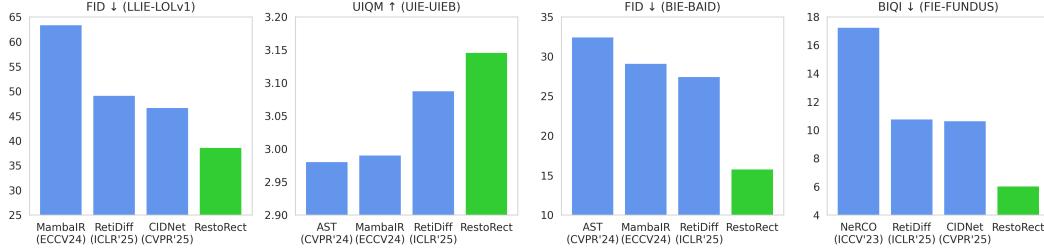


000 001 002 003 004 005 RESTORECT: DEGRADED IMAGE RESTORATION VIA 006 LATENT RECTIFIED FLOW & FEATURE DISTILLATION 007 008

009 **Anonymous authors**
010 Paper under double-blind review
011
012
013
014
015
016
017
018
019
020
021
022
023
024

ABSTRACT

010 Current approaches for restoration of degraded images face a critical trade-off:
011 high-performance models are too slow for practical use, while fast models produce
012 poor results. Knowledge distillation transfers teacher knowledge to students, but
013 existing static feature matching methods cannot capture how modern transformer
014 architectures dynamically generate features. We propose '**RestoRect**', a novel
015 Latent Rectified Flow Feature Distillation method for restoring degraded images.
016 We apply rectified flow to reformulate feature distillation as a generative process
017 where students learn to synthesize teacher-quality features through learnable tra-
018 jectories in latent space. Our framework combines Retinex theory for physics-
019 based decomposition with learnable anisotropic diffusion constraints, and trigono-
020 metric color space polarization. We introduce a Feature Layer Extraction loss for
021 robust knowledge transfer between different network architectures through cross-
022 normalized transformer feature alignment with percentile-based outlier detection.
023 RestoRect achieves better training stability, and faster convergence and inference
024 while preserving restoration quality. We demonstrate superior results across 15
025 image restoration datasets, covering 4 tasks, on 10 metrics.
026
027
028
029
030
031
032
033
034



035 Figure 1: RestoRect achieves superior performance on four image restoration tasks.
036
037

1 INTRODUCTION

040 Image restoration from degraded inputs including low-light (LLIE), underwater (UIE), backlit
041 (BAID), and fundus (FIE) enhancement, remains a key challenge in computer vision. Real-world
042 images often suffer from illumination degradation, noise, and compression artifacts that impair both
043 human perception and downstream tasks. Traditional optimization-based methods exploit phys-
044 ical priors but falter on images with complex degradations, while transformer-based deep learning
045 achieves strong restoration by learning rich multi-scale features. Generative approaches further
046 enhance quality, with diffusion models operating in latent spaces and integrating Retinex priors to
047 capture the complex distributions of natural images. However, such gains incur steep computational
048 costs, limiting real-time use. Knowledge distillation offers efficiency by transferring knowledge
049 from large teachers to compact students, but struggles with transformer-based restoration. Conven-
050 tional approaches compute static feature losses between teacher and student layers, neglecting the
051 dependency modeling, degrading student performance. Recent models such as Reti-Diff (He et al.,
052 2023) (Retinex priors) and HVI-CIDNet (Yan et al., 2024) (learnable color spaces) achieve good
053 restoration, but their distillation relies on static feature matching, which fails to capture generative
054 processes.

We propose **RestoRect**, which formulates knowledge distillation as a generative process through latent rectified flow. Instead of matching static features, student networks learn the dynamic synthesis of features through flow matching dynamics, using linear interpolation trajectories in latent space between noise and target features. This reduces sampling steps while preserving feature quality.

At the core of RestoRect is the **Feature Layer EXtraction (FLEX) Loss**, designed to address distribution mismatch in feature distillation. Unlike prior methods that assume teacher and student features share the same statistical space, FLEX normalizes both using student statistics, enabling meaningful comparison despite evolving feature distributions during training. To further stabilize learning, percentile-based outlier detection mitigates noisy or corrupted regions. Our framework integrates classical image processing with modern generative modeling: Retinex theory for physics-based decomposition, learnable anisotropic diffusion for structural consistency, and trigonometric color space polarization to eliminate the red discontinuity artifacts common in image transformations. Together, these components preserve both texture and color in restored images.

RestoRect employs a two-stage training paradigm for feature distillation. In Stage 1, the teacher network is trained with pixel, perceptual, and physics-based losses to achieve high-quality restoration. Stage 2 distills knowledge into the student via latent rectified flow. In its first phase, only rectified flow velocity predictors are trained while the main restoration network remains frozen. The pre-trained teacher extracts high-quality Retinex and image features from paired degraded and ground-truth inputs, which serve as targets for two rectified flow models. These models learn velocity fields that reproduce teacher-level features through learnable trajectories, enabling synthesis in only a few steps. In the second phase, the full restoration network is trained using these generative processes: velocity predictors dynamically generate student features, which are aligned with teacher features via our FLEX Loss that cross-normalizes multi-scale transformer representations and applies percentile-based outlier detection. This design allows the student to efficiently learn and generate teacher-quality features, achieving restoration performance comparable to diffusion-based methods while operating at significantly higher efficiency.

Our key technical contributions include: **1.** A novel framework modeling knowledge transfer as a generative process using latent rectified flow, where the student network learns velocity fields to synthesize teacher-quality features. **2.** A novel U-Net transformer architecture with Spatial Channel Layer Normalization (SCLN) and Query-Key normalization, for attention stability under degraded inputs. **3.** A novel Feature Layer EXtraction (FLEX) Loss using feature statistics to normalize both teacher and student representations for multi-scale alignment in transformers. **4.** Combining Retinex theory with learnable anisotropic diffusion constraints and trigonometric color space polarization to eliminate artifacts and boost restoration quality.

2 RELATED WORK

Degraded Image Restoration has evolved from classical signal processing to modern deep learning frameworks. Early approaches such as histogram equalization (Cheng & Shi, 2004), gamma correction (Huang et al., 2012), and Retinex theory (Edwin, 1977) provided interpretable solutions but failed to generalize across degradations. Retinex-based extensions (Fu et al., 2016; Li et al., 2018) incorporated physical priors for reflectance–illumination decomposition, yet remained constrained by hand-crafted assumptions. Deep learning enabled data-driven feature learning, with convolutional models by (Wei et al., 2018) and by (Wang et al., 2019) leveraging Retinex decomposition for improved color correction. Transformer-based methods further enhanced global illumination consistency (Zamir et al., 2022), while adaptive designs by (Xu et al., 2022) and state space models like by (Guo et al., 2024) advanced efficiency and context modeling. Specialized solutions addressed low-light enhancement (Guo et al., 2020; Jiang et al., 2021), underwater restoration (Naik et al., 2021; Guo et al., 2023), and backlit enhancement (Gaintseva et al., 2024; Jiang et al., 2021). Hybrid approaches such as by (He et al., 2025b) bridged optimization- and learning-based paradigms via deep unfolding, while by (Yan et al., 2024; 2025) introduced learnable color-space transformations to decouple brightness and chromaticity.

Image Generative Modeling aims to capture complex data distributions and synthesize realistic details. GAN-based methods (Cong et al., 2023; Jiang et al., 2021) achieved high-quality results but suffered from instability and mode collapse. Diffusion models improved fidelity through iterative denoising (Yi et al., 2023), though efficiency remained limited. Latent-space diffusion, such as Reti-

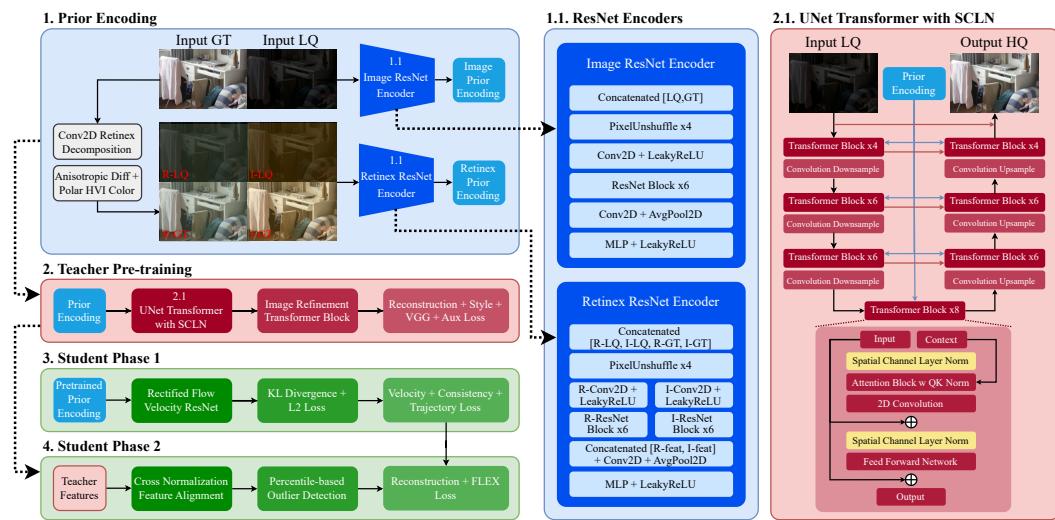
108 Diff (He et al., 2023), reduced overhead by incorporating Retinex priors. Flow-based approaches
 109 offered exact likelihoods and stable training (Kingma & Dhariwal, 2018), with rectified flow (Liu
 110 et al., 2022) enabling efficient straight-line sampling. Integrating generative priors into restoration
 111 networks has driven advances in knowledge distillation (Hinton et al., 2015), conditional and multi-
 112 scale generation (Saharia et al., 2022; Ho et al., 2022), and physics-informed restoration (Xia et al.,
 113 2023). Nonetheless, achieving real-time, high-fidelity restoration remains challenging due to the
 114 trade-off between generative quality and computational efficiency.

115 **Knowledge Distillation** enables compact models to inherit capabilities from larger teachers (Hin-
 116 ton et al., 2015). Early methods matched intermediate features (Romero et al., 2014) or atten-
 117 tion maps (Zagoruyko & Komodakis, 2016), using L2 losses (Heo et al., 2019) or attention trans-
 118 fer (Huang & Wang, 2017). For vision transformers, challenges from multi-head attention and
 119 positional encodings inspired approaches like distillation tokens in DeiT (Touvron et al., 2021) and
 120 attention matrix alignment (Wang et al., 2020). However, these strategies treat features as static tar-
 121 get, overlooking the dynamic generation in transformer architectures (Jiao et al., 2019). In image
 122 restoration, distillation is further complicated by multi-scale feature dependencies and complex dis-
 123 tributions (Zhang et al., 2022; Berrada et al., 2025). Architectural mismatches between teacher and
 124 student amplify these gaps, limiting transfer efficiency and degrading restoration quality, motivat-
 125 ing new paradigms that model feature generation as a learnable process rather than static matching
 126 (Bing et al., 2025).

3 METHODOLOGY

3.1 PROBLEM FORMULATION

132 We tackle efficient knowledge distillation for degraded image restoration, aiming to transfer knowl-
 133 edge from a powerful teacher \mathcal{F}_T to a lightweight student \mathcal{F}_S without sacrificing quality. Given a
 134 degraded input $I_{LQ} \in \mathbb{R}^{H \times W \times 3}$ and ground truth $I_{GT} \in \mathbb{R}^{H \times W \times 3}$, the objective is: $\mathcal{F}_S(I_{LQ}) \approx$
 135 $\mathcal{F}_T(I_{LQ}) \approx I_{GT}$. The main challenge is feature distribution mismatch between teacher and student.
 136 Standard distillation aligns features with simple distance metrics, which breaks down when dis-
 137 tributions differ significantly, especially in transformer-based networks where multi-head attention
 138 produces features with varying means, variances, and outlier characteristics.



