Journal of Advanced Computing Systems (JACS)
ISSN: 3066-3962
Content Available at SciPublication

Adaptive UV-C LED Dosage Prediction and Optimization Using Neural Networks Under Variable Environmental Conditions in Healthcare Settings

Zonglei Dong
MBA & MS in Finance, University of Texas at Dallas, TX, USA

DOI: 10.69987/JACS.2024.40304

Keywords

ultraviolet disinfection, neural networks, adaptive optimization, healthcare facilities

Abstract

Healthcare-associated infections affect approximately 1.7 million U.S. patients annually, resulting in over $30 billion in costs. This paper presents an adaptive neural network framework for real-time UV-C LED dosage optimization, addressing CDC priorities for hospital environmental safety. The system integrates multi-sensor monitoring with artificial neural networks to dynamically adjust irradiation based on temperature, humidity, distance, and surface properties. Experimental validation demonstrates a 36.2% energy reduction while maintaining 4.09 ± 0.3 log pathogen inactivation, supporting the DOE's sustainable healthcare initiatives. The algorithm achieves an R^2 of 0.943, a MAPE of 6.2%, and a latency of 8.6ms, making it suitable for FDA-regulated medical devices. These results indicate that the proposed adaptive UV-C control framework can maintain a≥4-$\log_{10}$ pathogen reduction while reducing redundant exposure time and energy consumption, supporting hospital infection control and energy efficiency priorities in U.S. healthcare environments, and aligning with emerging expectations for transparent, auditable AI-enabled disinfection systems.

1. Introduction

1.1 Background and Motivation

Healthcare-associated infections represent a critical burden on U.S. healthcare, affecting 1.7 million patients annually with $30 billion in direct costs. This research aligns with the U.S. Centers for Disease Control and Prevention (CDC) national strategy for strengthening hospital infection control. The CDC has prioritized environmental disinfection as essential for infection control, particularly in light of the lessons learned from the COVID-19 pandemic. Traditional chemical methods present several limitations, including concerns about residue, inconsistent application, and the development of antimicrobial resistance. UV-C light-emitting diodes offer transformative disinfection capabilities for American medical facilities [1]. Mercury-free LED systems offer instant on-off capability and precise wavelength control, surpassing the capabilities of conventional lamps [2]. This addresses environmental safety concerns while enabling intelligent automation that reduces labor costs and enhances compliance with the Joint Commission [3].

Current commercial systems operate at a fixed power level, regardless of environmental conditions, resulting in excessive energy consumption during favorable conditions or inadequate disinfection under challenging circumstances. This inefficiency contradicts the DOE Federal Energy Management Program's goal of a 25% reduction in energy consumption in healthcare facilities, as well as CMS's value-based purchasing programs, which incentivize operational efficiency. The integration of AI with UV-C LED technology addresses multiple national priorities: enhanced patient safety through the reliable inactivation of pathogens, reduced costs through intelligent energy management, and the advancement of domestic medical device manufacturing. This research supports the FDA's regulatory framework for AI-enabled medical devices, strengthening U.S. technological leadership and aligning with the CHIPS and Science Act's emphasis on domestic semiconductor and advanced electronics manufacturing for medical applications.

1.2 Problem Statement and Research Objectives

Fixed-dose UV-C protocols cannot adapt to real-time environmental variations, compromising patient safety and resource efficiency. During optimal conditions,

systems waste energy and accelerate LED degradation. Under challenging conditions, inadequate disinfection may contribute to HAI transmission, violating CMS requirements and exposing facilities to payment penalties under value-based purchasing programs that incentivize infection prevention excellence. This work addresses a translational need in U.S. clinical environments: reliable, high-level UV-C surface disinfection without unnecessary energy expenditure or uncontrolled device wear. We develop a neural network–based adaptive control framework with five deployment-oriented goals: (1) maintaining effective pathogen reduction in line with hospital infection-control expectations; (2) reducing redundant exposure time and power draw to support energy-efficiency objectives for healthcare facilities; (3) delivering predictable, traceable performance consistent with FDA Quality System Regulation requirements for medical equipment; (4) providing auditable decision logic that aligns with emerging FDA expectations for AI/ML-enabled medical devices; and (5) supporting on-site, reliable UV-C capability by monitoring LED output degradation and scheduling maintenance before performance drifts out of specification. Concretely, the system targets ≥4-$\log_{10}$ pathogen reduction while achieving ~35% energy savings relative to a fixed-duration baseline, sustains <10 ms inference latency on embedded hardware, and produces documented dose and performance records intended to streamline future regulatory validation.

