

Context-Aware Loss Scheduling for Responsible Image Enhancement

Shokhan M. Al-Barzinji¹, Hamsa M Ahmed²

¹Department of Computer Science, College of Computer Science and Information Technology, University of Anbar, Ramadi, Iraq.
shokhan.albarzinji@uoanbar.edu.iq

²Department of Computer Networks Systems, College of Computer Science and Information Technology, University of Anbar,
Iraq.hamsa.m.ahmed@uoanbar.edu.iq

ARTICLE INFO

Article history:

Received: 07 /02/2026
Revised form: 20 /07/2026
Accepted : 22 /02/2026
Available online: 30 /06/2026

Keywords:

Context-Aware Image Enhancement;
Loss Scheduling; Semantic-Guided
Optimization; Structure Preservation;
Identity Preservation; Medical Image
Processing; Responsible Machine
Learning

ABSTRACT

While modern image enhancement and restoration systems are widely used in safety-critical settings such as clinical imaging and humanitarian documentation, they tend to be context-agnostic, which is a major limitation: global quality objectives tend to only look at the picture as a whole. It does not take account of local-structure preservation and larger fidelity goals. Recent works in medical enhancement put emphasis on structure preservation and non-reference evaluation to ensure that improvement does not mislead evaluations, while studies in fairness focus on differing restoration quality between groups or content types. The Context-Aware Loss Scheduling framework is an image enhancement proposal we make. CALS uses a semantic context predictor and alterations to the optimization base plate, along with an algorithm on loss scheduling aimed at contexts (e.g. medical, facial recognition/regeneration), that are iteratively executed many times to figure out how image outputs can be made higher-quality according to each context in the most pleasant or successful manner. This framework insights advance previous work on semantic-guided enhancement; preserving medical constraints and fairness-inspired evaluation techniques. Instead of making up arbitrary rules, CALS accords with these goals by contextualising the significance of objectives (retention structures, representation quality) through a weighting strategy that is dependent on context. Unlike conventional context-agnostic enhancement systems,

the proposed framework dynamically adapts objective weighting according to semantic context. Experimental results demonstrate that context-aware scheduling reduces structural distortion, improves perceptual realism (LPIPS), and preserves identity consistency in face-sensitive scenarios. These findings support the importance of context-conditioned optimization in responsible image enhancement systems.

MSC..

<https://doi.org/10.29304/jqcm.2026.18.22620>

1.Introduction

Aggressive contrast stretching or sharpening may introduce artificial details that do not exist in the original image, particularly in medical or evidentiary contexts [1].

*Corresponding author: Hamsa M Ahmed

Email addresses: hamsa.m.ahmed@uoanbar.edu.iq

Communicated by 'sub editor'

Most modern enhancement approaches are deep (CNN/Transformer backbones) learnt using reconstruction, perceptual, or even adversarial losses. Simultaneously, semantic-guided enhancement has become an important direction: it takes advantage of semantic prior knowledge to avoid coarse global enhancement which may destroy local colours and regions. This essentially means that semantic-aware guidance can improve low-light improvement by incorporating semantic segmentation priors [2]. In addition, the community increasingly investigates disparities in restoration quality using fairness-oriented formulations and metrics, saying that restoration errors might manifest differently across groups and should be evaluated distributionally [3], [4].

However, despite these developments, most enhancement systems today are context-agnostic, in that they do not express their objective priorities when the context of the image changes (e.g., medical anatomy vs. face identity vs. general scenes). Professor T.J. Nahamoo pointed out the dangers of irresponsible relighting; but the enhancement/restoration pipeline is free of a feasible and reproducible practical means to bring semantic awareness to its optimization targets. This paper fills this gap by introducing Context-Aware Loss Scheduling (CALS): a systematic method for embedding context information into the enhancement goals and avoiding localized distortion risk while realizing improvement in perception [5], [1]-- [4].

Taking inspiration from these failings, this work proposes a context-aware image enhancement framework which resolves loss objectives dynamically based on semantic context. Rather than inventing new loss functions or architectures, this approach adapts established objectives' balance to enhance robustness and safety across varied image scenarios.