156 **Figure 2: Training framework flowchart for RestoRect.** Starting from top left (1. Prior Encoding)
 157 the inputs go through retinex decomposition and pass through encoders (1.1 ResNet Encoders) to
 158 prepare image and retinex prior encodings. Next these prior encodings are pre-trained (2. Teacher
 159 Pre-training) with the teacher model (2.2 UNet Transformer with SCLN) using a reconstruction
 160 loss. Finally the frozen prior encodings and teacher model are used for student phase 1 and phase 2
 161 training using rectified flow loss. Full architecture details in Appendix A.7.

162 3.2 TEACHER NETWORK TRANSFORMER PRETRAINING
163

164 Our method uses well-established Retinex theory to derive physics-informed features as priors for
165 knowledge distillation. Retinex models an image I as the product of reflectance R and illumination
166 L : $I = R \odot L$, where R encodes surface properties and L captures lighting. We use two
167 decomposition networks, \mathcal{D}_l (low-light) and \mathcal{D}_h (normal-light), each mapping $\mathcal{D}(I) \rightarrow (R, L)$ with
168 $R \in \mathbb{R}^{H \times W \times 3}$ and $L \in \mathbb{R}^{H \times W \times 1}$ (Wu et al., 2022; He et al., 2023). This dual setup ensures robust
169 decomposition under diverse lighting. The decomposed components are then encoded (Figure 2(1)): a
170 Retinex encoder extracts features from $[R; L]$ via reflectance (192-dim) and illumination (64-dim)
171 pathways, while an image encoder processes raw image features to preserve holistic appearance.
172 Our teacher network uses U-Net transformer architecture (Huang et al., 2020; Cao et al., 2022) with
173 key innovations for robust image restoration. The hierarchical transformer architecture processes
174 multi-scale representations through encoder-decoder structures with skip connections, incorporating
175 specialized normalization and attention mechanisms designed for degraded image inputs. Tra-
176 ditional layer normalization operates independently on spatial and channel dimensions, potentially
177 losing critical spatial correlations essential for restoration tasks.

178 **Spatial Channel Layer Normalization (SCLN)** is introduced that captures global image statistics:
179 $SCLN(x) = (x - \mu_{global}) / (\sqrt{\sigma_{global}^2 + \epsilon}) \cdot \gamma$, where the global statistics are computed across flat-
180 tened spatial-channel dimensions. This novel formulation ensures that normalization captures both
181 local spatial patterns and global image characteristics, with learnable channel-wise scaling $\gamma \in \mathbb{R}^C$
182 that adapts to different feature semantics. Transformer-based restoration suffers from attention in-
183 stability during training, particularly with degraded inputs which have irregular noise patterns and
184 missing information. We apply normalization to query and key representations before attention com-
185 putation, which prevents attention weight saturation in degraded regions, and ensures stable gradi-
186 ents throughout the attention mechanism: $\text{Attn}(Q, K, V) = \text{softmax} \left(\frac{\text{Norm}(Q) \cdot \text{Norm}(K)^T}{\sqrt{d_k}} \cdot \tau \right) V$. The
187 teacher network processes both raw images and their Retinex decompositions through separate path-
188 ways. This design allows queries from reflectance components to attend to illumination structure,
189 preserving intrinsic scene properties. Figure 3 shows in blue our SCLN with QK norm achieves more
190 stable training compared to vanilla layer normalization without QK norm in red. To our knowledge
191 no previous restoration method has used this transformer architecture.

192
193 Table 1: Computational overhead comparison between LayerNorm (LN) and Spatial Channel Layer
194 Norm (SCLN) across different precisions on an MLP network. Results averaged over 3 random
195 seeds on $512 \times 512 \times 64$ resolution.

Precision	Single Layer Performance (ms)				Multi-layer Performance (ms)			
	LayerNorm	SCLN	Overhead (%)		LayerNorm	SCLN	Overhead (%)	
FP32	0.7059 ± 0.0004	0.7113 ± 0.0000	0.0054	+0.76	3.05 ± 0.00	3.06 ± 0.00	0.01	+0.33
FP16	0.1899 ± 0.0000	0.1991 ± 0.0001	0.0092	+4.82	1.21 ± 0.00	1.17 ± 0.00	-0.05	-3.72
BF16	0.1935 ± 0.0001	0.2004 ± 0.0001	0.0069	+3.57	1.23 ± 0.00	1.17 ± 0.00	-0.06	-4.67

203 The benchmark results in Table 1 show that SCLN introduces minimal computational overhead due
204 to its efficient normalization strategy. At the single-layer level, SCLN incurs only 0.76% overhead
205 in FP32, as computing statistics across spatial-channel dimensions requires marginally more op-
206 erations than channel-wise normalization. The slightly higher percentage overhead in FP16/BF16
207 (4.82% and 3.57%) is primarily an artifact of the dramatically reduced absolute inference times, Lay-
208 erNorm executes so quickly in lower precision that even negligible absolute differences appear larger
209 percentagewise. At the full network level, SCLN’s overhead becomes negligible (0.33% in FP32)
210 or even negative (-3.72% in FP16, -4.67% in BF16), suggesting superior memory access patterns
211 and cache efficiency when operations are repeated across multiple layers. Standard LayerNorm’s
212 repeated reshape operations (to_3d/to_4d conversions) accumulate overhead, while SCLN’s direct
213 4D tensor operations benefit from better spatial locality and reduced memory bandwidth pressure.
214 This explains why SCLN actually becomes faster than LayerNorm in lower-precision full-network
215 scenarios, making the trade-off highly favorable with PSNR improvement with no speed penalty.
Note that we train our RestoRect models with FP32, and Table 1 is a toy MLP example.

Auxiliary Constraints like anisotropic diffusion (Perona et al., 1994) and polarized HVI color spaces (Yan et al., 2024) (Yan et al., 2025) are incorporated that enforce edge-preserving texture matching and eliminate artifacts. The anisotropic diffusion operator computes: $\mathcal{A}(I) = \nabla \cdot (c(|\nabla I|)\nabla I)$, with the diffusion coefficient defined as: $c(|\nabla I|) = \exp(-|\nabla I|^2/s^2)$, where s is a learnable sensitivity parameter initialized as $s = 0.1$ and constrained to $s \in [0.01, 1.0]$ to prevent numerical instability. The texture consistency loss enforces structural similarity between input and predicted reflectance: $L_{tex} = \|\mathcal{A}(I_{input}) - \mathcal{A}(R_{pred})\|_1$. This constraint preserves essential edge structures while suppressing noise, maintains texture coherence across different scales, and provides gradient-based supervision for fine-grained details. We additionally enforce illumination smoothness through gradient-aware weighting: $L_{lum} = \sum_{i,j} w_{i,j} (|\nabla_x L_{i,j}|^2 + |\nabla_y L_{i,j}|^2)$, where $w_{i,j} = \exp(-|\nabla L_{i,j}|)$ provides adaptive regularization based on local gradient magnitude. Standard image color spaces exhibit critical limitations for restoration like discontinuities at the red boundary ($H = 0^\circ$ and $H = 360^\circ$) and degenerate mappings in dark regions. To address these fundamental limitations, polarized HVI (Horizontal-Vertical-Intensity) color space is introduced that eliminates these artifacts through trigonometric parameterization. The polarized transformation maps hue to continuous coordinates: $H_{polar} = C_k \cdot S \cdot \cos(\pi H/3)$, $V_{polar} = C_k \cdot S \cdot \sin(\pi H/3)$, $I_{polar} = I_{max} = \max(R, G, B)$, where the adaptive intensity collapse factor is: $C_k = k \cdot \sin(\pi I_{max}/2) + \epsilon$, with learnable density parameter k initialized to 1.0 and constrained to $k \in [0.1, 5.0]$. This formulation eliminates red discontinuity through periodic parameterization, provides robustness through adaptive intensity collapse that prevents degenerate mappings in dark regions, and maintains color relationships under illumination changes. While (Yan et al., 2024) (Yan et al., 2025) frames HVI as a representation transformation, we define an explicit color loss in HVI space. The polarized color loss is computed as:

$$L_{col} = \|H_{polar}^{pred} - H_{polar}^{gt}\|_1 + \|V_{polar}^{pred} - V_{polar}^{gt}\|_1 + \|I_{polar}^{pred} - I_{polar}^{gt}\|_1 \quad (1)$$

The primary reconstruction objective employs pixel-wise supervision through L1 loss: $L_{rec} = \|I_{pred} - I_{gt}\|_1$, where I_{pred} represents the network's restored output and I_{gt} denotes the ground truth high-quality image. To capture perceptual similarity beyond pixel-level differences, we incorporate perceptual loss using pre-trained VGG (Simonyan & Zisserman, 2014) features. The perceptual loss extracts multi-scale feature representations that align with human visual perception: $L_{vgg} = \sum_l \lambda_l \|\phi_l(I_{pred}) - \phi_l(I_{gt})\|_2^2$, where ϕ_l represents VGG features at layer l , and λ_l denotes layer-specific weights that emphasize semantically important features. We additionally incorporate style loss that captures texture and artistic consistency through Gram matrix matching: $L_{sty} = \sum_l \|G_l(\phi_l(I_{pred})) - G_l(\phi_l(I_{gt}))\|_F^2$, where $G_l(\phi_l(I)) = \phi_l(I)\phi_l(I)^T$ computes the Gram matrix at layer l , and $\|\cdot\|_F$ denotes the Frobenius norm. This novel combination ensures that the restored images maintain both structural accuracy and perceptual realism. The complete teacher training objective combines these losses:

$$L_{teach} = L_{rec} + L_{vgg} + L_{sty} + \lambda_{tex} L_{tex} + \lambda_{col} L_{col} + \lambda_{lum} L_{lum} \quad (2)$$

with $\lambda_{tex} = 0.05$, $\lambda_{col} = 0.05$, and $\lambda_{lum} = 0.2$. Figure 3 shows in green our how our auxiliary constraints allow training of a stronger teacher model with faster convergence. (He et al., 2023) previously used reconstruction and style loss with perceptual VGG features. To our knowledge, we are the first to implement anisotropic diffusion texture and illumination smoothness constraints with explicit HVI color loss.

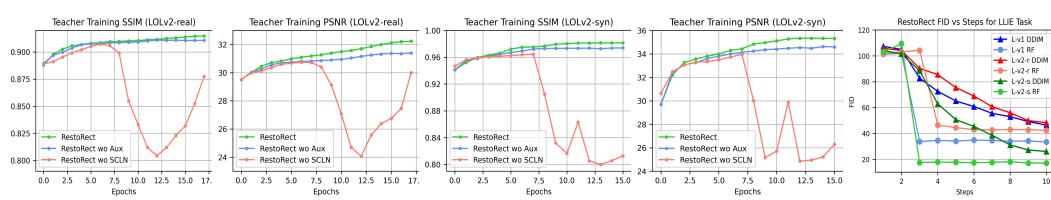


Figure 3: (1-4) Teacher model training with ablations of SCLN & QK Norm (red) and auxiliary losses (blue). (5) FID vs Steps inference performance show Rectified Flow (RF) student model producing high quality images in fewer steps compared to Denoising Diffusion Implicit Model (DDIM).

