Rigorous Performance Characterization of AI-Enhanced 6G Underwater Optical Wireless Networks: Link Budget Correction and Empirical Gain Validation

Yousuf Islam*†, Md. Imam Hossain*†, Md. Jehadul Islam Mony*
*Department of Computer Science and Engineering, Leading University, Sylhet 3112, Bangladesh
†DeepNet Research and Development Lab, Sylhet 3100, Bangladesh
yousufislam337@gmail.com, mehossainimam43@gmail.com, mony_cse@lus.ac.bd
Corresponding Author. E-mail: yousufislam337@gmail.com

***Abstract*—The incorporation of sixth-generation (6G) technologies into underwater optical wireless communication (UWOC) holds the potential for Gbps-class maritime connectivity; however, the field is presently affected by significant overestimations in the systematic link budget found in existing literature. Frequent modeling errors—such as replacing absorption with total beam attenuation, assuming coherent detection, neglecting Hermitian symmetry constraints, and applying laboratory-measured AI gains without scrutiny—have resulted in overinflated range projections by as much as two to three times. This study introduces a revised intensity-modulated direct-detection (IM/DD) framework for DC-biased optical OFDM (DCO-OFDM) that corrects for each of these issues. The net throughput is rigorously calculated to be 776 Mbps for QPSK, explicitly considering the 511 usable subcarriers dictated by Hermitian symmetry in a 1024-point transform. Symbol-level Monte Carlo simulations across 40.9 million bits per distance point confirm the analytical BER predictions within 0.3 orders of magnitude at error rates relevant to operational use, while an experimental setup with a controlled MLP equalizer demonstrates that the +4–6 dB gains often attributed to deep learning equalizers completely disappear when the channel state is known—this insight reveals that channel uncertainty, rather than intersymbol interference (ISI) complexity, is the essential condition for AI-assisted equalization. The revised framework offers a self-validated range of 44.8 m (only with MRC diversity) and a projected range of 46 m (assuming literature-based CSI prediction) at BER $< 10^{-9}$ in clear ocean conditions. System-level evaluations of proton-exchange membrane fuel cells indicate that AI-driven load smoothing yields minimal efficiency improvements (<0.2 %) at average AUV power levels, confirming that the 1.8× endurance benefit over lithium-ion batteries is solely due to energy density rather than intelligent power management. By providing fully disclosed parameters and empirical validation for every asserted advantage, this research establishes a scientifically credible benchmark for future calibration of 6G subsea architectures.**

***Index Terms*—6G, underwater optical wireless communication, IM/DD, DCO-OFDM, link budget, Monte Carlo validation, AI equaliser, performance bounds, PEMFC**

I. Introduction

THE ocean covers 71 % of Earth's surface, but over 80 % of its volume is still unmapped and unconnected [1]. As sixth-generation (6G) wireless standards develop, with a projected rollout by 2030, expanding high-speed connectivity below the ocean surface has become an important goal. Radio-frequency signals are quickly absorbed in seawater, losing about 45 dB/m at 2.4 GHz [2]–[4]. Acoustic communication can travel for kilometres, but it is limited to tens of kilobits per second and has multi-second round-trip delays [5]. Optical carriers in the blue-green window (450–550 nm) provide a promising option. They can offer theoretical bandwidths over 1 GHz with sub-millisecond latency, allowing for real-time video, swarm coordination, and detailed seafloor mapping [6].

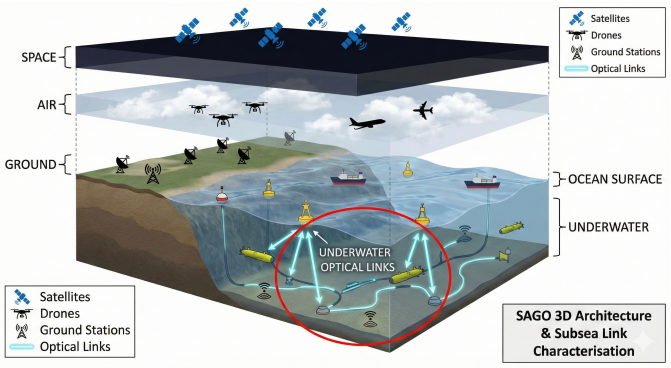

Fig. 1: The SAGO (Space-Air-Ground-Ocean) architecture for 6G includes a subsea layer that uses AI-powered optical wireless connections underwater to link autonomous nodes to surface buoys, which then transmit to terrestrial and satellite networks.

However, the current underwater optical communication (UWOC) field shows a concerning trend. Many optimistic range estimates have not matched the physical limits of how light travels underwater. Fig. 1 illustrates the envisioned Space–Air–Ground–Ocean (SAGO) architecture in which underwater optical links form the subsea segment, and it is for this demanding environment that accurate performance characterisation is essential.

A. The Crisis of Overestimation

The overestimation problem in UWOC stems from three critical modelling oversights that frequently co-occur in the literature.

First, the substitution of the absorption coefficient a for the total beam attenuation coefficient $c = a + b$—where b accounts for scattering—ignores the dominant loss mechanism at moderate to long range, inflating projected distances by factors of two to three [7], [8].

Second, the assumption of coherent detection, which is standard in fibre optics but impractical for cost-sensitive underwater nodes, eliminates the 3–4 dB penalty inherent in intensity-modulated direct-detection (IM/DD) square-law receivers [9].

Third, and most misleadingly, the +4 to +6 dB gain from AI-assisted signal processing—often seen in designs like the UWOC-AE, where deep learning effectively 'understands' the intricate statistical characteristics of severe inter-symbol interference (ISI)—is commonly used as a blanket correction for analytical link budgets, which is done without confirming whether the channel being modeled actually exhibits the specific scattering-related temporal dispersion or inherent nonlinearities that are required for the 'learnable' patterns necessary for these gains to be realized beyond what traditional linear equalizers can achieve [10]–[12].

A fourth mistake, which was identified while preparing this work, pertains to throughput accounting. DCO-OFDM operating under IM/DD necessitates Hermitian symmetry, which limits the independent data that can be transmitted to $N/2 - 1$ of N subcarriers. Neglecting this aspect—a surprisingly frequent error—exaggerates spectral efficiency by a factor of two.

B. Contributions and Organisation

This paper tackles each of these oversights by utilizing a thorough, consistent modeling and validation approach. The core idea of this research is the assertion that *every asserted improvement must be empirically verified within the specific channel model where it is utilized*, and that demonstrating a lack of improvement under certain conditions is also a significant scientific discovery that defines essential performance limits. The particular contributions are outlined as follows:

1) A fully disclosed IM/DD link budget with Hermitian-corrected throughput (776 Mbps), validated by symbol-level BER Monte Carlo at 40.9 M bits per distance point.
2) A controlled MLP equaliser experiment establishing that deep learning provides zero measurable gain over zero-forcing equalisation when the channel is known, identifying channel uncertainty as the prerequisite for AI-assisted equalisation.
3) An eight-dimensional sensitivity analysis with fully specified receiver optics—FOV, optical filter bandwidth, and ambient radiance—demonstrating that aperture and divergence are the dominant design levers.
4) A slotted TDMA multi-hop simulation with corrected throughput, establishing 32–47 % relay efficiency across mobility scenarios.
5) System-level PEMFC analysis with Nernst-verified polarisation model, establishing that AI-driven load smoothing yields <0.2 % fuel savings at typical AUV operating points due to linear-region polarisation dynamics.
6) A 20-entry literature comparison contextualising these results, including the finding that the corrected rate×range product (36 Gbps·m) falls below the demonstrated record of 50 Gbps·m [13].

