

(RESEARCH ARTICLE)



LABEL RELIABILITY BEFORE LABEL ACQUISITION: A COST-ACCOUNTED STUDY OF NEURAL OFDM ADAPTATION

Kabil Abderrahim ¹, Hudhaifa Mohammednawwar Qasim Alajar ² and Abdullah Al Mamun Sadi ^{2,*}

¹ School of Artificial Intelligence, Nanjing University of Information Science and Technology, 219 Ningliu Road, Nanjing, Jiangsu 210044, China.

² School of Electronics and Information Engineering, Nanjing University of Information Science and Technology, 219 Ningliu Road, Nanjing, Jiangsu 210044, China.

* Corresponding Author

ORCID Details

Kabil Abderrahim: <https://orcid.org/0009-0003-0509-961X>

Hudhaifa Mohammednawwar Qasim Alajar: <https://orcid.org/0009-0005-3422-2589>

Abdullah Al Mamun Sadi: <https://orcid.org/0009-0008-8945-3626>

Global Journal of Engineering and Technology Advances, 2026, 29(01), 205–214

Publication history: Received on 07 September 2026; revised on 12 October 2026; accepted on 14 October 2026

Article DOI: <https://doi.org/10.30574/gjeta.2026.29.1.0294>

Copyright © 2026 Author(s) retain the copyright of this article. This article is published under the terms of the Creative Commons Attribution License 4.0

ABSTRACT

Objective: Labels recovered through forward error correction (FEC) can adapt a neural orthogonal frequency-division multiplexing (OFDM) receiver without dedicated training symbols, but an oracle-label gap does not by itself justify transmitting new training frames. We ask whether the gap reflects label correctness, observation selection or label quantity, and whether paid known-payload probes remain worthwhile after full radio costs are charged.

Methodology: A coded 64-subcarrier OFDM link with a pretrained recurrent receiver was simulated. A development experiment (32 episodes) separated the three factors at matched update effort. A frozen confirmation (24 episodes at each of two nominal E_b/N_0 values, 7.25 and 8 dB) compared a codeword-level zero-syndrom mask, batch-level FEC adaptation, known-payload probes and conventional receivers by error-free user bits per transmitted sample, charging pilots, cyclic prefixes, probe frames and signaling.

Key findings: The oracle gap was chiefly a label-correctness effect. The mask, needing no new transmissions, improved efficiency over matched-update FEC adaptation by 0.0419 and 0.0352 bit/sample (95% intervals 0.0395–0.0443 and 0.0326–0.0378). Probes slightly lowered block error rate but lost net efficiency over the registered horizon. Charging measured serialized processing time put the mask below the frozen receiver and a colored-noise conventional receiver in throughput.

Conclusion: Ordinary labels should be filtered at codeword level before new training transmissions are purchased. This empirical control result is conditional on one checkpoint and a stable macro channel model, and does not imply that probes are unnecessary at longer horizons or under other impairments.

Keywords: Neural Receiver; OFDM; Online Adaptation; Recovered Labels; Supervision Cost.

1 INTRODUCTION

1.1 Background and motivation

Receiver-side label recovery can adapt a trained detector without dedicated training symbols: the decisions of the forward error correction (FEC) decoder serve as recovered labels for the neural detector in an orthogonal frequency-division multiplexing (OFDM) receiver. It can also produce an ambiguous diagnosis: a receiver improves with perfect simulator labels, but that improvement may reflect incorrect recovered labels, a biased observation set, or insufficient label quantity. These explanations imply different engineering responses. Buying new observations is justified only after useful ordinary supervision and its processing have been considered.

This paper studies that decision in one deliberately controlled OFDM fixture. Its contribution is an evidence chain connecting a matched-effort oracle decomposition to a deployable, zero-extra-frame control and a complete-frame cost comparison. The experiments are separated into development and fresh confirmation. We neither introduce FEC adaptation nor claim that reliability filtering itself is new. The result instead identifies which comparison changes the engineering conclusion in this implementation.

The original selection explanation does not survive the registered development test at the chosen partially eligible point. The subsequent study therefore tests codeword granularity before adding supervision. This negative branch matters: an oracle-trained model is not a deployable baseline, and a post-adaptation error improvement is not a net communication benefit.

