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

Matthew Broadbent, Ross Muers, Mazdak Salavati

Motivation

- Farms produce heterogeneous data (sensors, animals, machinery, drones) that remain siloed
- Data fragmentation prevents system-level analysis, near-real-time control and automation
- A standards-based digital twin can unify streams and enable predictive, prescriptive workflows

What is a Digital Twin?

- A digital representation of a real-world object or system
- Mature concept within engineering and manufacturing

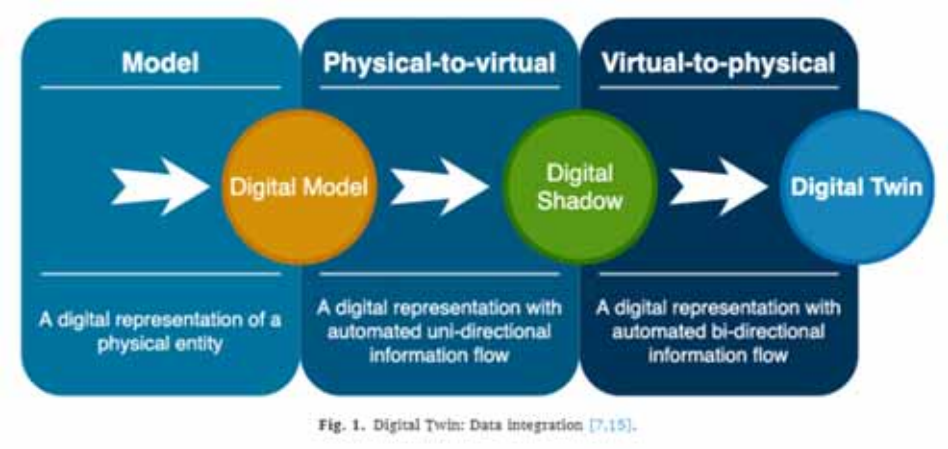

Fig. 1. Digital Twin: Data integration [7,15].

Purcell, W., & Neubauer, T. (2023). Digital Twins in Agriculture: A State-of-the-art review. SmartAgricultural Technology, 3, 100094. https://doi.org/10.1016/J.ATECH.2022.100094

Requirements

SRUC

- R1: Open standards & open-source to ensure interoperability and avoid vendor lock-in
- R2: High performance with predictable latency for API ops and high ingestion throughput
- R3: Modular, loosely coupled architecture so components can be swapped independently
- R4: Vertical & horizontal scalability from single node to federated deployments

What is { farm-twin }?

- A single point-of-truth for all data on a farm
- The **core** of a general-purpose Digital Twin
- Moving beyond single-purpose DTs towards reusability and modularity
- Application Programming Interface (API)-driven to enable programmatic interaction with the data
  - Both in and out of the DT
- Applications can then consume the data
  - Push back the results or actions (bi-directional flow)

