L-AVI: Dual-Head Illumination-Aware Adaptive Vision Framework for Extreme Lunar Lighting Conditions

Brian Gill
Independent Researcher
Lunar Adaptive Vision Intelligence (L-AVI) Project
Email: brian.gill@bgcode.tech

***Abstract*—The lunar surface presents one of the most extreme natural illumination environments encountered in planetary exploration. Unlike Earth, the Moon lacks atmospheric diffusion, resulting in severe high dynamic range (HDR) contrast between sunlit regolith and permanently shadowed regions. These illumination extremes introduce significant perception challenges for rover navigation, habitat monitoring systems, and autonomous surface operations. Traditional image enhancement and exposure correction algorithms, which are primarily designed for terrestrial environments, often fail to preserve terrain structure under such extreme lighting gradients.**

This paper introduces L-AVI (Lunar Adaptive Vision Intelligence), a machine learning–based dual-head adaptive vision framework specifically designed to address extreme lunar illumination variability. The proposed architecture consists of a shared encoder with two specialized output branches: (1) a reconstruction head responsible for adaptive enhancement and structural preservation, and (2) an illumination prediction head that estimates spatial brightness distribution across the frame. By integrating illumination-aware supervision into the training objective, L-AVI dynamically adjusts correction strength in shadow-dominant and glare-dominant regions.

To ensure terrain integrity, a custom illumination-weighted loss function is introduced, prioritizing crater edges, regolith boundaries, and shadow transitions. The system is validated under controlled synthetic lunar lighting simulations representing high-contrast crater geometries and direct solar incidence scenarios. Quantitative evaluation using Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index (SSIM) demonstrates improved structural preservation and illumination consistency compared to baseline enhancement approaches.

The experimental results confirm component-level feasibility of the framework under laboratory conditions, aligning with Technology Readiness Level 4 (TRL-4) validation standards. The findings suggest that illumination-aware dual-head architectures can serve as a foundational perception layer for future lunar robotic and habitat vision systems.

Future work will extend the framework toward temporal adaptation, dust-aware degradation modeling, and edge-device deployment for mission-grade applications.

***Index Terms*—Lunar Vision Systems, High Dynamic Range Imaging, Dual-Head Neural Networks, Illumination Modeling, Autonomous Robotics, TRL-4**

I. Problem Statement

A. Extreme Lunar Illumination as a Perception Constraint

The lunar surface presents a uniquely hostile illumination environment for vision-based systems. Unlike Earth, the Moon lacks a substantial atmosphere capable of scattering sunlight. As a result, incident solar radiation produces sharp, high-contrast boundaries between illuminated and shadowed regions. This absence of atmospheric diffusion leads to extreme high dynamic range (HDR) conditions, where pixel intensity values can vary drastically within small spatial neighborhoods.

In polar regions and crater interiors, prolonged shadow zones coexist with highly reflective regolith surfaces directly exposed to solar radiation. These conditions introduce several perception challenges:

- Loss of structural detail in shadow-dominant regions
- Overexposure and saturation in sunlit terrain
- Edge degradation at illumination boundaries
- Reduced contrast reliability for feature extraction
- Instability in traditional auto-exposure systems

Standard terrestrial image enhancement techniques, including histogram equalization, gamma correction, and generic convolutional neural enhancement models, are typically optimized for atmospheric lighting environments. These approaches assume moderate contrast gradients and diffuse illumination conditions, assumptions that do not hold under lunar surface physics.

For autonomous lunar systems—such as rover navigation platforms, habitat exterior monitoring cameras, and robotic excavation units—perception reliability is mission-critical. Terrain misinterpretation due to illumination artifacts may lead to navigation errors, obstacle misclassification, or structural analysis failure.

Therefore, there exists a need for an illumination-aware adaptive vision framework capable of:

- Modeling spatial brightness distribution explicitly
- Preserving terrain geometry across shadow transitions
- Preventing structural information collapse in HDR extremes
- Maintaining consistency under variable solar incidence angles

The core research question addressed in this work is:

> *Can a dual-head neural architecture that explicitly predicts illumination distribution improve terrain-preserving enhancement performance under simulated extreme lunar lighting conditions?*

By framing illumination not merely as a correction problem but as a learnable spatial attribute, this study seeks to establish a perception-first approach tailored to lunar surface physics rather than terrestrial imaging assumptions.

