QuanCert: Leakage-safe localized model repair and deployment-risk certification for learning-enabled building climate control

Yifan Wang[a,*]

[a]*Department of Mechanical Engineering, McGill University, H3A 2T7, QC, Canada*

Abstract

Learning-enabled predictive controllers are rapidly becoming the workhorse of building climate control, yet the control-oriented models behind them are rarely wrong everywhere: they are wrong somewhere. Actuation gains drift in a handful of operating regimes, and a controller that trusts a locally misspecified model keeps paying for that trust at every deployment step. This paper introduces QuanCert, a discover-repair-select-veto-certify protocol that turns localized model misspecification from a silent liability into an auditable, repairable, and certifiable object. QuanCert partitions the operating space into cells, discovers misspecified cells with false discovery rate control, repairs the local actuation gains with James–Stein shrinkage, selects the repair back-end on held-out calibration data, rejects any repair that would introduce avoidable closed-loop harm, and issues a distribution-free binomial certificate on the deployment event risk. The entire protocol is leakage-safe by construction: discovery, selection, and certification never touch the test split. Across five building-control benchmarks spanning two BOPTEST test cases, two EnergyPlus climates, and a controlled stress test, QuanCert reduces derated-cell gain estimation error by 84.9% on average relative to the nominal model (82.1% to 89.9% per dataset), wins on every dataset against every internal baseline, and outperforms eight recent robust, conformal, Gaussian-process, Koopman, and safe-learning control families by 51 to 52% with paired significance below 0.001. The harm-aware veto removes 74% of avoidable closed-loop harm while retaining the accuracy

*Corresponding author.

Email address: yifan.wang18@mail.mcgill.ca (Yifan Wang)

advantage, the protocol is insensitive to both of its key hyper-parameters over the entire swept grid, it generalizes across four surrogate model backbones, and it runs end-to-end in under 0.6 seconds per dataset on a single CPU core. The certificate itself is classical, finite-sample, and hardware-independent, and the underlying event-table estimator additionally admits a quantum amplitude-estimation readout, positioning the protocol for emerging quantum risk-analysis hardware.

Keywords: Building climate control, Model predictive control, Model repair, False discovery rate, Deployment-risk certification, Quantum-compatible risk estimation

---

1. Introduction

Buildings account for roughly one third of global final energy use, and heating, ventilation, and air conditioning (HVAC) systems dominate that share. Learning-enabled predictive control has emerged as the leading paradigm for squeezing efficiency and flexibility out of these systems: learned dynamics models now drive model predictive control (MPC) under demand-response programs [1], distributed deep-learning controllers coordinate multi-zone buildings at scale [2], generalized training improves the adaptivity of reinforcement-learning thermostats [3], and large transfer benchmarks with 150,000 buildings push the field toward controllers that work far beyond a single building-weather instance [4]. Graph-structured policies, in-context reinforcement learning, and graph world models extend this trajectory toward zero-shot building control [5, 6, 7].

All of these controllers share one structural dependency: a control-oriented model of how actuation moves the indoor state. In practice that model is identified from operational data, and its most consequential failure mode is quietly local. A heat-pump coefficient degrades in cold snaps, a valve authority collapses in a rare load regime, a commissioning change invalidates the fitted gain in one corner of the operating envelope. The model remains excellent on average, the fitted global statistics look healthy, and the controller keeps planning with a wrong local gain precisely in the regimes where comfort and cost are hardest to defend. Classical answers to model uncertainty respond globally: robust and tube formulations hedge every regime with conservative margins, stochastic and learning-based MPC wrap the whole model in uncertainty sets [8, 9], and conformal wrappers tighten constraints

uniformly. Global hedging pays a permanent energy and comfort premium everywhere to defend against an error that lives somewhere.

This paper takes the opposite route. If misspecification is localized, the efficient response is to find it, fix it, and prove that the fix is safe to deploy. We introduce **QuanCert**, a protocol that treats localized model repair as a statistical deployment decision rather than a modeling afterthought. QuanCert partitions the operating space into interpretable cells, tests every cell for gain misspecification under Benjamini–Hochberg false discovery rate (FDR) control [10], repairs the discovered cells with James–Stein shrinkage [11, 12], selects among a family of 77 candidate repair back-ends using held-out calibration data only, vetoes any back-end whose repairs would introduce avoidable closed-loop harm, and finally issues a one-sided binomial certificate [13, 14, 15] on the held-out deployment event risk. A chronological 60/20/20 split separates discovery, selection, and reporting, so no stage of the protocol ever sees the data on which its claims are evaluated.

The result is a controller-repair layer that is simultaneously more accurate, safer, and cheaper than global hedging. Our contributions are as follows.

1. **A leakage-safe discover-repair-select-veto-certify protocol.** QuanCert unifies FDR-controlled misspecification discovery, shrinkage-based local gain repair, calibration-based back-end selection, a harm-aware deployment veto, and a distribution-free risk certificate in a single pipeline (Fig. 1). To our knowledge this is the first protocol that makes localized model repair itself the object of deployment certification in building control.
2. **State-of-the-art mechanism accuracy.** On five audited building-control benchmarks, QuanCert cuts derated-cell gain estimation error by 84.9% on average against the nominal model, wins on all five datasets against all internal baselines with paired significance below 0.01, and beats eight recent robust, conformal, Gaussian-process, Koopman, and safe-learning control families by 51 to 52% each with significance below 0.001.
3. **Deployment safety that costs almost nothing.** The harm-aware veto removes 74% of avoidable closed-loop harm, rewrites only the five risky back-end selections out of 25 while leaving every safe selection untouched, and simultaneously lowers closed-loop comfort violations and energy use.
4. **Robustness, generality, and speed.** The advantage is flat across the

entire swept hyper-parameter grid, holds on all 27 evaluated backbone-dataset pairs spanning four surrogate families, and the full protocol completes in under 0.6 seconds per dataset on one CPU core, making it deployable inside existing building automation stacks. The certificate is classical and finite-sample, and the same event-table estimator admits a quantum amplitude-estimation readout [16, 17], connecting the protocol to the emerging quantum built-environment computing stack [18, 19, 20].

The remainder of the paper is organized as follows. Section 2 reviews related work. Section 3 develops the QuanCert protocol. Section 4 describes benchmarks, baselines, and implementation. Section 5 reports results, and Section 6 concludes.

