Supplementary Information

Objective-Driven Efficient Sparse Metering for Machine Learning-Based Fault Diagnosis and Self-Healing Restoration in Active Distribution Networks

Omar Gamal Abdelmohsen, Tamer Amin Kawady, Tamer Fatouh Atia, Eman Salah Ali

This file contains four supplementary figures and seven supplementary tables referenced in the main article. Together they document the dataset construction, the model configurations, the evaluation-protocol ablation, the instrument-error study, the agreement of three independent estimates of the ± 1-hop accuracy, the meter-count sensitivity at selected budgets, and the localization accuracy by fault resistance and fault type.

Supplementary Figures

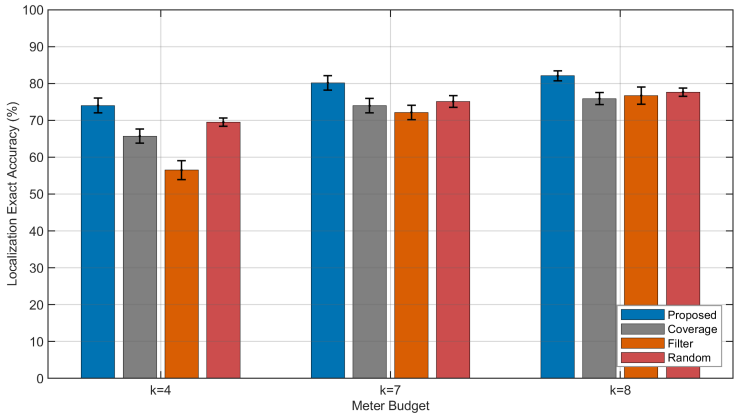


Figure S1. Placement methods at three meter budgets (mean $\pm$ s.d., five seeds). The objective-driven wrapper leads at all three budgets, including the seven-meter operating point, with the widest margin at $k = 4$.

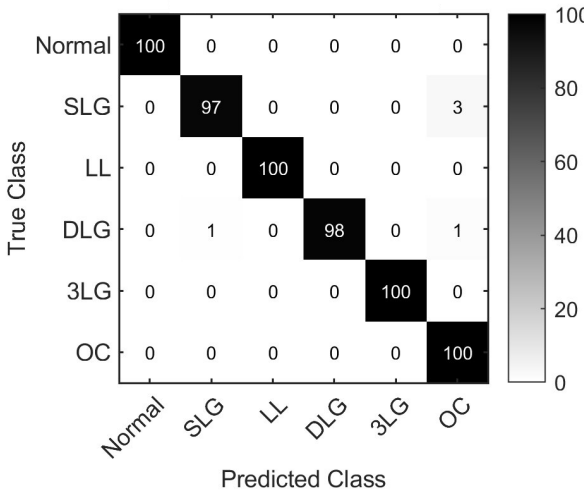


Figure S2. Six-class confusion of the deployed random forest at the three-meter classification placement, row-normalized. Errors occur only for ground faults: 3% of SLG faults are read as open-conductor (OC) events and 1% of DLG faults as SLG or OC; every other class is recognized without error.