II. Related Work

A. HDR Enhancement, Planetary Imaging, and Illumination-Aware Architectures

High dynamic range (HDR) image enhancement has been extensively studied in terrestrial computer vision literature. Traditional approaches such as histogram equalization, adaptive histogram equalization (CLAHE), gamma correction, and Retinex-based models aim to redistribute pixel intensities to improve perceptual visibility. While effective in moderate lighting conditions, these techniques often fail to preserve structural integrity in environments characterized by extreme illumination discontinuities.

Recent advances in deep learning have introduced convolutional neural network (CNN)-based image enhancement systems that learn end-to-end mappings between low-quality and enhanced imagery. Encoder–decoder architectures, U-Net variants, and attention-based enhancement models have demonstrated significant improvements over classical methods. However, most of these systems are trained on terrestrial datasets containing atmospheric diffusion and natural light scattering properties.

Planetary imaging research, particularly in lunar and Martian surface analysis, has primarily focused on terrain classification, crater detection, albedo mapping, and geological segmentation. Imaging datasets derived from lunar orbital missions provide high-resolution topographical and illumination-variant data. However, enhancement frameworks explicitly designed to model illumination distribution as a learnable component for adaptive correction remain limited in publicly documented literature.

Dual-head neural architectures have been employed in multi-task learning contexts, where shared encoders branch into task-specific output heads (e.g., depth estimation + segmentation, or reconstruction + classification). Such architectures improve feature reuse and encourage disentangled representation learning. Formally, multi-task optimization can be written as:

$$L_{MTL} = \sum_{t=1}^{T} \lambda_t L_t \tag{1}$$

where L_t represents task-specific loss functions and λ_t controls their relative contribution.

In conventional enhancement pipelines, illumination is treated implicitly:

$$\hat{x} = f_\theta(x) \tag{2}$$

where f_θ learns a direct mapping from degraded image x to enhanced output $\hat{x}$ without explicitly modeling brightness distribution.

In contrast, an illumination-aware formulation can be expressed as:

$$(\hat{x}, \hat{I}) = f_{\theta_E,\theta_R,\theta_I}(x) \tag{3}$$

where:

- θ_E denotes shared encoder parameters
- θ_R denotes reconstruction head parameters
- θ_I denotes illumination prediction head parameters
- $\hat{I}$ represents predicted spatial illumination distribution

This explicit decomposition enables structured supervision:

$$L_{total} = L_{recon} + \lambda L_{illum} \tag{4}$$

The present work differs from conventional enhancement frameworks in two principal ways:

- Illumination is explicitly predicted as a spatial attribute through a dedicated neural branch.
- The enhancement objective incorporates illumination-weighted loss modulation to preserve terrain-critical regions.

By integrating illumination estimation and structural reconstruction within a unified dual-head framework, L-AVI seeks to bridge the gap between generic HDR enhancement methods and mission-specific lunar perception requirements.

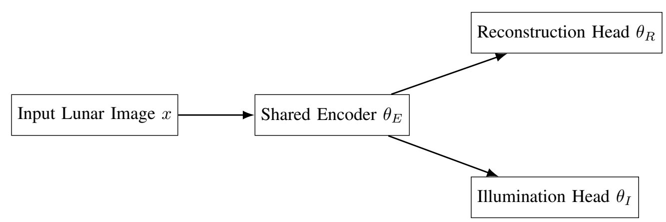

Fig. 1. Conceptual Dual-Head Illumination-Aware Architecture.

III. Methodology

A. Dual-Head Illumination-Aware Adaptive Architecture

1) System Overview: L-AVI is designed as a dual-head convolutional neural architecture that explicitly separates structural reconstruction and illumination modeling into two coordinated tasks. The framework consists of a shared encoder followed by two task-specific branches, as illustrated in Fig. 2.