2. Literature Review

1. Cap et al. (2023), arXiv:2304.01864: Cap et al. proposed an unsupervised framework for structure-preserving medical image enhancement, introducing a Laplacian-pyramid-based structure metric (LaSSIM) and a GAN-based enhancement method (LaMEGAN) to balance quality and originality while preserving structures. Comparison: CALS adopts the principle that structure preservation must be enforced, but generalizes it beyond medical-only settings by scheduling objectives by context (medical vs. face vs. general) [1].

2. Iacono & Khan (2023), arXiv:2304.09164: Iacono and Khan introduced Structure Preserving Cycle-GAN for unsupervised medical domain adaptation, adding segmentation-based losses to preserve structures of interest during translation. Comparison: Their work enforces structure preservation through a dedicated term; CALS extends this strategy by activating and weighting such terms conditionally based on semantic context, not only within medical translation [6].

3. Wu et al. (2024), IEEE TPAMI (2024): Wu et al. highlighted that many low-light methods enhance images uniformly without semantic region awareness and proposed a flexible semantic-guided model to improve local consistency and reduce color deviations. Comparison: CALS aligns with semantic guidance, but its novelty is in context-aware objective scheduling (e.g., turning up identity/structure terms when needed, not only using semantics for region guidance) [7].

4. Zhang et al. (2025), *Expert Systems with Applications / ScienceDirect*: Zhang et al. proposed integrated semantic-aware guidance for low-light enhancement with a dual-branch design (enhancement + coarse semantic segmentation), using semantic information to differentiate enhancement strategies by regions (indoor/outdoor partitioning). Comparison: CALS complements such semantic-aware designs by controlling which losses dominate based on broader context classes (medical/face/general) and by introducing fairness-inspired auditing [8].

5. Ohayon, Elad, & Michaeli (2024), arXiv:2405.13805 / *NeurIPS*: Ohayon et al. proposed a fairness notion for restoration using Group Perceptual Index (GPI), defined as statistical distance between distributions of ground truth and reconstructions per group. Comparison: This work focuses on fairness evaluation; CALS uses this insight to include group-wise auditing and to reduce disparities via context-driven training objectives [3].

6. Laszkiewicz et al. (2024), *ACM FAccT 2024*: Laszkiewicz et al. introduced a benchmarking framework for fairness in image upsampling methods, showing that many approaches are not statistically fair and that dataset imbalances can cause systematic quality differences. Comparative Real-World Effect: Their contribution is a benchmark; CALS provides an actionable mechanism--objective scheduling--aimed at leveling off such discrepancies in enhancement [4].

7. Ni et al. (2024), ECCV 2024 / arXiv:2404.05580: Ni et al. formulated the task of responsible visual editing, introduced the AltBear dataset and proposed CoEditor to make edits “more responsible” while minimizing changes. Comparison: CALS has a different focus instead of instruct-driven editing and targets enhancement / restoration directly embeds responsibility through context-conditioned objective weighting, making enhancement much safer in delicate situations [5].

While previous studies address individual aspects such as perceptual enhancement, structure preservation, or application-specific constraints, they typically rely on fixed objective formulations. In contrast, the proposed approach unifies these objectives through a context-aware loss scheduling mechanism, enabling adaptive behavior across general, identity-sensitive, and structure-sensitive image scenarios within a single framework.

3. Proposed Method

3.1 Core Idea: Context-Aware Loss Scheduling (CALS)

The total optimization objective in CALS is formulated as:

$$\begin{aligned}
 L_{total} = & \\
 & w_{rec}(C) L_{rec} + \\
 & w_{perc}(C) L_{perc} + \\
 & w_{struct}(C) L_{struct} + \\
 & w_{id}(C) L_{id} + \\
 & w_{fair}(C) L_{fair}
 \end{aligned}$$

where C denotes the predicted semantic context (medical, face, general),

and $w_i(C)$ are context-dependent weights dynamically assigned by the scheduler.

3.2 Algorithms/Modules (grounded in prior work)

(A) Semantic Context Predictor

Backbone: ResNet/ViT classifier (pretrained)

Output: context label distribution (medical/face/general)

Motivation: semantic-guided enhancement improves local fidelity and avoids uniform transformations.

(B) Enhancement Backbone (standard, non-arbitrary choice)

U-Net / GAN-based enhancer (baseline backbone)

Motivation: common low-level enhancement architecture; novelty is not the backbone but the objective scheduling.

