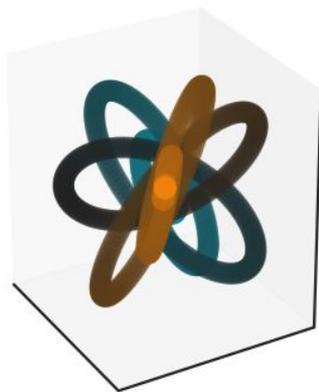


Legendre Memory Units (LMUs)



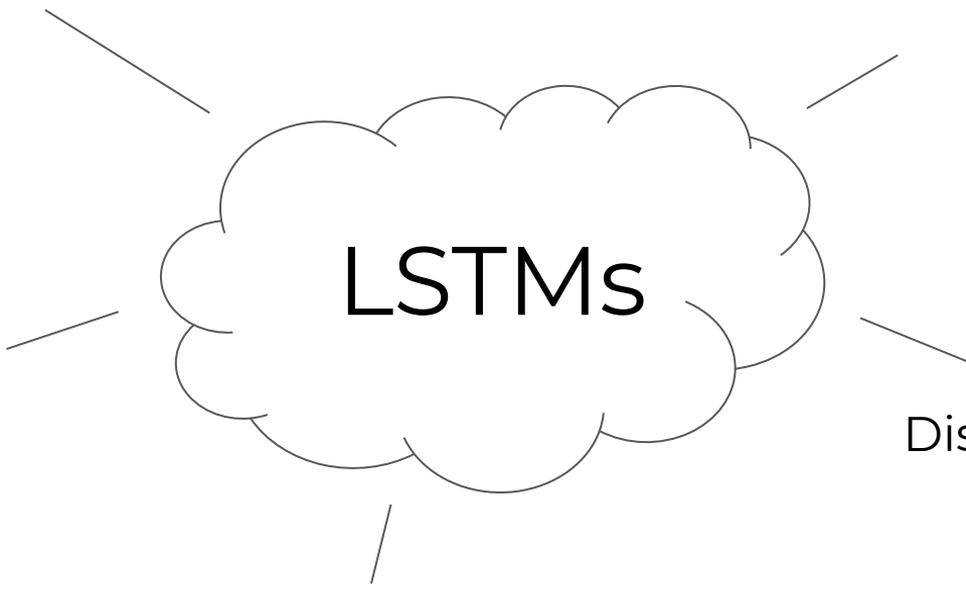
Continuous-Time Representation in
Recurrent Neural Networks