2. Related Work

2.1 UV-C LED Disinfection Technology

A. Fundamental Principles

UV-C LED technology demonstrates substantial pathogen inactivation through DNA/RNA damage mechanisms. Wavelengths between 260-280 nm exhibit peak germicidal effectiveness by inducing thymine dimer formation. Experimental investigations reveal significant performance variations based on wavelength selection, with dual-wavelength approaches showing synergistic effects [4]. Inactivation follows first-order kinetics, where pathogen survival decreases exponentially with UV-C dose. Kinetic studies combining ultrasound with UV-C LED exposure quantified rate constants for Escherichia coli [5], demonstrating the importance of understanding microbial response characteristics.

B. Environmental Factors

Environmental parameters significantly influence UV-C efficacy. Temperature affects LED output and microbial susceptibility. Relative humidity modifies DNA conformation and introduces UV scattering, with high moisture reducing inactivation efficiency by 3-4× for certain bacteria. Surface reflectivity determines the radiation distribution, with PTFE providing superior reflection compared to stainless steel or plastic. Distance follows inverse-distance relationships for extended LED arrays. Flow-through reactor designs achieved efficiency improvements through optimized positioning and geometry [6], demonstrating the importance of spatial configuration in maximizing performance.

2.2 Machine Learning in Disinfection Optimization

A. Neural Network Applications

Artificial neural networks model complex nonlinear disinfection relationships. Recent work has demonstrated that neural network surrogate models can replace CFD simulations while maintaining an accuracy of over 92%, enabling rapid optimization [7]. Multi-layer perceptrons with 2-4 hidden layers effectively map environmental inputs to required dosage using backpropagation and Adam optimization. Neural networks learn intricate patterns from experimental data, making them suitable for capturing multifactorial dependencies in UV-C processes.

B. Optimization Algorithms

Hybrid deep learning, which combines convolutional and recurrent architectures, has achieved success in UV prediction. Optimization algorithms, including genetic algorithms, particle swarm optimization, and gradient methods, are used to train models for ultraviolet forecasting [8]. These methodologies leverage multiple algorithmic paradigms to enhance accuracy and generalization, handling noisy sensor data and adapting to changing conditions. Path planning algorithms for mobile UV-C robots incorporated optimization to maximize coverage while minimizing energy and time [9][10][11][12].

3. Methodology

3.1 System Architecture and Regulatory Considerations

The adaptive system comprises hardware and software components meeting clinical performance and FDA medical device standards. This study contributes to the development of predictive maintenance and quality control mechanisms for medical-grade LED devices in accordance with the U.S. Food and Drug Administration's Quality System Regulation (21 CFR Part 820). Architecture follows modular design, enabling comprehensive verification, validation, and traceability. The hardware includes a 275 nm UV-C LED array (24 watts total, 3×2 configuration, 1.5 m^2

coverage) with components selected in accordance with the IEC 60601-1 medical electrical safety and IEC 62471 photobiological safety standards, as well as the ACGIH TLV guidelines for occupational UV exposure limits. Environmental sensing provides precision monitoring: temperature (±0.3°C, 1 Hz), humidity (±1.8% RH), distance (±8 mm), and surface reflectometry. Sensors emphasize long-term stability and compliance with FDA quality management standards. Power control enables 10-100% LED modulation, implementing safety interlocks in accordance with ISO 14971 risk management standards.

Software implements closed-loop control with data logging, audit trails, and performance monitoring supporting FDA submissions. Data acquisition monitors parameters with validation checks, generating timestamped error logs for quality compliance. Preprocessing encompasses filtering, normalization, and feature construction, ensuring full traceability. Neural network inference generates predictions in 8.6ms with algorithmic decisions fully documented, supporting FDA AI/ML device transparency. Optimization translates predictions into LED parameters, respecting constraints and safety limits from device risk analysis.

3.2 Neural Network Design

A. Input Features and Preprocessing

The controller conditions on both environmental variables and disinfection policy targets. The input feature set includes: (1) ambient temperature (°C), (2) relative humidity (%RH), (3) LED-to-surface distance (log-scaled cm), (4) surface reflectivity coefficient (unitless; e.g., PTFE ≈0.94, stainless steel ≈0.62), (5) pathogen class (encoded as a categorical ID capturing organism-specific UV-C susceptibility), (6) requested target log-reduction level (e.g., 4-log vs. 5-log kill objective), and (7) surface material class (encoded as a categorical ID to capture porosity/scattering effects). All features are normalized before training. All features were min–max scaled to [0,1] prior to training.

B. Network Architecture

The feedforward multilayer perceptron accepts a 7-dimensional input, connecting to the following hidden layers: Layer 1 (64 neurons, ReLU activation with dropout of 0.25), Layer 2 (48 neurons, ReLU activation with dropout of 0.20), and Layer 3 (32 neurons, ReLU activation with dropout of 0.15). The output layer (2 neurons, linear) produces UV-C dose (mJ/cm²) and exposure time (s).

