SEQUENTIAL MODELS FOR PREDICTIVE DIAGNOSTICS

van der Westhuizen J., Lasenby J. - Cambridge University

UNIVERSITY OF CAMBRIDGE

Abstract

Sequential machine learning techniques allow the extraction of critical information embedded in hours of continuously monitored signals in the intensive care unit (ICU). The aim of this project was to provide clinicians with a continuous predictive probability of a patient's physiological status belonging to a range of conditions. Experimentation was done with Hidden Markov Models (HMMs), and Recurrent Neural Networks (RNNs). To date, both models have performed similarly, but more experimentation with deep RNNs is needed.

Datasets

- 3 NICU patients (Dead, Good, IVH)
  - ECG, ABP, SpO_2 (300 Hz)
- 4 ICU patients (Dead, Good)
  - ECG, ABP, ICP (300 Hz)
- 270 TBI patients (GOSm [1])
  - 9 Input channels (5s avg)

All of the sequences in the datasets were separated into smaller segments as shown below. A random split of 50:10:40 (training:validation:testing) was used for performance evaluation.

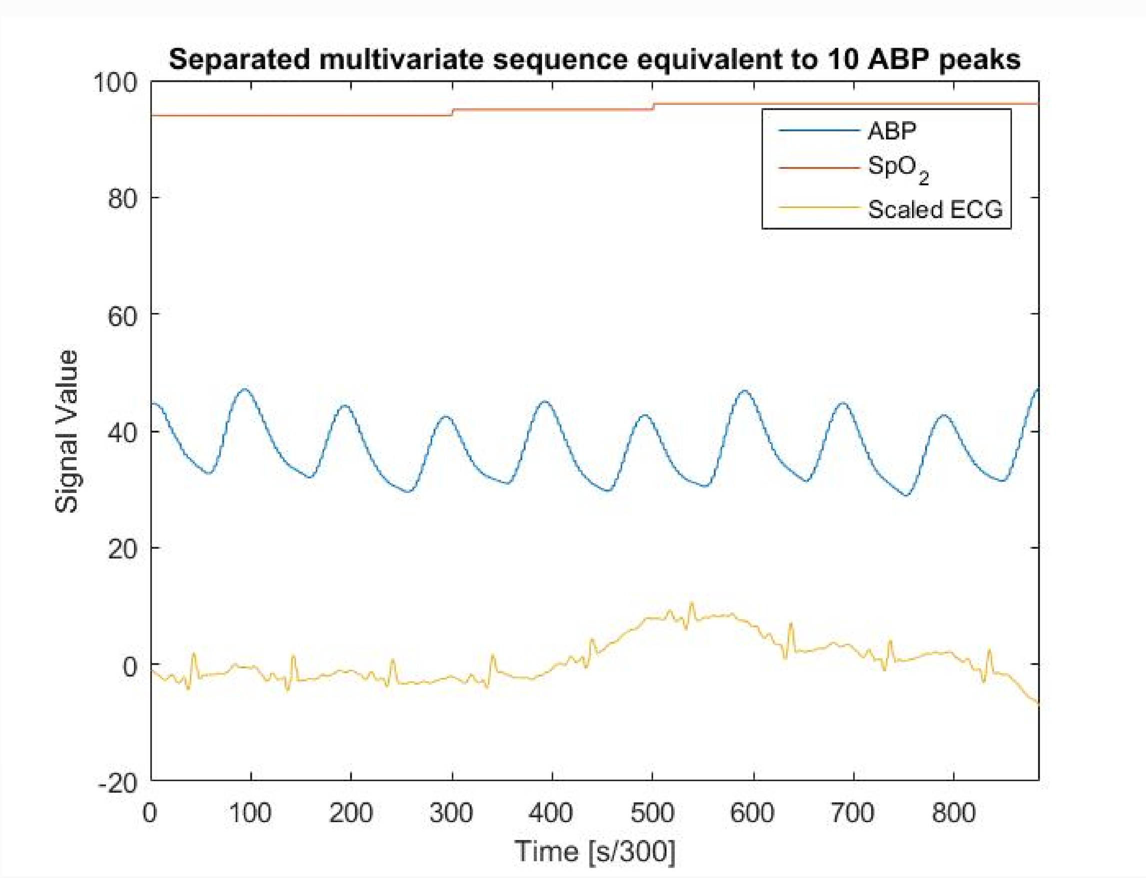

Long Short-Term Memory (LSTM)

The LSTM variant [2] of RNNs was used in this study. Illustrated below is the LSTM unit architecture on the left, followed by the two labeling structures experimented with, in this study.

