Edge Implementation of an LSTM Error Compensator for Accurate Vibration Measurements on UAV-Deployable SHM Sensor Packages

Joud N. Satme, Alexandra Giese, Daniel Coble, Hung-Tien Huang, Austin R.J. Downey

UNIVERSITY OF South Carolina

Molinaroli College of Engineering and Computing

Outline

- **Rapid Structural health monitoring**
    - Process definition
    - Challenges attributed to current methods
    - Proposed solutions
- **Nonlinear error compensation**
    - Long Short-Term Memory Networks
    - Previous challenges and our proposed approach
- **Hardware and edge computation**
    - Rapid vs Real-time
    - Types of processing hardware
    - System adopted in this work
- **Error compensator development**
    - Benchtop experiment
    - Training data
    - Model progression
- **Results and Discussion**
    - Model choice and training
    - Performance metrics
- **Conclusion and future work**

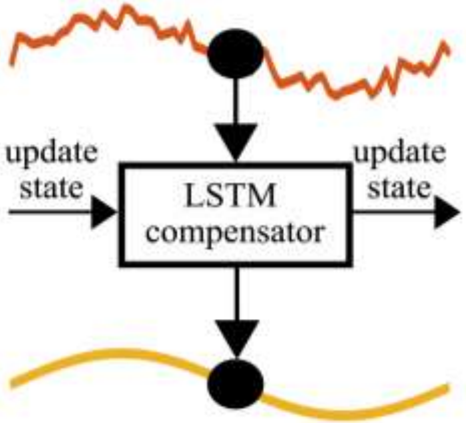

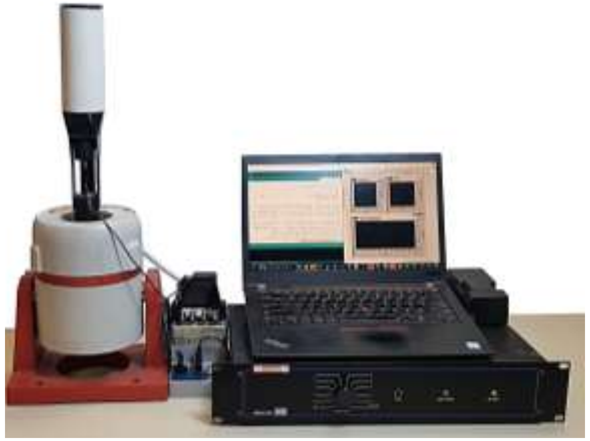

Rapid Structural health monitoring

- Is the process of rapid data acquisition, signal processing, and inferring the state of a structure given its response, enabling timely assessment and prognostics.
- Challenges attributed to current methods
  - Tedious process
  - Costly and dangerous
  - Requires highly skilled personnel
  - Specialized equipment and vehicles
- Proposed solutions
  - Minimally invasive sensor
  - Deployed via UAVs
  - For rapid data acquisition
  - Intelligent signal processing
  - Reducing the time from measurement to prognostics

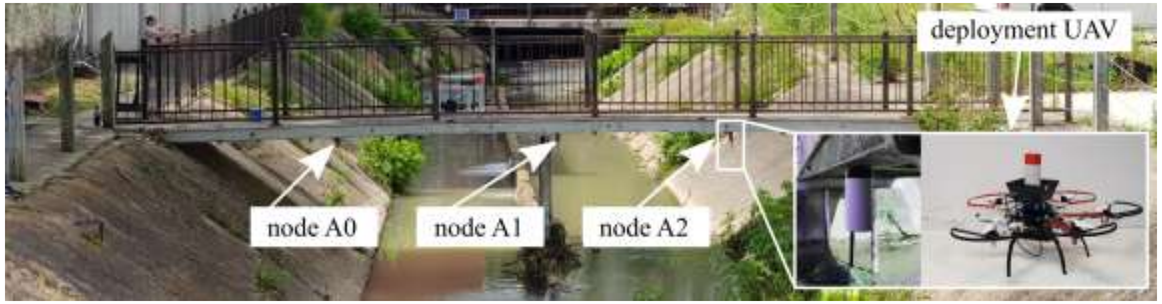

Rapid Structural health monitoring

- System breakdown
    - Electropermanent magnet (EPM)
    - MEMS accelerometer
    - Microcontroller
    - Lithium polymer battery
    - Memory storage
    - RF wireless communication
- Enables rapid deployment of vibration-sensing nodes onto structures via a hexacopter UAV for rapid assessment.

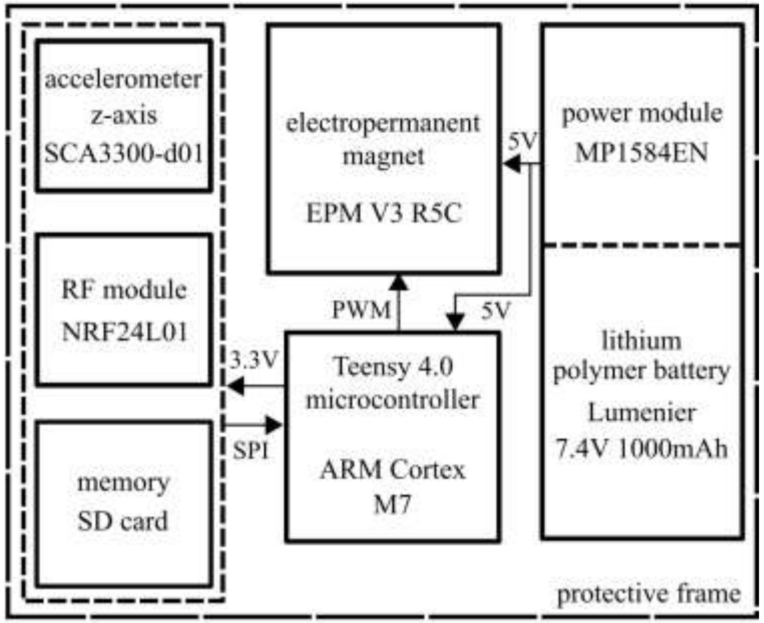

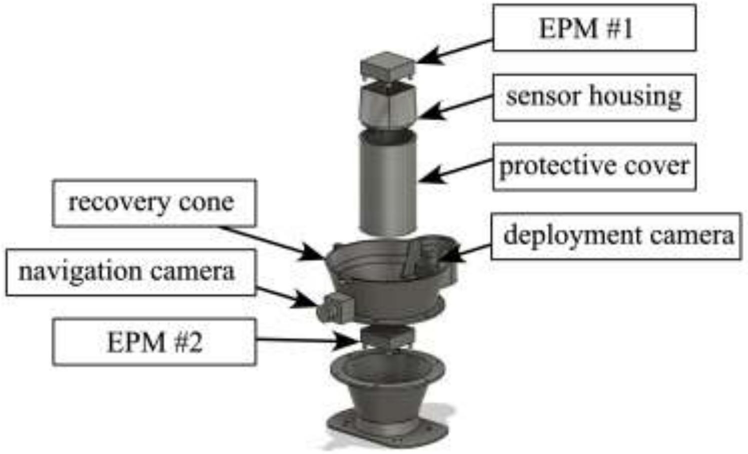

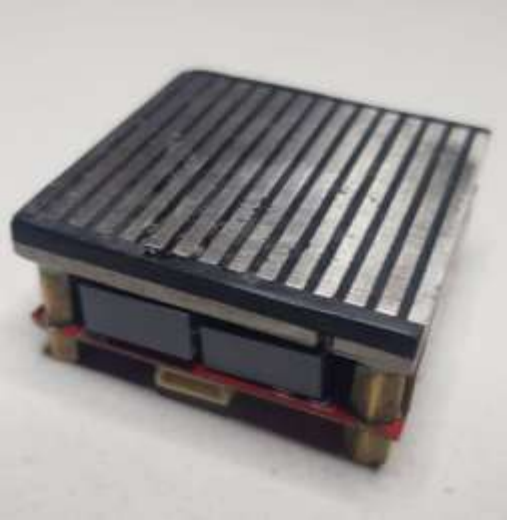

Rapid Structural health monitoring

- System breakdown
  - Electropermanent magnet (EPM)
  - MEMS accelerometer
  - Microcontroller
  - Lithium polymer battery
  - Memory storage
  - RF wireless communication
- Enables rapid deployment of vibration-sensing nodes onto structures via a hexacopter UAV for rapid assessment.

