Data-Driven Voltage Constraints for Aggregator Dispatch Under Partial Visibility: A Quantile Regression Approach

Y. Merrad and S. Ben Elghali, M Mesri

***Abstract*—Distribution system operators (DSOs) impose conservative, topology-blind active-power export bounds to guarantee voltage feasibility, but such bounds lead to unnecessary photovoltaic (PV) curtailment. This paper proposes a data-driven approach in which the aggregator exploits its partial voltage visibility to identify linear voltage surrogates via quantile regression (QR) and enforces them as dispatch constraints. The quantile formulation controls the underapproximation budget on training data; the out-of-sample calibration gap is characterized. QR is compared against ordinary least squares (OLS) and Ridge regression baselines with out-of-sample validation using pinball loss and calibration rate. Simulations on the IEEE 33-bus and 69-bus radial feeders under three visibility levels (100%, 80%, 60%) and three underapproximation budgets ($\epsilon \in \{0.08, 0.06, 0.05\}$) show that QR achieves a significantly higher injection cap on the 33-bus feeder ($+20$–23 percentage points), while on the 69-bus feeder the gains are configuration-dependent. Notably, OLS at full visibility can yield *lower* utilization than the baseline, demonstrating that the conservative quantile bias is essential. The analysis also reveals that under partial visibility, the AC power-flow bisection, rather than the surrogate, becomes the binding constraint.**

***Index Terms*—Distributed energy resources, voltage constraints, aggregator dispatch, quantile regression, partial observability, distribution networks.**

I. Introduction

THE increasing penetration of distributed energy resources (DERs), especially photovoltaic (PV) generation, creates operational challenges in radial distribution networks. Voltage rise caused by active-power injection from downstream buses may exceed admissible limits even when thermal constraints are not binding [1], [2].

When DERs are coordinated by an aggregator, the aggregator generally lacks access to the full network topology and has only *partial voltage visibility*: measurements are available at a subset of buses [3], [4]. As a consequence, the aggregator relies on DSO-imposed export bounds as the primary protective mechanism. In the most conservative case, a single common active-power limit is applied to all DERs, all buses, and all time periods. Such a rule guarantees voltage security but is typically overly restrictive [5].

The literature on voltage control shows that bus voltage can often be approximated by a linear function of active and reactive injections [2], [6]. Such surrogates, once identified from data, can be used by the aggregator without explicit knowledge of the feeder topology. The purpose of this study is to make the worst-case philosophy less conservative by using partial visibility to learn voltage-related constraints that enable more intelligent dispatch. The method can be viewed as a dynamic, dispatch-integrated form of hosting capacity allocation, where the injection cap is adapted based on learned voltage sensitivities rather than fixed by worst-case topology analysis.

Since the surrogate is used inside the dispatch problem, it must not be overly optimistic. The key insight is that the contribution lies not in learning a linear voltage model per se, but in controlling its conservative bias: a surrogate that fits the conditional mean (OLS, Ridge) can be counter-productive, yielding *more* curtailment than the topology-blind baseline. The proposed approach controls the proportion of underapproximation points through quantile regression (QR) [7], testing budgets $\epsilon \in \{0.08, 0.06, 0.05\}$. Surrogate quality is evaluated on both training and held-out test data (30% of samples) using pinball loss and calibration rate, the proper metrics for quantile models.

The remainder of this paper is organized as follows. Section II reviews related work. Section III states the problem. Section IV presents the methods. Section V presents numerical results. Section VI discusses findings. Section VII concludes.

II. Related Work

The linearized DistFlow equations [1] provide a tractable approximation of power flow in radial networks. Extensions to sensitivity-based representations, where voltage magnitude is a linear function of nodal injections, are developed in [6], and the convex relaxation framework of [2] further justifies linear surrogates for optimization. Data-driven methods that learn voltage sensitivities directly from measurements have recently been proposed [8]–[11]. In particular, [11] introduces a mean-corrected recursive estimator for online sensitivity updates in PV-rich feeders, demonstrating that adaptive surrogate maintenance improves dispatch robustness under time-varying operating conditions—a complementary direction to the offline identification pursued here. In contrast to global feasibility-learning approaches such as [12], which learn the full AC-feasible set via differentiable optimization layers, the present work focuses on conservative per-bus surrogate constraints enforceable by aggregators without full network knowledge. Direct numerical comparison with [8]–[10] is outside the scope of this work: those methods operate in a closed-loop real-time control setting at the DSO level with full network access, whereas the proposed approach is an offline aggregator-side dispatch framework under partial visibility. The two paradigms are complementary rather than competing.

Y. Merrad and S. Ben Elghali are with Aix-Marseille Université, CNRS, LIS UMR 7020, Marseille, France (e-mail: yacine.merrad@lis-lab.fr).

Manuscript received XX, 2026.

Distributionally robust and chance-constrained OPF formulations guarantee $\Pr(V_i > V^{max}) \leq \alpha$ with robustness to distributional uncertainty [13], [14]. Unlike these approaches, which require access to the full network model and operate at the DSO level, our method operates at the aggregator level under partial visibility. Data-driven hosting capacity estimation [15] determines the maximum DER penetration a feeder can accommodate; the recent equitable curtailment framework of [16] further addresses overvoltage mitigation in PV-rich active distribution networks through fair injection allocation, which the present dispatch-integrated approach complements by learning voltage-aware caps online.

