Understanding Circulatory Responsiveness to Fluids in Children with Shock

James Stainer; Supervisors: Mark Peters[1], Samiran Ray[2], Simon Arridge[3]

[1]Great Ormond Street Hospital, NHS;
[2]Dept of Computer Science, UCL

Background

When a patient goes into shock, there is insufficient blood flow to tissues around the body. This can lead to organ damage and sometimes death, particularly if untreated. Fluids are used to treat shock by artificially increasing blood volume, which should improve cardiac output, therefore increasing organ perfusion.

Recent research found that fluid boluses significantly increased mortality in critically ill children with shock in resource-limited settings, which opposes the common belief that fluids are essential during shock.[1]

This research aims to characterise the response to fluid resuscitation, through applying artificial intelligence methods to time-series data, in order to enable greater decision making on whether fluids should be administered.

Data

Data for this project has been provided by GOSH and includes deidentified data for every patient admitted to the PICU between 2016-2018. The dataset includes monitor data recorded every five seconds for the entirety of a patients stay in the PICU at GOSH, although the variables recorded for each patient differ due to differing needs during their stay. For example, end-tidal CO_2 levels are only measured when a child is intubated. Each patient in the dataset typically share the same core variables measured such as HR, BP, and respiratory rate.

Electronic health record (EHR) data will be used to link patients that have received a fluid bolus to monitoring data. Fluid bolus administration is not always immediately recorded as clinicians prioritise giving urgent care, therefore it is necessary to implement a mapping algorithm between the monitoring data and EHR data - this can be done by utilising change-point detection algorithms or using machine learning classifiers.

Patients who are being treated on the cardiac ward are excluded from the dataset. Patients on the cardiac ward will typically have artificially adjusted HR, so the relationship between their HR and BP would be atypical and generate unwanted bias.

Figure 1: Example of Time-Series Data Available

Data Pre-Processing

The raw data contains unwanted artefacts due to either malfunctioning recording software or medical procedures requiring disconnect of equipment. This means that outliers and missing data are common, as shown in Figure 1.

The data pre-processing pipeline taken was as follows:

- Identify Missing Data
- Impute Missing Data
- Remove Outliers
- Resample onto Larger Timeframe

Outliers

Data that is anomalous can bias results and therefore affect the reliability of methods applied. To minimise the effect of outliers a Hampel filter is used to detect and replace these values. A Hampel filter is a rolling window that replaces the central value with the median where the central value is significantly larger than the median.

Impute Missing Data

For missing data with variables such as HR and BP, there exists measurements from which the data can be derived. For example alongside HR, each patient's pulse rate is measured alongside their HR. Although, there may be minor discrepancies between the two measurements, the pulse rate measurement provides an accurate representation of the HR.

Mean arterial pressure (MAP) is the average BP in an individual during a single cardiac cycle and is measured alongside systolic and diastolic BP in the dataset. MAP can be calculated as

$$MAP = \frac{\text{systolic} + 2 \cdot \text{diastolic}}{3},$$

Therefore, when one variable is missing, we can compute the missing value and use this for imputation.

Methods of Analysis

The structure of the research being undertaken is as follows:

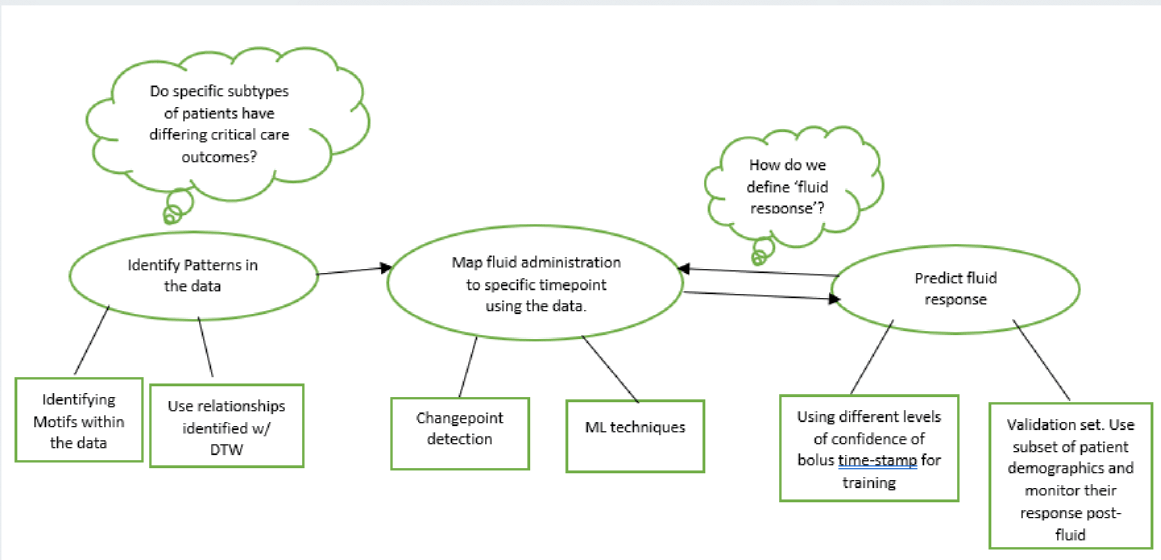

Figure 2: Research Overview - Pattern identification has been researched in Year 1 through use of Dynamic Time Warping (DTW). Prediction of fluid response using supervised methods is currently being researched as part of Years 2-4.

