Pair Recurrent Neural Network (PaiRNN)

Zhymir Thompson, Austin R.J. Downey, and Jason Bakos

UNIVERSITY OF SOUTH CAROLINA

RNN

- What is an RNN?
  - Recurrent Neural Network
  - Takes in sequence of inputs
  - Maintains ‘state’ to learn
- How does it differ from a neural network?
  - Takes one input at a time
  - Inputs are independent (current input is not affected by previous input)

RNN

- How do RNN’s work?
- On forward propagation:
  - Input and hidden state passed in initially
  - Output prediction and hidden state passed back into RNN until desired iterations complete
- On backpropagation:
  - Backpropagates ‘through time’
  - Compounded loss with decay over state

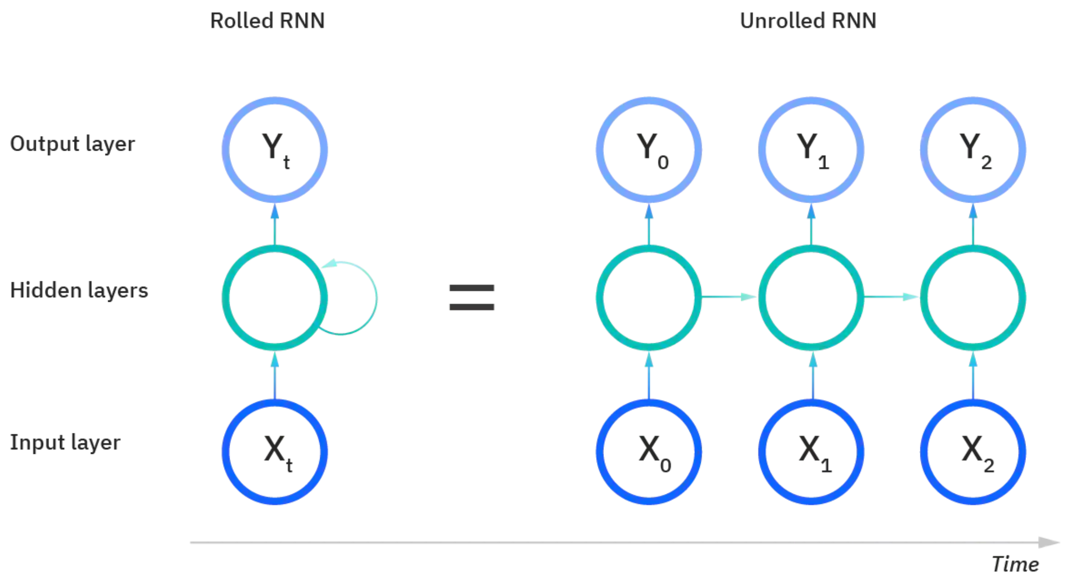

IBM RNN image[1]

PaiRNN Design

- Identical twin RNN models
  - Predictor – Makes inferences on future given current data
  - Learner – Makes inferences on current data given historical data; uses inferences to adjust weights
- Stream of data
- Queues – Containers to receive data from stream for model inferencing
- Predictor and learner both inference on respective data
- Learner adjust weights, and sends new weights to predictor

