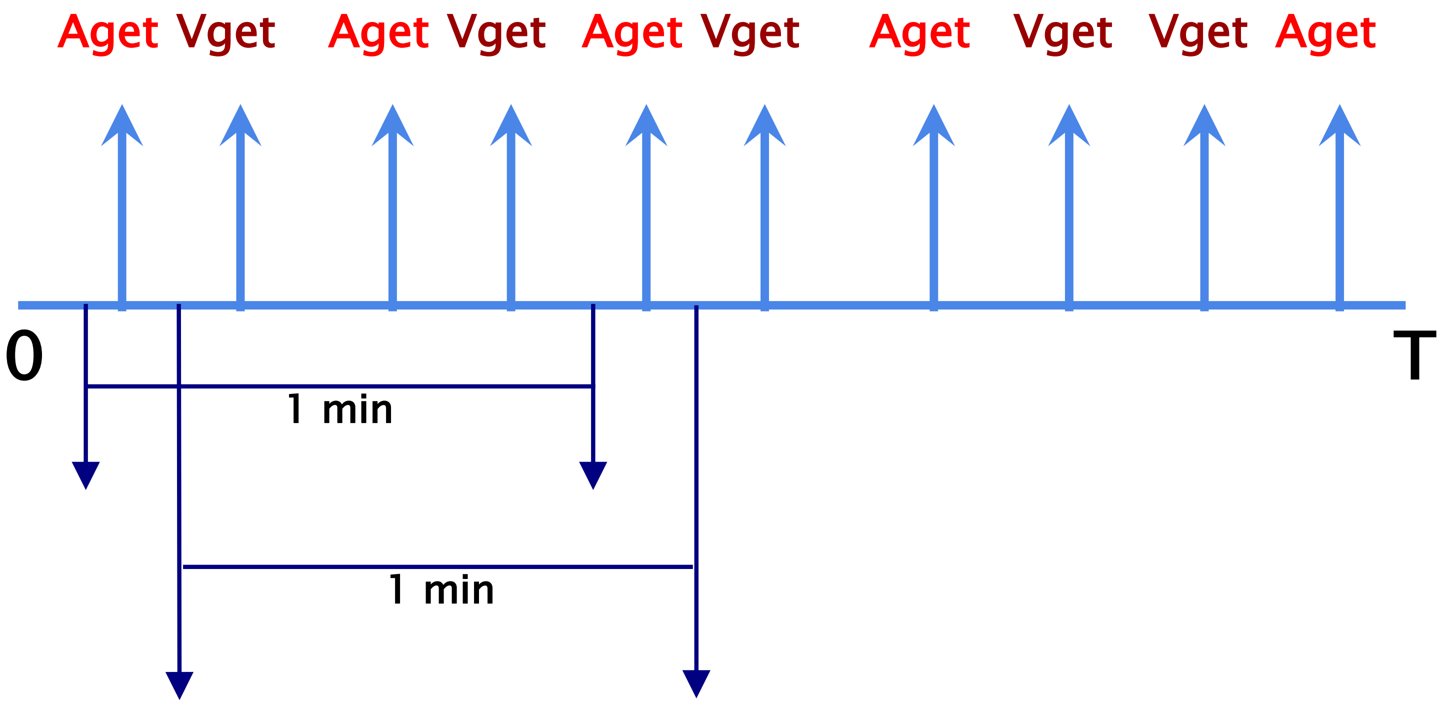

$$\Box^{[0,\tau]}(\#_0^\tau \mathtt{Vget} \geqslant B_1 \wedge \#_0^\tau \mathtt{Vget} \leqslant B_2)$$

Safety 'for any 1 minute window, heart rate is in the interval [60,100]"

Event counting not expressible in MTL (Metric Temporal Logic)

Framework functionality

- Broad range of techniques
  - Monte-Carlo simulation of composed models
    - with (confidence level) guarantees for non-linear flows
  - (approximate) quantitative verification against variants of MTL
    - to ensure property is satisfied
  - parametric analysis
    - for in silico evaluation, to reduce need for testing on patients
  - automated synthesis of optimal timing parameters
    - to determine delays between paces so that energy usage is optimised for a given patient
  - patient-specific parameterisation
  - hardware-in-the-loop simulation
    - parameter optimisation with respect to real energy measurements

- See http://www.veriware.org/pacemaker.php

Correction of Bradycardia

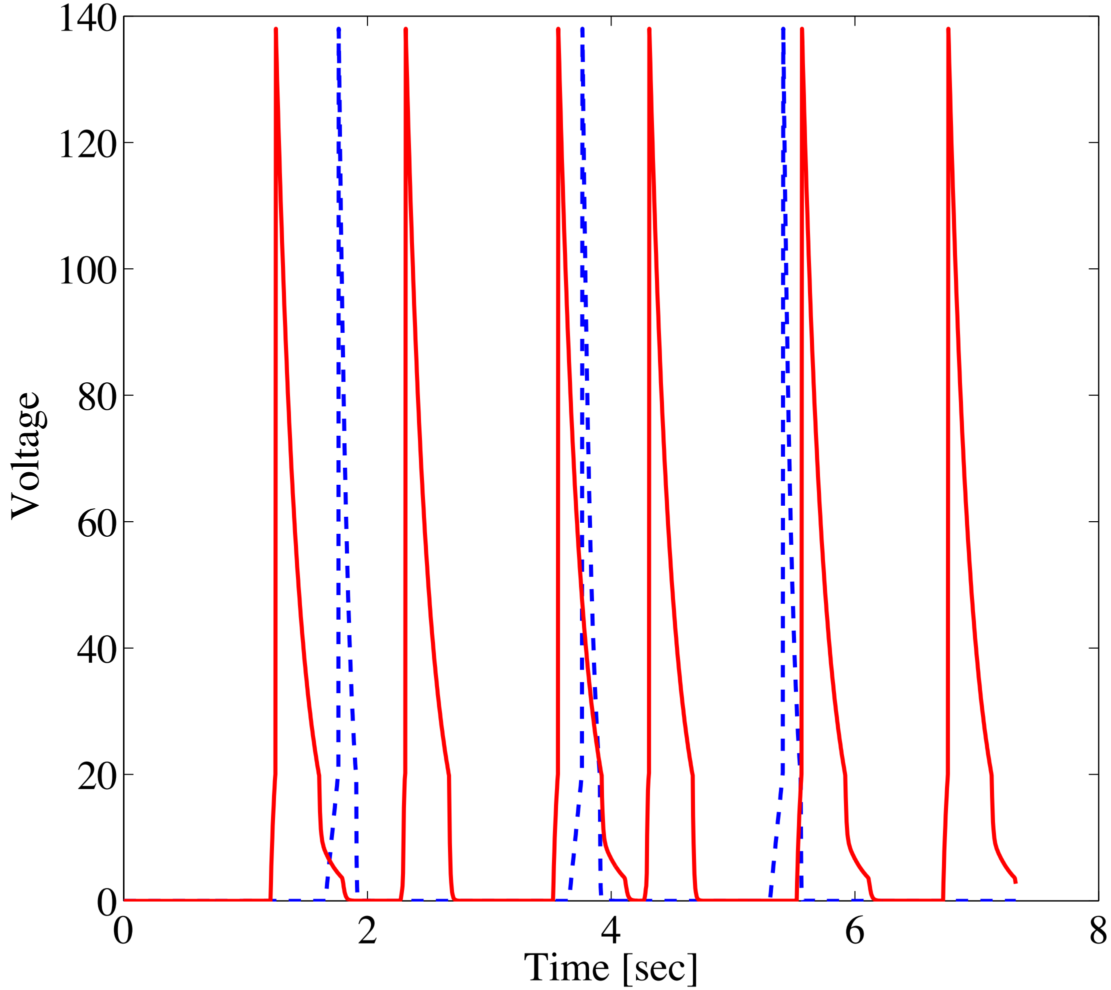

Blue lines original (slow) heart beat, red are induced (correcting)

