{farm-twin}: An Open-Source ICAR ADE–Aligned Farm-Scale Digital Twin

Broadbent Matthew*[1], Muers Ross[1], Salavati Mazdak[1]

[1] Scotland's Rural College

Digital Twin (DT) technology can synchronise cyber-physical systems through continuous data exchange and real-time analytics, but agriculture is limited by fragmented, proprietary data that hinders interoperability and reduces the value to farms. {farm-twin} aims to provide an open-source, farm-scale digital twin aligned with the ICAR Animal Data Exchange (ADE) standard to support near real-time decision-making in dairy and livestock systems. The specific objectives are to implement an ICAR-compliant data backbone, demonstrate a modular architecture, and show that this design can support scalable, near real-time farm monitoring and integration of advanced analytics.

{farm-twin} implements a three-layer system. The data ingestion layer aggregates streams from animal wearables, environmental sensors, machinery telemetry, and farm management systems, mapping inputs to the ICAR ADE for livestock entities. The core engine layer uses a web framework and database to maintain current and historical farm state and exposes a secure API as a single source of truth, containerised for deployment at edge or cloud. A service layer hosts visualisation, analytics, and automation modules that consume the API for "what-if" simulations, alerts, and closed-loop control. Performance and scalability were explored using a realistic dairy-farm load profile with concurrent API access and measurement ingestion; descriptive performance metrics were used to assess near real-time suitability.

Under the dairy-farm scenario, {farm-twin} provided low-latency access to read/write endpoints and absorbed high-rate measurement streams without compromising responsiveness, supporting near real-time human-in-the-loop use while maintaining fidelity of the virtual farm. The main technical limitation observed was database write saturation at extreme ingestion rates, indicating clear optimisation paths such as sharding or complementary time-series storage. Keeping analytics and simulation outside the core, and enforcing ICAR ADE semantics at the ingestion layer, proved effective for interoperability and extensibility, allowing AI/ML modules (for example, health alerts, greenhouse gas indicators) to be deployed as independent services. Overall, {farm-twin} operationalises ICAR ADE as the backbone of a live digital twin and offers an open, extensible foundation for predictive, prescriptive and federated livestock applications.