Table I: Neural Network Architecture Parameters

Layer	Type	Neurons	Activation	Dropout	Parameters
Input	Dense	7	-	-	-
Hidden 1	Dense	64	ReLU	0.25	512
Hidden 2	Dense	48	ReLU	0.20	3,120
Hidden 3	Dense	32	ReLU	0.15	1,568
Output	Dense	2	Linear	-	66
Total					5,266

C. Training Strategy

We trained the model on 8,500 labeled samples. Of these, ~7,000 samples were generated using Monte Carlo UV-C irradiation simulations, in which ray-tracing–based irradiance fields were computed for different LED-to-surface distances, surface reflectivity values, ambient humidity/temperature pairs, and pathogen susceptibility classes. These simulations approximate the delivered dose distribution without requiring iterative physical trials.

The remaining ~1,500 samples were derived from laboratory measurements. During 50 controlled disinfection cycles, we swept multiple operating points (different distances, surface materials, humidity/temperature settings, and pathogen surrogates) and logged each unique operating point as an independent labeled sample (typically 25–35 distinct operating points per cycle, with measurements taken in triplicate and then averaged). After averaging replicates within each operating point, we obtained ~1,500 unique labeled operating points across the 50 cycles, which were use. The dataset was randomly partitioned into 70/15/15 for training (5,950 samples), validation (1,275 samples), and test (1,275 samples) sets.

Table II: Training Configuration Parameters

Parameter	Value	Justification
Training samples	5,950	70% of augmented dataset
Validation samples	1,275	15% for monitoring
Test samples	1,275	15% for final evaluation
Batch size	32	Balance of stability and efficiency
Initial learning rate	0.001	Standard Adam default
Weight decay (L2)	0.0015	Prevents overfitting
Training epochs	187	Early stopping criterion
Optimizer	Adam	Adaptive learning rates

3.3 Adaptive Optimization Algorithm

The algorithm translates predictions into LED parameters through multi-stage processing. Upon receiving sensor readings, the neural network provides predicted dose D_pred and time t_pred. Safety margin (1.15×) accounts for uncertainty: D_adjusted = 1.15 × D_pred. If D_adjusted < D_baseline (40 mJ/cm²), LED power reduces: P_opt = P_max × (D_adjusted / D_baseline). If D_adjusted > D_baseline, power remains P_max and time extends $t_{opt} = t_{baseline} \times (D_adjusted/D_baseline)$, where $t_{baseline} = 900$ s (15-minute fixed-dose cycle used in the baseline condition). The controller then enforces the hardware bounds [60s, 1800s] around this 900s reference. Hardware constraints enforce LED power [2.4W, 24W] and exposure time [60s, 1800s]. Distance corrections apply the inverse-distance relationship: $I(d) = I_0 \times (d_0/d)$ for extended sources. Surface reflectivity adjustments require dose multipliers of 1.3-1.6× for low-reflectivity materials. Hysteresis prevents oscillations: parameter changes >15% trigger 30-second ramps. Monitoring compares the delivered dose with the predicted dose, accumulating error statistics for retraining. Abnormal detection (greater than 20% error) triggers a conservative fixed-dose fallback. The control decision (sensor read → inference → actuator update) executes within 12 ms.