Energy consumption

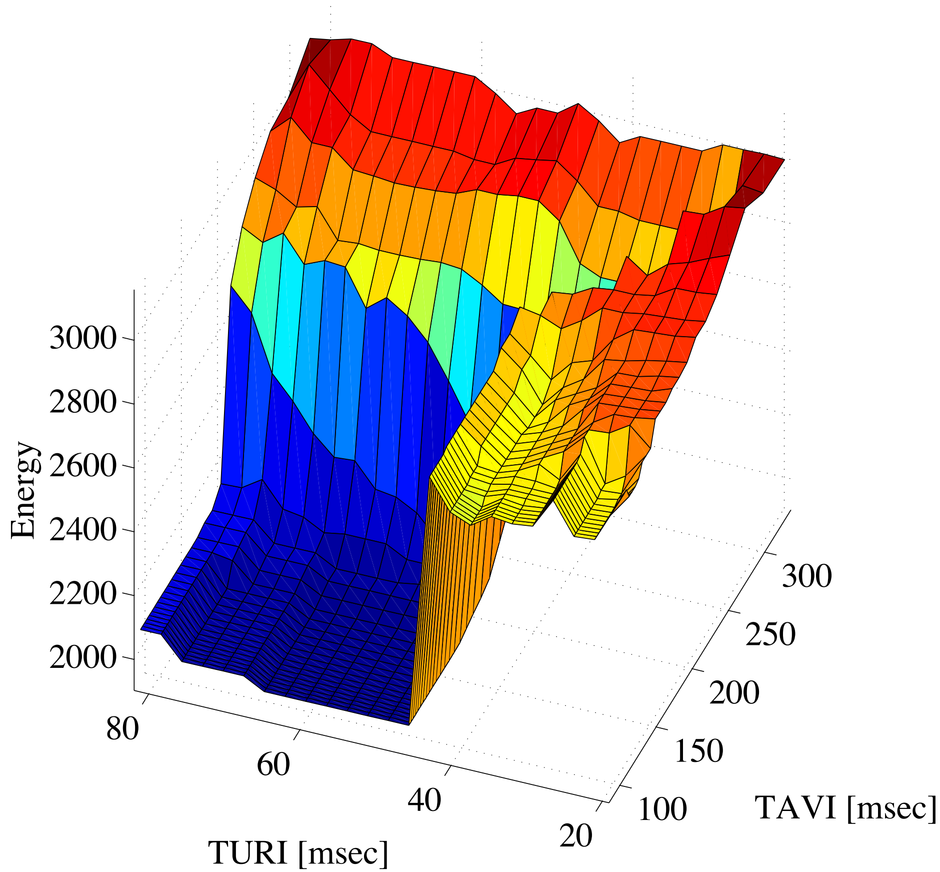

Efficiency "energy consumed must be below some fixed level"

Battery charge in 1 min under Bradycardia, varying timing parameters

Based on real power measurements

Hardware-in-the-loop simulation and energy optimization of cardiac pacemakers.
Barker *et al*, In *Proc EMBC*, 2015

Modulation during physical activity

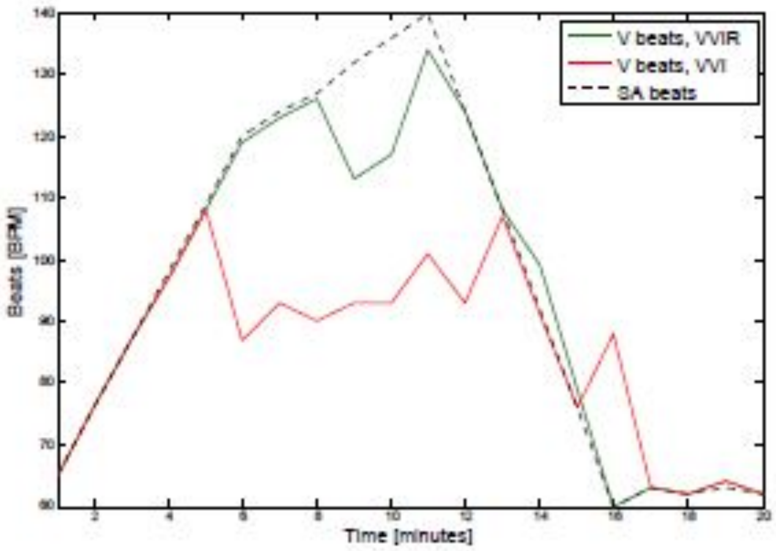

(a) Young patient

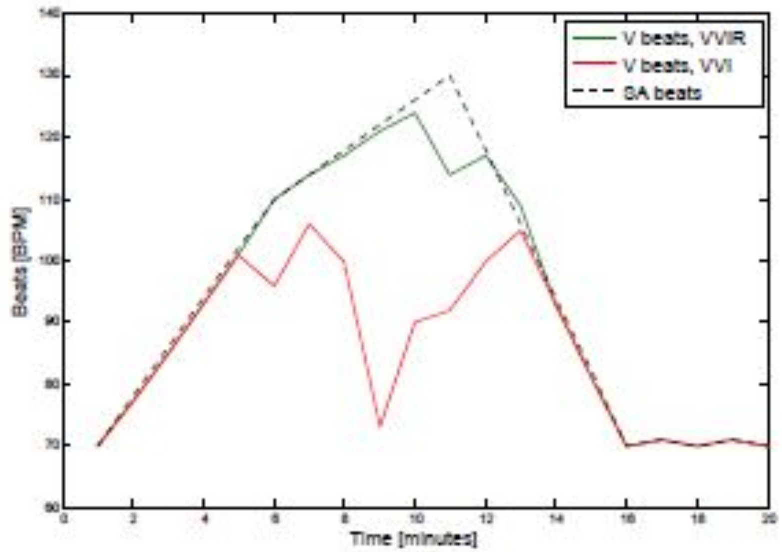

(b) Old patient

Rate modulation during exercise. Black dashed line indicates metabolic demand, and the green and red curves show rate-adaptive VVIR and fixed-rate VVI pacemakers.

Formal Modelling and Validation of Rate-Adaptive Pacemakers, Kwiatkowska *et al*. In *IEEE International Conference on Healthcare Informatics*, ACM. 2014

From verification to synthesis…

- Automated verification aims to establish if a property holds for a given model
- Can we find a model so that a property is satisfied?
  - difficult…
- The parameter synthesis problem is
  - given a parametric network of timed I/O automata, set of controllable and uncontrollable parameters, CMTL property ϕ and length of path n
  - find the optimal controllable parameter values, for any uncontrollable parameter values, with respect to an objective function O, such that the property ϕ is satisfied on paths of length n, if such values exist
- Objective function
  - maximise cardiac output, or ensure robustness

Synthesising Optimal Timing Delays for Timed I/O Automata. Diciolla et al. In *14th International Conference on Embedded Software (EMSOFT'14)*, ACM. To appear. 2014

Optimal timing delays

- Bi-level optimisation problem
- Safe heart rhythm CMTL property (inner problem)

$$\phi = \Box^{[0,T]} \left(vPeriod \in [500, 1000] \right)$$

  - at any time in [0,T] any two consecutive ventricular beats are between 500 and 1000 ms, i.e. heart rate of 60 and 120 BPM
- Cost function (outer problem)

$$2 \cdot \#_0^{60000} (act = AP) + 3 \cdot \#_0^{60000} (act = VP)$$

  - energy consumption in 1 minute

$$\frac{\sum_{(\mathbf{q},\eta) \in Vbeat(\rho')} |\eta(CO) - \overline{CO}|}{|Vbeat(\rho')|}$$

  - mean difference between cardiac output and reference value

Synthesis results

- Solved through SMT encoding (inner problem) combined with evolutionary computation (outer problem)
- Pacemaker parameters:
  - TLRI: time the PM waits before pacing atrium
  - TURI: time before pacing ventricle after atrial event
- Significant improvement (>50%) over default values
  - path 20
- A (exact),B (evo) energy
- C (exact),D (evo) CO
  - evo faster, less precise