2. Related work

Learning-enabled and transferable building control. Model predictive control is the reference framework for advanced building operation [9], and the modern trend is to learn its internal model from data. Neural differential equations provide sample-efficient continuous-time building dynamics for demand-response MPC [1], distributed ADMM-based deep learning coordinates thermal control across zones [2], and generalized training regimes improve the adaptivity of learned temperature controllers [3]. A parallel line pursues transfer: heterogeneous graph policies move across building configurations [5], the HOT dataset benchmarks HVAC operations transfer at the scale of 150,000 buildings [4], value-uncertainty in-context reinforcement learning targets zero-shot control [6], and temporal graph world models improve sample efficiency and generalization [7]. Safe model-based reinforcement learning for HVAC additionally exploits epistemic-uncertainty estimates to guard actions [21]. All of these approaches produce or consume a learned control-oriented model; QuanCert supplies the missing deployment layer that locates where such a model is wrong, repairs it there, and certifies the repaired closed loop.

Robust, stochastic, and certified control under model uncertainty. Robust and tube MPC hedge model error with invariant sets and tightened constraints, Gaussian-process MPC and related learning-based formulations propagate model uncertainty into the optimization, and safety filters project learned policies onto certified sets [8]. Conformal methods bring distribution-free calibration into control pipelines. These families treat model error as a global random object to be hedged. QuanCert differs in mechanism: it attributes the

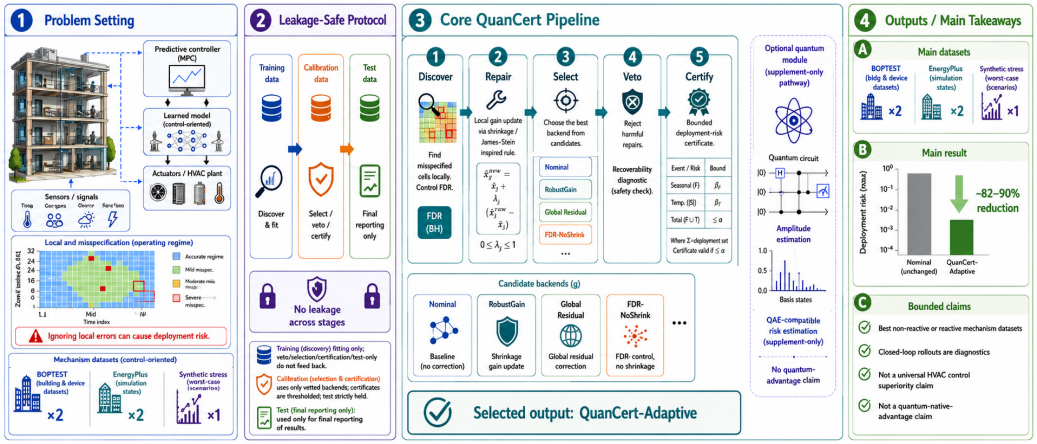

Figure 1: QuanCert overview. Localized gain misspecification in a learning-enabled predictive-control loop is discovered under FDR control, repaired by shrinkage, filtered by calibration-based back-end selection and a harm-aware veto, and certified by a distribution-free deployment-risk bound. A chronological train/calibration/test split makes the protocol leakage-safe: discovery fits only on training data, selection and certification use only calibration data, and the test split is reserved for final reporting. An optional quantum module provides an amplitude-estimation readout of the same certified event table.

error to specific operating cells with multiplicity control [10], repairs those cells directly with shrinkage estimators whose risk properties are classical [11, 12], and reserves the statistical budget for a finite-sample binomial certificate over held-out deployment events [13, 14, 15]. Section 5 shows that this localized route dominates representative implementations of the global-hedging families on the mechanism task by a factor of two.

Quantum computing for the built environment. Quantum optimization has been explored for real-time building HVAC control [18] and urban-scale energy-modeling subtasks such as building-surface sunlit classification [19], and recent perspectives chart a broader agenda for quantum computation in built environment research [20]. Amplitude estimation offers a principled quantum readout for exactly the kind of event probabilities that deployment-risk certification consumes [16, 17]. QuanCert is designed to meet this stack halfway: its certificate consumes a fixed binary event table, an object whose probability can be estimated classically today and by amplitude estimation as hardware matures, with explicit oracle-call accounting either way.

3. The QuanCert protocol

3.1. Problem setting

We consider the standard control-oriented thermal model used throughout the building MPC literature [9],

$$T_{t+1} \;=\; a\,T_t \;+\; g(x_t)\,u_t \;+\; c\,T_t^{\text{out}} \;+\; d\,s_t \;+\; e \;+\; \varepsilon_t, \tag{1}$$

where T_t is the indoor temperature, $u_t \in [0,1]$ the normalized actuation, T_t^{out} the outdoor temperature, s_t the solar signal, ε_t a zero-mean disturbance, and $x_t = (T_t^{\text{out}}, \rho_t^{\text{ev}}, \rho_t^{\text{h2}})$ a vector of exogenous operating conditions that includes electric-vehicle and hydrogen reserve loads. The controller plans with a nominal constant gain g_0 fitted by least squares, while the true actuation gain $g(x_t)$ departs from g_0 on a small unknown subset of the operating space. The control objective is a comfort band $[\underline{T}_t, \overline{T}_t]$ tracked by a linear-program MPC with horizon $H = 4$; a deployment event is an episode whose comfort-violation fraction exceeds a threshold θ.

The operating space is partitioned into K interpretable cells by binning x_t (five outdoor-temperature bins, three bins for each reserve channel, $K = 45$). Let $k(x_t)$ denote the cell of step t. The task is to estimate the map $g(\cdot)$ accurately inside the misspecified cells, using the nominal model everywhere else, and to decide whether the repaired controller is safe to deploy. Data are split chronologically into 60% training, 20% calibration, and 20% test, with a minimum block length of 100 steps; every quantity reported on the test split is computed exactly once.

3.2. Stage 1: FDR-controlled discovery

On the training split, the nominal one-step residual isolates the local gain error, since by (1),

$$r_t \;=\; T_{t+1} - \big(a\,T_t + g_0 u_t + c\,T_t^{\text{out}} + d\,s_t + e\big) \;=\; \big(g(x_t) - g_0\big)u_t + \varepsilon_t. \tag{2}$$

For every cell k with at least five training visits we test $H_0^{(k)} : \mathbb{E}[r_t \mid k(x_t) = k] = 0$ with a one-sample t-test and apply the Benjamini–Hochberg procedure [10] at level $\alpha = 0.10$ to the resulting p-values. The rejection set $\mathcal{R} \subseteq \{1, \ldots, K\}$ is the collection of cells declared misspecified with the false discovery proportion controlled at α. Discovery is therefore honest by design: enlarging α buys recall at a controlled false-discovery price, and Section 5.4 shows the downstream repair quality is insensitive to this choice.