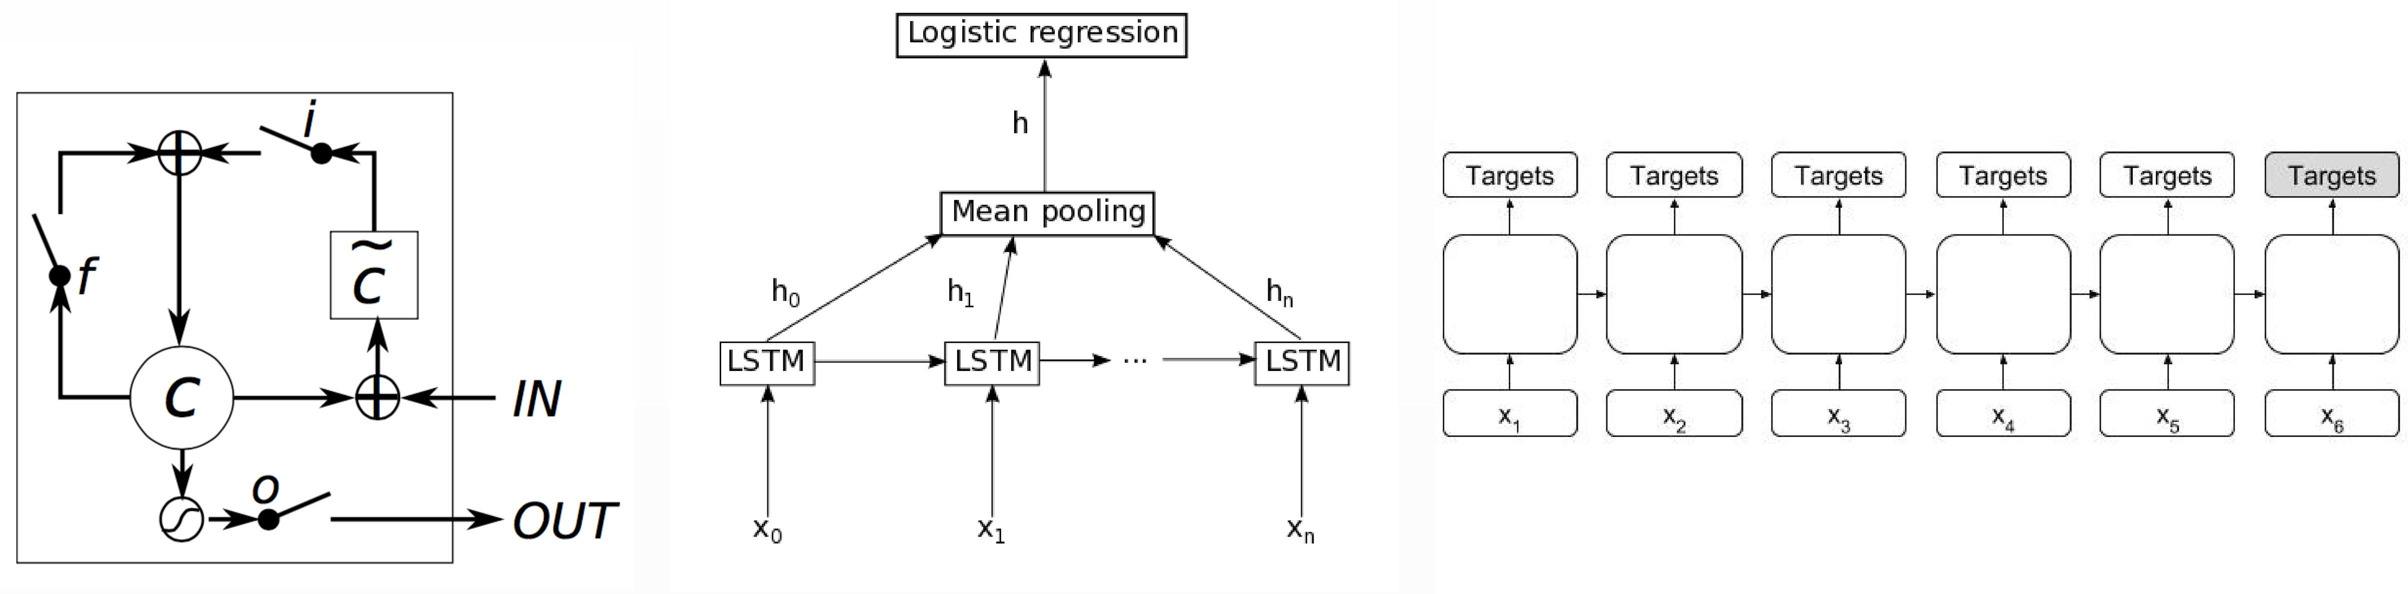

The middle figure illustrates the graph of the mean pooling labeling structure. The target replication structure, where the sequence label is repeated for each time step, is illustrated on the right. The dropout technique was employed in the models as a regularizer, where randomly selected weights are omitted from training during each iteration. The models were trained for 100 epochs, which is the number of times the model is presented with the entire dataset.

HMM

To determine whether HMMs perform better when trained on extracted features, Gaussian Processes (GPs) were used to smooth the input data. Best accuracies were achieved with the GP mean and standard deviation as input. The figure below illustrates the HMM log probabilities of 3 patients, with GP feature inputs, 7 hidden states, and 3 GMMs for output emissions.