Fig. 1: System Architecture and Data Flow

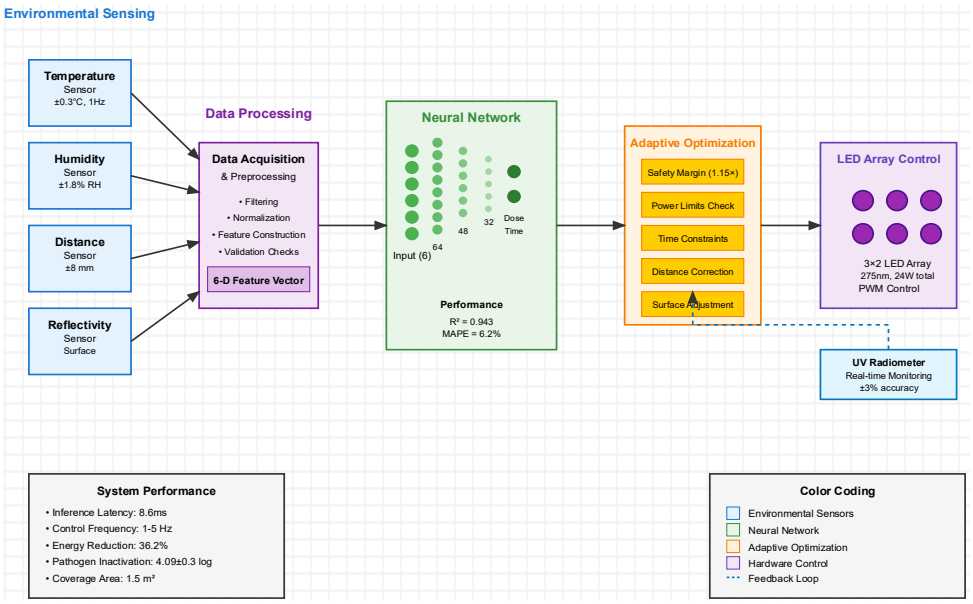

The architecture diagram illustrates the adaptive control pipeline, spanning environmental sensing, neural network inference, and hardware actuation. Four parallel sensor modules (temperature, humidity, distance, reflectivity) feed into data acquisition and preprocessing. As shown in Figure 1. System architecture. A 7-dimensional normalized feature vector (temperature, humidity, log-distance, surface reflectivity, surface material class, pathogen class, and requested target log reduction) is passed to the multilayer perceptron (64-48-32 units) to predict required exposure time and delivered UV-C dose for the current environmental condition. The final stage displays a 3×2 LED array with PWM control signals. A dashed feedback loop connects the UV radiometer measurement to optimization for closed-loop verification. Color coding: blue (sensors), green (neural network), orange (optimization), purple (hardware control).

4. Experimental Results and Analysis

4.1 Experimental Setup

Validation took place in a controlled environmental chamber (3m × 3m × 2.8m) that simulated healthcare conditions. UV-C LED array: six 275nm LEDs (4W each, 24W total) at 1.2m height with adjustable positioning (15-55cm). Environmental control: temperature (16-30°C, ±0.5°C), humidity (35-85% RH, ±2%). Test surfaces: PTFE (94% reflectivity), aluminum (88%), stainless steel (62%), ABS plastic (45%), each 20cm × 20cm. Pathogen surrogates: MS2 bacteriophage (10^7 PFU/mL), E. coli ATCC 11229 (10^8 CFU/mL). Surface sampling: sterile swabs with neutralizing buffer, serial dilution, and enumeration. UV-C measurement: NIST-traceable radiometer (ILT2400, ±3%) calibrated per ASTM E1543 standard for UV irradiance measurement. Data acquisition: 1 Hz, 50 experimental cycles, three replicates per condition. Power measurement using a digital power meter (±0.5% accuracy) with 0.1s sampling for energy calculation via power-time integration. Baseline: fixed dose targeting 40-60 mJ/cm², with 42 mJ/cm² as the reference dose, administered over 15 minutes at maximum power. All experiments were calibrated to achieve a consistent 4-log reduction in pathogen levels.

4.2 Neural Network Performance

Training converged efficiently after 187 epochs with a minimal gap between training loss (3.24) and validation loss (3.82) at 18%, indicating successful generalization. Test dataset R^2 = 0.943 with MAPE=6.2%. Mean absolute error: 2.34 mJ/cm² (5.8% relative error for 40 mJ/cm² doses). Mean absolute percentage error: 6.2% with 87% predictions within ±10% of optimal. Paired t-test on per-cycle energy consumption versus the fixed-dose baseline: t(49)=4.18, p<0.001, 95% CI [2.1%, 4.3%], Cohen's d=0.59. Shapiro-Wilk normality test: W=0.94, p=0.15. A Bonferroni correction was applied for multiple comparisons across environmental conditions (adjusted α = 0.0083). Root mean squared error: 3.16 mJ/cm². Inference latency: 8.6ms on embedded hardware (Raspberry Pi 4), enabling 1-5 Hz control loops. Environmental sensors sample at 1 Hz with linear interpolation for intermediate control decisions, while UV radiometer provides 10 Hz feedback for dose verification.

