Journal of Innovations in Computer Science and Engineering
Shahid Beheshti Unversity

January 2026, Volume 3, Issue 2

A Spiking Neural Network–Driven Neuromorphic Digital Twin Architecture for Real-Time Resilient Resource Allocation in Smart Sports Clubs

Mohsen Esmaeili Sani✉, **CODE ORCID:** 0000-0002-2963-5021
Ph.D. Sport Management, Department of Sports Management, Faculty of Sports Sciences, University of Mazandaran, Babolsar, Iran E-mail: mhx71@yahoo.com

ABSTRACT
Smart sports clubs combine large conditioned volumes, volatile member demand and growing IoT exposure, yet their energy, space and storage resources are still allocated by static rules. We present NeuroTwin, a spiking-neural-network–driven neuromorphic digital twin that allocates these resources hour by hour, with decisions cheap and fast enough for real-time operation, and stays resilient to grid outages, equipment faults, demand surges and cyber attacks while protecting member privacy. A seven-layer cyber-physical architecture couples IoT sensing with local differential privacy, event-driven spiking perception at the edge (occupancy sensing, the NeuroCast forecaster and a spiking intrusion detector), a Spike-Gated Trust layer that quarantines falsified sensors and substitutes twin estimates, and seven decentralized population-coded spiking agents for zone cooling, space allocation and battery storage. The agents are trained by twin-in-the-loop privileged distillation: a receding-horizon MPC teacher on the twin labels the states the spiking students visit (DAgger). The twin is built exclusively from measured public data: 63 sports buildings of the Building Data Genome 2 project, 62 184 UC Berkeley gym-occupancy counts, the UCI occupancy dataset and the TON_IoT telemetry dataset, with checksummed provenance and leakage-controlled splits. The twin is a realistic composite of these measured traces with simulated disruptions and is evaluated offline, so all results are relative comparisons under identical conditions. The spiking perception modules reached 97.79%/98.33% occupancy accuracy and a mean intrusion-detection macro-F1 of 0.95 across seven IoT devices. On 15 held-out weeks, NeuroTwin improved the weekly operating objective of the incumbent rule-based controller by 22.5% under grid outages (paired Wilcoxon, Holm-corrected $p = 0.009$) and by 17.8% under compound disruptions (better in 12/15 weeks, $p = 0.075$), and was on par under nominal operation. It also outperformed the best of five strengthened model-free MARL baselines under outages (10.1%, $p = 0.027$), while a forecast-driven MPC remained stronger whenever its inputs could be trusted. Without retraining, it also improved compound-disruption performance at a second club in another climate, and it tolerated $\varepsilon = 2$ local differential privacy at no measurable cost ($\varepsilon = 1$: a few percent). Its spiking actors need 386.1 nJ per hourly decision for all seven agents, 1.7× less than matched ANN actors (an advantage that vanishes beyond $T \approx 8$ simulation steps) and many orders of magnitude less than online MPC.

CC BY

INDEX TERMS
Neuromorphic computing, spiking neural networks, digital twin, cyber-physical systems, multi-agent reinforcement learning, edge computing, Internet of Things, resilience, differential privacy, smart buildings, sports facilities.

I. INTRODUCTION

Sports facilities are among the most energy-intensive building types in the non-residential stock. Large conditioned volumes, high internal gains from exercising members, swimming-pool water treatment and long, demand-driven opening hours produce load profiles that differ markedly from offices or schools [1–3]. A recent systematic review of building automation for sports facilities found that only a small fraction of the operation-optimization literature addresses this building class, that most studies rely on simulation, and that facilities in hot climates are under-represented [1]. At the same time, sports clubs are being turned into instrumented cyber-physical systems (CPS): Internet of Things (IoT) meters, occupancy counters and networked building-management system (BMS) points stream data continuously [4, 5], and the facility can increasingly be represented by a *digital twin*, a synchronized virtual replica used for monitoring, what-if analysis and closed-loop control [6–9].

Turning such a twin into a real-time resource allocator raises four coupled requirements. **(i) Energy and cost.** Space conditioning, lighting, storage and on-site photovoltaics (PV) must be co-ordinated against time-of-use tariffs and demand charges, while members are served and thermal comfort is kept within standard bands [10]. **(ii) Resilience.** Grid outages, equipment derating and demand surges are operational realities; the allocation must degrade gracefully and recover quickly [11–13]. **(iii) Security and privacy.** Networked sensors can be spoofed by false-data injection (FDI) [14, 15], and member counts are personal data that should not leave the site unprotected [16, 17]. **(iv) Edge efficiency.** Decisions must be produced locally, continuously and cheaply, ideally on low-power hardware co-located with the BMS [18, 19].

Neuromorphic computing addresses the fourth requirement directly. Spiking neural networks (SNNs) communicate with sparse binary events, replace dense multiply-accumulate (MAC) operations by accumulate (AC) operations, and map naturally onto

Submit Date: 2026-10-04
Revise Date: 2026-11-09
Accept Date: 2026-11-16
✉ Corresponding author

event-driven processors such as TrueNorth, SpiNNaker and Loihi [20–24]. With surrogate-gradient training, deep SNNs now approach the accuracy of conventional networks on many tasks [25–28], and population-coded spiking actors have been shown to match deep actors in continuous control at one to two orders of magnitude lower inference energy [29, 30]. Neuromorphic systems are therefore being positioned as the substrate for always-on edge intelligence in CPS [31–34].

Yet three gaps separate these strands. **G1:** spiking reinforcement learning has been demonstrated almost exclusively on single-agent robotics benchmarks [29, 35–38], not on multi-agent resource allocation in buildings. **G2:** neuromorphic digital twins have been proposed conceptually, e.g. for cyber-defense [39] or wireless split computing [40], but not as an operational twin that closes the loop over measured building data. Recent sports-venue twins rely on synthetic data [41] or focus on spectator experience [42]. **G3:** multi-agent deep reinforcement learning (MARL) for HVAC and buildings [43–47] rarely evaluates resilience to compound disruptions, cyber attacks and privacy constraints jointly, and is seldom validated on measured data from more than one site [48, 49].

This paper closes these gaps with **NeuroTwin**, a spiking-neural-network–driven neuromorphic digital twin for real-time resilient resource allocation in smart sports clubs. Throughout, real time refers to the hourly decision cycle of the building-management system: every allocation is computed from the observations available at that hour, without a solver or look-ahead. The twin itself is validated offline on recorded data (scope in Section V-C). The study is driven by four research questions:

RQ1. Can event-driven SNNs provide the twin's perception functions (occupancy sensing, demand forecasting and intrusion detection) on measured data at a fraction of the inference energy of conventional networks?

RQ2. Can decentralized spiking agents allocate energy, space and storage in a multi-zone sports club better than the incumbent rule-based controller and than standard MARL, under nominal operation and under disruptions?

RQ3. How much do cyber-resilience (trust gating) and privacy (local differential privacy) cost or gain in allocation quality?

RQ4. Do the learned policies generalize to unseen weeks and to a different site and climate without retraining?

The contributions are:

C1. **A seven-layer neuromorphic CPS digital-twin architecture** (Fig. 1) that integrates IoT sensing, edge SNN perception, a trust layer, a physics-and-data twin core, spiking multi-agent decision making and an operator console, with explicit privacy and security boundaries.

C2. **A data-driven twin built exclusively from measured, public and citable datasets**: 63 sports buildings of the Building Data Genome 2 project [50], 62 184 gym-occupancy counts from UC Berkeley's Recreational Sports Facility [51], the UCI occupancy dataset [52] and the TON_IoT telemetry dataset [53, 54], with MD5-verified provenance, logged cleaning and leakage-controlled splits. No generative or synthetic data are used.

C3. **Spiking perception modules**: a parametric-LIF occupancy sensor, *NeuroCast*, a recurrent spiking forecaster with delta-event encoding and a learnable seasonal anchor, and a spiking intrusion detector, all evaluated against ten classical and modern baselines including DLinear and PatchTST [55, 56].

C4. **Spike-Gated Trust (SGT)**, an observation-repair mechanism in which sensor channels flagged by the spiking IDS are quarantined and replaced by twin estimates or spiking forecasts before they reach the agents.

C5. **Twin-in-the-loop privileged distillation for spiking MARL**: a receding-horizon MPC teacher that runs on the twin with privileged state supervises seven decentralized, population-coded spiking actors through DAgger [57, 58], with residual learning on the incumbent controller [59], cost-asymmetric pinball losses and an activity regularizer. We also strengthen five model-free MARL baselines with two twin-enabled devices, a shadow counterfactual reward and round-robin exploration, and compare all of them under identical observations and actuators.

C6. **A rigorous evaluation protocol** covering 24 controllers, six disruption families, 15 held-out weeks, multiple seeds, zero-shot transfer to a second club in another climate, an energy model for neuromorphic inference, sensitivity analyses of the trust layer and of the time-step budget, a power analysis, paired non-parametric tests with Holm correction [60–62] and an open operator console.

Relation to established methods. The ingredients of contribution 5, receding-horizon MPC [88], DAgger [57], learning by cheating [58], residual policies [59] and population-coded spiking actors [29], are established, and we use them as published. The methodological contribution lies in three elements that, to our knowledge, have not been combined or used before. (a) The digital twin itself is the privileged teacher, and its privileges (true thermal state, the full exogenous trajectory and the disruption schedule) are exactly the information that is missing or untrustworthy at run time; distillation therefore transfers disruption foresight into decentralized spiking students that see only local, possibly falsified and privatized observations after SGT. (b) The student losses mirror the cost structure of the objective (asymmetric pinball losses for space and cooling, residual on the incumbent), which keeps the students on the safe side of the teacher. (c) Spike-Gated Trust closes the loop from spiking intrusion detection to control by repairing observations with twin estimates, whereas earlier spiking IDS work stops at detection [71, 72]. The comparison with five model-free MARL baselines on identical observations and actuators isolates the effect of (a) and (b), and the ablations and sensitivity analyses in Tables IX and X quantify (c). Contributions 1, 2 and 6 are integration and evaluation contributions.

Section II reviews related work. Section III formalizes the allocation problem, threat and privacy models. Section IV presents the NeuroTwin architecture and algorithms. Section V describes data, twin construction and protocol. Section VI reports results, Section VII discusses implications and limitations, and Section VIII concludes.

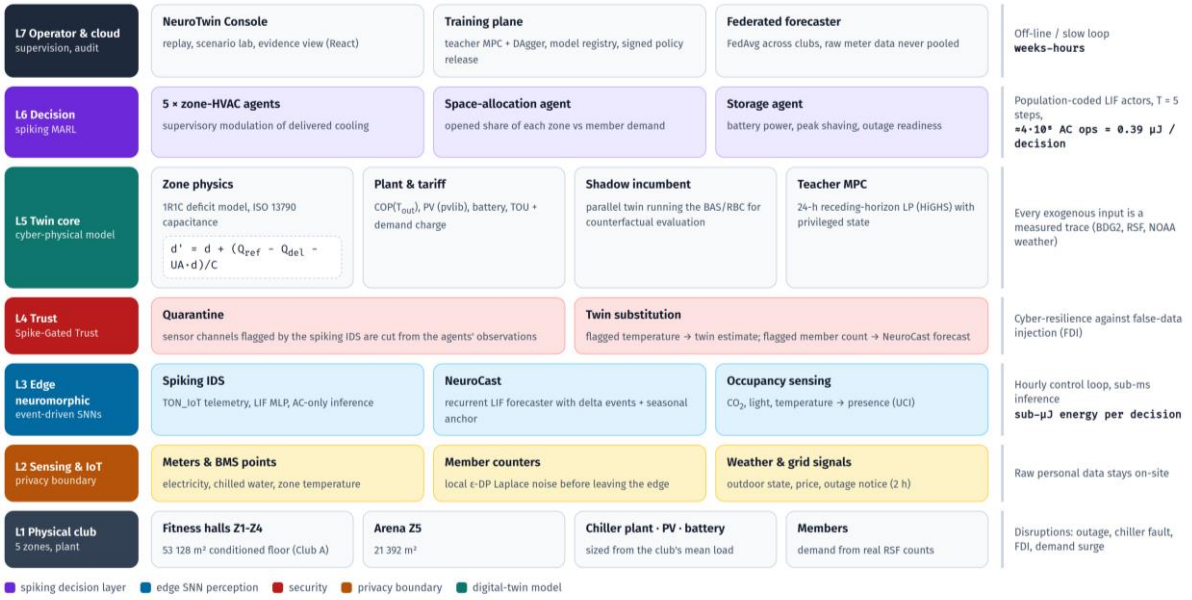


Fig. 1. NeuroTwin architecture. Seven layers from the physical club (L1) to the operator console (L7). IoT telemetry crosses a local-differential-privacy boundary (L2), is processed by event-driven spiking networks at the edge (L3) and filtered by the Spike-Gated Trust layer (L4) before reaching the twin core (L5) and the spiking multi-agent controllers (L6).

II. Related Work

A. Spiking Neural Networks and Neuromorphic Hardware

SNNs are often called the third generation of neural networks [20]. Neurons integrate input currents into a membrane potential and emit a binary spike when a threshold is crossed. Digital neuromorphic processors such as IBM TrueNorth [21], SpiNNaker [22] and Intel Loihi/Loihi 2 [23, 24], and mixed-signal chips such as DYNAP-SE2 [63], execute such networks event by event, so energy scales with spike activity rather than with network size. Direct training became practical with surrogate gradients [25, 26, 64], learnable membrane time constants [65], open frameworks [27, 28] and online learning rules such as e-prop. Spiking transformers [66] and spiking state-space models extend the model zoo. Recent perspectives argue that scaling neuromorphic systems to real applications requires end-to-end co-design with sensing and control [31–33].

