

Design and Validation of a Containerized IoT Architecture for Wearable Sensor Ingestion, Activity Recognition, and Multi-Source Data Visualization

Arash Omid Hosseinabadi

1. Summary

Wearable sensing devices generate continuous streams of motion and physiological data, but the infrastructure around them is just as important as the sensors themselves. This thesis builds and validates Wearable IoT Activity Pipeline, a fully containerized platform that takes raw sensor data all the way from acquisition through cleaning, storage, AI-based activity recognition, and live dashboard visualization. The system runs as twelve independent Docker services communicating over MQTT, stores time-series data in InfluxDB 3, runs an ONNX model for Human Activity Recognition (HAR), and surfaces results in Grafana.

Three distinct data flows were built and validated. The first replays a publicly available IMU dataset (SIDDHA) through the full pipeline to confirm correctness and reproducibility under controlled conditions. The second connects a real MetaWear wristband over Bluetooth, streaming live accelerometer and gyroscope data through cleaning, storage, and live HAR inference. The third integrates EEG and ECG signals from an open neuroimaging dataset as fake sensors, proving that the architecture can absorb fundamentally different sensor modalities, different schemas, sampling rates, and data shapes, without touching any existing service.

2. Architecture and Design

The platform is built around a clear, stage-by-stage pattern: each data source feeds into a protocol adapter or simulator, which publishes raw messages to the EMQX MQTT broker. A dedicated cleaner service validates, normalizes, and republishes canonical rows. The ingest service then stores clean sensor rows in InfluxDB, while the HAR service consumes clean IMU rows to run inference and write predictions to its own tables. Grafana reads from InfluxDB and displays everything in live dashboards.




Figure 1: Processing pattern used by the Wearable IoT Activity Pipeline platform.

The most important design rule is ownership separation: the ingest service owns all clean sensor storage, while the HAR service owns all prediction storage. This prevents the ingest layer from accumulating responsibilities it should not hold, and makes debugging far simpler, if a prediction is wrong, you know exactly which service to look at.

The HAR service supports two modes, switchable via a single environment variable. In database-polling mode it queries InfluxDB periodically and re-evaluates stored rows, useful for reproducible dataset evaluation where deterministic ordering matters. In MQTT-streaming mode it reacts to each incoming clean row as it arrives, enabling near-real-time inference on live watch data. Both modes build 40-sample sliding windows with a 20-sample stride, using a gyro-first channel layout and sum-based score aggregation, the configuration validated as optimal for the supplied ONNX model.

The system was built incrementally across seven phases, each validated before the next was introduced: MQTT transport, ingest and storage, dataset replay, HAR inference, live wearable integration, Grafana visualization, and finally the EEG/ECG extensibility demonstration.




Figure 2: Incremental design-and-validation methodology used in the project.

3. Visualization and Observability

Grafana connects directly to InfluxDB as its source of truth, so every panel reflects data that has already been cleaned, validated, and persisted, not transient MQTT traffic. The live watch dashboard updates every second and shows accelerometer and gyroscope signals from the MetaWear bracelet, the current predicted activity, model confidence, and a rolling prediction history. Row counters for both the clean IMU table and the predictions table are visible at a glance, making it easy to verify that data is flowing correctly during a live session.

The EEG/ECG dashboard follows the same structure: separate panels for EEG channel waveforms, ECG signal, row counts per table, and sampling-rate metadata. Figure 3 below shows the EEG/ECG monitoring dashboard; the live watch dashboard is laid out identically, replacing physiological signal panels with IMU and HAR panels.




Figure 3: Grafana EEG/ECG Monitoring Dashboard

Figure 3: Grafana EEG/ECG fake-sensor monitoring dashboard. The live MetaWear watch dashboard follows the same layout, showing IMU signals and HAR predictions in place of EEG/ECG panels.

4. Results and Validation

The HAR model reached 85% accuracy (119 out of 140 windows correct) on the watch-device, seven-activity subset, Typing, Brushing Teeth, Playing Catch, Dribbling, Writing, Clapping, and Folding Clothes. An unfiltered run across all devices and activities gave 15%, which reflects the expected drop from mixing phone-device data and unsupported activity classes into the evaluation, not a pipeline failure.

The live MetaWear session (recording phase4_live_validation_001, approximately 385 seconds) stored 9,599 clean IMU rows against an expected ~9,625, a difference of less than 0.3%, attributable to BLE startup behavior. The HAR service produced 463 predictions at roughly one every 0.83 seconds. Throughout the entire run, the ingest writer queue drained to zero, with no failed batches, no retries, and no dropped lines.

For the EEG/ECG extension, replaying 600 seconds at 100 Hz produced approximately 60,000 rows in each of the eeg_clean and ecg_clean tables, with the same zero-failure write health as the watch pipeline. New sensor families were integrated without changing a single line of existing service code.

5. Contributions

- A reusable IoT microservice architecture that cleanly separates protocol adaptation, validation, ingestion, inference, and visualization into independently deployable services.

- A unified platform supporting both real-time streaming and reproducible offline evaluation through two interchangeable HAR execution modes.
- Clear storage ownership, sensor rows belong to the ingest service, prediction rows belong to the HAR service, keeping schema responsibilities well-defined.
- Quantitative reliability validation through throughput, queue depth, retry counts, and data-loss monitoring across all phases.
- A practical extensibility pattern: EEG and ECG were added as new sensor families without modifying the existing watch or HAR pipelines.

6. Limitations and Future Work

The 85% HAR figure applies to the seven general-purpose Siddha activities on watch-device data, it does not represent tennis-specific stroke recognition, which would require a purpose-built labeled dataset. Live MetaWear sessions had no ground-truth labels, so the live pipeline was validated for reliability rather than classification accuracy. The EEG and ECG paths cover ingestion and visualization only; no machine-learning models or clinical interpretations are claimed. The platform is validated for thesis-scale live testing, not production deployment.

The most natural next steps are training a tennis-specific HAR model, adding a live annotation tool so coaches can label activities in real time, extending physiological signal processing for EEG/ECG, and exploring edge deployment to bring inference closer to the sensor.

7. Conclusion

This thesis shows that a well-structured IoT microservice architecture can handle real wearable acquisition, reproducible dataset replay, AI-based activity recognition, time-series storage, live visualization, and heterogeneous sensor extensibility, all within a single containerized platform. The key result is not any single accuracy number but the validated infrastructure pattern itself: a clean pipeline from sensor to dashboard that proved robust across three fundamentally different data sources and six months of incremental development.