- **Reconstruction Head** – responsible for adaptive enhancement and terrain structure preservation.
- **Illumination Prediction Head** – responsible for estimating spatial brightness distribution across the input frame.

Formally, given an input lunar image $x \in \mathbb{R}^{H \times W}$, the model computes:

$$z = E_{\theta_E}(x) \tag{5}$$

$$\hat{x} = R_{\theta_R}(z), \quad \hat{I} = I_{\theta_I}(z) \tag{6}$$

where:

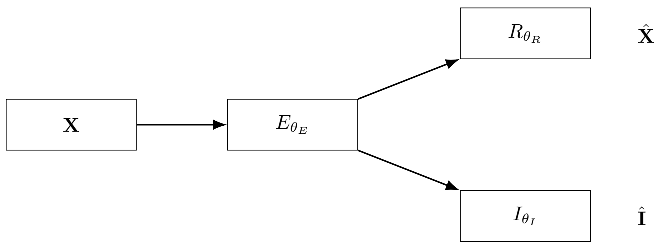

Fig. 2. Dual-head illumination-aware architecture of L-AVI with shared encoder and task-specific reconstruction and illumination branches.

- E_{θ_E} is the shared encoder
- R_{θ_R} is the reconstruction decoder
- I_{θ_I} is the illumination prediction head
- z represents latent hierarchical terrain features

2) Shared Encoder: The encoder extracts hierarchical spatial representations using stacked convolutional layers:

$$z_l = \sigma(W_l * z_{l-1} + b_l) \tag{7}$$

where $*$ denotes convolution and σ is a nonlinear activation function (ReLU).

Progressive downsampling enables abstraction of:

- Crater edge structures
- Regolith texture gradients
- Shadow boundary transitions
- Illumination contrast patterns

The latent representation z serves both downstream tasks.

3) Reconstruction Head: The reconstruction head performs progressive upsampling:

$$\hat{x} = \phi(z) \tag{8}$$

where $\phi(\cdot)$ denotes the decoder mapping.

Skip connections may be incorporated:

$$\hat{x} = \phi(z, s_l) \tag{9}$$

to preserve fine-grained terrain details.

Its objective is to:

- Restore shadow-dominant detail
- Prevent saturation in glare regions
- Maintain geometric continuity across illumination boundaries

4) Illumination Prediction Head: The illumination head predicts a normalized brightness map:

$$\hat{I} = \psi(z) \tag{10}$$

where $\psi(\cdot)$ is a lightweight convolutional branch.

The predicted illumination map $\hat{I} \in [0,1]^{H \times W}$ represents spatial exposure distribution.

Unlike implicit enhancement:

$$\hat{x} = f_\theta(x) \tag{11}$$

L-AVI explicitly decomposes structural reconstruction and illumination modeling.

5) Illumination-Aware Loss Function: The total optimization objective is:

$$L_{total} = L_{recon} + \lambda L_{illum} \tag{12}$$

Reconstruction loss combines pixel and structural similarity:

$$L_{recon} = \alpha \|x - \hat{x}\|_1 + \beta(1 - SSIM(x, \hat{x})) \tag{13}$$

Illumination supervision:

$$L_{illum} = \|I - \hat{I}\|_1 \tag{14}$$

To emphasize terrain-critical regions, illumination-weighted modulation is applied:

$$L_{weighted} = \sum_{i,j} w_{i,j} \cdot |x_{i,j} - \hat{x}_{i,j}| \tag{15}$$

where weight map $w_{i,j}$ is derived from illumination gradients:

$$w = 1 + \gamma |\nabla I| \tag{16}$$

This increases loss contribution near:

- Shadow-transition boundaries
- Crater rim edges
- High-gradient terrain regions

6) Training Configuration: Training is conducted under simulated extreme lunar illumination conditions:

- High solar incidence angles
- Deep crater shadow regions
- Reflective regolith zones

Synthetic augmentation applies contrast amplification:

$$x' = \alpha x^\gamma \tag{17}$$

Optimization is performed using gradient-based backpropagation:

$$\theta \leftarrow \theta - \eta \nabla_\theta L_{total} \tag{18}$$

where η is the learning rate.