B. Spiking Reinforcement Learning

Hybrid training of a spiking actor with a deep critic was introduced for mapless navigation [30] and generalized by PopSAN, whose population coding of observations and actions enabled continuous control comparable to deep actors with 140× lower energy on Loihi [29]. Directly trained deep spiking Q-networks reached human-level Atari control [35], multi-scale dynamic coding improved spiking actors [36], fully spiking actors removed floating-point read-outs [37], and event-based obstacle avoidance has been deployed on neuromorphic accelerators [38]. Policy distillation into spiking students has recently been shown to cut training cost for neuromorphic control [67]. All of these works are single-agent robotics studies; to our knowledge, spiking actors have not been used for multi-agent energy-resource allocation in buildings.

C. Spiking Time-Series Forecasting and Intrusion Detection

SNN forecasters encode real-valued series by pulse-width modulation [68], derivative (delta) spikes [69] or learned temporal encoders [70]. They reach accuracy competitive with recurrent networks at lower operation cost on load-forecasting benchmarks. For IoT security, convolutional SNNs [71] and SNN/decision-tree hybrids [72] detect attacks on resource-constrained devices, and a neuromorphic *cyber-twin* has been proposed conceptually for cognitive defense [39]. None of these studies feeds the spiking detector's output back into a control loop.

D. Digital Twins for Buildings and Sports Venues

Digital-twin surveys distinguish digital models, shadows and twins by the direction of data flow [8, 9, 73], and stress security and privacy as open issues [74]. For buildings, twins are used for monitoring, anomaly detection, operational optimization and predictive maintenance [75], with validated real-time architectures [76], district-scale CPS [77] and city-scale benchmarking [78, 79]. Cognitive twins that diagnose drift under extreme heat have recently been proposed [80]. Sports-venue twins remain scarce. Existing work addresses spectator experience [42], facility management [81] and, most recently, a diffusion-model twin trained largely on synthetic samples [41]. Energy studies of sports centres focus on audits and retrofits [2, 3, 82, 83]. Neuromorphic twins appear only in communications [40, 84] and in a conceptual security architecture [39].

E. Multi-Agent Reinforcement Learning and MPC for Building Energy

Reviews of RL for building control highlight sample inefficiency, safety and the gap between simulation and deployment [48, 49, 85–87]. Model predictive control (MPC) remains the reference for optimal operation when a model is available [88]. Multi-agent formulations decompose zones or subsystems into cooperating agents: attention-based MADRL for multi-zone HVAC [43], MARCO for AHU/chiller coordination [44], grid-aware multi-building control [45], cooling-water systems [46], occupant-centric control [47], federated acceleration [89] and hierarchical DER coordination [90]. Generic CTDE algorithms include MADDPG [91], QMIX [92], FACMAC [93], IPPO [94] and MAPPO [95]. Imitation of an expert controller with DAgger [57] and *learning by cheating* from a privileged teacher [58] are established in robotics but rarely used to train decentralized building agents.

F. Security, Privacy and Resilience

False-data injection can bypass residual-based bad-data detection [14], and CPS surveys catalogue sensor spoofing as a primary threat [15]. Occupancy data reveal individual behaviour [17]; differential privacy (DP) provides formal guarantees [16, 96], and federated learning keeps raw data local [97, 98]. Resilience is commonly quantified by the area under a performance curve during and after a disruption [11–13].

TABLE I

Positioning of NeuroTwin against representative related work (✓ addressed, ◐ partially, – not addressed)

Work	SNN	Multi-agent	Digital twin	Measured data	Resilience	Security	Privacy	Edge energy
PopSAN [29]	✓	–	–	–	–	–	–	✓
ILC-SAN [37]	✓	–	–	–	–	–	–	✓
MADRL-HVAC [43]	–	✓	–	◐	◐	–	–	–
Grid-aware MADRL [45]	–	✓	–	✓	–	–	◐	–
FA-MADRL [89]	–	✓	–	◐	–	–	✓	–
Building DT [76]	–	–	✓	✓	–	–	–	–
Cognitive DT [80]	–	–	✓	◐	✓	◐	–	–
Neuromorphic cyber-twin [39]	✓	–	✓	–	–	✓	–	◐
Sports-venue DT [41]	–	–	✓	◐	–	–	–	–
Conv-SNN IDS [71]	✓	–	–	✓	–	✓	–	✓
NeuroTwin (this work)	✓	✓	✓	✓	✓	✓	✓	✓

III. System Model and Problem Formulation

A. Cyber-Physical Club Model

The club consists of $Z = 5$ conditioned activity zones z (four fitness halls and one arena in Club A), a central chiller plant, a battery of capacity B_E (kWh) and power B_P (kW), and rooftop PV of P_{pv} kWp. Time advances in hourly steps t. For each zone, the twin receives the measured non-HVAC electricity $E_z(t)$ and the measured thermal demand $Q_z(t)$ of a real BDG2 building, the outdoor temperature $T_{out}(t)$, and the member demand $D_z(t)$. The state of zone z is its thermal deficit d_z (K) relative to the set-point, positive when the zone is warmer than intended. A lumped 1R1C model with ISO 13790 parameters [99] gives

$$d_z(t+1) = d_z(t) + \frac{Q_z\^{ref}(t) - Q_z\^{del}(t) - UA_z d_z(t)}{C_z} \quad (1)$$

where $C_z = 165$ kJ m^{-2} K^{-1} $\times A_z$ (ISO 13790 "medium" class), $UA_z = C_z/24$ h, $Q_z^{ref} = Q_z(1 - 0.15(1 - u_z))$ is the demand corrected for closed space, and Q_z^{del} is the cooling actually delivered. Chiller electricity is $Q^{del}/\mathrm{COP}(T_{out})$ with $\mathrm{COP} = \mathrm{clip}(6.2 - 0.09(T_{out} - 20), 2.5, 6.2)$. Plug and lighting load scales with the opened share $u_z \in [0.1, 1]$ as $E_z(\rho_z + (1 - \rho_z)u_z)$, where ρ_z is the 5th-percentile-to-mean ratio of the zone's own meter. Members served are $\min(D_z, u_z \cdot \mathrm{Cap}_z)$ with $\mathrm{Cap}_z = A_z/10$ m² (ASHRAE 62.1 weight-room density [100]). Absolute comfort values are model-based; all controller comparisons share the same model and are therefore relative (Section VII-B).

B. Decentralized Partially Observable Markov Game

Seven agents act every hour: five zone-HVAC agents $i = 1..5$ choose $a_i \in [-1, 1]$, which sets $Q^{del}_i = \mathrm{clip}(Q^{ref}_i(1 + 0.6a_i), 0, Q^{cap}_i)$; a space-allocation agent chooses the opened shares $\mathbf{u} \in [0.1, 1]^5$; and a storage agent chooses the battery power $b \in [-1, 1]$. Agent i observes only a local vector o_i (23–37 features: calendar, tariff and next-three-hour prices, outdoor temperature, PV, outage notice, state of charge, running peak, its zone's deficit and demand, the quarantine flag and spiking forecasts). The team shares the reward

$$r(t) = -[\pi(t)g_t + DC \cdot \Delta peak + \lambda_c \sum_z viol_z + \lambda_u U_t + VoLL \cdot S_t] \quad (2)$$

where $\pi(t)$ is a three-period time-of-use tariff (0.065/0.10/0.22 \$ kWh^{-1}), g_t the grid import, DC = 3.5 \$ kW^{-1} per week (15 \$ kW^{-1} month^{-1}) applied to increments of the running peak, $\mathrm{viol}_z = \max(|d_z| - \mathrm{band}, 0)$ in occupied zones with band = 1.5 K (4 K unoccupied) following ASHRAE 55 [10], $\lambda_c = 25$ \$ K^{-1} h^{-1}, U_t the members turned away ($\lambda_u = 4$ \$ per member-hour) and S_t the unserved energy during outages (VoLL = 10 \$ kWh^{-1}). The objective is to maximize the expected weekly return. The formulation is a Dec-POMDP; centralized training with decentralized execution (CTDE) is used throughout.

C. Disruption, Threat and Privacy Models

- **Grid outage**: 2–6 h islanding with a 2-h utility notice; PV and the battery must cover the load, otherwise load is shed proportionally.
- **Equipment fault**: the chiller serving one zone is derated to 50% for 6–24 h.
- **Cyber attack (FDI)**: for 6–24 h an adversary adds ±3 K to one zone's temperature sensor and, with probability 0.5, scales the member counter by 0.2 to induce closures [14].
- **Demand surge**: member demand rises by 80% for 3–6 h (tournament or event day).
- **Privacy**: member counts are released through an ε-locally-differentially-private Laplace mechanism with sensitivity of one member [16].

D. Resilience Metrics

Following the resilience-triangle literature [11, 12], we define the instantaneous service performance as the product of the served-member ratio, a comfort factor and the served-energy ratio:

$$Q(t) = \frac{served(t)}{D(t)} \cdot (1 - min(\Sigma viol/5, 1)) \cdot r_{serve}(t), \quad RI = \frac{1}{T}\sum_{t=1}^{T} Q(t) \tag{3}$$

where RI ∈ [0, 1] is the resilience index (normalized area under the performance curve), and the recovery time is the number of hours from the first drop of Q below 0.95 until it stays above 0.95 for three consecutive hours.

IV. The NeuroTwin Architecture

A. Layered Design

NeuroTwin separates concerns into seven layers (Fig. 1). The **physical layer** (L1) is the club with its zones, plant and members. The **sensing layer** (L2) contains meters, BMS points and member counters; it is also the privacy boundary, where member counts are perturbed by the local DP mechanism before leaving the edge. The **edge neuromorphic layer** (L3) hosts three event-driven SNNs: an occupancy sensor, the NeuroCast forecaster and a spiking intrusion detector. The **trust layer** (L4) implements Spike-Gated Trust. The **twin core** (L5) holds the zone physics, plant, tariff and disruption models, a *shadow* copy of the incumbent controller for counterfactual evaluation, and the MPC teacher. The **decision layer** (L6) contains the seven spiking agents. The **operator layer** (L7) hosts training, model release and the NeuroTwin Console. Fast loops (L2–L6) run hourly at the edge; slow loops (L7) run off-line.

B. Spiking Neuron Model and Training

All spiking networks use parametric leaky integrate-and-fire (PLIF) neurons with soft reset [65]:

$$v_t = \beta \odot v_{t-1} + Ws_t\text{^}\{in\}, s_t = H(v_t - \theta), v_t \leftarrow v_t - \theta s_t \tag{4}$$

with a per-neuron learnable leak $\beta = \sigma(w_\beta)$ and threshold $\theta = 1$. The Heaviside step H is differentiated with the fast-sigmoid surrogate $\partial s/\partial v \approx (1 + k|v - \theta|)^{-2}$, $k = 5$ [25, 64]. Networks are trained by backpropagation through time with Adam [101].

Energy model

Following common practice for neuromorphic estimates [29, 102], inference energy is computed from counted operations: a synaptic event triggered by a binary spike costs one accumulate (AC, 0.9 pJ), while real-valued inputs, such as the first layer of a directly encoded network, cost multiply-accumulates (MAC, 4.6 pJ, 32-bit, 45 nm):

$$E = N_{AC} \cdot 0.9\,pJ + N_{MAC} \cdot 4.6\,pJ \tag{5}$$

Every SNN forward pass in our code counts $N_{AC} = \Sigma$ (spikes × fan-out) and N_{MAC} exactly. Conventional baselines are charged their analytic MAC counts. Memory access and static power are not modelled; the estimate is a lower bound that is consistent across models.

C. NeuroCast: Spiking Forecaster with Seasonal Anchor

NeuroCast unrolls a two-layer recurrent PLIF network (128 neurons each, recurrent synapses on the first layer) over an input window of L hours, presenting each hour for two simulation sub-steps. Inputs are the normalized target, weather and calendar features plus send-on-delta ON/OFF events of the target (threshold 0.15 σ), which are binary by construction [69]. A leaky non-spiking read-out integrates output spikes. Because pure rate codes struggle to represent levels precisely [70], the forecast for horizon h is anchored on two seasonal references with learnable weights:

$$\hat{y}_{t+h} = w_{1h}x_t + w_{2h}x_{t+h-24} + SNN_h(x) \tag{6}$$

The anchor costs $2H$ MACs per forecast and lets the spiking core learn deviations from persistence and from the daily cycle rather than absolute levels. Inside the twin, occupancy forecasts are *cross-fitted*: weeks are split into four folds aligned with the control split, and each fold is predicted by a model trained on the other three, so every forecast seen by an agent is out-of-sample.

D. Spiking Intrusion Detection and Spike-Gated Trust

The spiking IDS is a two-layer PLIF network (128–64, $T = 8$ steps) on causal features of each telemetry stream: current and three lagged readings, first difference, five-sample rolling standard deviation and the logarithm of the inter-arrival time. Absolute timestamps are excluded to avoid temporal leakage. In the twin, the detector's empirical true- and false-positive rates measured on TON_IoT thermostat and weather devices (Section VI-C) drive per-channel flags.

Spike-Gated Trust (Fig. 4) repairs the agents' observations: if channel k is flagged at time t, its value is quarantined and replaced by a twin estimate. For zone temperatures, the estimate propagates the 1R1C model of (1) from the last trusted reading using the plant's commanded cooling, with deliberately mis-specified parameters (±20% on C and UA) so that substitution is not an oracle. For member counts, the replacement is the NeuroCast one-hour-ahead forecast. The quarantine flag itself is part of the observation, so agents can learn to act conservatively while a sensor is distrusted.

E. Population-Coded Spiking Agents

Each agent type (HVAC, allocation, storage) has one spiking actor; the five HVAC agents share parameters and receive a one-hot zone identifier. Following PopSAN [29], every observation dimension o_j is encoded by eight Gaussian receptive fields with learnable centres μ and widths σ, $a_{jk} = \exp(-(o_j - \mu_{jk})^2/2\sigma_{jk}^2)$. The stimuli drive deterministic integrate-and-fire encoder neurons, so every synaptic layer of the actor receives binary spikes and costs accumulates only, as in PopSAN. The encoder spikes feed two PLIF layers of 128 neurons. Each action dimension is decoded from the firing rates of eight output neurons over $T = 5$ steps by a learnable linear read-out followed by tanh. The actor outputs a *residual* added to the incumbent rule-based action, clipped to [−1, 1] [59]. The space-allocation action is projected onto the safe set $u_z \geq \hat{D}_z/\mathrm{Cap}_z + 0.02$, where $\hat{D}_z$ is the observed demand. The same projection applies to every controller, so it adds no advantage to ours.