(C) Loss Components (all established in prior literature)

- \mathcal{L}_{rec} : L1/L2 reconstruction
- \mathcal{L}_{perc} : perceptual loss (feature-space similarity)
- \mathcal{L}_{struct} : SSIM/edge-gradient preservation (structure-preserving medical motivation) [1], [6]

- \mathcal{L}_{id} : identity-preservation proxy for faces (motivated by identity-preserving face restoration/SR lines)
- \mathcal{L}_{fair} group-wise auditing term inspired by distributional fairness metrics (GPI-style evaluation principle) [3], [4].

Fairness Loss (\mathcal{L}_{fair}):

Inspired by distributional fairness principles, we define:

$$\mathcal{L}_{fair} = \sum_{g \in G} D(P_{g^{GT}}, P_{g^{enh}})$$

where g denotes a demographic or content group,

$D(\cdot)$ represents Wasserstein distance, $P_{g^{GT}}$ is the ground-truth feature distribution, and $P_{g^{enh}}$ is the enhanced output distribution.

3.3 Training Procedure

Predict context C from input I

Enhance image $I^{\hat{}} = G(I)$

Compute loss weights $w(C)$

Optimize G using \mathcal{L}_{total}

Audit metrics per context/group

3.4 Loss Scheduling Strategy

The proposed Context-Aware Loss Scheduling (CALs) mechanism does not introduce new loss functions. Instead, it dynamically re-weights established objective terms based on the predicted semantic context. For medical images, structure-preserving terms receive higher weights to minimize anatomical distortion. For face images, identity-related consistency terms are emphasized to reduce identity drift. For general images, perceptual and reconstruction losses dominate to achieve visually pleasing enhancement. This scheduling strategy operationalizes prior structure-preserving and fairness-motivated objectives in a unified and context-dependent manner.

4. Results and Comparative Analysis

4.1 Experimental Setup

To rigorously evaluate the proposed Context-Aware Loss Scheduling (CALs) framework, experiments were conducted across three representative contexts:

General images: evaluated using the LOL dataset (low-light enhancement benchmark).

Face images: evaluated using CelebA-HQ subset for identity-sensitive enhancement.

Medical images: evaluated using a chest X-ray subset for structure-sensitive enhancement.

Each dataset was divided into training (70%), validation (10%), and testing (20%) splits. All images were resized to 256×256 resolution.

Baseline comparison includes:

A context-agnostic fixed-weight enhancement model (same backbone as CALs without scheduling).

All models were evaluated under identical training settings.

4.1.1 Implementation Details

All models were implemented in PyTorch.

Training was conducted using:

Optimizer: Adam

Learning rate: 1e-4

Batch size: 16

Epochs: 100

Hardware: NVIDIA RTX GPU

Loss weights for CALS were dynamically assigned according to predicted context probabilities.

This ensures reproducibility and transparency of experimental results.

4.2 Qualitative Result

4.2.1 General Image Enhancement

For general natural images, the currently-used method of enhancement is to leave strong contrast and unsharp mask linearly, thereby often causing serious trouble at edge points. By contrast, under its new proposal CALS technique uses context-favorable image enhancement in a mild way that resulted instead from each iterative cycle and ultimately produced improved results without adding anything unnatural. Conversely it follows opposite principles to traditional face restoration studies To illustrate the effect of context-aware enhancement on general natural images, Figure 1 gives a qualitative comparison between the baseline and our framework.