270
271

3.3 STUDENT NETWORK TRAINING WITH LATENT RECTIFIED FLOW

272

Traditional knowledge distillation treats feature transfer as static matching between teacher and student representations. This approach suffers from several limitations including assuming compatible feature distributions between architectures, lacking flexibility in handling multi-modal feature distributions, and being unable to adapt to varying complexity of restoration tasks. We reformulate knowledge distillation as a generative process using rectified flow, which models feature synthesis through straight-line paths in latent space. Given teacher features $\mathbf{f}_{\text{teach}} \in \mathbb{R}^d$ and noise $\mathbf{z} \sim \mathcal{N}(0, I)$, rectified flow defines the interpolation path: $\mathbf{x}_t = (1-t)\mathbf{z} + t\mathbf{f}_{\text{teach}}$, $t \in [0, 1]$. The velocity field represents the direction of optimal transport: $\mathbf{v}(\mathbf{x}_t, t) = \frac{d\mathbf{x}_t}{dt} = \mathbf{f}_{\text{teach}} - \mathbf{z}$. We train separate velocity prediction networks $\epsilon_\theta^{\text{ref}}$ and $\epsilon_\theta^{\text{img}}$ for reflectance and image features using the velocity matching objective: $L_{\text{vel}} = \mathbb{E}_{t, \mathbf{z}, \mathbf{f}_{\text{teach}}} [\|\epsilon_\theta(\mathbf{x}_t, t, \mathbf{c}) - \mathbf{v}(\mathbf{x}_t, t)\|_2^2]$, where \mathbf{c} represents conditioning information from the input image. Each velocity predictor implements a Residual MLP architecture. During inference, we solve the ODE using Euler's method with adaptive step sizing: $\mathbf{x}_{t+\Delta t} = \mathbf{x}_t + \Delta t \cdot \epsilon_\theta(\mathbf{x}_t, t, \mathbf{c})$. This requires only 1-4 integration steps compared to 10+ steps for DDIM models, providing significant computational advantages. Standard knowledge distillation losses (KL divergence, L2 distance) assume that teacher and student features exist in compatible statistical distributions. This assumption fails for complex transformer architectures, and when fine-tuning on different datasets, leading to suboptimal knowledge transfer (Lin et al., 2022).

288

FLEX (Feature Layer EXtraction) Loss addresses feature distribution mismatch through cross-normalization for distribution alignment, percentile-based outlier detection for robust training, and dynamic resolution-aware weighting for multi-scale importance. Unlike (Berrada et al., 2025) which is specialized for diffusion autoencoders, FLEX provides a general-purpose distillation loss that transfers feature distributions across heterogeneous teacher-student architectures. The key method is cross-normalization using student statistics. For each layer l , FLEX normalizes both teacher and student features using student statistics:

295

$$\mu_{\text{stud}}^l = \text{mean}(\mathbf{f}_{\text{stud}}^l), \quad \sigma_{\text{stud}}^l = \text{std}(\mathbf{f}_{\text{stud}}^l) + \epsilon, \quad \mathbf{f}_{\text{teach}}^l, \quad \mathbf{f}_{\text{stud}}^l = \frac{\mathbf{f}_{\text{stud}}^l - \mu_{\text{stud}}^l}{\sigma_{\text{stud}}^l}$$

297

This aligns both features to the student's distribution, enabling meaningful comparison across architecture capacity differences. FLEX incorporates fast percentile-based outlier detection to handle extreme values that destabilize training. This masking strategy prioritizes training stability over complete spatial coverage, as extreme outliers generate destabilizing gradients that outweigh their informational value. The outlier mask identifies reliable spatial locations: $M_{\text{reliable}}^{l,c,h,w} = \mathbb{I}[\|\mathbf{f}_{\text{stud}}^{l,c,\text{norm},h,w}\| \leq \tau_p^{l,c}]$, where $\tau_p^{l,c}$ is the p -th percentile of normalized feature magnitudes for layer l , channel c , with $p=95\%$ by default. FLEX computes dynamic resolution-based weights:

304

$$w_l^{\text{res}} = \max \left((H_{\text{base}} W_{\text{base}} / H_l W_l)^{0.25}, 0.1 \right)$$

306

where $(H_{\text{base}}, W_{\text{base}}) = (64, 64)$ ensures appropriate weighting across resolutions. The complete FLEX loss combines masked feature matching with dual weighting:

309

$$L_{\text{FLEX}} = \sum_l w_l^{\text{layer}} \cdot w_l^{\text{res}} \cdot \frac{\sum_{c,h,w} M_{\text{reliable}}^{l,c,h,w} \cdot \|\mathbf{f}_{\text{teach}}^{l,c,\text{norm},h,w} - \mathbf{f}_{\text{stud}}^{l,c,\text{norm},h,w}\|^2}{\sum_{c,h,w} M_{\text{reliable}}^{l,c,h,w} + \epsilon} \quad (3)$$

312

where w_l^{layer} represents predefined layer weights and the denominator normalizes by reliable elements. FLEX includes SNR-aware application, activating only when $t/T < \tau_{\text{SNR}} = 0.4$, focusing distillation on cleaner intermediate states. Cross-normalization enables stable transfer between different architectures, outlier detection prevents training instability, dynamic weighting balances multi-scale contributions, and streaming processing optimizes memory usage. Standard KD methods lack these capabilities, assuming compatible distributions and uniform spatial weighting.

313

Trajectory Consistency Regularization is introduced to ensure smooth and semantically consistent rectified flow trajectories, which prevents erratic feature generation and maintains coherence throughout the ODE integration process (Yang et al., 2024). We enforce smooth transitions between consecutive ODE steps: $L_{\text{trans}} = \sum_{i=1}^{N-1} \|\mathbf{f}_{\text{pred}}^{i+1} - \mathbf{f}_{\text{pred}}^i\|_2^2$, where $\mathbf{f}_{\text{pred}}^i$ represents predicted features at the i -th integration step. We ensure final generated features align with teacher targets: $L_{\text{target}} = \|\mathbf{f}_{\text{pred}}^{\text{final}} - \mathbf{f}_{\text{teach}}\|_2^2$. We enforce consistency in semantic feature representations across

324 the trajectory: $L_{cons} = \sum_{i=1}^N \text{cos_dist}(\mathbf{f}_{pred}^i, \mathbf{f}_{teach})$. The complete trajectory consistency loss is:
 325 $L_{traj} = \alpha_{trans} L_{trans} + \alpha_{target} L_{target} + \alpha_{cons} L_{cons}$, with $\alpha_{trans} = 0.1$, $\alpha_{target} = 0.5$, and
 326 $\alpha_{cons} = 0.2$. Our training protocol addresses the challenge of jointly learning velocity prediction
 327 and restoration quality through a principled two-phase approach. We first train rectified flow
 328 components while freezing the main restoration network:

$$329 \quad L_{phase1} = L_{vel}^{rex} + L_{vel}^{img} + \lambda_{KD} L_{KD} + \lambda_{traj} L_{traj} \quad (4)$$

330 This phase establishes stable velocity prediction capabilities without interference from restoration
 331 objective gradients. We use separate optimizers for reflectance and image velocity predictors with
 332 learning rates $lr_{rex} = 2 \times 10^{-4}$ and $lr_{img} = 2 \times 10^{-4}$. The complete network is then trained using
 333 features generated by learned velocity predictors, where $\lambda_{FLEX} = 0.15$, $\lambda_{vel} = 0.05$:

$$334 \quad L_{phase2} = L_{rec} + \lambda_{FLEX} L_{FLEX} + \lambda_{vel} (L_{vel}^{rex} + L_{vel}^{img}) \quad (5)$$

338 4 EXPERIMENTS

340 **Experimental Setup.** We implement our model in PyTorch and trained it on 8 NVIDIA H100 GPUs.
 341 Teacher pretraining is performed for 15-20 epochs depending on dataset convergence, while student
 342 phases I and II are each trained for 10 epochs. We use Adam optimizer with momentum terms (0.9,
 343 0.999). For fair comparison with prior work (He et al., 2023), we adopt the same configuration
 344 of transformer blocks, attention heads, and channel dimensions: [3, 3, 3, 3], [1, 2, 4, 8], and [64,
 345 128, 256, 512] from levels 1-4. During inference, we make 4 function evaluation calls for rectified
 346 flow generation, yielding faster generation and higher-quality outputs compared to state-of-the-art
 347 methods. Training follows the methodology of Reti-Diff and CIDNet across datasets and tasks.

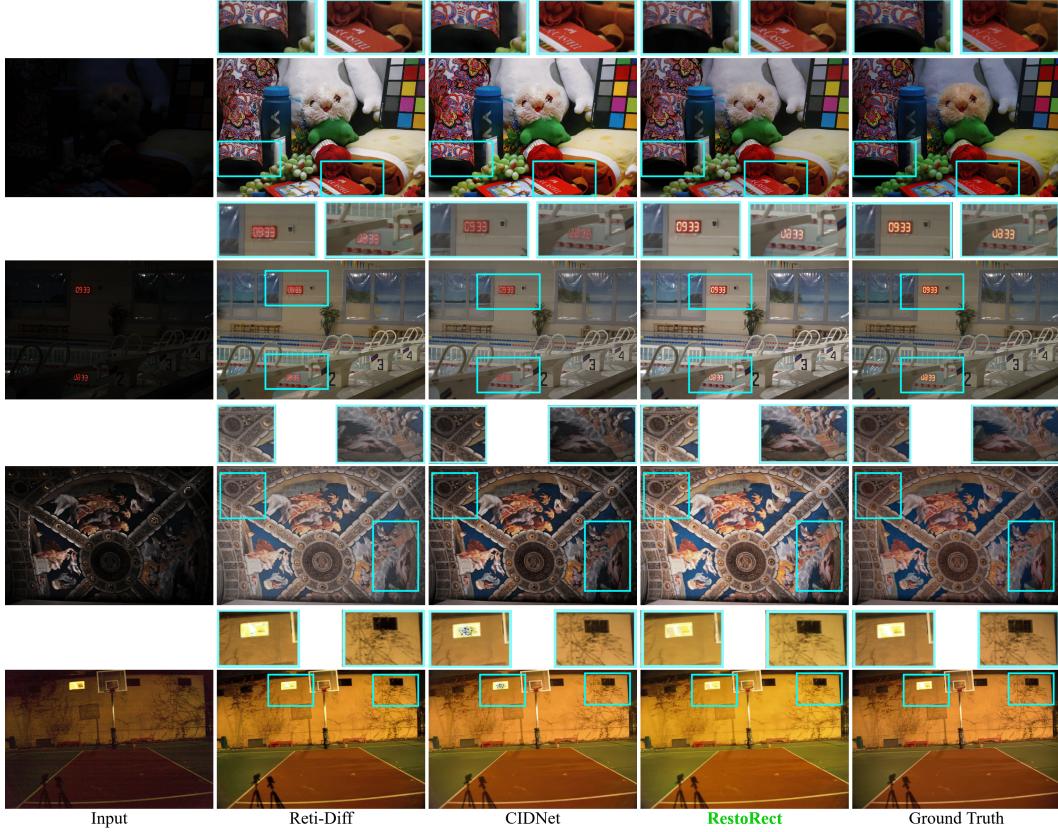
348 **Quantitative Evaluation.** For the low-light image enhancement (LLIE) task, we conduct experiments
 349 on LOL-v1 (Wei et al., 2018), LOL-v2-real, LOL-v2-syn (Yang et al., 2021), and SID (Chen
 350 et al., 2019). Performance is evaluated with PSNR, SSIM, FID, and BIQI (Hore & Ziou, 2010;
 351 Moorthy & Bovik, 2010), where higher PSNR/SSIM and lower FID/BIQI indicate better results.
 352 RestoRect achieves state-of-the-art performance across all datasets shown in Table 2, with improve-
 353 ments on almost every metric over the second-best methods (RetiDiff and CIDNet). The visual
 354 results shown in Figure 4 highlight clear improvements in fine grained details shown in cyan boxes
 355 (please zoom in for clarity).