Nonlinear error compensation

- Sensors suffer from degradation of performance in the low-energy low-frequency range.
- This being the window of interest native to many civil structures (>20Hz).
- Neural Network-based error compensation immerges as a solution to combat signal degradation.
- With Long Short-Term Memory (LSTM) networks showing potential processing temporal sensor data.

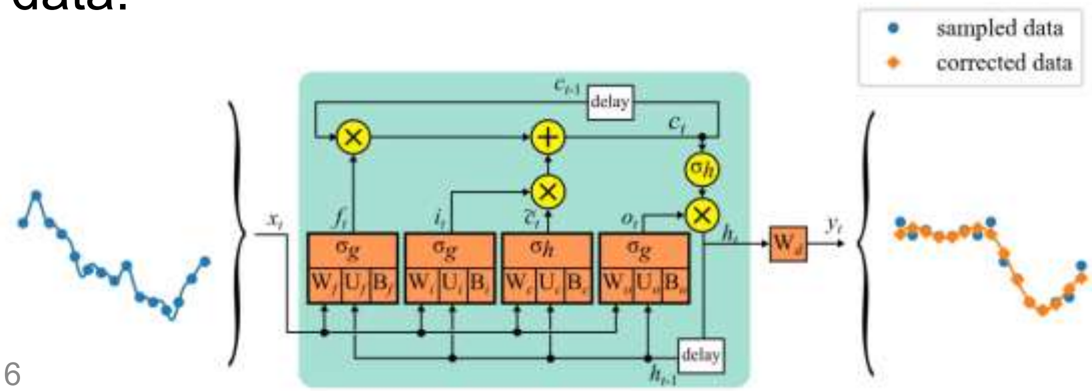

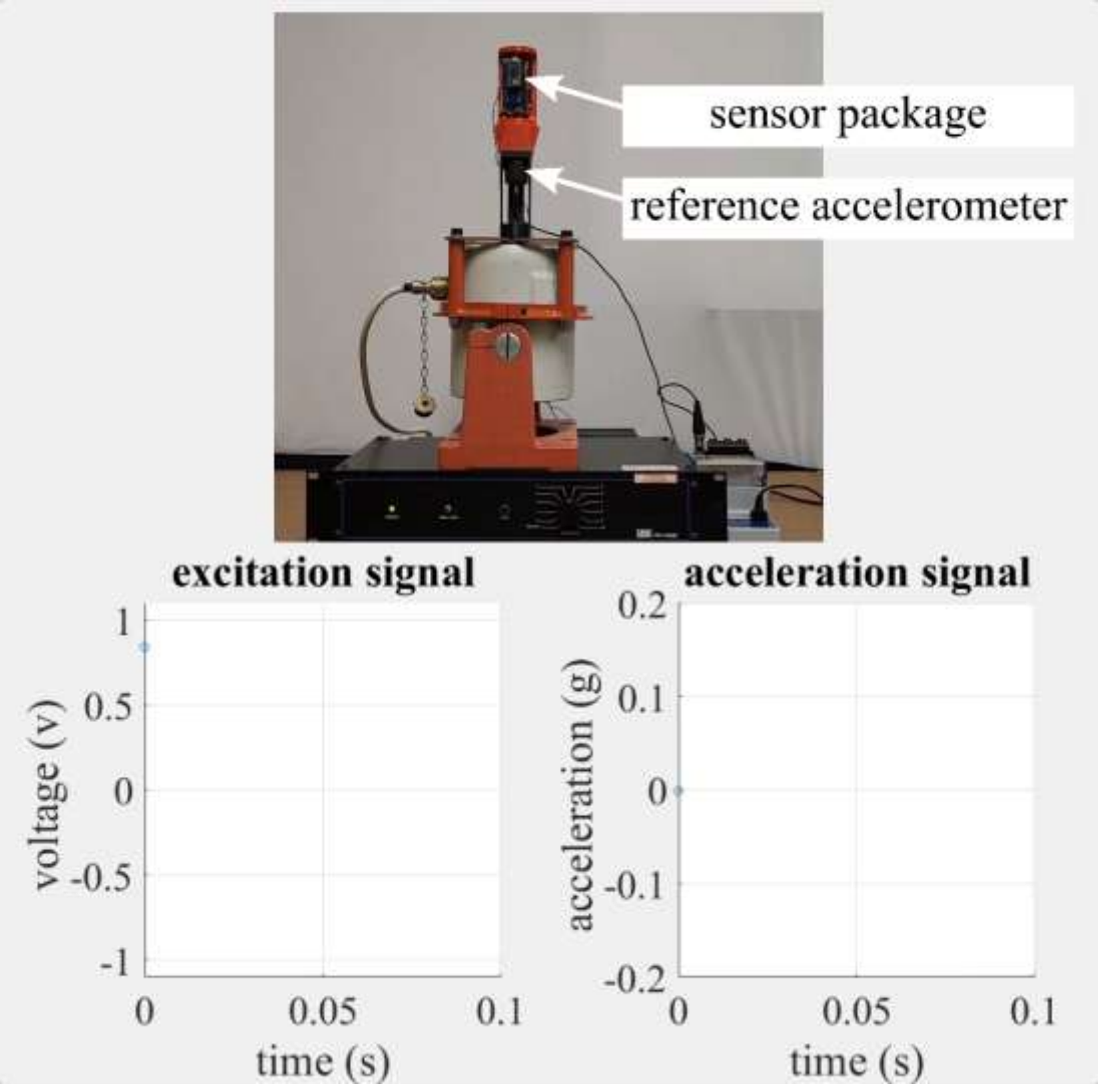

Hardware and edge computation

- Important
  - Real-tir
  - Determ
  - Proces ing it real-tim .
  - This dis ng where c

Anchorman: The Legend of Ron Burgundy

Hardware and edge computation

- Processing hardware come in all shapes and sizes!
  - GPUs (Graphics Processing Units)
    - Optimized for parallel computing
    - Great for machine learning and graphics
    - Power-intensive
  - FPGAs (Field-Programmable Gate Arrays)
    - Application-specific custom hardware acceleration
    - More energy-efficient than GPUs and CPUs
    - Harder to program
  - Micro-processor
    - Power and cost efficient
    - Does not require an operating system
    - Lacks the memory and peripherals
  - Micro-controller
    - Integrate a CPU, memory, and peripherals
    - Low-power and cost-effective
    - Lacks a dedicated operating system
    - Useful for real-time applications
    - Limited in memory and computational power

GPU

FPGA

Micro-processor

Micro-controller

Error compensator Development

- Goal: develop and train an edge-implemented LSTM error compensator model that takes in low-fidelity vibration data and produce a higher fidelity signal directly on device.
- Parameters:
  - Enhance the signal-to-noise ratio of low fidelity drone-deployable accelerometers
  - With a small enough model to fit on the Teensy 4.0 microcontroller (ARM Cortex M7)

sensor deployment

accelerometer measurement

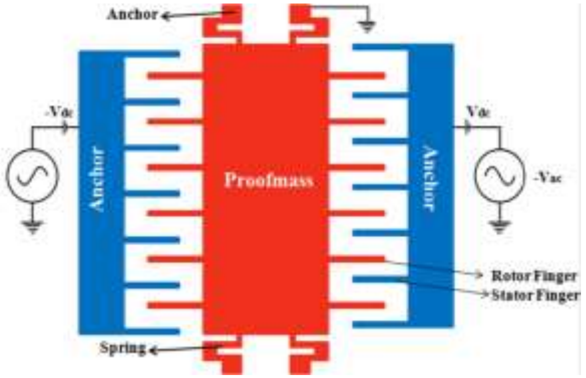

error compensation

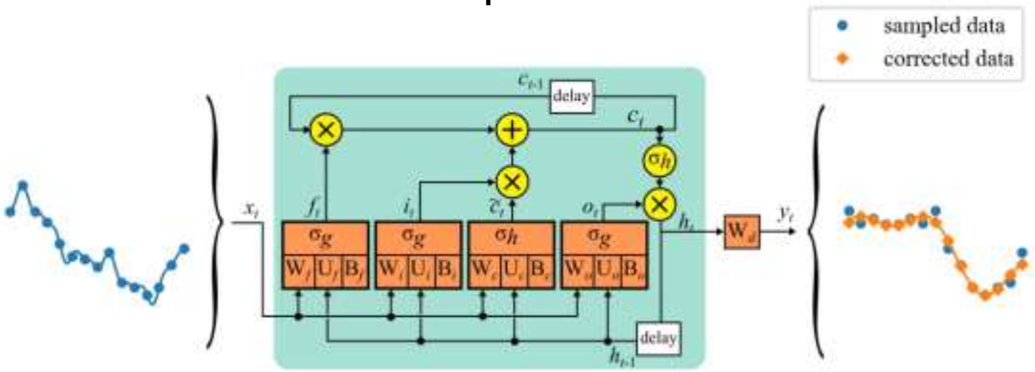

Error compensator Development

- Data is collected while using an electromagnetic shaker
- Driven by a known control signal in the range of interest
- A digital trigger is used to synchronize measurements and excitation
- A high pass filter is imposed on the high fidelity sensor signal to simulate a low fidelity sensor with degraded performance in the range >10 Hz