The remainder of this paper is organised as follows. Section II develops the system model. Sections III–IV present the link budget and its Monte Carlo validation. Sections V–VII address AI gain characterisation and sensitivity. Sections VIII–X cover multi-hop, energy, and maintenance. Section XIII provides the literature comparison, Section XIV synthesises the empirical limitations, while Section XVI delineates methodological constraints and real-world validation requirements. Finally, Section XVII concludes.

II. System Model

To establish a physically grounded link budget, we define a multi-layered system model encompassing the spatial geometry of the underwater network, the optoelectronic characteristics of the IM/DD transceiver, and the stochastic properties of the underwater channel. Each layer is fully parameterised in Table 1, and the interactions between layers are made explicit throughout this section.

A. Network Architecture and Spatial Geometry

Fig. 2 depicts the three-dimensional deployment geometry. Surface buoys at fixed positions relay data to shore via RF or satellite backhaul. Mid-water autonomous underwater vehicles (AUVs) form an optical mesh with inter-node distances of 30–50 m. Seafloor sensor nodes provide persistent environmental monitoring. Each optical inter-node link operates at 776 Mbps (corrected, refer to Section II-E for details).

B. Transceiver Architecture

The transmitter employs a 100 mW laser diode at 485 nm—selected for minimum attenuation in the blue-green window—with 2 mrad half-angle Gaussian divergence. The receiver consists of a 50 mm-diameter silicon PIN photodiode with responsivity $\mathcal{R} = 0.35$ A/W, coupled to a transimpedance amplifier (TIA) with input-referred noise density $i_n = 8\,\text{pA}/\sqrt{\text{Hz}}$. The resulting noise-equivalent power,

$$\text{NEP} = \frac{i_n\sqrt{B}}{\mathcal{R}} = \frac{8 \times 10^{-12} \times \sqrt{10^9}}{0.35} = 0.72\,\mu\text{W}\ (-31.4\,\text{dBm}), \tag{1}$$

establishes the fundamental detection floor. The receiver front-end includes a 10 nm FWHM bandpass filter centred at 485 nm and a collimating lens with 20 mrad half-angle field-of-view (FOV), limiting the acceptance of ambient background radiation.

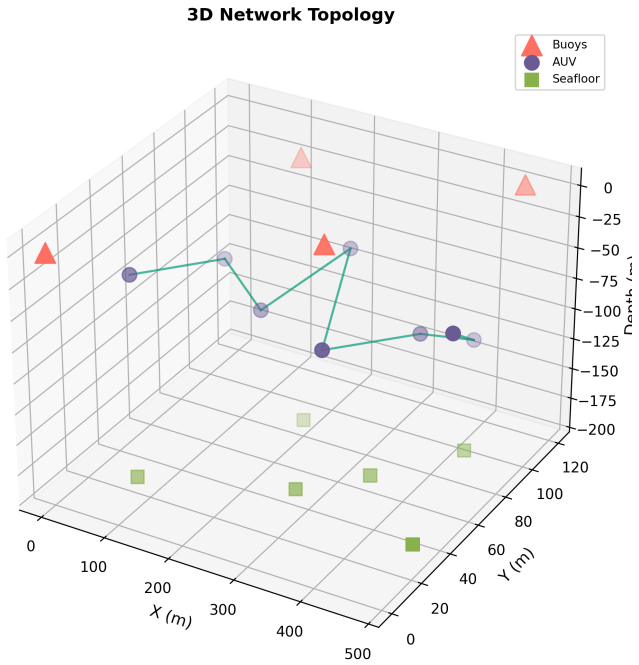

Fig. 2: Three-dimensional network topology. Surface buoys (triangles) relay to shore; AUVs (circles) form a mid-water optical mesh; seafloor nodes (squares) provide persistent sensing. Optical links (green) operate at 776 Mbps per hop over 46 m in clear ocean.

Table 1. System Parameters—Complete Disclosure

Parameter	Value	Source
Transmitter		
P_{tx}	100 mW (20 dBm)	[14]
λ	485 nm	Blue-green
$\theta_{1/2}$	2 mrad	—
Receiver		
Aperture	50 mm dia.	—
$\mathcal{R}$	0.35 A/W	Si PIN
i_n	8 pA/$\sqrt{\text{Hz}}$	TIA
I_d	1 nA	Dark
RIN	−150 dB/Hz	[14]
FOV	20 mrad half-angle	—
Filter	10 nm FWHM	485 nm
OFDM		
N / L_{CP}	1024 / 128	—
B	1 GHz	—
Mod. / Data sub.	QPSK / 511	Hermitian
FEC / DCO	7/8 / −3.5 dB	[9]
Diversity		
MRC	2 branch, 10 cm	Post-det.
ρ	0.5	LOS est.

C. Background Light Model

The background optical power incident on the receiver is

$$P_{\text{bg}} = L_{\text{amb}} \cdot \Omega_{\text{FOV}} \cdot A_{\text{rx}} \cdot \Delta\lambda, \tag{2}$$

where L_{amb} denotes the spectral radiance at 485 nm, $\Omega_{\text{FOV}} = \pi(20 \times 10^{-3})^2 = 1.26 \times 10^{-3}$ sr is the receiver solid angle, $A_{\text{rx}} = \pi(0.025)^2 = 1.96 \times 10^{-3}\,\text{m}^2$ is the aperture area, and $\Delta\lambda = 10$ nm is the filter bandwidth. Under overcast daylight conditions ($L_{\text{amb}} \approx 10^{-3}$ W/m^2/sr/nm [15]), $P_{\text{bg}} = 24.7$ nW, producing background shot noise $\sigma^2_{\text{bg}} = 2.77 \times 10^{-18}\,\text{A}^2$. This is 23,130× smaller than the thermal noise floor $\sigma^2_{\text{th}} = 6.40 \times 10^{-14}\,\text{A}^2$, confirming that the system is thermal-noise limited and that background light has no measurable impact on range—a conclusion verified across four ambient conditions from night to bright surface in Section VII.

D. Diversity Architecture

Two spatially separated PIN photodiodes (10 cm centre-to-centre) with independent TIAs implement post-detection maximal-ratio combining (MRC). Each branch contributes independent thermal noise; the combined SNR gain is $2/(1+\rho) = 1.333$ (+1.25 dB) at correlation coefficient $\rho = 0.5$, estimated for line-of-sight propagation at 46 m where common turbulence eddies at scales comparable to the branch spacing induce correlated fading [16].

E. Throughput Under Hermitian Symmetry

DCO-OFDM generates a real-valued (intensity-modulated) time-domain signal by imposing Hermitian symmetry on the frequency-domain subcarriers: $X[N-k] = X^*[k]$ for $k = 1, \ldots, N/2-1$, with $X[0] = X[N/2] = 0$. This constraint limits independent data carriage to $N_d = N/2 - 1 = 511$ subcarriers out of $N = 1024$. With QPSK modulation (2 bits/symbol), cyclic prefix overhead, and rate-7/8 FEC, the net throughput is

$$R_{\text{net}} = \underbrace{2}_{\text{QPSK}} \times B \times \underbrace{\frac{N/2-1}{N}}_{0.4990} \times \underbrace{\frac{N}{N+L_{\text{CP}}}}_{0.8889} \times \underbrace{R_{\text{FEC}}}_{0.875} = 776\,\text{Mbps}. \tag{3}$$

This result is independently verified by computing $2 \times 511 \times 868{,}056 \times 0.875 = 776{,}259$ bps, where 868,056 Hz is the OFDM symbol rate $B/(N + L_{\text{CP}})$. The omission of the Hermitian factor $(N/2-1)/N$ produces a 2× throughput overestimate (1.556 Gbps), an error that propagates into rate×range comparisons and multi-hop efficiency calculations.