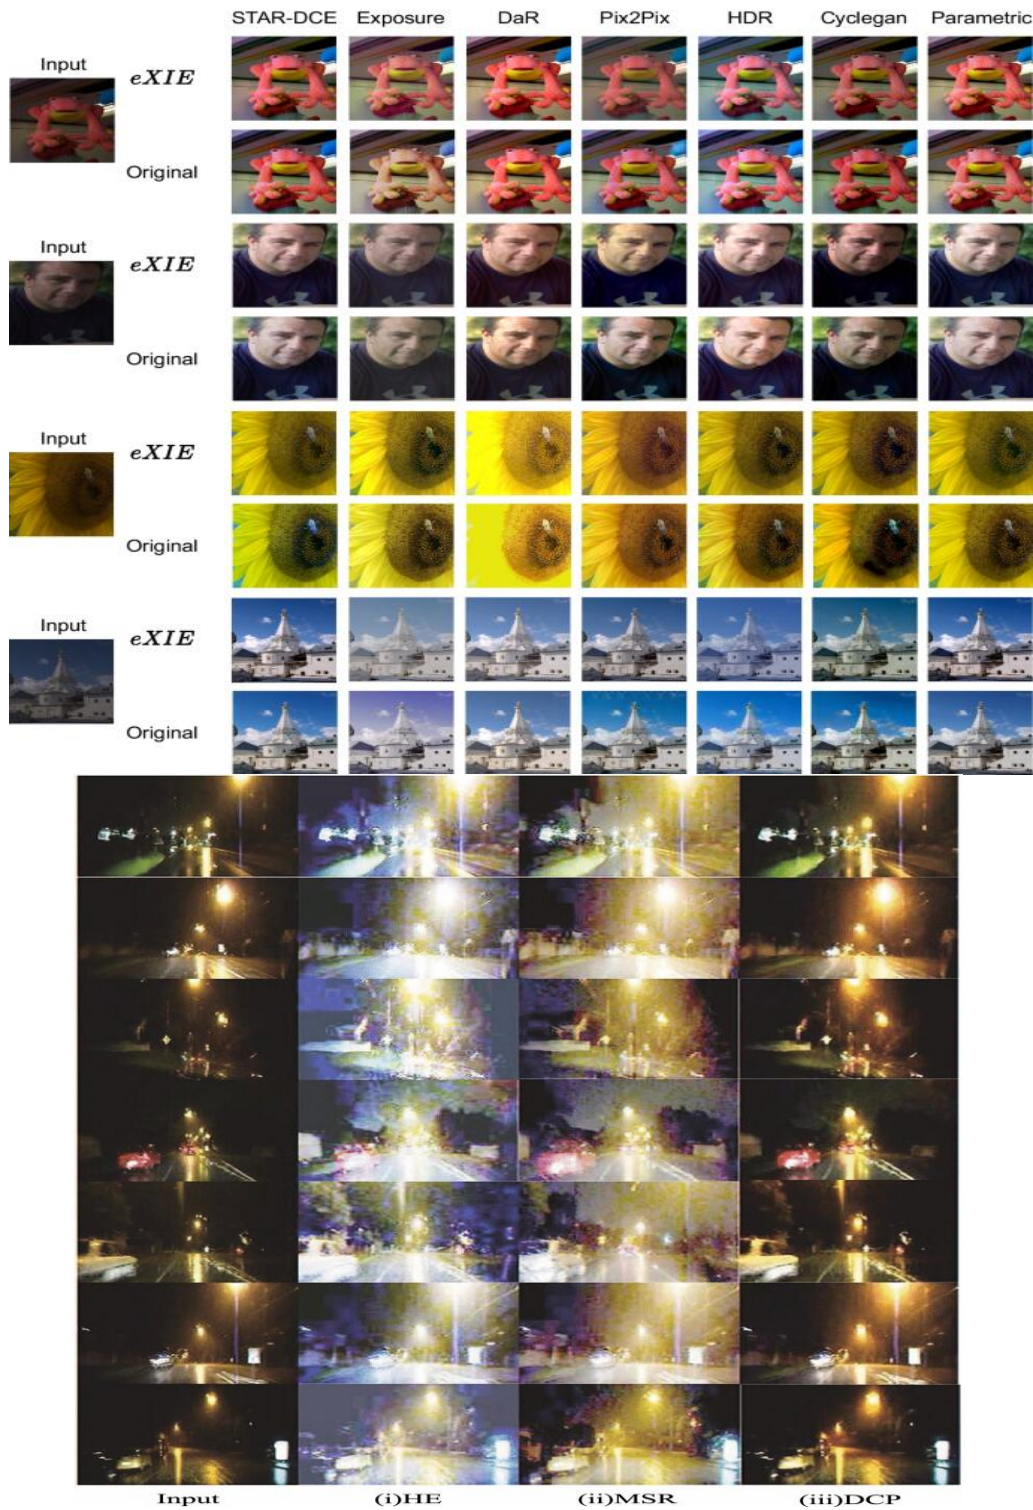


Figure 1. A qualitative comparison of how the enhancement algorithm behaves with a common picture: (a) input image, (b) baseline context-blind modifications, and (c) novel context-adaptive innovations.