1.2 Related work

Schibisch et al. recover training labels through error-correcting codes and already compare corrupted, correctly received and corrected observations using equal-size training sets [1]. Fischer et al. apply recovered labels to recurrent OFDM adaptation and explicitly discuss erroneous labels and conservative eligibility [2]. Our decomposition is thus a controlled application of established questions, not their first formulation.

Confidence-based symbol selection balances reliability and diversity in Finish et al. [3]. Teng and Chen use code constraints in syndrome losses and discuss the limitations of imperfect label recovery [4]. The hard codeword mask used here is a transparent baseline alongside these ideas; we do not compare its performance across their different waveforms.

Pilot acquisition and adaptation cost are also established concerns. Park et al. combine few-pilot meta-learning with adaptive pilot quantity [5]. Cohen et al. actively choose simulator channel tasks for meta-learning, which differs from acquiring new on-air observations [6]. Uzlaner et al. reduce retraining through modular drift detection [7]. Obeed and Jian redesign demodulation reference signals (DMRS) to support continual OFDM learning and consider parallel and forward-pass-reuse architectures [8]. Our fixed waveform and paused adaptation do not implement those service-continuity mechanisms. Neural and pilotless waveform design can also change throughput [9]; here the waveform and ordinary pilot allocation remain fixed.

The specific distinction is a fresh-episode test of whether the measured FEC oracle gap still motivates paid supervision after a free codeword control, with explicit user-payload accounting and both white- and colored-noise a posteriori probability (APP) receivers. This is a narrow empirical contribution, not a claim of priority for cost-aware learning.

1.3 Objective of the study

The objective is to determine, in one controlled fixture, (i) whether the measured gap between recovered-label and oracle-label adaptation is attributable to label correctness, observation selection or label quantity; (ii) whether a codeword-level zero-syndrome mask, which requires no new transmissions, improves complete-stream radio efficiency over batch-level FEC adaptation; and (iii) whether paid known-payload probes remain worthwhile once their transmission, setup and signaling costs are charged. The hypothesis for (ii), frozen before the final data were generated, was a practical gain of at least 0.01 bit/sample.

2 MATERIALS AND METHODS

2.1 Simulation environment and resources

All results come from numerical simulation; no chemicals, biological materials or physical radio equipment were used. The executed environment comprised Python 3.12.14, PyTorch 2.9.1 (CUDA 12.8), Sionna-no-rt 2.0.1, NumPy 2.5.2, SciPy 1.18.1 and Matplotlib 3.11.1, installed from official prebuilt PyPI and PyTorch wheels, on a local workstation with an NVIDIA RTX 5070 graphics processor; Sionna supplies the 5G low-density parity-check (LDPC) encoder. All timing results refer to this device. Every neural arm starts from one validated universal checkpoint, so results are conditional on a single training seed.

2.2 System and channel model

After the fast Fourier transform (FFT), the simulator implements the diagonal model

$$Y_{t,k} = H_{t,k}X_{t,k} + W_{t,k} \quad (1)$$

where Y , H , X and W are the received symbol, channel coefficient, transmitted symbol and noise at OFDM symbol index t and subcarrier index k . Each frame has 36 OFDM symbols and 64 subcarriers, unit-energy 16-point quadrature amplitude modulation (16-QAM), and 36 pilots at time indices $\{0, 15, 30\}$ and frequency indices $\{2, 7, \dots, 57\}$. The remaining 2268 resource elements carry seven rate-1/2 LDPC codewords, each with 648 information and 1296 coded bits. The implementation uses the Sionna 5G LDPC encoder and 20 belief-propagation (BP) iterations. No cyclic redundancy check (CRC) is modeled. Block error rate (BLER) means any information-bit error in the complete 4536-bit user frame, rather than in one constituent codeword.

The sample rate is 10 MHz, the useful transform is 64 samples and the cyclic prefix (CP) is 16 samples. One frame therefore costs 2880 complex samples and lasts 288 μ s. The channel uses eight taps separated by 100 ns, an exponential power profile whose last-to-first ratio is -13 dB, and a normalized sum of eight random-phase sinusoids per tap. Doppler follows 100 km/h at 5.9 GHz. This finite-sinusoid model approximates fading; we do not assume exact Gaussian tap statistics. Fading initialization is independent between frames.