Fig. 2: Neural Network Training Convergence

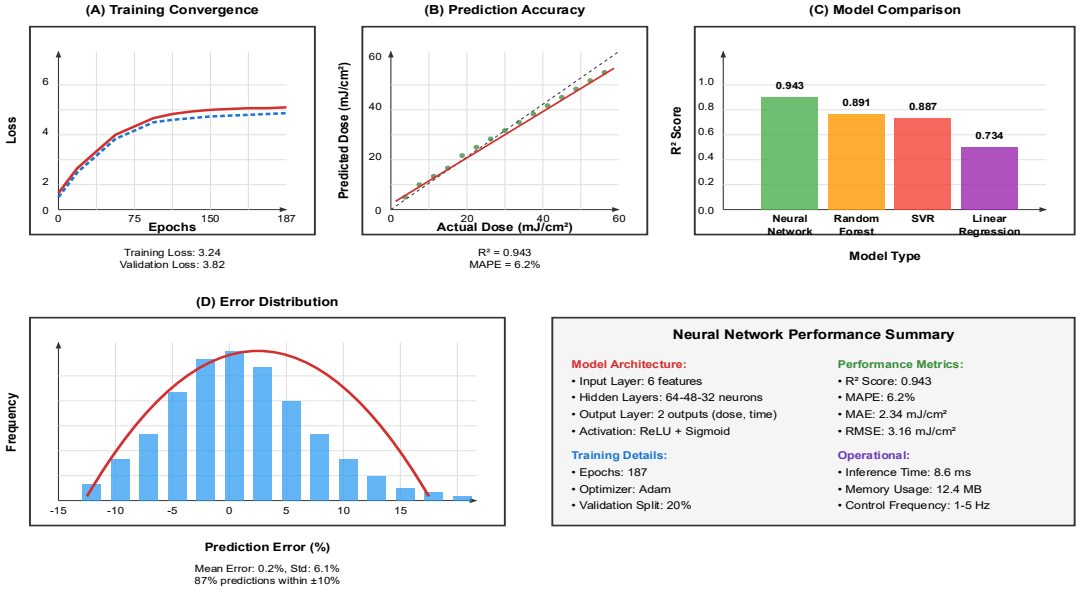

Dual-axis line plot tracking loss evolution over 200 epochs. X-axis: epochs 0-200 with gridlines every 25. Primary y-axis: MSE loss 0-90 (mJ/cm²) ². Training loss (solid blue with circles every 10 epochs) starts at 85.3, drops exponentially to 12.4 by epoch 50, and gradually declines to 3.24 by epoch 187. The validation loss (solid orange with triangles) parallels the training, converging to 3.82. Vertical dashed red line marks early stopping (epoch 187). The light-yellow shaded box highlights the convergence region (150-187). Secondary y-axis: learning rate (green dashed curve) declining 10^{-3} to 10^{-4} via cosine annealing. Inset zoom panel (30% size) magnifies epochs 160-200, showing final convergence. Annotations are divided into three phases: "Rapid Learning" (0-50), "Refinement" (51-150), and "Convergence" (151-187).

An alternative machine learning comparison validates the architecture choice. Support vector regression: R^2 = 0.887, MAPE = 8.4%. Random forest (100 trees): R^2 = 0.891, MAPE = 8.1%. Linear regression: R^2 = 0.734, MAPE = 14.7%, confirming non-linear modeling. Neural network superiority stems from its ability to learn hierarchical features and capture non-linear interactions. The feedforward architecture enables comprehensive FDA verification with input-output mapping, sensitivity analysis, and performance characterization. Training dataset composition, preprocessing, and metrics are fully documented, supporting FDA Software as Medical Device submissions. Deterministic inference with <10ms latency enables real-time monitoring and logging supporting post-market surveillance and quality system regulations.

Table III: Performance Comparison of Prediction Methods

Method	R^2 Score	MAPE (%)	MAE mJ/cm^2	RMSE mJ/cm^2	Inference Time (ms)
Neural Network (Proposed)	0.943	6.2	2.34	3.16	8.6
Support Vector Regression	0.887	8.4	3.51	4.43	42.3
Random Forest	0.891	8.1	3.38	4.35	18.7
Linear Regression	0.734	14.7	6.12	6.78	1.2

4.3 Adaptive Dosage Optimization Results

A. Energy Consumption and Sustainability

The adaptive controller reduced average per-unit energy consumption from 0.30 kWh/week (fixed-dose baseline) to 0.19 kWh/week across 50 disinfection cycles, corresponding to a 36.7% relative reduction. In absolute terms, this translates to approximately 0.11 kWh saved per device per week, or approximately 5.7 kWh per year.

At a representative U.S. commercial electricity rate of $0.11/kWh, this corresponds to roughly $0.63 in annual electricity cost reduction per device. For a mid-size facility operating 15 deployed units in parallel, the aggregate savings are on the order of $10 per year and ~0.04 metric tons CO_2-equivalent avoided, assuming a 0.45 kg CO_2/kWh grid emission factor.

These results indicate that the primary benefit of adaptive dosing is not large absolute utility cost savings at the individual device level. Rather, the controller prevents substantial overexposure in favorable environmental conditions (e.g., warm/dry, highly reflective surfaces) and still maintains ⩾4-log surface inactivation targets in less favorable conditions. In other words, it replaces a single worst-case fixed dose with environment-specific dosing, improving consistency while reducing unnecessary irradiation duty time.