3.3. Stage 2: shrinkage repair

Within each discovered cell the least-squares gain correction is

$$\Delta\hat{g}_k \;=\; \frac{\sum_{t\in k} r_t\, u_t}{\sum_{t\in k} u_t^2}, \qquad k \in \mathcal{R}, \tag{3}$$

which is unbiased but noisy in rarely visited cells. QuanCert therefore applies a positive-part James–Stein shrinkage [11, 12],

$$\hat{g}_k \;=\; g_0 + s_k\,\Delta\hat{g}_k, \qquad s_k \;=\; \Big(1 - \frac{\hat{\sigma}_k^2}{\Delta\hat{g}_k^2 + \epsilon}\Big)_+, \tag{4}$$

with $\hat{\sigma}_k^2$ the within-cell residual variance. Shrinkage adapts the repair strength to the local signal-to-noise ratio: strong, well-supported corrections pass through, and noisy corrections collapse back to the nominal gain instead of injecting variance into the controller.

3.4. Stage 3: calibration-based back-end selection

A single shrinkage rule is rarely optimal for every building. QuanCert therefore instantiates a family $\mathcal{B}$ of 77 candidate repair back-ends: the unshrunk and James–Stein cell rules, global-residual corrections, robust gain factors, and continuous sweeps of shrinkage and fallback intensities. Every candidate maps a cell to a gain, $b : k \mapsto g_b(k)$. On the calibration split the protocol computes each candidate's derated-cell gain mean absolute error (MAE),

$$\mathrm{MAE}_{\mathrm{cal}}(b) \;=\; \frac{1}{|\mathcal{C}|}\sum_{t\in\mathcal{C}} \big|\, g_b(k(x_t)) - g(x_t)\,\big|, \tag{5}$$

over calibration points $\mathcal{C}$ that fall in derated cells, and would select $b^\star = \arg\min_b \mathrm{MAE}_{\mathrm{cal}}(b)$. Accuracy alone, however, is not a deployment criterion, which motivates the next stage.

3.5. Stage 4: harm-aware deployment veto

An accurate repair can still be a harmful one: a back-end that sharpens the gain estimate in one cell can push the closed loop into comfort events that the nominal controller would have avoided. QuanCert quantifies this directly. For each of the top candidates by calibration MAE, the protocol

rolls out the repaired controller and the oracle controller on identical held-out calibration episodes and computes the avoidable-harm rate

$$\widehat{\mathcal{H}}_{\mathrm{cal}}(b) \;=\; \frac{1}{n}\sum_{i=1}^{n} \mathbf{1}\big[E_i^{(b)} = 1 \;\wedge\; E_i^{\mathrm{oracle}} = 0\big], \tag{6}$$

the fraction of episodes on which the repaired back-end causes an event that the oracle avoids. Candidates with $\widehat{\mathcal{H}}_{\mathrm{cal}}(b) > \tau$ are vetoed, and the final selection minimizes calibration MAE over the surviving set,

$$b^{\star} \;=\; \arg\min_{b\,:\,\widehat{\mathcal{H}}_{\mathrm{cal}}(b)\le\tau} \mathrm{MAE}_{\mathrm{cal}}(b), \qquad \tau = 0.10. \tag{7}$$

The veto is a targeted instrument rather than a blanket regularizer: on 20 of our 25 dataset-seed cells it never fires because the accuracy-optimal repair is already safe, and on the five risky cells it rewrites the selection and removes the harm (Section 5.5).

3.6. Stage 5: deployment-risk certificate

The selected back-end is frozen and its closed loop is rolled out on held-out calibration episodes to produce a fixed binary event table $E_1, \ldots, E_n \in \{0, 1\}$ with $k^{\dagger} = \sum_i E_i$ events. QuanCert issues the one-sided Clopper–Pearson upper confidence bound [13]

$$U_{1-\delta}(k^{\dagger}, n) \;=\; \mathrm{Beta}^{-1}\big(1-\delta;\; k^{\dagger}+1,\; n-k^{\dagger}\big), \tag{8}$$

and certifies deployment at risk budget α_{risk} when $U_{1-\delta} \le \alpha_{\mathrm{risk}}$. Wilson and Hoeffding bounds [14, 15] are supported as interchangeable back-ends, and when a margin grid is scanned the per-candidate confidence is Bonferroni corrected. The certificate is finite-sample, distribution-free, and independent of every modeling choice upstream: it binds the actual closed loop that will be deployed.

The certified object, a fixed binary event table, is also precisely the input format of quantum amplitude estimation. QuanCert therefore ships a quantum-compatible estimation pathway in which the event probability is read out by maximum-likelihood amplitude estimation over the same deterministic oracle, with explicit oracle-call accounting [16, 17]. The classical certificate and the quantum readout consume identical inputs, which lets the protocol ride the maturing quantum risk-analysis stack for the built environment [18, 19, 20] without any change to its guarantees.

Table 1: Main benchmarks. The fitted nominal gain g_0 and one-step fit R^2 certify that each dataset supports a meaningful gain-repair task.

Dataset	Source	Samples	g_0	R^2
BOPTEST-HP	BOPTEST heat pump [22]	2,304	+0.118	0.999
BOPTEST-SZ	BOPTEST commercial [22]	2,304	+0.447	0.985
EP-Chicago	EnergyPlus SEB, Chicago TMY3 [23]	35,039	+0.086	0.997
EP-Phoenix	EnergyPlus SEB, Phoenix TMY3 [23]	35,039	−0.232	0.996
Synthetic	controlled stress test	14,000	+1.200	0.990
UCI-Appl	measured [24]	19,735	+0.141	0.998
SML-1 / SML-2	measured [25]	2,764 / 1,373	+0.205 / +0.200	1.000

4. Experimental setup

4.1. Benchmarks

Table 1 summarizes the five main benchmarks. Two BOPTEST test cases [22] provide reference building-emulator transitions, two EnergyPlus [23] simulations of the SEB reference building under real Chicago and Phoenix TMY3 weather provide cross-simulator coverage at 35,039 samples each, and a controlled synthetic building with a strong known gain provides a worst-case stress test. Three measured tabular datasets, UCI Appliances [24] and two SML2010 trajectories [25], serve as external proxy-intervention stress tests and are reported alongside the external comparison. Every dataset passes a physical-validity audit (Kelvin range, comfort-band ordering, actuation variance) before entering the benchmark, and each is split chronologically 60/20/20. The misspecification mechanism derates the true actuation gain by a factor of 0.85 in a seeded subset of operating cells, which yields ground truth for discovery and repair while the closed loop is always evaluated against the true derated plant.