Nominal $g = E_b/N_0$ in dB sets $N_0 = 1/(2 \cdot 10^{g/10})$ for data symbols. This axis does not absorb pilot or CP energy. The side-noise condition multiplies the noise variance by $10^{0.6}$ on the first four and last four subcarriers. The disturbance is Gaussian, not an explicit adjacent transmitter. Within-symbol Doppler inter-carrier interference (ICI), synchronization error and radio-frequency (RF) nonlinearities are absent. These are implementation assumptions, not an exact replication of the simulator in [2].

The recurrent detector processes received samples, the known pilot grid, sparse least-squares (LS) pilot information and configured nominal noise. It has three bidirectional long short-term memory (LSTM) stages scanning frequency, time and frequency, with 64 units per direction, followed by an eight-unit rectified-linear (ReLU) layer and four bit logits. All arms start from the same validated 100,000-update checkpoint. Universal training used batch size 128, a nominal signal-to-noise ratio (SNR) of 8–30 dB and velocity 0–200 km/h; tap counts broadened from 4–10 to 1–14 halfway through training. Inference at 7.25 dB is consequently below the universal training SNR range.

2.3 Problem formulation

Let $\hat{c}_{b,j}$ denote BP hard coded-bit decisions for codeword j in batch b , and $s_{b,j}$ the fraction of satisfied mother-code checks. With S_b the mean of $s_{b,j}$ over all codewords in batch b , the inherited batch eligibility rule [2] is

$$e_b = 1\{g > 7\}1\{S_b \geq 0.82\} \quad (2)$$

where $1\{\cdot\}$ is the indicator function. Eligibility is a real syndrome calculation, not a comparison to transmitted bits. A passing batch can contain incorrect labels. The simulator's ordinary payload and realized channel are excluded from deployable label decisions.

Our primary utility is the number of error-free user bits per transmitted complex sample,

$$G = \frac{4536 \sum_{i \in U} 1\{\hat{u}_i = u_i\}}{2880(N_{user} + N_{known} + N_{control})} \quad (3)$$

where U is the set of user frames, u_i and $\hat{u}_i$ are the transmitted and detected information bits of frame i , and N_{user} , N_{known} and $N_{control}$ count user frames, known-payload probe frames and control (setup and request) frames. In the registered

protocol $N_{user} = 2600$ (1600 collection and 1000 post-adaptation frames). Truth enters this evaluator only. The numerator counts entirely error-free payloads; it is not a claim of CRC-certified delivery. Pilots and CP are included in every denominator frame. Known-payload transmissions never contribute user bits and never reveal labels of earlier failed observations.

2.4 Experimental procedures

2.4.1 Development decomposition

An eligibility-only diagnostic fixes the operating point without comparing adapted methods. The subsequent decomposition uses 32 independent episodes, each collecting 32 batches of 50 frames and evaluating 1000 fresh frames. The four adapted arms are: recovered labels on eligible batches; correct oracle labels on exactly those batches; correct labels on an independently uniform count-matched subset; and correct labels on all candidate batches. All use 32 updates, with complete coverage before repeated permutations. The first pair isolates correctness, the second selection at matched quantity, and the third available quantity at fixed update effort. Repetition is part of the last estimand. A fourth contrast, labeled adaptation, compares recovered-label adaptation with the frozen receiver. A frozen recurrent neural network (RNN), APP and a qualified soft iterative estimation, demapping and decoding (IEDD) control [2] are also evaluated. Oracle arms are nondeployable diagnostics throughout.

2.4.2 Adaptation controls

The MASK arm sets $m_{b,j} = 1\{s_{b,j} = 1\}$ and minimizes

$$L_b = \frac{\sum_{j,q} m_{b,j} \ell(z_{b,j,q} \hat{c}_{b,j,q})}{\sum_{j,q} m_{b,j}} \quad (4)$$

where ℓ is binary cross-entropy (BCE), $z_{b,j,q}$ are the detector bit logits and q indexes transmitted code bits. Empty masks skip the update. The rule uses each ordinary batch once and no new radio frames. Zero syndrome is a constraint, not guaranteed correctness. MASK-ORACLE uses the same masks with true labels, while ALL-ORACLE uses all ordinary labels. Neither is a deployable method.