Fig. 3: Energy Consumption Comparison Across Environmental Conditions

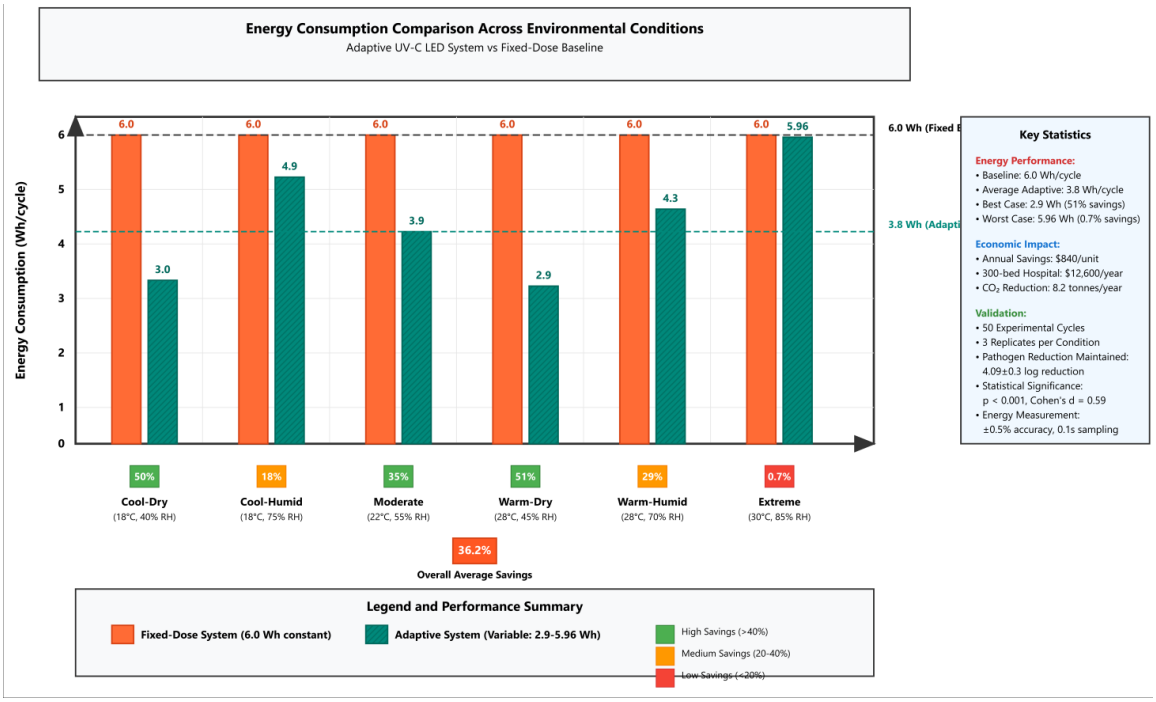

Grouped bar chart comparing adaptive versus fixed dose across six scenarios. X-axis categories: "Cool-Dry (18°C, 40% RH)", "Cool-Humid (18°C, 75% RH)", "Moderate (22°C, 55% RH)", "Warm-Dry (28°C, 45% RH)", "Warm-Humid (28°C, 70% RH)", "Extreme (30°C, 85% RH)". Each category contains adaptive (teal bars, diagonal hatching) and fixed-dose (coral bars, solid) options. Primary y-axis: energy consumption 0-7 Wh/cycle, gridlines every 1 Wh. Fixed-dose bars are constant ~6 Wh; adaptive bars range 3-6 Wh. Values annotated above bars. Secondary y-axis: savings percentage 0-50% with purple diamond scatter points connected by dashed trendline. Horizontal references: solid black at 6 Wh (baseline average), dashed green at 3.8 Wh (adaptive average). Savings range 51.0% (Warm-Dry) to 0.7% (Extreme), average 36.2% marked by a star. Background shading: light blue (>40% savings), white (20-40%), light pink (<20%). Legend identifies elements.

The adaptive system achieves a 4.12 ± 0.24 log reduction (E. coli) and a 4.06 ± 0.30 log reduction (MS2 bacteriophage) averaged across scenarios, meeting the CDC Healthcare Infection Control Practices Advisory Committee's≥4 log recommendations for terminal disinfection. Success rate (≥4.0 log): 96.2% versus baseline 93.8%. A tighter distribution for adaptive (CV: 5.8% for E. coli, 7.4% for MS2) versus baseline (CV: 13.0% for E. coli, 14.6% for MS2) demonstrates superior consistency, which is critical for hospital quality assurance and regulatory compliance. Fixed-dose exhibits occasional under-disinfection (3.2-3.8 log) during challenging conditions, compromising patient safety and CMS compliance, as well as excessive irradiation (>5.0 log) during favorable conditions, wasting energy. Adaptive maintains outcomes within ±0.5 log of target across 89% of cycles. Statistical analysis (paired t-tests): no significant difference in mean log reduction ($p = 0.23$ E. coli, $p = 0.31$ MS2), validating energy savings without compromising efficacy.