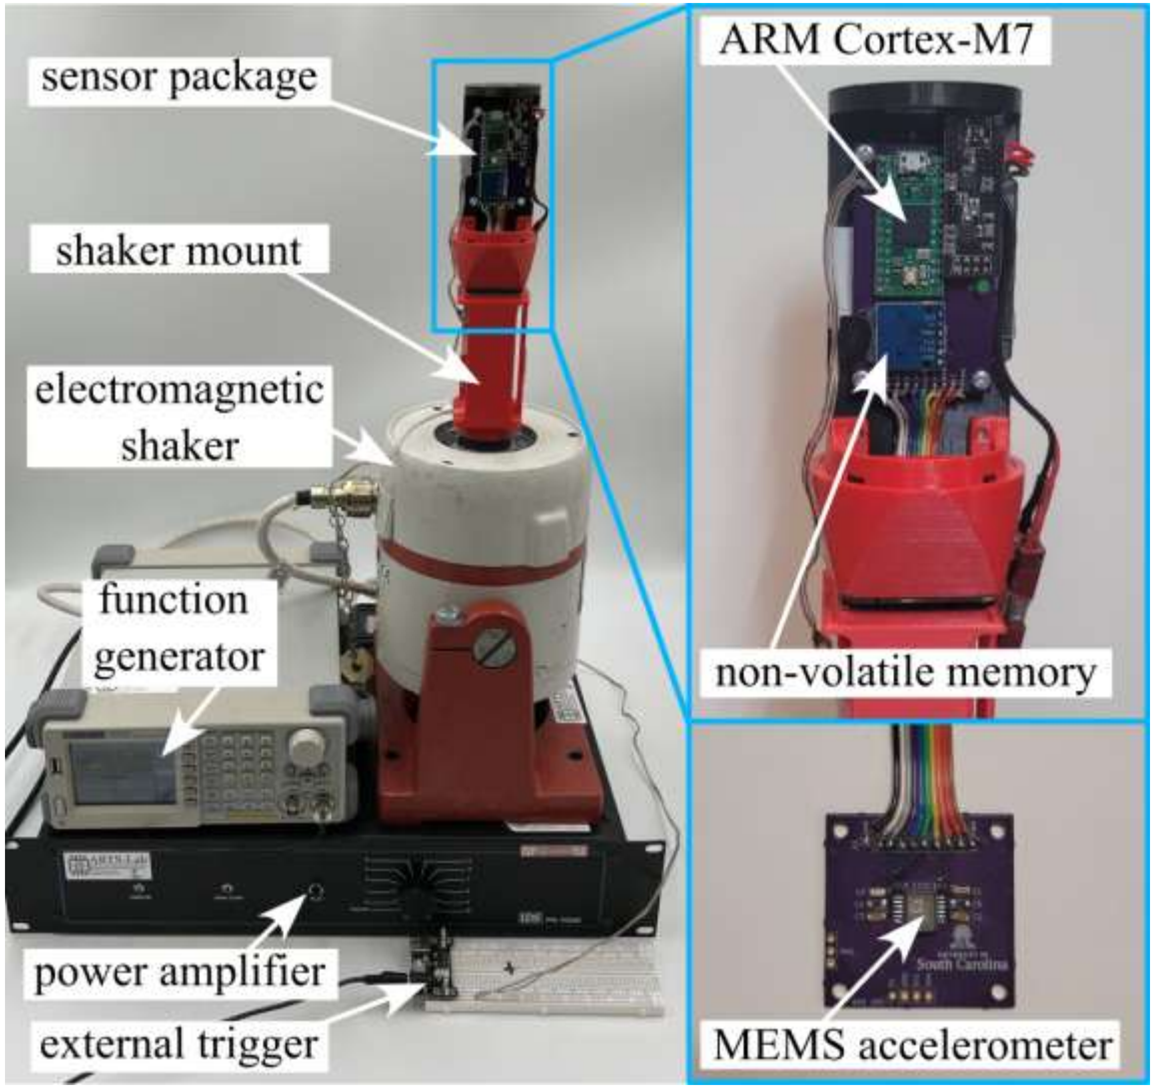

Error compensator Development

- Chirp excitation and single frequency
- Bandwidth of interest 0.1 - 20.9 HZ
- Synthetically degraded signal
  - HPF of the 5th order
  - Cutt-off of 10 Hz
  - Applied to the reference signal