The coordination of DER aggregators under DSO-imposed constraints has received increasing attention [3], [4]. Conservative power-injection caps are commonly used to translate voltage-security requirements [5]. The present work follows this philosophy while relaxing the conservatism through voltage-awareness at the aggregator level. Similar to safety-filter approaches in learning-based OPF [17] and safe multi-agent reinforcement learning for decentralized DER control [18], our AC power-flow bisection provides a second protection layer beyond statistical calibration, forming a two-layer safety architecture (statistical + physics-based). Unlike [18], which requires online interaction with the environment to learn safety guarantees, our approach is fully offline and does not assume access to real-time feedback from the distribution network.

Quantile regression [7], [19] provides a convex framework to fit linear models with asymmetric loss. Unlike standard least-squares methods, QR naturally biases the surrogate toward conservative (overapproximating) predictions, which is precisely what is needed when the surrogate serves as a dispatch constraint. The closest methodological antecedent is [20], which applies deep quantile regression as a surrogate model for joint chance-constrained OPF with renewable generation. The present work differs in three respects: it targets the aggregator's offline dispatch problem rather than the DSO's online OPF; it employs a linear (interpretable) surrogate rather than a neural network; and it explicitly characterizes partial-visibility effects. Conformal prediction methods [21] offer distribution-free coverage guarantees that could further strengthen the statistical layer; this extension is discussed in Section VI.

III. Problem Formulation

A. Notation and Feasibility Metric

Let $\mathcal{N}$ denote the set of buses and $\mathcal{N}_{DER} \subseteq \mathcal{N}$ the subset hosting DERs. The available PV generation is $\bar{P}_i(t)$. Bus net injections are $P_i^{net}(t) = P_i^{DER}(t) - P_i^L(t)$ and $Q_i^{net}(t) = Q_i^{DER}(t) - Q_i^L(t)$. The admissible voltage interval is $[V^{min}, V^{max}] = [0.95, 1.05]$ p.u. The global voltage-feasibility rate is

$$\phi = \frac{1}{T}\sum_{t=1}^{T} \mathbf{1}\big(\forall i \in \mathcal{N},\ V_i(t) \in [V^{min}, V^{max}]\big),$$

and a dispatch is acceptable if $\phi \geq 0.9$.

B. Compared Approaches

Four approaches are compared:

1) *Conservative baseline*: a single common worst-case injection fraction $\gamma \in [0,1]$, independent of bus and time.
2) *OLS surrogate*: ordinary least squares regression of voltage on net injections, incorporated into the dispatch LP.
3) *Ridge surrogate*: ℓ_2-regularized regression ($\alpha = 1.0$), same dispatch structure.
4) *Proposed QR surrogate*: quantile regression at $\tau = 1-\epsilon$, providing a built-in conservative bias.

The aggregator's incentive is economic: revenue is proportional to dispatched energy. A tighter cap reduces revenue; voltage-aware dispatch may justify a relaxed cap.

C. Portfolio and Partial Visibility

The aggregator observes voltage, knows load, and controls DER injection only at *portfolio buses* $\mathcal{N}_{V,vis}$. Non-portfolio buses are opaque. The visibility percentages 100%, 80%, and 60% correspond to the share of feeder buses in the portfolio. Under reduced visibility, the surrogate is fitted and enforced only on portfolio buses; non-portfolio injections act as unmodeled disturbances absorbed by the intercept and the budget ϵ.

IV. Methods

A. Baseline: Common Worst-Case Envelope

All DERs are constrained by a fixed injection fraction γ: $P_i^{DER}(t) = \gamma \bar{P}_i(t)$. The largest feasible γ is found by bisection: $\max_\gamma\ \gamma$ s.t. $\phi \geq 0.9$, evaluated by full AC power flow over all $T = 360$ hours.

B. OLS and Ridge Baselines

To assess the added value of quantile regression, two mean-regression baselines are incorporated into the same dispatch LP structure. Both OLS and Ridge fit the conditional mean, producing surrogates where approximately 50% of predictions underestimate the true voltage—precisely the opposite of the conservative bias required for safe dispatch. The Ridge regularization parameter $\alpha = 1.0$ is fixed rather than cross-validated; its purpose is to provide a regularized mean-regression baseline rather than an optimally tuned predictor. Cross-validated Ridge would improve R^2 on the IEEE 69-bus feeder but would preserve $\pi^{under} \approx 50\%$, leaving the qualitative comparison with QR unchanged.

C. Proposed: QR-Based Voltage Constraints

The voltage at a portfolio bus is approximated by the affine model

$$\hat{V}_i(t) = a_i + \sum_{j\in\mathcal{N}_{V,vis}} b_{ij} P_j^{net}(t) + \sum_{j\in\mathcal{N}_{V,vis}} c_{ij} Q_j^{net}(t).$$

The DER injection is $P_i^{DER}(t) = \alpha_i(t)\bar{P}_i(t)$ with $0 \leq \alpha_i(t) \leq \rho$.

1) Offline Identification: Controlling underapproximation at level ϵ is equivalent to fitting the $(1-\epsilon)$-quantile. For each visible bus i:

$$\min_{\theta_i} \sum_{s=1}^{N_s} \rho_\tau\big(V_i^{(s)} - \hat{V}_i^{(s)}(\theta_i)\big), \quad \tau = 1-\epsilon,$$

with pinball loss $\rho_\tau(u) = \tau\, u\, \mathbf{1}(u \geq 0) + (\tau - 1)\, u\, \mathbf{1}(u < 0)$. This is a linear program (LP) solvable by standard solvers. On a standard workstation using HiGHS [22], the wall-clock time per bus is below 0.1 s for $N_{train} = 800$ samples and $2|\mathcal{N}|$ features. The QR identification is performed once offline; the online dispatch LP is then solved independently for each of the $T = 360$ hourly intervals.