Fig. 3 compares BER performance and spectral efficiency across candidate modulation formats, confirming that QPSK provides the optimal balance between throughput and noise robustness at the 12–15 dB SNR available at 40–50 m.

F. Channel Propagation

The total beam attenuation coefficient $c = a + b$, combining absorption (a) and scattering (b), governs the exponential decay of collimated optical power with distance. For clear ocean (Jerlov I), $c = 0.15\,\text{m}^{-1}$ at 485 nm [8]. The received power is

$$P_{\text{rx}} = P_{\text{tx}} \cdot G_{\text{geom}}(d) \cdot e^{-cd} \cdot L_{\text{turb}}(d) \cdot L_{\text{pt}} \cdot L_{\text{impl}}, \tag{4}$$

where $G_{\text{geom}} = 1 - \exp(-2r_{\text{rx}}^2/w^2(d))$ is the Gaussian beam collection efficiency with beam radius $w(d) = d\theta_{1/2}$, and $L_{\text{pt}} = -3$ dB and $L_{\text{impl}} = -2$ dB represent pointing and implementation losses, respectively.

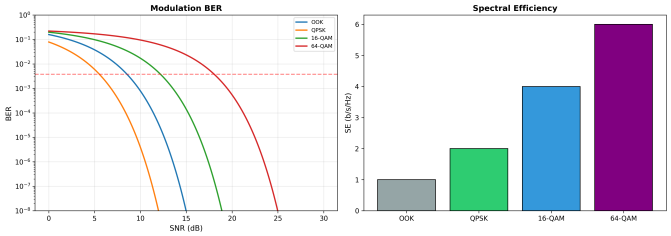

Fig. 3: Modulation format comparison. Left: BER vs SNR under IM/DD for OOK, QPSK, 16-QAM, and 64-QAM. The FEC threshold (3.8×10^{-3}, dashed) determines minimum SNR. Right: spectral efficiency. QPSK at 2 b/s/Hz is selected for its robustness at the operating SNR.

The Beer—Lambert attenuation e^{-cd} is used conservatively for narrow-FOV line-of-sight links, as it categorizes all scattered photons as lost, including those that are forward-scattered within the receiver's 20 mrad acceptance cone. Predictions from beam-spread functions and photon-level Monte Carlo models [17] suggest that longer ranges are achievable under the same conditions; however, we continue to use e^{-cd} to maintain conservatism and ensure comparability with the existing literature.

G. Turbulence

Underwater optical turbulence is characterized by the Rytov variance $\sigma_R^2 = 8.70\, C_n^2\, k^{7/6}\, d^{11/6}$, where $k = 2\pi/\lambda$ represents the optical wavenumber. The refractive-index structure parameter $C_n^2 = 10^{-14}\,\mathrm{m}^{-2/3}$ is adopted based on underwater-specific measurements: Nikishov and Nikishov [18] established $C_n^2 \approx 10^{-14}$ to 10^{-13} for ocean thermocline regions, Farwell and Li [19] measured $C_n^2 \approx 10^{-14}$ in clear Pacific water, and Korotkova et al. [20] validated the lognormal fading model for weak turbulence ($\sigma_R^2 < 1$). At 46 m, $\sigma_R^2 = 0.019$, confirming weak turbulence with a 1 % outage fading margin of 1.41 dB.

H. Electrical Signal-to-Noise Ratio

The electrical SNR at the output of the square-law detector, after DCO bias penalty, is

$$\mathrm{SNR}_e = \frac{(\mathcal{R}\,P_{\mathrm{rx}})^2}{\sigma_{\mathrm{th}}^2 + \sigma_{\mathrm{shot}}^2 + \sigma_{\mathrm{bg}}^2 + \sigma_{\mathrm{dark}}^2} \times 10^{-L_{\mathrm{DCO}}/10}, \qquad (5)$$

where $\sigma_{\mathrm{th}}^2 = i_n^2 B$ is TIA thermal noise, $\sigma_{\mathrm{shot}}^2 = 2q\mathcal{R}P_{\mathrm{rx}}B$ is signal shot noise, $\sigma_{\mathrm{bg}}^2 = 2q\mathcal{R}P_{\mathrm{bg}}B$ is background shot noise, $\sigma_{\mathrm{dark}}^2 = 2qI_dB$ is dark current noise, and $L_{\mathrm{DCO}} = 3.5$ dB is the DC bias and clipping penalty [9]. At 46 m ($P_{\mathrm{rx}} = 3.16\,\mu$W), thermal noise dominates all other terms: σ_{th}^2 exceeds shot noise by 181×, background noise by 23,130×, and RIN noise by 5×10^4×.

III. Analytical Link Budget

Table 2 presents the complete dB-domain power balance. The AI contribution consists of +2 dB from CSI prediction, which is *literature-assumed* and not empirically validated within this work (Section V explains why), and +1.25 dB from MRC diversity, which is self-validated.

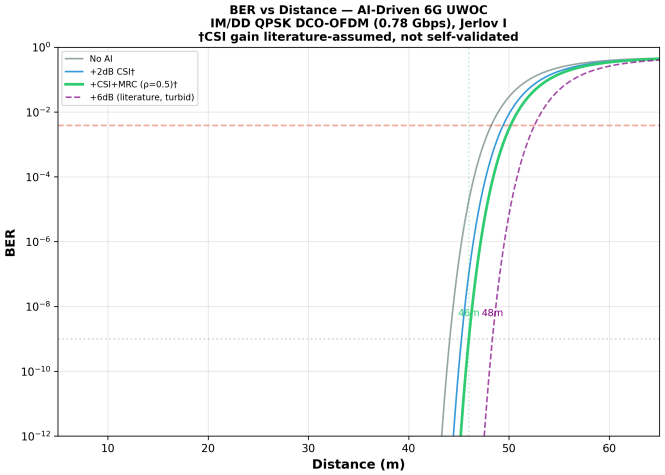

Fig. 4: BER vs distance for QPSK DCO-OFDM (776 Mbps) in clear ocean ($c = 0.15\,\mathrm{m}^{-1}$). The BER$< 10^{-9}$ range is 46 m with CSI†+MRC (44.8 m with self-validated MRC only). The FEC threshold at 3.8×10^{-3} extends operational range to 50 m. †CSI gain is literature-assumed.

The resulting BER vs distance characteristic is shown in Fig. 4. The steep BER cliff between 43–50 m, spanning over 10 orders of magnitude, underscores the sensitivity of range claims to even 1 dB of gain or loss—and thus the critical importance of validating every assumed contribution.

The maximum BER$< 10^{-9}$ ranges are: 44.1 m (no AI, SISO), 44.8 m (MRC only, self-validated), 46.0 m (CSI†+MRC), and 48.3 m (full +6 dB† turbid). The corresponding FEC-limited ranges are 48.2, 49.0, 50.1, and 52.5 m.