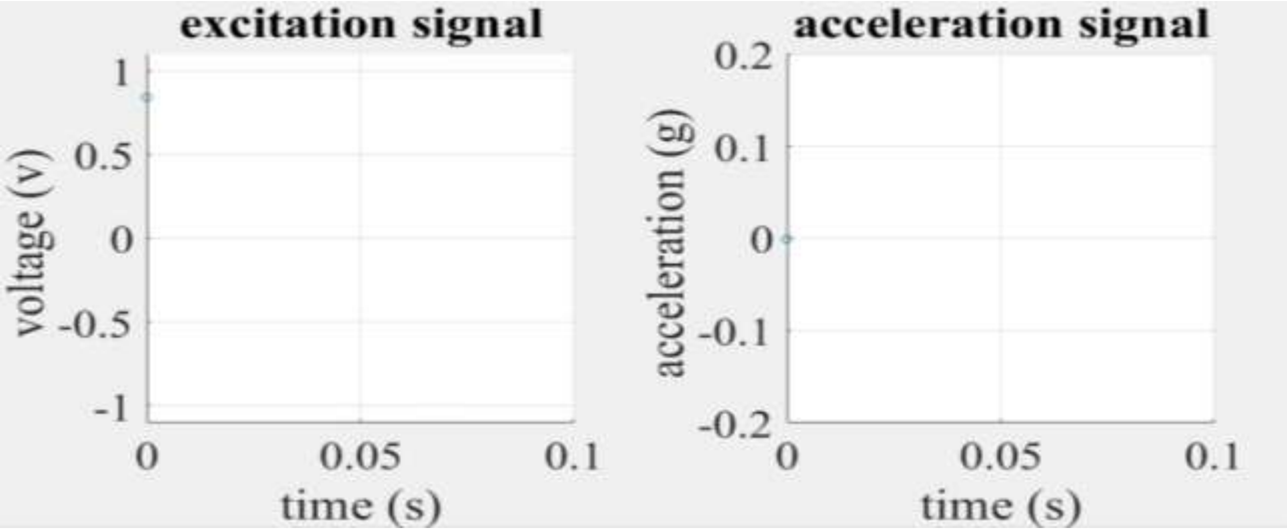

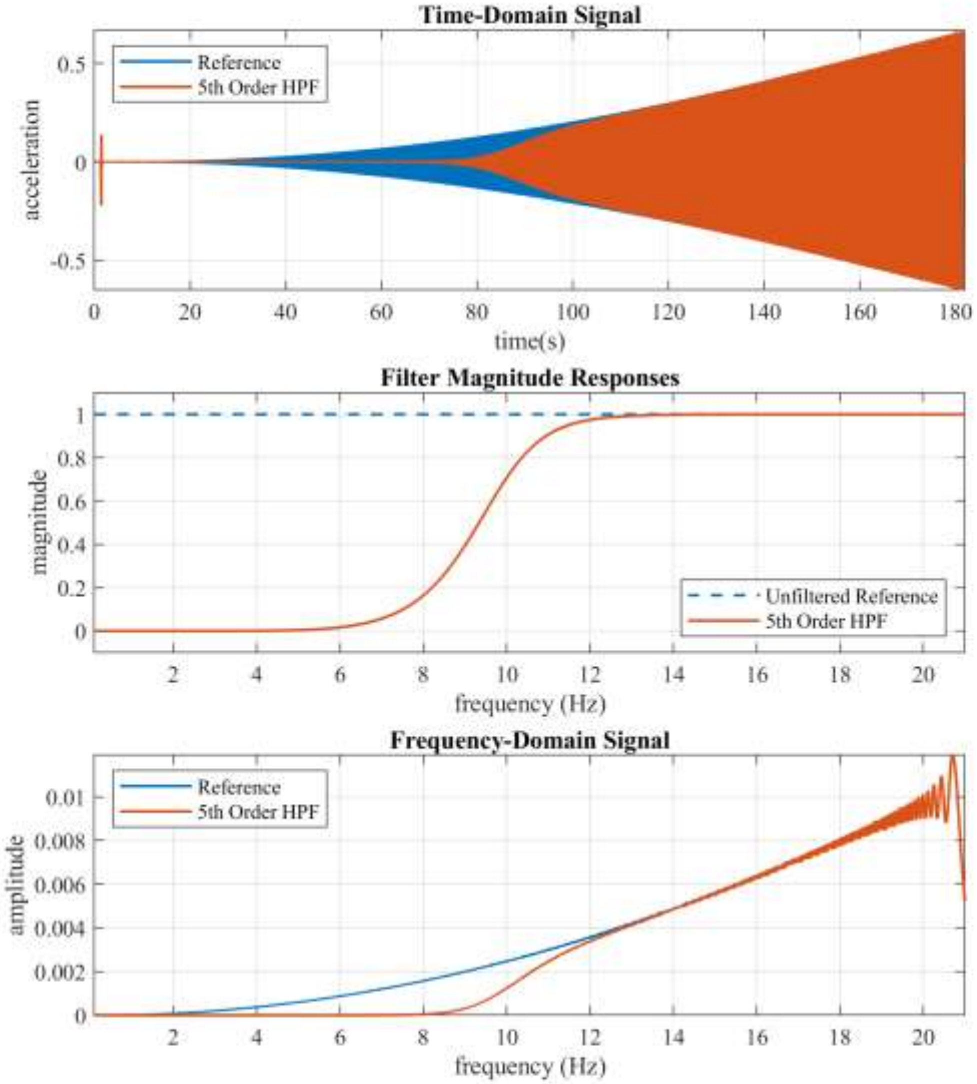

Error compensator Development

- Chirp excitation and single frequency
- Bandwidth of interest 0.1 - 20.9 HZ
- Synthetically degraded signal
  - HPF of the 5th order
  - Cutt-off of 10 Hz
  - Applied to the reference signal

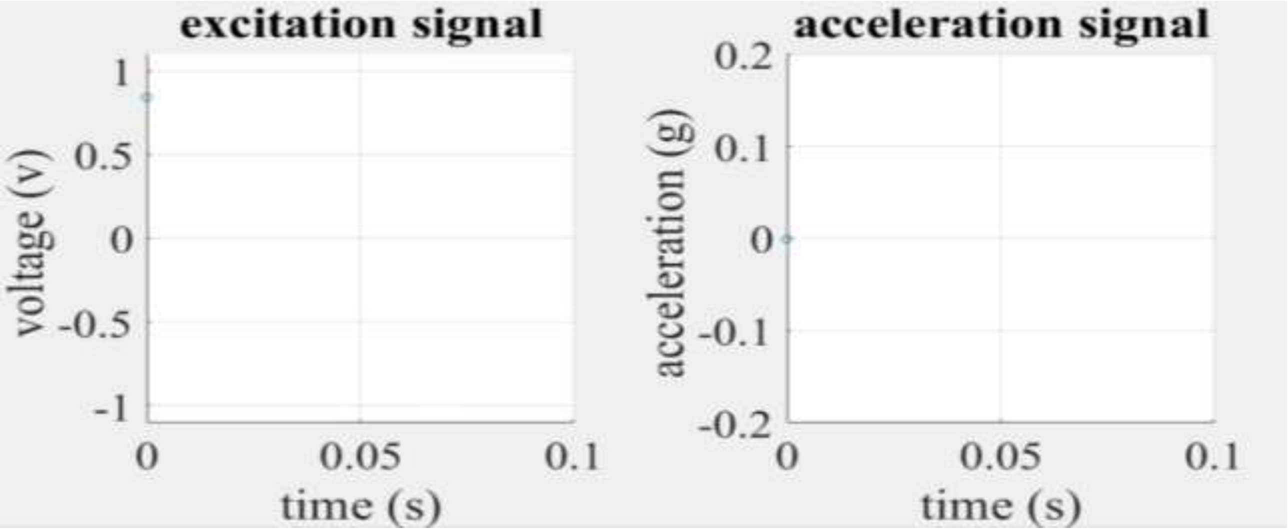

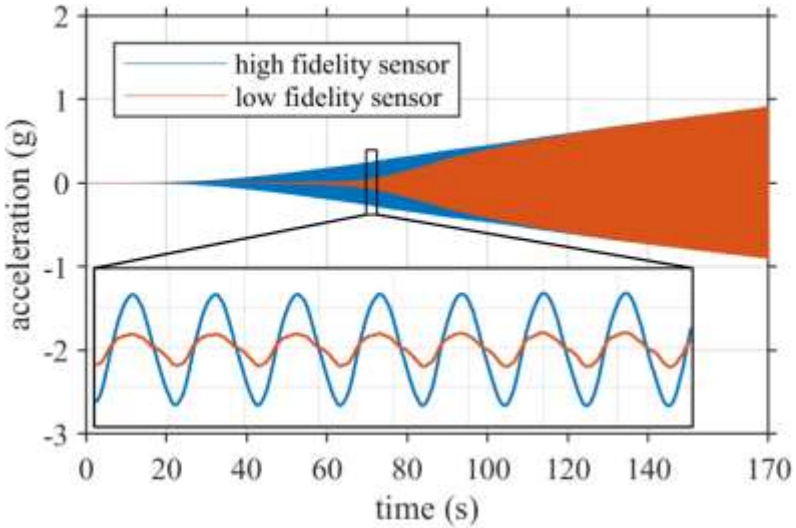

(a)

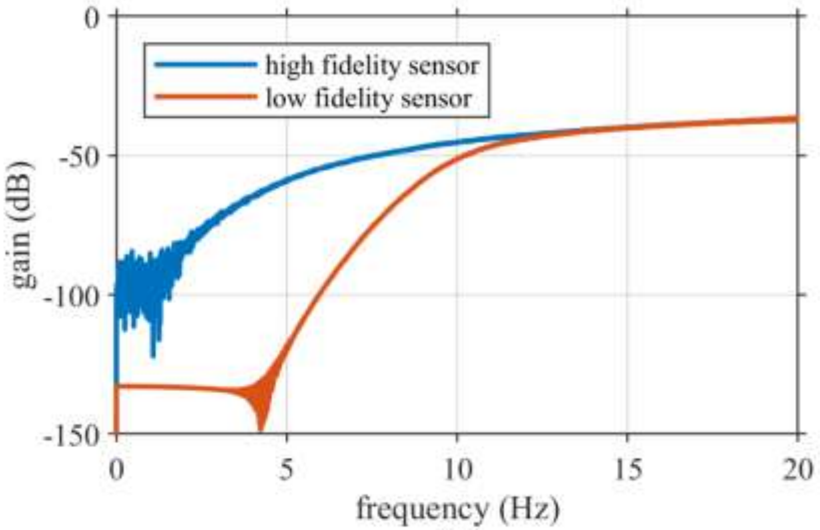

(b)

Results and Discussion

- Model size was progressively investigated with a single layer of 1, 2, 4, 6, 8,10,15, 25, 50,100,150,170 units.
- Training data consisted of frequency sweeps and single frequency excitation in the range of 0.1 – 20.9 Hz.
- Supervised training framework
- Model performance is evaluated through
  - signal-to-noise ratio (SNR_{dB})
  - root mean squared error (RMSE)
  - time response assurance criterion (TRAC)
  - memory footprint
  - forward-pass latency
  - energy / inference