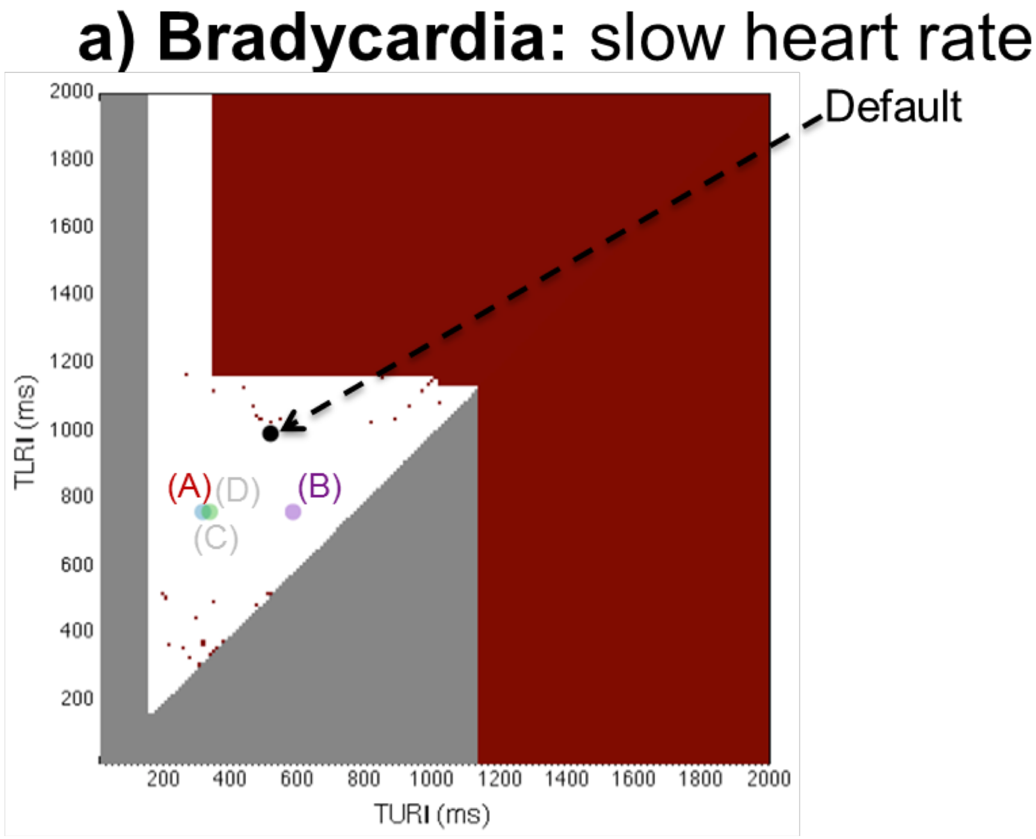

Synthesising robust and optimal parameters for cardiac pacemakers using symbolic and evolutionary computation techniques, Kwiatkowska et al.,In *Proc* HSB 2015

Case study: Personalisation

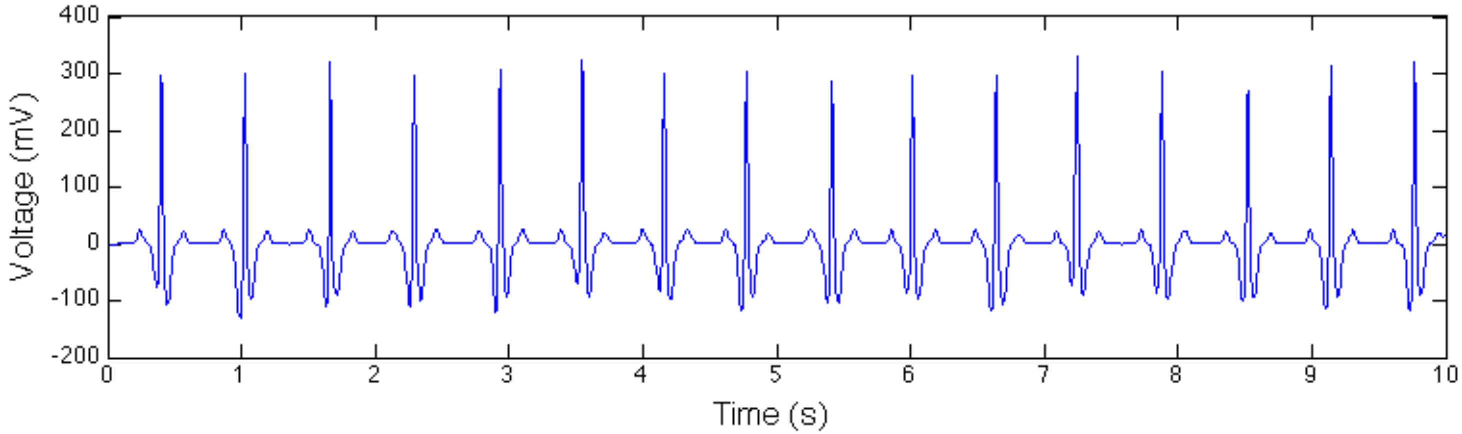

- Personalisation of wearable devices
  - estimate parameters for a heart model based on ECG data
  - generate synthetic ECG
  - useful for model-based development of personalised devices
- Developed HeartVerify based on Simulink/Stateflow
  - variety of tools and techniques
  - http://www.veriware.org/pacemaker.php

Estimation from ECG data

- Method for personalisation of parameters
  - filtering and analysis of the input ECG
  - detection of characteristic waves, P, QRS, T
  - mapping of intervals: explicit parameters
  - implicit parameters, eg conduction delays, use Gaussian Process optimisation
  - compare synthetic ECG with real ECG using statistical distance

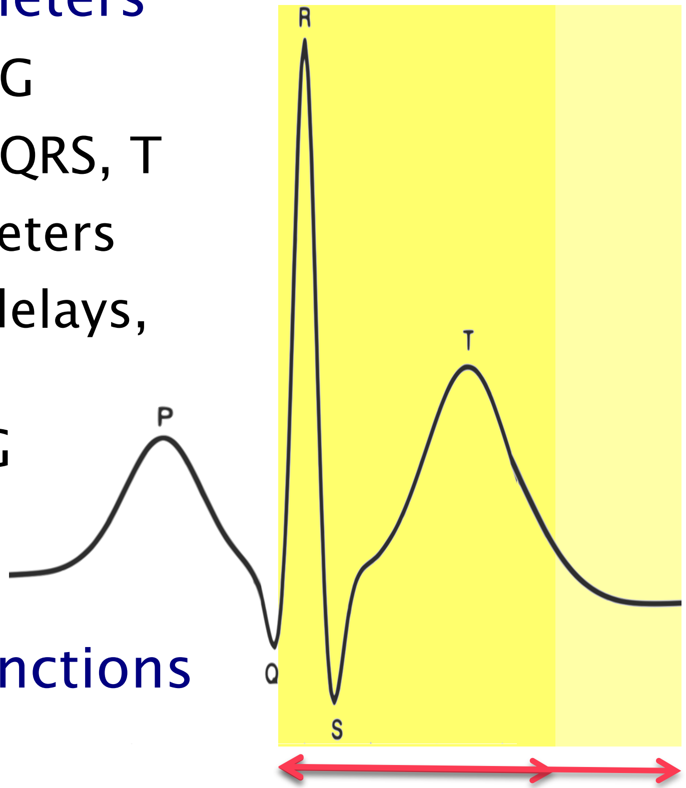

- Synthetic ECG = sum of Gaussian functions centred at each wave l_i

$$\mathsf{synthECG}(t) = \sum_{i \in \{P,Q,R,S,T\}} \sum_{l_i \in \mathsf{Peaks}_i} a_i \cdot \exp\left(-\frac{(t-l_i)^2}{2c_i^2}\right).$$

Statistical distance

- Computed between the filtered and synthetic ECG
- How similar are two signals?
  - returns value between 0 (identical) and 1
- Works by phase assignment
  - discretise the wave forms into discrete distributions,
  - then compute total variation distance

$$d(\mu_{i,p}, \mu_{j,p}) = \frac{1}{2} \sum_{x \in X} |\mu_{i,p} - \mu_{j,p}|.$$