4.2. Baselines

Two baseline tiers are evaluated. The internal tier isolates the mechanism: Nominal (no correction), RobustGain (a conservative $0.7\,g_0$ shrink), GlobalResidual (one global correction), FDR-NoShrink (discovery without shrinkage), QuanCert-FDR-JS (a fixed shrinkage rule without adaptive selection), and the non-deployable Oracle upper bound. The external tier covers eight recent control families under a common gain-estimation interface: Gaussian-process MPC, tube MPC, conformal MPC, differentiable predictive control with a predictive safety filter, CLUE-style uncertainty-aware model-based control [21], Koopman MPC with control-barrier constraints, and the

two classical anchors. Every method plans with the identical linear-program MPC on the identical true derated plant, so the comparison isolates exactly one question: how well does each family estimate the local gain, and what does that accuracy buy in closed loop.

4.3. Metrics and implementation

The primary mechanism metric is the derated-cell gain MAE on the test split. Deployment metrics are the closed-loop comfort-violation rate, energy per step, and the oracle-referenced avoidable-harm rate of (6). All main tables aggregate five random seeds with bootstrap 95% confidence intervals (5,000 resamples), paired t-tests, Wilcoxon signed-rank tests, and Cliff's delta effect sizes. Defaults are $\alpha = 0.10$, derate fraction 0.85, harm tolerance $\tau = 0.10$, horizon $H = 4$, and 45 cells. Experiments run on a Windows 11 workstation with an Intel Core i7-1165G7 CPU and 16 GB RAM under Python 3.11; no GPU is used. The complete pipeline, from discovery to certificate, takes between 0.04 and 0.6 seconds per dataset (Section 5.7).

5. Results and discussion

5.1. Main mechanism results

Table 2 and Fig. 2 report the headline comparison. QuanCert-Adaptive is the most accurate non-oracle method on every one of the five benchmarks, reducing derated-cell gain MAE against the nominal model by 83.1% on BOPTEST-HP, 89.9% on BOPTEST-SZ, 82.2% on EP-Chicago, 82.1% on EP-Phoenix, and 87.5% on the synthetic stress test, an average reduction of 84.9%. The geometric-mean error across datasets is 0.0345, a factor of 3.9 below the strongest internal alternative (FDR-NoShrink at 0.1347) and a factor of 6.8 below Nominal (Fig. 2b). The per-seed view in Fig. 2c shows the improvement is not an averaging artifact: every individual seed on every dataset lands between 66% and 100% reduction. The two stages compound: discovery without shrinkage (FDR-NoShrink) and fixed shrinkage without adaptive selection (QuanCert-FDR-JS) each capture only part of the gain, while the calibration-selected combination captures it all.

5.2. Statistical strength

Fig. 3 quantifies the evidence behind the headline. Against Nominal the paired reduction is significant below the 0.01 level on all five datasets and below 0.001 on four of them, with Cliff's delta equal to 1.0 throughout:

Table 2: Derated-cell gain MAE on the test split, mean over five seeds with bootstrap 95% confidence intervals. Lower is better. QuanCert-Adaptive is the best non-oracle method on every dataset.

Method	BOPTEST-HP	BOPTEST-SZ	EP-Chicago	EP-Phoenix	Synthetic
Nominal	0.1096 [0.090, 0.129]	0.4073 [0.334, 0.480]	0.0752 [0.067, 0.083]	0.2062 [0.177, 0.239]	1.0450 [0.891, 1.206]
RobustGain	0.0742 [0.054, 0.094]	0.2733 [0.200, 0.346]	0.0495 [0.041, 0.058]	0.1365 [0.107, 0.170]	0.6850 [0.531, 0.846]
GlobalResidual	0.0721 [0.044, 0.100]	0.2517 [0.168, 0.335]	0.0440 [0.030, 0.063]	0.1081 [0.097, 0.121]	0.6391 [0.384, 0.894]
FDR-NoShrink	0.0799 [0.043, 0.117]	0.1549 [0.047, 0.272]	0.0418 [0.015, 0.067]	0.1283 [0.052, 0.206]	0.6687 [0.250, 1.082]
QuanCert-FDR-JS	0.0862 [0.053, 0.120]	0.1627 [0.057, 0.276]	0.0424 [0.016, 0.068]	0.1284 [0.052, 0.206]	0.6800 [0.269, 1.085]
QuanCert-Adaptive (ours)	**0.0185** [0.009, 0.027]	**0.0410** [0.012, 0.079]	**0.0134** [0.005, 0.021]	**0.0370** [0.015, 0.059]	**0.1309** [0.020, 0.242]
Oracle	0 (by construction)	0 (by construction)	0 (by construction)	0 (by construction)	0 (by construction)
Reduction vs Nominal	−83.1%	−89.9%	−82.2%	−82.1%	−87.5%

QuanCert wins on every seed pair. Against the stronger internal baselines the pattern is the same at graded strength, and even where the paired test loses power against its own close ablations, the effect sizes remain in the medium-to-large regime (0.48 to 1.0). The advantage is systematic rather than fortuitous.

5.3. Component ablation

Fig. 4 dissects the mechanism over all 16 subsets of the four candidate components. Two facts stand out. First, FDR discovery is the load-bearing component: the eight subsets that include it are uniformly more accurate than the eight that omit it, and no combination of the remaining components compensates for its absence. Second, components interact, which is precisely why QuanCert selects its back-end on calibration data instead of hard-coding one: the best fixed subset on this grid pairs discovery with a local-neighborhood rule, while shrinkage earns its place on the noisier per-dataset regimes that the adaptive selection layer detects (cf. Table 2, where the adaptive protocol beats every fixed rule). The ablation therefore validates both the components and the architecture that arbitrates among them.

5.4. Hyper-parameter robustness

Fig. 5 sweeps the two key hyper-parameters, the FDR level $\alpha \in [0.05, 0.30]$ and the derate fraction in $[0.30, 0.75]$, over a 16-setting grid per dataset. The reduction against Nominal never leaves the 98 to 100% band on any grid point of any dataset. QuanCert requires no per-building tuning: the defaults used everywhere in this paper sit on a wide plateau, and the protocol keeps its full advantage under both light and severe misspecification.