FEC-ONCE trains once on each eligible batch. FEC-REPLAY uses those same labels for 32 updates. Three paid controls acquire 50-frame bursts: random chooses eight of 32 collection slots uniformly; periodic uses every fourth slot; rejection requests a burst after each rejected batch until eight bursts. Each transmitted payload is independently generated from a shared transmitter/receiver seed. A probe arm with acquired bursts trains only on those new frames for 32 updates. A rejection arm requesting no burst falls back to FEC replay. Different realized costs are reported rather than treated as a matched-cost selection advantage.

The Adam optimizer uses learning rate 0.001, moments 0.9/0.999 and epsilon 10^{-7} in full precision. All neural arms score collection frames with the frozen receiver before a single adaptation round. Probe requests use only already observed outcomes. Evaluation follows adaptation; no future observations inform it. The radio is paused during adaptation. This is causal adaptation within stable macro conditions, not demonstrated continuous channel tracking.

2.4.3 Conventional controls

The APP receiver uses pilot linear minimum mean-square error (LMMSE) estimates with configured power delay profile (PDP), Doppler and nominal noise. A stronger colored APP also knows the configured side-noise profile, not realized channel or noise values. It uses that covariance in estimation and demapping. These are favorable statistical knowledge assumptions and must not be confused with learned adaptation. The validated complex Wiener orientation and channel-estimation error variance are retained. Soft IEDD uses four five-iteration decoding/refinement stages; its Gaussian soft-symbol approximation and dense computation limit its status as a strong comparator. It remains a secondary control on 100 post frames per final episode, whereas both APP receivers use every ordinary evaluation frame.

2.5 Experimental design and statistical analysis

The final protocol freezes two nominal SNRs, 7.25 and 8 dB, with all other physical parameters fixed. Each condition has 24 fresh independent adaptation episodes, 1600 ordinary collection frames and 1000 post frames per episode; episode seeds derive from two fixed seed bases (96,000,000 and 97,000,000) plus 10,000 times the episode index. Every probe strategy pays up to 400 known frames, one full setup frame if used, and one reverse request frame per triggered burst. Control reliability is assumed ideal. No retransmission mechanism is simulated.

The two primary episode-paired contrasts compare the mask and FEC replay in G , using Bonferroni-adjusted two-sided Student intervals with familywise 95% coverage for that pair of tests. A practical gain requires a mean of at least 0.01 bit/sample and a lower bound above zero. Planning inflated the larger development paired standard deviation to 0.010;

24 episodes give approximately 0.0049 bit/sample half-width under that assumption. Sample size was not changed after final observations. Episodes, not coded bits, are the inferential units.

Fourteen comparisons to the other seven realistic controls form a separate secondary family. Four oracle BLER contrasts and six probe BLER contrasts form separately adjusted diagnostic families. Marginal arm intervals are descriptive. These are separate families, not one joint confidence guarantee. Results condition on one universal training seed. All exact seeds, protocols, source hashes and per-frame error indicators are provided in the reproduction package, which is available from the corresponding author on request.

3 RESULTS AND DISCUSSION

3.1 Development: availability versus reliability

The eligibility screen selected 7.25 dB with 9 of 12 batches eligible; all 12 batches at 8 dB passed. A nominal 16-tap diagnostic was invalid because of a 14-tap implementation cap and is excluded, without changing selection. At the frozen partially eligible point, fresh decomposition episodes had 12–25 eligible batches out of 32. Fig. 1 shows the four registered contrasts. Correcting the labels yields a 7.85-point BLER reduction, whereas the upper interval bounds for selection and quantity are below 0.23 percentage points. This excludes their registered two-point benefit in this fixture, not exact equality or a universal selection effect.




Figure 1: Development decomposition at 7.25 dB: BLER reduction with familywise 95% intervals (32 episodes); dashed line, practical threshold.