  - finally compute the mean of the distances for each point

$$d(w_i, w_j) = \frac{\sum_{p \in P} d(\mu_{i,p}, \mu_{j,p})}{|P|}$$

- Method not affected by the heart rate

Raw ECG signal

- Real data

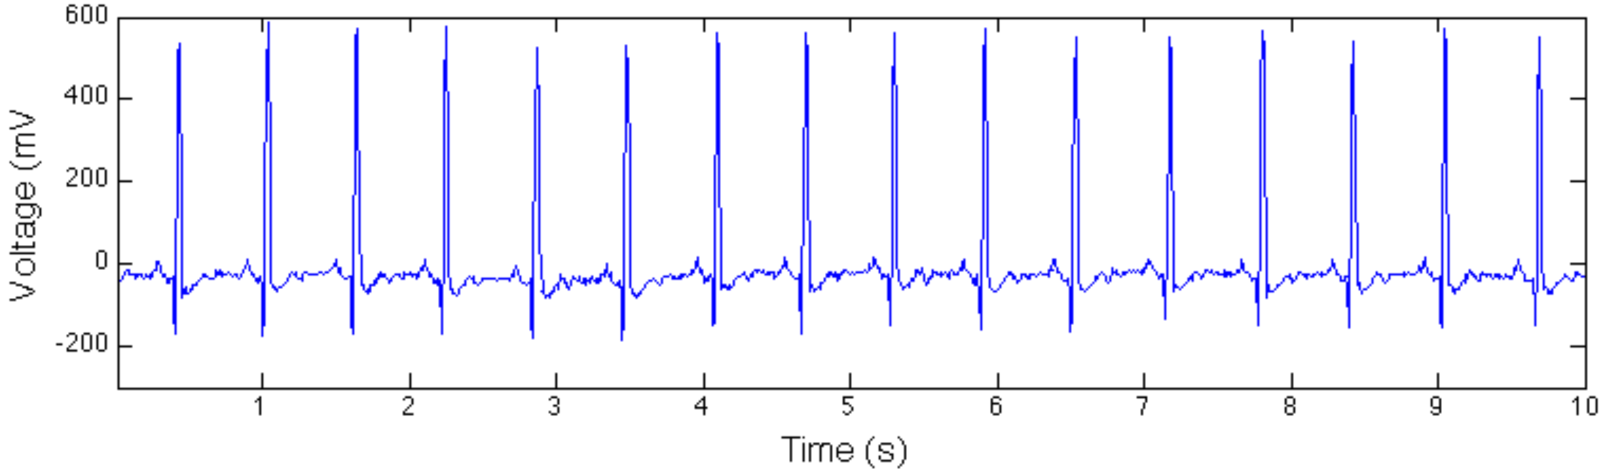

Filtered signal

- P,Q,R,S,T waves identified

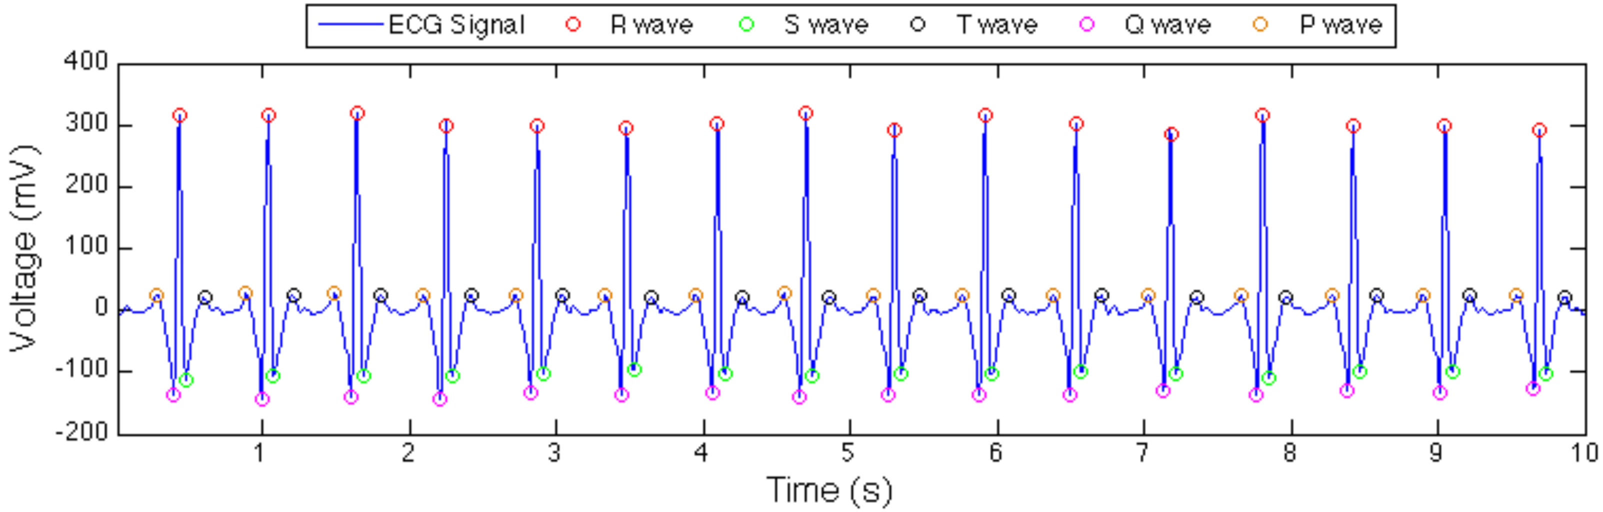

Synthetic ECG

- Produced by the personalised model

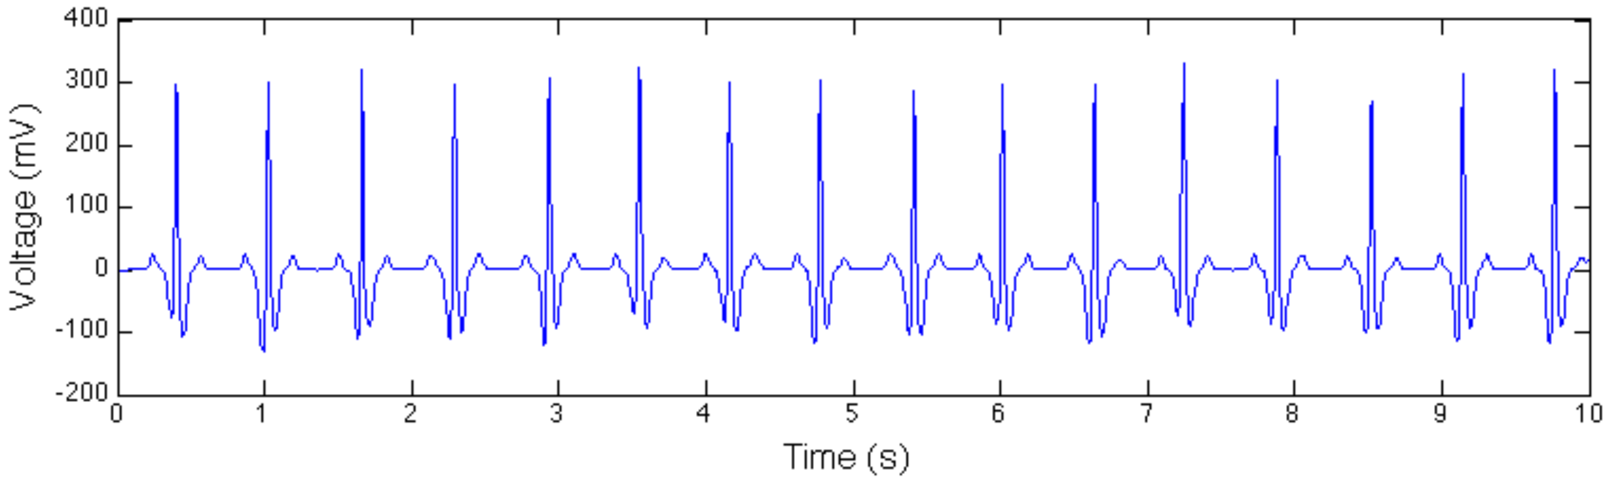

Wearable authentication devices

- Nymi band
  - ECG (Electrocardiogram) used as a biometric identifier
  - first creates biometric template
  - compares with real ECG signal when required
  - difficult to copy
- Can be paired with devices
  - with an app companion
- Proposed uses
  - for access into buildings and restricted spaces
  - for payment, etc