2) Out-of-Sample Validation: The dataset of $N_s = 1142$ operating points (each generated by randomly sampling an hour and a scaling factor, then running AC power flow) is split into training (70%, $N_{train} = 800$) and test (30%, $N_{test} = 342$). These points are *independent* of the $T = 360$ simulation hours used for dispatch evaluation.

For QR, the proper scoring metric is the pinball loss at $\tau = 1-\epsilon$, averaged over visible buses and test samples. The *calibration rate* measures the fraction of predictions where $\hat{V}_i \geq V_i$; for a well-calibrated τ-quantile model this should be $\approx \tau$. For OLS and Ridge, R^2 is reported alongside pinball loss at $\tau = 0.5$.

3) Dispatch and Tightening: The aggregator solves, for each hour t: $\max_{\alpha(t)} \sum_i \alpha_i(t)\bar{P}_i(t)$ subject to $0 \leq \alpha_i(t) \leq \rho$ and $V^{min} \leq \hat{V}_i(t) \leq V^{max}$ for all $i \in \mathcal{N}_{V,vis}$. The dispatch is validated via full AC power flow on all buses. If $\phi < 0.9$, ρ is tightened by bisection until the feasibility target is met.

V. Numerical Results

A. Simulation Setup

Two standard radial feeders are modeled in pandapower [22]: the IEEE 33-bus [1] and 69-bus [23] feeders, with ≈40% of buses hosting DERs. The simulation horizon is 15 days with hourly resolution ($T = 360$). The slack bus voltage is $V_0 = 1.045$ p.u. Load profiles combine residential, commercial, and mixed daily shapes with seasonal and weekend variations. PV profiles follow a Gaussian day shape with cloud noise. PV capacity is auto-calibrated so that the baseline γ falls in $[0.45, 0.75]$.

Three visibility levels (100%, 80%, 60%) and three hidden-bus selections (seeded for reproducibility) are tested. After each dispatch, the full AC power flow is run on the complete network and all bus voltages are checked.

B. Surrogate Quality

Table I reports R^2 scores at 100% visibility. All methods achieve excellent fit (OLS and QR: $R^2 > 0.996$; Ridge on IEEE 69: $R^2 \approx 0.92$ due to over-regularization with $\alpha = 1.0$), with negligible train-test gap confirming that the linear approximation does not overfit. However, R^2 is a mean-prediction metric and is not appropriate for evaluating quantile models. Table II reports the proper QR metrics: pinball loss at $\tau = 1-\epsilon$ and calibration rate.

TABLE I
SURROGATE QUALITY (R^2) AT 100% VISIBILITY.

Network	Method	R^2_{train}	R^2_{test}
5*IEEE 33	OLS	0.9996	0.9994
	Ridge	0.9973	0.9972
	QR (ϵ=0.08)	0.9987	0.9986
	QR (ϵ=0.06)	0.9984	0.9983
	QR (ϵ=0.05)	0.9983	0.9981
5*IEEE 69	OLS	0.9993	0.9990
	Ridge	0.9235	0.9217
	QR (ϵ=0.08)	0.9979	0.9977
	QR (ϵ=0.06)	0.9972	0.9970
	QR (ϵ=0.05)	0.9968	0.9966

TABLE II
PINBALL LOSS AND CALIBRATION RATE AT 100% VISIBILITY. FOR QR, $\tau = 1-\epsilon$; FOR OLS/RIDGE, $\tau = 0.5$.

Network	Method	Pinball ($\times 10^{-4}$)		Calibration	
(lr)3-4(lr)5-6		Train	Test	Train	Test
5*IEEE 33	OLS (τ=0.5)	1.15	1.33	0.581	0.554
	Ridge (τ=0.5)	2.39	2.28	0.576	0.579
	QR (ϵ=0.08)	0.52	0.58	0.922	0.885
	QR (ϵ=0.06)	0.42	0.46	0.942	0.917
	QR (ϵ=0.05)	0.36	0.40	0.952	0.928
5*IEEE 69	OLS (τ=0.5)	0.47	0.56	0.594	0.545
	Ridge (τ=0.5)	1.41	1.42	0.606	0.575
	QR (ϵ=0.08)	0.25	0.27	0.921	0.874
	QR (ϵ=0.06)	0.20	0.21	0.941	0.915
	QR (ϵ=0.05)	0.18	0.19	0.951	0.931

On training data, the calibration rate closely matches τ. On test data, it drops by 3–5 percentage points (e.g., 0.885 vs. target 0.92 for $\epsilon = 0.08$). This finite-sample gap is expected [19] and does not compromise dispatch safety, since the bisection provides a second protection layer via AC power-flow validation.

C. Dispatch Performance

Table III summarizes the key metrics for both networks. For each configuration, γ (baseline) or ρ (proposed) is the maximum common injection fraction maintaining $\phi \geq 0.9$. The PV utilization U, curtailed energy E^{curt} (MWh), and empirical underapproximation rate π^{under} are reported.

D. Graphical Results

Fig. 1 compares injection cap and utilization across methods at 100% visibility. On the IEEE 33-bus feeder, QR achieves $\rho \approx 0.999$ while OLS and Ridge remain at $\rho \approx \gamma$, demonstrating the critical role of the quantile bias. On the 69-bus feeder, the advantage is smaller and depends on ϵ (see Discussion).

