LSTM Modeling of Perinatal Fetal Heart Rate

Philip A Warrick, Emily F Hamilton

PeriGen Inc.
Montreal, Canada

Objectives: Characterizing the fetal heart rate (FHR) through modelling can give insights into the fetal cardiac state and overall fetal well-being. In this preliminary study we modelled FHR using a Long Short Term Memory (LSTM) cell.

Methods: Because there are several non-linearities within an LSTM cell, we took some precautions to avoid overfitting. We did this in two ways. By restricting the model to a single LSTM cell, the model was limited to the 14 LSTM weight and bias terms. Secondly, we applied several degrees of dropout to the inputs to force learning to adapt to missing input data. The data consisted of FHR perinatal recordings from 12 fetuses with durations ranging from a few hours to 14 hours. The data was median filtered then low-pass filtered, and finally subsampled to 0.25 Hz. We used ten recordings for training, one for validation and one for testing. We trained for 1000 epochs. To compare prediction accuracy, we used the normalized measure % variance accounted for (VAF). VAF is defined as $100 \times (1 - \sigma_e^2/\sigma_y^2)$ where σ_e^2 and σ_y^2 are the mean-square error and the target signal variance, respectively.

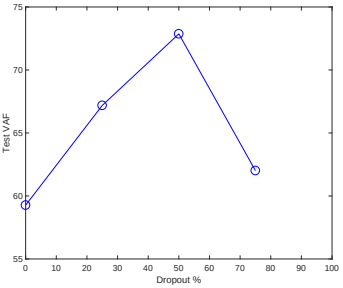

Variance accounted for (VAF) vs. dropout percentage for test case.

Results: The figure compares the VAF of the test file for % dropout values of 0, 25, 50 and 75, respectively. The highest VAF occurred for % dropout = 50 and all non-zero values of dropout performed better than applying no dropout at all.

Conclusions: Dropout was effective at reducing overfitting, resulting in a better predictor for the case studied. In the next phase of this work, we will apply this technique to a larger dataset and examine the effect of using multiple LSTM cells.