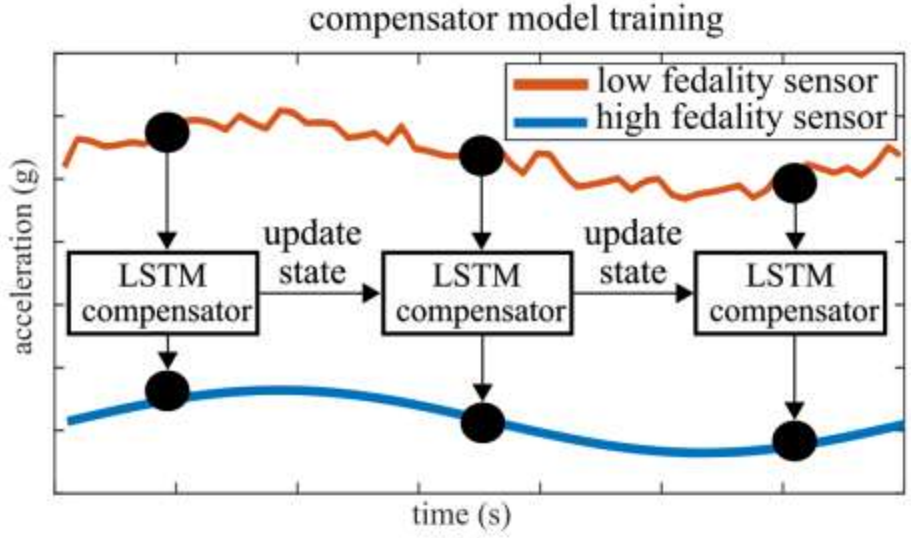

(a)

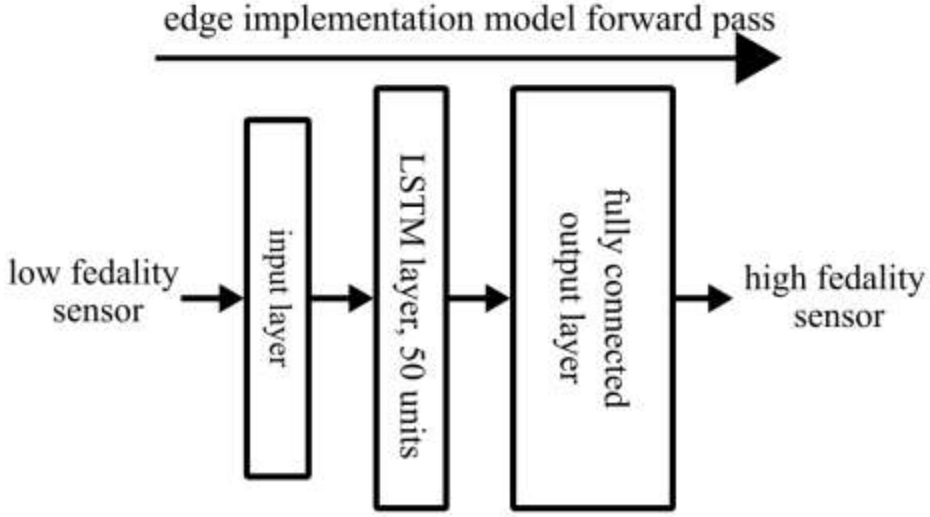

(b)

Results and Discussion

- SNR_{dB} is observed to have a logarithmic relation to the number of units.
- After 50-units the SNR_{dB} has negligible enhancement compared to the added computational complexity

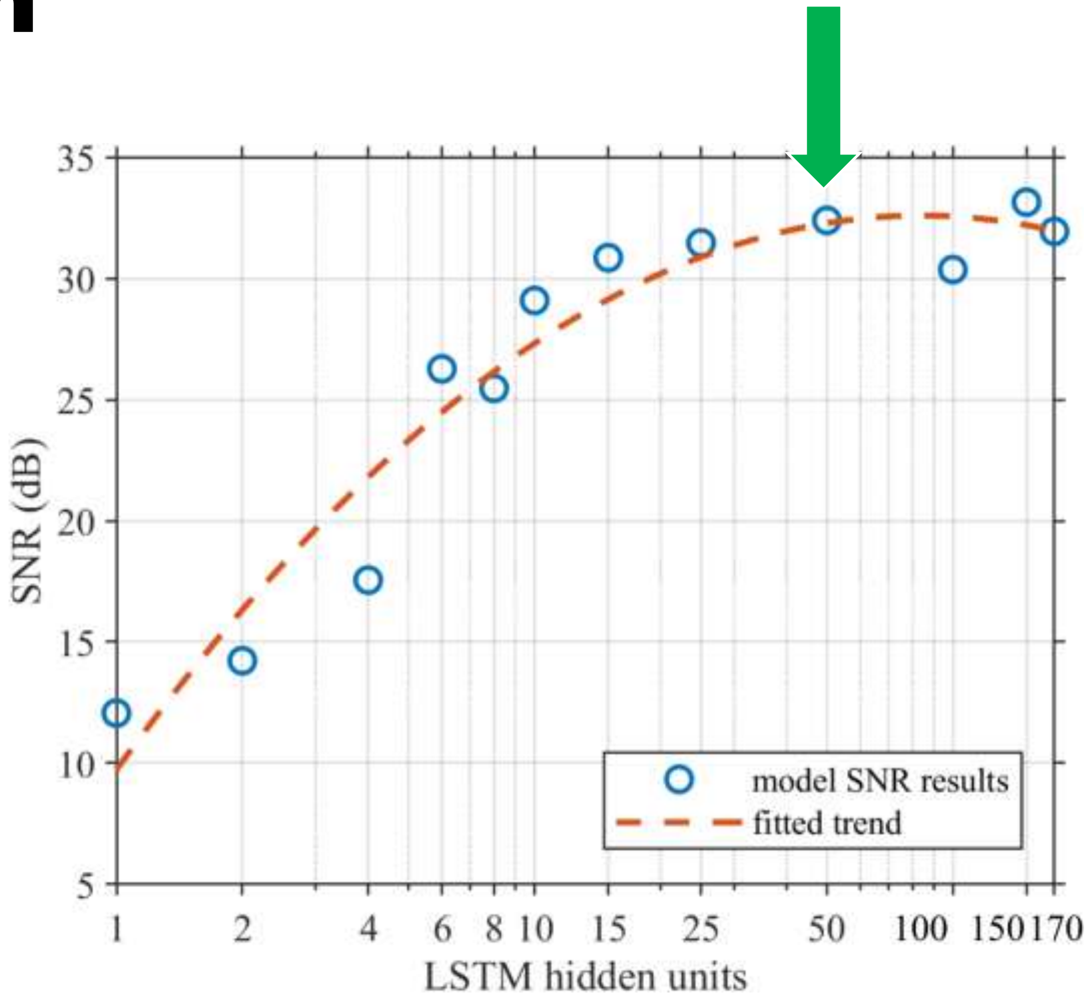

Results and Discussion

- Current draw is measured for 6 separate stages:
  1. Modules' initialization and startup
  2. LSTM Model parameter loading
  3. Input data loading into RAM
  4. LSTM gets constructed and performs warm-up forward passes before timing.
  5. Forward-pass benchmarking performs 1000 repeated timed LSTM forward passes.
  6. Results writing writes the completed benchmark results into nonvolatile memory
- Energy per inference is computed using stage 5 and 1000 iterations