Aaron Voelker, Ivana Kajić, Chris Eliasmith



Image caption
generation

Music composition

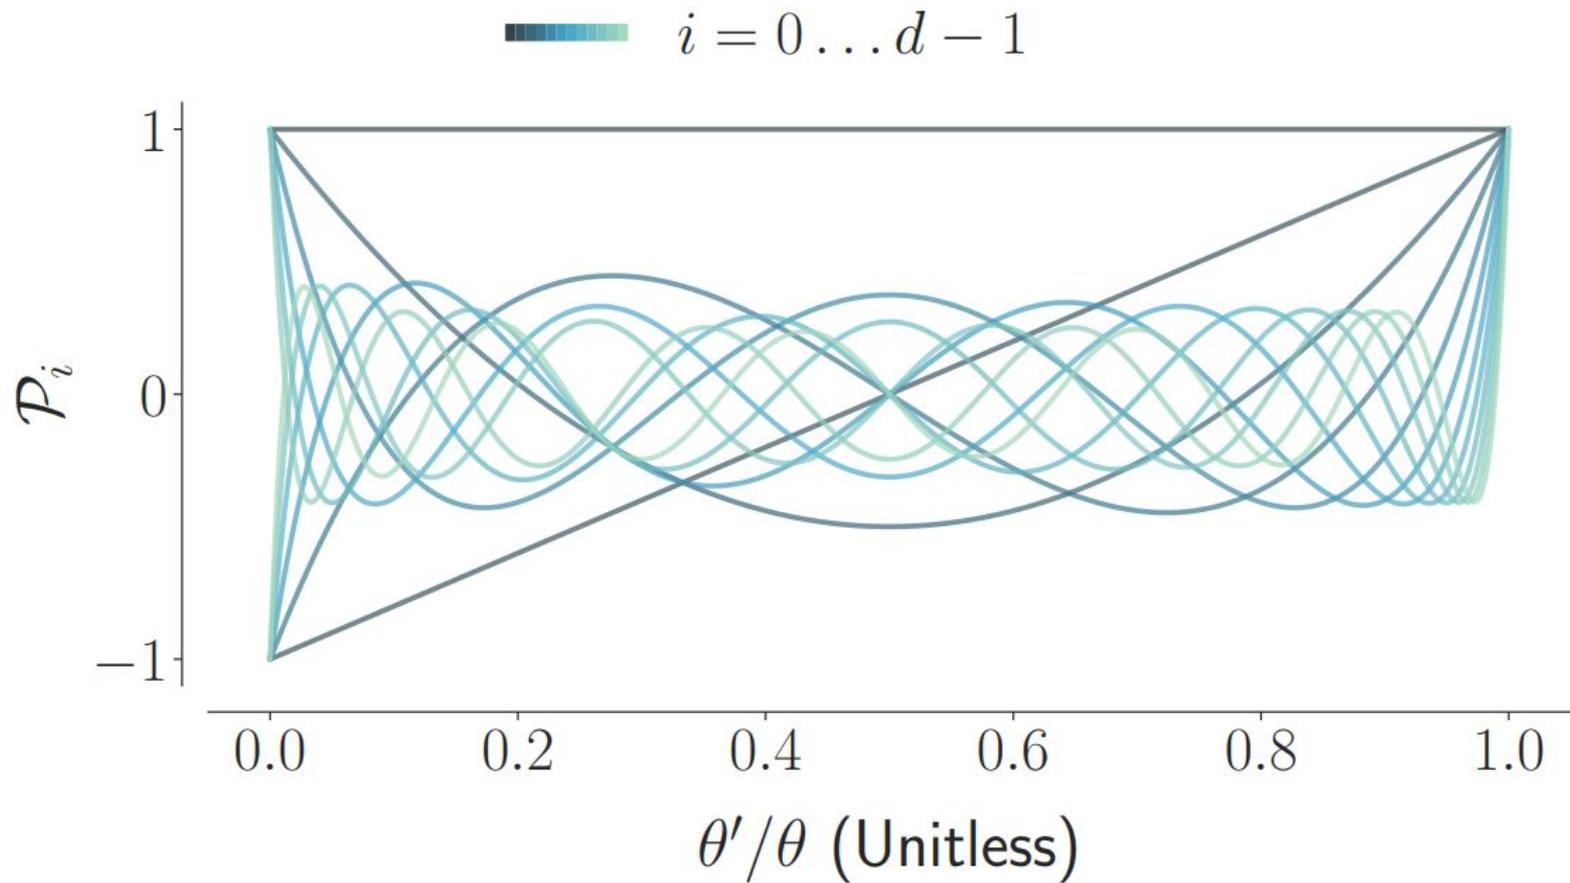


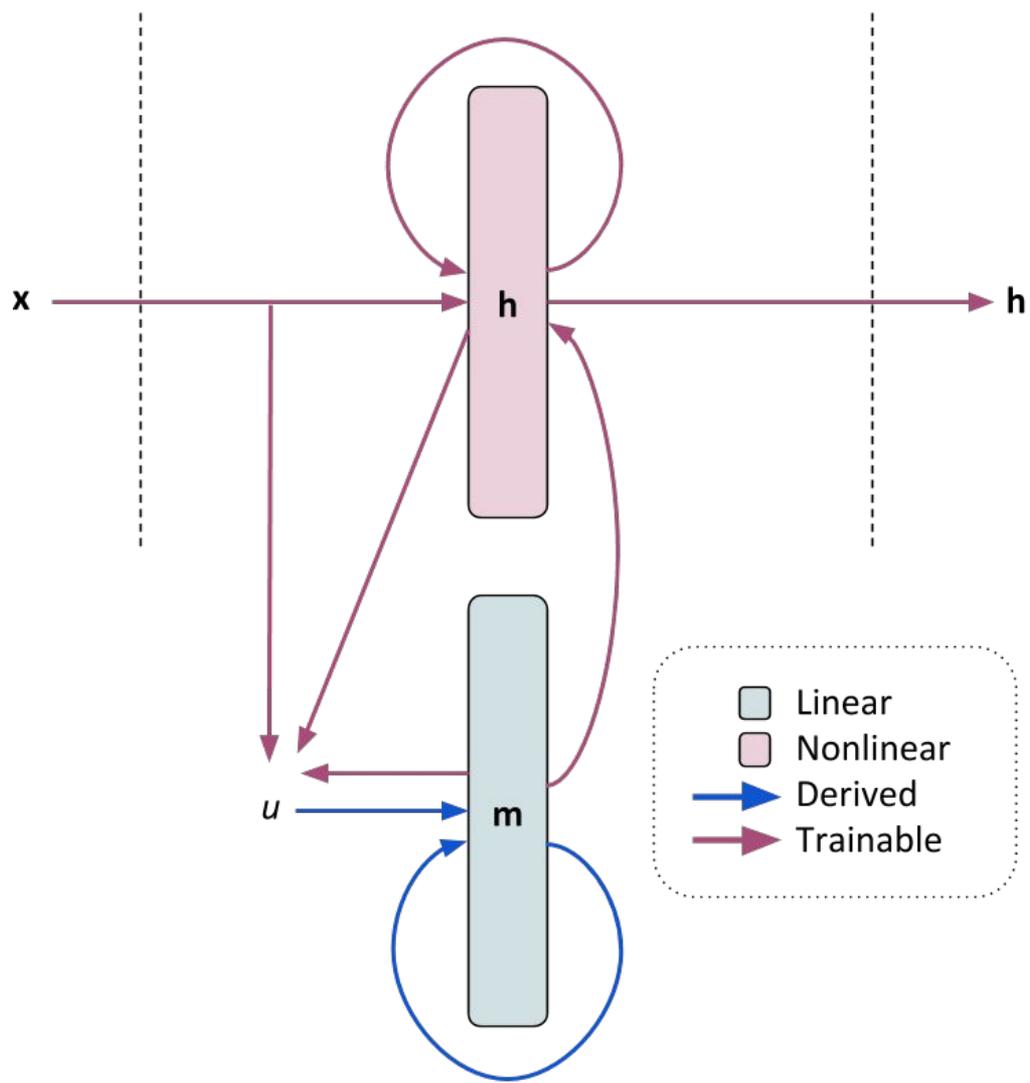