3.2 Frozen confirmation: complete-stream radio efficiency

Table 1 lists the complete-stream efficiency G and the post-adaptation BLER of every arm, and Table 2 gives the paired contrasts that carry the inference. In the primary comparison, the mask improved on matched-update FEC replay at both SNRs and met the registered practical criterion in both cases (Table 2, first row; Fig. 2). All fourteen secondary comparisons favored the mask with adjusted lower bounds above zero.

Table 1: Mean complete-stream efficiency G and post-adaptation BLER (24 episodes per SNR).

	G (bit/sample)		Post-adaptation BLER	
	7.25 dB	8 dB	7.25 dB	8 dB
Arm				
Frozen	0.4054	0.5834	0.7409	0.6330
FEC once	0.4281	0.6313	0.7033	0.5540
FEC replay	0.4371	0.6312	0.6885	0.5541
Mask	0.4790	0.6664	0.6193	0.4960
Mask oracle	0.4790	0.6664	0.6193	0.4960
All oracle	0.4827	0.6693	0.6132	0.4912
Random probes	0.4180	0.5799	0.6135	0.4912
Periodic probes	0.4178	0.5791	0.6140	0.4927
Rejection probes	0.4176	0.6329	0.6135	0.5487
White APP	0.2671	0.4080	0.8291	0.7427
Colored APP	0.4071	0.5746	0.7401	0.6366

Note: FEC, forward error correction; APP, a posteriori probability receiver; BLER, block error rate.

Table 2: Paired difference in G, mask minus comparator (bit/sample), with Bonferroni-adjusted 95% intervals.

Comparator	Family	7.25 dB	8 dB
FEC replay	Primary	0.0419 [0.0395, 0.0443]	0.0352 [0.0326, 0.0378]
Frozen RNN	Secondary	0.0737 [0.0694, 0.0779]	0.0830 [0.0788, 0.0872]
FEC once	Secondary	0.0509 [0.0468, 0.0550]	0.0351 [0.0321, 0.0382]
White APP	Secondary	0.2119 [0.2059, 0.2179]	0.2584 [0.2512, 0.2656]
Colored APP	Secondary	0.0719 [0.0673, 0.0766]	0.0918 [0.0853, 0.0983]
Random probes	Secondary	0.0610 [0.0589, 0.0631]	0.0865 [0.0847, 0.0882]
Periodic probes	Secondary	0.0612 [0.0593, 0.0631]	0.0873 [0.0854, 0.0891]
Rejection probes	Secondary	0.0615 [0.0595, 0.0635]	0.0335 [0.0289, 0.0382]

Note: RNN, recurrent neural network; FEC, forward error correction; APP, a posteriori probability receiver.




Figure 2: Paired improvement of the mask over FEC replay: episodes (dots) and familywise 95% intervals; dashed line, 0.01 bit/sample threshold.

Dedicated probes achieved slightly lower post-adaptation BLER than the mask in most comparisons (Table 1; Fig. 3), but this did not translate into net efficiency. Each probe strategy transmits up to 400 known frames plus a setup frame, none of which carries user bits, and the mask retains a clear advantage in G over all three probe strategies (Table 2). At 7.25 dB the rejection-triggered strategy requested nearly all of its eight bursts. At 8 dB it requested only 6.25 known frames on average, so it behaves largely like FEC replay and shows neither the BLER gain nor the cost of the other two (supplementary material). Better post-adaptation BLER alone therefore need not improve net utility.




Figure 3: Post-adaptation BLER versus complete-stream efficiency for selected arms (episode means).

Among the conventional receivers, the white-noise APP performed worst, whereas the colored APP, which exploits knowledge of the side-noise profile, performed close to the frozen RNN; the mask exceeded both in G at both SNRs. The

colored APP therefore materially changes the conventional comparison and is retained despite its favorable knowledge assumption.

3.3 Label reliability diagnostics

Table 3: Label-retention audit and BLER differences, mask minus comparator, with adjusted 95% intervals.