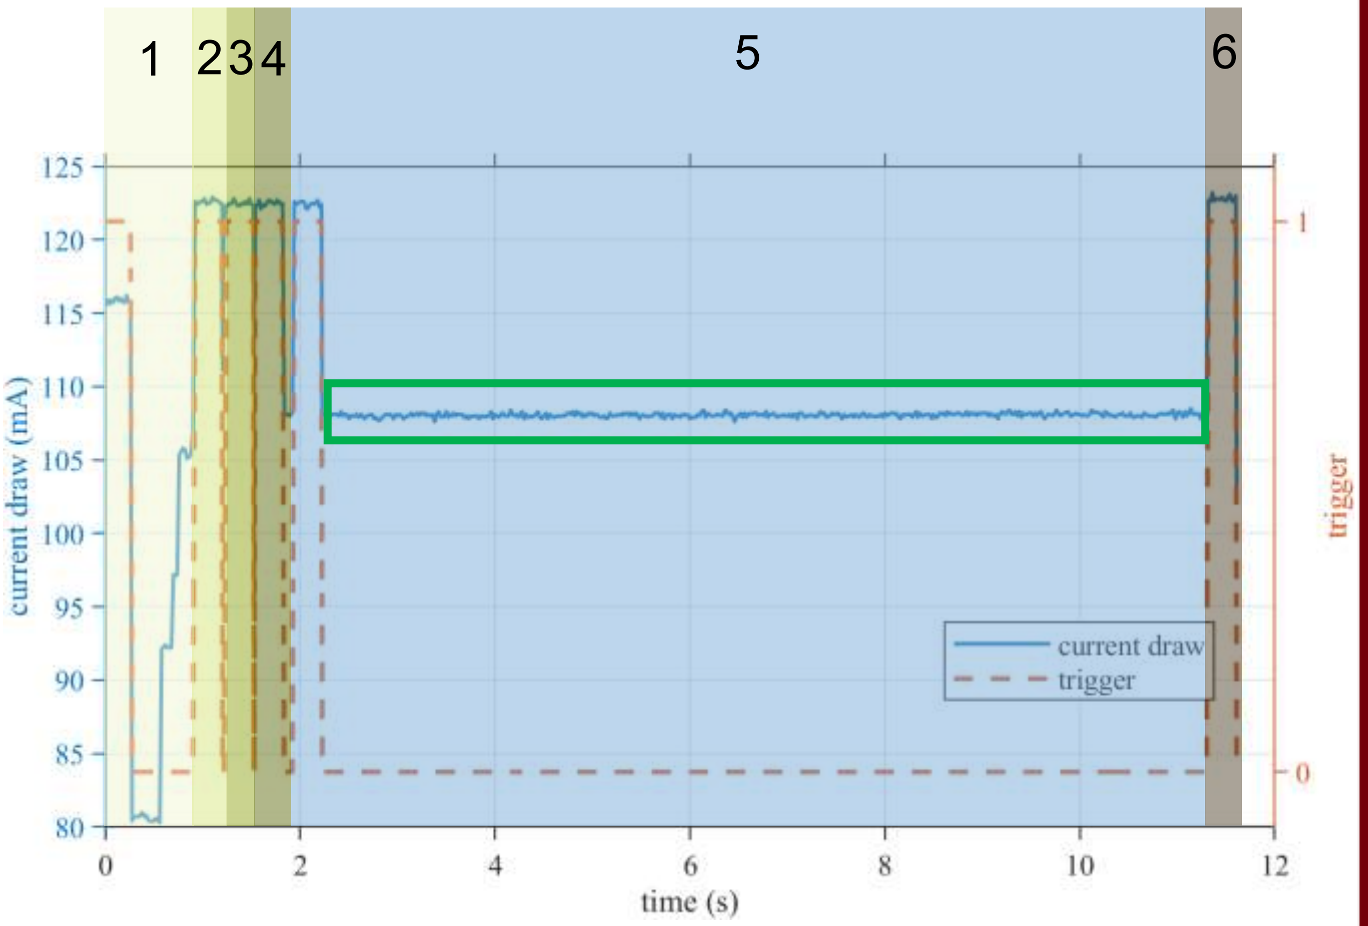

$$E_{\text{Stage5}} = \int_{t_5}^{t_6} 7.4\, I(t)\, dt \qquad E_{\text{inference}} = \frac{E_{\text{Stage5}}}{1000}$$

Results and Discussion

- All-encompassing spider plot showing
  1. TRAC
  2. Energy per inference
  3. Memory footprint
  4. Forward-pass latency
  5. RMSE

  for all 12 models examined
- All results are extracted from the hardware implementation on the Teensy 4.0 microcontroller
- 50-unit model is chosen for its well-rounded performance in all examined metrics.

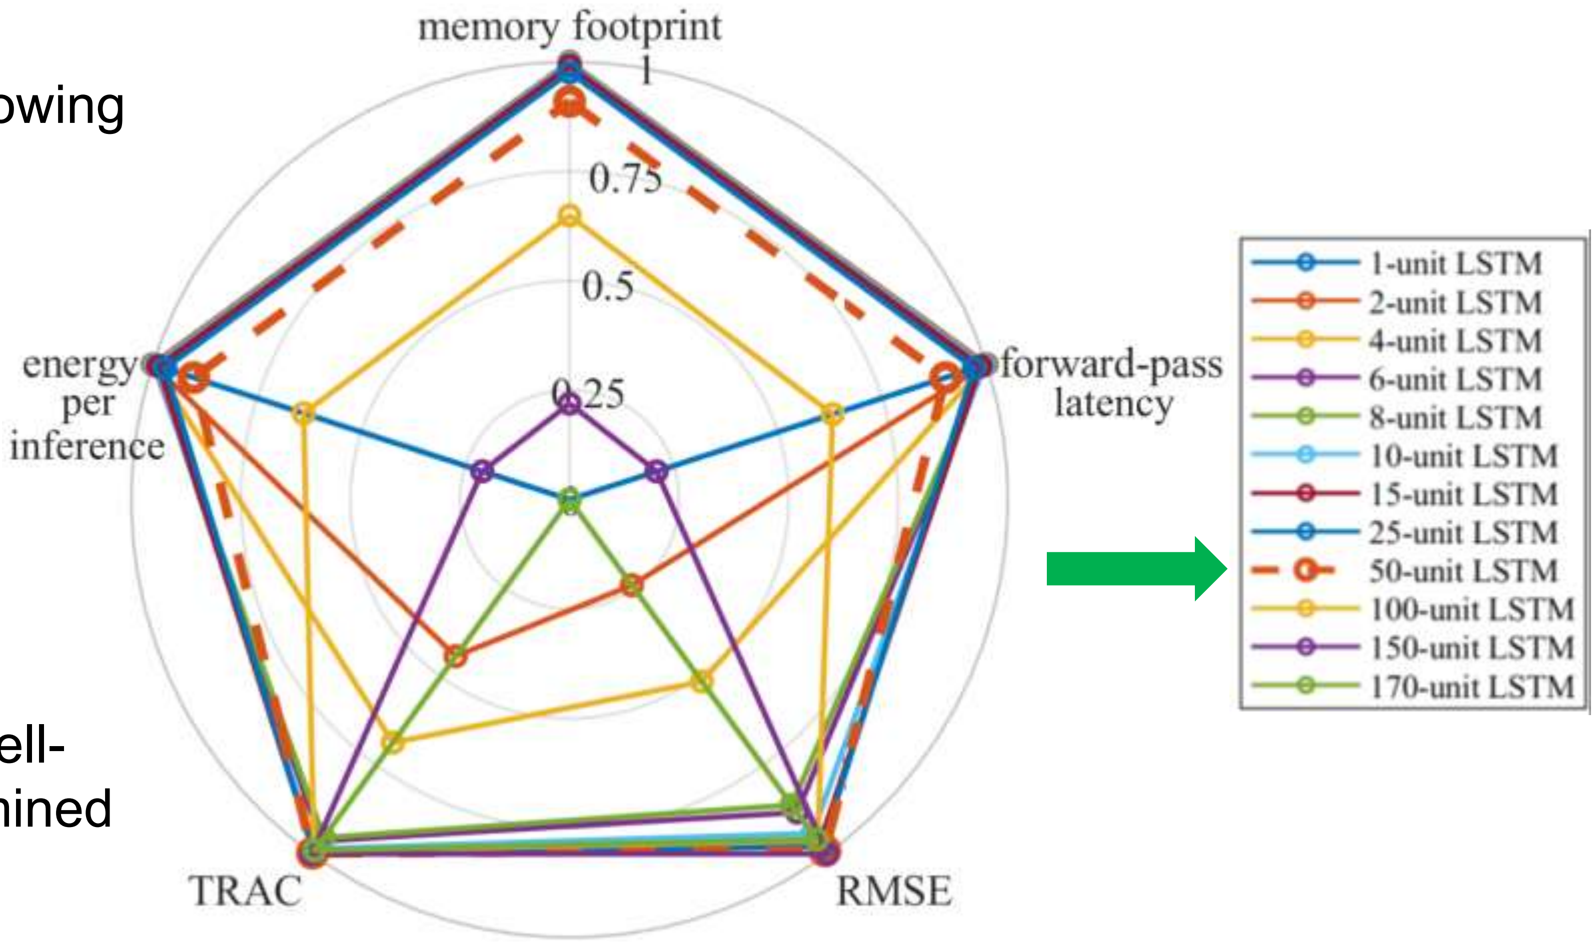

Results and Discussion

The Chosen 50-unit model specs:

- Latency per sample: 266 μs
- Runtime memory footprint: 44604 bytes
- Energy per inference: 89.45 mJ
- RMSE: 0.0081723 g
- TRAC: 0.99943
- SNR_{dB}: 32.39 dB

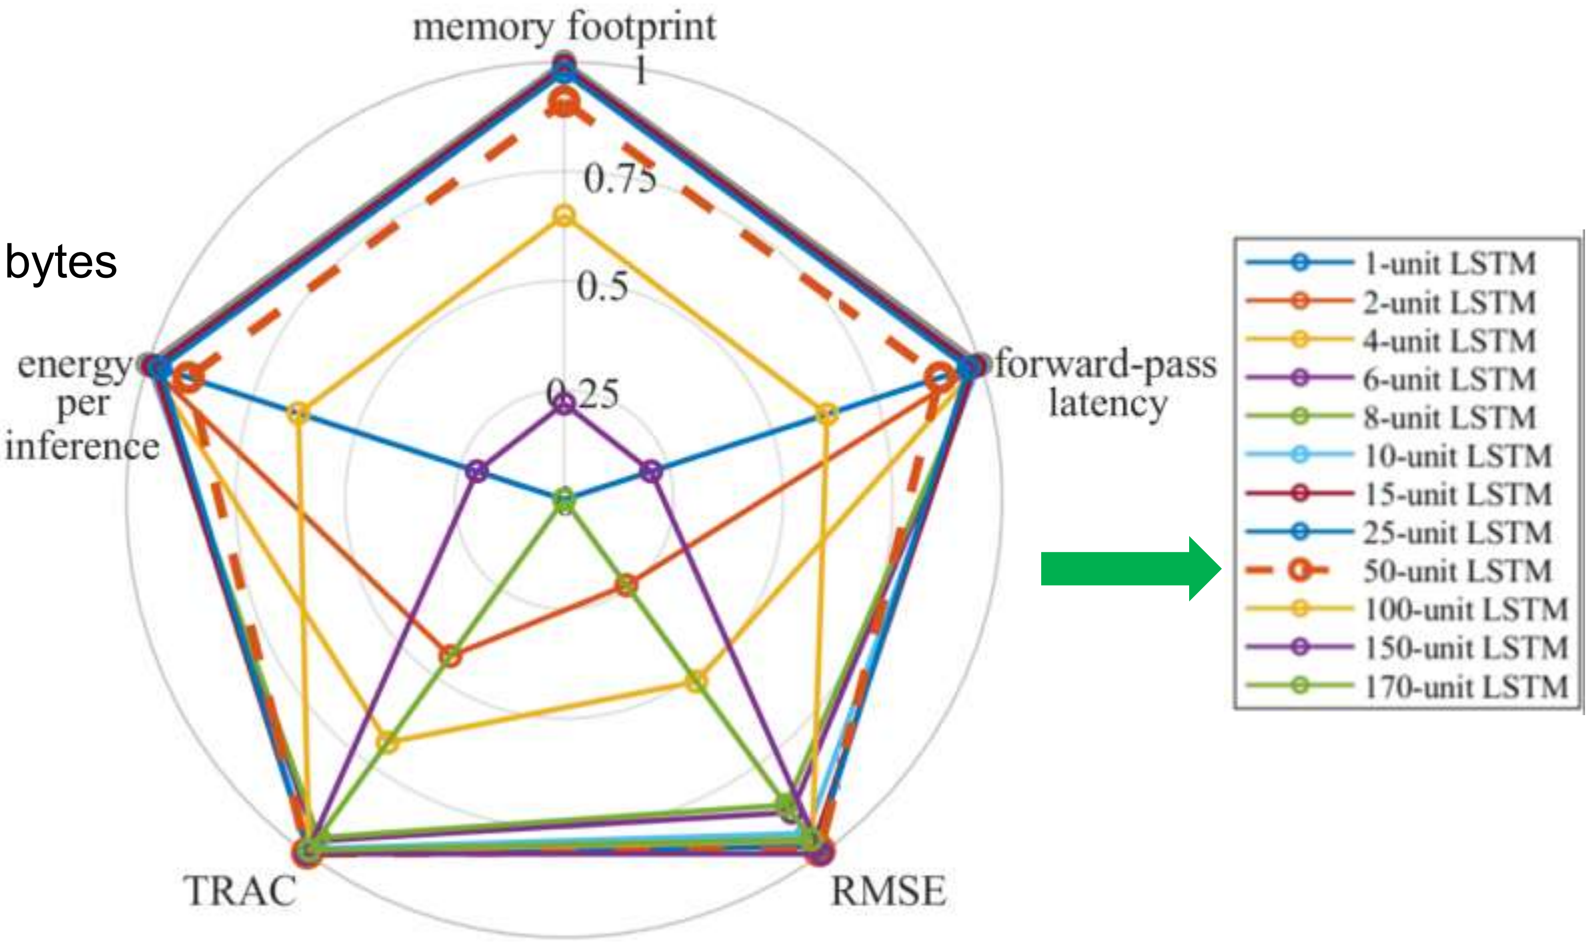

Results and Discussion

- 50-unit model is deployed
- compensation is done on-device
- With the time and frequency domain plots shown