LSTMs

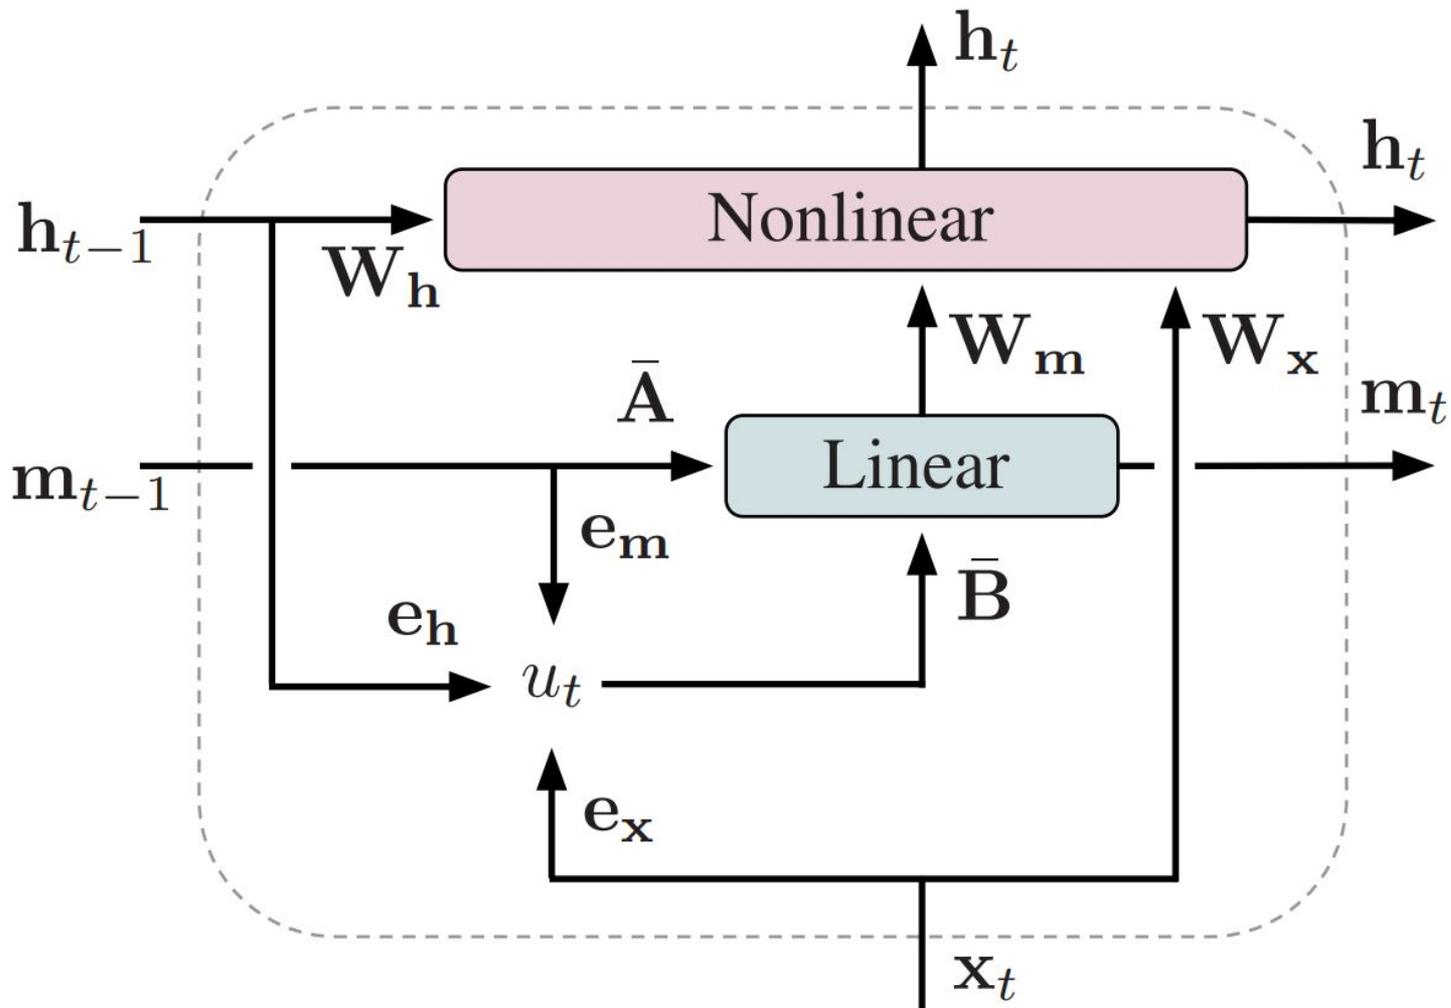
Speech
recognition

Disease prediction

Machine translation



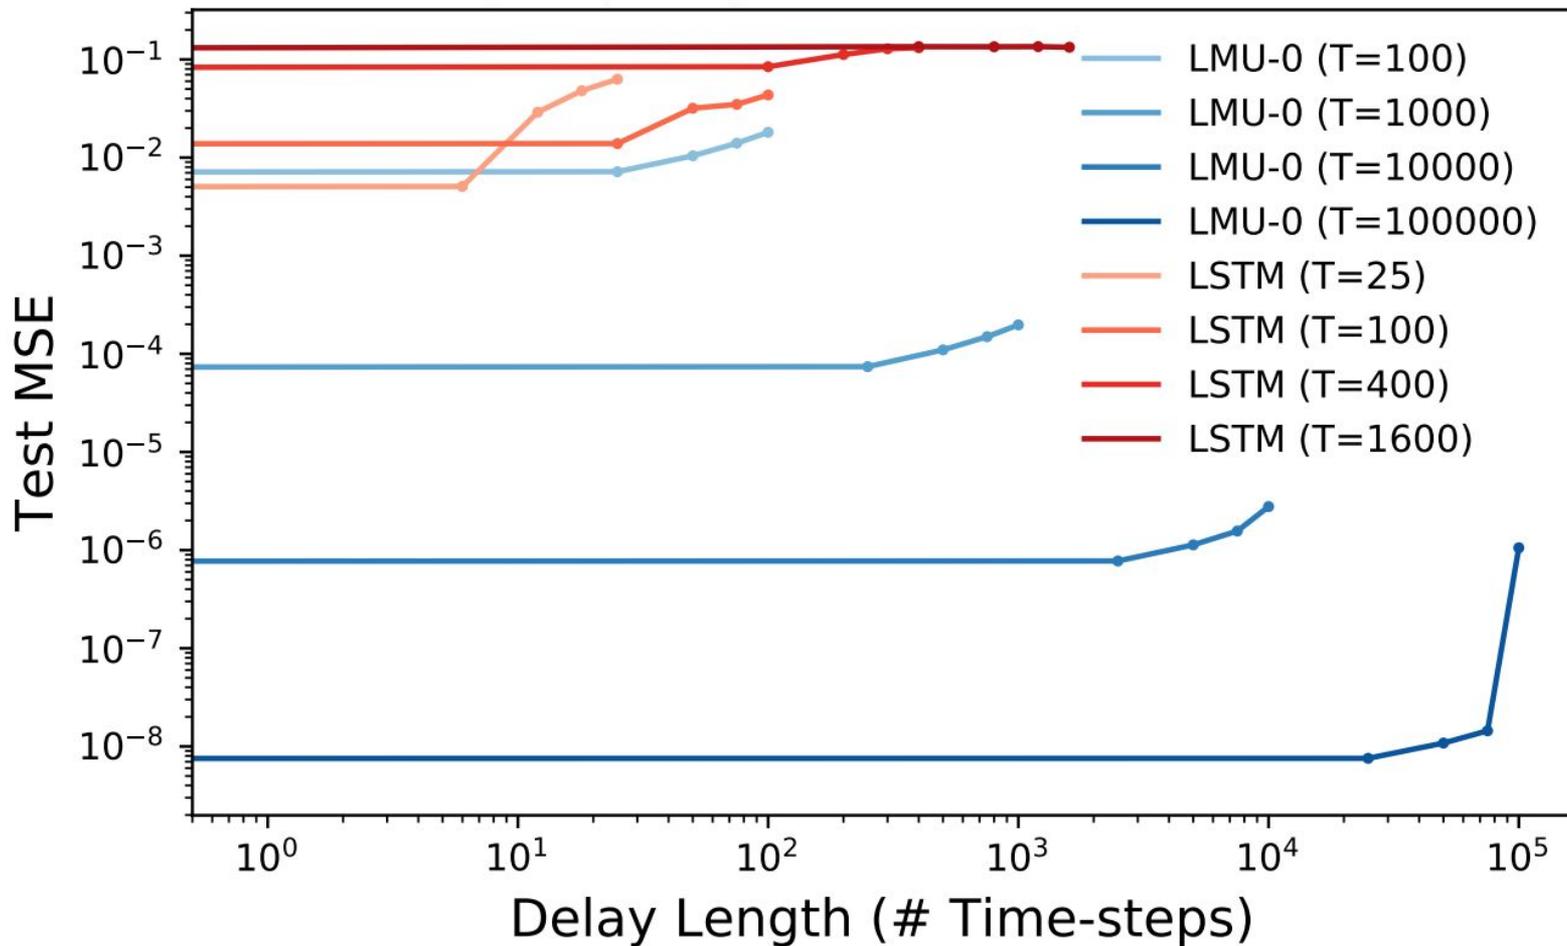