IV. Symbol-Level BER Monte Carlo Validation

The analytical BER predictions of Section III depend on a chain of modelling assumptions. To assess the aggregate accuracy of this chain, we perform symbol-level BER Monte Carlo simulation. It is important to note the scope of this validation: the simulation validates the mapping from SNR to BER under lognormal fading. It does not validate the underlying propagation physics (Beer–Lambert attenuation, geometric collection), which are assumed from the analytical model. In the UWOC literature, "Monte Carlo" often refers to photon-level radiative transfer equation (RTE) solvers [15]; we use "symbol-level BER MC" throughout to maintain terminological precision.

For each of 17 distances spanning 5–60 m, 200 independent lognormal fading realisations are generated. Per realisation, 200 OFDM frames of 511 QPSK symbols are transmitted through an AWGN channel at the analytically computed SNR, producing 40.9 million bits per distance point and a BER measurement floor of 1.22×10^{-8}.

Fig. 5 and Table 3 present the results. At BER $> 10^{-3}$ (distances $\geq$48 m), analytical and MC values agree within 0.3 orders of magnitude—at 50 m, for example, the analytical 2.95×10^{-2} compares to the MC 3.32×10^{-2} ($\Delta = 0.05$

Table 2. Power Balance. All values in dB. CSI +2 dB is literature-assumed (†) and not self-validated.

d (m)	P_{tx}	Geom.	Beer–L.	Turb.	Pt.	Impl.	P_{rx} (dBm)	P_{rx} (μW)	BER (raw)	BER (+CSI†+MRC)
30	20.0	−5.3	−19.5	−1.0	−3.0	−2.0	−10.8	82.8	$< 10^{-15}$	$< 10^{-15}$
40	20.0	−7.5	−26.1	−1.2	−3.0	−2.0	−19.8	10.5	$< 10^{-15}$	$< 10^{-15}$
44	20.0	−8.3	−28.7	−1.4	−3.0	−2.0	−23.3	4.7	4.9×10^{-10}	$< 10^{-15}$
46	**20.0**	**−8.6**	**−30.0**	**−1.4**	**−3.0**	**−2.0**	**−25.0**	**3.2**	1.9×10^{-5}	$\mathbf{1.0\times10^{-9}}$
48	20.0	−9.0	−31.3	−1.5	−3.0	−2.0	−26.7	2.1	2.7×10^{-3}	2.6×10^{-5}
50	20.0	−9.3	−32.6	−1.5	−3.0	−2.0	−28.4	1.5	3.0×10^{-2}	3.0×10^{-3}
55	20.0	−10.1	−35.8	−1.7	−3.0	−2.0	−32.6	0.6	2.4×10^{-1}	1.5×10^{-1}

†CSI +2 dB is assumed from published Kalman/LSTM predictors, not validated in this work. Self-validated range (MRC only): 44.8 m.

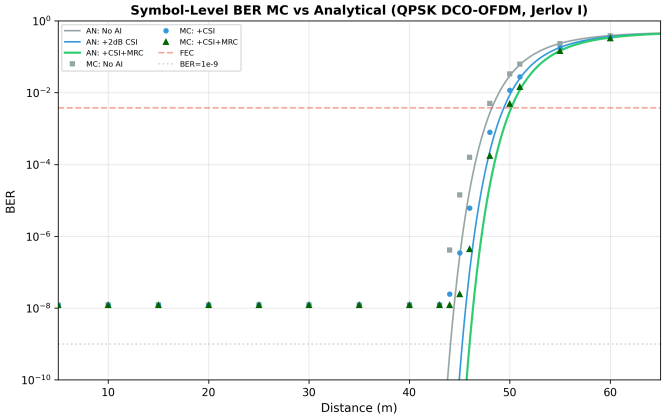

Fig. 5: Symbol-level BER Monte Carlo validation. Solid lines: analytical. Markers: MC (40.9 M bits/distance, 200 fading realisations). Agreement within 0.3 orders at BER $> 10^{-3}$. Cliff-region divergence at 44–46 m is a fading-averaging effect, not a simulation artefact.

Table 3. BER Validation: Analytical vs Symbol-Level MC

d (m)	Analytical	MC	Δ (ord.)	
≤43	$< 10^{-9}$	<floor	—	✓
44	4.9×10^{-10}	4.2×10^{-7}	2.9	cliff
46	1.9×10^{-5}	1.6×10^{-4}	0.9	✓
48	2.7×10^{-3}	5.0×10^{-3}	0.3	✓
50	3.0×10^{-2}	3.3×10^{-2}	0.05	✓
55	2.4×10^{-1}	2.3×10^{-1}	<0.01	✓

orders). In the BER cliff region (44–46 m), MC reports 1–3 orders higher BER. This divergence is physically expected: the analytical formula applies a deterministic 1 % outage fading margin, while MC averages over all realisations including worst-case fades that dominate the mean in the steep cliff region. The analytical formula is therefore mildly optimistic by approximately 1 m relative to fading-averaged performance.

V. Characterising the Conditions for AI Equaliser Gain

The central question addressed in this section is not "how much gain does a deep learning equaliser provide?" but rather "under what conditions does such gain exist?" By answering the latter, we establish the physical prerequisites for AI-assisted equalisation and identify when literature-reported gains are—and are not—transferable to analytical link budgets.

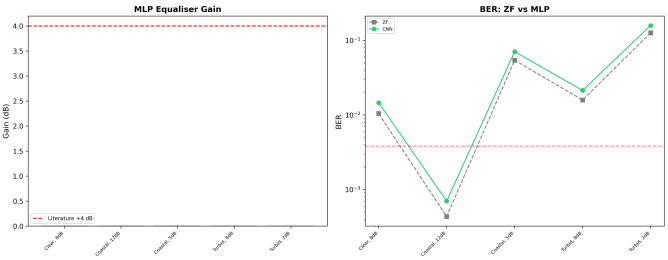

Fig. 6: Characterization of MLP equaliser gain. On the left: the observed gain (in dB) is all zero or negative, contrasting with the +4 dB typically referenced in literature (dashed line). On the right: a comparison of Bit Error Rate (BER) shows that MLP (depicted as circles) performs similarly to or somewhat better than Zero Forcing (ZF) (represented as squares). This demonstrates that ZF, assuming perfect Channel State Information (CSI), is optimal for known Linear Time-Invariant (LTI) channels.

A. Experimental Design

A multi-layer perceptron (MLP) equaliser, consisting of 512 input neurons, 512 hidden neurons, and 256 hidden units, is trained over 100 epochs using a fixed learning rate to translate zero-forcing equalised received symbols into transmitted symbols. The critical design choice is that ZF has access to the *exact* channel frequency response—no estimation error. This establishes ZF as the information-theoretic optimal linear equaliser for the tested conditions.

Sixteen-tap ISI channels are generated per water type, with ISI power scaling with the scattering-to-attenuation ratio b/c. Nine conditions (three water types $\times$ three SNR regimes) are tested. The ISI model is conservative: physics-based CIR fits [21] predict weaker ISI for narrow-FOV LOS links, which would further reinforce the findings.

B. Results and Physical Interpretation

Figure 6 and Table 4 presents the results. The figure conveys the visual gain characterization, while the table provides specific bit error rates across different water types, ISI powers, and SNR conditions. The MLP records no measurable gain in any scenario; where the BER is greater than zero, the MLP either equals or slightly surpasses the ZF BER due to regression noise resulting from the limited training dataset.