Architecture

Layer 1: Mapping

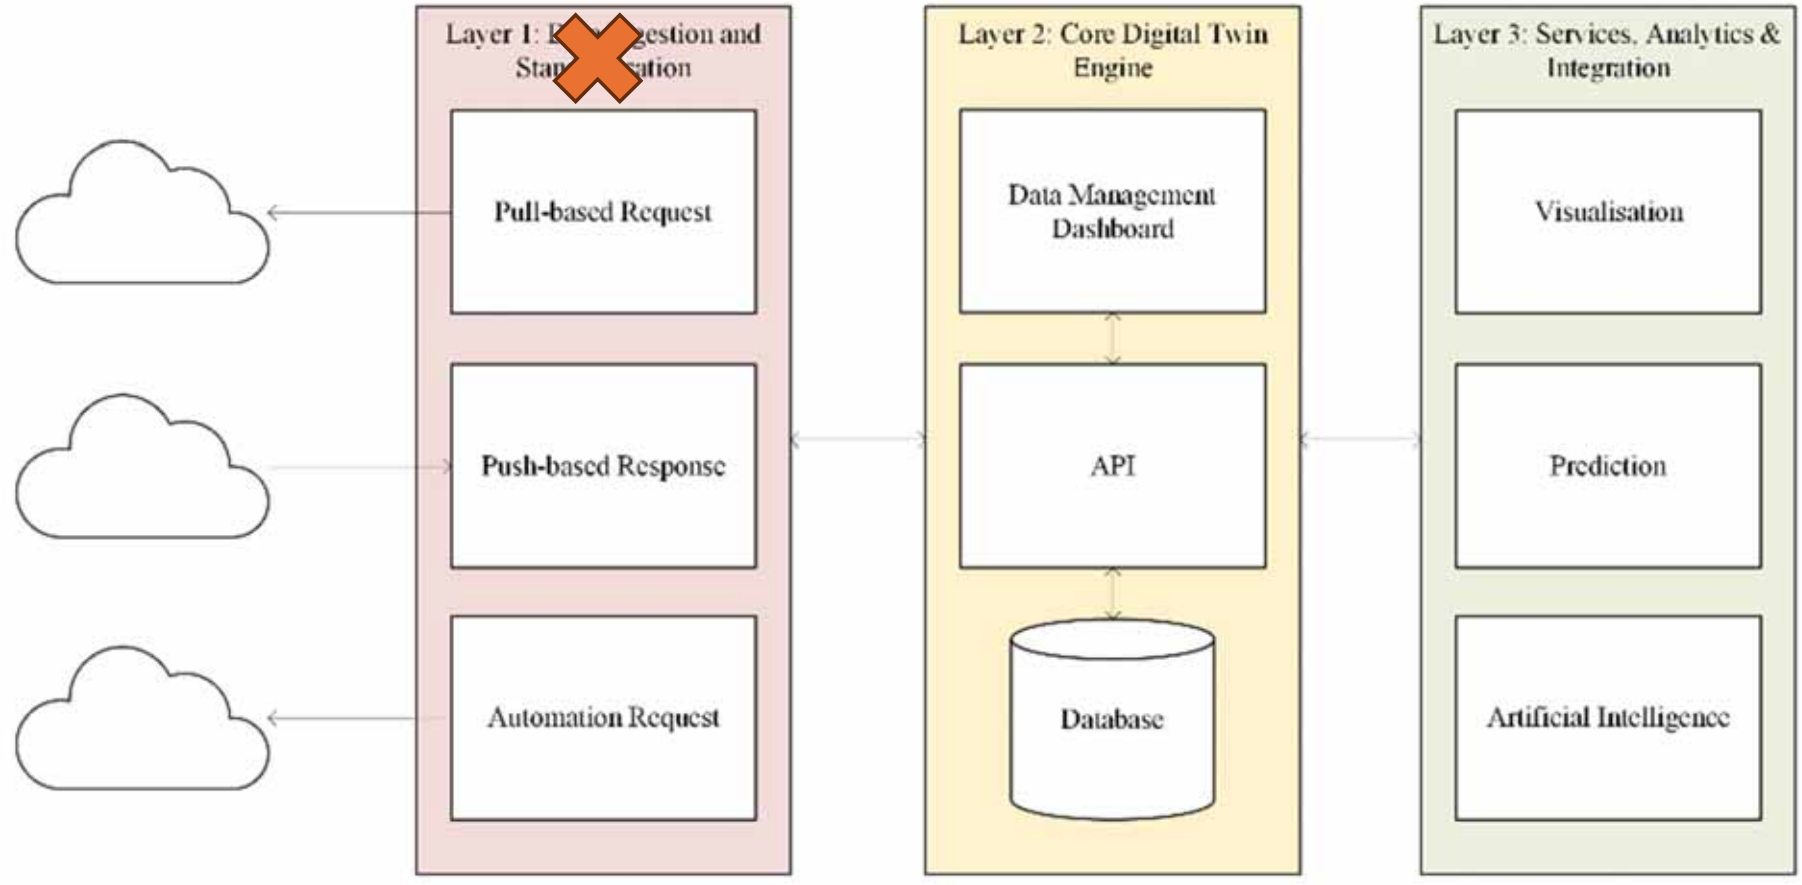

Prototype & Tech Stack

- Implemented in Python 3.11; FastAPI for API; MongoDB for data storage; Docker for deployment
- Aligned with ICAR ADE v1.5.0; enforced via Pydantic models
- Released free & open-source under AGPL-3.0
- Reproducible deployment, auto-generated API docs and pytest test suites
- Flexible deployment: local or cloud