356 For the underwater image enhancement (UIEB) task, we evaluate on UIEB (Li et al., 2019) and LSUI
 357 (Peng et al., 2023), using PSNR, SSIM, and UIQM (Panetta et al., 2015). Higher values across all
 358 metrics indicate better performance. RestoRect outperforms RetiDiff by 1.76dB PSNR on UIEB
 359 and matches its performance on LSUI while achieving superior SSIM scores shown in Table 3. For
 360 the backlit image enhancement (BAID) task, experiments are performed (Lv et al., 2022), with eval-
 361 uation on PSNR, SSIM, and FID. RestoRect demonstrates substantial improvements with 4.48dB
 362 PSNR gain over RetiDiff and 11.65 FID reduction shown in Table 4. Additionally, we test on real-
 363 world fundus image enhancement (FIE) (Shen et al., 2020) images using the LOL-v2-syn pretrained
 364 model, evaluating with BIQI and CLIPQA (Wang et al., 2023b), where higher CLIPQA values
 365 indicate better performance. RestoRect achieves the lowest BIQI score of 6.033, outperforming
 366 SNRNet shown in Table 5. The visual results shown in Figure 5 highlight our performance with
 367 details shown in yellow boxes (please zoom in for clarity). **We note that Reti-Diff baseline im-**
 368 **ages for UIEB and LSUI in middle row very closely match the ground truth while the scores are**
 369 **marginally worse than ours. This makes us believe that the publicly available checkpoints provided**
 370 **on Reti-Diff’s github might be overfitted to the validation set, unlike our model which has not seen**
 371 **the validation set.**

372 For real-world image restoration, we test on five unpaired datasets: DCIM (Lee et al., 2013), LIME
 373 (Guo et al., 2016), MEF (Wang et al., 2013), NPE (Ma et al., 2015), and VV (He et al., 2025a). Using
 374 the LOL-v2-syn pretrained model for inference, we evaluate with BRISQUE (Mittal et al., 2012),
 375 where lower values are better. RestoRect consistently outperforms CIDNet across most datasets,
 376 achieving the best scores on DCIM (16.56), LIME (16.12), and VV (24.42) as shown in Table
 377 6. We further evaluate on single image contrast enhancement (SICE) (Cai et al., 2018), which
 378 contains underexposed and overexposed images, training on the resized SICE training set and test
 379 on the datasets SICE-Mix and SICE-Grad (Zheng et al., 2022) with metrics PSNR, SSIM, LPIPS.

378 Table 2: LLIE task results. Best result shown in Green and second best shown in Blue. We also
 379 include our performance on RestoRect reconstruction teacher model.

380 Methods	381 LOL-v1				382 LOL-v2-real				383 LOL-v2-syn				384 SID			
	385 PSNR↑	386 SSIM↑	387 FID↓	388 BIQI↓	389 PSNR↑	390 SSIM↑	391 FID↓	392 BIQI↓	393 PSNR↑	394 SSIM↑	395 FID↓	396 BIQI↓	397 PSNR↑	398 SSIM↑	399 FID↓	400 BIQI↓
MRQ (Liu et al., 2023)	25.24	0.855	53.32	22.73	22.37	0.854	68.89	33.61	25.54	0.940	21.56	25.09	24.80	0.688	63.72	29.53
LAGC (Wang et al., 2023c)	24.53	0.866	59.73	25.50	22.20	0.863	70.34	31.70	25.58	0.941	21.58	30.32	23.17	0.640	78.80	30.56
DiffIR (Xia et al., 2023)	23.15	0.828	70.13	26.38	21.15	0.816	72.33	29.15	24.76	0.921	21.36	27.74	23.17	0.640	78.80	30.56
CUE (Zheng et al., 2023)	21.86	0.841	69.83	27.15	21.19	0.829	67.05	28.83	24.41	0.917	31.34	33.83	23.25	0.652	77.38	28.85
GSAD (Hou et al., 2023)	20.33	0.852	51.64	19.96	20.90	0.847	46.77	28.85	24.22	0.927	19.24	25.76	—	—	—	—
AST (Zhou et al., 2024)	21.09	0.858	87.67	21.23	21.68	0.857	91.81	25.17	22.25	0.927	19.20	20.78	—	—	—	—
Mamba (Guo et al., 2024)	22.33	0.863	63.39	20.17	21.97	0.840	56.09	24.46	25.75	0.958	17.95	20.37	21.14	0.656	154.76	32.72
RetiDiff (He et al., 2023)	25.35	0.866	49.14	17.75	22.97	0.857	43.18	23.66	27.53	0.951	13.82	15.77	25.53	0.692	51.66	25.58
CIDNet (Yan et al., 2024)	23.50	0.900	46.69	14.77	24.11	0.871	48.04	18.45	25.71	0.942	18.60	15.87	22.90	0.676	55.29	29.12
RestoRect (teacher only)	22.18	0.862	63.77	26.50	20.11	0.833	65.84	29.21	23.15	0.911	28.72	28.13	22.60	0.717	68.42	27.13
RestoRect	27.84	0.945	38.67	8.35	22.97	0.911	42.80	10.47	27.69	0.968	16.75	11.67	26.19	0.923	54.23	19.57



418 Figure 4: LLIE task visual results (Top to Bottom: LOL-v1, v2-real, v2-syn, SID). Clear improvements
 419 in fine grained details are shown in cyan boxes (please zoom in for clarity).

423 RestoRect achieves superior PSNR and SSIM performance over CIDNet by 1.6dB and 0.031 on
 424 SICE-Mix, and 2.0dB and 0.077 on SICE-Grad, as shown in Table 7.

425 **Qualitative Evaluation.** We conduct a user study to evaluate low-light image enhancement. Eight
 426 participants are shown 20 low-light images alongside enhanced outputs from RestoRect, Reti-Diff,
 427 and CIDNet (RAVE included for BAID dataset). In a blind comparison, subjects are asked to select
 428 the result that appears closest to the ground truth. Figure 8 presents the preference distributions,
 429 showing that RestoRect consistently achieves the highest preference across all five datasets, high-
 430 lighting its ability to generate visually appealing results perceived as closest to the ground truth.
 431 Figure 8 and Table 11 shows comparison of RestoRect’s student model parameter size (M) and
 432 GFLOPs against other transformer architecture baselines demonstrating efficiency.

432

Table 3: UIEB task results

433

Methods	UIEB			LSUI		
	PSNR↑	SSIM↑	UIQM↑	PSNR↑	SSIM↑	UIQM↑
SUnet (Naik et al., 2021)	18.28	0.855	2.942	20.89	0.875	2.746
PUIE (Fu et al., 2022)	21.38	0.882	3.021	23.70	0.902	2.974
UShape (Peng et al., 2023)	22.91	0.905	2.896	24.16	0.917	3.022
PUGAN (Cong et al., 2023)	23.05	0.897	2.902	25.06	0.916	3.106
ADP (Zhou et al., 2023)	22.90	0.892	3.005	24.28	0.913	3.075
NU2Net (Guo et al., 2023)	22.38	0.903	2.936	25.07	0.908	3.112
AST (Zhou et al., 2024)	22.19	0.908	2.981	27.46	0.916	3.107
Mamba (Guo et al., 2024)	22.60	0.939	2.991	27.68	0.916	3.118
RetiDiff (He et al., 2023)	24.12	0.910	3.088	28.10	0.929	3.208
RestoRect	25.88	0.950	3.121	28.10	0.937	3.229

434

Table 4: BAID task results

435

Methods	BAID		
	PSNR↑	SSIM↑	FID↓
EnGAN (Jiang et al., 2021)	17.96	0.819	43.55
URetinex (Wu et al., 2022)	19.08	0.845	42.26
CLIPPLIT (Liang et al., 2023)	21.13	0.853	37.30
DiffRet (Yi et al., 2023)	22.07	0.861	38.07
DiffIR (Xia et al., 2023)	21.10	0.835	40.35
AST (Zhou et al., 2024)	22.61	0.851	32.47
Mamba (Guo et al., 2024)	23.07	0.874	29.13
RAVE (Gaintseva et al., 2024)	21.26	0.872	64.89
RetiDiff (He et al., 2023)	23.19	0.876	27.47
RestoRect	27.67	0.965	15.82

436

Table 5: FIE task results

437

Methods	Fundus	
	BIQI↓	CLIPQ↑
SNRNet (Xu et al., 2022)	6.144	0.557
URetinex (Wu et al., 2022)	12.158	0.561
SCI (Ma et al., 2022)	23.527	0.552
MIRNet (Zamir et al., 2022)	14.925	0.527
FourLLE (Wang et al., 2023a)	7.741	0.508
CUE (Zheng et al., 2023)	11.721	0.448
NeRCO (Yang et al., 2023)	17.256	0.451
RetiDiff (He et al., 2023)	10.788	0.525
CIDNet (Yan et al., 2024)	10.663	0.529
RestoRect	6.033	0.503

438

439

440

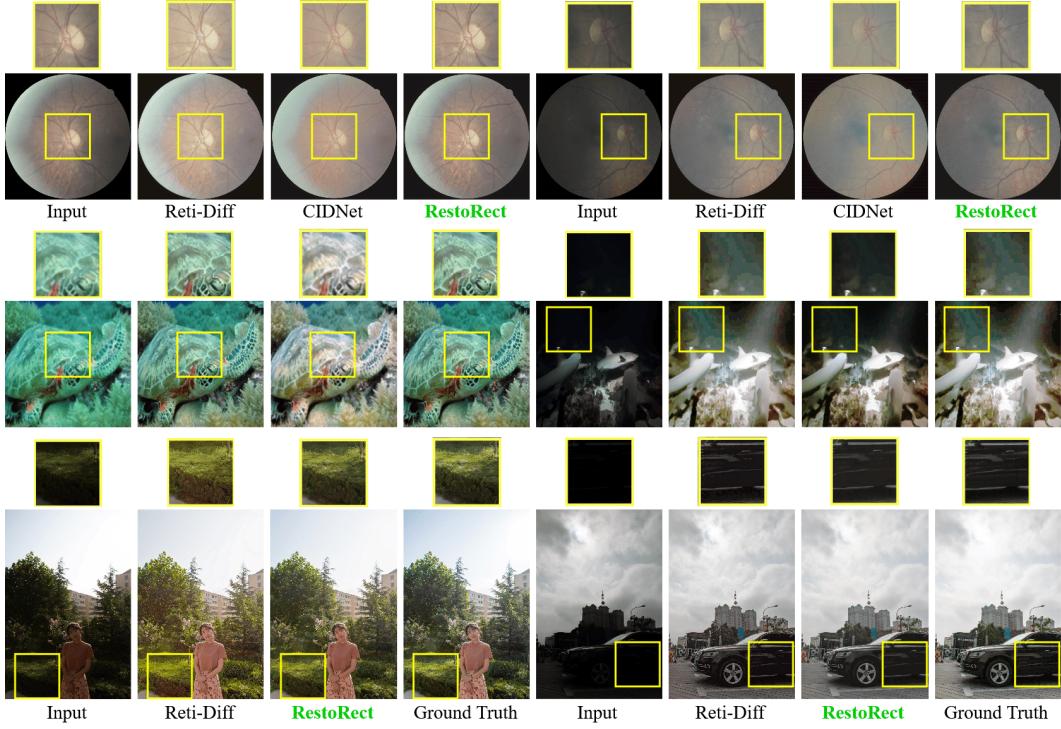


Figure 5: FIE (Top), UIEB (Middle Left), LSUI (Middle Right), BAID (Bottom) task visual results. Clear improvements in fine grained details are shown in yellow boxes (please zoom in for clarity).

441

442

443

444

445

446

447

448

449

450

451

452

453

454

455

456

457

458

459

460

461

462

463

464

465

466

467

468

469

Ablation and Generalizability. Figure 3 presents the results of teacher model training under different ablation settings. The removal of auxiliary constraints, such as anisotropic diffusion and the polarized HVI color space loss, is shown in blue. In contrast, the ablation of SCLN and QK normalization from the transformer block is shown in red, where a standard layer normalization and vanilla QK computation are used instead, following (He et al., 2023). As illustrated in green, the teacher model achieves the best performance with RestoRect when all proposed components are included. Table 8 further reports student model performance across different training and testing conditions on the LOL-v1, LOL-v2-real, and LOL-v2-synthetic datasets. In the table, '-FLEX' denotes models trained on the same dataset as the test set but without the FLEX loss. The FLEX training strategy demonstrates substantial improvements, with gains across all metrics compared to the full model results shown in Table 2. Subsequent rows in Table 8 evaluate cross-dataset transfer, where models trained on one dataset are tested on another, highlighting their strong generalization capacity. These results demonstrate that models trained for a given task can effectively transfer knowledge and serve as strong initialization points for fine-tuning on other datasets. Figure 7 shows visual results for LLIE task with ablation of SCLN and QK Norm, Auxiliary constraints, FLEX loss, compared to full RestoRect architecture and Ground Truth. Figure 3 also demonstrates the FID performance of RestoRect across different inference steps for the LLIE task. Our rectified flow formulation consistently outperforms (He et al., 2023) DDIM across all LLIE datasets, generating restored image within 3-4 steps, making it ideal for real time applications.