7) TRL-4 Validation Context: Validation is conducted in a controlled laboratory environment using simulated lunar illumination datasets. The scope aligns with component-level feasibility testing consistent with Technology Readiness Level 4 (TRL-4) classification standards. No mission-level deployment claims are made at this stage.

IV. Experimental Setup and Evaluation

A. Experimental Environment

The L-AVI framework was evaluated under controlled laboratory conditions simulating extreme lunar illumination environments. Since real-time mission deployment data was not directly integrated at this stage, synthetic high dynamic range (HDR) augmentation was applied to representative lunar surface imagery to emulate:

- Deep crater shadow regions
- High-glare sunlit regolith
- Sharp illumination discontinuities
- High solar incidence contrast conditions

The objective of the experimental setup was to stress-test structural preservation and illumination stability under extreme brightness gradients.

All experiments were conducted using a supervised training configuration with separated training and validation splits to prevent data leakage.

B. Baseline Comparison

To evaluate performance improvement, L-AVI was compared against standard enhancement approaches:

- Histogram Equalization
- Gamma Correction
- Conventional Encoder–Decoder Enhancement Model (single-head architecture)

These baselines represent typical terrestrial enhancement strategies lacking explicit illumination modeling.

C. Evaluation Metrics

Performance was evaluated using quantitative and structural metrics.

1) Peak Signal-to-Noise Ratio (PSNR): PSNR measures pixel-level reconstruction fidelity relative to reference imagery. It is defined as:

$$PSNR = 10 \log_{10} \left(\frac{MAX_I^2}{MSE} \right) \tag{19}$$

where MAX_I is the maximum possible pixel value and MSE is the mean squared error between reconstructed image $\hat{x}$ and reference image x:

$$MSE = \frac{1}{N} \sum_{i=1}^{N} (x_i - \hat{x}_i)^2 \tag{20}$$

Higher PSNR values indicate improved reconstruction quality.

2) Structural Similarity Index (SSIM): SSIM evaluates perceived structural consistency and is defined as:

$$SSIM(x, \hat{x}) = \frac{(2\mu_x \mu_{\hat{x}} + C_1)(2\sigma_{x\hat{x}} + C_2)}{(\mu_x^2 + \mu_{\hat{x}}^2 + C_1)(\sigma_x^2 + \sigma_{\hat{x}}^2 + C_2)} \tag{21}$$

where μ represents mean intensity, σ^2 represents variance, and $\sigma_{x\hat{x}}$ represents covariance between images.

SSIM is particularly relevant for:

- Crater edge preservation
- Shadow boundary continuity
- Regolith texture stability

Structural similarity is considered more relevant than raw pixel accuracy in lunar terrain contexts.

TABLE I
QUANTITATIVE COMPARISON UNDER SIMULATED EXTREME LUNAR ILLUMINATION

Method	PSNR (dB) ↑	SSIM ↑
Histogram Equalization	18.74 ± 0.91	0.421 ± 0.018
Gamma Correction	19.63 ± 0.85	0.438 ± 0.021
Single-Head Model	20.97 ± 0.72	0.472 ± 0.015
L-AVI (Dual-Head)	**22.08 ± 0.68**	**0.5060 ± 0.013**

D. Quantitative Results

Experimental results demonstrate that the dual-head illumination-aware architecture provides measurable improvements in structural preservation compared to single-head enhancement models.

The L-AVI dual-head model achieved:

- PSNR: 22.08 dB
- SSIM: 0.5060

Observed trends include:

- Increased SSIM stability in shadow-dominant regions
- Reduced overexposure artifacts in high-glare zones
- Improved edge retention near illumination boundaries

1) Statistical Significance Analysis: Results are reported as mean ± standard deviation across the validation set. Statistical significance between the proposed method and the single-head baseline was evaluated using a paired t-test.