F. Twin-in-the-Loop Privileged Distillation

Model-free MARL faces three difficulties here: the comfort term is cliff-shaped, credit assignment across seven agents is noisy, and exogenous variation (weather, attendance, disruptions) dominates the return. In preliminary runs, unmodified MAPPO did not improve on its rule-based anchor within 150 updates. NeuroTwin therefore exploits the twin itself. A **privileged teacher** solves, at every visited state, a 24-hour receding-horizon linear program over the twin's linearized dynamics (1), power balance, battery, demand-charge and comfort slacks. The LP is solved with HiGHS [103] in about 13 ms and uses the *true* state and exogenous trajectory, including the disruption schedule. The **students** are the seven spiking actors, which see only their local,

possibly attacked and privatized observations after SGT. Training follows DAgger [57] in the spirit of *learning by cheating* [58], as summarized below.

Algorithm 1. *Twin-in-the-loop privileged DAgger for spiking agents*

```
Input: twin 𝒯 with disruption sampler, teacher π* (24-h MPC, privileged), students {π_θ_{k}} (spiking, local),
incumbent π_RBC, rounds R = 16, decay β = 0.7.
Initialize datasets 𝒟_k ← ∅ for each agent type k.
for r = 0, …, R − 1 do
    Sample 16 training weeks and disruption schedules; reset 𝒯.
    for t = 1, …, 168 do
        a ← π(s_t) (privileged LP); a^b ← π_RBC(o_t); add (o_{i,t}, a*_i − a^b_i) to 𝒟_{type(i)} for all agents i.
        Execute a* with probability β^r, otherwise clip(a^b + π_θ(o_t)); step 𝒯.
    Train each π_θ_{k} on 𝒟_k: pinball loss with τ = 0.8 (allocation) and τ = 0.6 (HVAC), smooth-L1 (storage).
    Evaluate on 8 fixed validation weeks; keep the best parameters.
return spiking students for edge deployment
```

The asymmetric pinball losses encode the cost asymmetry of the objective: under-allocating space or under-cooling is costlier than the opposite error, so students are biased toward the safe side of the teacher's action.

G. Enhanced Model-Free MARL Baselines

For a fair comparison, all model-free baselines (centralized PPO, IPPO, MAPPO with ANN or spiking actors, and MATD3) receive the same residual parametrization, safety projection and observations. On-policy methods are further strengthened with two twin-enabled devices: a **shadow counterfactual reward**, in which the twin replays the incumbent controller on the same week and disruptions and agents learn from $r - r^{\text{shadow}}$ (an action-independent baseline that leaves the optimal policy unchanged), and **round-robin exploration**, in which only one agent group explores per update. MATD3 uses a centralized twin-Q critic over the joint residual action, delayed policy updates and target smoothing [91, 93, 104].

H. Complexity

Per decision, a spiking actor with n_{in} inputs, population size p, hidden width h and T steps needs $n_{\text{in}} \cdot p$ MAC-equivalents for the receptive fields and $\approx T \cdot (\rho_0 n_{\text{in}} p \cdot h + \rho_1 h^2 + \rho_2 h \cdot m \cdot p)$ ACs, where ρ_0, ρ_1 and ρ_2 are the firing rates of the encoder, hidden and output populations and m is the action dimension. A dense ANN of the same width needs $n_{\text{in}} h + h^2 + h \cdot m$ MACs. With measured firing rates of 0.363 this is about 420,194 ACs and 1,720 MACs per hour for all seven agents (386.1 nJ), versus 660.6 nJ for the ANN actors and roughly 0.2 J of CPU energy for one 24-h MPC solve (≈13 ms at ≈15 W). The teacher is therefore used only during training.

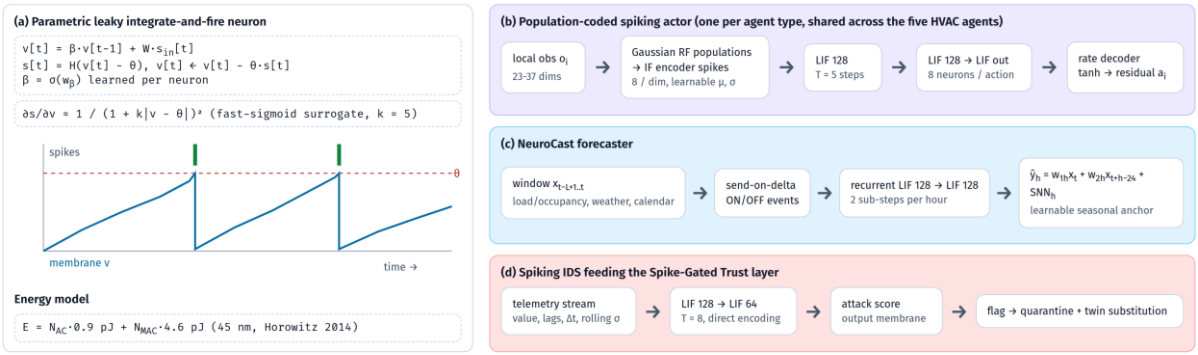


Fig. 2. Spiking building blocks: (a) PLIF neuron with surrogate gradient and the energy model; (b) population-coded spiking actor; (c) NeuroCast forecaster with send-on-delta events and seasonal anchor; (d) spiking IDS feeding the Spike-Gated Trust layer.

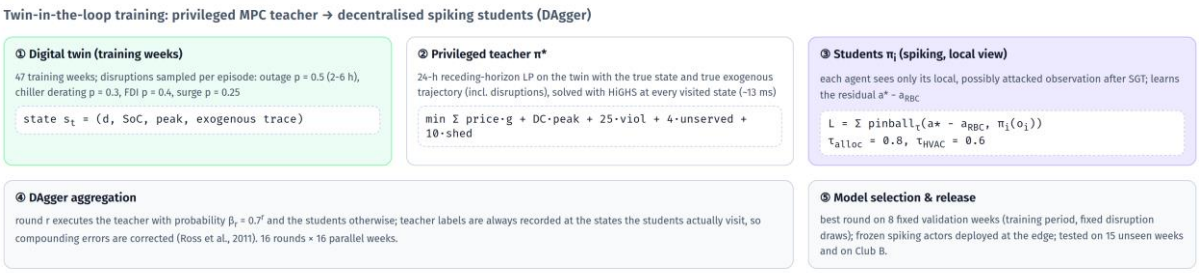


Fig. 3. Twin-in-the-loop training. A privileged 24-h MPC teacher runs on the digital twin and labels the states visited by the decentralized spiking students (DAgger). Students learn residuals on the incumbent controller with cost-asymmetric losses.

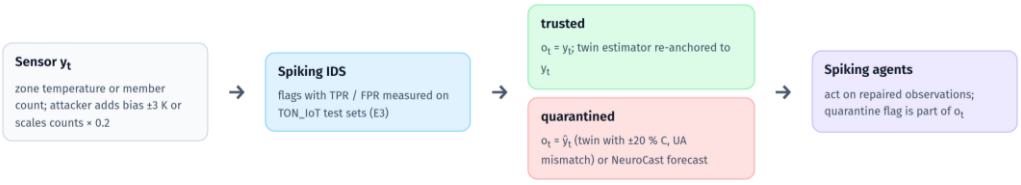


Fig. 4. Spike-Gated Trust. Readings flagged by the spiking IDS are quarantined and replaced by twin estimates (zone temperature) or NeuroCast forecasts (member counts).

V. Experimental Setup

A. Datasets

All data are measured, public and citable (Table II). Raw files are downloaded by a script and verified against MD5 checksums. No generative or synthetic data augmentation is used anywhere, which keeps the findings tied to observed operation and lets them generalize.

TABLE II
Measured datasets used in this study

Dataset	Source	Content used	Size	Role
BDG2 [50]	ASHRAE GEPIII / Sci. Data	hourly electricity & chilled water of 63 sports buildings (gymnasia, fitness centres, stadia, pools) at 12 sites; NOAA weather	2 years × 8 784 h	twin loads, E2-B
UC Berkeley RSF [51]	Kaggle	gym people counts every 10 min with calendar flags	62 184 records	member demand, E2-A
UCI Occupancy [52]	UCI / authors	temperature, humidity, light, CO_2, ground-truth presence	20 560 records	E1
TON_IoT telemetry [53, 54]	UNSW Canberra	7 IoT devices incl. thermostat & weather, 7 attack classes	401 119 records	E3, SGT rates

B. Data Quality, Cleaning and Validity

Cleaning is deterministic and logged. A *seasonal* Hampel filter compares each value with the median of the same clock hour on the surrounding ±7 days (8 scaled MADs), so genuine daily peaks such as evening events are preserved. Gaps of up to 6 h are linearly interpolated, and longer gaps are filled with the hour-of-week median. RSF hours without records between 00:00 and 06:00 (facility closed) are set to zero. Across the ten Club A/B meter series, raw missingness ranged from 0.01% to 7.60% and flagged outliers from 0.07% to 4.16%. Construct validity is supported by the physics: daily chilled-water demand of Club A correlates with outdoor temperature at Pearson r = 0.90 (Fig. 5c), and member demand shows the expected weekday-evening peak (Fig. 5a).

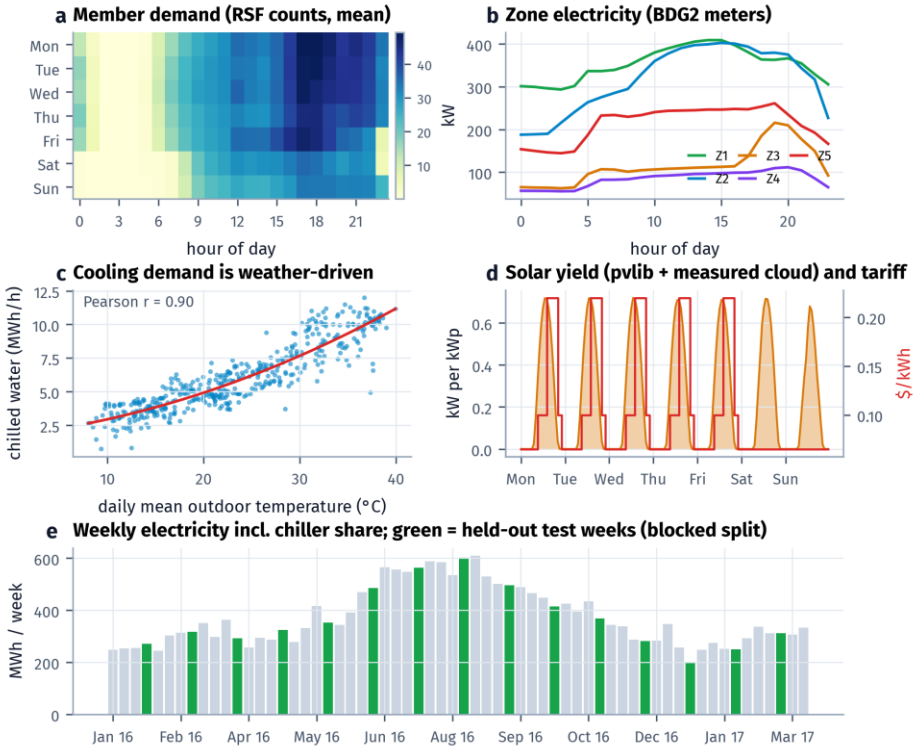


Fig. 5. Measured data behind the Club A twin: (a) mean RSF member counts by weekday and hour; (b) mean daily electricity profile of the five zones; (c) daily chilled-water demand vs outdoor temperature; (d) PV yield computed with pvlib from measured cloud cover, and the TOU tariff; (e) weekly electricity including the chiller share, with held-out test weeks in green.

C. Twin Construction

Club A is assembled from five BDG2 buildings at the *Fox* site (Tempe, Arizona, ASHRAE climate 2B): four fitness centres and one sports arena, 74 520 m² in total. Club B, used only for zero-shot transfer, is assembled from five buildings at the *Bull* site (Austin, Texas, climate 2A): a gymnasium, two stadia, an athletic building and a swimming pool. One Fox building, a pool, was excluded because its hot-water meter averages ≈40 MW, which is physically implausible and indicates a unit inconsistency in the source. The simulation period is 4 January 2016 to 12 March 2017, the span over which RSF counts exist. The battery is sized at 1 h of mean club load (energy) and 0.5 h (power), PV at 0.6 kWp per kW of mean load. PV output is computed with pvlib [105] from Ineichen clear-sky irradiance, the Kasten–Czeplak cloud model [106] with measured cloud cover, the Erbs

decomposition, isotropic transposition (20° tilt, south), the Faiman cell-temperature model and PVWatts with 14% system losses. Table III lists all parameters.

Scope of the twin. The twin is a synchronized, executable model of a club driven by measured traces, but in this study it is operated offline: recorded 2016–17 data are replayed hour by hour as if they arrived live, and disruptions are injected by the twin's scenario generator. Three consequences follow. First, real time refers to the decision cycle: each hourly allocation is computed from the observations available at that hour by actors that need 386.1 nJ and no solver per decision; the closed loop with a physical BMS, its communication latencies and actuator dynamics were not tested. Second, the club is a realistic composite (loads of five BDG2 buildings at one site, member demand of the UC Berkeley RSF, IDS operating points from TON_IoT), so the results describe a plausible sports club rather than a specific facility. Third, all controllers are evaluated on the same twin, weeks and disruption draws, so the conclusions concern relative performance under identical conditions; absolute costs and comfort values are model-based (Section VII-B). The zero-shot results on Club B, built from another site and climate, test whether these relative conclusions survive a change of data source; operational validity still requires a pilot deployment.