486 Table 6: Unpaired task results

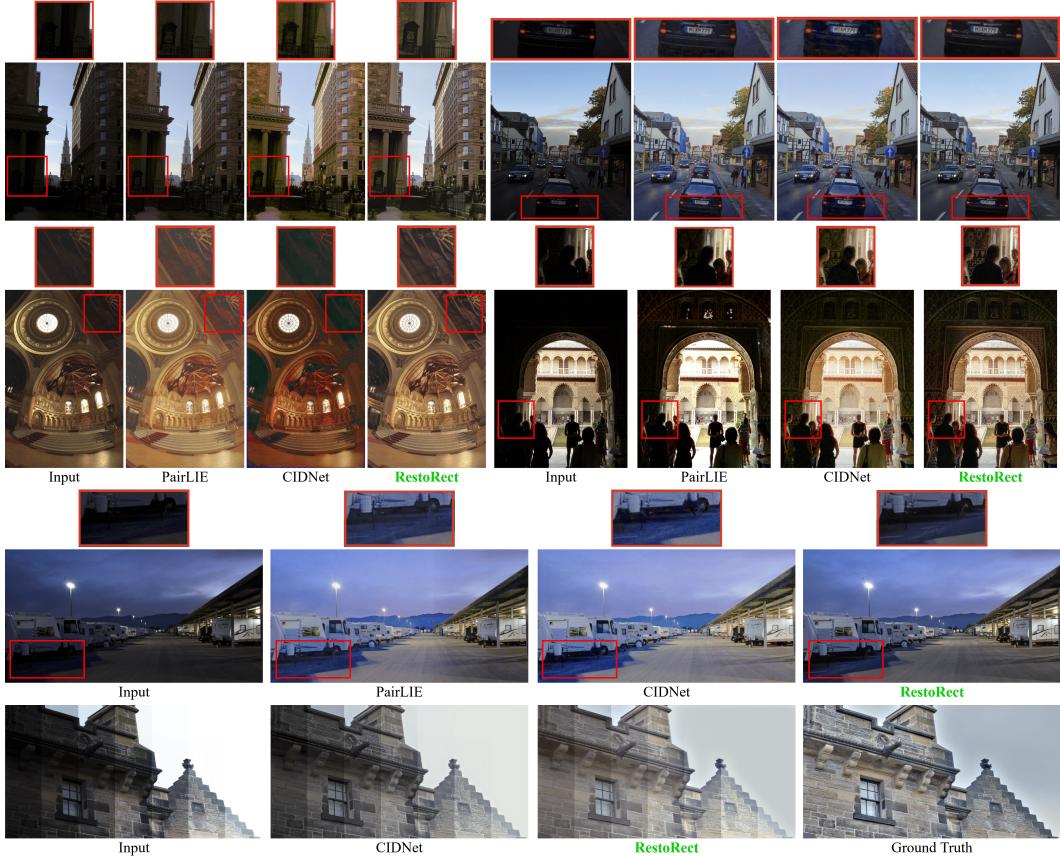
Methods	DCIM	LIME	MEF	NPE	VV	BRISQUE \downarrow
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	
KinD (Zhang et al., 2019)	48.72	39.91	49.94	36.85	50.56	12.397
ZeroDCE (Guo et al., 2020)	27.56	20.44	17.32	24.72	34.66	0.606
RUAS (Liu et al., 2021)	38.75	27.59	23.68	47.85	38.37	0.407
LLFlow (Wang et al., 2022)	26.36	27.06	30.27	28.86	31.67	12.428
SNRAware (Xu et al., 2022)	37.35	39.22	31.28	26.63	78.72	0.633
PairLIE (Fu et al., 2023)	33.31	25.23	27.53	28.27	39.13	0.382
CIDNet (Yan et al., 2024)	21.47	16.25	13.77	18.99	30.63	12.450
RestoRect	16.56	16.12	14.69	23.91	24.42	0.619

487 Table 7: SICE task results

Methods	SICE-Mix			SICE-Grad		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
RetiNet (Wei et al., 2018)	12.397	0.606	0.407	12.450	0.619	0.364
ZeroDCE (Guo et al., 2020)	12.428	0.633	0.382	12.475	0.644	0.334
URatinet (Wu et al., 2022)	10.903	0.600	0.402	10.894	0.610	0.356
RUAS (Liu et al., 2021)	8.684	0.493	0.525	8.628	0.494	0.499
LLFlow (Wang et al., 2022)	12.737	0.617	0.388	12.737	0.617	0.388
LEDNet (Zhou et al., 2022)	12.668	0.579	0.412	12.551	0.576	0.383
CIDNet (Yan et al., 2024)	13.425	0.636	0.362	13.446	0.648	0.318
RestoRect	15.041	0.667	0.393	15.447	0.715	0.354

488 Table 8: Ablation

Test	Train	PSNR \uparrow	SSIM \uparrow	FID \downarrow	BIQI \downarrow
		-FLEX	24.27	0.891	44.75
v1	v2-s	18.32	0.827	99.36	18.74
	v2-r	17.57	0.827	111.66	21.68
v2-r	v1	22.27	0.874	48.92	18.57
	v2-s	21.15	0.837	106.29	22.92
v2-s	v1	27.89	0.942	17.93	11.95
	v2-r	17.18	0.768	117.84	25.26



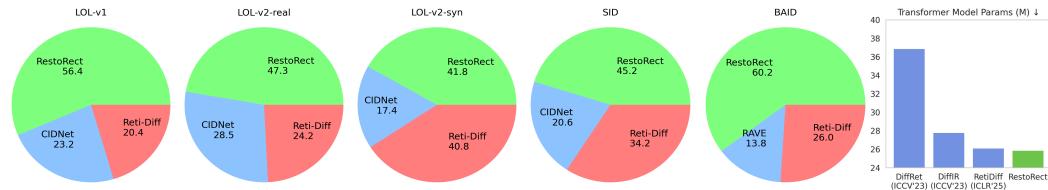
511 Figure 6: DCIM (Row 1 Left), LIME (Row 1 Right), MEF (Row 2 Left), NPE (Row 2 Right), VV (Row 3), SICE-Grad (Row 4) task visual results. Clear improvements in fine grained details are shown in red boxes (please zoom in for clarity).

525 The student model ablation study on LOL-v1 dataset in Table 9 evaluates the contribution of each
 526 component in the knowledge distillation framework. The metrics demonstrate that the combination
 527 of Rectified Flow (RF) with Trajectory Consistency (RF+TC) and FLEX loss components yields
 528 substantial improvements over traditional KD methods of (He et al., 2023), with the full configura-
 529 tion achieving 27.84 PSNR, 0.945 SSIM, 38.67 FID, and 8.35 BIQI on LOL-v1. The incremental
 530 addition of FLEX components shows that Cross-Normalization (CN) alone provides modest gains,
 531 while adding Percentile Masking (PM) delivers significant improvements (27.82 PSNR), and the
 532 complete FLEX formulation with Resolution Weighting (RW) achieves optimal performance. Hy-
 533 perparameter analysis reveals that the percentile threshold of 95% for outlier detection and SNR
 534 threshold of 0.4 provide the best balance between training stability and performance.

536 **Failure Cases.** Despite strong performance across many restoration tasks, RestoRect is limited in
 537 extreme degradation scenarios seen in Figure 9. Failure occurs when input images contain severe
 538 overexposure or underexposure with complete information loss in large spatial regions. The model
 539 produces unrealistic artifacts including color bleeding, checkerboard-like noise patterns, and hallu-
 540 cinatory textures that deviate from natural image statistics. These artifacts are highlighted in the



564
565 Figure 7: LLIE task visual results with ablation of SCLN and QK Norm, Auxiliary constraints,
566 FLEX loss, compared to full RestoRect model and Ground Truth.
567



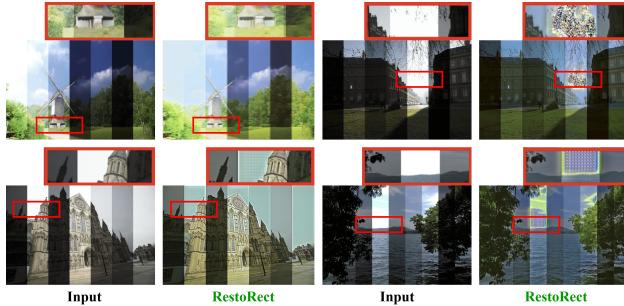
577
578 Figure 8: Qualitative human evaluation user study on LLIE and BAID datasets. Student model
579 parameter size (M) comparison against other transformer architecture baselines showing efficiency.
580

581 zoomed regions in red boxes, where recovered areas show synthetic-looking patterns rather than co-
582 herent scene content. These cases reveal that when the signal-to-noise ratio falls below a threshold,
583 velocity prediction networks struggle to generate meaningful features, resulting in the model relying
584 on learned priors that introduce perceptually implausible reconstructions.

585 Table 9: Student ablation study (L-v1).

Configuration	PSNR↑	SSIM↑	FID↓	BIQI↓
RetiDiff (Traditional KD)	25.35	0.866	49.14	17.75
RF(no TC) + no FLEX	24.10	0.855	49.28	17.95
RF(TC) + no FLEX	24.27	0.891	44.75	9.02
RF(TC) + FLEX(CN)	25.64	0.913	42.18	8.73
RF(TC) + FLEX(CN+PM)	27.82	0.947	39.25	8.42
RF(TC) + FLEX(CN+PM+RW)	27.84	0.945	38.67	8.35
Percentile Threshold $p = 85\%$	26.21	0.898	39.84	8.58
Percentile Threshold $p = 90\%$	27.58	0.946	39.12	8.43
Percentile Threshold $p = 95\%$	27.84	0.945	38.67	8.35
Percentile Threshold $p = 99\%$	27.67	0.943	38.91	8.41
SNR Threshold $\tau_{SNR} = 0.2$	27.15	0.937	40.12	8.64
SNR Threshold $\tau_{SNR} = 0.4$	27.84	0.945	38.67	8.35
SNR Threshold $\tau_{SNR} = 0.6$	27.89	0.941	39.28	8.47
SNR Threshold $\tau_{SNR} = 0.8$	26.92	0.934	40.53	8.71

585 Figure 9: Failure cases on SICE-Mix dataset.



594
 595 **Downstream Application.** For downstream application, RestoRect demonstrates strong transfer
 596 capability on the ExDark (Loh & Chan, 2019) low-light object detection task shown in Table 10.
 597 Following (Cai et al., 2023) and (He et al., 2023), low light images from the ExDark dataset were
 598 restored and object-detection task was performed using YOLOv3 model. Images enhanced by Re-
 599 stoRect achieve 77.1% mean average precision across 12 object categories, outperforming RetiDiff
 600 (He et al., 2023) (75.8%) and other methods. Improvements are observed in categories like Bicy-
 601 cle (85.8%), Bottle (79.5%), and Motor (77.9%), indicating that the restoration quality translates to
 602 real-world vision tasks.