Fig. 2 shows results across all configurations and visibility levels. On the 33-bus feeder, QR maintains $\rho \approx 0.999$ regardless of visibility, with utilization ranging from 0.657 to 0.696. On the 69-bus feeder, the gains are configuration-dependent.

Fig. 3 shows the minimum monitored-bus voltage over the first week. The proposed method maintains all bus voltages within the $[0.95, 1.05]$ p.u. band while allowing significantly

TABLE III
SUMMARY OF RESULTS: INJECTION CAP (γ/ρ), FEASIBILITY RATE (ϕ), PV UTILIZATION (U), CURTAILED ENERGY (E^{curt}, MWH), AND EMPIRICAL UNDERAPPROXIMATION RATE (π^{under}).

Network	Method	Vis. (%)	ϵ	γ/ρ	ϕ	U	E^{curt} (MWh)	π^{under}
16*IEEE 33	Baseline	all	n/a	0.462	0.90	0.462	335.2	n/a
	OLS	100	n/a	0.462	0.90	0.453	340.6	41.9%
	OLS	80	n/a	0.999	1.00	0.674	202.9	41.6%
	OLS	60	n/a	0.999	0.99	0.665	208.7	43.3%
	Ridge	100	n/a	0.462	0.90	0.453	340.9	42.4%
	Ridge	80	n/a	0.485	0.90	0.463	334.7	42.0%
	Ridge	60	n/a	0.999	0.94	0.673	204.0	44.1%
	QR	100	0.08	0.999	1.00	0.696	189.2	7.8%
	QR	100	0.06	0.999	1.00	0.696	189.7	5.8%
	QR	100	0.05	0.999	1.00	0.695	189.9	4.8%
	QR	80	0.08	0.999	1.00	0.668	206.9	7.7%
	QR	80	0.06	0.999	1.00	0.667	207.3	5.8%
	QR	80	0.05	0.999	1.00	0.667	207.4	4.8%
	QR	60	0.08	0.999	1.00	0.658	213.0	8.0%
	QR	60	0.06	0.999	1.00	0.658	213.3	6.0%
	QR	60	0.05	0.999	1.00	0.657	213.5	5.0%
16*IEEE 69	Baseline	all	n/a	0.483	0.90	0.483	209.2	n/a
	OLS	100	n/a	0.483	0.90	0.474	213.1	40.6%
	OLS	80	n/a	0.483	0.90	0.474	213.1	41.2%
	OLS	60	n/a	0.999	1.00	0.664	136.3	40.3%
	Ridge	100	n/a	0.483	0.90	0.476	212.4	39.4%
	Ridge	80	n/a	0.483	0.90	0.476	212.4	40.6%
	Ridge	60	n/a	0.563	0.90	0.518	195.1	38.3%
	QR	100	0.08	0.493	0.90	0.478	211.4	7.9%
	QR	100	0.06	0.526	0.90	0.497	203.5	5.9%
	QR	100	0.05	0.999	0.96	0.671	133.2	4.9%
	QR	80	0.08	0.493	0.90	0.478	211.4	8.0%
	QR	80	0.06	0.526	0.90	0.497	203.5	6.0%
	QR	80	0.05	0.999	0.96	0.671	133.2	5.0%
	QR	60	0.08	0.999	1.00	0.655	139.7	7.8%
	QR	60	0.06	0.999	1.00	0.655	139.9	5.9%
	QR	60	0.05	0.999	1.00	0.654	140.2	4.9%

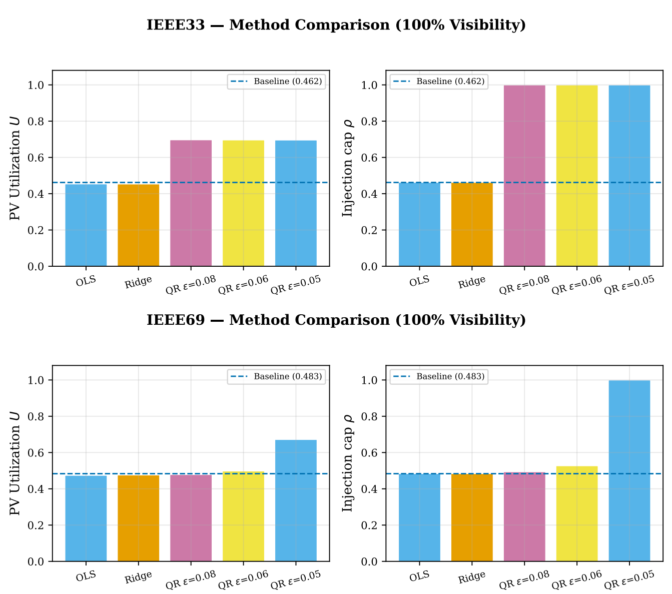

Fig. 1. Method comparison at 100% visibility: injection cap and utilization for OLS, Ridge, and QR on IEEE 33-bus (top) and 69-bus (bottom).