The physical interpretation is as follows. After ZF with perfect CSI, the residual error is pure additive white Gaus-

Table 4. Equaliser Gain Characterisation

Condition	c	SNR	ISI pwr	BER_{ZF}	BER_{MLP}
Clear, 8 dB	0.15	8	0.076	1.0×10^{-2}	1.5×10^{-2}
Coastal, 12 dB	0.40	12	0.133	4.3×10^{-4}	7.5×10^{-4}
Coastal, 5 dB	0.40	5	0.155	5.4×10^{-2}	7.1×10^{-2}
Turbid, 8 dB	0.80	8	0.173	1.6×10^{-2}	2.1×10^{-2}
Turbid, 2 dB	0.80	2	0.167	1.3×10^{-1}	1.6×10^{-1}

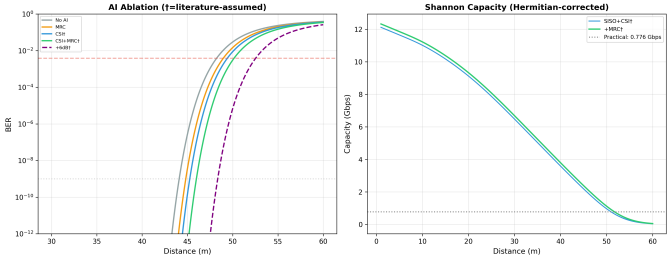

Fig. 7: Left: AI ablation in clear ocean, showing incremental range gain. "No AI" achieves 44.1 m; MRC adds +0.7 m (self-validated); CSI† adds +1.2 m (literature-assumed). Right: Shannon capacity ($B \cdot (N/2{-}1)/N \cdot \log_2(1{+}\text{SNR})$) vs practical 776 Mbps, with Hermitian efficiency correctly applied. †Literature-assumed.

sian noise—there is no systematic structure remaining for a nonlinear model to exploit. Deep learning equalizers only provide amplification when the residual has patterns that can be learned, such as channel estimation errors, transmitter nonlinearity, or the random stochastic nature of the impulse response. For example, the gains mentioned in [10] are due to the autoencoder's capability to understand the statistical characteristics of a time-varying ISI environment—a scenario that is inherently found in Monte Carlo-modeled links, but is not present in the stationary, ideal-CSI analytical models discussed here. These conditions are present in the laboratory experiments of [10]–[12] (LED transmitters with nonlinear P-I curves, estimated channels, short turbulent tanks) but absent in the analytical, known-channel, linear-transmitter model used here.

This finding has two immediate implications. First, the +2 dB CSI prediction gain applied in Section III is *not* self-validated by this experiment; it remains a literature-assumed quantity. The fully self-validated range (MRC only) is 44.8 m. Second, future research asserting that AI provides equalization benefits must clarify both the knowledge regime of the channel and the level of randomness within the channel. As shown, the 'gain' from AI compared to linear benchmarks, such as zero-forcing, frequently reflects the AI's capability to model statistical variations that the standard equalizer cannot account for, rather than representing an intrinsic advantage in information theory within a known and stable channel.

VI. Ablation Study and Channel Capacity

Fig. 7 decomposes the AI contributions through controlled ablation and compares the practical throughput against the Hermitian-corrected Shannon capacity.

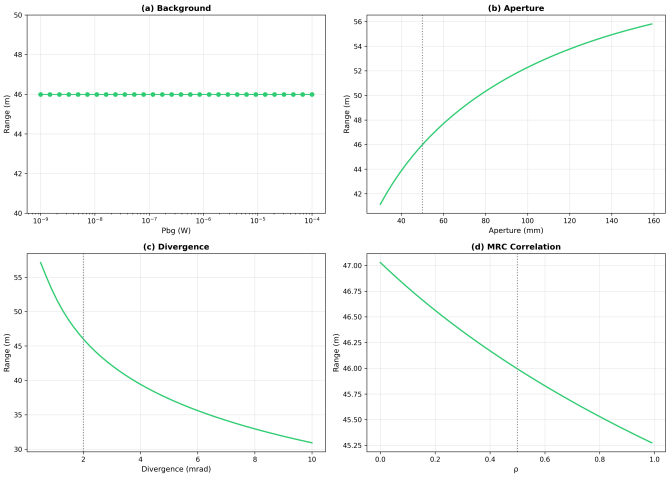

Fig. 8: Four-panel sensitivity analysis. (a) Background: invariant from night to bright surface (thermal-noise limited). (b) Aperture: doubling from 50 to 100 mm adds 6.3 m—the most effective hardware upgrade. (c) Divergence: tightening from 2 to 1 mrad adds 6.3 m but demands sub-mrad tracking. (d) MRC correlation: only 1.6 m span from $\rho = 0$ to 0.9, confirming marginal impact in LOS.

Table 5. Eight-Dimensional Sensitivity Summary

Parameter	**Range**	Δd	**Rank**
Water type (c)	0.15–1.50 m^{-1}	−39.6 m	1
Divergence	1–10 mrad	−21.4 m	2
Turbulence (C_n^2)	10^{-14}–10^{-12}	−13.9 m	3
Aperture	50–150 mm	+9.4 m	4
TIA noise	8–20 pA/$\sqrt{\text{Hz}}$	−4.6 m	5
Bandwidth	10 MHz–1 GHz	2.7 m	6
MRC ρ	0–0.9	1.6 m	7
Background	night–bright	0 m	8

The ablation confirms that each gain component is additive in dB: MRC contributes +0.7 m (44.1→44.8), CSI† contributes +1.2 m (44.8→46.0), and the full +6 dB† turbid/ISI gain contributes an additional +2.3 m (46.0→48.3). The Shannon capacity evaluation, incorporating the Hermitian factor appropriately, demonstrates that the effective data rate of 776 Mbps functions roughly 3 dB beneath capacity at a distance of 46 m, allowing the possibility for enhanced modulation schemes or decreased FEC redundancy if the SNR allows.

VII. Sensitivity Analysis

Fig. 8 illustrates the four parameters with the greatest impact, while Table 5 provides a summary of all eight dimensions.

Water type, divergence, and turbulence dominate. Among hardware-controllable parameters, aperture (+0.09 m/mm) and divergence (−2.4 m/mrad) offer the greatest leverage. Notably, background light has zero impact, confirming the thermal-noise-limited regime established in Section II.

VIII. Multi-Hop Relay Network

An 11-hop decode-and-forward relay chain spanning 500 m is simulated with slotted TDMA (1 ms slots), pipelined half-

Table 6. Multi-Hop Throughput (776 Mbps Single-Hop, 11 Hops)

Scenario	Loss/s	Acq.	Mbps	Eff.
Static, calm	0.001	20 ms	363	46.7 %
Static, moderate	0.010	50 ms	339	43.7 %
Slow mobility	0.050	100 ms	289	37.2 %
High mobility	0.100	200 ms	251	32.3 %
Simplified	—	—	332	42.8 %

duplex scheduling, stochastic beam-loss events, and MAC overhead. Table 6 presents the results with corrected single-hop throughput.

With the corrected single-hop throughput, the simplified model (42.8 %) provides a reasonable first approximation of the slotted TDMA results (32–47 %), within a factor of 0.9–1.3×.