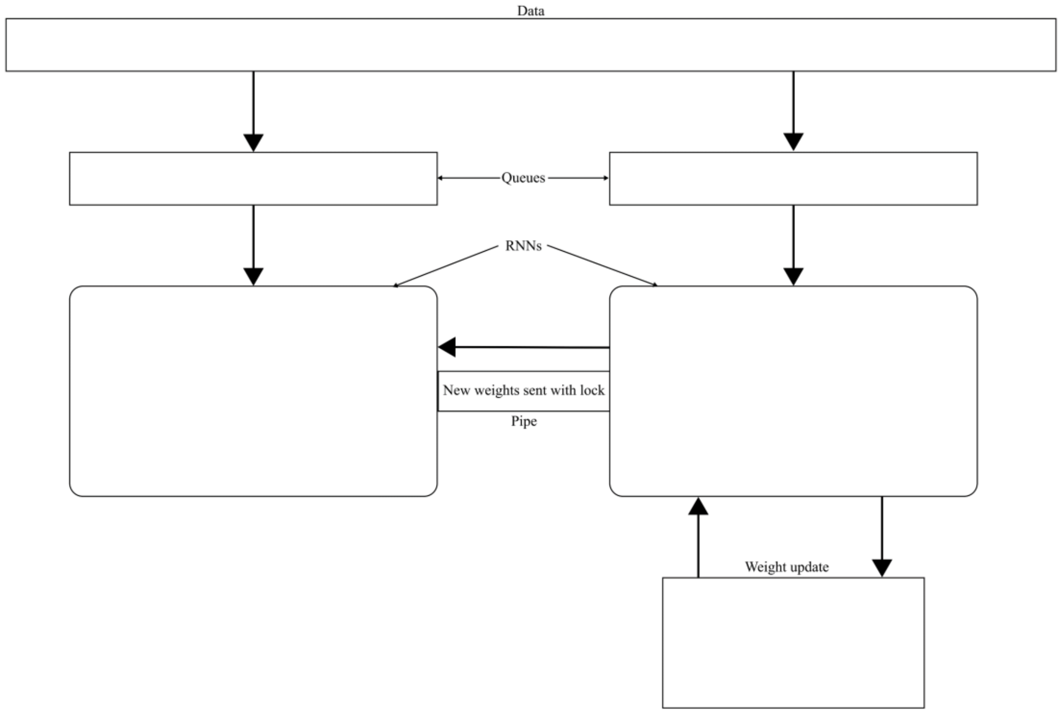

PAIRNN PERFORMANCE

- Compared prediction vs observed signal at various time slices
- Converges over 3(s)
- Predicts well once converged
- Adjusts well albeit imperfectly to signal post-nonstationarity event