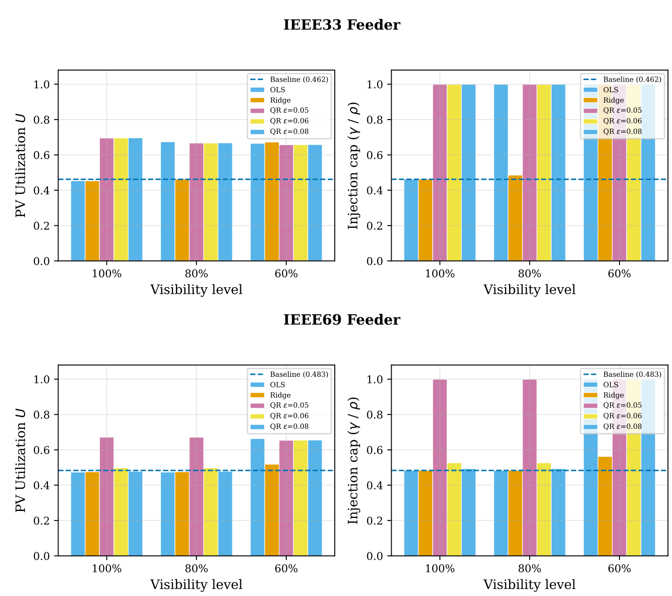

Fig. 2. Injection cap and utilization for all configurations on IEEE 33-bus (top) and 69-bus (bottom).

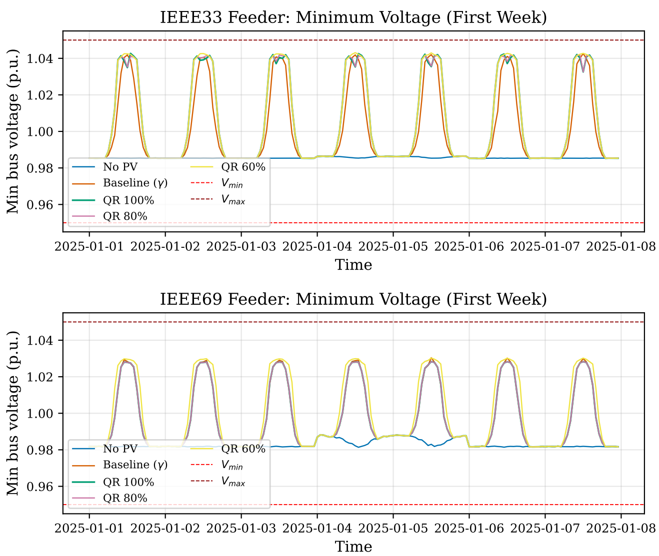

Fig. 3. Minimum bus voltage (first week) for IEEE 33-bus (top) and 69-bus (bottom).

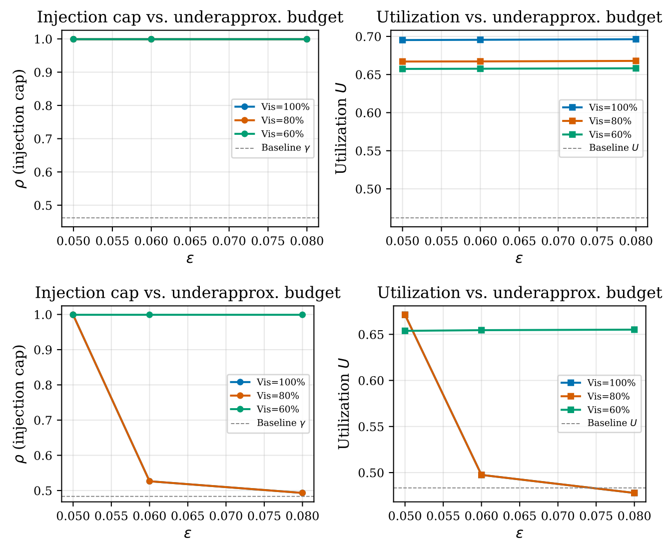

Fig. 5. Trade-off between ϵ and injection cap ρ for IEEE 33-bus (top) and 69-bus (bottom).

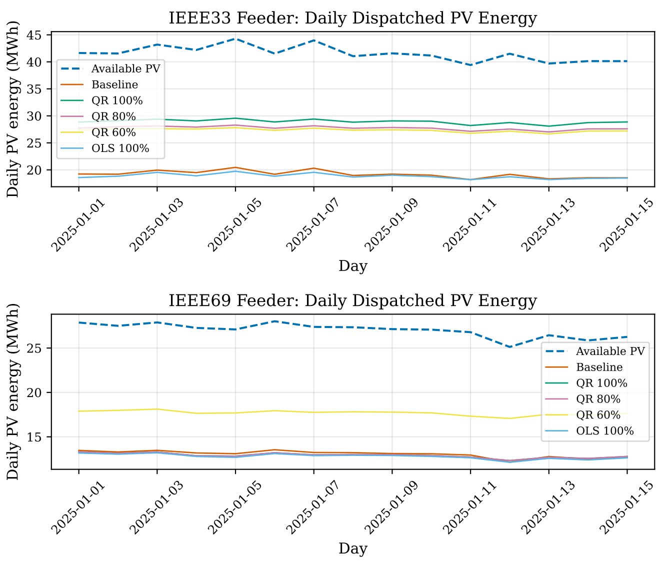

Fig. 4. Daily dispatched PV energy for IEEE 33-bus (top) and 69-bus (bottom).