Case study: Attack on ECG biometrics

- ECG biometrics
  - increasing in popularity
  - Nymi band
  - are they secure?
- Synthetic ECGs
  - model-based: build model from data, 41 volunteers
  - inject synthetic signals to break authentication
  - 80% success rate
- Results
  - serious weakness
  - discuss countermeasures

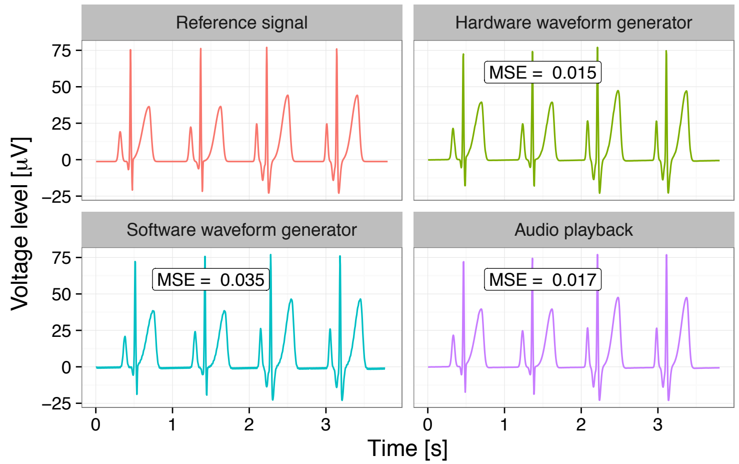

Broken Hearted: How to Attack ECG Biometrics, Ebertz et al., In *Proc* NDSS 2017

Case study: Transferability of attack

- Predicting how easy it is to attack biometrics when collecting data from different sources
  - ECG, eye movements, mouse movements, touchscreen dynamics, gait
- Model-based framework
- Features
  - amplitude for ECG
  - curvature for mouse
- Human study
  - easy for eye movements
  - ECG more chaotic

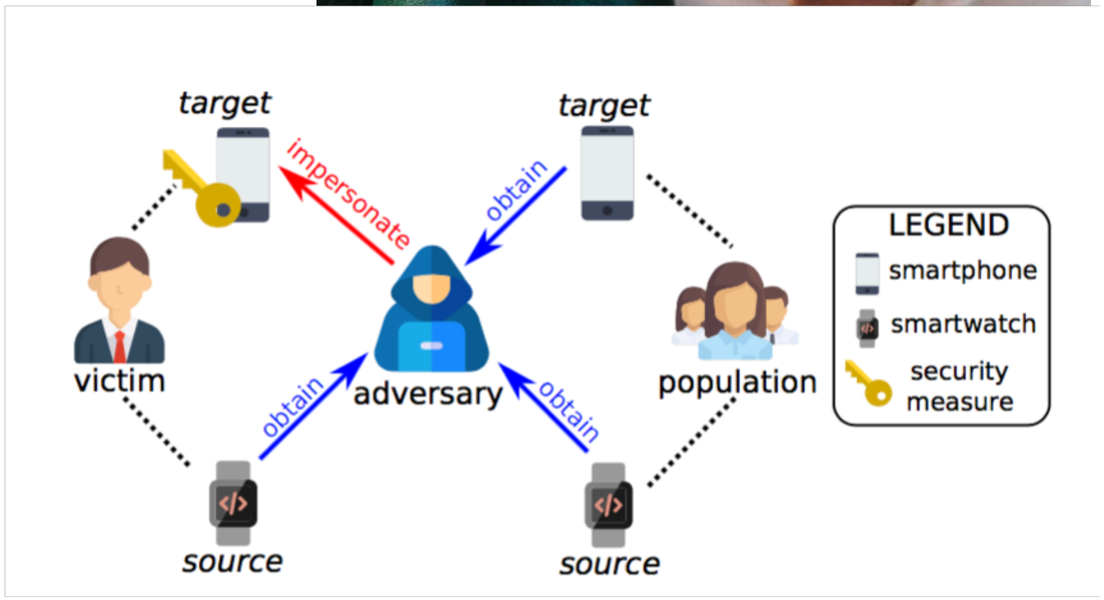

When your fitness tracker betrays you, Ebertz et al., In *Proc* S&P 2018

Case study: Intention anticipation

- Gaze tracking can reveal human intention
  - driver assistance
  - semi-autonomous driving
  - handover
- Predictive framework
  - model-based: build model from data, 124 cases from 75 drivers
- Model (ML+HMM)
  - anticipates intention 3.64 seconds before a real action was carried out
  - with 93.5% accuracy

(a) G3P14E2's intention from gaze.

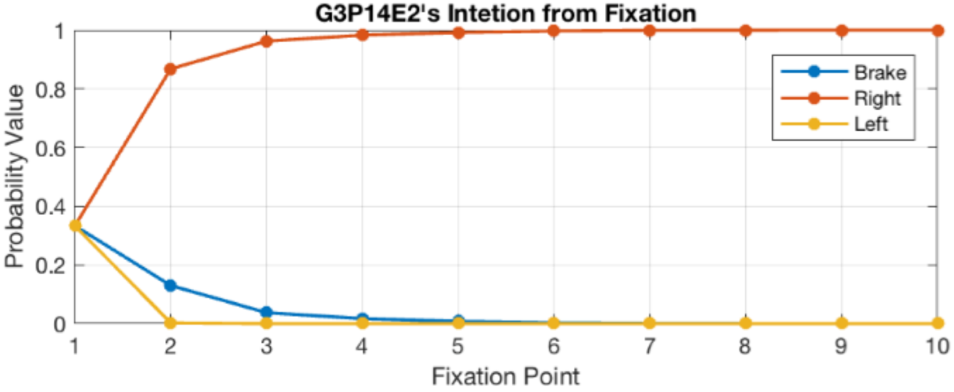

Gaze-Based Intention Anticipation over Driving Manoeuvres, Wu et al., submitted

Capturing emotion

- Affective analysis of physiological signals
  - electrocardiogram (ECG), electrodermal activity (EDA), breath rhythm, skin temperature, etc
  - single-source or multi-sensor fusion
- ECG signal attractive
  - unobtrusive, low cost, widespread, high sensitivity
- Conventional approach
  - multi-step process
  - extract heart rate (HR) and apply heart-rate variability (HRV) analysis
  - feature extraction, selection and calibration difficult

- Here propose end-to-end deep learning solution
  - supports personalization and feature calibration

AffecTech approach

Project topic: Personalisation and Verification for Affective Modelling.