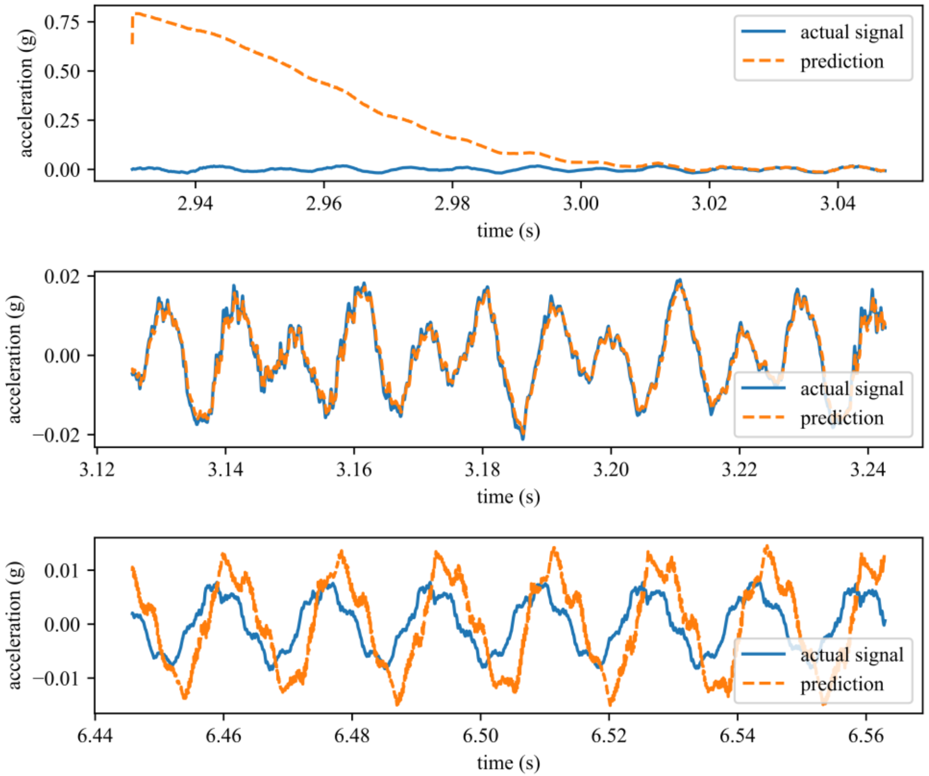

COMPARISON

MLP

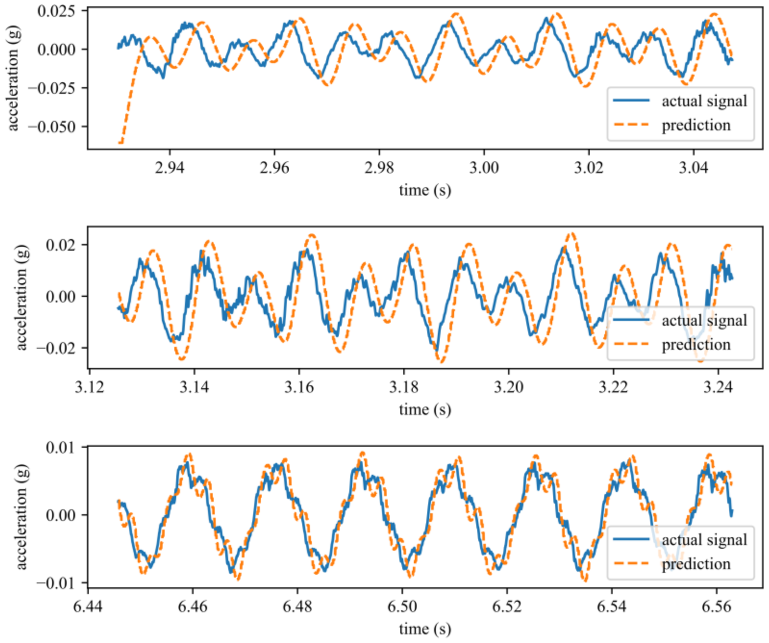

RNN

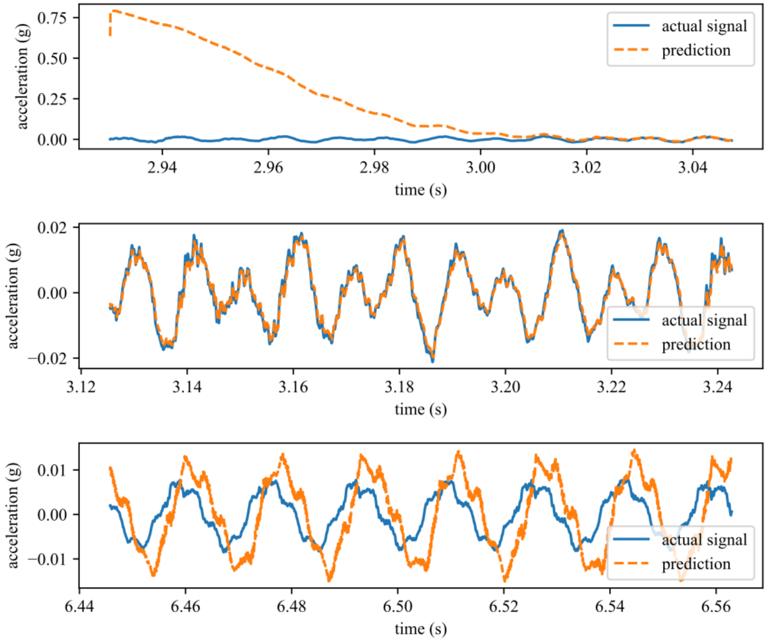

References

- https://www.ibm.com/cloud/learn/recurrent-neural-networks