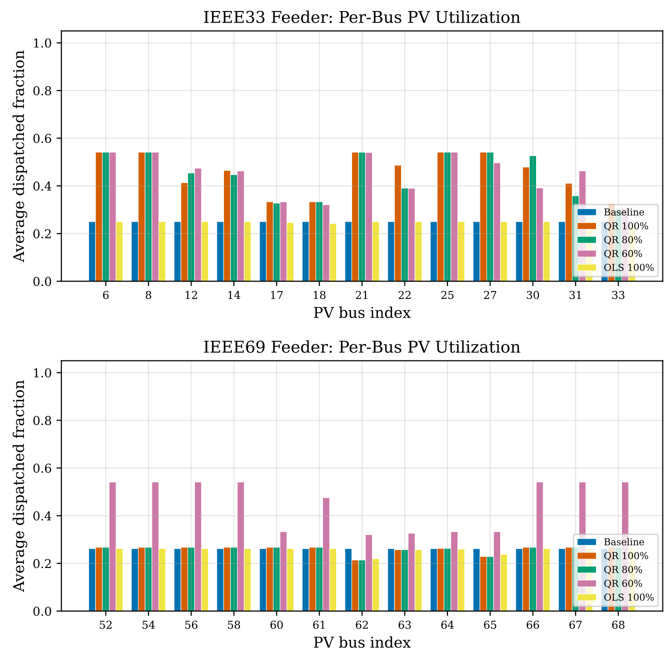

Fig. 6. Per-bus average PV utilization for IEEE 33-bus (top) and 69-bus (bottom).

more PV injection than the baseline. The voltage time series confirms that the QR surrogate does not trade off energy dispatch for hidden voltage violations.

Fig. 4 shows the daily dispatched PV energy versus the available PV over the 15-day period. QR consistently recovers a larger share of available energy than the baseline, with the gap most pronounced during high-generation days when the baseline cap becomes most restrictive.

Fig. 5 illustrates the ϵ–ρ trade-off. On the 33-bus feeder, ρ is insensitive to ϵ in the tested range. On the 69-bus feeder, a non-monotone pattern emerges at 100% and 80% visibility (discussed below).

Fig. 6 shows the average dispatched fraction per PV bus. QR distributes curtailment more efficiently across buses than the uniform baseline, concentrating injection on electrically less sensitive buses and reducing curtailment on the most constrained ones.

Fig. 7 compares in-sample and out-of-sample R^2 for all methods, confirming negligible overfitting. Fig. 8 and Fig. 9 show pinball loss and calibration rate; on test data, calibration drops by 3–5 pp due to finite-sample effects.

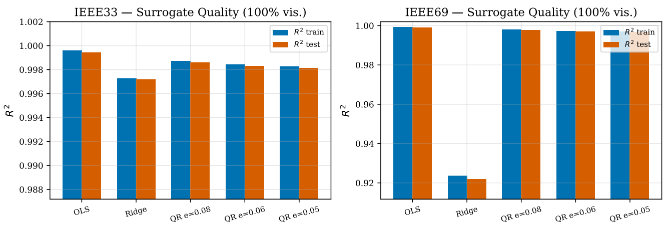

Fig. 7. R^2 (train vs. test) for OLS, Ridge, and QR at 100% visibility.

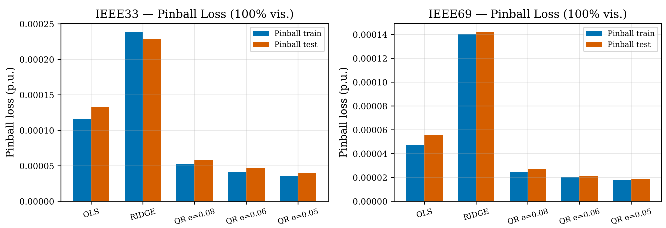

Fig. 8. Pinball loss (train vs. test) at 100% visibility.

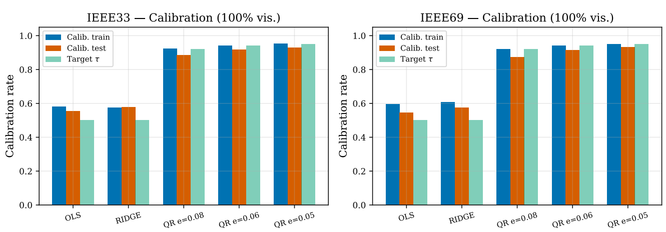

Fig. 9. Calibration rate (train vs. test). The target τ bar shows the ideal rate for a perfectly calibrated model.

VI. Discussion

The results reveal five key findings.

The conservative quantile bias is essential, not optional. Table III shows that OLS at 100% visibility on the IEEE 33-bus feeder yields $U = 0.453$, which is *lower* than the baseline $U = 0.462$. The OLS surrogate's optimistic predictions ($\approx$50% underestimation, Table II) cause the dispatch LP to trigger more voltage violations than the uniform baseline, forcing the bisection to tighten ρ below γ. By contrast, QR at the same configuration achieves $U = 0.696$ (+23 pp). The value of the proposed method therefore lies not in using a linear surrogate per se, but specifically in the quantile bias that prevents the surrogate from being counter-productive.

QR and OLS can invert at partial visibility. A subtler effect appears at reduced visibility on the IEEE 33-bus feeder: OLS at 80% achieves $U = 0.674$, marginally above QR at $U = 0.668$ (Table III). This apparent reversal is consistent with QR's conservative bias: at partial visibility, the surrogate constraints restrict per-bus injections below the global cap $\rho = 0.999$. Since QR overapproximates more consistently than OLS, it imposes tighter per-bus limits on the visible buses, marginally reducing aggregate utilization. OLS, fitting the conditional mean, produces looser per-bus constraints that allow marginally more energy through once the bisection has verified global AC feasibility. This does not invalidate QR's safety advantage—its lower π^{under} is precisely what prevents undetected voltage violations—but it shows that utilization and safety are not monotonically aligned across methods and visibility levels.