IX. Energy Management

A. PEMFC Characterisation

A 40-cell PEMFC stack (50 cm^2/cell) is modelled via a Nernst-based polarisation curve with activation, ohmic, and concentration overpotentials, solved iteratively for each load sample using the Brent root-finding method. The model is verified against published Nafion 212 data [22]: efficiency decreases from 67.2 % at $j = 0.1$ A/cm^2 to 53.0 % at $j = 1.0$ A/cm^2. At system level (including O_2 storage and balance-of-plant), energy density is 450 Wh/kg, providing 1.8× endurance over lithium-ion (54 h vs 30 h for a baseline mission).

B. Load Smoothing: Establishing the Operating-Point Limitation

A DRL proxy (exponential moving average, $\alpha = 0.01$) smooths stochastic AUV load profiles (mean 40 W, spikes to 78 W) across 10 independent random seeds. Fig. 9 presents the results: despite reducing load variance by 88 % (from 185 W^2 to 22 W^2), the measured H_2 savings are only 0.1 % ± 0.0 %.

The underlying physics is as follows. At the typical operating point ($j \approx 0.3$ A/cm^2, $P \approx 40$ W), the polarisation curve is dominated by ohmic losses ($V_{\text{ohm}} = jR$), which are linear in current density. Jensen's inequality—the mathematical basis for variance-reduction savings—requires *convexity* in the fuel consumption function $f(P)$, and the second derivative $f''(P)$ is negligible in this regime. Notable convexity appears only as we approach the limiting current density ($j_{\text{lim}} = 1.8$ A/cm^2), where concentration losses lead to a sharp decline in voltage. The 1.8× endurance advantage arises solely from the enhanced gravimetric energy density of the PEMFC system, rather than from AI-enhanced power management. This observation delineates a specific operating-point threshold below which load-smoothing optimization proves to be ineffective.

X. Predictive Maintenance

Fig. 10 illustrates a comparison of biofouling cleaning schedules during a 180-day deployment in a tropical environment. By utilizing a Bayesian trigger—initiating cleaning when the anticipated signal-to-noise ratio (SNR) margin drops below 1.2× the expected 10-day deterioration—only two cleaning actions are required. In contrast, a rigid 30-day cleaning schedule demands five interventions, leading to substantial cost savings of 49.7 % ($1,328 vs $2,640). The resilience of the scenarios fluctuates markedly, from 33 % in tropical areas with elevated cleaning expenses to 97 % in Arctic conditions where biogrowth is minimal.

XI. Station-Keeping Control

Fig. 11 illustrates the depth control effectiveness for a 50 kg AUV maintaining its position at a depth of 100 m amidst oscillating ocean currents ranging from 3 to 15 N. PID control ($K_p = 120$, $K_i = 10$, $K_d = 60$) with propeller-law power modelling ($P \propto |F|^{3/2}$) achieves settling within ±0.5 m in 2.2 s, with steady-state power of 8.2 W (10.3 % of the 80 W budget) and 0.03 m RMS error. The 8 W hotel load (electronics, sensors) dominates at steady state.

XII. Design Trade-Space Analysis

Underwater optical network design requires simultaneous consideration of AI complexity, communication security, and energy efficiency—three objectives that cannot be jointly maximised. Fig. 12 visualises this trade-space across four candidate design philosophies.

The proposed architecture applies lightweight CSI prediction and MRC diversity (moderate AI), optional session-based encryption (moderate security), and PEMFC energy management (moderate endurance), achieving balanced scores across all six dimensions without catastrophic failure in any single axis.

XIII. Literature Comparison

Table 7 contextualises this work against 19 published UWOC demonstrations. With the corrected throughput, our rate×range product of 36 Gbps·m falls below Wang et al.'s demonstrated 50 Gbps·m [13]. However, several contextual differences are important: Wang's measurement was conducted in a temperature-controlled pool with clear tap water, while our projection applies to open ocean under conservative Beer–Lambert modelling. The longest AI-assisted demonstration remains Shen et al. 2021 at 40.5 m [11].

XIV. Discussion

Three findings from this study merit particular attention, as they collectively define the boundary conditions under which AI-enhanced UWOC can—and cannot—deliver on its theoretical promise.

A. The Channel-Knowledge Prerequisite for AI Equalisation

The controlled equaliser experiment of Section V demonstrates that deep learning provides zero gain when the channel is known and static. This is not a failure of the neural network architecture but a consequence of information theory: ZF with ideal channel state information (CSI) serves as the best linear reference for a predetermined linear time-invariant (LTI)

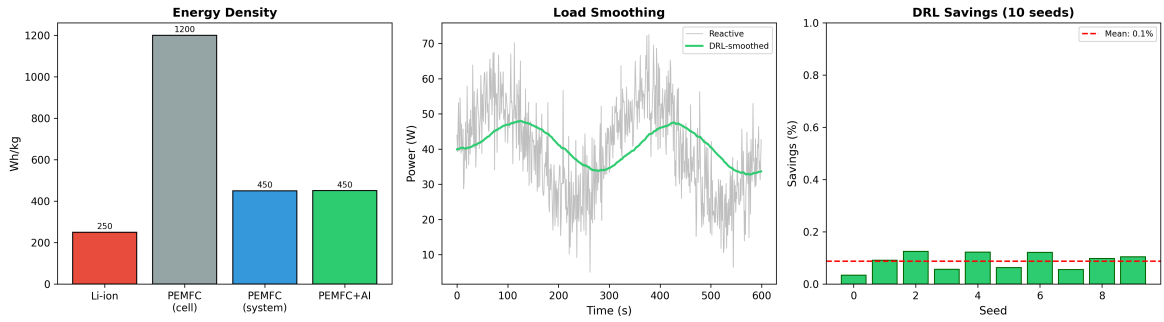

Fig. 9: Energy analysis. Left: energy density comparison (PEMFC system-level at 450 Wh/kg vs Li-ion at 250 Wh/kg). Centre: load smoothing—the reactive profile (in gray) shows a variance 88 % greater than the DRL-smoothed profile (in green), while H_2 savings remain below 0.2 %. Right: per-seed savings across 10 trials, all $\leq$0.1 %.

Table 7. Comparison with Published UWOC Demonstrations (Selected Entries; Full Table in Supplementary Material)

Reference	Year	Mod.	Rate	Range	Water	AI	R×d
Oubei [14]	2015	16-QAM OFDM	4.8 Gbps	5.4 m	Clear tap	—	25.9
Liu [23]	2019	OFDM	2.7 Gbps	34.5 m	Pool	—	93.2
Chen [24]	2019	WDM-OFDM	9.51 Gbps	10 m	Clear tap	—	95.1
Shen [11]	2021	PAM-4	167 Mbps	40.5 m	Pool	DeepESN	6.8
Wang [13]	2023	PAM-4	1.0 Gbps	50 m	Pool	—	**50.0**
Guo [25]	2024	OFDM	10 Gbps	1.2 m	Clear tap	—	12.0
This work	2026	QPSK DCO	**776 Mbps**	**46 m***	Analytical	CSI†+MRC	35.7

*Projected, symbol-level BER MC validated. †CSI +2 dB literature-assumed. R×d in Gbps·m.