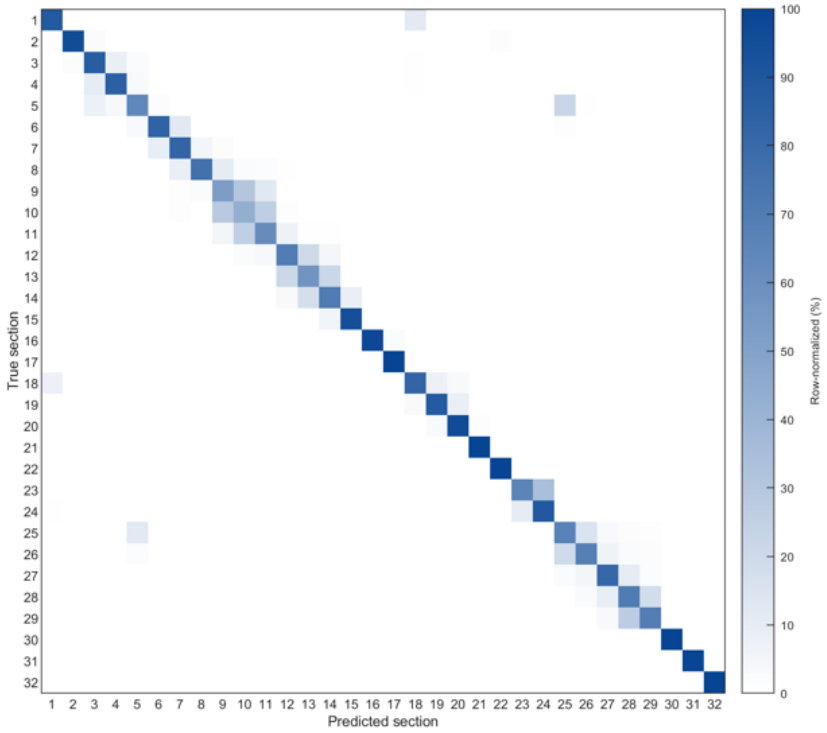


Figure S3. Section-identification confusion of the deployed random forest at the cost-optimal placement ($k = 7$), row-normalized, on held-out events aggregated over five seeds. Most errors fall on adjacent sections, which is what the ± 1-hop accuracy reported in the main article records.

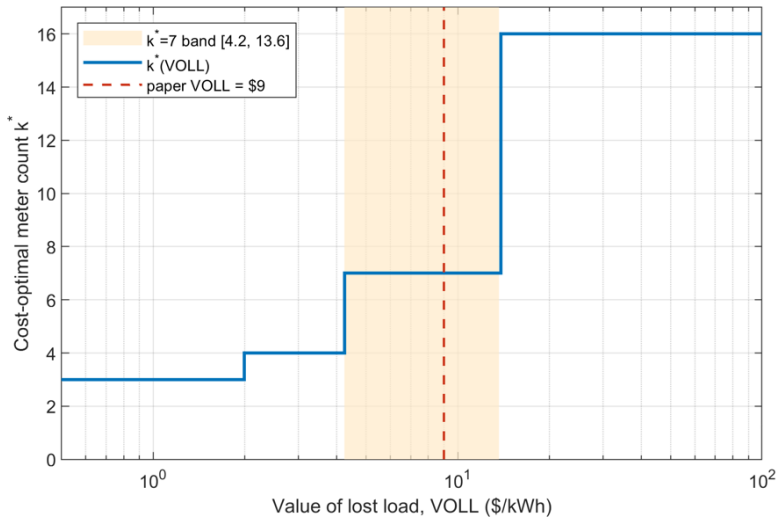


Figure S4. Robustness of the cost-optimal meter count $k^\star$ to the assumed value of lost load. The shaded band marks the \$4.2–\$13.6/kWh range over which the point minimum stays at seven meters, bracketing the assumed \$9/kWh; outside it the minimum moves to four meters below and to sixteen above, reflecting how little the total cost varies across that range.

Supplementary Tables

Table S1. Dataset composition. Each shunt variant contributes one sample for every combination of the 32 faulted sections, twelve fault resistances (0.001–100 Ω, logarithmic) and five loading levels, that is $32 \times 12 \times 5 = 1{,}920$ samples. Open-conductor events carry no fault resistance and are recorded with $R_f = 0$; each section and loading level holds three open-phase variants with four measurement-noise replicas, which again gives 1,920 samples. The 990 normal-operation samples cover 33 loading levels, five source voltages and six noise replicas, with one event per loading level. The distinct (section, resistance, loading) points therefore number $32 \times 12 \times 5 + 32 \times 5 = 2{,}080$, with no case lost to non-convergence. Grouping reduces these to 1,728: the lowest bin of the grouping key holds $R_f = 0$, 0.001 and 0.0028 Ω, which accounts for 320 of the reduction, and a further 32 points, 1.5% of the sweep, share an event index through the same quantization. Merging is never looser than splitting, so the protocol remains conservative. With the 33 normal-operation events the held-out total is 1,761.