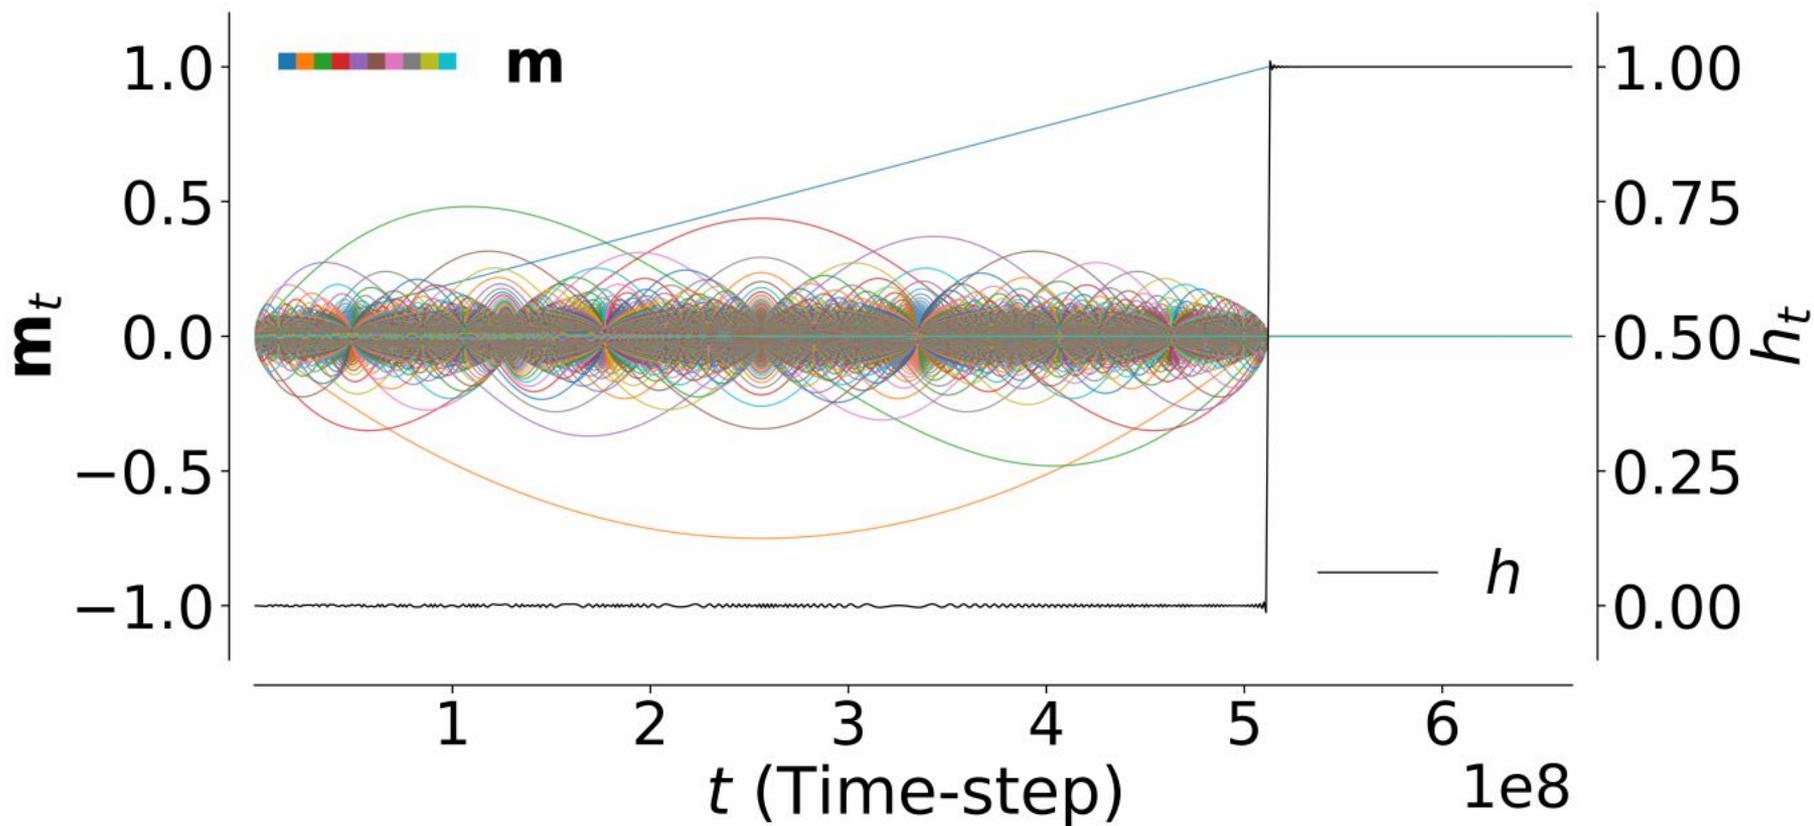
Model	Validation	Test
RNN-orth	88.70	89.26
RNN-id	85.98	86.13
LSTM	90.01	89.86
LSTM-chrono	88.10	88.43
GRU	92.16	92.39
JANET	92.50	91.94
SRU	92.79	92.49
GORU	86.90	87.00
NRU	95.46	95.38
Phased LSTM	88.76	89.61
LMU	96.97	97.15
FF-baseline	92.37	92.65

Extended from
Chandar et al. (2019)