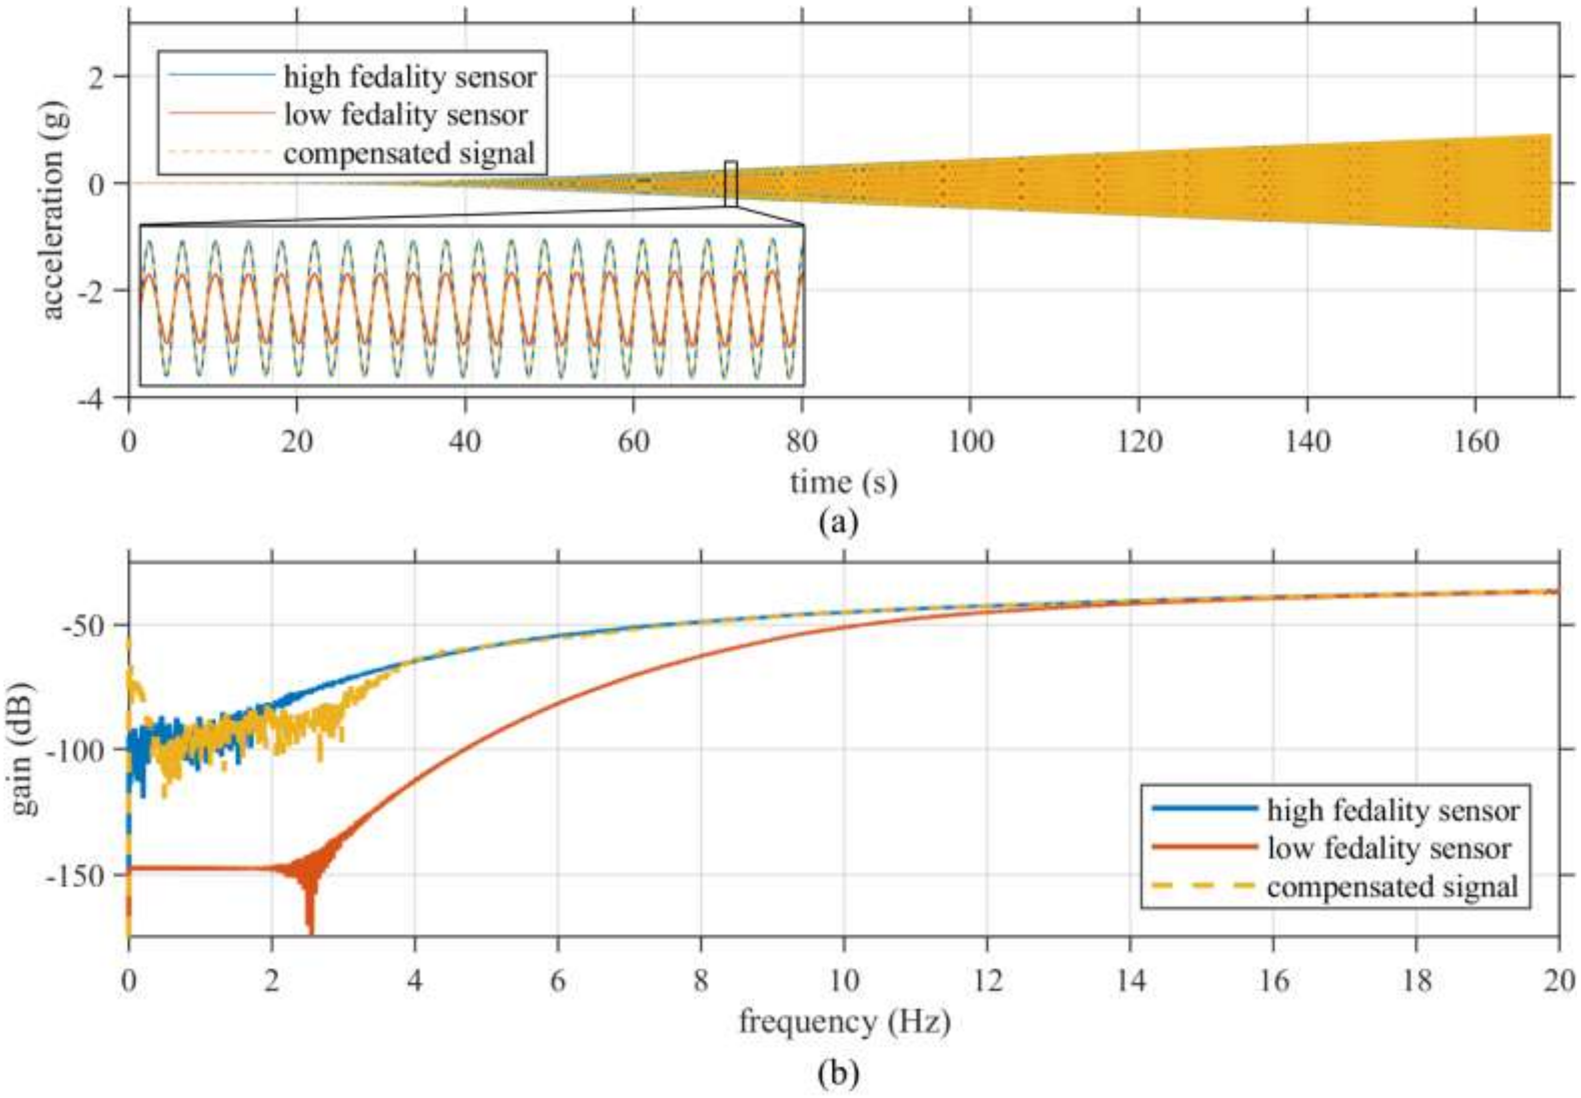

Results and Discussion

- Significant improvement is shown in the lower frequency signal fidelity >10 Hz
- Shown in
  - Frequency Response Function (FRF)
  - Power Spectral Density (PSD)

FRF

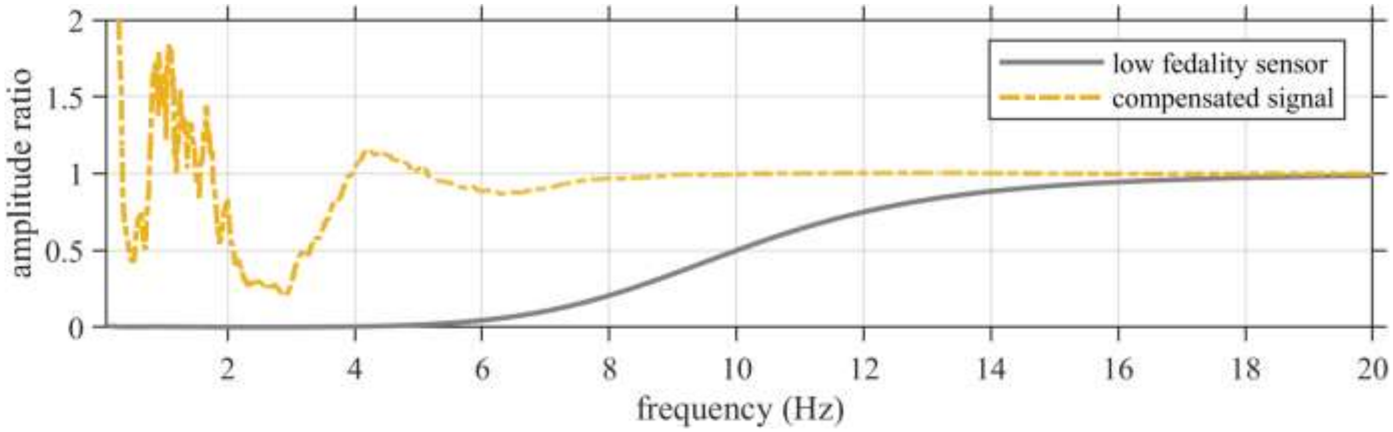

PSD

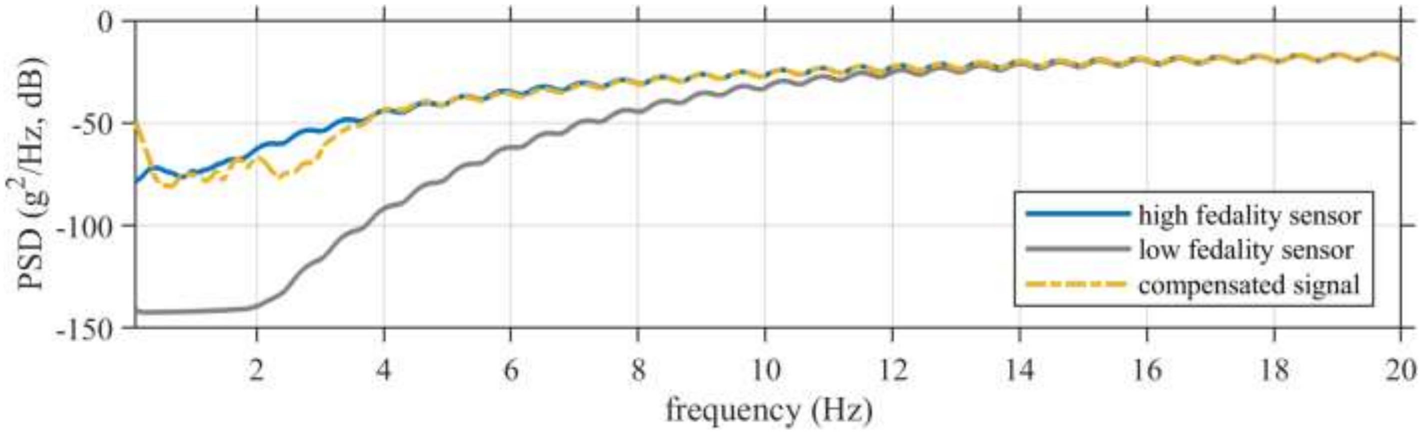

Conclusion and future work

- LSTM error-compensating networks have shown great potential at enhancing the fidelity of minimally invasive sensing systems that suffer from signal degradation in the low frequency range (> 20 Hz)
- This work shows a successful implementation of the compensating network on an ARM Cortex M7 microprocessor onboard a UAV-deployable sensor package

Conclusion and future work

- Future work will focus on enhancing model performance
- Further reducing the energy per inference and latency
- Optimizing the system for long term field deployment

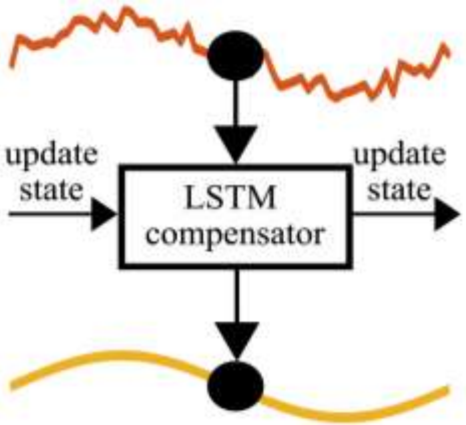

Thank You for Your Time

Any Questions?

Name: Joud satme
Email: Jsatme@email.sc.edu
Lab GitHub: github.com/arts-laboratory

Molinaroli College of Engineering and Computing
UNIVERSITY OF SOUTH CAROLINA