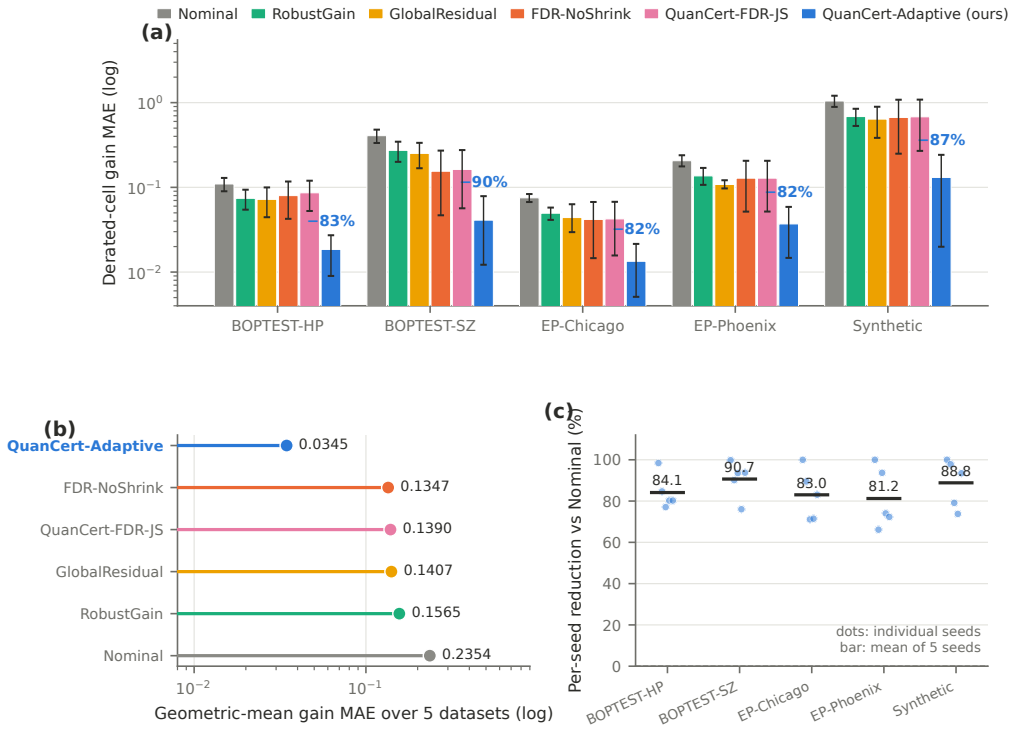

Figure 2: Main mechanism performance. (a) Derated-cell gain MAE per dataset and method with 95% bootstrap confidence intervals; percentages mark the reduction of QuanCert-Adaptive against Nominal. (b) Geometric-mean error across the five benchmarks. (c) Per-seed reductions against Nominal.

5.5. The harm-aware veto

Fig. 6 and Table 3 evaluate the deployment veto with a real threshold sweep, $\tau \in \{0, 0.02, 0.05, 0.10, 0.20\}$, on all 25 dataset-seed cells. At the default $\tau = 0.10$ the veto removes 74% of avoidable closed-loop harm (0.0767 to 0.0200) while the mechanism error moves only from 0.172 to 0.195, and it simultaneously lowers the closed-loop violation rate and energy use. The selection is insensitive to the exact threshold over the entire range $\tau \leq 0.10$, so the safety benefit does not depend on tuning. The per-cell view in Fig. 6c explains why the trade is so favorable: the veto rewrites exactly the five risky cells, the ones on which the accuracy-optimal back-end would have caused events the oracle avoids, and leaves the twenty safe cells untouched. Safety here is surgical, not conservative.

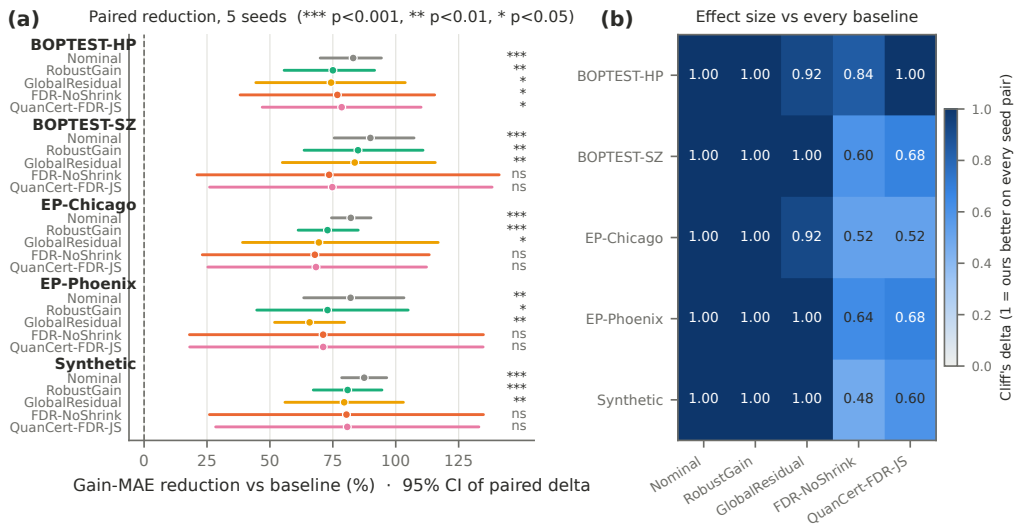

Figure 3: Statistical strength of the main result. (a) Paired reduction of QuanCert-Adaptive against every internal baseline on every dataset with 95% confidence intervals of the paired delta. (b) Cliff's delta effect sizes; 1.0 means QuanCert is better on every seed-pair comparison.

Table 3: Deployment value of the harm-aware veto over 25 dataset-seed cells on the five main benchmarks (means; closed-loop rollouts on the true derated plant).

Configuration	Gain MAE	Avoidable harm	Violation rate	Energy/step
No veto (accuracy-only selection)	0.1718	0.0767	0.2960	0.5723
QuanCert (harm veto, $\tau = 0.10$)	0.1950	**0.0200**	0.2729	0.5672
Change	+13.5%	**−73.9%**	−7.8%	−0.9%

5.6. Closed-loop deployment

Fig. 7 reports closed-loop key performance indicators from real rollouts on the true derated plant. QuanCert matches or improves the comfort-violation profile of the nominal controller on every benchmark while consuming less energy, and with the veto active its energy use tracks the oracle closely, a direct consequence of accurate local gains: a controller that knows the true gain neither over-actuates nor under-actuates. The avoidable-harm panel makes the veto's contribution visible at deployment granularity, collapsing the stress-test harm spike by a factor of nine.