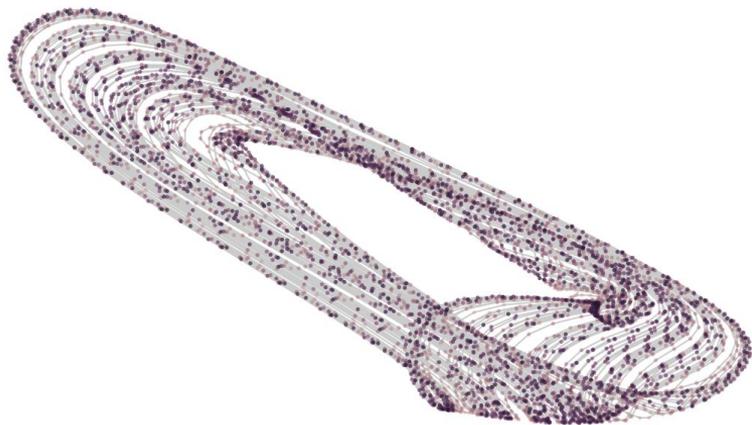
Memory Capacity of Recurrent Architectures



Scalable in Time

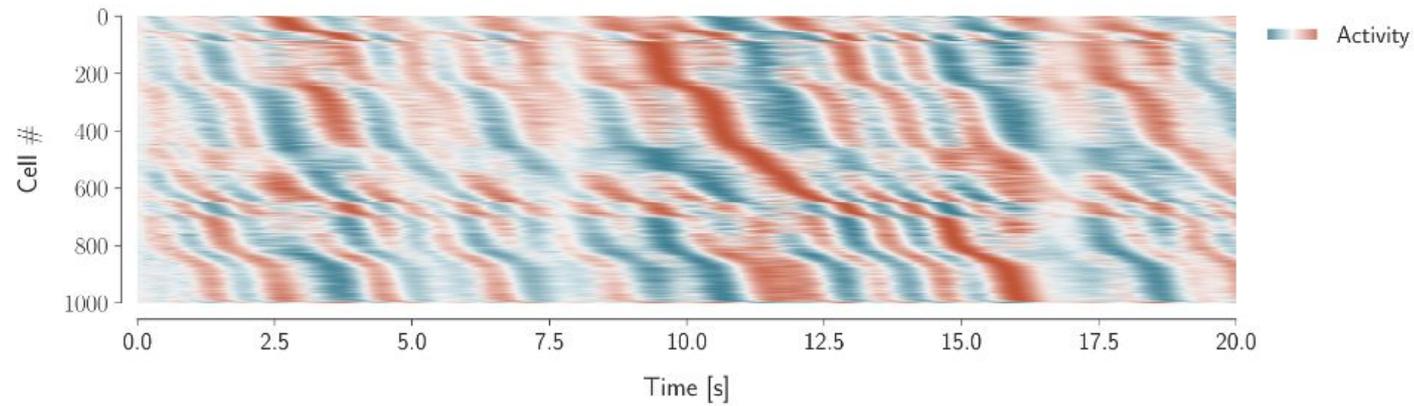
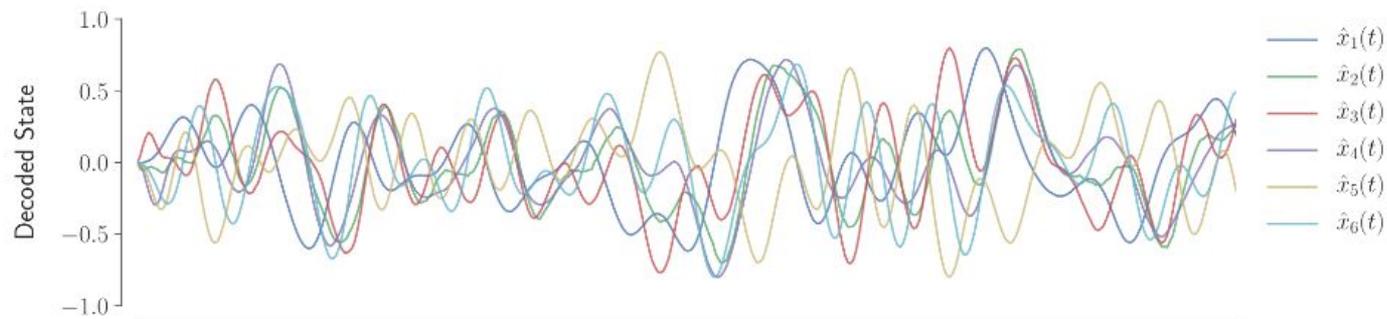
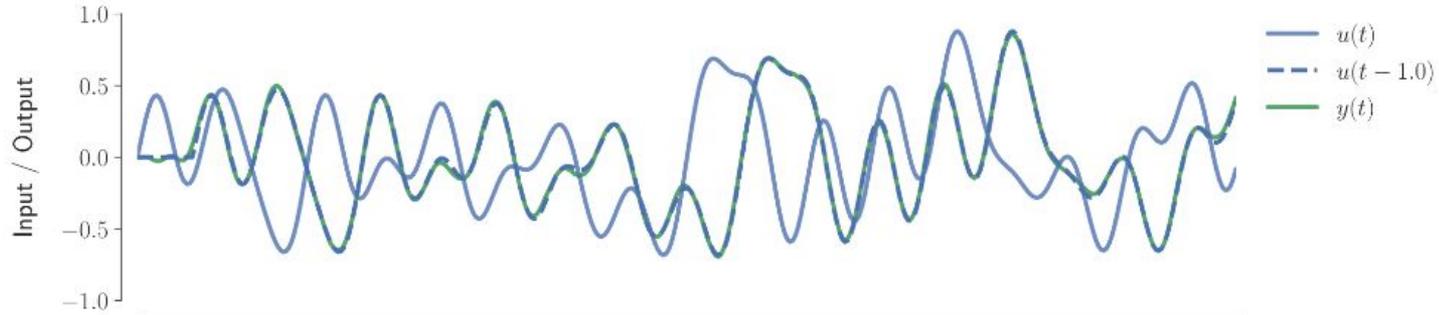