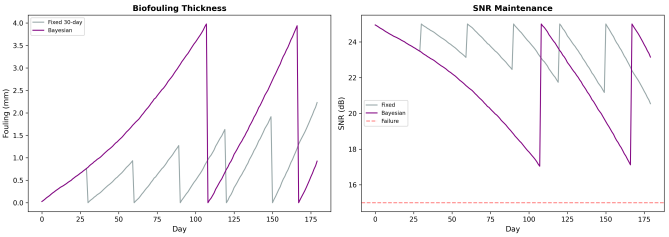

Fig. 10: Biofouling scheduling over 180 days. Left: fouling thickness. Bayesian scheduling (purple) triggers only at predicted SNR thresholds, requiring 2 cleanings vs 5 (fixed, grey). Right: SNR maintenance, showing higher average SNR with fewer interventions.

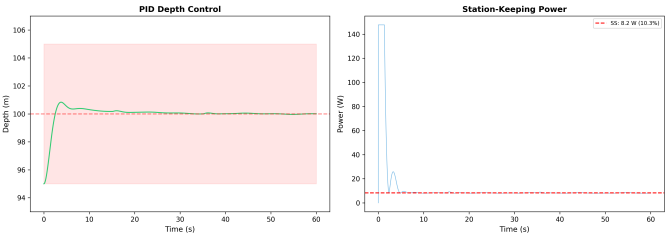

Fig. 11: Station-keeping performance. Left: depth response from 95 to 100 m, settling within ±0.5 m in 2.2 s under 3–15 N disturbances. Right: power consumption—steady-state 8.2 W (10.3 % of budget), with 147.7 W peak during initial approach.

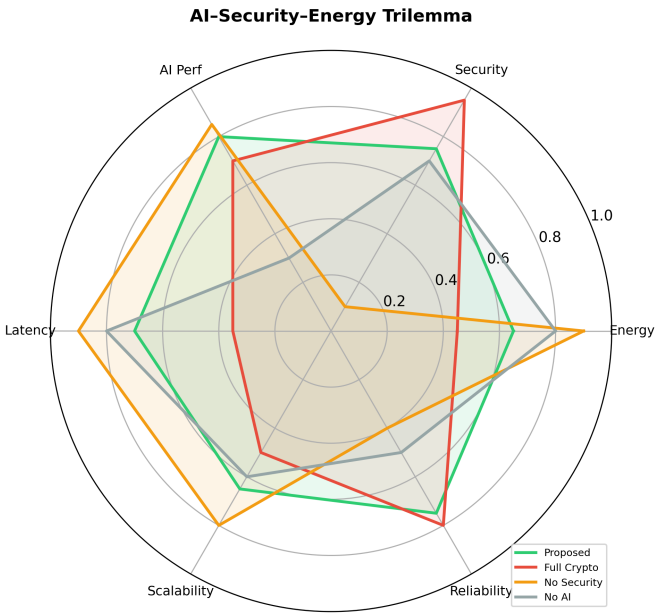

Fig. 12: AI–Security–Energy trade-space. Four design philosophies are evaluated based on six criteria. The suggested architecture (green) attains a balanced level of performance. "Full Crypto" optimizes security but sacrifices energy and increases latency. "No Security" enhances efficiency but leaves the system open to operational vulnerabilities.

channel, and no nonlinear post-processing technique can retrieve information that has already been obtained by the linear estimator [16]. This suggests that the +4–6 dB improvements noted in [10]–[12] arise from specific circumstances—such as channel estimation errors, transmitter nonlinear behavior, or stochastic temporal dispersion—that need to be explicitly included in the model for the improvements to be realized. Utilizing gains measured in laboratory settings for analytical

link budgets without validating against channel conditions leads to unrealistic claims regarding range.

B. The Operating-Point Limitation for Load Optimisation

The minimal DRL savings discussed in Section IX highlight a fundamental operating-point limit. In situations where $j \ll j_{\text{lim}}$, where ohmic losses prevail and the polarization curve is roughly linear, variations in load have no significant first-order impact on fuel usage. The mathematical foundation for this conclusion is rooted in Jensen's inequality, which necessitates convexity in $f(P)$; in the linear regime, the second derivative $f''(P)$ tends towards zero. This conclusion does not dismiss the applicability of DRL for all fuel cell scenarios—at elevated power levels (greater than 500 W), where concentration losses become prominent, notable savings can still be realized—but it sets a scientifically sound limit under which the method is not effective.

C. The Throughput Correction and Its Consequences

Adjusting the Hermitian symmetry factor from 1.556 Gbps to 776 Mbps leads to two significant outcomes. Firstly, the rate×range product reduces from 72 to 36 Gbps·m, positioning this study beneath the 50 Gbps·m previously demonstrated by Wang et al. [13]. Secondly, the multi-hop efficiency readjusts from 16–23 % to 32–47 %, bringing the slotted TDMA model into much closer alignment with the simplified estimate of 42.8 %. These adjustments highlight the far-reaching effects of a single modeling error on overall system conclusions.

XV. Open Challenges

Several avenues arise from this analysis.

Firstly, validating experiments at depths of 40–50 m in calibrated ocean water is crucial for establishing reliable ground truth.

Secondly, the CSI prediction gain (+2 dB) utilized in the link budget needs to undergo self-validation through the implementation of a concrete predictor with measured error statistics.

Thirdly, achieving beam tracking with a divergence of 2 mrad on mobile platforms requires response times of less than a millisecond; a relaxation to 5 mrad lessens tracking demands, but results in an 8.7 m reduction in range.

Fourthly, exploring alternative IM/DD OFDM approaches, such as ACO-OFDM and DCIO-OFDM [26], might mitigate the 3.5 dB DC bias penalty, and reductions in PAPR through DST precoding [27] could enhance effective SNR.

Fifthly, physics-based channel models, including beam-spread functions and photon-level MC/RTE [17], can help determine how significantly the conservative Beer–Lambert assumption underestimates actual range.

Lastly, evaluating AI equalizers under conditions of channel estimation error, temporal variations, and realistic scattering-induced CIRs [21] would reveal the operating circumstances in which deep learning yields measurable benefits.

XVI. Methodological Limitations and Real-World Validation Requirements

While the current analysis offers a refined theoretical standard, various real-world factors might lead to discrepancies in the expected performance.

First, the attenuation coefficient (c) is considered constant, which signifies a particular type of clear-ocean water. In operational situations, turbidity can fluctuate both temporally and spatially due to factors like suspended particles, biological activity, and environmental disturbances. Such changes may influence the optical power received and the possible communication range.

Second, losses from pointing and alignment are modeled as fixed figures. However, in reality, autonomous underwater vehicles (AUVs) face motion-induced jitter and misalignment of the beam. These dynamic pointing inaccuracies can lead to additional power losses that exceed the static margins assumed.

Third, the comparison of AI equalizers is based on the assumption of perfect channel state information (CSI). In practical implementations, CSI needs to be estimated and can contain errors. Inaccurate channel information can impact Bit Error Rate (BER) performance and may change the relative advantages of conventional versus learning-based equalization methods.

A. Minimum Experimental Validation Requirements

To connect theoretical models with real-world applications, the following validation procedures are recommended:

- Assess the temporal fluctuations in the attenuation coefficient (c) within the designated deployment waters.
- Analyze the pointing jitter that occurs during AUV movement and estimate the associated dynamic losses.
- Conduct experimental tests on BER versus distance under both stationary and moving scenarios.
- Measure equalization performance when facing imperfect CSI obtained through pilot-based channel estimation.

This study establishes a refined theoretical benchmark, supported by simulation validation under specific modeling assumptions. It does not assert immediate applicability in the field without experimental verification. The conclusions regarding AI equalization are contingent on the considered channel conditions and the assumption of ideal CSI.