Class	Variants	Cases/variant	Samples
Normal	—	—	990
SLG	3 (a,b,c–g)	1,920	5,760
LL	3 (ab,bc,ca)	1,920	5,760
DLG	3 (ab,bc,ca–g)	1,920	5,760
3LG	1	1,920	1,920
OC	3 (a,b,c)	640	1,920
Total			**22,110**

Table S2. Model configurations and tuning. Hyperparameters were selected by random search followed by Bayesian optimization on the training partition only.

Model	Configuration	Selection
RF (cls.)	100 trees, min leaf 3, 5 predictors/split ($\lfloor\sqrt{p}\rfloor$, $p = 30$)	Bayesian search
GBM	AdaBoost.M2, 400 cycles, rate 0.1, $\leq$64 splits/tree	Grid over rate/depth
SVM	RBF kernel, $C = 50$, kernel scale 5.48, one-vs-one ECOC	Random + Bayesian
ANN	MLP, hidden layers 256–128–64, SMOTE-balanced	Random search
RF (loc.)	100 trees, min leaf 3, $\lfloor\sqrt{p}\rfloor$ predictors/split, 32-class output	Reuses the tuning above

Table S3. Effect of the evaluation protocol. Top: classification (random forest, three meters); the bias column is the optimistic offset of each protocol relative to the event-held-out reference, and both offsets are smaller than the seed-to-seed spread of the reference (± 0.41). Bottom: localization (random forest, seven meters, five seeds); a random split over samples inflates the exact accuracy by 10.7 pp (paired $p = 0.0002$), whereas holding out whole buses tests spatial extrapolation.

Protocol	**Macro-F_1 (%)**	κ	**Bias (pp)**
Random split (70/15/15)	98.98	0.987	$+0.27$
10-fold CV	99.04 ± 0.21	0.988	$+0.33$
Event-held-out (used here)	98.71 ± 0.41	0.985	— (reference)

Protocol (localization)	**Exact (%)**	**± 1-hop (%)**	**MAE (hop)**
Random sample split	90.73 ± 0.17	98.83 ± 0.15	0.257 ± 0.002
Event-held-out (used here)	80.00 ± 1.84	97.33 ± 0.40	0.320 ± 0.015
Location-held-out (whole buses)	—	—	0.622 ± 0.155

Table S4. Robustness to instrument errors (three meters, accuracy in %, mean $\pm$ s.d. over five seeds, ten error draws per condition). Every condition is added to the base measurement-error model of the dataset. Ratio errors and offsets are fixed per channel, on top of the random base gain error; quantization rounds each channel over its range across the dataset, whereas the base model quantizes over each sample's full feature range. Bold marks the better of the two augmented ensembles. The GBM clean-trained accuracy is the five-seed event-split value from the classification study, a different split realization from the other cells; the RF augmentation cost is measured within a single run, the GBM cost across the two. The last row is the clean-data accuracy given up in exchange for noise robustness, which is what separates the two models.

	RF		**GBM**	
Condition	Clean-trained	Augmented	Clean-trained	Augmented
No added error	99.5 ± 0.2	$\mathbf{99.1 \pm 0.4}$	99.9 ± 0.04	96.7 ± 0.4
Ratio error $\pm 0.5\%$	97.7 ± 0.6	97.5 ± 0.7	—	—
Ratio error $\pm 1\%$ (class 1)	96.1 ± 0.8	$\mathbf{96.3 \pm 0.5}$	—	94.4 ± 0.7
Ratio error $\pm 2\%$	93.8 ± 1.3	94.3 ± 1.0	—	—
Ratio error $\pm 5\%$	87.8 ± 2.5	$\mathbf{91.9 \pm 0.9}$	—	91.0 ± 0.7
Offset $\pm 0.5\%$ of channel s.d.	95.4 ± 0.7	95.7 ± 0.6	—	—
Offset $\pm 1\%$ of channel s.d.	93.2 ± 0.8	94.1 ± 0.5	—	—
Offset $\pm 2\%$ of channel s.d.	88.7 ± 1.2	$\mathbf{92.0 \pm 0.4}$	—	91.1 ± 0.8
Quantization, 12-bit	99.0 ± 0.4	98.5 ± 0.4	—	—
Quantization, 10-bit	97.6 ± 0.8	97.5 ± 0.9	—	—
Quantization, 8-bit	91.0 ± 1.2	$\mathbf{94.0 \pm 0.7}$	—	93.4 ± 0.6
Combined + 30 dB AWGN	80.8 ± 1.3	$\mathbf{90.8 \pm 0.6}$	—	90.1 ± 0.7
Cost of augmentation	-0.4 pp		-3.2 pp	