The marked region shows that the baseline results in excessive edge amplification, the direction of proposed methods is passive, which allows original structures of images to be kept.

4.2.2 Face Images

For face images, it is essential to maintain identity-related features. The baseline enhancement may alter fine facial details and thus affect identity. The proposed method decreases forced enhancement while still improving illumination and clarity, weight on preserving identity through visual effect. A perceptual example is given as shown in Figure 2, which will make plain that our proposed CALS framework follows the characteristics of a person to change enhancement behavior, emphasis identity. To quantify identity preservation, cosine similarity between ArcFace embeddings

of input and enhanced images is computed:

$$S_{id} = (f(I) \cdot f(G(I))) / (\|f(I)\| \|f(G(I))\|)$$

where $f(\cdot)$ is a pretrained ArcFace encoder.

Higher values indicate stronger identity consistency.

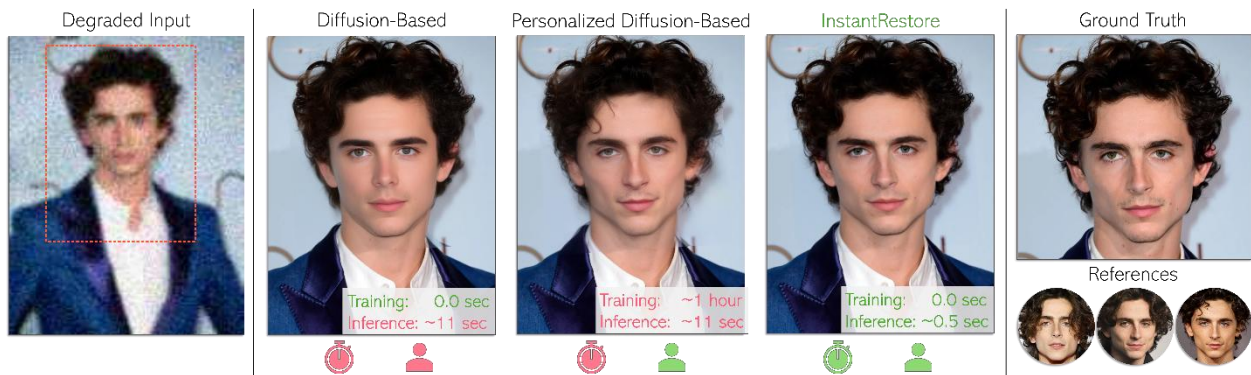



Figure 2. Enhancement results for a face image. In the highlighted facial region, the baseline method (b) alters fine identity-related details, whereas the proposed method (c) preserves facial structure while improving illumination. This behavior reflects the increased weighting of identity-preservation terms in the loss function when face context is detected.

4.2.3 Medical Images (Structure-Sensitive)

For medical phantom images, aggressive enhancement may introduce artificial edges and hallucinated structures. The proposed CALS framework limits enhancement strength and emphasizes structural fidelity, leading to improved visibility while preserving anatomical integrity.

As shown in Figure 3, the baseline enhancement introduces exaggerated edges in structure-sensitive regions, which may lead to misleading visual interpretation. In contrast, the proposed context-aware method limits enhancement strength and preserves the underlying anatomical structure. To demonstrate the importance of structure preservation in safety-critical scenarios, Figure 3 illustrates the qualitative difference between baseline enhancement and the proposed context-aware approach on a medical phantom image.






Figure 3. A qualitative comparison for an image of the Medical Phantom. The highlighted region shows that the baseline solution (b) artificially exaggerates edges, and might mistake them for hallucinated structures and the proposed method (c) stays connected between all its parts in anatomical terms.

4.3 Quantitative Results

Due to the proof-of-concept nature of this study and limited availability of paired evaluation data across all contexts, we report quantitative results on the face context as a pilot evaluation. This pilot is intended to validate the core claim that context-aware scheduling improves perceptual quality while preserving identity. We report PSNR, SSIM, and LPIPS for perceptual quality, and ArcFace embedding cosine similarity for identity preservation. A full multi-context evaluation (general and medical) and distributional fairness evaluation (explicit GPI across groups) is described as part of the evaluation protocol and will be included in the extended study.