TABLE III

Digital-twin parameters (Club A)

Parameter	Value	Basis
Zones / floor area	5 / 74 520 m²	BDG2 metadata
Thermal capacitance C_z	165 kJ m^{-2} K^{-1} × A_z	ISO 13790 medium class [99]
Time constant C/UA	24 h	commercial buildings
Comfort band (occ. / unocc.)	±1.5 K / ±4 K	ASHRAE 55 [10]
Capacity Cap_z	A_z / 10 m² per member	ASHRAE 62.1 [100]
Chiller COP	$6.2 - 0.09(T_{out} - 20)$, [2.5, 6.2]	air-cooled plant curve
Battery	2,440 kWh / 1,220 kW, η = 0.95	1 h / 0.5 h of mean load
PV	1,464 kWp	0.6 kWp per kW mean load
Tariff	0.065 / 0.10 / 0.22 (0.15 winter) \$ kWh^{-1}	3-period TOU, peak 14–20 h weekdays
Demand charge	15 \$ kW^{-1} $month^{-1}$	commercial structure
Comfort / unserved / VoLL	25 \$ $K^{-1}h^{-1}$ / 4 \$ / 10 \$ kWh^{-1}	objective weights
Episode / split	168 h; 47 train, 15 test weeks	blocked: week mod 4 = 3 is test

D. Experiments, Baselines and Metrics

- **E1 – Occupancy sensing** (UCI): the original train/test1/test2 protocol [52]; logistic regression, random forest, gradient boosting, MLP and the spiking MLP; accuracy, F1, MCC, AUC; five seeds for neural models.
- **E2 – Day-ahead forecasting** (24 hourly values issued at midnight from a 48-h window): (A) RSF occupancy with a chronological split (train to Aug 2016, validation Sep–Oct, test Nov 2016–Mar 2017); (B) electricity of sports buildings with a *leave-sites-out* split: models are trained on 2016 data of ten sites and tested on 2017 data of the 23 buildings of the Fox and Bull sites, which are never seen in training. Baselines: seasonal naive, Ridge, MLP, LSTM [107], GRU [108], TCN [109], Transformer [110], DLinear [55], PatchTST [56], NeuroCast without anchor, and federated NeuroCast (FedAvg over the ten training sites [97]). Metrics: MAE, RMSE, CV(RMSE) and NMBE per ASHRAE Guideline 14 [111]; Diebold–Mariano tests [112]; three seeds.
- **E3 – Intrusion detection** (TON_IoT): binary and multi-class detection for all seven devices with stratified 70/30 splits over three seeds; random forest, histogram GBDT, MLP and the spiking MLP; macro-F1, TPR, FPR and AUC.
- **E4 – Resilient resource allocation**: 24 controllers. Non-learning: as-operated BAS (recorded operation), an expert rule-based controller (RBC: pre-cooling, setback, demand-following space use, TOU battery with demand limiting), RBC + SGT, a realistic 24-h MPC with forecasts (and + SGT), the privileged MPC teacher and an actuator-consistent perfect-information weekly LP (upper bound). Model-free MARL: centralized PPO, IPPO, MAPPO (ANN and SNN actors) and MATD3. Ours: NeuroTwin, its ablations (− SGT, − NeuroCast forecasts, − population coding, $T \in \{1, 3, 5, 8\}$), a distilled ANN counterpart and DP variants $\varepsilon \in \{2, 1, 0.5, 0.1\}$. Evaluation covers 15 held-out weeks × six scenario families (nominal, outage, chiller fault, cyber FDI, surge, compound) with fixed disruption draws shared by all controllers, plus zero-shot transfer to Club B (nominal and compound).

Statistical protocol: learning methods are trained with three seeds (one seed for sweeps). For every scenario we report the mean over weeks and seeds with 95% bootstrap confidence intervals, and we test the seed-averaged per-week objective of NeuroTwin against RBC, MPC-24h and the best model-free MARL method with paired Wilcoxon signed-rank tests, Holm-corrected across scenarios [60–62]. This follows recommendations for statistically reliable RL evaluation [113, 114]. Power: with n = 15 paired weeks, the exact two-sided Wilcoxon signed-rank test has 80% power for standardized paired effects $d_z \geq 0.85$ at $\alpha = 0.05$ and $d_z \geq 1.15$ at the strictest Holm step ($\alpha = 0.05/8$), estimated with 4,000 Monte-Carlo replications per effect size; the smallest attainable p-value is $6.1 \cdot 10^{-5}$.

E. Implementation and Reproducibility

Everything is implemented in Python 3.11 with PyTorch 2.x (SNN layers and operation counters written from scratch), NumPy/SciPy (HiGHS LP), scikit-learn and pvlib. All experiments run on a 4-core CPU without GPU; wall-clock times are reported. The operator console is a React/TypeScript application whose data bundle is exported directly from the experiment logs. Code, cleaned data subsets, checksums, logs and the console are released in the project repository (see Data Availability).

VI. Results

A. E1 – Neuromorphic Occupancy Sensing

Table IV summarizes E1. The spiking occupancy sensor reached 97.79% accuracy (F1 0.971) on test set 1 and 98.33% (F1 0.961) on test set 2, the recording with the door mostly open. That is 4.9 and 4.3 percentage points above an MLP of identical width, with a five-seed standard deviation below 0.3 points. Logistic regression remained the strongest model on this small, nearly linearly separable benchmark, consistent with the original study [52]. The SNN closes most of the gap between flexible neural models and the linear reference while using 9.0 nJ per inference against 20.3 nJ for the MLP; its mean firing rate is 0.319.

TABLE IV

E1: occupancy detection on the UCI benchmark (mean ± s.d.; neural models 5 seeds, ensembles 3 seeds)

Model	Acc. test 1 (%)	F1 test 1	Acc. test 2 (%)	F1 test 2	MCC test 2	Energy (nJ)
LogReg	97.86 ± 0.00	0.971	99.19 ± 0.00	0.981	0.976	–
RandomForest	95.22 ± 0.16	0.934	97.63 ± 0.24	0.945	0.931	–
GBDT	95.16 ± 0.05	0.933	94.94 ± 0.00	0.869	0.843	–
MLP	92.86 ± 0.47	0.897	94.01 ± 0.59	0.845	0.813	20.3
Spiking MLP (ours)	97.79 ± 0.08	0.971	98.33 ± 0.28	0.961	0.951	9.0

B. E2 – Spiking Day-Ahead Forecasting

Gym occupancy (Task A). Day-ahead member demand is noisy: semester breaks and holidays fall inside the test period. NeuroCast reached CV(RMSE) 49.3%, on par with the best deep baseline, PatchTST (49.0%; Diebold–Mariano $p = 0.785$), and significantly more accurate than MLP, LSTM, TCN, Transformer, NeuroCast-noAnchor (DM $p < 0.05$). For occupancy, the learnable seasonal anchor is essential: without it the spiking model degrades to 58.5%. Two simple references, seasonal naive (44.6%) and Ridge (44.8%), remained best on this holiday-heavy period, a known pattern for short, regime-shifting series [55]. NeuroCast needs 1,170 nJ per 24-h forecast versus 4,190 nJ for the LSTM and 10,286 nJ for the Transformer.

Unseen buildings (Task B). Under the leave-sites-out protocol, PatchTST was the most accurate model (CV(RMSE) 29.9%, NMBE 1.9%), followed by the Transformer (32.1%) and the recurrent networks (≈33–34%). Here the seasonal anchor *hurts*: weekly schedule changes between buildings make yesterday's profile a poor reference, and the un-anchored spiking model (37.7%) beats the anchored one (43.1%). The best spiking forecaster costs 1,498 nJ per 24-h forecast: 1.4× less than PatchTST, 2.7× less than the LSTM and 6.8× less than the Transformer, at a 7.8-point accuracy cost. The linear DLinear model is both cheaper (10.6 nJ) and slightly less accurate (38.8%), so spiking forecasting pays off only against non-linear deep models. Federated NeuroCast, trained without pooling any site's meter data, reached 50.9%; the federation over ten heterogeneous sites costs accuracy that personalization would have to recover. Only PatchTST approaches the ASHRAE Guideline 14 hourly criterion (CV(RMSE) ≤ 30%, |NMBE| ≤ 10%) [111] on unseen sites; all models satisfy the NMBE criterion.

TABLE V

E2: day-ahead (24 h) forecasting, mean of 3 seeds. DM p: Diebold–Mariano test vs NeuroCast. Task A MAE in members, Task B MAE in kW

Task · model	MAE	CV(RMSE) %	NMBE %	Energy (nJ)	DM p
A · SeasonalNaive	4.6	44.6	0.1	–	0.019
A · Ridge	5.6	44.8	10.5	53	0.008
A · MLP	6.4	52.9	9.5	895	0.019
A · LSTM	6.8	54.2	12.0	4,190	0.004
A · GRU	6.6	51.9	7.9	3,144	0.123
A · TCN	6.8	53.4	6.7	4,902	0.022
A · Transformer	7.4	57.7	13.6	10,286	< 0.001
A · DLinear	7.1	52.4	8.2	11	0.971
A · PatchTST	6.5	49.0	16.7	2,050	0.785
A · NeuroCast (SNN)	6.3	49.3	15.9	1,170	–
A · NeuroCast-noAnchor	7.4	58.5	9.6	1,247	< 0.001
B · SeasonalNaive	21.6	64.1	1.2	–	0.068
B · Ridge	29.6	44.9	4.2	37	0.122
B · MLP	24.6	33.7	2.8	725	< 0.001
B · LSTM	26.4	34.1	4.1	4,020	0.002
B · GRU	25.7	33.4	4.0	3,017	< 0.001
B · TCN	26.2	36.0	3.7	4,806	0.002
B · Transformer	23.7	32.1	2.0	10,254	< 0.001
B · DLinear	23.1	38.8	3.2	11	0.002
B · PatchTST	21.4	29.9	1.9	2,050	< 0.001
B · NeuroCast (SNN)	28.4	43.1	5.8	983	–

Task · model	MAE	CV(RMSE) %	NMBE %	Energy (nJ)	DM p
B · NeuroCast-noAnchor	30.0	37.7	5.5	1,498	0.014
B · Fed-NeuroCast (SNN)	30.1	50.9	5.7	–	–

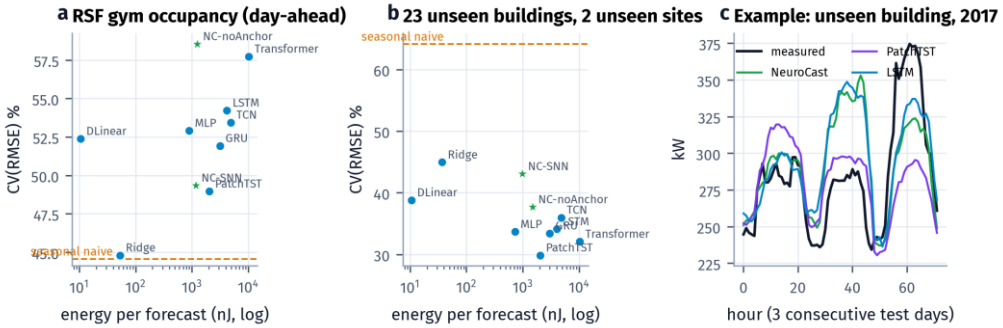


Fig. 6. E2 forecasting: accuracy versus inference energy for (a) RSF gym occupancy and (b) 23 buildings at two sites never seen in training; (c) example day-ahead forecasts for an unseen building in 2017.

C. E3 – Spiking Intrusion Detection

Across the 7 TON_IoT devices, the spiking IDS reached a mean binary macro-F1 of 0.948, versus 0.947 for the MLP and 0.959 for the random forest (Table VI). On the two devices that matter for the twin, the thermostat and the weather station, it detected 96.9% and 100.0% of attack samples at false-positive rates of 6.63% and 0.05%. These operating points parameterize the SGT layer in E4. Multi-class attribution is harder for all models on single-variable devices (motion light, thermostat), because several attack types leave identical traces in one sensor value. Ensembles retain a small edge in multi-class F1, while the SNN uses 24.0 nJ per decision.

TABLE VI

E3: TON_IoT telemetry intrusion detection (macro-F1, mean of 3 splits)

Device	Records	RF bin.	MLP bin.	SNN bin.	RF multi	MLP multi	SNN multi	SNN TPR / FPR
Thermostat	52,774	0.954	0.943	0.941	0.706	0.650	0.634	0.969 / 0.066
Weather	59,260	1.000	1.000	1.000	0.998	0.993	0.992	1.000 / 0.001
Motion Light	59,488	0.874	0.860	0.860	0.377	0.379	0.372	0.997 / 0.236
Fridge	59,944	0.969	0.970	0.970	0.636	0.629	0.608	1.000 / 0.050
Garage Door	59,587	0.979	0.978	0.978	0.724	0.667	0.668	1.000 / 0.036
Modbus	51,106	0.942	0.876	0.887	0.734	0.614	0.616	0.864 / 0.082
GPS Tracker	58,960	1.000	0.999	0.999	0.993	0.945	0.936	1.000 / 0.001

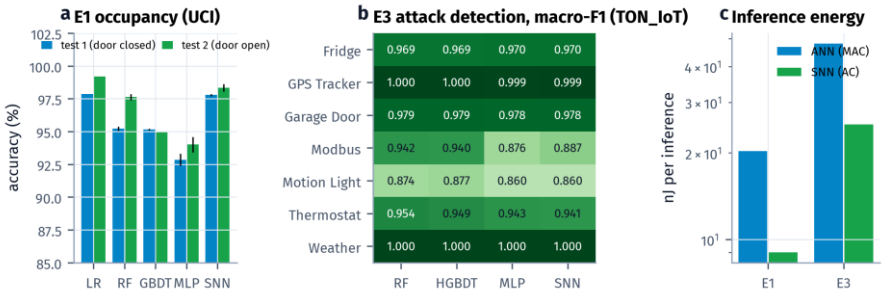


Fig. 7. Edge perception: (a) E1 occupancy accuracy on the two UCI test recordings; (b) E3 binary macro-F1 per TON_IoT device and model; (c) inference energy of ANN and SNN detectors (log scale).