RESTful HTTP APIs

SRUC

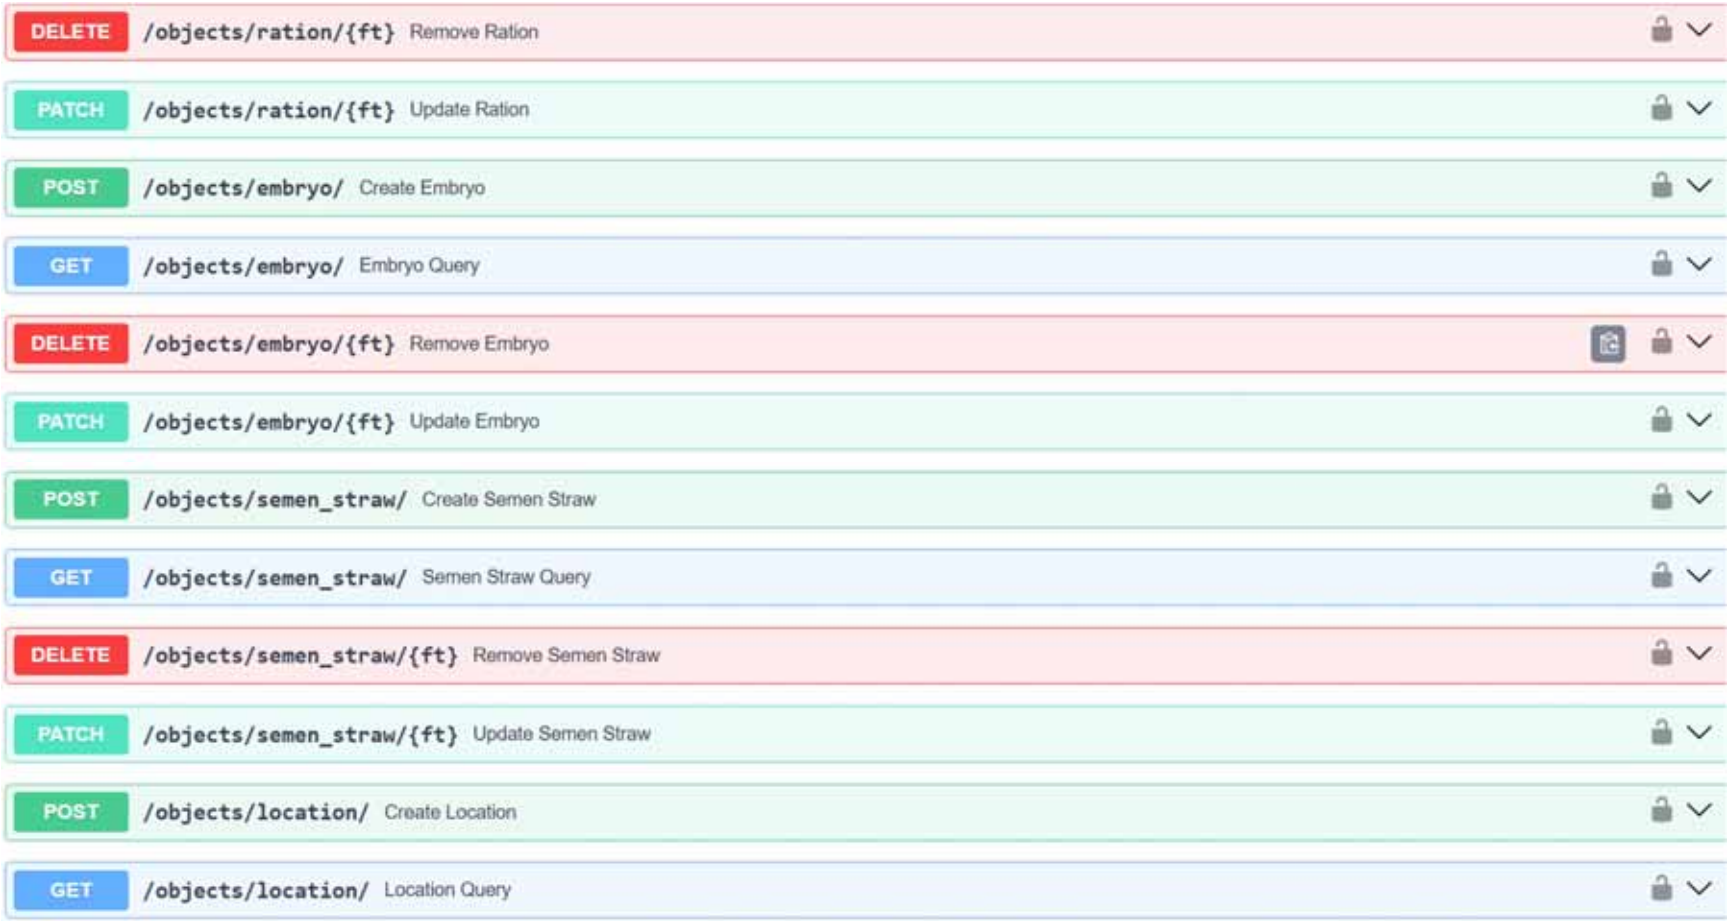

Evaluation

Endpoint	Concurrent Users	Median Latency (ms)	p95 Latency (ms)
GET /objects/animals/\{id\}	1	45	52
GET /objects/animals/\{id\}	10	64	94
GET /objects/animals/\{id\}	50	118	192
POST /measurements/samples	1	62	78
POST /measurements/samples	10	71	101
POST /measurements/samples	50	85	145

Test Machine: 4 vCPUs, 8GB RAM, and a solid-state drive, simulating a typical edge computing node or a small cloud instance.

Progress

- ICAR ADE ~70% complete
  - Set and Summary objects missing
  - As per WG discussions, we have made the 'meta' field required, and the 'source', 'sourceId' and 'modified' fields within that, also required
- JWT and OAuth Scopes for user authentication and permissions
- Non-animal related data integration

Conclusion & Next Steps

- {farm-twin} shows an ICAR ADE-aligned, API-centred core can support near-real-time farm monitoring
- Next: complete ADE alignment, open-source analytics & visualisations, additional technology ~~connectors~~ mappings, ….

- IEEE ICC Paper
- GitHub: digitaldairychain/farm-twin

Acknowledgements

- Funded by :
  - UKRI Digital Dairy Value Chain (99890)
  - Borderlands Inclusive Growth Deal (Dairy Nexus)

- Wider Team:
  - Ross Muers
  - Mazdak Salavati

Thanks for listening!

matt.broadbent@sruc.ac.uk