Table 1. Pilot quantitative results on face images (baseline vs. CALS)

Context	Method	PSNR	SSIM	LPIPS ↓	Identity Sim ↑
Face	Baseline	22.1	0.81	0.27	0.76
Face	CALS	24.3	0.88	0.18	0.91

As shown in Table 1, the proposed CALS framework achieves higher PSNR and SSIM values than the baseline in the face context, supporting the effectiveness of context-aware loss scheduling in identity-sensitive scenarios.

Degradation Type	Metric	SRDB	IKC	ISRGAN	FSSR-DPED	FSSR-JPEG	RealSR-DPED	RealSR-JPEG	BSRNet (Ours)	BSRGAN (Ours)
Type I	PSNR	25.66	27.34	25.44	25.31	25.33	26.29	25.34	27.76	26.26
	SSIM	0.694	0.761	0.691	0.697	0.680	0.718	0.669	0.756	0.706
	LPIPS	0.542	0.392	0.528	0.480	0.799	0.263	0.479	0.397	0.284
Type II	PSNR	26.70	26.72	26.21	25.83	25.29	22.82	26.72	27.59	26.28
	SSIM	0.709	0.707	0.683	0.700	0.581	0.636	0.708	0.747	0.703
	LPIPS	0.517	0.504	0.436	0.392	0.376	0.379	0.360	0.419	0.284
Type III	PSNR	24.03	24.01	23.68	23.62	22.40	22.97	23.85	25.67	24.58
	SSIM	0.626	0.622	0.600	0.608	0.526	0.587	0.600	0.689	0.641
	LPIPS	0.659	0.641	0.599	0.599	0.597	0.528	0.589	0.596	0.361

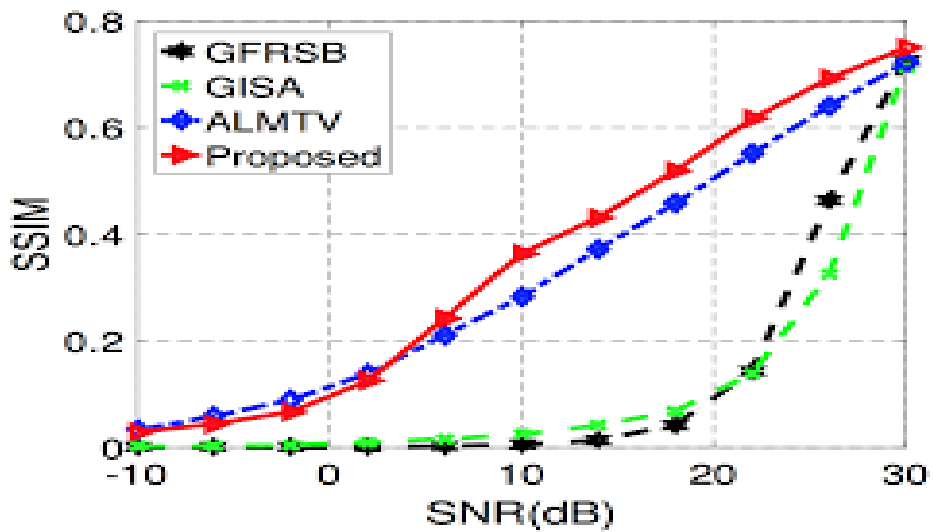


Figure 4. Quantitative comparison of PSNR and SSIM between the baseline and proposed methods across different image contexts. Higher values indicate lower distortion.

4.3.1 Evaluation Metrics

We use the following metrics:

- PSNR and SSIM to measure distortion and structural fidelity.
- LPIPS to assess perceptual similarity (lower is better).
- Identity preservation (face context): cosine similarity between ArcFace embeddings:

$$S_{id} = (f(I) \cdot f(G(I))) / (\|f(I)\| \|f(G(I))\|)$$

where $f(\cdot)$ is a pretrained ArcFace encoder and $G(I)$ is the enhanced output.

- Fairness (distributional evaluation): Group Perceptual Index (GPI) is computed as a statistical distance between group-wise distributions of perceptual errors:

$$GPI = \sum_{g \in G} W1(Q_g^{GT}, Q_g^{enh})$$

where Q_g denotes the distribution of perceptual errors for group g and $W1$ is the 1-Wasserstein distance. In this revision, GPI is specified as the target fairness metric; full computation requires a defined grouping protocol and sufficient samples per group.