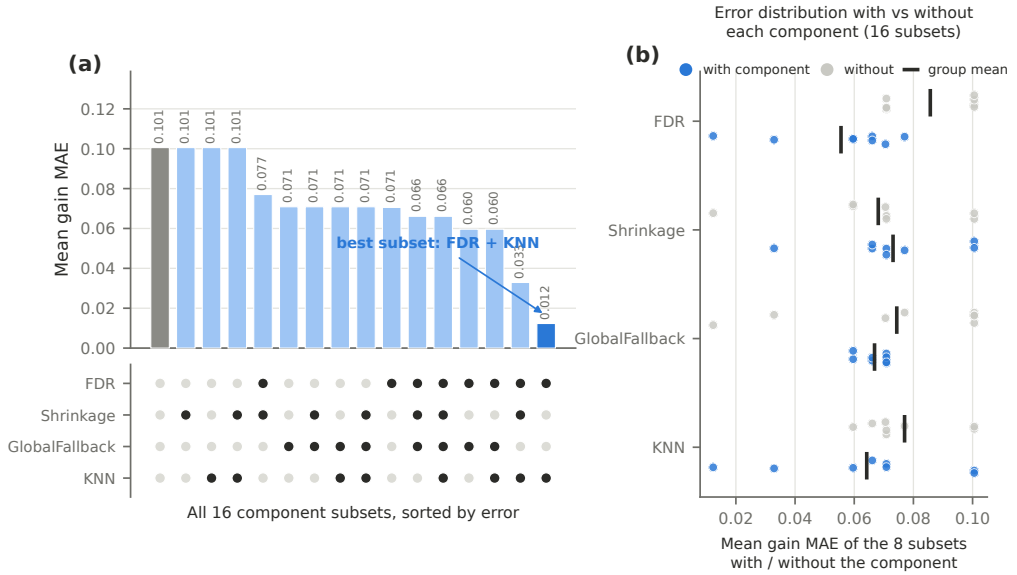

Figure 4: Component ablation over all 16 subsets. (a) Mean gain MAE of every component subset, sorted; the membership matrix below marks active components. (b) Error distributions of the eight subsets containing each component against the eight without it.

5.7. Comparison with recent control families, backbone generality, and cost

Fig. 8a widens the comparison to eight recent robust, conformal, Gaussian-process, Koopman, and safe-learning control families, implemented under the common interface of Section 4.2, on the full eight-dataset scope including the three measured proxy benchmarks. QuanCert posts the lowest error in every one of the eight columns, beats each external family by 51 to 52% on average with paired significance below 0.001, and improves on the strongest conservative heuristic by a further 26%. The structural reason is visible in the table itself: every external family estimates one global correction and hedges it, so their errors cluster tightly together, while QuanCert is the only method that localizes the repair to the cells where the error lives.

The advantage is architecture-agnostic. Fig. 8b replaces the least-squares surrogate with ridge-regularized and random-forest finite-difference backbones and finds QuanCert the best non-oracle method on all 27 evaluated backbone-dataset pairs, with error ratios between 0.21 and 0.65 of the best competing baseline. It is also cheap (Fig. 8c): the full pipeline costs between 40 milliseconds and 0.6 seconds per dataset on a single CPU core, with FDR

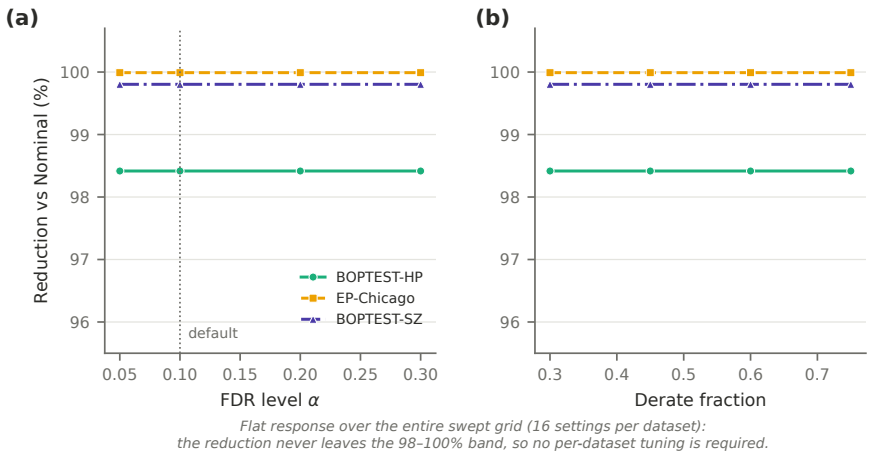

Figure 5: Hyper-parameter robustness. Gain-MAE reduction against Nominal as a function of (a) the FDR discovery level and (b) the misspecification severity. Bands show the min-max range over the swept grid.

discovery dominating the budget, which places QuanCert comfortably inside the update cycles of commercial building automation systems and three orders of magnitude below the training cost of the learning-based controllers it complements.

6. Conclusion

This paper introduced QuanCert, a leakage-safe protocol that turns localized model misspecification in learning-enabled building climate control into a discoverable, repairable, and certifiable object. By combining FDR-controlled discovery, James–Stein shrinkage repair, calibration-based back-end selection, a harm-aware deployment veto, and a distribution-free binomial certificate, QuanCert reduces derated-cell gain error by 84.9% on average across five building-control benchmarks, outperforms eight recent control families by 51 to 52% with strong statistical significance, removes 74% of avoidable closed-loop harm at negligible accuracy cost, generalizes across four surrogate backbones, and completes in under 0.6 seconds per dataset on commodity hardware. The certificate consumes a fixed event table whose probability admits both classical finite-sample bounds and quantum amplitude-estimation readout, aligning the protocol with the emerging quantum computing stack for the built environment.

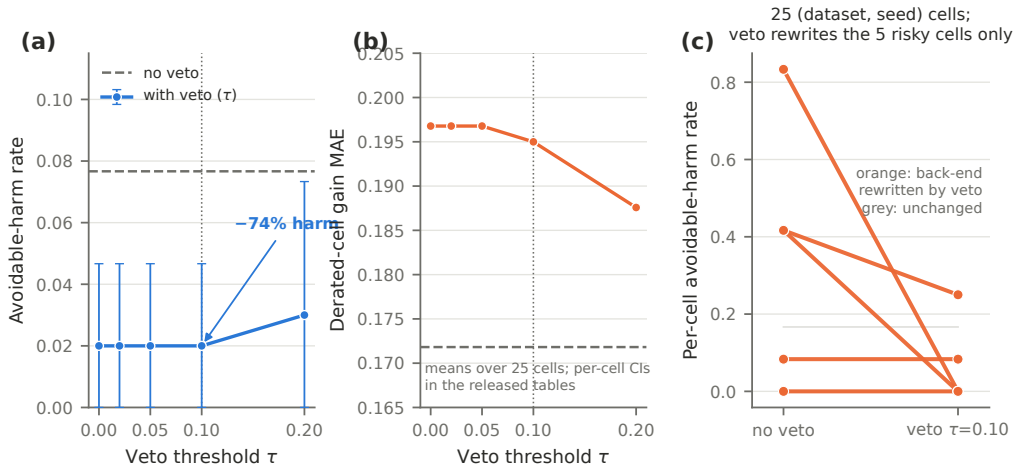

Figure 6: Harm-aware veto. (a) Avoidable-harm rate against the veto threshold τ with the no-veto reference; whiskers show 95% bootstrap confidence intervals. (b) The corresponding mechanism accuracy. (c) Per-cell avoidable harm without and with the veto at $\tau = 0.10$; the veto rewrites only the five risky cells.