Quantity	7.25 dB	8 dB
Candidate coded labels retained by the mask (%)	44.78	57.18
Coded labels retained (bits)	155,985,264	199,182,240
Erroneous retained labels observed by the evaluator	0	0
Mask minus mask-oracle BLER	0.0000 [0.0000, 0.0000]	0.0000 [0.0000, 0.0000]
Mask minus all-oracle BLER	0.0060 [0.0039, 0.0082]	0.0048 [0.0028, 0.0068]
Mask minus random-probe BLER	0.0058 [0.0028, 0.0087]	0.0049 [0.0026, 0.0072]
Mask minus periodic-probe BLER	0.0053 [0.0028, 0.0078]	0.0034 [0.0011, 0.0057]
Mask minus rejection-probe BLER	0.0058 [0.0035, 0.0080]	−0.0527 [−0.0629, −0.0425]

Note: BLER, block error rate.

Table 3 summarizes the evaluator-only label audit and the oracle and probe BLER gaps. The mask retained fewer than three-fifths of the candidate coded labels at both SNRs, and the evaluator observed no erroneous retained label; the mask therefore reproduced the mask-oracle result exactly in these runs. This is an empirical observation, not a correctness certificate. A small residual gap to the all-oracle arm remains, with adjusted upper bounds below 0.01. Positive probe entries show the post-adaptation BLER advantage of probes over the mask.

The once-versus-replay comparison separates the inherited accepted-batch update count from a fixed 32-update budget. The mask changes label quantity as well as reliability, so its comparison with FEC alone is not a pure correctness experiment; the matched-mask oracle diagnostic supplies that check. Decoder and optimizer timing, label counts, BCE, coded bit error rate (BER) and payload BER are reported in the supplementary tables.

3.4 Processing cost changes the comparison

Table 4 reports the measured adaptation interval and the serialized throughput, defined as error-free delivered bits divided by radio time plus the measured inference, decoder and adaptation intervals. These descriptive local timings show that a radio-efficiency gain does not establish an end-to-end throughput gain: at both SNRs the mask has lower serialized throughput than the frozen RNN and the colored APP.

Table 4: Measured adaptation time and serialized throughput.

	7.25 dB		8 dB	
Arm	Adapt. (s)	Mbit/s	Adapt. (s)	Mbit/s
Frozen RNN	0.000	1.365	0.000	1.965
FEC replay	0.792	1.084	0.792	1.568
Mask	0.825	1.174	0.827	1.637
Colored APP	0.000	1.630	0.000	2.296

Note: Adapt., adaptation time; RNN, recurrent neural network; FEC, forward error correction; APP, a posteriori probability receiver.

The full mask adaptation interval exceeds the 0.4608 s of radio time that 1600 ordinary frames occupy. The implementation therefore does not establish a net wall-time benefit or real-time service; its positive primary finding concerns transmitted-sample efficiency.

3.5 Interpretation and finite-horizon caveat

The useful distinction is between an imperfect batch label rule and a need for new observations. A batch can satisfy an average-check threshold while individual failed codewords still supply harmful targets. Filtering at the codeword level is therefore a useful simple control before interpreting a large oracle gap as motivation for additional acquisition.

The cost finding has a finite horizon. Let C be the number of collection frames, a the number of successfully delivered collection frames, p_m and p_p the stationary post-success probabilities for the mask and probe arms, and P the number of paid frames. At post-adaptation horizon n , a probe beats the free mask precisely when

$$n(p_p - p_m)(C + n) > P(a + np_m) \quad (5)$$

Thus even a small positive post-success difference can eventually amortize probe costs under indefinite stationarity. Our experiment measures only $n = 1000$ and cannot validate such extrapolation. It also does not search for a longer horizon that would make a favored strategy win. Conversely, the weak performance of the qualified IEDD implementation cannot justify a broad neural advantage. The recommendation supported here is to test a strong free-label control and resource accounting before designing more elaborate acquisition policies.

3.6 Limitations

The study uses one checkpoint, one finite-sinusoid simulator and two adjacent SNRs. It does not measure training-seed robustness, field channels, continuous inter-frame fading, macro drift, ICI, RF impairments or hybrid automatic repeat request (HARQ) label recovery. Known statistical covariance is available to the conventional controls; estimating it under unknown interference is a different question. The final frame-error rates remain high, so the fixture is a stress test rather than evidence of a production-quality link. A zero observed masked-label error count is not a correctness certificate or an independent-bit confidence bound. The modeled control signaling is reliable and waveform-level synchronization is idealized. Hardware timing measures a local execution path and excludes file I/O and transmitter processing; no real-time throughput is established.