Chaotic Time-Series Prediction

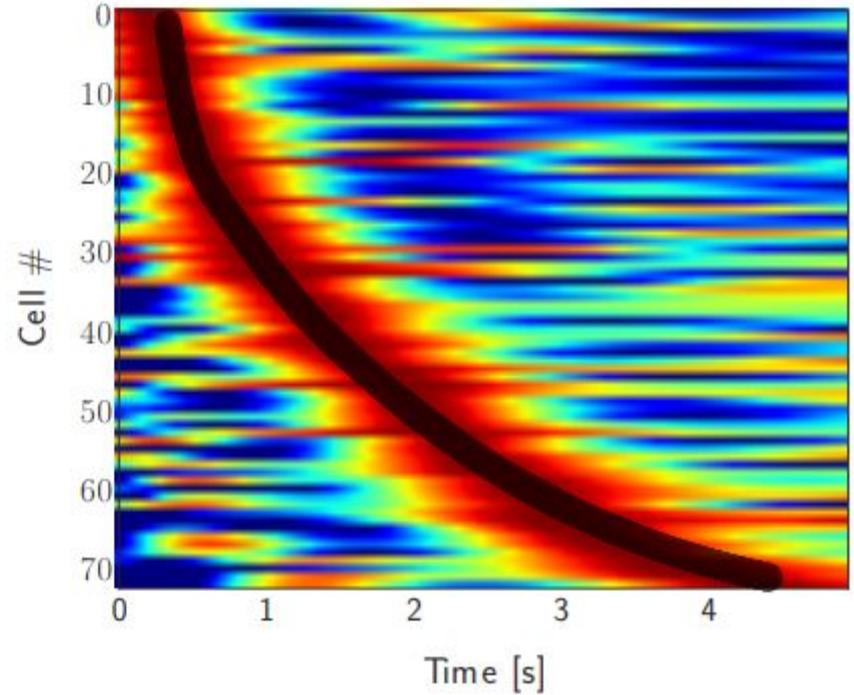
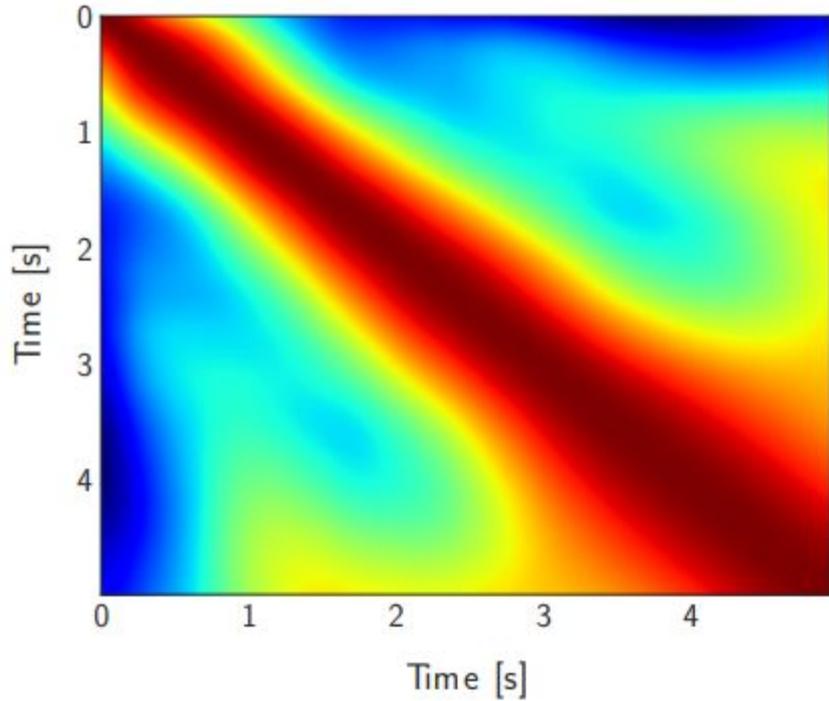