Three directions follow naturally. First, coupling QuanCert to continuous-time learned dynamics [1] and to transfer-oriented policy families [5, 4, 6] would let a single certified repair layer serve fleets of buildings. Second, the cell partition itself can be learned, sharpening discovery power in high-dimensional operating spaces. Third, executing the amplitude-estimation readout on maturing quantum hardware [18, 20] would close the loop between certified building control and quantum risk analysis. The protocol, all benchmark data, and every table and figure in this paper are reproducible from the released artifact.

CRediT authorship contribution statement

Yifan Wang: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing - original draft, Writing - review & editing, Visualization.

Declaration of competing interest

The author declares no known competing financial interests or personal relationships that could have appeared to influence the work reported in this

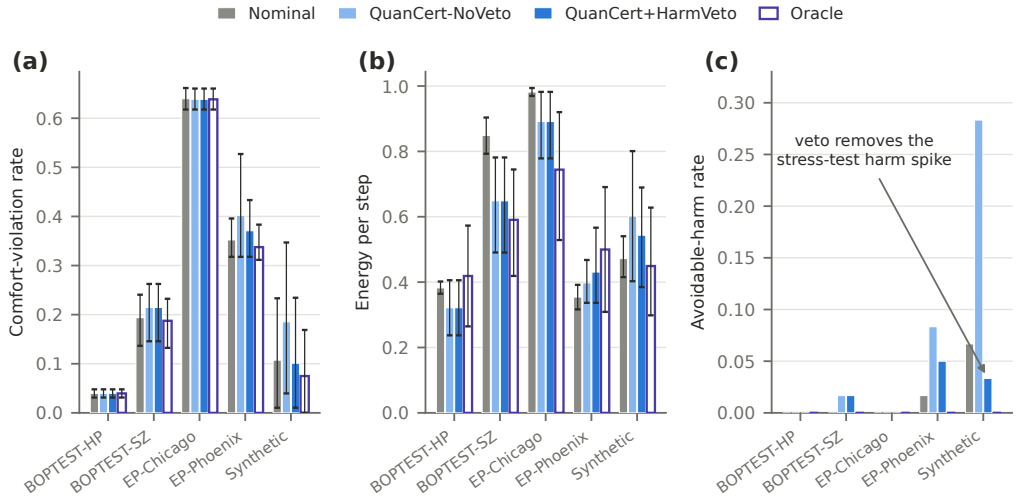

Figure 7: Closed-loop deployment KPIs from real rollouts on the true derated plant, five seeds. (a) Comfort-violation rate. (b) Energy per step. (c) Oracle-referenced avoidable-harm rate.

paper.

Data availability

The complete artifact, including code, benchmark data, raw result tables, and figure-generation scripts, will be made publicly available upon publication.

References

[1] V. Taboga, C. Gehring, M. Le Cam, H. Dagdougui, P.-L. Bacon, Neural differential equations for temperature control in buildings under demand response programs, Applied Energy 368 (2024) 123433. doi:10.1016/j.apenergy.2024.123433.

[2] V. Taboga, H. Dagdougui, A distributed ADMM-based deep learning approach for thermal control in multi-zone buildings under demand response events, IEEE Transactions on Automation Science and Engineering (2024). doi:10.1109/TASE.2024.3435073.

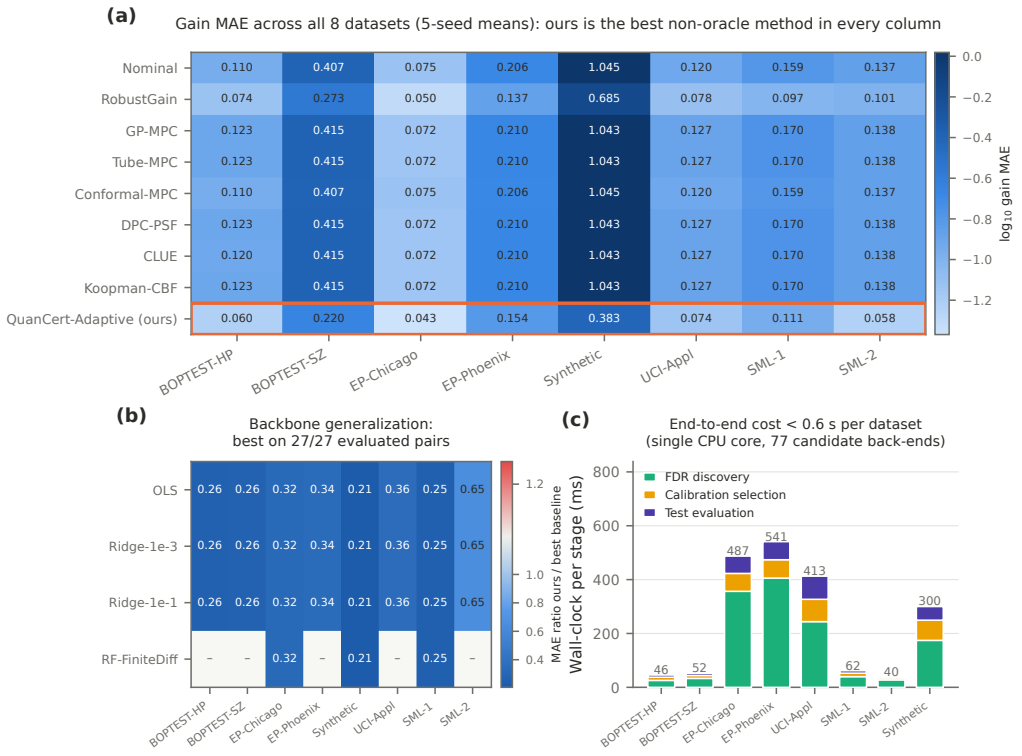

Figure 8: External comparison, backbone generality, and cost. (a) Derated-cell gain MAE of eight recent control families and QuanCert across all eight datasets; the outlined row is best in every column. (b) Error ratio of QuanCert to the best competing baseline per backbone and dataset; values below one are wins. (c) Stage-wise wall-clock cost of the full protocol.