D. E4 – Resilient Resource Allocation

TABLE VII

E4: weekly operating objective (negative cost-equivalent in normalized units; higher is better), mean over 15 held-out weeks and seeds. Bold green: best deployable controller per scenario. Oracles use perfect information and are not deployable. LP-Oracle: actuator-consistent weekly LP replayed open loop (corrected in the revision, Section VI-D).

Controller	Nominal	Outage	Chiller fault	Cyber FDI	Surge	Compound	Club B nom.	Club B comp.	Seeds
BAS	-154.0	-455.3	-155.2	-154.0	-174.5	-477.4	-147.1	-446.1	1
RBC	-132.7	-356.5	-132.9	-151.8	-153.5	-399.0	-133.7	-379.8	1
RBC+SGT	-135.7	-359.3	-136.0	-149.4	-156.6	-396.9	-136.0	-379.5	1
MPC-24h	**-122.1**	**-208.0**	**-122.9**	-160.4	**-143.0**	**-273.8**	-132.0	-266.6	1
MPC-24h+SGT	-135.3	-221.1	-136.1	-167.1	-158.9	-283.8	-141.6	-273.6	1
PPO-central-ANN	-127.1	-288.8	-127.5	**-141.2**	-147.8	-324.2	-133.7	-329.1	3

Controller	Nominal	Outage	Chiller fault	Cyber FDI	Surge	Compound	Club B nom.	Club B comp.	Seeds
IPPO-ANN	-128.8	-296.7	-129.1	-149.1	-149.4	-342.5	-131.8	-338.4	3
MAPPO-ANN	-129.0	-300.5	-129.4	-148.5	-149.7	-344.4	-131.9	-336.6	3
MAPPO-SNN	-130.8	-337.9	-131.1	-159.6	-151.6	-392.2	-133.5	-379.7	3
MATD3-ANN	-134.3	-329.6	-134.8	-151.3	-155.0	-372.2	-137.7	-366.6	3
Distill-ANN	-125.8	-250.0	-125.9	-150.1	-148.6	-303.7	-133.1	-298.6	3
NeuroTwin (ours)	-126.5	-253.7	-126.9	-152.1	-149.6	-309.5	-136.1	-305.2	3
Teacher-MPC-oracle	-112.9	-193.5	-112.9	-112.7	-133.7	-213.5	-119.9	-217.6	1
LP-Oracle	-106.6	-187.7	-106.7	-106.6	-127.4	-208.0	-113.9	-213.9	1

Table VII and Fig. 8 report the main results. Under **nominal** operation NeuroTwin scored -126.5 against -132.7 for the incumbent RBC (mean per-week change +3.9%, Wilcoxon $p_{Holm} = 0.108$, better in 12/15 weeks; not significant) and -154.0 for the recorded as-operated baseline. The realistic 24-h MPC with forecasts scored -122.1. The best model-free MARL method, PPO-central-ANN, reached -127.1.

Under **disruptions** the differences grow. During grid outages NeuroTwin scored -253.7 versus -356.5 for RBC (mean per-week gain +22.5%, $p_{Holm} = 0.009$, better in 13/15 weeks): its storage and allocation agents inherit the teacher's outage behaviour of holding charge when the notice appears and shedding discretionary space. Under compound disruptions it scored -309.5 versus -399.0 (+17.8% per week; breakdown by disruption type in Table XI, $p_{Holm} = 0.075$) and versus -324.2 for the best model-free MARL method ($p_{Holm} = 0.473$). Under **false-data injection**, the forecast-driven MPC collapsed from -122.1 to -160.4 because it trusts the falsified state, while NeuroTwin scored -152.1 (vs MPC $p_{Holm} = 0.406$).

The strengthened model-free MARL baselines improved clearly on RBC under outages (best: PPO-central-ANN, -288.8) and compound disruptions (PPO-central-ANN, -324.2). Under false-data injection, PPO-central-ANN, whose single policy observes the global state, was the best deployable controller (-141.2); decentralized policies cannot cross-check a falsified zone sensor against the others. Privileged distillation produced the best decentralized learning controllers under outages and compound disruptions: NeuroTwin -253.7 and -309.5 versus -288.8 and -324.2 for the best MARL method (compound: +2.2% per week, $p_{Holm} = 0.473$). It needed about 19 min of single-core training versus 2 min for MAPPO. Seeds agreed closely (standard deviation of the compound objective across seeds 3.2). The realistic forecast-driven MPC remains stronger whenever its inputs can be trusted (compound: -273.8), at the price of a central solver, live forecasts and roughly eight orders of magnitude more computation per decision (Section VI-F). The weekly LP oracle (-208.0) and the privileged teacher (-213.5) bound what perfect foresight could achieve.

Oracle consistency. In the submitted version the weekly LP oracle scored below the receding-horizon teacher (-234.6 vs -213.5 under compound disruptions), although it has more information. The cause was a modelling inconsistency in the LP, not in the twin: its actuator model omitted the lower limit of the cooling action ($Q_{del} \geq 0.4\ Q_{ref}$, i.e. $a \geq -1$), so 22% of the planned zone-hours requested less cooling than the plant can deliver. Because the weekly plan is replayed open loop, these hours were clipped, the zones were over-cooled and comfort violations accumulated (214 K·h per compound week). The teacher's 24-h LP has the same simplification but re-plans every hour from the true state, so feedback absorbs the error. Imposing the twin's service projection on the LP did not change its score, which isolates the actuator limit as the cause. With the limit added, the weekly LP scores -106.6 (nominal) and -208.0 (compound) and lies above the teacher in all eight scenarios (Table VII), so it is a consistent perfect-information upper bound. No other result is affected.

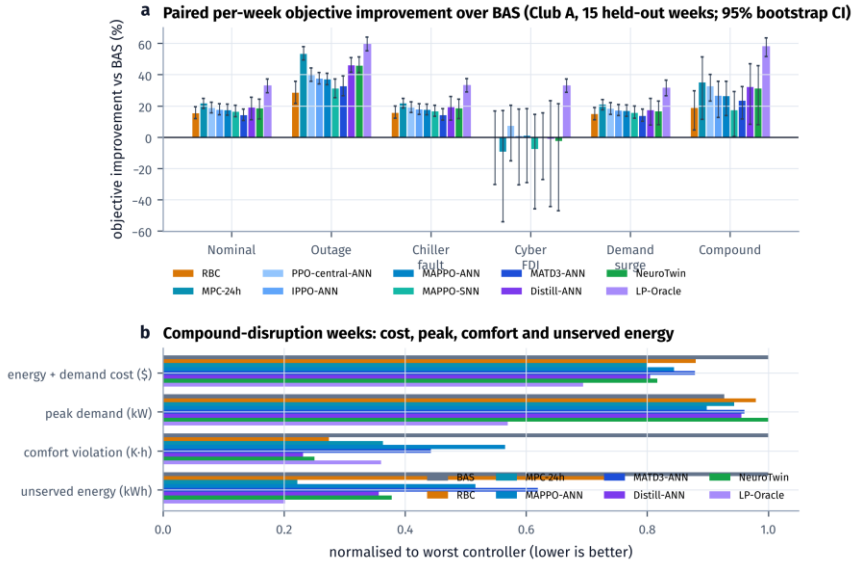


Fig. 8. E4: (a) improvement of the weekly objective over the as-operated BAS for every controller and disruption family (Club A, 15 held-out weeks; 95% bootstrap CI); (b) breakdown under compound disruptions, normalized to the worst controller.

TABLE VIII

E4 compound-disruption weeks (Club A): physical and economic breakdown per week. Comfort values are model-based and comparable only between controllers (Sections III-A, VII-B)

Controller	Cost ($)	Grid energy (MWh)	Peak (kW)	Comfort (K·h)	Unserved members	Unserved energy (kWh)	Service level (%)	Resilience index	Recovery (h)
BAS	42,379	312.2	3,122	239.6	6,847	6,403	98.23	0.891	2.7
RBC	37,305	261.8	3,296	65.7	7,485	4,795	98.14	0.962	4.4
MPC-24h	33,866	243.4	3,176	87.1	7,475	1,421	98.20	0.938	4.1
MAPPO-ANN	35,793	253.8	3,025	135.5	7,152	3,304	98.21	0.925	5.1
Distill-ANN	34,097	242.8	3,216	55.5	7,654	2,282	98.18	0.961	2.0
NeuroTwin	34,598	242.3	3,366	60.0	7,573	2,422	98.18	0.956	2.3
LP-Oracle	29,406	239.4	1,918	86.3	4,100	1,294	99.06	0.933	1.5

E. Resilience and Cyber-Resilience

Fig. 9 traces the service performance Q(t) during one held-out summer week. Spike-Gated Trust improved NeuroTwin's objective by 4.5 under FDI and by 1.4 under compound disruptions, but cost 4.3 in nominal weeks because of IDS false alarms (FPR ≈ 6.6% on the thermostat), when trusted readings are needlessly replaced by imperfect twin estimates. For the rule-based controller the corresponding numbers were +2.4 under FDI and −3.0 nominally. For the forecast-driven MPC, SGT did not help (−6.8 under FDI), because the MPC integrates the substituted estimate over its whole horizon. Trust gating is therefore a net benefit only when attacks are frequent or the IDS false-positive rate is low; lowering the thermostat detector's FPR is the most direct lever. Resilience indices and recovery times are reported in Table VIII.

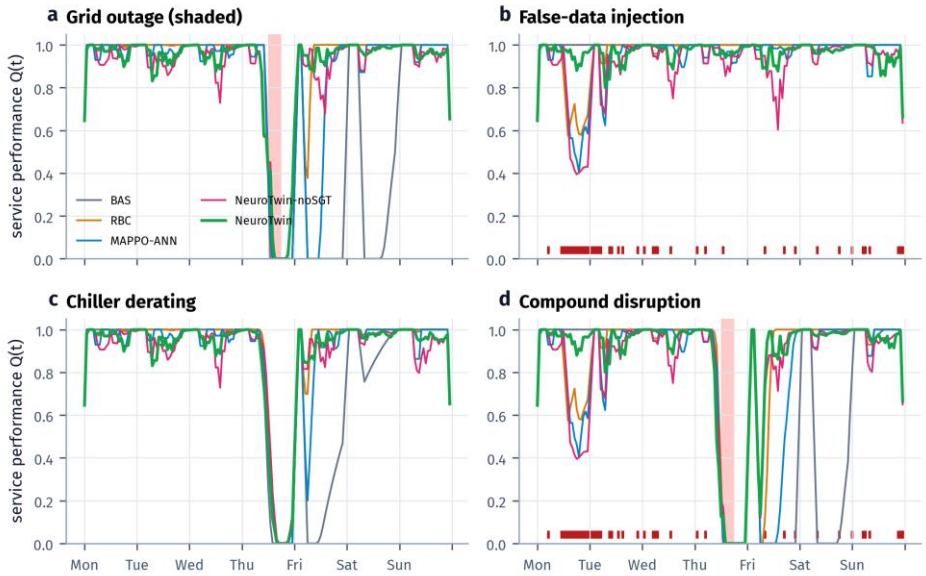


Fig. 9. Service performance Q(t) (3-h moving average) during one held-out summer week under (a) a grid outage (shaded), (b) false-data injection (red ticks: SGT quarantine events), (c) chiller derating and (d) compound disruption.

F. Ablations, Sensitivity Analyses and Neuromorphic Efficiency

TABLE IX

Ablation study and neuromorphic cost (Club A objective; energy per hourly decision for all seven agents)

Variant	Nominal	Outage	Cyber FDI	Compound	Energy (nJ)	AC ops	MAC ops	Firing rate
NT (full)	-126.5	-253.7	-152.1	-309.5	386.1	420,194	1,720	0.363
NT − SGT	-122.2	-249.5	-156.6	-310.8	382.5	416,183	1,720	0.361
NT − NeuroCast	-128.0	-256.2	-153.0	-311.9	374.6	407,485	1,720	0.349
NT − population code	-126.1	-254.6	-150.7	-312.1	247.0	133,806	27,520	0.242
NT, T = 1	-137.0	-340.1	-152.3	-380.8	9.1	1,290	1,720	–
NT, T = 3	-128.2	-261.2	-155.0	-319.7	184.3	196,021	1,720	0.280
NT, T = 8	-123.4	-262.9	-147.9	-316.3	681.8	748,733	1,720	0.394
Distill-ANN	-125.8	-250.0	-150.1	-303.7	660.6	0	143,616	–

The spiking actors fired at 0.363 spikes per neuron per step, giving 386.1 nJ per hourly decision for all seven agents, versus 660.6 nJ for architecturally matched ANN actors trained by the same distillation (1.7×). The ANN student scored -125.8 (nominal) and -303.7 (compound) against -126.5 and -309.5 for the spiking student, so spiking costs 5.8 objective units under compound disruptions. The only MAC-equivalents (1,720 per decision) are the Gaussian receptive fields; every synaptic event is an accumulate. Two findings temper the case for population coding: the directly encoded spiking actor reached similar quality

(-312.1 compound) at 247.0 nJ, because its 8× smaller input layer outweighs the cost of its analog first layer. With T = 1, the hidden layers never reached threshold and the actor degenerated to a constant residual (-380.8). Energy grows linearly with the time-step budget (≈99 nJ per additional step for all seven agents, fitted over T = 3–8), so the energy advantage over the matched ANN actors (660.6 nJ) disappears at T ≈ 7.8: at T = 8 the spiking actors need 681.8 nJ, 3% more than the ANN, without being more accurate (compound -316.3 vs -303.7). T = 3 needs 3.6× less energy than the ANN at a 16.0-point compound loss, and T = 5 1.7× less at 5.8 points (Fig. 10b). The design rule for these actors is therefore T ≤ 5; beyond T ≈ 8 the spiking actor offers no energy benefit over its ANN counterpart. Compared with a receding-horizon MPC that solves an LP of ≈5 000 variables every hour (≈13 ms of CPU time, on the order of 0.1–0.2 J), the spiking agents are roughly eight orders of magnitude cheaper per decision, and they run without a central solver or a forecast service.