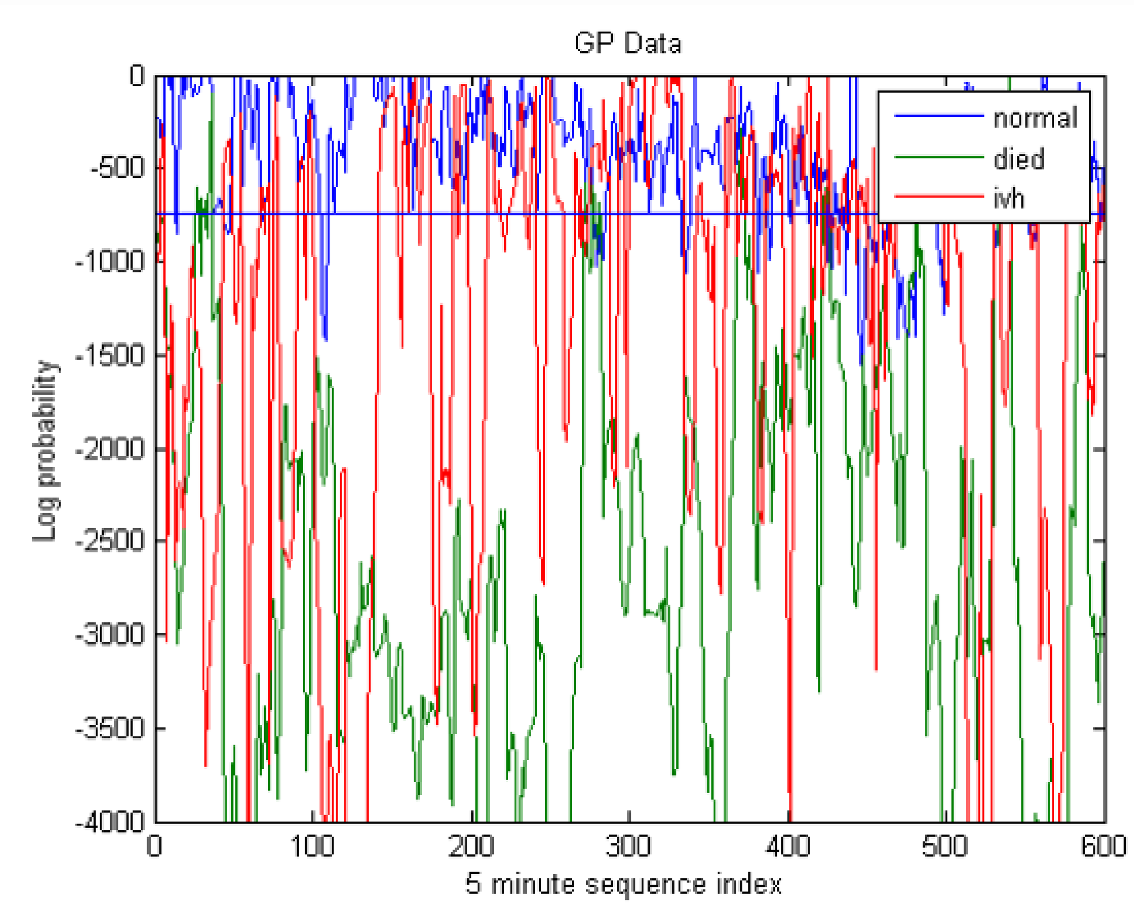

Results

Model Accuracies

Dataset	HMM	LSTM
NICU	90%	91%
ICU	N/A	100%
TBI	87%	85%

Best RNN performance was achieved with 3 layers of 100 LSTM units, 50% dropout, and target replication. The training curves for this model with the NICU dataset are illustrated below.

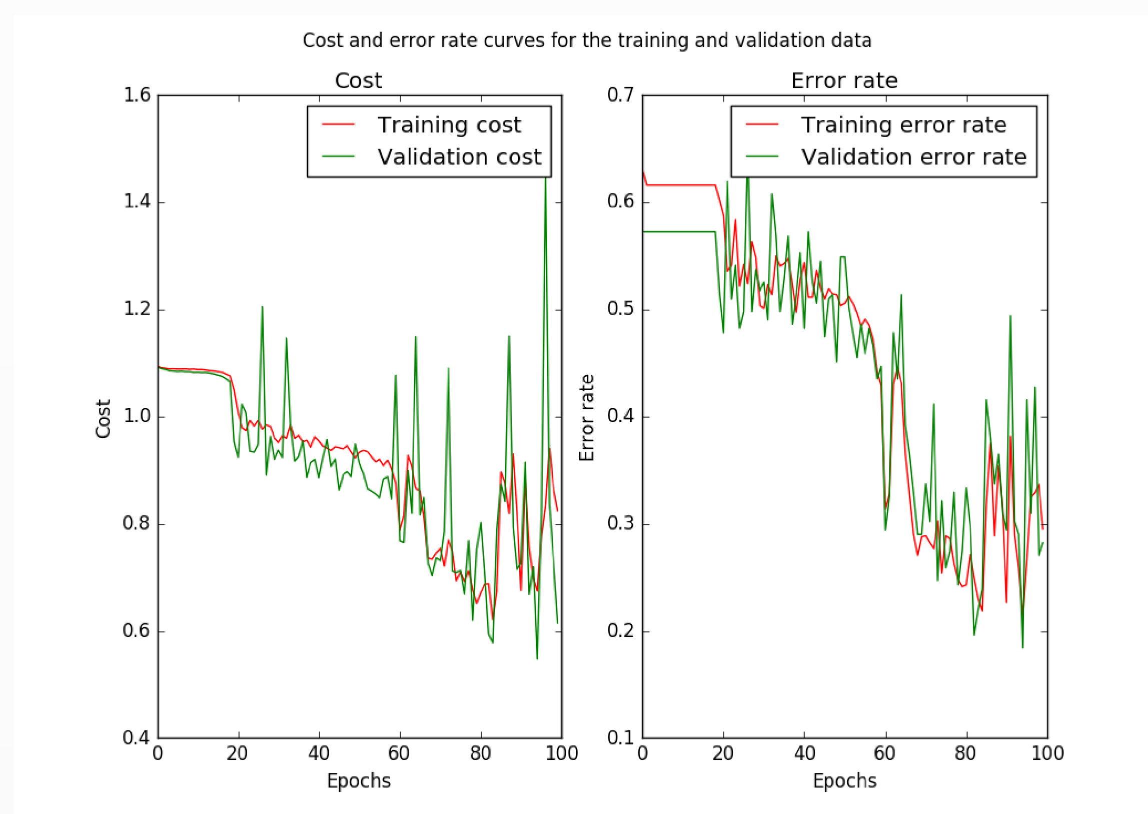

References

[1] B., Jennett, M., Bond, Assessment of outcome after severe brain damage: a practical scale, in *The Lancet*, 1975

[2] A., Graves, Generating sequences with recurrent neural networks, in *arXiv preprint arXiv:1308.0850*, 2013

Acknowledgements

The authors would like to thank the Sky Cambridge Trust for funding the study. The W G Collins Endowment Fund and Christ's College is thanked for funding provided to attend the ICVSS.

Future Work

- Apply LSTM to other datatypes
- Evaluate performance on larger medical datasets
- Reverse engineer model to understand underlying biological processes