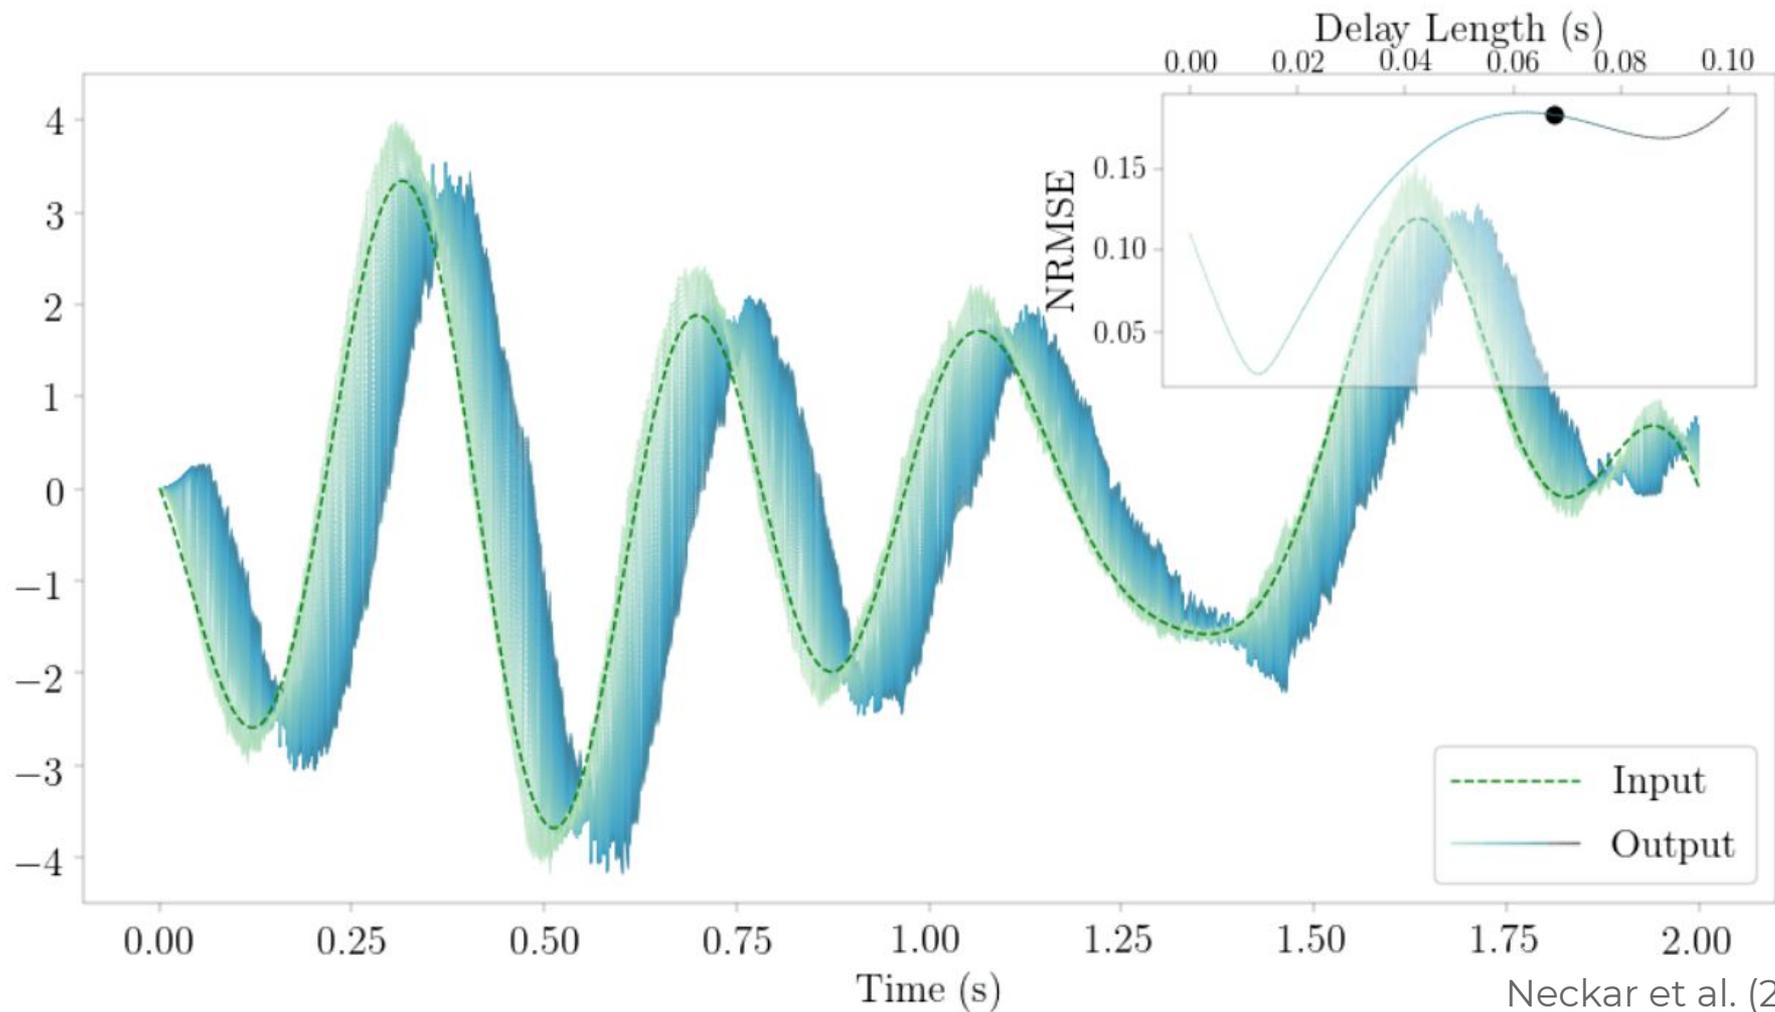
Mackey-Glass Dataset

Model	Test NRMSE	Training Time (s/epoch)
LSTM	0.079	20.34 s
LMU	0.054	12.89 s
Hybrid	0.050	16.21 s



Time Cell Data





Neural Implementation