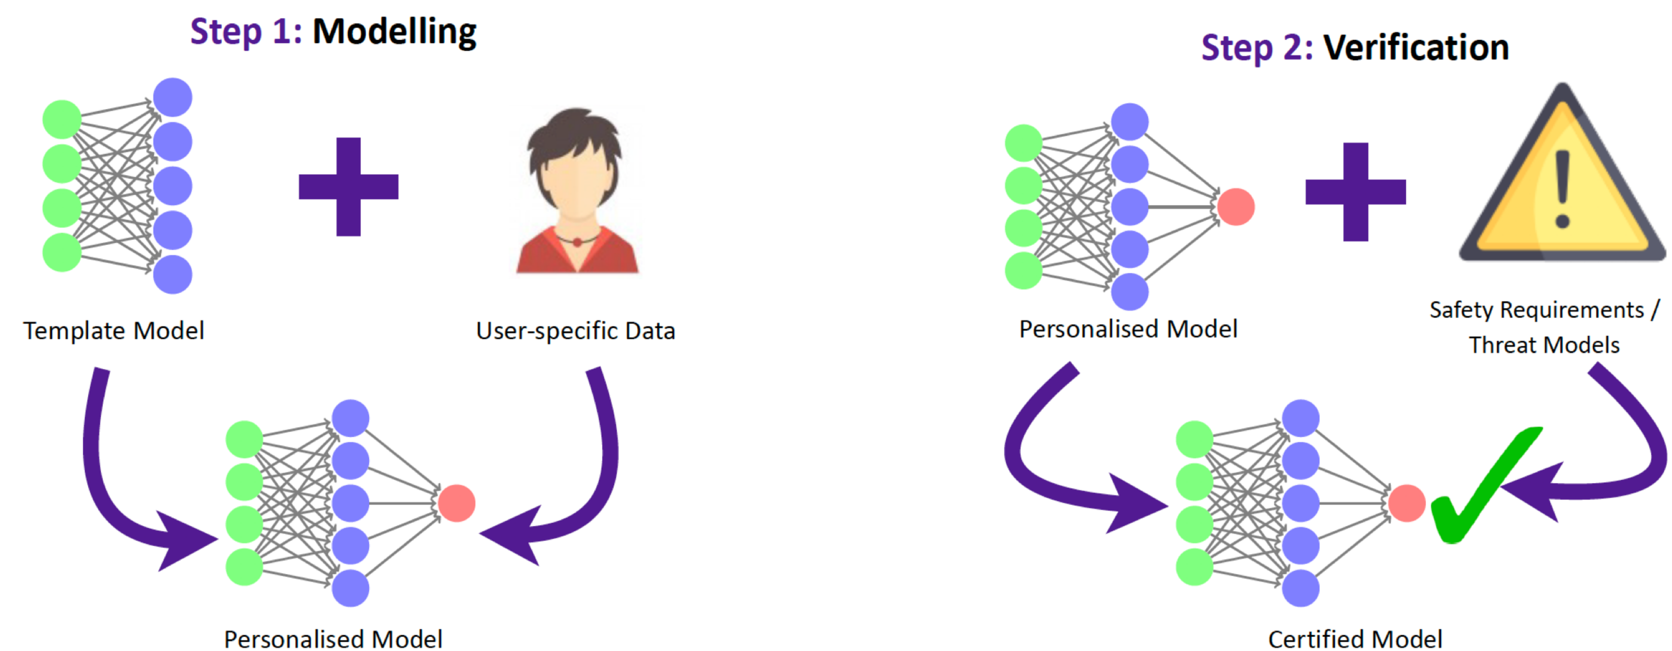

Deep learning solution

- Convolutional recurrent network (CRNN), classifier
  - ECG only, maxpooling to extract salient features
  - denoising of ECG signal
  - data augmentation, in view of sparsity
  - re-balancing, to deal with overrepresentation
- Key novelty: Siamese architecture (S-CRNN) to implement feature calibration
  - two copies of CRNN, sharing parameters
  - process user-specific template and data
  - template learn before experiment
- Evaluation on dataset for arousal in driving
  - binary low/high arousal

CRNN architecture

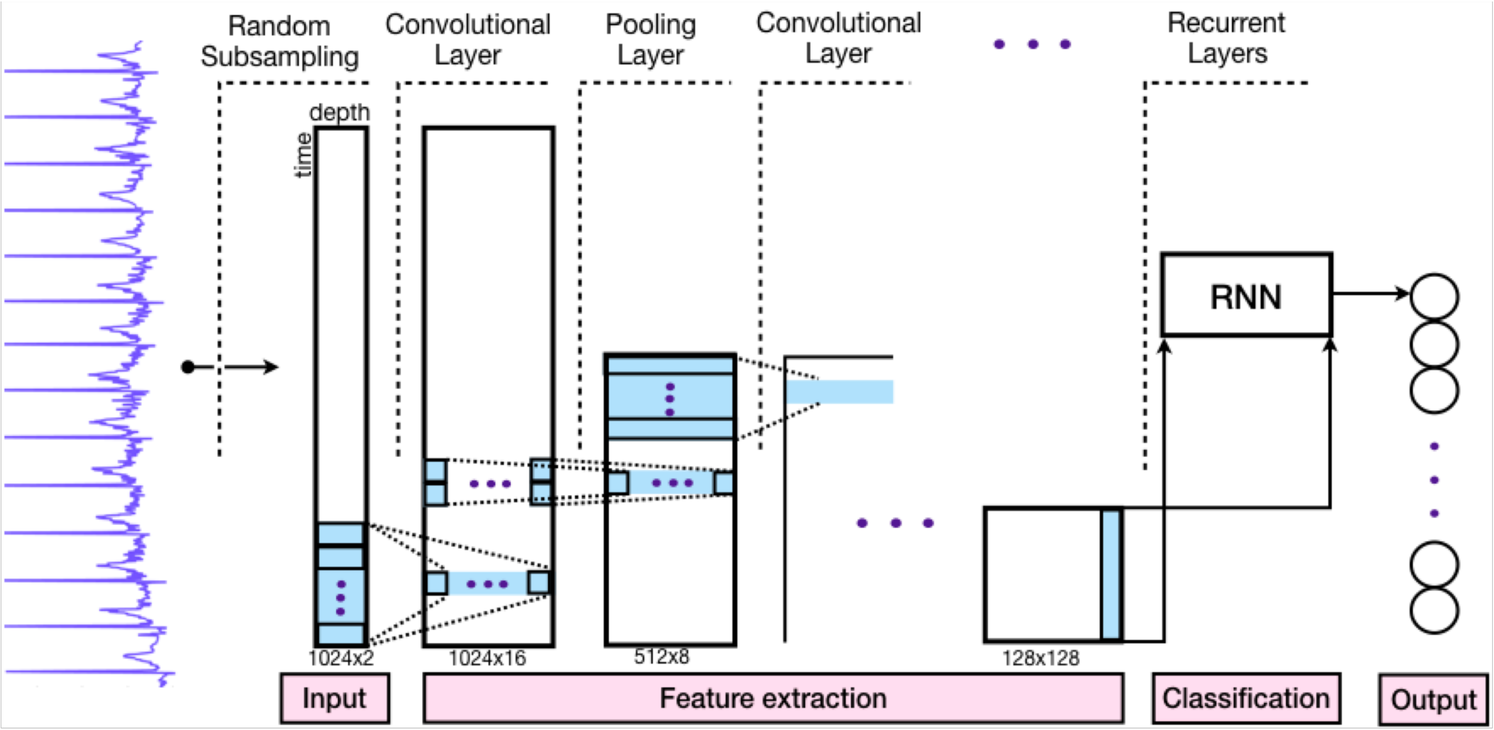

- Convolutional Recurrent Neural Network for arousal recognition

Siamese architecture

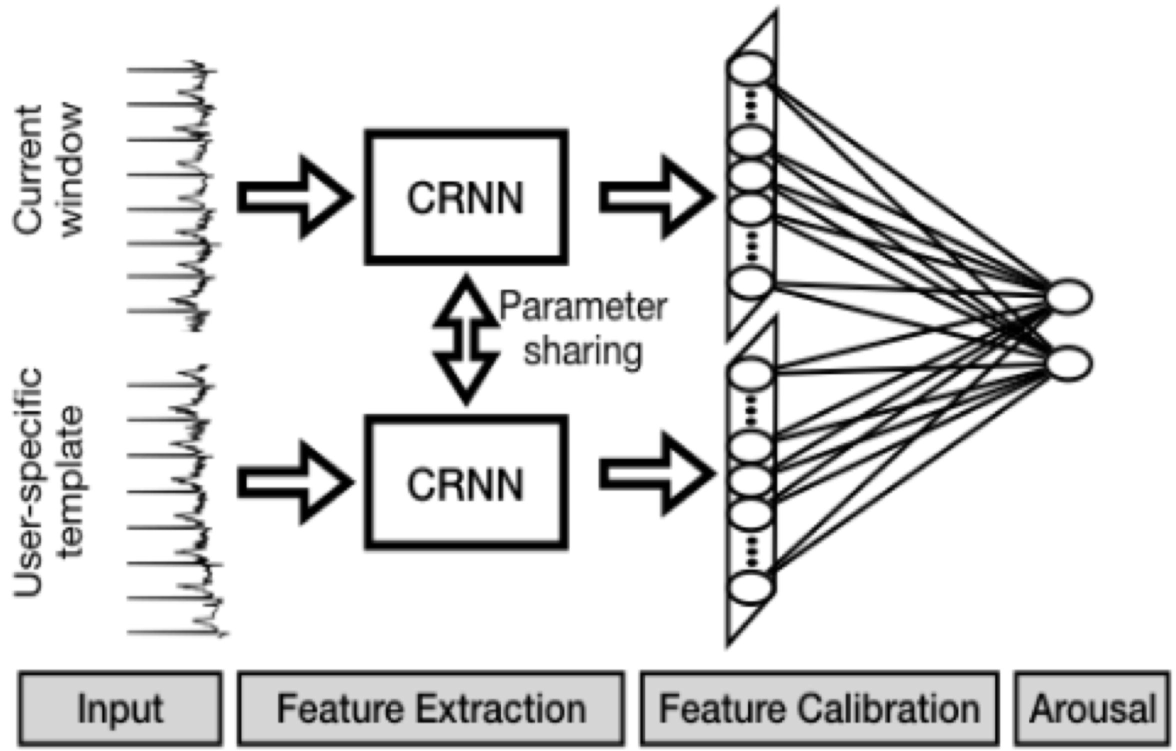

- Feature calibration (relative feature saliency)