602 Table 10: Downstream application on low-light image detection
 603 task on ExDark dataset.

Methods	Bicycle	Boat	Bottle	Bus	Car	Cat	Chair	Cup	Dog	Motor	People	Table	Mean
Restormer	77.0	71.0	68.8	91.6	77.1	62.5	57.3	68.0	69.6	69.2	74.6	49.7	69.7
SCI	73.4	68.0	69.5	86.2	74.5	63.1	59.5	61.0	67.3	63.9	73.2	47.3	67.2
SNR-Net	78.3	74.2	74.5	89.6	82.7	66.8	66.3	62.5	74.7	63.1	73.3	57.2	71.9
Retformer	78.1	74.5	74.2	91.2	82.2	65.0	63.3	67.0	75.4	68.6	75.3	55.6	72.5
RetiDiff	82.0	77.9	76.4	92.2	83.3	69.6	67.4	74.4	75.5	74.3	78.3	57.9	75.8
RestoRect	85.8	79.1	79.5	91.5	83.5	70.2	68.8	74.1	78.2	77.9	78.2	57.4	77.1

Table 11: Transformer model complexity comparison

Methods	Source	Params M	GFLOPs
Restormer	CVPR'22	26.13	144.25
Diff-Reti	ICCV'23	56.88	198.16
DiffIR	ICCV'23	27.80	35.32
GSAD	NIPS'23	17.17	670.33
Reti-Diff	ICLR'25	26.11	87.63
RestoRect	Ours	25.87	49.50

Table 12: RestoRect Image Quality Evaluation Results for all 15 datasets across 10 metrics

Dataset	PSNR	SSIM	FID	NIQE	LPIPS	BRISQ	BIQI	UCIQE	UIQM	CLIPQ
LOL-v1	27.85	0.94	38.67	7.47	0.11	27.16	8.35	0.52	2.60	0.499
LOL-v2 Real	22.97	0.91	42.81	7.74	0.13	28.44	10.48	0.51	2.88	0.500
LOL-v2 Syn	27.70	0.97	16.75	5.74	0.06	15.68	11.68	0.55	2.77	0.498
SID	26.19	0.92	54.23	5.87	0.15	20.05	19.57	0.85	2.38	0.498
UIEB	25.89	0.95	20.26	6.49	0.11	17.89	13.99	0.58	3.12	0.501
LSUI	28.10	0.94	17.83	5.04	0.18	21.82	16.06	0.57	3.23	0.499
BAID	27.68	0.97	15.83	8.11	0.06	34.39	10.49	0.56	2.87	0.501
Fundus	20.45	0.92	37.04	8.27	0.06	27.54	6.03	0.60	2.06	0.503
DCIM	19.92	0.82	72.72	6.36	0.17	16.57	10.26	0.57	2.33	0.499
LIME	18.36	0.76	101.31	6.12	0.21	16.13	11.76	0.59	2.19	0.497
MEF	17.20	0.69	74.06	6.13	0.26	14.70	11.23	0.56	2.83	0.499
NPE	16.28	0.77	63.75	7.10	0.18	23.91	12.91	0.53	2.64	0.498
VV	17.45	0.80	91.08	7.55	0.20	24.42	9.81	0.63	2.20	0.498
SICE (mix)	15.04	0.67	125.23	6.60	0.39	21.88	11.51	0.54	3.02	0.496
SICE (grad)	15.45	0.72	80.86	6.32	0.35	21.98	11.16	0.54	2.95	0.497

5 CONCLUSION

633 We present RestoRect, a generative knowledge distillation framework that reformulates degraded
 634 image restoration through latent rectified flow. Unlike traditional approaches that rely on static feature
 635 matching, RestoRect models feature transfer through learnable trajectories and introduces the
 636 FLEX loss for principled distribution alignment. Combined with a specialized U-Net transformer
 637 architecture and physics-based constraints, our method achieves state-of-the-art results across 15
 638 datasets covering low-light, underwater, backlit, and fundus enhancement. RestoRect delivers better
 639 perceptual quality with only 4 inference steps, making it both effective and computationally
 640 efficient. Beyond restoration, this generative distillation method highlights new opportunities for ef-
 641 ficient model compression and cross-architecture transfer in computer vision, establishing potential
 642 foundation for broader advances in fast high-quality image, and video restoration for future work.

6 LIST OF ACRONYMS

Table 13: List of acronyms used in this paper

Acronym	Full Form	Acronym	Full Form		
Tasks & Methods					
LLIE	Low-Light Image Enhancement	PSNR	Peak Signal-to-Noise Ratio		
UIE	Underwater Image Enhancement	SSIM	Structural Similarity Index Measure		
FIE	Fundus Image Enhancement	FID	Fréchet Inception Distance		
SCLN	Spatial Channel Layer Normalization	NIQE	Natural Image Quality Evaluator		
FLEX	Feature Layer EXtraction Loss	LPIPS	Learned Perceptual Image Patch Similarity		
Datasets					
LOL	Low-light Outdoor Lighting	BRISQUE	Blind/Referenceless Image Spatial Quality Evaluator		
SID	See in the Dark	BIQI	Blind Image Quality Index		
UIEB	Underwater Image Enhancement Bench	UCIQE	Underwater Color Image Quality Evaluation		
LSUI	Large Scale Underwater Image	UIQM	Underwater Image Quality Measure		
BAID	Backlit Image Dataset	CLIPQA	CLIP-based Image Quality Assessment		
DCIM	Digital Camera Image				
LIME	Low-light Image Enhancement				
MEF	Multi-Exposure Fusion				
NPE	Naturalness Preserved Enhancement				
VV	Video Visibility				
SICE	Single Image Contrast Enhancement				
Evaluation Metrics					
Other Technical Terms					

REFERENCES

668 Tariq Berrada, Pietro Astolfi, Melissa Hall, Marton Havasi, Yohann Benchetrit, Adriana Romero-
669 Soriano, Kartek Alahari, Michal Drozdzal, and Jakob Verbeek. Boosting latent diffusion with
670 perceptual objectives. In *The Thirteenth International Conference on Learning Representations*,
671 2025.

672 Zhaodong Bing, Linze Li, and Jiajun Liang. Optimizing knowledge distillation in transformers: En-
673 abling multi-head attention without alignment barriers. *arXiv preprint arXiv:2502.07436*, 2025.

674 Jianrui Cai, Shuhang Gu, and Lei Zhang. Learning a deep single image contrast enhancer from
675 multi-exposure images. *IEEE Transactions on Image Processing*, 27(4):2049–2062, 2018.

676 Yuanhao Cai, Hao Bian, Jing Lin, Haoqian Wang, Radu Timofte, and Yulun Zhang. Retinexformer:
677 One-stage retinex-based transformer for low-light image enhancement. In *Proceedings of the
678 IEEE/CVF international conference on computer vision*, pp. 12504–12513, 2023.

679 Hu Cao, Yueyue Wang, Joy Chen, Dongsheng Jiang, Xiaopeng Zhang, Qi Tian, and Manning Wang.
680 Swin-unet: Unet-like pure transformer for medical image segmentation. In *European conference
681 on computer vision*, pp. 205–218. Springer, 2022.

682 Chen Chen, Qifeng Chen, Minh N Do, and Vladlen Koltun. Seeing motion in the dark. In *Proceed-
683 ings of the IEEE/CVF International conference on computer vision*, pp. 3185–3194, 2019.

684 Heng-Da Cheng and XJ Shi. A simple and effective histogram equalization approach to image
685 enhancement. *Digital signal processing*, 14(2):158–170, 2004.

686 Runmin Cong, Wenyu Yang, Wei Zhang, Chongyi Li, Chun-Le Guo, Qingming Huang, and Sam
687 Kwong. Pugan: Physical model-guided underwater image enhancement using gan with dual-
688 discriminators. *IEEE Transactions on Image Processing*, 32:4472–4485, 2023.

689 Land Edwin. The retinex theory of color vision. *Scientific american*, 237:108–128, 1977.

690 Xueyang Fu, Delu Zeng, Yue Huang, Xiao-Ping Zhang, and Xinghao Ding. A weighted variational
691 model for simultaneous reflectance and illumination estimation. In *Proceedings of the IEEE
692 conference on computer vision and pattern recognition*, pp. 2782–2790, 2016.

693 Zhenqi Fu, Wu Wang, Yue Huang, Xinghao Ding, and Kai-Kuang Ma. Uncertainty inspired under-
694 water image enhancement. In *European conference on computer vision*, pp. 465–482. Springer,
695 2022.

702 Zhenqi Fu, Yan Yang, Xiaotong Tu, Yue Huang, Xinghao Ding, and Kai-Kuang Ma. Learning a sim-
 703 ple low-light image enhancer from paired low-light instances. In *Proceedings of the IEEE/CVF*
 704 *conference on computer vision and pattern recognition*, pp. 22252–22261, 2023.

705 Tatiana Gaintseva, Martin Benning, and Gregory Slabaugh. Rave: Residual vector embedding for
 706 clip-guided backlit image enhancement. In *European Conference on Computer Vision*, pp. 412–
 707 428. Springer, 2024.

708 Chunle Guo, Chongyi Li, Jichang Guo, Chen Change Loy, Junhui Hou, Sam Kwong, and Runmin
 709 Cong. Zero-reference deep curve estimation for low-light image enhancement. In *Proceedings of*
 710 *the IEEE/CVF conference on computer vision and pattern recognition*, pp. 1780–1789, 2020.

711 Chunle Guo, Ruiqi Wu, Xin Jin, Linghao Han, Weidong Zhang, Zhi Chai, and Chongyi Li. Under-
 712 water ranker: Learn which is better and how to be better. In *Proceedings of the AAAI conference*
 713 *on artificial intelligence*, volume 37, pp. 702–709, 2023.

714 Hang Guo, Jinmin Li, Tao Dai, Zhihao Ouyang, Xudong Ren, and Shu-Tao Xia. Mambair: A simple
 715 baseline for image restoration with state-space model. In *European conference on computer*
 716 *vision*, pp. 222–241. Springer, 2024.

717 Xiaojie Guo, Yu Li, and Haibin Ling. Lime: Low-light image enhancement via illumination map
 718 estimation. *IEEE Transactions on image processing*, 26(2):982–993, 2016.

719 Chunming He, Chengyu Fang, Yulun Zhang, Tian Ye, Kai Li, Longxiang Tang, Zhenhua Guo, Xiu
 720 Li, and Sina Farsiu. Reti-diff: Illumination degradation image restoration with retinex-based
 721 latent diffusion model. *arXiv preprint arXiv:2311.11638*, 2023.

722 Chunming He, Yuqi Shen, Chengyu Fang, Fengyang Xiao, Longxiang Tang, Yulun Zhang, Wang-
 723 meng Zuo, Zhenhua Guo, and Xiu Li. Diffusion models in low-level vision: A survey. *IEEE*
 724 *Transactions on Pattern Analysis and Machine Intelligence*, 2025a.

725 Chunming He, Rihan Zhang, Fengyang Xiao, Chengyu Fang, Longxiang Tang, Yulun Zhang, and
 726 Sina Farsiu. Unfoldir: Rethinking deep unfolding network in illumination degradation image
 727 restoration. *arXiv preprint arXiv:2505.06683*, 2025b.

728 Byeongho Heo, Jeesoo Kim, Sangdoo Yun, Hyojin Park, Nojun Kwak, and Jin Young Choi. A
 729 comprehensive overhaul of feature distillation. In *Proceedings of the IEEE/CVF international*
 730 *conference on computer vision*, pp. 1921–1930, 2019.

731 Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv*
 732 *preprint arXiv:1503.02531*, 2015.

733 Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P
 734 Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition
 735 video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022.

736 Alain Hore and Djemel Ziou. Image quality metrics: Psnr vs. ssim. In *2010 20th international*
 737 *conference on pattern recognition*, pp. 2366–2369. IEEE, 2010.

738 Jinhui Hou, Zhiyu Zhu, Junhui Hou, Hui Liu, Huanqiang Zeng, and Hui Yuan. Global structure-
 739 aware diffusion process for low-light image enhancement. *Advances in Neural Information Pro-
 740 cessing Systems*, 36:79734–79747, 2023.

741 Huimin Huang, Lanfen Lin, Ruofeng Tong, Hongjie Hu, Qiaowei Zhang, Yutaro Iwamoto, Xianhua
 742 Han, Yen-Wei Chen, and Jian Wu. Unet 3+: A full-scale connected unet for medical image
 743 segmentation. In *ICASSP 2020-2020 IEEE international conference on acoustics, speech and*
 744 *signal processing (ICASSP)*, pp. 1055–1059. Ieee, 2020.

745 Shih-Chia Huang, Fan-Chieh Cheng, and Yi-Sheng Chiu. Efficient contrast enhancement using
 746 adaptive gamma correction with weighting distribution. *IEEE transactions on image processing*,
 747 22(3):1032–1041, 2012.

748 Zehao Huang and Naiyan Wang. Like what you like: Knowledge distill via neuron selectivity
 749 transfer. *arXiv preprint arXiv:1707.01219*, 2017.