4 CONCLUSION

Within one controlled coded OFDM fixture, a large gap between recovered-label and oracle-label adaptation should not by itself be read as a case for purchasing new training transmissions. The gap was dominated by whether labels were correct, not by which batches were used or how many labels were available, and a codeword-level constraint filter that costs no extra frames reached post-adaptation error rates close to those of paid probes while avoiding their transmission and signaling overhead.

The defensible result is a bounded empirical comparison: before acquiring additional supervision, distinguish label correctness from selection and quantity, test codeword-level recovery controls, and charge all new radio resources. This does not establish a universal limit on the value of probes, and the measured serialized processing comparison precludes claiming an end-to-end throughput advantage for the filter. Repeating the comparison across independently trained checkpoints, longer horizons, continuous channel drift and field channels is the natural next step.

STATEMENTS ON COMPLIANCE WITH ETHICAL STANDARDS

Disclosure of conflict of interest

The authors declare that they have no conflict of interest. Individual disclosures: Kabil Abderrahim declares no financial or non-financial conflict of interest; Hudhaifa Mohammednawwar Qasim Alajar declares no financial or non-financial conflict of interest; Abdullah Al Mamun Sadi declares no financial or non-financial conflict of interest.

Data Availability Statement

The simulation data, analysis scripts and reproduction package supporting the findings of this study are available from the corresponding author upon reasonable request.

REFERENCES

- [1] Schibisch S, Cammerer S, Dörner S, Hoydis J, ten Brink S. Online label recovery for deep learning-based communication through error correcting codes. In: Proceedings of the 15th International Symposium on Wireless Communication Systems (ISWCS); 2018. p. 1–5. doi:10.1109/ISWCS.2018.8491189
- [2] Fischer MB, Dörner S, Cammerer S, Shimizu T, Lu H, ten Brink S. Adaptive neural network-based OFDM receivers. In: Proceedings of the IEEE 23rd International Workshop on Signal Processing Advances in Wireless Communication (SPAWC); 2022. p. 1–5. doi:10.1109/SPAWC51304.2022.9833970

- [3] Finish RA, Cohen Y, Raviv T, Shlezinger N. Symbol-level online channel tracking for deep receivers. In: Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP); 2022. p. 8897–901. doi:10.1109/ICASSP43922.2022.9747026
- [4] Teng CF, Chen YL. Syndrome-enabled unsupervised learning for neural network-based polar decoder and jointly optimized blind equalizer. IEEE J Emerg Sel Top Circuits Syst. 2020;10(2):177–88. doi:10.1109/JETCAS.2020.2992593
- [5] Park S, Jang H, Simeone O, Kang J. Learning to demodulate from few pilots via offline and online meta-learning. IEEE Trans Signal Process. 2021;69:226–39. doi:10.1109/TSP.2020.3043879
- [6] Cohen KM, Park S, Simeone O, Shamai S. Bayesian active meta-learning for reliable and efficient AI-based demodulation. IEEE Trans Signal Process. 2022;70:5366–80. doi:10.1109/TSP.2022.3220035
- [7] Uzlaner N, Raviv T, Shlezinger N, Todros K. Asynchronous online adaptation via modular drift detection for deep receivers. IEEE Trans Wirel Commun. 2025;24(5):4454–68. doi:10.1109/TWC.2025.3533959
- [8] Obeed M, Jian M. Learning during detection: continual learning for neural OFDM receivers via DMRS [Preprint]. arXiv:2602.20361. 2026 Feb 23 [cited 2026 Oct 12]. doi:10.48550/arXiv.2602.20361
- [9] Ait Aoudia F, Hoydis J. End-to-end learning for OFDM: from neural receivers to pilotless communication. IEEE Trans Wirel Commun. 2022;21(2):1049–63. doi:10.1109/TWC.2021.3101364