Results for S-CRNN

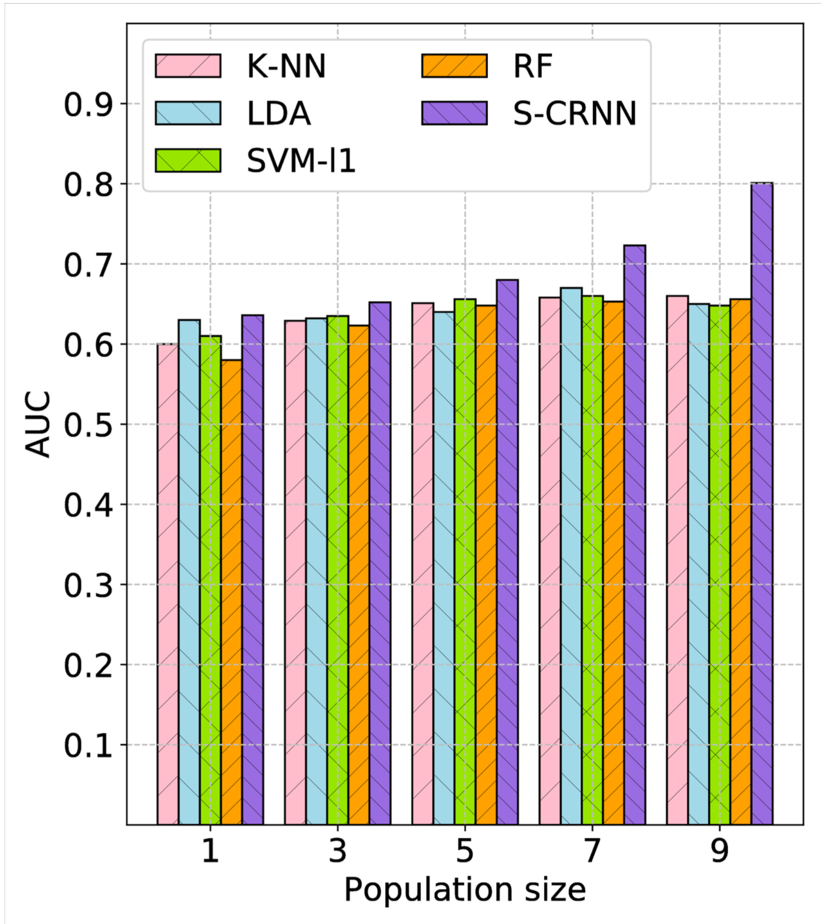

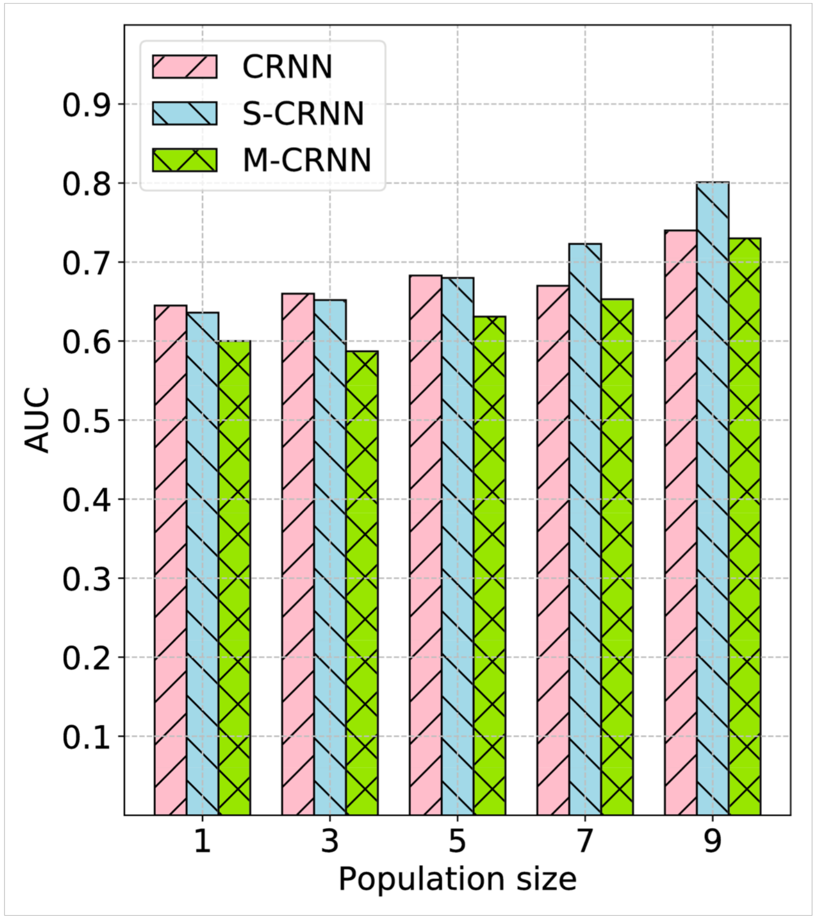

- 21% improvement over HRV analysis

Calibrating the Classifier: Siamese Neural Network Architecture for End-to-End Arousal Recognition from ECG, Patane and K., In *Proc* LOD 2018