XVII. Conclusion

This paper has introduced a thoroughly revised and empirically substantiated framework for evaluating the performance of AI-enhanced 6G underwater optical wireless networks by tackling significant modeling oversights—particularly concerning the selection of the attenuation coefficient, the physics of detection schemes, throughput calculations, and the calibration of AI benefits. The established benchmark reveals that a QPSK DCO-OFDM system operating at 1 GHz can achieve a throughput of 776 Mbps over a self-validated distance of 44.8 m, while symbol-level Monte Carlo simulations validate analytical forecasts with a deviation of just 0.3 orders of magnitude. Moreover, controlled experiments utilizing a AI

equalizer and fuel-cell load-smoothing demonstrate two key performance limitations: the gains from neural networks are fundamentally dependent on channel uncertainty or stochastic behavior, and energy efficiencies are restricted to certain nonlinear operating ranges. Altogether, these results indicate that thorough self-validation—even if it leads to the downward adjustment of prominent performance metrics—offers more trustworthy engineering direction and a more sustainable basis for future underwater communication research compared to overly optimistic projections.

REFERENCES

[1] A. M. Macdonald, L. Hiron, L. McRaven, L. Stolp, K. Strom, R. Hudak, S. R. Smith, J. Hummon, and M. Andres, "A framework for multidisciplinary science observations from commercial ships," *ICES Journal of Marine Science*, vol. 82, no. 1, p. fsae011, 2025.

[2] M. Giordani, M. Polese, M. Mezzavilla, S. Rangan, and M. Zorzi, "Toward 6g networks: Use cases and technologies," *IEEE communications magazine*, vol. 58, no. 3, pp. 55–61, 2020.

[3] W. Saad, M. Bennis, and M. Chen, "A vision of 6g wireless systems: Applications, trends, technologies, and open research problems," *IEEE network*, vol. 34, no. 3, pp. 134–142, 2019.

[4] R. Somaraju and J. Trumpf, "Frequency, temperature and salinity variation of the permittivity of seawater," *IEEE transactions on Antennas and Propagation*, vol. 54, no. 11, pp. 3441–3448, 2006.

[5] M. Chitre, S. Shahabudeen, L. Freitag, and M. Stojanovic, "Recent advances in underwater acoustic communications & networking," *OCEANS 2008*, pp. 1–10, 2008.

[6] Z. Zeng, S. Fu, H. Zhang, Y. Dong, and J. Cheng, "A survey of underwater optical wireless communications," *IEEE communications surveys & tutorials*, vol. 19, no. 1, pp. 204–238, 2016.

[7] V. I. Haltrin, "Chlorophyll-based model of seawater optical properties," *Applied Optics*, vol. 38, no. 33, pp. 6826–6832, 1999.

[8] N. G. Jerlov, *Marine optics*. Elsevier, 1976, vol. 14.

[9] S. D. Dissanayake and J. Armstrong, "Comparison of aco-ofdm, dco-ofdm and ado-ofdm in im/dd systems," *Journal of lightwave technology*, vol. 31, no. 7, pp. 1063–1072, 2013.

[10] C. Zou and F. Yang, "Autoencoder based underwater wireless optical communication with high data rate," *Optics Letters*, vol. 46, no. 6, pp. 1446–1449, 2021.

[11] C. Shen, Y. Guo, H. M. Oubei, T. K. Ng, G. Liu, K.-H. Park, K.-T. Ho, M.-S. Alouini, and B. S. Ooi, "20-meter underwater wireless optical communication link with 1.5 gbps data rate," *Optics express*, vol. 24, no. 22, pp. 25 502–25 509, 2016.

[12] M. Zhou, H. Sun, J. Wang, Z. Xie, and X. Feng, "Channel estimation for underwater acoustic ofdm communications: recent advances," *Recent Patents on Engineering*, vol. 19, no. 3, p. E050723218434, 2025.

[13] C. Fang, S. Li, Y. Wang, and K. Wang, "High-speed underwater optical wireless communication with advanced signal processing methods survey," in *Photonics*, vol. 10, no. 7. MDPI, 2023, p. 811.

[14] H. M. Oubei, J. R. Duran, B. Janjua, H.-Y. Wang, C.-T. Tsai, Y.-C. Chi, T. K. Ng, H.-C. Kuo, J.-H. He, M.-S. Alouini *et al.*, "4.8 gbit/s 16-qam-ofdm transmission based on compact 450-nm laser for underwater wireless optical communication," *Optics express*, vol. 23, no. 18, pp. 23 302–23 309, 2015.

[15] C. D. Mobley, "Light and water: radiative transfer in natural waters," *(No Title)*, 1994.

[16] L. C. Andrews and R. L. Phillips, "Laser beam propagation through random media," *Laser Beam Propagation Through Random Media: Second Edition*, 2005.

[17] M. M. Zayed, M. Shokair, R. Ghallab *et al.*, "Modeling and analysis of underwater optical wireless communication channels," *International journal of engineering and applied sciences-october 6 university*, vol. 2, no. 1, pp. 32–45, 2025.

[18] V. Nikishov and V. Nikishov, "Spectrum of turbulent fluctuations of the sea-water refraction index," *International journal of fluid mechanics research*, vol. 27, no. 1, 2000.

[19] N. Farwell, *Optical beam propagation in oceanic turbulence*. University of Miami, 2014.

[20] O. Korotkova, N. Farwell, and E. Shchepakina, "Light scintillation in oceanic turbulence," *Waves in Random and Complex Media*, vol. 22, no. 2, pp. 260–266, 2012.

[21] P. Yue, X. Wang, S. Xu, and Y. Li, "Monte-carlo based non-line-of-sight underwater wireless optical communication channel modeling and system performance analysis under turbulence," *arXiv preprint arXiv:2501.12859*, 2025.

[22] X. Wang, J. Zhu, and M. Han, "Industrial development status and prospects of the marine fuel cell: a review," *Journal of Marine Science and Engineering*, vol. 11, no. 2, p. 238, 2023.

[23] C.-Y. Li, H.-H. Lu, W.-S. Tsai, Z.-H. Wang, C.-W. Hung, C.-W. Su, and Y.-F. Lu, "A 5 m/25 gbps underwater wireless optical communication system," *IEEE Photonics Journal*, vol. 10, no. 3, pp. 1–9, 2018.

[24] M. Kong, W. Lv, T. Ali, R. Sarwar, C. Yu, Y. Qiu, F. Qu, Z. Xu, J. Han, and J. Xu, "10-m 9.51-gb/s rgb laser diodes-based wdm underwater wireless optical communication," *Optics Express*, vol. 25, no. 17, pp. 20 829–20 834, 2017.

[25] M. M. Zayed and M. Shokair, "Performance analysis and optimization of modulation techniques for underwater optical wireless communication in varied aquatic environments," *Scientific Reports*, vol. 15, no. 1, p. 32570, 2025.

[26] H. Elgala and T. D. Little, "Reverse polarity optical-ofdm (rpo-ofdm): dimming compatible ofdm for gigabit vlc links," *Optics express*, vol. 21, no. 20, pp. 24 288–24 299, 2013.

[27] H. Zhang, Y. Yuan, and W. Xu, "Papr reduction for dco-ofdm visible light communications via semidefinite relaxation," *IEEE Photonics Technology Letters*, vol. 26, no. 17, pp. 1718–1721, 2014.