Under partial visibility, the surrogate may not be the binding constraint. On the IEEE 33-bus feeder, QR achieves $\rho \approx 0.999$ at all visibility levels, raising utilization by +20–23 pp. On the IEEE 69-bus feeder, at 100% and 80% visibility with $\epsilon = 0.08$, QR achieves only $\rho \approx 0.493$, comparable to the baseline $\gamma = 0.483$. The dispatch LP enforces the surrogate only on visible buses; after dispatch, the full AC power flow is evaluated on *all* buses. If any hidden bus violates the voltage limit, the bisection tightens ρ, and the method degenerates to the baseline. The apparent gains at 60% visibility on the 69-bus feeder reflect the particular random selection of hidden buses, which happen to be electrically less sensitive. This clarifies the scope: the QR surrogate adds value primarily when visible buses capture voltage-critical locations.

The ϵ–ρ relationship can be non-monotone. On the IEEE 69-bus feeder (Fig. 5, Table III), $\epsilon = 0.06$ yields $\rho = 0.526$, while the more conservative $\epsilon = 0.05$ achieves $\rho = 0.999$. At $\tau = 0.95$, the quantile shift predicts voltages sufficiently above true values that the surrogate constraint $\hat{V}_i \leq V^{max}$ becomes non-binding in the dispatch LP; the LP dispatches at maximum ρ, and AC power-flow validation remains feasible ($\phi = 0.96$). At $\tau = 0.94$, the shift is smaller and the surrogate remains partially binding, capping ρ at 0.526. The $\phi = 0.96$ result for $\epsilon = 0.05$ indicates that 14 of the 360 simulation hours experience at least one bus voltage violation; these violations are the direct consequence of the surrogate being non-binding, allowing an overly aggressive dispatch that the bisection does not tighten further since $\phi \geq 0.9$ is still satisfied. A stricter feasibility target (e.g., $\phi \geq 0.99$) would force the bisection to tighten ρ in this regime. On the 33-bus feeder, where the surrogate is consistently binding, ρ is insensitive to ϵ.

The calibration guarantee degrades slightly out of sample. Table II shows that test-data calibration rates are systematically 3–5 pp below the target τ (e.g., 0.885 vs. 0.92 for $\epsilon = 0.08$). This finite-sample gap results from variability in the quantile estimate and distributional shift between training and test operating points. It does not compromise dispatch safety because the bisection on AC power flow provides a second protection layer, consistent with the two-layer safety architecture advocated in [17]. For safety-critical deployments, conformal prediction techniques [21] could provide distribution-free coverage guarantees without distributional assumptions.

A. Limitations and Future Work

The present study has several limitations: (i) only active power is controlled; extending to reactive power dispatch would increase flexibility and reduce the reliance on the scalar cap; (ii) the common bound ρ is a scalar applied uniformly to all DERs; a bus-dependent cap ρ_i is a natural extension that would improve the allocation of injection rights; (iii) the feasibility criterion $\phi \geq 0.9$ counts violation hours but does not account for violation severity or temporal clustering; replacing it with a CVaR-based or chance-constrained formulation [13], [14] would provide stronger probabilistic guarantees; (iv) only a single 15-day period is simulated; multi-seasonal validation would test surrogate robustness under distributional shift, which could widen the out-of-sample calibration gap and degrade the conservation guarantee—motivating online surrogate updates as proposed in [11]; (v) only balanced single-phase equivalent models of the IEEE 33-bus and 69-bus feeders are considered; validation on unbalanced three-phase networks with OLTC regulation—where coupled voltage violations across phases introduce additional nonlinear interactions [24]—and on real DSO data would substantially strengthen applicability claims; (vi) the non-monotone ϵ–ρ behavior observed on the 69-bus feeder is an empirical finding dependent on the residual distribution; a theoretical characterization of the conditions under which the surrogate becomes non-binding remains an open question.

VII. Conclusion

This study has compared a conservative topology-blind injection bound with a data-driven voltage-aware dispatch formulation under partial visibility. The contributions are: (1) a portfolio-aware formulation of partial visibility; (2) a convex reformulation as a quantile regression LP; (3) a systematic comparison against OLS and Ridge baselines; (4) out-of-sample validation using pinball loss and calibration rate; (5) an empirical characterization of the ϵ–visibility–ρ trade-off.

On the IEEE 33-bus feeder, QR raises utilization from $U \approx 0.46$ to $U \approx 0.66$–0.70 (+20–23 pp) while maintaining $\phi \geq 0.9$ across all visibility levels. On the IEEE 69-bus feeder, the gains are configuration-dependent, with full benefits observed only when visible buses capture voltage-critical locations. The analysis reveals that under partial visibility, the AC power-flow bisection, rather than the surrogate, may become the binding constraint, causing the method to degenerate to the baseline.

The comparison with OLS and Ridge, which achieve $\pi^{under} \approx 40\%$, demonstrates that the conservative quantile bias ($\pi^{under} \leq \epsilon$ on the dispatch simulation) is essential for safe dispatch. In particular, OLS at full visibility yields *lower* utilization than the baseline, confirming that a linear surrogate alone is insufficient without the quantile shift.

Acknowledgment

This work has been carried out within the framework of the European Union's Horizon Europe research and innovation programme under Grant Agreements No. 101096946 (Flex-CHESS) and No. 101096836 (MASTERPIECE).

References

[1] M. E. Baran and F. F. Wu, "Network reconfiguration in distribution systems for loss reduction and load balancing," *IEEE Trans. Power Del.*, vol. 4, no. 2, pp. 1401–1407, Apr. 1989.