The three main goals of this research are:

- Identify patterns within critical illness in PICU patients in critical care
- Characterise response to fluids
- Predict exact physiological response to fluids through supervised learning techniques

Methods of Analysis

Dynamic Time Warping

DTW is a method of identifying similarities between two time-series, where each sequence varies in length [2]. The package TSLearn in python implements k-means and Kernel k-means Dynamic Time Warping clustering estimators which has been used to explore relationships in BP and HR between patients [3]. Using this method, we found a subgroup of patients who were likely to have worse health outcomes on average.

DTW Results

Without age-standardising each variable, DTW poorly clustered the data. However, after standardising for age the algorithm identified subgroups that were particularly vulnerable to worse outcomes. The clustering task performed on shock index ($\frac{\text{HR}}{\text{systolic BP}}$) identified a cluster of 29 patients all with significantly worse patient outcome than the dataset average - 31% mortality rate, mean length of ventilation 10 days and 2 hours, and mean length of stay 8 days and 12 hours.

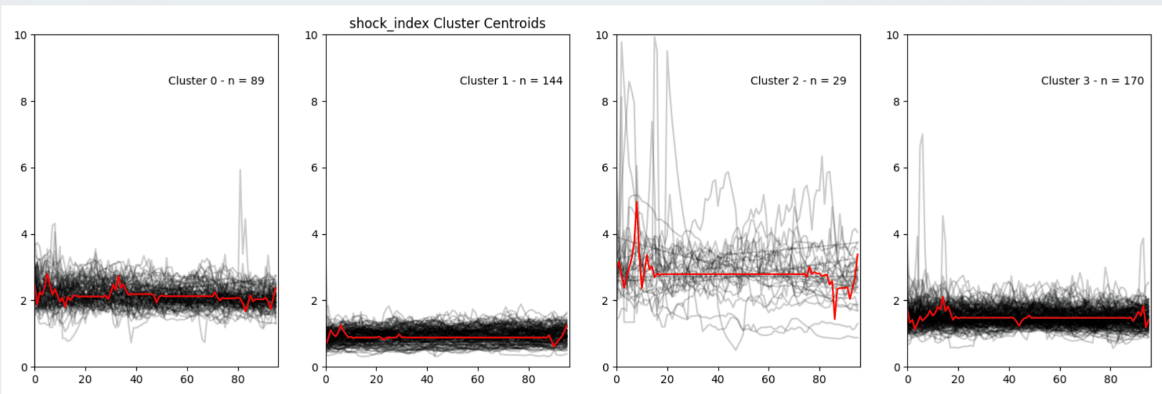

Figure 3: Shock Index Centroids using DTW K-means Clustering

Recurrent Neural Networks

Supervised learning methods can be applied to the dataset used in this research to predict the exact physiological response to fluids. For example, using methods such as recurrent neural networks (RNNs), it is possible to predict HR and BP at some future timepoint after fluid administration. RNN are a type of neural network that uses time series data. Their ability to 'memorise' past events means that predictions can be made through utilising key information from previous timepoints.

Our current research aims to harness RNN to accurately predict the exact physiological response to fluids, at some time point in the future. If successful, this tool will be able to assist clinicians in determining whether fluid administration is likely to cause an unwanted response e.g. a patient having only a transitory positive response to fluids before deteriorating into a worse state and would have been better off with an alternative to fluids.

References

- Kathryn Maitland et al. "Mortality after Fluid Bolus in African Children with Severe Infection". In: *New England Journal of Medicine* 364.26 (June 2011). PMID: 21615299, pp. 2483–2495. DOI: 10.1056/NEJMoa1101549. URL: https://doi.org/10.1056/NEJMoa1101549.
- "Dynamic Time Warping". In: *Information Retrieval for Music and Motion*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 69–84. ISBN: 978-3-540-74048-3. DOI: 10.1007/978-3-540-74048-3_4.
- Romain Tavenard et al. "Tslearn, A Machine Learning Toolkit for Time Series Data". In: *Journal of Machine Learning Research* 21.118 (Apr. 2020), pp. 1–6.

james.stainer.20@ucl.ac.uk