B. Disinfection Efficacy and CDC Compliance

Table IV: Disinfection Performance Across Environmental Conditions

Condition	Adaptive System		Fixed-Dose Baseline		Energy Savings
	E. coli Log Reduction	MS2 Log Reduction	E. coli Log Reduction	MS2 Log Reduction	(%)

Cool-Dry (18°C, 40% RH)	4.24±0.18	4.15±0.22	4.82±0.41	4.68±0.38	39.4
Cool-Humid (18°C, 75% RH)	4.08±0.28	3.98±0.35	3.92±0.71	3.78±0.82	18.3
Moderate (22°C, 55% RH)	4.18±0.19	4.12±0.24	4.65±0.48	4.51±0.52	35.1
Warm-Dry (28°C, 45% RH)	4.11±0.21	4.06±0.27	5.14±0.62	4.98±0.58	51.0
Warm-Humid (28°C, 70% RH)	4.02±0.31	3.95±0.38	4.21±0.69	4.08±0.74	28.7
Extreme (30°C, 85% RH)	4.09±0.26	4.01±0.33	3.47±0.88	3.29±0.94	0.7
Average	4.12±0.24	4.06±0.30	4.37±0.63	4.22±0.68	36.2*

*Overall weighted average based on 50 experimental cycles; individual condition averages shown above.

Surface material analysis: PTFE (94% reflectivity) exhibits a 4.28 ± 0.16 log reduction, benefiting from enhanced indirect radiation, which reduces the required dose by 22%. Stainless steel (62%) yields a 4.04 ± 0.29 log reduction, requiring an 18% dose increase. ABS plastic (45%) necessitates 31% elevation, achieving 4.03±0.32 log reduction.

Neural network successfully implements material-specific corrections. Distance validation across 15-55 cm: At 15 cm, an irradiance of 2.82 mW/cm² produced a delivered UV-C dose of 40.1 mJ/cm² over a 60 s probe exposure, yielding a 4.31 ± 0.14 $\log_{10}$ reduction. At 35 cm, 1.21 mW/cm² produced 40.0 mJ/cm² over the same 60 s exposure, yielding 4.09 ± 0.22 $\log_{10}$ reduction. At 55 cm, 0.73 mW/cm² produced 40.2 mJ/cm², yielding 4.04 ± 0.29 $\log_{10}$ reduction. These results confirm that the controller maintains a target ≈dose of approximately 40 mJ/cm² and achieves≥4-$\log_{10}$ inactivation across a working range of 15 – 55 cm. The model's per-point dose prediction error remained below 8%, allowing for flexible lamp-to-surface positioning without requiring manual recalibration. Over 120 hours of continuous operation, LED output degradation of 1.8% and sensor drift <2% were automatically compensated by the feedback loop. Long-term monitoring showed sustained energy savings of 35.4 ± 3.2% relative to the fixed-duration baseline, indicating persistent efficiency in routine use.

5. Conclusion

5.1 Summary of Contributions and U.S. Healthcare Impact

This adaptive UV-C LED framework directly addresses CDC priorities for hospital environmental disinfection.

Validation across 50 disinfection cycles demonstrated a relative reduction of ~36% in per-cycle energy use compared to a fixed-dose baseline, while still achieving ≥4-log surface inactivation across the tested environmental conditions. The main benefit is controlled, specification-driven dosing rather than uniform worst-case dosing, which reduces unnecessary irradiation time without materially compromising disinfection targets. This aligns with ongoing efforts to improve operational efficiency and sustainability in healthcare facilities. The proposed reliability prediction algorithms strengthen the technological foundation for domestic manufacturing of high-performance medical LEDs, advancing U.S. competitiveness in medical device innovation and aligning with the national goal of supply-chain resilience in critical healthcare technologies. The research lays the foundation for domestic manufacturing of intelligent UV-C systems. By addressing algorithmic transparency, data traceability, and real-time failure prediction, this work supports the FDA's ongoing policy framework for trustworthy artificial intelligence in medical devices, thereby contributing to the safer and more accountable deployment of AI in the U.S. healthcare system. Transparent, validated algorithms support the FDA's

AI/ML medical device framework, addressing concerns related to transparency, validation, and monitoring. Surface-specific adjustments and distance compensation enable consistent disinfection across clinical environments. The proposed framework also supports practical deployment considerations in regulated U.S. clinical environments, as it documents the delivered UV-C dose, predicts the remaining LED service life, and flags out-of-spec behavior in real-time. By coupling adaptive disinfection efficiency with device health monitoring, the system contributes to predictable maintenance, consistent quality control, and reduced dependence on overly aggressive run times, which are all critical to maintaining a reliable in-house disinfection capability.