TABLE X

Sensitivity of Spike-Gated Trust to the mis-specification e of the twin estimator (C and UA scaled by 1 + e); NeuroTwin students trained with e = +20%, mean of 3 seeds, Club A, 15 held-out weeks

Estimator error e	NT nominal	NT cyber FDI	NT compound	Compound gain vs RBC (weeks better)	RBC+SGT cyber FDI
-40%	-126.4	-151.7	-309.3	+17.9% (12/15)	-149.0
-20%	-126.7	-152.2	-308.6	+17.9% (12/15)	-149.3
-10%	-126.6	-152.6	-310.0	+17.9% (12/15)	-149.3
+10%	-126.5	-152.3	-308.8	+18.0% (12/15)	-149.5
+20% (paper)	-126.5	-152.1	-309.5	+17.8% (12/15)	-149.4
+40%	-127.0	-153.0	-308.2	+18.4% (12/15)	-149.1
NT − SGT	-122.2	-156.6	-310.8	–	–
RBC (no SGT)	-132.7	-151.8	-399.0	–	-151.8

Sensitivity of Spike-Gated Trust. The twin estimator that replaces quarantined zone temperatures uses deliberately mis-specified parameters (+20% on C and UA). Table X varies this error from −40% to +40% at test time, with the students trained at +20%. Under false-data injection NeuroTwin's objective stayed between -153.0 and -151.7 (worst at e = +40%), against -156.6 without SGT, and under compound disruptions between -310.0 and -308.2; the compound gain over RBC remained +17.8% to +18.4%. Performance therefore degrades gracefully rather than collapsing: across the whole range, no objective moved by more than 1.8 units, and the benefit of SGT over the same students without it remained +3.6 to +4.9 under false-data injection. The reason is that quarantine episodes are short (6–24 h) and the estimator restarts from the last trusted reading, so a gain error of up to ±40% accumulates only a small temperature error. Because C and UA are scaled jointly, the estimator's time constant is preserved; a mis-specified time constant and adaptive attacks on the estimator itself were not tested. The nominal cost of SGT (false alarms, Section VI-E) is likewise unaffected by e.

TABLE XI

Which disruptions drive the compound-disruption gain? Single-family scenarios (from Table VII) and compound scenarios with one disruption type removed; gain = mean per-week improvement of NeuroTwin over RBC

Scenario	NeuroTwin	RBC	MPC-24h	Gain vs RBC	Weeks better	Wilcoxon p
Nominal	-126.5	-132.7	-122.1	+3.9%	12/15	0.108 (Holm)
Grid outage only	-253.7	-356.5	-208.0	+22.5%	13/15	0.009 (Holm)
Chiller fault only	-126.9	-132.9	-122.9	+3.6%	11/15	0.124 (Holm)
Cyber FDI only	-152.1	-151.8	-160.4	+2.8%	11/15	0.146 (Holm)
Demand surge only	-149.6	-153.5	-143.0	+2.5%	12/15	0.108 (Holm)
Compound (all four)	**-309.5**	**-399.0**	**-273.8**	**+17.8%**	**12/15**	**0.075 (Holm)**
Compound without outage	-182.3	-175.8	-189.2	-1.1%	10/15	0.359
Compound without chiller fault	-309.3	-398.5	-272.9	+17.8%	12/15	0.015
Compound without cyber FDI	-276.6	-378.1	-229.7	+21.7%	13/15	0.001
Compound without surge	-279.8	-375.2	-245.8	+20.5%	13/15	0.003

What drives the compound gain. Table XI shows that the compound-disruption gain is concentrated in one disruption type. Of the single-family scenarios only the grid outage produced a large and significant gain (+22.5%), whereas chiller faults, cyber FDI and demand surges gave +2.5% to +3.6%, similar to nominal weeks. Removing the outage from the compound scenario removed the gain entirely (-1.1%, better in 10/15 weeks, p = 0.36), while removing the chiller fault (+17.8%), the cyber attack (+21.7%) or the surge (+20.5%) left it unchanged or larger. The compound benefit of NeuroTwin is therefore an outage-preparation benefit (holding battery charge after the 2-h notice and shedding discretionary space) that persists when other disruptions coincide, not a broad improvement across disruption types; the cyber attack, if anything, dilutes it.

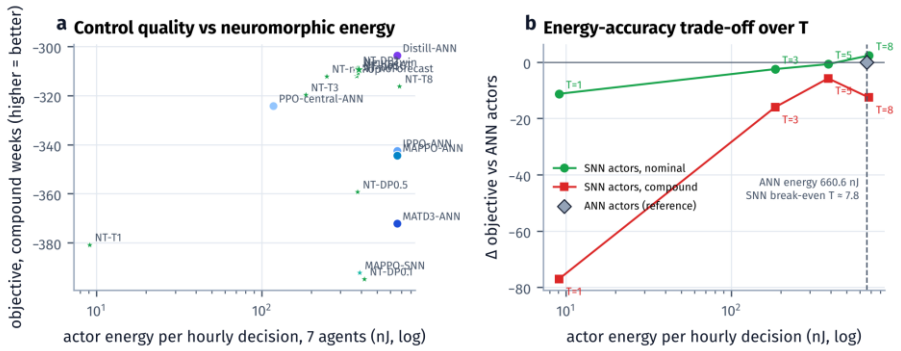


Fig. 10. (a) Control quality under compound disruptions versus actor energy per decision (log scale); stars: spiking actors. (b) Energy–accuracy trade-off over the number of SNN simulation steps T, as the change of the objective relative to the matched ANN actors; the dashed line marks the ANN energy (break-even T ≈ 7.8).

G. Privacy and Generalization

With local differential privacy on member counts, NeuroTwin's nominal objective moved from -126.5 (no DP) to -126.2 at ε = 2, -130.4 at ε = 1, -130.8 at ε = 0.5, -149.2 at ε = 0.1. Under compound disruptions the values were -308.7 (ε = 2), -311.3 (ε = 1), -359.3 (ε = 0.5), -394.8 (ε = 0.1) against -309.5. A budget of ε = 2 is therefore free and ε = 1 costs a few percent in nominal weeks. At ε ≤ 0.5 the Laplace noise (scale ≥ 2 members per count, amplified by the club scaling) starts to mislead space allocation and outage preparation, and at ε = 0.1 the controller falls below the incumbent (Fig. 11a). Because the agents also rely on meters, forecasts and the safety projection, moderate formal privacy is affordable.

Zero-shot transfer. Without retraining, the policies were deployed on Club B in Austin, a different climate (2A, humid) with a different building mix. NeuroTwin scored -136.1 (nominal) and -305.2 (compound), versus -133.7 and -379.8 for RBC and -147.1 and -446.1 for the as-operated baseline (Table VII, Fig. 11b). Because all observations are normalized by each zone's own meter statistics, the agents run without site-specific tuning. Their benefit carries over where it matters most, under disruptions, while nominal operation at the new site is on par with or slightly below the local incumbent and would profit from brief fine-tuning.

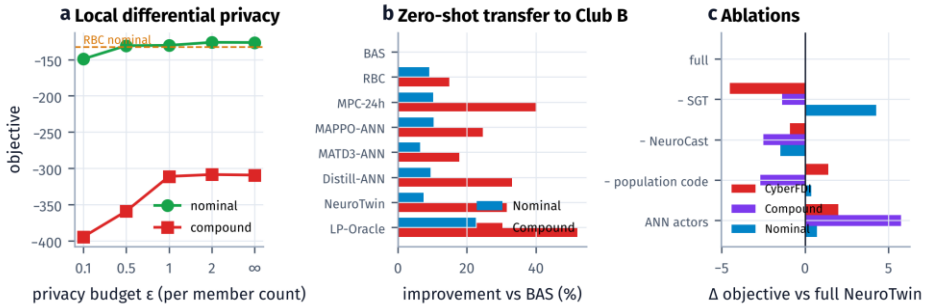


Fig. 11. (a) Privacy–utility trade-off under ε-local differential privacy on member counts; (b) zero-shot transfer to Club B (improvement over BAS); (c) ablations, change of the objective relative to the full NeuroTwin.

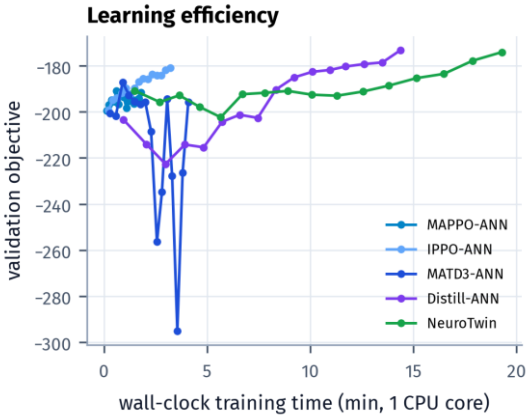


Fig. 12. Learning efficiency: validation objective versus single-core wall-clock time for model-free MARL and twin-in-the-loop distillation.

H. NeuroTwin Console

The operator console (Fig. 13) is a React/TypeScript application designed with the *ui-ux-pro-max* design-intelligence toolkit (real-time operations pattern, dense dashboard spacing, Fira Sans/Fira Code, WCAG-contrast dark and light themes). It replays held-out weeks for every controller and scenario: a plan view of the club with zone thermal deficits, opened space, member load and per-sensor trust shields; grid-import, state-of-charge and service-performance charts; a resilience laboratory comparing controllers; the spike raster of a live decision; and an evidence page with every perception result and dataset checksum. All numbers come from the experiment logs.

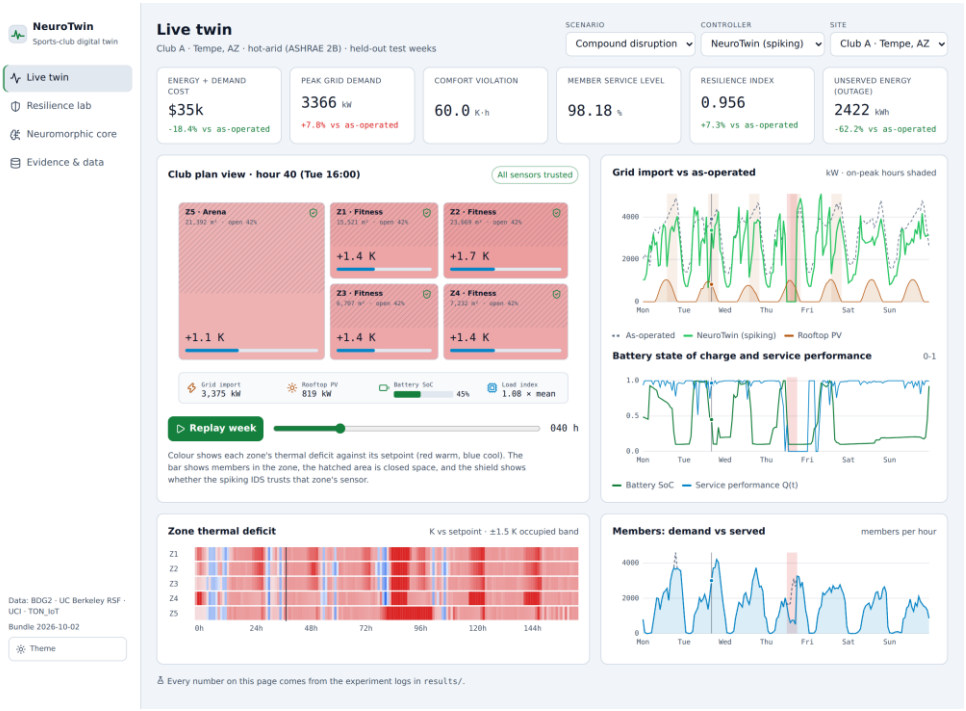


Fig. 13. NeuroTwin Console: replay of a held-out week under compound disruption (plan view with trust shields, grid import vs as-operated, SoC and service performance, zone deficit heat map).

VII. Discussion

A. Implications

Three lessons generalize beyond sports clubs. First, **the digital twin is most valuable as a teacher, not only as a simulator**. With an explicit physics-and-tariff model, a privileged MPC can label any state, and DAgger turns those labels into decentralized policies that need no solver, forecasts or global state at run time. Model-free MARL, given the same twin, observations and twin-enabled variance reduction, also improved on the incumbent, but needed more computation and was weaker under outages and compound disruptions; its centralized variant was the most robust to sensor falsification, which points to hybrid designs. Second, **neuromorphic actors make the twin's decisions cheap enough to embed** in BMS controllers or sensor gateways: about 386 nJ for all seven agents per hourly decision, 1.7× below matched ANN actors and many orders of magnitude below an online MPC. The margin over ANNs is modest, however: it shrinks linearly with the time-step budget and vanishes at $T \approx 8$ (Section VI-F), and direct input coding is cheaper than population coding at equal quality. Third, **security and privacy must be designed into the observation pipeline**. Trust gating pays off only when its false-alarm cost is lower than the attack cost it prevents, and local DP on member counts is affordable down to $\varepsilon \approx 1$ because the policy is anchored by meters, forecasts and safety constraints.