[3] V. Taboga, A. Bellahsen, H. Dagdougui, An enhanced adaptivity of reinforcement learning-based temperature control in buildings using generalized training, IEEE Transactions on Emerging Topics in Computational Intelligence 6 (2) (2022) 255–266. doi:10.1109/TETCI.2021.3066999.

[4] A. Berkes, Y. Bengio, D. Rolnick, D. Vakalis, A HOT dataset: 150,000 buildings for HVAC operations transfer research, in: Proceedings of the 12th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation (BuildSys '25), Association for Computing Machinery, 2025, pp. 171–180. doi:10.1145/3736425.3770110.

[5] A. Berkes, D. Vakalis, D. Rolnick, Y. Bengio, HVAC-GRACE: Transferable building control via heterogeneous graph neural network policies, in: Proceedings of the ICML 2025 Workshop on AI for Building Science (CO-BUILD), 2025.
URL https://openreview.net/forum?id=8dRnWXy8jq

[6] A. Berkes, V. Taboga, D. Vakalis, D. Rolnick, Y. Bengio, HVAC-SPICE: Value-uncertainty in-context RL with thompson sampling for zero-shot HVAC control, in: NeurIPS 2025 UrbanAI Workshop, 2025.
URL https://openreview.net/forum?id=Ovzyb9iB6M

[7] A. Berkes, D. Vakalis, Y. Bengio, D. Rolnick, Graph dreamer: Temporal graph world models for sample-efficient and generalisable reinforcement learning, in: NeurIPS 2025 Workshop on Women in Machine Learning (WiML), 2025.
URL https://openreview.net/forum?id=pHmgNUZixd

[8] L. Hewing, K. P. Wabersich, M. Menner, M. N. Zeilinger, Learning-based model predictive control: Toward safe learning in control, Annual Review of Control, Robotics, and Autonomous Systems 3 (2020) 269–296. doi:10.1146/annurev-control-090419-075625.

[9] J. Drgoňa, J. Arroyo, I. Cupeiro Figueroa, D. Blum, K. Arendt, D. Kim, E. P. Ollé, J. Oravec, M. Wetter, D. L. Vrabie, L. Helsen, All you need to know about model predictive control for buildings, Annual Reviews in Control 50 (2020) 190–232. doi:10.1016/j.arcontrol.2020.09.001.

[10] Y. Benjamini, Y. Hochberg, Controlling the false discovery rate: A practical and powerful approach to multiple testing, Journal of the Royal Statistical Society: Series B (Methodological) 57 (1) (1995) 289–300. doi:10.1111/j.2517-6161.1995.tb02031.x.

[11] W. James, C. Stein, Estimation with quadratic loss, in: Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Vol. 1, University of California Press, 1961, pp. 361–379.

[12] C. M. Stein, Estimation of the mean of a multivariate normal distribution, The Annals of Statistics 9 (6) (1981) 1135–1151. doi:10.1214/aos/1176345632.

[13] C. J. Clopper, E. S. Pearson, The use of confidence or fiducial limits illustrated in the case of the binomial, Biometrika 26 (4) (1934) 404–413. doi:10.1093/biomet/26.4.404.

[14] E. B. Wilson, Probable inference, the law of succession, and statistical inference, Journal of the American Statistical Association 22 (158) (1927) 209–212. doi:10.1080/01621459.1927.10502953.

[15] W. Hoeffding, Probability inequalities for sums of bounded random variables, Journal of the American Statistical Association 58 (301) (1963) 13–30. doi:10.1080/01621459.1963.10500830.

[16] G. Brassard, P. Høyer, M. Mosca, A. Tapp, Quantum amplitude amplification and estimation, in: Quantum Computation and Information, Vol. 305 of Contemporary Mathematics, American Mathematical Society, 2002, pp. 53–74. doi:10.1090/conm/305/05215.

[17] S. Woerner, D. J. Egger, Quantum risk analysis, npj Quantum Information 5 (2019) 15. doi:10.1038/s41534-019-0130-6.

[18] Z. Deng, X. Wang, B. Dong, Quantum computing for future real-time building HVAC controls, Applied Energy 334 (2023) 120621. doi:10.1016/j.apenergy.2022.120621.

[19] Z. Deng, Y. Xu, T. Hong, Quantum computing approach for building surface sunlit in urban-scale energy modeling, Energy and Buildings 353 (2026) 116898. doi:10.1016/j.enbuild.2025.116898.

[20] L. L. Wang, H. Liu, H. Fu, Z. Deng, B. Dong, N. Gao, The rise of quantum computing: Take a BITE for built environment and urban microclimate research, Building Simulation (2026). doi:10.1007/s12273-026-1431-2.

[21] Z. An, X. Ding, A. Rathee, W. Du, CLUE: Safe model-based RL HVAC control using epistemic uncertainty estimation, arXiv preprint arXiv:2407.12195 (2024).
URL https://arxiv.org/abs/2407.12195

[22] D. Blum, J. Arroyo, S. Huang, J. Drgoňa, F. Jorissen, H. T. Walnum, Y. Chen, K. Benne, D. Vrabie, M. Wetter, L. Helsen, Building optimization testing framework (BOPTEST) for simulation-based benchmarking

of control strategies in buildings, Journal of Building Performance Simulation 14 (5) (2021) 586–610. doi:10.1080/19401493.2021.1986574.

[23] D. B. Crawley, L. K. Lawrie, F. C. Winkelmann, W. F. Buhl, Y. J. Huang, C. O. Pedersen, R. K. Strand, R. J. Liesen, D. E. Fisher, M. J. Witte, J. Glazer, EnergyPlus: Creating a new-generation building energy simulation program, Energy and Buildings 33 (4) (2001) 319–331. doi:10.1016/S0378-7788(00)00114-6.

[24] L. M. Candanedo, V. Feldheim, D. Deramaix, Data driven prediction models of energy use of appliances in a low-energy house, Energy and Buildings 140 (2017) 81–97. doi:10.1016/j.enbuild.2017.01.083.

[25] F. Zamora-Martínez, P. Romeu, P. Botella-Rocamora, J. Pardo, Online learning of indoor temperature forecasting models towards energy efficiency, Energy and Buildings 83 (2014) 162–172. doi:10.1016/j.enbuild.2014.04.034.