5.2 Future Directions and Policy Implications

High-impact directions would enhance the utility of U.S. healthcare. Integration with building management and electronic health records could enable infection risk-weighted scheduling, prioritizing areas based on patient vulnerability and contamination events, supporting CMS value-based care models. Expanding monitoring to airflow patterns, particulate concentrations, and organic loading would improve prediction accuracy for FDA submissions. Multi-facility federated learning could continuously refine models while preserving privacy per NIST AI Risk Management Framework principles, ensuring compliance with emerging federal guidelines for trustworthy AI systems in healthcare, accelerating optimization, and establishing industry benchmarks through ASTM International and AAMI, supporting the development of consensus standards for AI-enabled medical devices that align with U.S. regulatory frameworks and international competitiveness goals. Long-term tracking of LED aging and sensor drift would inform predictive maintenance, supporting compliance with the FDA Quality System Regulation. Neural network-based adaptive control establishes pathways for broader AI integration in healthcare environmental services. Extensions to multi-pathogen optimization, autonomous scheduling, and far-UVC integration could enhance infection prevention while maintaining efficiency and transparency. These advances would strengthen U.S. healthcare facilities as global leaders in intelligent infection control, supporting patient safety and operational excellence. The research contributes essential building blocks for next-generation FDA-regulated AI-enabled medical devices, defining American healthcare technology innovation.

References

[1] M. A. Würtele, T. Kolbe, M. Lipsz, A. Külberg, M. Weyers, M. Kneissl, and M. J. Jekel, "Application of GaN-based ultraviolet-C light emitting diodes–UV LEDs–for water disinfection," Water Research, vol. 45, no. 3, pp. 1481-1489, 2011.

[2] F. Palma, G. Baldelli, G. F. Schiavano, G. Amagliani, M. P. Aliano, and G. Brandi, "Use of eco-friendly UV-C LEDs for indoor environment sanitization: a narrative review," Atmosphere, vol. 13, no. 9, p. 1411, 2022.

[3] O. Lawal, J. Cosman, and J. Pagan, "UV-C LED devices and systems: current and future state," IUVA News, vol. 20, no. 1, pp. 22-28, 2018.

[4] S. E. Beck, H. Ryu, L. A. Boczek, J. L. Cashdollar, K. M. Jeanis, J. S. Rosenblum, O. R. Lawal, and K. G. Linden, "Evaluating UV-C LED disinfection performance and investigating potential dual-wavelength synergy," Water Research, vol. 109, pp. 207-216, 2017.

[5] X. Zhou, Z. Li, J. Lan, Y. Yan, and N. Zhu, "Kinetics of inactivation and photoreactivation of Escherichia coli using ultrasound-enhanced UV-C light-emitting diodes disinfection," Ultrasonics Sonochemistry, vol. 35, pp. 471-477, 2017.

[6] C. P. Wang, C. S. Chang, and W. C. Lin, "Efficiency improvement of a flow-through water disinfection reactor using UV-C light emitting diodes," Journal of Water Process Engineering, vol. 40, p. 101819, 2021.

[7] L. Tiseni, D. Chiaradia, M. Gabardi, M. Solazzi, D. Leonardis, and A. Frisoli, "UV-C mobile robots with optimized path planning: Algorithm design and on-field measurements to improve surface disinfection against SARS-CoV-2," IEEE Robotics & Automation Magazine, vol. 28, no. 1, pp. 59-70, 2021.

[8] S. Dogru and L. Marques, "Path and trajectory planning for UV-C disinfection robots," IEEE Robotics and Automation Letters, vol. 8, no. 7, pp. 4099-4106, 2023.

[9] J. Conroy, C. Thierauf, P. Rule, E. Krause, H. Akitaya, A. Gonczi, D. Lee, T. Tjanaka, M. Yim, and M. Scheutz, "Robot development and path planning for indoor ultraviolet light disinfection," in 2021 IEEE International Conference on Robotics and Automation (ICRA), pp. 7795-7801, IEEE, 2021.

[10] C. Zhang, H. Qin, S. Sun, Y. Pan, K. Liu, T. Li, and X. Zhao, "JPMDP: Joint base placement and multi-configuration path planning for 3D surface disinfection with a UV-C robotic system," Robotics and Autonomous Systems, vol. 174, p. 104644, 2024.

[11] M. Seyedin, A. Hassanpour, A. Jalali, and M. Raisee, "Robust optimization of a novel ultraviolet (UV) photoreactor for water disinfection: A neural network approach," Chemosphere, vol. 362, p. 142788, 2024.

[12] A. M. Ahmed, M. H. Ahmed, S. K. Saha, O. Ahmed, and A. Sutradhar, "Optimization algorithms as training approach with hybrid deep learning methods to develop an ultraviolet index forecasting model," Stochastic Environmental Research and Risk Assessment, vol. 36, no. 10, pp. 3011-3039, 2022.