PhysioNet Challenge 2018

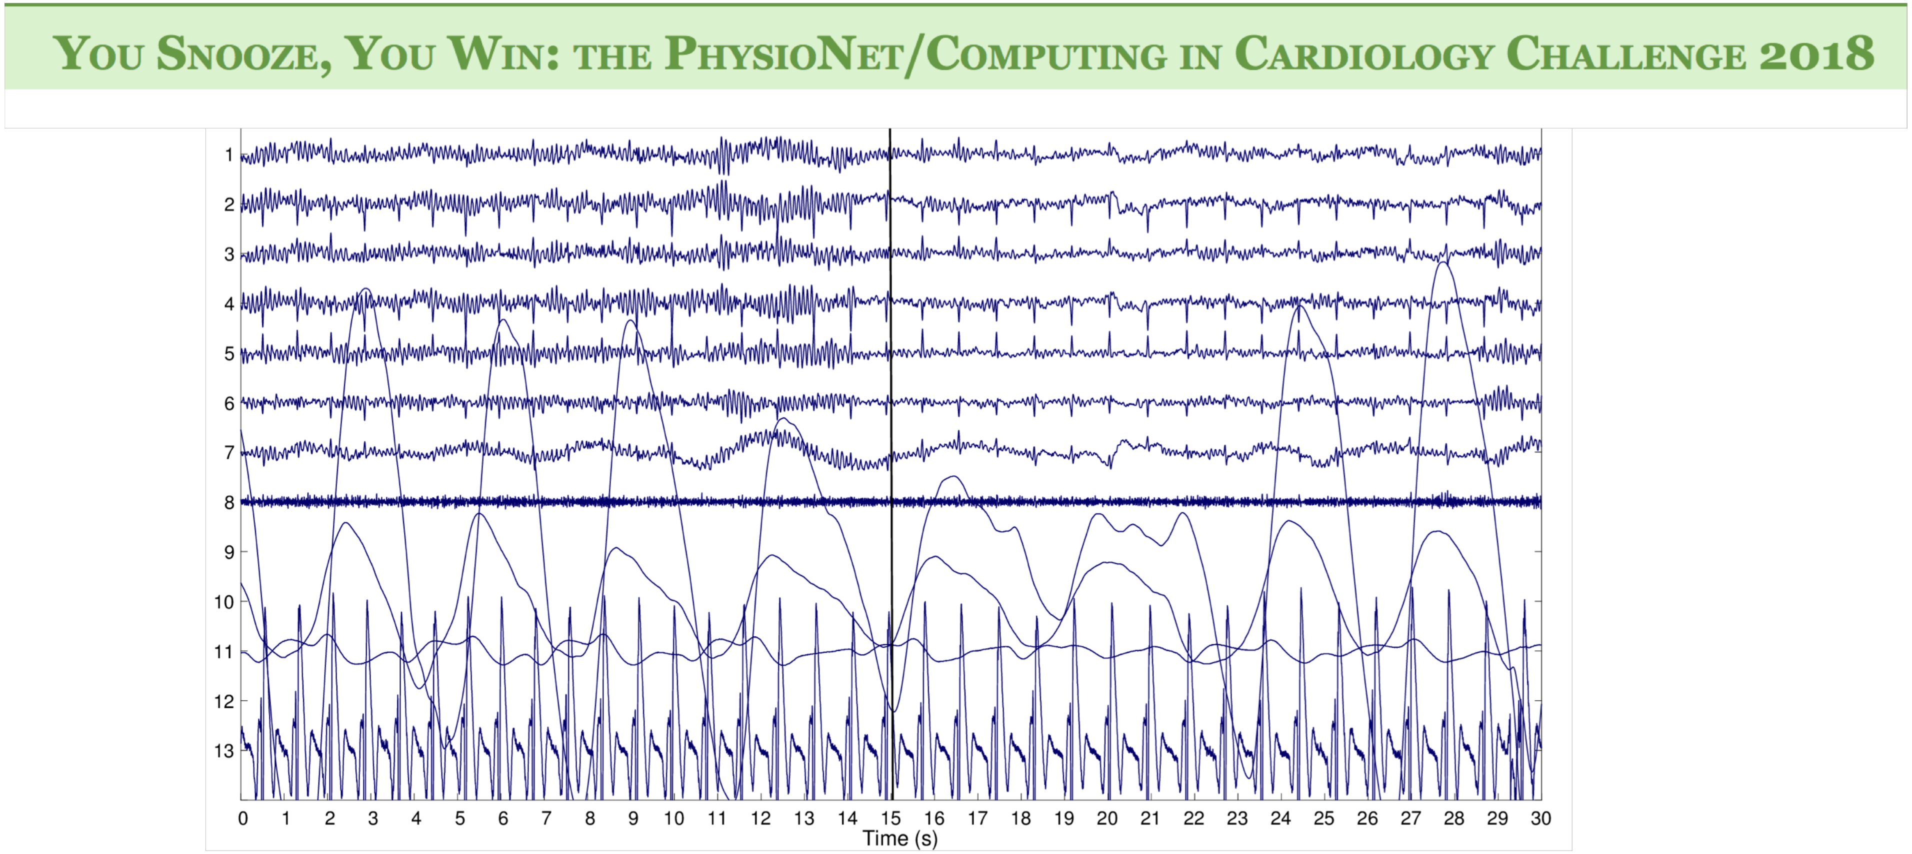

- Annual challenge to address significant unsolved clinical problems: classifying sleep arousals from EEG
  - Siamese architecture successful, placed 5th in competition

Automated Recognition of Sleep Arousal using Multimodal and Personalized Deep Ensembles of Neural Networks., Patane *et al*, In *Proc* CinC 2018

Next steps: probabilistic guarantees

- Need to probabilistic guarantees: probability that local perturbations result in predictions that are close to original

- Work with Bayesian inference and
- Gaussian processes (GPs)
- Define safety with prob $1-\varepsilon$

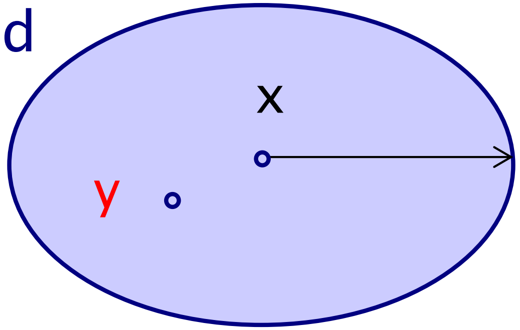

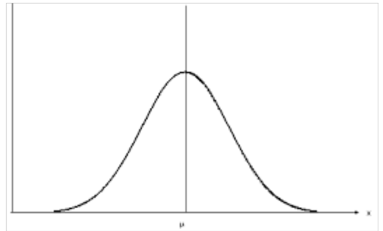

$$Prob(\exists y \in \eta \text{ s.t. } ||f(x)-f(y)||>\delta \mid D) \leq \varepsilon$$

- i.e. conditioned on training data D
- NB differs from pointwise thresholding in Bayesian deep learning

Robustness Guarantees for Bayesian Inference with Gaussian Processes., Cardelli *et al.*, In *Proc* AAAI 2019

Probabilistic guarantees for GPs

- Computation for general stochastic processes intractable
- For GPs, can obtain tight upper bounds by
  - approximating extrema of mean and variance for a test point
  - using Borell-TIS inequality
  - and solving optimization problems (analytical or convex opt)
- Applies to fully-connected (and convolutional) neural networks in the limit of infinitely many neurons…

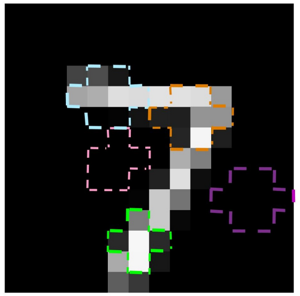

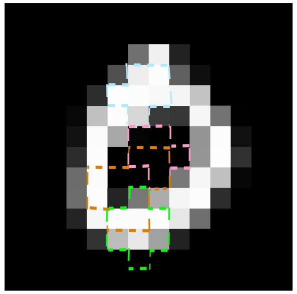

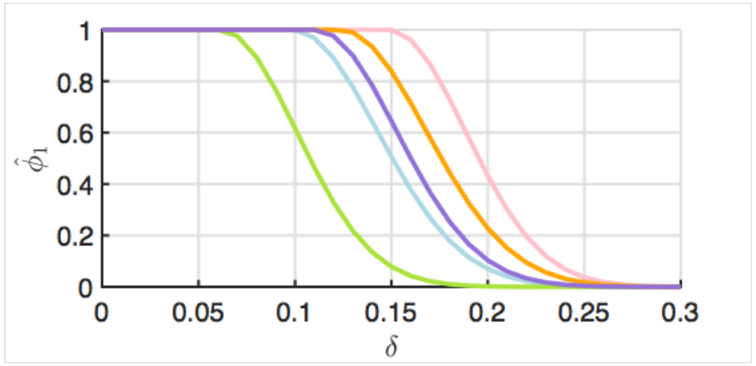

- Scalability continues to be an issue

Looking to the future…

- Progress towards emotion recognition from physiological signals
  - end-to-end deep learning architecture
  - personalisation and feature calibration
  - generalises to other contexts, good performance

- Future directions
  - robustness guarantees
  - synthesis of personalised intervention strategies
  - multi-modal sensor fusion
  - incorporation of contextual data
  - more complex disorders
  - intention prediction, biofeedback
  - brain machine interfaces, connection to neuroscience…

Acknowledgements

- My group and collaborators in this work
- Project funding
  - AFFECTech: Personal Technologies for Affective Health http://qav.comlab.ox.ac.uk/projects/erc-affectech/
  - VERIWARE www.veriware.org
  - PRISM www.prismmodelchecker.org
  - Mobile Autonomy Programme Grant: Safety, Trust and Integrity http://qav.comlab.ox.ac.uk/projects/epsrc-mobaut/
  - New ERC Advanced Grant FUN2MODEL – positions!!!

    "From FUNction-based TO MOdel-based automated probabilistic reasoning for DEep Learning"