756 Yifan Jiang, Xinyu Gong, Ding Liu, Yu Cheng, Chen Fang, Xiaohui Shen, Jianchao Yang, Pan Zhou,
 757 and Zhangyang Wang. Enlightengan: Deep light enhancement without paired supervision. *IEEE*
 758 *transactions on image processing*, 30:2340–2349, 2021.

759

760 Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu.
 761 Tinybert: Distilling bert for natural language understanding. *arXiv preprint arXiv:1909.10351*,
 762 2019.

763 Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions.
 764 *Advances in neural information processing systems*, 31, 2018.

765

766 Chulwoo Lee, Chul Lee, and Chang-Su Kim. Contrast enhancement based on layered difference
 767 representation of 2d histograms. *IEEE transactions on image processing*, 22(12):5372–5384,
 768 2013.

769

770 Chongyi Li, Chunle Guo, Wenqi Ren, Runmin Cong, Junhui Hou, Sam Kwong, and Dacheng Tao.
 771 An underwater image enhancement benchmark dataset and beyond. *IEEE transactions on image*
 772 *processing*, 29:4376–4389, 2019.

773

774 Mading Li, Jiaying Liu, Wenhan Yang, Xiaoyan Sun, and Zongming Guo. Structure-revealing low-
 775 light image enhancement via robust retinex model. *IEEE transactions on image processing*, 27
 776 (6):2828–2841, 2018.

777

778 Zhexin Liang, Chongyi Li, Shangchen Zhou, Ruicheng Feng, and Chen Change Loy. Iterative
 779 prompt learning for unsupervised backlit image enhancement. In *Proceedings of the IEEE/CVF*
 International Conference on Computer Vision, pp. 8094–8103, 2023.

780

781 Sihao Lin, Hongwei Xie, Bing Wang, Kaicheng Yu, Xiaojun Chang, Xiaodan Liang, and Gang
 782 Wang. Knowledge distillation via the target-aware transformer. In *Proceedings of the IEEE/CVF*
 conference on computer vision and pattern recognition, pp. 10915–10924, 2022.

783

784 Risheng Liu, Long Ma, Jiaao Zhang, Xin Fan, and Zhongxuan Luo. Retinex-inspired unrolling with
 785 cooperative prior architecture search for low-light image enhancement. In *Proceedings of the*
 786 *IEEE/CVF conference on computer vision and pattern recognition*, pp. 10561–10570, 2021.

787

788 Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and
 transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022.

789

790 Yunlong Liu, Tao Huang, Weisheng Dong, Fangfang Wu, Xin Li, and Guangming Shi. Low-light
 791 image enhancement with multi-stage residue quantization and brightness-aware attention. In
 792 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 12140–12149,
 2023.

793

794 Yuen Peng Loh and Chee Seng Chan. Getting to know low-light images with the exclusively dark
 795 dataset. *Computer vision and image understanding*, 178:30–42, 2019.

796

797 Xiaoqian Lv, Shengping Zhang, Qinglin Liu, Haozhe Xie, Bineng Zhong, and Huiyu Zhou. Backl-
 798 itnet: A dataset and network for backlit image enhancement. *Computer Vision and Image Under-
 799 standing*, 218:103403, 2022.

800

801 Kede Ma, Kai Zeng, and Zhou Wang. Perceptual quality assessment for multi-exposure image
 fusion. *IEEE Transactions on Image Processing*, 24(11):3345–3356, 2015.

802

803 Long Ma, Tengyu Ma, Risheng Liu, Xin Fan, and Zhongxuan Luo. Toward fast, flexible, and robust
 804 low-light image enhancement. In *Proceedings of the IEEE/CVF conference on computer vision*
 805 and pattern recognition, pp. 5637–5646, 2022.

806

807 Anish Mittal, Anush Krishna Moorthy, and Alan Conrad Bovik. No-reference image quality assess-
 808 ment in the spatial domain. *IEEE Transactions on image processing*, 21(12):4695–4708, 2012.

809

Anush Krishna Moorthy and Alan Conrad Bovik. A two-step framework for constructing blind
 image quality indices. *IEEE Signal processing letters*, 17(5):513–516, 2010.

810 Ankita Naik, Apurva Swarnakar, and Kartik Mittal. Shallow-uwnet: Compressed model for under-
 811 water image enhancement (student abstract). In *Proceedings of the AAAI Conference on Artificial*
 812 *Intelligence*, volume 35, pp. 15853–15854, 2021.

813

814 Karen Panetta, Chen Gao, and Sos Agaian. Human-visual-system-inspired underwater image quality
 815 measures. *IEEE Journal of Oceanic Engineering*, 41(3):541–551, 2015.

816 Lintao Peng, Chunli Zhu, and Liheng Bian. U-shape transformer for underwater image enhance-
 817 ment. *IEEE transactions on image processing*, 32:3066–3079, 2023.

818

819 Pietro Perona, Takahiro Shiota, and Jitendra Malik. Anisotropic diffusion. In *Geometry-driven*
 820 *diffusion in computer vision*, pp. 73–92. Springer, 1994.

821

822 Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and
 823 Yoshua Bengio. Fitnets: Hints for thin deep nets. arxiv 2014. *arXiv preprint arXiv:1412.6550*,
 824 2014.

825 Chitwan Saharia, William Chan, Huiwen Chang, Chris Lee, Jonathan Ho, Tim Salimans, David
 826 Fleet, and Mohammad Norouzi. Palette: Image-to-image diffusion models. In *ACM SIGGRAPH*
 827 *2022 conference proceedings*, pp. 1–10, 2022.

828

829 Ziyi Shen, Huazhu Fu, Jianbing Shen, and Ling Shao. Modeling and enhancing low-quality retinal
 830 fundus images. *IEEE Transactions on Medical Imaging*, 40(3):996–1006, 2020.

831 Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image
 832 recognition. *arXiv preprint arXiv:1409.1556*, 2014.

833

834 Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and
 835 Hervé Jégou. Training data-efficient image transformers & distillation through attention. In
 836 *International conference on machine learning*, pp. 10347–10357. PMLR, 2021.

837

838 Chenxi Wang, Hongjun Wu, and Zhi Jin. Fourllie: Boosting low-light image enhancement by fourier
 839 frequency information. In *Proceedings of the 31st ACM International Conference on Multimedia*,
 840 pp. 7459–7469, 2023a.

841

842 Jianyi Wang, Kelvin CK Chan, and Chen Change Loy. Exploring clip for assessing the look and
 843 feel of images. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pp.
 844 2555–2563, 2023b.

845

846 Ruixing Wang, Qing Zhang, Chi-Wing Fu, Xiaoyong Shen, Wei-Shi Zheng, and Jiaya Jia. Underex-
 847 posed photo enhancement using deep illumination estimation. In *Proceedings of the IEEE/CVF*
 848 *conference on computer vision and pattern recognition*, pp. 6849–6857, 2019.

849

850 Shuhang Wang, Jin Zheng, Hai-Miao Hu, and Bo Li. Naturalness preserved enhancement algorithm
 851 for non-uniform illumination images. *IEEE transactions on image processing*, 22(9):3538–3548,
 852 2013.

853

854 Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. Minilm: Deep self-
 855 attention distillation for task-agnostic compression of pre-trained transformers. *Advances in neu-*
 856 *ral information processing systems*, 33:5776–5788, 2020.

857

858 Yinglong Wang, Zhen Liu, Jianzhuang Liu, Songcen Xu, and Shuaicheng Liu. Low-light image en-
 859 hancement with illumination-aware gamma correction and complete image modelling network. In
 860 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 13128–13137,
 861 2023c.

862

863 Yufei Wang, Renjie Wan, Wenhan Yang, Haoliang Li, Lap-Pui Chau, and Alex Kot. Low-light
 864 image enhancement with normalizing flow. In *Proceedings of the AAAI conference on artificial*
 865 *intelligence*, volume 36, pp. 2604–2612, 2022.

866

867 Chen Wei, Wenjing Wang, Wenhan Yang, and Jiaying Liu. Deep retinex decomposition for low-light
 868 enhancement. *arXiv preprint arXiv:1808.04560*, 2018.

864 Wenhui Wu, Jian Weng, Pingping Zhang, Xu Wang, Wenhan Yang, and Jianmin Jiang. Uretinex-net:
 865 Retinex-based deep unfolding network for low-light image enhancement. In *Proceedings of the*
 866 *IEEE/CVF conference on computer vision and pattern recognition*, pp. 5901–5910, 2022.

867 Bin Xia, Yulun Zhang, Shiyin Wang, Yitong Wang, Xinglong Wu, Yapeng Tian, Wenming Yang,
 868 and Luc Van Gool. Diffir: Efficient diffusion model for image restoration. In *Proceedings of the*
 869 *IEEE/CVF international conference on computer vision*, pp. 13095–13105, 2023.

870 Xiaogang Xu, Ruixing Wang, Chi-Wing Fu, and Jiaya Jia. Snr-aware low-light image enhancement.
 871 In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp.
 872 17714–17724, 2022.

873 Qingsen Yan, Yixu Feng, Cheng Zhang, Pei Wang, Peng Wu, Wei Dong, Jinqiu Sun, and Yanning
 874 Zhang. You only need one color space: An efficient network for low-light image enhancement.
 875 *arXiv preprint arXiv:2402.05809*, 2024.

876 Qingsen Yan, Kangbiao Shi, Yixu Feng, Tao Hu, Peng Wu, Guansong Pang, and Yanning Zhang.
 877 Hvi-cidnet+: Beyond extreme darkness for low-light image enhancement. *arXiv preprint*
 878 *arXiv:2507.06814*, 2025.

879 Ling Yang, Zixiang Zhang, Zhilong Zhang, Xingchao Liu, Minkai Xu, Wentao Zhang, Chenlin
 880 Meng, Stefano Ermon, and Bin Cui. Consistency flow matching: Defining straight flows with
 881 velocity consistency. *arXiv preprint arXiv:2407.02398*, 2024.

882 Shuzhou Yang, Moxuan Ding, Yanmin Wu, Zihan Li, and Jian Zhang. Implicit neural representation
 883 for cooperative low-light image enhancement. In *Proceedings of the IEEE/CVF international*
 884 *conference on computer vision*, pp. 12918–12927, 2023.

885 Wenhuan Yang, Wenjing Wang, Haofeng Huang, Shiqi Wang, and Jiaying Liu. Sparse gradient
 886 regularized deep retinex network for robust low-light image enhancement. *IEEE Transactions on*
 887 *Image Processing*, 30:2072–2086, 2021.

888 Xunpeng Yi, Han Xu, Hao Zhang, Linfeng Tang, and Jiayi Ma. Diff-retinex: Rethinking low-
 889 light image enhancement with a generative diffusion model. In *Proceedings of the IEEE/CVF*
 890 *international conference on computer vision*, pp. 12302–12311, 2023.

891 Sergey Zagoruyko and Nikos Komodakis. Paying more attention to attention: Improving the per-
 892 formance of convolutional neural networks via attention transfer. *arXiv preprint arXiv:1612.03928*,
 893 2016.

894 Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-
 895 Hsuan Yang, and Ling Shao. Learning enriched features for fast image restoration and enhance-
 896 ment. *IEEE transactions on pattern analysis and machine intelligence*, 45(2):1934–1948, 2022.

897 Yonghua Zhang, Jiawan Zhang, and Xiaojie Guo. Kindling the darkness: A practical low-light
 898 image enhancer. In *Proceedings of the 27th ACM international conference on multimedia*, pp.
 899 1632–1640, 2019.

900 Zizhao Zhang, Han Zhang, Long Zhao, Ting Chen, Sercan Ö Arikan, and Tomas Pfister. Nested
 901 hierarchical transformer: Towards accurate, data-efficient and interpretable visual understanding.
 902 In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pp. 3417–3425,
 903 2022.

904 Naishan Zheng, Man Zhou, Yanmeng Dong, Xiangyu Rui, Jie Huang, Chongyi Li, and Feng Zhao.
 905 Empowering low-light image enhancer through customized learnable priors. In *Proceedings of*
 906 *the IEEE/CVF International Conference on Computer Vision*, pp. 12559–12569, 2023.