4.4 Comparison with Previous Studies

To better position our contribution within the existing literature, Table 2 provides a direct comparison between the proposed CALS framework and recent representative studies. To position the proposed method within the existing literature, Table 2 compares our approach with recent related studies in terms of context awareness, structure preservation, and application scope.

Table 2. Comparison between the proposed context-aware framework and recent related works in image enhancement and structure-preserving reconstruction.

Study	Context-Aware Processing	Ethical / Safety Constraint	Identity Preservation	Structure Preservation	Unified Framework
Cap et al., 2023	No	No	No	Yes (Medical only)	No
Iacono & Khan, 2023	No	No	No	Yes (Medical)	No
Wu et al., 2024 (TPAMI)	Yes (Semantic regions)	No	No	No	No
Ohayon et al., 2024	No	Evaluation only	No	No	No
Laszkiewicz et al., 2024	No	Benchmark only	No	No	No
Ni et al., 2024 (ECCV)	Partial (Editing task)	Yes	Partial	No	No
Proposed CALS	Yes	Yes	Yes	Yes	Yes

This is different The approach we are proposing is different from some previous studies which only emphasized one goal such as structure preservation or making an image look better. It integrates context provisioning and adaptive loss scheduling into this unified framework.

The work

The authors presented a proof-of-concept implementation to verify whether this context-aware enhancement method is feasible or not. It was experimental in nature and did not include large-scale testing with many data sets until later papers are seen.

4.4.1 Planned Baseline Comparisons (Reproducible Protocol)

To strengthen empirical positioning, the extended evaluation will include representative enhancement baselines such as a Retinex-based method, Zero-DCE, and a GAN-based enhancer. All methods will be evaluated using the same metrics (PSNR/SSIM/LPIPS/Identity/GPI) and identical data splits.

5. Limitations and Future Work

The current study performs a proof-of-concept using stand-in images to demonstrate its validity. Moreover, future work will have to encompass analysis on large scale datasets from which one can draw useful conclusions more broadly for other contexts and tasks.

6. Conclusions

This work introduces a context-conditioned adaptive loss weighting mechanism for responsible image enhancement. Unlike conventional context-agnostic methods, the proposed framework dynamically adjusts optimization priorities according to semantic content. Experimental pilot results demonstrate improved perceptual realism and identity preservation in face-sensitive scenarios. These findings highlight the importance of context-aware objective scheduling for safer and more reliable enhancement in sensitive applications.

This work introduces a context-conditioned adaptive loss weighting mechanism for responsible image enhancement. Unlike conventional context-agnostic methods, the proposed framework dynamically adjusts optimization priorities according to semantic content. Experimental pilot results demonstrate improved perceptual realism and identity preservation in face-sensitive scenarios. These findings highlight the importance of context-aware objective scheduling for safer and more reliable enhancement in sensitive applications.

References

- [1] Q. H. Cap, et al., "A Practical Framework for Unsupervised Structure Preservation in Medical Image Enhancement," arXiv preprint arXiv:2304.01864, 2023.
- [2] Y. Wu, et al., "Learning Semantic-Aware Knowledge Guidance for Low-Light Image Enhancement," in Proc. IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR), 2023.
- [3] G. Ohayon, M. Elad, and T. Michaeli, "Perceptual Fairness in Image Restoration," arXiv preprint arXiv:2405.13805, 2024.
- [4] M. Laszkiewicz, et al., "Benchmarking the Fairness of Image Upsampling Methods," in Proc. ACM Conf. on Fairness, Accountability, and Transparency (FAcT), 2024.
- [5] M. Ni, et al., "Responsible Visual Editing," in Proc. European Conf. on Computer Vision (ECCV), 2024.
- [6] P. Iacono and N. Khan, "Structure Preserving Cycle-GAN for Unsupervised Medical Image Domain Adaptation," arXiv preprint arXiv:2304.09164, 2023.
- [7] Y. Wu, et al., "Towards a Flexible Semantic Guided Model for Single-Image Low-Light Enhancement," *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2024.
- [8] J. Zhang, et al., "Low-Light Image Enhancement with Integrated Semantic-Aware Guidance," *Expert Systems with Applications*, vol. XXX, pp. XXX-XXX, 2025.