[2] M. Farivar and S. H. Low, "Branch flow model: Relaxations and convexification – Part I," *IEEE Trans. Power Syst.*, vol. 28, no. 3, pp. 2554–2564, Aug. 2013.

[3] K. Christakou, J.-Y. Le Boudec, M. Paolone, and D.-C. Tomozei, "Efficient computation of sensitivity coefficients of node voltages and line currents in unbalanced radial electrical distribution networks," *IEEE Trans. Smart Grid*, vol. 4, no. 2, pp. 741–750, Jun. 2013.

[4] E. Dall'Anese, S. S. Guggilam, A. Simonetto, Y. C. Chen, and S. V. Dhople, "Optimal power flow pursuit," *IEEE Trans. Smart Grid*, vol. 9, no. 2, pp. 1155–1165, Mar. 2018.

[5] F. Capitanescu, L. F. Ochoa, H. Margossian, and N. D. Hatziargyriou, "Assessing the potential of network reconfiguration to improve distributed generation hosting capacity in active distribution systems," *IEEE Trans. Power Syst.*, vol. 30, no. 1, pp. 346–356, Jan. 2014.

[6] S. Bolognani and S. Zampieri, "On the existence and linear approximation of the power flow solution in power distribution networks," *IEEE Trans. Power Syst.*, vol. 31, no. 1, pp. 163–172, Jan. 2016.

[7] R. Koenker and G. Bassett, "Regression quantiles," *Econometrica*, vol. 46, no. 1, pp. 33–50, Jan. 1978.

[8] R. Gupta, F. Sossan, and M. Paolone, "Counterfactual-based data-driven voltage control in distribution networks," *IEEE Trans. Smart Grid*, vol. 12, no. 5, pp. 4274–4284, Sep. 2021.

[9] X. Zhou, Z. Chen, Z. Zhou, et al., "Data-driven real-time voltage control in distribution networks," *IEEE Trans. Power Syst.*, vol. 36, no. 5, pp. 4525–4539, Sep. 2021.

[10] X. Zhou, L. Fang, and M. Paolone, "Adaptive data-driven voltage regulation in active distribution networks," *IEEE Trans. Smart Grid*, vol. 14, no. 2, pp. 1421–1432, Mar. 2023.

[11] H. Wang, Y. Yao, J. Zhao, and F. Ding, "Data-driven mean-corrected recursive estimation-based optimal DER dispatch for distribution system voltage control," *IEEE Trans. Sustainable Energy*, vol. 16, no. 4, 2025.

[12] P. L. Donti, D. Rolnick, and J. Z. Kolter, "DC3: A learning method for optimization with hard constraints," in *Proc. Int. Conf. Learn. Representations (ICLR)*, 2021.

[13] D. Bienstock, M. Munoz, and S. Oren, "Chance-constrained optimal power flow: Risk-aware network control under uncertainty," *SIAM Rev.*, vol. 56, no. 3, pp. 461–495, 2014.

[14] C. Zhao and Y. Guan, "Data-driven risk-averse stochastic optimization with Wasserstein metric," *Oper. Res. Lett.*, vol. 46, no. 2, pp. 262–267, 2018.

[15] S. Wang, L. F. Ochoa, and G. P. Harrison, "DER hosting capacity of distribution networks: A data-driven approach," *IEEE Trans. Sustainable Energy*, vol. 13, no. 2, pp. 1235–1245, Mar. 2022.

[16] E. E. Ahmed, A. Demirci, G. Poyrazoglu, and S. D. Manshadi, "An equitable active power curtailment framework for overvoltage mitigation in PV-rich active distribution networks," *IEEE Trans. Sustainable Energy*, 2024.

[17] A. Venzke, G. Qu, S. Low, and S. Chatzivasileiadis, "Learning optimal power flow: Worst-case guarantees for neural networks," in *Proc. IEEE Int. Conf. Commun., Control, Comput. Technol. Smart Grids (SmartGridComm)*, pp. 1–7, 2020.

[18] M. Zhang, G. Guo, and S. Magnússon, "Data-driven decentralized control of inverter-based renewable energy sources using safe guaranteed multi-agent deep reinforcement learning," *IEEE Trans. Sustainable Energy*, vol. 15, no. 2, 2024.

[19] R. Koenker, *Quantile Regression*. Cambridge, U.K.: Cambridge Univ. Press, 2005.

[20] G. Chen, H. Zhang, and H. Hui, "Deep-quantile-regression-based surrogate model for joint chance-constrained optimal power flow with renewable generation," *IEEE Trans. Sustainable Energy*, vol. 14, no. 1, 2023.

[21] Y. Chen, Y. Shi, and B. Zhang, "Optimal power flow with chance constraints: A conformal prediction approach," *IEEE Trans. Power Syst.*, vol. 38, no. 2, pp. 1193–1206, Mar. 2023.

[22] L. Thurner, A. Scheidler, F. Schäfer, et al., "pandapower – An open-source Python tool for convenient modeling, analysis, and optimization of electric power systems," *IEEE Trans. Power Syst.*, vol. 33, no. 6, pp. 6510–6521, Nov. 2018.

[23] D. Das, "Optimal placement of capacitors in radial distribution system using a fuzzy-GA method," *Int. J. Elect. Power Energy Syst.*, vol. 30, no. 6–7, pp. 361–367, 2008.

[24] C. Zhao, G. Li, and L. Zhang, "A regulation method considering coupled voltage violation and unbalance in three-phase four-wire hybrid AC/DC LVDNs with high penetration PVs," *IEEE Trans. Sustainable Energy*, vol. 15, no. 2, 2024.