B. Threats to Validity

- *Construct validity.* The twin uses measured loads but a reduced-order thermal model and a supervisory action that modulates delivered cooling relative to the measured demand. Absolute comfort numbers are therefore model-based; relative comparisons share the same model.
- *Data combination.* Member demand comes from the UC Berkeley RSF, while loads come from BDG2 buildings in Arizona and Texas; all traces are real and time-aligned over 2016–17, but they are not from one facility. Results should be read as a realistic composite, not as a case study of a specific club.
- *Operational validity.* The twin is run offline on replayed data with simulated disruptions (Section V-C). The results therefore support statements about relative controller performance, not about the savings of a deployed system; a pilot with a live BMS is needed to measure latency, communication faults and actuator dynamics.
- *Energy estimates.* Neuromorphic energy is an operation-count estimate (45 nm constants) that ignores memory traffic, I/O and static power. Hardware measurements on Loihi 2 or similar platforms are future work.
- *Attack model.* FDI is modelled as bias and scaling attacks with detector rates taken from TON_IoT; adaptive adversaries who target the IDS itself are out of scope.
- *Statistical power.* With 15 paired weeks the design detects $d_z \geq 0.85$ ($\alpha = 0.05$) or ≥ 1.15 (strictest Holm step) with 80% power. With the observed standard deviation of per-week gains over RBC (11–22 percentage points), the minimum detectable mean gain is about 9–18% at $\alpha = 0.05$ and 12–25% after Holm correction. The outage (+22.5%) and compound (+17.8%) gains are at or above this threshold; nominal-week differences (≈4%) and differences between learning methods are below it (NeuroTwin vs PPO-central-ANN under compound disruptions: +2.2%, detectable 21%; vs the distilled ANN: -1.4%, detectable 4.2%) and must not be interpreted as evidence of equivalence or superiority.

C. Limitations and Future Work

Future work will (i) deploy the spiking actors on neuromorphic hardware with on-chip learning for continual adaptation; (ii) extend the twin with pool hydronics once trustworthy hot-water metering is available; (iii) add carbon-intensity signals and grid services such as demand response; (iv) study adaptive adversaries and certified robustness of SGT; and (v) federate the distillation across clubs so that each site contributes teacher labels without sharing raw data.

VIII. Conclusion

This paper presented NeuroTwin, a spiking-neural-network–driven neuromorphic digital twin for real-time resilient resource allocation in smart sports clubs, built and evaluated exclusively on measured, public data. Spiking perception matched or approached conventional networks on occupancy sensing and intrusion detection at about half their inference energy, and traded accuracy for energy in forecasting. Seven decentralized population-coded spiking agents, trained by twin-in-the-loop privileged distillation, improved on the incumbent rule-based controller under grid outages (+22.5% per week, $p_{Holm} = 0.009$) and, less conclusively, under compound disruptions (+17.8%, $p_{Holm} = 0.075$). They outperformed the best of five strengthened model-free MARL baselines under outages (+10.1%, $p_{Holm} = 0.027$), carried their disruption benefit to a second club without retraining, and tolerated $\varepsilon = 2$ local differential privacy at no measurable cost, all at 386.1 nJ per hourly decision. Forecast-driven MPC remains the stronger controller when its inputs are trustworthy, and centralized PPO was the most robust to false-data injection. Combining the twin teacher with centralized cross-checking of sensors is therefore the most promising next step. The twin was operated offline on replayed measured data with simulated disruptions, so these findings are relative comparisons under identical conditions; a pilot deployment with a live BMS is the necessary next step. The code, cleaned data subsets, logs and operator console are released to support replication.

Appendix: Hyperparameters

TABLE XII
Hyperparameters of the learning components

Component	Setting
PLIF neurons	$\theta = 1$, soft reset, learnable β (init 0.8–0.9), fast-sigmoid surrogate $k = 5$
Spiking actor	8 Gaussian RF / input dim → IF encoder; LIF 128–128; 8 output neurons / action; $T = 5$; residual on RBC
NeuroCast	window 48 h (twin: 24 h), horizon 24 h (twin: 6 h); recurrent LIF 128 + LIF 128; 2 sub-steps/h; delta threshold 0.15; Huber loss; AdamW $2 \cdot 10^{-3}$, one-cycle; 8–20 epochs
Spiking IDS / occupancy SNN	LIF 128–64 (IDS) or 64–64 (occupancy), $T = 8$, BatchNorm, AdamW $2 \cdot 10^{-3}$, class-balanced CE
DAgger distillation	16 rounds × 16 parallel weeks; $\beta_r = 0.7^r$; ≤ 400 minibatches (1 024) per round and type; Adam 10^{-3}; pinball $\tau = 0.8$ (allocation), 0.6 (HVAC)
Teacher MPC	24-h horizon LP, HiGHS; terminal SoC value $0.08 \cdot B_E$; privileged true state and exogenous trajectory
PPO / IPPO / MAPPO	150 updates × 16 weeks; $\gamma = 0.9$, $\lambda = 0.9$; clip 0.2; 5 epochs, 4 minibatches; actor $3 \cdot 10^{-4}$, critic 10^{-3}; shadow reward; round-robin exploration
MATD3	75 rounds × 96 gradient steps; batch 512; $\gamma = 0.95$; $\tau = 0.005$; policy delay 2; noise 0.1; BC regularizer 0.1; 8 critic warm-up rounds
Validation / selection	8 fixed training-period weeks with fixed disruption draws; best checkpoint by validation objective

Declarations

Funding: Competing interests: the authors declare none. *Ethics:* only public, de-identified aggregate data were used; no human-subject data were collected.

References

[1] M. Elnour et al., "Performance and energy optimization of building automation and management systems: Towards smart sustainable carbon-neutral sports facilities," Renew. Sustain. Energy Rev., vol. 162, 112401, 2022.
[2] E. Trianti-Stourna et al., "Energy conservation strategies for sports centers: Part B. Swimming pools," Energy Build., vol. 27, no. 2, pp. 123-135, 1998.
[3] P. Artuso and A. Santiangeli, "Energy solutions for sports facilities," Int. J. Hydrogen Energy, vol. 33, no. 12, pp. 3182-3187, 2008.
[4] A. Al-Fuqaha et al., "Internet of things: A survey on enabling technologies, protocols, and applications," IEEE Commun. Surveys Tuts., vol. 17, no. 4, pp. 2347-2376, 2015, doi: 10.1109/COMST.2015.2444095.
[5] E. A. Lee, "Cyber physical systems: Design challenges," in Proc. IEEE ISORC, 2008, pp. 363-369, doi: 10.1109/ISORC.2008.25.
[6] M. Grieves and J. Vickers, "Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems," in Transdisciplinary Perspectives on Complex Systems. Cham: Springer, 2017, pp. 85-113, doi: 10.1007/978-3-319-38756-7_4.
[7] F. Tao, H. Zhang, A. Liu, and A. Y. C. Nee, "Digital twin in industry: State-of-the-art," IEEE Trans. Ind. Informat., vol. 15, no. 4, pp. 2405-2415, 2019, doi: 10.1109/TII.2018.2873186.
[8] A. Fuller, Z. Fan, C. Day, and C. Barlow, "Digital twin: Enabling technologies, challenges and open research," IEEE Access, vol. 8, pp. 108952-108971, 2020, doi: 10.1109/ACCESS.2020.2998358.
[9] A. Rasheed, O. San, and T. Kvamsdal, "Digital twin: Values, challenges and enablers from a modeling perspective," IEEE Access, vol. 8, pp. 21980-22012, 2020, doi: 10.1109/ACCESS.2020.2970143.
[10] ANSI/ASHRAE Standard 55-2020, Thermal Environmental Conditions for Human Occupancy. Atlanta, GA: ASHRAE, 2020.
[11] M. Bruneau et al., "A framework to quantitatively assess and enhance the seismic resilience of communities," Earthquake Spectra, vol. 19, no. 4, pp. 733-752, 2003, doi: 10.1193/1.1623497.
[12] S. Hosseini, K. Barker, and J. E. Ramirez-Marquez, "A review of definitions and measures of system resilience," Reliab. Eng. Syst. Saf., vol. 145, pp. 47-61, 2016, doi: 10.1016/j.ress.2015.08.006.
[13] M. Panteli and P. Mancarella, "The grid: Stronger, bigger, smarter? Presenting a conceptual framework of power system resilience," IEEE Power Energy Mag., vol. 13, no. 3, pp. 58-66, 2015, doi: 10.1109/MPE.2015.2397334.
[14] Y. Liu, P. Ning, and M. K. Reiter, "False data injection attacks against state estimation in electric power grids," ACM Trans. Inf. Syst. Secur., vol. 14, no. 1, 13, 2011, doi: 10.1145/1952982.1952995.
[15] A. Humayed, J. Lin, F. Li, and B. Luo, "Cyber-physical systems security - A survey," IEEE Internet Things J., vol. 4, no. 6, pp. 1802-1831, 2017, doi: 10.1109/JIOT.2017.2703172.

[16] C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," Found. Trends Theor. Comput. Sci., vol. 9, no. 3-4, pp. 211-407, 2014, doi: 10.1561/0400000042.
[17] R. Jia, R. Dong, S. S. Sastry, and C. J. Spanos, "Privacy-enhanced architecture for occupancy-based HVAC control," in Proc. ACM/IEEE ICCPS, 2017, pp. 177-186.
[18] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, "Edge computing: Vision and challenges," IEEE Internet Things J., vol. 3, no. 5, pp. 637-646, 2016, doi: 10.1109/JIOT.2016.2579198.
[19] M. Satyanarayanan, "The emergence of edge computing," Computer, vol. 50, no. 1, pp. 30-39, 2017, doi: 10.1109/MC.2017.9.
[20] W. Maass, "Networks of spiking neurons: The third generation of neural network models," Neural Networks, vol. 10, no. 9, pp. 1659-1671, 1997, doi: 10.1016/S0893-6080(97)00011-7.
[21] P. A. Merolla et al., "A million spiking-neuron integrated circuit with a scalable communication network and interface," Science, vol. 345, no. 6197, pp. 668-673, 2014, doi: 10.1126/science.1254642.
[22] S. B. Furber, F. Galluppi, S. Temple, and L. A. Plana, "The SpiNNaker project," Proc. IEEE, vol. 102, no. 5, pp. 652-665, 2014, doi: 10.1109/JPROC.2014.2304638.
[23] M. Davies et al., "Loihi: A neuromorphic manycore processor with on-chip learning," IEEE Micro, vol. 38, no. 1, pp. 82-99, 2018, doi: 10.1109/MM.2018.112130359.
[24] G. Orchard et al., "Efficient neuromorphic signal processing with Loihi 2," in Proc. IEEE Workshop on Signal Processing Systems (SiPS), 2021, pp. 254-259.
[25] E. O. Neftci, H. Mostafa, and F. Zenke, "Surrogate gradient learning in spiking neural networks," IEEE Signal Process. Mag., vol. 36, no. 6, pp. 51-63, 2019, doi: 10.1109/MSP.2019.2931595.
[26] F. Zenke and T. P. Vogels, "The remarkable robustness of surrogate gradient learning for instilling complex function in spiking neural networks," Neural Comput., vol. 33, no. 4, pp. 899-925, 2021, doi: 10.1162/neco_a_01367.
[27] J. K. Eshraghian et al., "Training spiking neural networks using lessons from deep learning," Proc. IEEE, vol. 111, no. 9, pp. 1016-1054, 2023, doi: 10.1109/JPROC.2023.3308088.
[28] W. Fang et al., "SpikingJelly: An open-source machine learning infrastructure platform for spike-based intelligence," Sci. Adv., vol. 9, no. 40, eadi1480, 2023, doi: 10.1126/sciadv.adi1480.
[29] G. Tang, N. Kumar, R. Yoo, and K. P. Michmizos, "Deep reinforcement learning with population-coded spiking neural network for continuous control," in Proc. Conf. Robot Learning (CoRL), PMLR 155, 2021, pp. 2016-2029.
[30] G. Tang, N. Kumar, and K. P. Michmizos, "Reinforcement co-learning of deep and spiking neural networks for energy-efficient mapless navigation with neuromorphic hardware," in Proc. IEEE/RSJ IROS, 2020, pp. 6090-6097.
[31] K. Roy, A. Jaiswal, and P. Panda, "Towards spike-based machine intelligence with neuromorphic computing," Nature, vol. 575, pp. 607-617, 2019, doi: 10.1038/s41586-019-1677-2.
[32] C. D. Schuman et al., "Opportunities for neuromorphic computing algorithms and applications," Nat. Comput. Sci., vol. 2, pp. 10-19, 2022, doi: 10.1038/s43588-021-00184-y.
[33] D. Kudithipudi et al., "Neuromorphic computing at scale," Nature, vol. 637, pp. 801-812, 2025, doi: 10.1038/s41586-024-08253-8.
[34] M. Nilsson et al., "Integration of neuromorphic AI in event-driven distributed digitized systems: Concepts and research directions," Front. Neurosci., vol. 17, 1074439, 2023.
[35] G. Liu, W. Deng, X. Xie, L. Huang, and H. Tang, "Human-level control through directly trained deep spiking Q-networks," IEEE Trans. Cybern., vol. 53, no. 11, pp. 7187-7198, 2023.
[36] D. Zhang, T. Zhang, S. Jia, and B. Xu, "Multi-scale dynamic coding improved spiking actor network for reinforcement learning," in Proc. AAAI, vol. 36, no. 1, 2022, pp. 59-67.
[37] D. Chen, L. Peng, T. Li, et al., "Fully spiking actor network with intralayer connections for reinforcement learning," IEEE Trans. Neural Netw. Learn. Syst., 2024.
[38] L. Zanatta et al., "Directly-trained spiking neural networks for deep reinforcement learning: Energy efficient implementation of event-based obstacle avoidance on a neuromorphic accelerator," Neurocomputing, vol. 562, 126885, 2023.
[39] N. Nasir and S. Al Hamadi, "Towards the neuromorphic Cyber-Twin: An architecture for cognitive defense in digital twin ecosystems," Front. Big Data, 2025.
[40] J. Chen, S. Park, P. Popovski, H. V. Poor, and O. Simeone, "Neuromorphic split computing with wake-up radios: Architecture and design via digital twinning," IEEE Trans. Signal Process., 2024.
[41] X. Xi, "Research on the construction of a digital twin system for sports venues through the collaboration of multimodal diffusion models and Internet of Things data," Discov. Artif. Intell., 2026.
[42] E. Glebova, M. Book, Y. Su, M. Peric, and J. Heller, "Sports venue digital twin technology from a spectator virtual visiting perspective," Front. Sports Act. Living, vol. 5, 1289140, 2023.
[43] L. Yu et al., "Multi-agent deep reinforcement learning for HVAC control in commercial buildings," IEEE Trans. Smart Grid, vol. 12, no. 1, pp. 407-419, 2021, doi: 10.1109/TSG.2020.3011739.
[44] S. Nagarathinam, V. Menon, A. Vasan, and A. Sivasubramaniam, "MARCO - Multi-agent reinforcement learning based control of building HVAC systems," in Proc. ACM e-Energy, 2020, pp. 57-67.
[45] B. Zhang, W. Hu, A. M. Y. M. Ghias, X. Xu, and Z. Chen, "Multi-agent deep reinforcement learning-based coordination control for grid-aware multi-buildings," Appl. Energy, vol. 328, 120215, 2022.
[46] Q. Fu et al., "Optimal control method of HVAC based on multi-agent deep reinforcement learning," Energy Build., vol. 270, 112284, 2022.
[47] X.-B. Liu et al., "Enhancing HVAC energy management through multi-zone occupant-centric approach: A multi-agent deep reinforcement learning solution," Energy Build., vol. 303, 113770, 2024.
[48] Z. Wang and T. Hong, "Reinforcement learning for building controls: The opportunities and challenges," Appl. Energy, vol. 269, 115036, 2020, doi: 10.1016/j.apenergy.2020.115036.
[49] Z. Nagy et al., "Ten questions concerning reinforcement learning for building energy management," Build. Environ., vol. 241, 110435, 2023, doi: 10.1016/j.buildenv.2023.110435.
[50] C. Miller et al., "The Building Data Genome Project 2, energy meter data from the ASHRAE Great Energy Predictor III competition," Sci. Data, vol. 7, 368, 2020, doi: 10.1038/s41597-020-00712-x.
[51] N. Rose, "Crowdedness at the Campus Gym" (UC Berkeley RSF people counts, 2015-2017), Kaggle dataset, 2017. [Online]. Available: https://www.kaggle.com/datasets/nsrose7224/crowdedness-at-the-campus-gym
[52] L. M. Candanedo and V. Feldheim, "Accurate occupancy detection of an office room from light, temperature, humidity and CO2 measurements using statistical learning models," Energy Build., vol. 112, pp. 28-39, 2016, doi: 10.1016/j.enbuild.2015.11.071.
[53] A. Alsaedi, N. Moustafa, Z. Tari, A. Mahmood, and A. Anwar, "TON_IoT telemetry dataset: A new generation dataset of IoT and IIoT for data-driven intrusion detection systems," IEEE Access, vol. 8, pp. 165130-165150, 2020, doi: 10.1109/ACCESS.2020.3022862.
[54] N. Moustafa, "A new distributed architecture for evaluating AI-based security systems at the edge: Network TON_IoT datasets," Sustain. Cities Soc., vol. 72, 102994, 2021, doi: 10.1016/j.scs.2021.102994.