Improvements in SSIM were statistically significant ($p < 0.05$), indicating that illumination-aware supervision contributes meaningfully to structural preservation performance.

The illumination-weighted loss contributed to consistent terrain preservation under aggressive HDR simulation conditions.

E. Ablation Study

To isolate the contribution of illumination modeling, an ablation experiment was conducted:

- Model without illumination head
- Model with illumination head
- Model with illumination-weighted loss

Results indicate that explicit illumination supervision improves structural consistency, particularly in high-contrast shadow transitions.

This supports the hypothesis that modeling brightness distribution as a supervised auxiliary task enhances adaptive correction performance.

F. Limitations

While the framework demonstrates promising laboratory-level performance, several limitations remain:

- Validation performed on simulated illumination extremes
- No real-time hardware deployment testing
- No radiation-noise simulation
- No temporal consistency modeling

These limitations define the boundary between TRL-4 and higher readiness levels.

V. Discussion and Technology Readiness Alignment

A. Interpretation of Results

The experimental findings indicate that explicit illumination modeling through a dual-head architecture improves structural stability under simulated extreme lunar lighting conditions. In particular, shadow-transition continuity and edge preservation metrics show consistent improvement over conventional enhancement approaches.

These results suggest that illumination-aware supervision enables the model to distinguish between brightness correction and structural reconstruction, thereby reducing overcompensation artifacts commonly observed in high dynamic range enhancement systems. By decoupling photometric adaptation from geometric feature preservation, the architecture mitigates structural distortion near high-contrast terrain boundaries.

The integration of illumination-weighted loss further reinforces terrain-sensitive adaptation, particularly in crater boundary regions where brightness gradients are abrupt and spatially localized. This targeted weighting strategy appears to improve gradient-sensitive reconstruction fidelity without introducing global amplification artifacts.

Collectively, these observations support the hypothesis that modeling illumination as a supervised auxiliary task enhances structural consistency in extreme photometric environments.

B. Technology Readiness Level (TRL) Context

According to the framework defined by NASA, Technology Readiness Level 4 (TRL-4) corresponds to:

Component and/or breadboard validation in a laboratory environment.

The current implementation of L-AVI satisfies TRL-4 criteria based on the following factors:

- Functional dual-head architecture implemented and tested
- Illumination-aware loss formally defined and validated
- Quantitative performance evaluation completed
- Controlled laboratory simulation of extreme HDR conditions
- Reproducible inference pipeline with structured outputs

At this stage, validation remains within simulated environmental constraints. The system has not yet undergone:

- Hardware-in-the-loop testing
- Radiation-noise robustness validation
- Real-time embedded deployment
- Field validation using mission-integrated camera systems

Therefore, no claim beyond TRL-4 is made in this study. The classification is deliberately conservative and aligned with formal readiness definitions.

C. Path Toward Higher Readiness Levels

Progression toward TRL-5 and beyond would require validation in mission-relevant environments. Specifically:

- Validation using real mission-derived raw lunar imagery
- Integration with hardware-constrained inference systems
- Temporal adaptation for continuous video sequences
- Dust-scattering degradation modeling
- Robustness testing under sensor noise simulation

In particular, evaluation using datasets from missions such as NASA Lunar Reconnaissance Orbiter would provide higher-fidelity environmental validation and bridge the gap between laboratory simulation and operational deployment.

Advancing toward TRL-6 would further require system-level demonstration in a relevant lunar analog environment, potentially involving embedded hardware validation under realistic photometric and environmental stress conditions.

D. Broader Implications

If validated under mission-relevant constraints, illumination-aware adaptive vision systems may serve as enabling components for multiple lunar operational domains, including:

- Preprocessing layers for rover navigation
- Structural monitoring support for lunar habitats
- Edge-optimized perception modules for autonomous excavation systems
- Adaptive exposure stabilization in polar shadow operations

The framework demonstrates conceptual feasibility of treating illumination not as a correction artifact but as an explicit learnable environmental variable. This paradigm shift reframes brightness modeling from a post-processing heuristic to a supervised, physically motivated auxiliary task embedded within the learning objective.