Table S5. Exact and ± 1-hop accuracy at the seven-meter operating point from three runs that share the data, the model and the event-held-out protocol but differ in their random splits. The spreads between runs, 0.4 pp in exact accuracy and 0.3 pp in ± 1-hop accuracy, are within the seed spread of a single run and move SAIDI by under 0.01 h. The reliability chain uses the meter-count sweep.

Run	**Exact (%)**	**± 1-hop (%)**	**Role in the paper**
Placement comparison	80.00 ± 1.84	97.33 ± 0.40	reported with the distance MAE
Meter-count sweep	80.1 ± 2.0	97.38 ± 0.48	drives SAIDI, the cost curve and the economics
Learner comparison	80.4 ± 1.2	97.6 ± 0.6	compares RF, ANN, GBM and SVM on one footing

Table S6. Meter-count sensitivity at selected budgets (five seeds; the cost spread is across seeds). Because only the ±1-hop criterion enters the reliability model, the cost follows the ±1-hop column rather than the exact-accuracy column: from seven to sixteen meters the exact accuracy gains 12.8 points while ±1-hop gains 1.9, and the total cost does not fall. The two bold budgets cannot be separated by a paired test over the shared seeds ($p = 0.084$); after a Holm correction for fourteen comparisons, every other budget costs more.

Meters k	**Exact (%)**	±1**-hop (%)**	**SAIDI (h)**	**Total cost ($/yr)**
2	64.0 ± 1.5	86.1 ± 1.0	0.792	$27{,}089 \pm 833$
3	69.7 ± 1.9	93.8 ± 0.4	0.602	$20{,}993 \pm 307$
4	74.0 ± 2.0	95.3 ± 0.4	0.566	$20{,}041 \pm 325$
6	77.9 ± 2.0	96.4 ± 0.3	0.540	$19{,}698 \pm 254$
7	$\mathbf{80.1 \pm 2.0}$	$\mathbf{97.4 \pm 0.5}$	**0.515**	$\mathbf{19{,}142 \pm 393}$
8	$\mathbf{82.1 \pm 1.3}$	$\mathbf{97.5 \pm 0.3}$	**0.511**	$\mathbf{19{,}275 \pm 280}$
12	87.4 ± 1.6	98.2 ± 0.4	0.496	$19{,}822 \pm 350$
14	89.0 ± 1.7	98.5 ± 0.4	0.489	$20{,}118 \pm 308$
16	92.9 ± 1.0	99.3 ± 0.2	0.468	$19{,}964 \pm 181$

Table S7. Localization accuracy at the seven-meter operating point by fault resistance and by fault type (event-held-out, five seeds pooled, 32,020 test predictions). Shunt faults are swept over twelve resistances; open-conductor (OC) events carry no fault resistance. The OC rows rest on 480 distinct scenarios (32 sections, five loadings, three open phases), each with four measurement-noise replicas.

R_f (Ω)	n	**Exact (%)**	±1**-hop (%)**	**Type**	n	**Exact (%)**	±1**-hop (%)**
0.001	2,500	98.44	100.00	SLG	8,706	85.38	99.09
0.0028	2,500	98.48	100.00	LL	8,706	86.62	99.04
0.0081	2,400	99.13	100.00	DLG	8,706	85.76	99.14
0.023	2,450	99.02	100.00	3LG	2,902	74.22	91.87
0.066	2,460	99.27	100.00	OC	3,000	34.10	87.33
0.19	2,660	97.97	100.00	**All**	32,020	80.00	97.33
0.53	2,320	89.22	99.83				
1.52	2,300	79.13	97.78				
4.33	2,280	74.87	98.03				
12.3	2,460	67.40	97.11				
35.1	2,200	54.73	94.86				
100	2,490	54.58	92.37				
OC (no R_f)	3,000	34.10	87.33				