


Supplementary Information for DF-PLNet: Budgeted Periodic Filtering for Industrial Control System Process-Record Diagnostics

Yipeng Zhang, Jiahui Chen, Yue Li, and Yun Shi

Supplementary Figure S1: Passive observation setting



Supplementary Figure S1. Passive observation point for DF-PLNet in a layered ICS network. The monitor observes chronological process-record streams for offline residual-stream diagnostics without changing plant operation.

Supplementary Note 1: Relation to official SWaT and HAI benchmark protocols

This supplementary note summarizes how the processed-record protocol evaluated in the paper differs from common usage of the official SWaT and HAI benchmarks. The comparison defines the experimental context for the reported numbers. Supplementary Table S1 lists the differences.

The table also sets the comparison boundary for the paper: model families, PFM budgets, and coverage audits are evaluated under the same processed-record protocol, while native-feature benchmark studies use different inputs, splits, and evaluation sets.

Supplementary Table S1. Differences between common official-benchmark usage and the processed-record protocol in this paper.

Aspect	Common official-benchmark usage	This paper
Input features	Native multivariate channels (51 SWaT and 86 HAI numeric process variables), often modeled per channel	One scalar per record: mean absolute z-score across numeric columns with training-normal statistics, plus the record timestamp
Training labels	Often unsupervised or semi-supervised training on attack-free normal data	Supervised training; training partitions include labeled attack records
Split definition	Released train/test files used as provided	Custom chronological splits: SWaT early records (including early attacks) for training and tail records for test; HAI hai-train1-hai-train4 plus hai-test1 for training and hai-test2 for test
Threshold selection	Varies across studies; sometimes tuned with test-side information or point-adjust protocols	Selected on training scores and labels by F1 maximization, then frozen for the test partition
Evaluation set	All test records	Complete-window value positions; PFM-enabled rows evaluate retained subsets and report the coverage-aware audit
Evaluation granularity	Often event-level or point-adjusted metrics	Record-position metrics without point adjustment
Deployment mode	Online or streaming detection commonly targeted	Offline batch diagnostics within already-collected partitions

Note. The processed-record protocol supports conclusions about budgeted residual-stream diagnostics on released records.

Supplementary Note 2: Baseline input scope

Supplementary Table S2 records the input requirements that determine which detector families are comparable in the released process-record setting.

Supplementary Table S2. Input requirements of detector families in the released process-record setting.

Family	Representative	Required inputs	Status in this setting
Payload deep detectors	Liu et al.	Packet payload bytes	Input unavailable
Protocol-field and sequence-aware IDS	Caselli et al.; DPI	Parsed protocol fields, function codes	Input unavailable
Invariant and process-semantics methods	Feng et al.	Multivariate process variables, system semantics	Input unavailable
Native-feature stream ensembles	Kitsune	Packet-level statistical feature vectors	Input unavailable
Unsupervised time-series detectors	USAD; LSTM	Generic time series	Requires matched adaptation to the supervised record-position protocol
Supervised record-position scorers	CNN; Transformer; DF-PLNet variants	Value/interval stream	Compared in the main-text diagnostic results tables

Note. Status is defined relative to the shared value/interval input stream, supervised record-position scoring task, and train-threshold protocol used in this study.

Supplementary Note 3: Representation sanity check

Supplementary Table S3 reports the SWaT no-PFM sanity check that replaced the mean-absolute-z scalar with the raw FIT101 feature, a single flow-sensor channel.

Supplementary Table S3. SWaT representation sanity check without PFM.

Input value	AUC	PR-AUC	F1	Recall	FPR	MCC
Mean-absolute-z scalar	0.9671±0.0009	0.9696±0.0021	0.8866±0.0067	0.7993±0.0124	0.0036±0.0022	0.8117±0.0081
FIT101 only	0.7819±0.0044	0.7320±0.0045	0.7770±0.0104	0.7941±0.0222	0.2494±0.0082	0.5454±0.0158

Note. Values are mean $\pm$ s.d. over seeds 777, 778, and 779 under the same no-PFM protocol. The FIT101-only run is a diagnostic check, not a claim that FIT101 is the strongest single-feature baseline.

Supplementary Note 4: PFM budget sensitivity

Supplementary Table S4 reports the 0 %, 5 %, 10 %, and 20 % PFM budget settings. The 0 % setting reuses DF-PLNet without PFM, and the 10 % setting reuses the main DF-PLNet setting.

Supplementary Table S4. PFM budget sensitivity on SWaT and HAI.

Dataset	Budget	Eval. records	AUC	F1	Recall	FPR	Speed
SWaT	0 %	54,600	0.9671±0.0009	0.8866±0.0067	0.7993±0.0124	0.0036±0.0022	97132
SWaT	5 %	51,800	0.9687±0.0039	0.8959±0.0037	0.8159±0.0081	0.0056±0.0028	23338
SWaT	10 %	49,000	0.9505±0.0150	0.8966±0.0122	0.8260±0.0415	0.0157±0.0252	67248
SWaT	20 %	43,600	0.9638±0.0025	0.8713±0.0300	0.7758±0.0452	0.0041±0.0029	103985
HAI	0 %	230,200	0.6466±0.0004	0.3118±0.0106	0.2300±0.0165	0.0093±0.0018	79883
HAI	5 %	218,800	0.6471±0.0008	0.3061±0.0054	0.2562±0.0020	0.0158±0.0012	37027
HAI	10 %	207,200	0.6486±0.0029	0.3112±0.0091	0.2452±0.0051	0.0126±0.0022	80293
HAI	20 %	184,200	0.6418±0.0073	0.1320±0.1653	0.0948±0.1286	0.0063±0.0036	56631

Note. Values are mean $\pm$ s.d. over seeds 777, 778, and 779. Eval. records are retained complete-window value positions after pruning and tail exclusion. Budgets denote maximum PFM drop ratios. Speed is reported in evaluated record positions per second.