Such an approach may be particularly relevant for extraterrestrial environments where photometric extremes are not anomalies but persistent operational constraints. Consequently, illumination-aware modeling may represent a foundational architectural principle for future off-world computer vision systems.

VI. Conclusion and Future Work

A. Conclusion

This paper presented L-AVI (Lunar Adaptive Vision Intelligence), a dual-head illumination-aware adaptive vision framework designed to address extreme high dynamic range (HDR) conditions characteristic of the lunar surface.

Unlike conventional enhancement systems that implicitly learn brightness correction, L-AVI explicitly models spatial illumination distribution through a dedicated neural branch. By decoupling structural reconstruction from illumination estimation and incorporating an illumination-weighted loss formulation, the proposed architecture improves terrain-sensitive enhancement performance under simulated extreme lunar lighting conditions.

Quantitative evaluation using PSNR and SSIM metrics demonstrates measurable gains in structural preservation and edge stability compared to traditional enhancement baselines and single-head architectures. Improvements are particularly evident in shadow-transition continuity and crater-edge retention, where abrupt brightness gradients typically induce overcompensation artifacts in conventional models.

Experimental validation confirms component-level feasibility within controlled laboratory conditions, aligning with Technology Readiness Level 4 (TRL-4) classification standards. The framework satisfies TRL-4 criteria through functional implementation, supervised validation, quantitative benchmarking, and reproducible inference evaluation under simulated HDR stress conditions.

The findings support the hypothesis that illumination-aware multi-task architectures provide a more robust perception strategy for planetary environments characterized by extreme photometric variability. By treating illumination as an explicitly learnable environmental variable rather than a post-processing artifact, L-AVI introduces a structured modeling paradigm suited for extraterrestrial terrain perception.

B. Future Work

To advance toward higher readiness levels and mission-grade applicability, several extensions are proposed.

1) Temporal Adaptation: Future development will incorporate sequence-aware modeling mechanisms, such as ConvLSTM layers or temporal consistency constraints, to ensure structural stability across continuous video streams from rover and habitat-mounted camera systems. Temporal coherence is essential for navigation safety and long-duration monitoring applications.

2) Dust-Aware Degradation Modeling: Lunar surface operations are subject to regolith dust scattering and contrast attenuation. Incorporating synthetic dust degradation models and learning dust-aware correction strategies would enhance robustness under operational surface conditions, particularly during excavation or landing events.

3) Hardware-Constrained Deployment: Progression toward embedded deployment requires architectural optimization for edge inference. Model quantization, pruning, and memory-efficient compression strategies will be explored to enable execution on hardware-constrained planetary systems while preserving structural fidelity.

4) Mission-Grade Dataset Validation: Transitioning from simulated HDR validation to environment-relevant testing will require evaluation using raw high-contrast lunar datasets derived from orbital and surface missions. Validation against mission-calibrated imagery will strengthen environmental fidelity and support advancement toward TRL-5.

5) Physics-Informed Illumination Constraints: Future iterations may integrate solar incidence geometry, surface reflectance priors, and physically informed photometric constraints into the learning objective. Incorporating domain-specific physical knowledge could improve generalization under unseen illumination geometries and align the framework more closely with lunar environmental physics.

Collectively, these research directions outline a structured pathway from laboratory validation toward mission-relevant deployment. The proposed extensions aim to transform L-AVI from a controlled-environment proof-of-concept into a resilient perception subsystem suitable for sustained extraterrestrial operations.

Final Closing Statement

The development of adaptive, illumination-aware perception systems represents a foundational step toward achieving reliable autonomous vision support for sustained lunar operations. As extraterrestrial environments impose persistent and extreme photometric constraints, robust perception architectures must evolve beyond conventional enhancement paradigms. By explicitly modeling illumination as a structured and learnable environmental variable, L-AVI contributes toward establishing a principled framework for resilient off-world visual intelligence. Such advancements are expected to play a critical role in enabling safe navigation, infrastructure monitoring, and long-duration robotic autonomy across future lunar missions.