[55] A. Zeng, M. Chen, L. Zhang, and Q. Xu, "Are transformers effective for time series forecasting?," in Proc. AAAI, vol. 37, no. 9, 2023, pp. 11121-11128.
[56] Y. Nie, N. H. Nguyen, P. Sinthong, and J. Kalagnanam, "A time series is worth 64 words: Long-term forecasting with transformers," in Proc. ICLR, 2023.
[57] S. Ross, G. Gordon, and D. Bagnell, "A reduction of imitation learning and structured prediction to no-regret online learning," in Proc. AISTATS, PMLR 15, 2011, pp. 627-635.
[58] D. Chen, B. Zhou, V. Koltun, and P. Krahenbuhl, "Learning by cheating," in Proc. CoRL, PMLR 100, 2020, pp. 66-75.
[59] T. Johannink et al., "Residual reinforcement learning for robot control," in Proc. IEEE ICRA, 2019, pp. 6023-6029, doi: 10.1109/ICRA.2019.8794127.
[60] F. Wilcoxon, "Individual comparisons by ranking methods," Biometrics Bull., vol. 1, no. 6, pp. 80-83, 1945, doi: 10.2307/3001968.
[61] S. Holm, "A simple sequentially rejective multiple test procedure," Scand. J. Stat., vol. 6, no. 2, pp. 65-70, 1979.
[62] J. Demsar, "Statistical comparisons of classifiers over multiple data sets," J. Mach. Learn. Res., vol. 7, pp. 1-30, 2006.
[63] O. Richter et al., "DYNAP-SE2: A scalable multi-core dynamic neuromorphic asynchronous spiking neural network processor," Neuromorph. Comput. Eng., vol. 4, 014003, 2024.
[64] F. Zenke and S. Ganguli, "SuperSpike: Supervised learning in multilayer spiking neural networks," Neural Comput., vol. 30, no. 6, pp. 1514-1541, 2018, doi: 10.1162/neco_a_01086.
[65] W. Fang et al., "Incorporating learnable membrane time constant to enhance learning of spiking neural networks," in Proc. IEEE/CVF ICCV, 2021, pp. 2661-2671.
[66] Z. Zhou et al., "Spikformer: When spiking neural network meets transformer," in Proc. ICLR, 2023.
[67] F. Moro et al., "Neuromorphic principles meet reinforcement learning for virtual robotic control," Neuromorph. Comput. Eng., 2026.
[68] S. Lucas and E. Portillo, "Methodology based on spiking neural networks for univariate time-series forecasting," Neural Netw., vol. 173, 106171, 2024.
[69] D. L. Manna, A. Vicente-Sola, P. Kirkland, T. J. Bihl, and G. Di Caterina, "Time series forecasting via derivative spike encoding and bespoke loss functions for spiking neural networks," Computers, vol. 13, no. 8, 202, 2024.
[70] C. Lv et al., "Efficient and effective time-series forecasting with spiking neural networks," in Proc. ICML, PMLR 235, 2024.
[71] Z. Wang et al., "An efficient intrusion detection model based on convolutional spiking neural network," Sci. Rep., vol. 14, 7054, 2024.
[72] A. R. Zarzoor, N. A. S. Al-Jamali, and D. A. A. Qader, "Intrusion detection method for internet of things based on the spiking neural network and decision tree method," Int. J. Electr. Comput. Eng., vol. 13, no. 2, pp. 2278-2288, 2023.
[73] S. Mihai et al., "Digital twins: A survey on enabling technologies, challenges, trends and future prospects," IEEE Commun. Surveys Tuts., vol. 24, no. 4, pp. 2255-2291, 2022.
[74] Y. Wang et al., "A survey on digital twins: Architecture, enabling technologies, security and privacy, and future prospects," IEEE Internet Things J., vol. 10, no. 17, pp. 14965-14987, 2023.
[75] A. S. Cespedes-Cubides and M. Jradi, "A review of building digital twins to improve energy efficiency in the building operational stage," Energy Informatics, vol. 7, 11, 2024.
[76] L. Hadjidemetriou et al., "A digital twin architecture for real-time and offline high granularity analysis in smart buildings," Sustain. Cities Soc., vol. 98, 104795, 2023.
[77] S. Agostinelli, F. Cumo, G. Guidi, and C. Tomazzoli, "Cyber-physical systems improving building energy management: Digital twin and artificial intelligence," Energies, vol. 14, no. 8, 2338, 2021.
[78] A. Francisco, N. Mohammadi, and J. E. Taylor, "Smart city digital twin-enabled energy management: Toward real-time urban building energy benchmarking," J. Manage. Eng., vol. 36, no. 2, 04019045, 2020.
[79] Y. Zhou et al., "Advances in emerging digital technologies for energy efficiency and energy integration in smart cities," Energy Build., vol. 315, 114289, 2024.
[80] S. R. Samaei et al., "Cognitive digital twins for climate-resilient building energy systems: Diagnosis-informed control under extreme thermal stress," Artif. Intell. Sustain. Cities, 2026.
[81] O. Hakimi, H. Liu, and O. Abudayyeh, "Digital twin-enabled smart facility management: A bibliometric review," Front. Eng. Manag., vol. 11, pp. 32-49, 2024.
[82] F. Zuccari et al., "Energy analysis of swimming pools for sports activities: Cost effective solutions for efficiency improvement," Energy Procedia, vol. 126, pp. 123-130, 2017.
[83] D. Al. Katsaprakakis et al., "Rational use of energy in sports centres to achieve net zero: The SAVE project (Part A)," Energies, vol. 16, no. 10, 4040, 2023.
[84] J. Chen, N. Skatchkovsky, and O. Simeone, "Neuromorphic wireless cognition: Event-driven semantic communications for remote inference," IEEE Trans. Cogn. Commun. Netw., vol. 9, no. 2, pp. 252-265, 2023.
[85] J. R. Vazquez-Canteli and Z. Nagy, "Reinforcement learning for demand response: A review of algorithms and modeling techniques," Appl. Energy, vol. 235, pp. 1072-1089, 2019, doi: 10.1016/j.apenergy.2018.11.002.
[86] M. Biemann, F. Scheller, X. Liu, and L. Huang, "Experimental evaluation of model-free reinforcement learning algorithms for continuous HVAC control," Appl. Energy, vol. 298, 117164, 2021.
[87] T. Wei, Y. Wang, and Q. Zhu, "Deep reinforcement learning for building HVAC control," in Proc. ACM/IEEE DAC, 2017, doi: 10.1145/3061639.3062224.
[88] J. Drgona et al., "All you need to know about model predictive control for buildings," Annu. Rev. Control, vol. 50, pp. 190-232, 2020, doi: 10.1016/j.arcontrol.2020.09.001.
[89] Y. Xia et al., "Federated accelerated deep reinforcement learning for multi-zone HVAC control in commercial buildings," IEEE Trans. Smart Grid, 2025.
[90] Y.-S. Pei et al., "Multi-agent hierarchical deep reinforcement learning for HVAC control with flexible DERs," IEEE Trans. Smart Grid, 2025.
[91] R. Lowe et al., "Multi-agent actor-critic for mixed cooperative-competitive environments," in Proc. NeurIPS, 2017.
[92] T. Rashid et al., "Monotonic value function factorisation for deep multi-agent reinforcement learning," J. Mach. Learn. Res., vol. 21, no. 178, pp. 1-51, 2020.
[93] B. Peng et al., "FACMAC: Factored multi-agent centralised policy gradients," in Proc. NeurIPS, 2021.
[94] C. Schroeder de Witt et al., "Is independent learning all you need in the StarCraft multi-agent challenge?," arXiv:2011.09533, 2020.
[95] C. Yu et al., "The surprising effectiveness of PPO in cooperative multi-agent games," in Proc. NeurIPS Datasets and Benchmarks, 2022.
[96] M. Abadi et al., "Deep learning with differential privacy," in Proc. ACM CCS, 2016, pp. 308-318, doi: 10.1145/2976749.2978318.
[97] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in Proc. AISTATS, PMLR 54, 2017, pp. 1273-1282.
[98] M. Khalil, M. Esseghir, and L. Merghem-Boulahia, "Privacy-preserving federated learning: An application for big data load forecast in buildings," Comput. Secur., vol. 131, 103211, 2023.
[99] ISO 13790:2008, Energy performance of buildings - Calculation of energy use for space heating and cooling. Geneva: ISO, 2008.
[100] ANSI/ASHRAE Standard 62.1-2019, Ventilation for Acceptable Indoor Air Quality. Atlanta, GA: ASHRAE, 2019.

[101] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in Proc. ICLR, 2015.
[102] M. Horowitz, "Computing's energy problem (and what we can do about it)," in Proc. IEEE ISSCC, 2014, pp. 10-14, doi: 10.1109/ISSCC.2014.6757323.
[103] Q. Huangfu and J. A. J. Hall, "Parallelizing the dual revised simplex method," Math. Program. Comput., vol. 10, no. 1, pp. 119-142, 2018, doi: 10.1007/s12532-017-0130-5.
[104] S. Fujimoto, H. van Hoof, and D. Meger, "Addressing function approximation error in actor-critic methods," in Proc. ICML, PMLR 80, 2018, pp. 1587-1596.
[105] W. F. Holmgren, C. W. Hansen, and M. A. Mikofski, "pvlib python: A python package for modeling solar energy systems," J. Open Source Softw., vol. 3, no. 29, 884, 2018, doi: 10.21105/joss.00884.
[106] F. Kasten and G. Czeplak, "Solar and terrestrial radiation dependent on the amount and type of cloud," Solar Energy, vol. 24, no. 2, pp. 177-189, 1980, doi: 10.1016/0038-092X(80)90391-6.
[107] S. Hochreiter and J. Schmidhuber, "Long short-term memory," Neural Comput., vol. 9, no. 8, pp. 1735-1780, 1997, doi: 10.1162/neco.1997.9.8.1735.
[108] K. Cho et al., "Learning phrase representations using RNN encoder-decoder for statistical machine translation," in Proc. EMNLP, 2014, pp. 1724-1734.
[109] S. Bai, J. Z. Kolter, and V. Koltun, "An empirical evaluation of generic convolutional and recurrent networks for sequence modeling," arXiv:1803.01271, 2018.
[110] A. Vaswani et al., "Attention is all you need," in Proc. NeurIPS, 2017.
[111] ASHRAE Guideline 14-2014, Measurement of Energy, Demand, and Water Savings. Atlanta, GA: ASHRAE, 2014.
[112] F. X. Diebold and R. S. Mariano, "Comparing predictive accuracy," J. Bus. Econ. Stat., vol. 13, no. 3, pp. 253-263, 1995.
[113] P. Henderson et al., "Deep reinforcement learning that matters," in Proc. AAAI, 2018, pp. 3207-3214, doi: 10.1609/aaai.v32i1.11694.
[114] R. Agarwal, M. Schwarzer, P. S. Castro, A. C. Courville, and M. G. Bellemare, "Deep reinforcement learning at the edge of the statistical precipice," in Proc. NeurIPS, 2021.

Auther Biography:

Mohsen Esmaeili Sani**, CODE ORCID:** 0000-0002-2963-5021
Ph.D. Sport Management, Department of Sports Management, Faculty of Sports Sciences, University of Mazandaran, Babolsar, Iran E-mail: mhx71@yahoo.com