907 Shen Zheng, Yiling Ma, Jinqian Pan, Changjie Lu, and Gaurav Gupta. Low-light image and video
 908 enhancement: A comprehensive survey and beyond. *arXiv preprint arXiv:2212.10772*, 2022.

909 Jingchun Zhou, Qian Liu, Qiuping Jiang, Wenqi Ren, Kin-Man Lam, and Weishi Zhang. Under-
 910 water camera: Improving visual perception via adaptive dark pixel prior and color correction.
 911 *International Journal of Computer Vision*, pp. 1–19, 2023.

918 Shangchen Zhou, Chongyi Li, and Chen Change Loy. Lednet: Joint low-light enhancement and
919 deblurring in the dark. In *European conference on computer vision*, pp. 573–589. Springer, 2022.
920

921 Shihao Zhou, Duosheng Chen, Jinshan Pan, Jinglei Shi, and Jufeng Yang. Adapt or perish: Adaptive
922 sparse transformer with attentive feature refinement for image restoration. In *Proceedings of the*
923 *IEEE/CVF conference on computer vision and pattern recognition*, pp. 2952–2963, 2024.

924

925

926

927

928

929

930

931

932

933

934

935

936

937

938

939

940

941

942

943

944

945

946

947

948

949

950

951

952

953

954

955

956

957

958

959

960

961

962

963

964

965

966

967

968

969

970

971

972 A APPENDIX
973974 A.1 ETHICS STATEMENT
975976 LLMs were only used for editorial assistance and polishing grammar for the manuscript, with no
977 participation in technical interpretation, or content development.
978979 A.2 REPRODUCIBILITY STATEMENT
980981 Code and pretrained model weights will be released upon acceptance.
982983 A.3 BROADER IMPACT
984985 Efficient image restoration has positive applications in medical imaging, autonomous systems, and
986 accessibility. No significant negative societal impacts are identified by us.
987988 A.4 THEORETICAL JUSTIFICATION OF FLEX LOSS
989990 We provide theoretical justification for FLEX’s key design choices to ensure stable optimization
991 dynamics.
992993 **Assumption 1** (Feature Boundedness): Teacher and student features are bounded during training:
994 $\|\mathbf{f}_{\text{teach}}^l\|, \|\mathbf{f}_{\text{stud}}^l\| \leq M$ for some constant $M > 0$.
995996 **Assumption 2** (Non-degeneracy): Student feature standard deviations satisfy $\sigma_{\text{stud}}^l \geq \sigma_{\min} > 0$ to
997 prevent division by zero in normalization.
998999 **Claim 1** Cross-normalization using student statistics prevents gradient explosion when teacher and
1000 student features have different scales.
10011002 *Justification:* Standard feature matching $L = \|\mathbf{f}_{\text{teach}} - \mathbf{f}_{\text{stud}}\|^2$ produces gradients proportional to
1003 $(\mathbf{f}_{\text{teach}} - \mathbf{f}_{\text{stud}})$. When teacher features are much larger than student features, this difference can be
arbitrarily large, causing unstable training.
10041005 FLEX cross-normalization ensures both normalized features have the same scale:
1006

1007
$$\mathbf{f}_{\text{teach}}^{\text{norm}} = \frac{\mathbf{f}_{\text{teach}} - \mu_{\text{stud}}}{\sigma_{\text{stud}}}, \quad \mathbf{f}_{\text{stud}}^{\text{norm}} = \frac{\mathbf{f}_{\text{stud}} - \mu_{\text{stud}}}{\sigma_{\text{stud}}} \quad (6)$$

1008

1009 Both normalized features have bounded variance, preventing gradient explosion regardless of the
1010 original scale mismatch.
10111012 **Claim 2** Percentile-based masking provides robustness to feature corruption.
10131014 *Justification:* By masking extreme values above the p -th percentile (default $p = 95\%$), FLEX fo-
1015 cuses learning on reliable feature regions. If corruption affects only a small fraction of spatial
1016 locations, most corrupted features will exceed the percentile threshold and be masked out. By ex-
1017 cluding the top 5% extreme activations, FLEX prevents gradient dominance by outliers, ensuring
that meaningful feature patterns rather than numerical instabilities drive the optimization.
10181019 For corruption affecting $\alpha < (100 - p)/100$ of spatial locations, the outlier detection will identify
and exclude most corrupted regions, limiting their impact on the overall loss.
10201021 **Claim 3** The resolution weighting $w_l^{\text{res}} = \max\left((H_{\text{base}}W_{\text{base}}/H_lW_l)^{0.25}, 0.1\right)$ balances multi-scale
1022 contributions.
10231024 *Justification:* Higher resolution features contain more spatial elements, potentially dominating the
1025 loss. The inverse relationship with spatial resolution prevents this dominance. The 0.25 exponent
provides gradual rather than aggressive down-weighting, preserving fine-grained information while
1026 preventing over-emphasis on high-resolution layers.
1027

1026
1027

A.5 THEORETICAL JUSTIFICATION OF RECTIFIED FLOW FOR KNOWLEDGE DISTILLATION

1028
1029

We provide theoretical grounding for reformulating knowledge distillation as a rectified flow process.

1030
1031

Rectified Flow Formulation: For teacher features $\mathbf{f}_{\text{teach}}$ and noise $\mathbf{z} \sim \mathcal{N}(0, \mathbf{I})$, we define the linear interpolation path:

1032
1033
1034

$$\mathbf{x}_t = (1 - t)\mathbf{z} + t\mathbf{f}_{\text{teach}}, \quad t \in [0, 1] \quad (7)$$

The corresponding velocity field is:

1035
1036

$$\mathbf{v}(\mathbf{x}_t, t) = \frac{d\mathbf{x}_t}{dt} = \mathbf{f}_{\text{teach}} - \mathbf{z} \quad (8)$$

1037

The student network learns a velocity predictor $\mathbf{v}_\theta(\mathbf{x}_t, t)$ by minimizing:

1038
1039
1040

$$L_{\text{vel}} = \mathbb{E}_{t, \mathbf{z}, \mathbf{f}_{\text{teach}}} [\|\mathbf{v}_\theta(\mathbf{x}_t, t) - \mathbf{v}(\mathbf{x}_t, t)\|_2^2] \quad (9)$$

1041
1042
1043

Claim 4 (Transport Cost Reduction): Learning rectified flow reduces convex transport costs compared to arbitrary couplings between noise and target features.

1044
1045
1046
1047
1048
1049

Justification: Following Liu et al. (2022), the rectified flow procedure transforms an arbitrary coupling between distributions into a deterministic coupling with provably non-increasing convex transport costs. Specifically, for any convex cost function c , the expected cost $\mathbb{E}[c(\mathbf{f}_{\text{teach}} - \mathbf{z})]$ is reduced through the straight-line parameterization. This property ensures that the student learns an efficient transport map from noise to teacher-quality features, which is more sample-efficient than curved trajectories used in diffusion-based distillation methods like DDIM.

1050
1051

Claim 5 (Exact Simulation without Discretization): Straight-line ODEs can be solved exactly with Euler integration, enabling few-step inference.

1052
1053
1054
1055
1056
1057
1058
1059

Justification: The velocity field $\mathbf{v}(\mathbf{x}_t, t) = \mathbf{f}_{\text{teach}} - \mathbf{z}$ is constant along the trajectory, making the ODE solution exact: $\mathbf{x}_1 = \mathbf{x}_0 + \int_0^1 \mathbf{v}(\mathbf{x}_t, t) dt = \mathbf{z} + (\mathbf{f}_{\text{teach}} - \mathbf{z}) = \mathbf{f}_{\text{teach}}$. This means Euler integration with step size Δt incurs zero discretization error for perfectly straight paths. In practice, the learned velocity predictor \mathbf{v}_θ approximates this constant field, allowing accurate simulation with as few as 1-4 steps. In contrast, curved trajectories (e.g., DDIM’s probability flow ODE) require many more steps to achieve similar accuracy, as shown empirically in Figure 3 where our method achieves lower FID with 4 steps than DDIM with 10 steps.

1060
1061

A.6 ARCHITECTURE CHOICE JUSTIFICATION

1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072

Using separate networks for low-light and normal-light images ensures robust Retinex decomposition across illumination conditions. As reported by Reti-Diff and Diff-IR, a single adaptive network would require more complex conditioning mechanisms. The two-phase student training approach addresses fundamental optimization challenges in generative knowledge distillation. Phase separation prevents objective conflicts as simultaneously learning velocity prediction and image reconstruction creates competing gradients. The velocity predictor tries to match teacher feature distributions while the reconstruction network optimizes for pixel-level accuracy. These objectives can work against each other, leading to suboptimal solutions. Feature space stabilization where Phase 1 establishes stable feature generation capabilities before introducing reconstruction complexity. This ensures the velocity predictors learn meaningful feature trajectories rather than shortcuts that minimize reconstruction error. Only the student network is deployed during inference, with no additional computational overhead compared to baseline restoration networks.

1073
1074
1075
1076
1077
1078
1079

For SNR threshold (0.4), we notice performance remains stable within ± 0.2 range. The threshold determines when FLEX loss is applied - too low (0.2) restricts learning, too high (0.8) includes noisy states. For outlier percentile value, we found that lower percentiles (90%) are more aggressive in outlier detection but may remove useful information. Higher percentiles (99%) retain more data but include potential artifacts. For resolution weighting exponent value (0.25), we notice values from 0.125-0.5 show similar performance. This parameter balances multi-scale contributions as lower values provide gentler weighting while higher values more aggressively down-weight high-resolution features.

1080 Standard layer normalization operates on channel dimensions independently, losing spatial correlations
1081 crucial for restoration tasks. SCLN computes global statistics across both spatial and channel
1082 dimensions, capturing holistic image characteristics while maintaining learnable channel-wise scal-
1083 ing. Degraded images contain irregular noise patterns that can cause attention weight saturation.
1084 Normalizing Q and K before attention computation prevents extreme attention weights and ensures
1085 stable gradient flow. The "RestoRect w/o SCLN" ablation (red curve in Figure 3) essentially repre-
1086 sents the RetiDiff baseline architecture using standard layer normalization, providing direct compar-
1087 ison between our architectural innovations and existing methods. FLEX loss becomes more critical
1088 for cross-domain scenarios, as feature distribution mismatches are more severe between different
1089 datasets than within-dataset variations. On modern GPUs (RTX 4090/H100), the difference between
1090 3-step (156ms) and 5-step (198ms) inference is minimal compared to the quality improvement. The
1091 5-step choice during inference optimizes the quality-practicality trade-off for real-world deployment
1092 across different types of datasets.
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133

1188
1189**Illumination Branch (E_I):**

$$\begin{aligned}
 1190 \quad E_I(X_I \oplus X_{I,gt}) = & \text{AdaptiveAvgPool2d} \\
 1191 \quad & \text{LeakyReLU}_{0.1}(\text{Conv2d}_{128 \rightarrow 64}^{3 \times 3}) \\
 1192 \quad & \text{LeakyReLU}_{0.1}(\text{Conv2d}_{128 \rightarrow 128}^{3 \times 3}) \\
 1193 \quad & \text{LeakyReLU}_{0.1}(\text{Conv2d}_{64 \rightarrow 128}^{3 \times 3}) \\
 1194 \quad & \text{ResBlock}^6(\text{LeakyReLU}_{0.1}(\text{Conv2d}_{32 \rightarrow 64}^{3 \times 3}) \\
 1195 \quad & (X_I \oplus X_{I,gt}))))))) \\
 1196 \quad & \\
 1197 \quad & \\
 \end{aligned} \tag{14}$$

1198
1199**Feature Fusion:**

$$\text{feat}_R = \text{MLP}_R(E_R(\text{output})) \in \mathbb{R}^{192} \tag{15}$$

$$\text{feat}_I = \text{MLP}_I(E_I(\text{output})) \in \mathbb{R}^{64} \tag{16}$$

$$\text{IPR}_{rex} = [\text{feat}_R; \text{feat}_I] \in \mathbb{R}^{256} \tag{17}$$

1204

1205 **Image ResNet Encoder (IRE)** The IRE processes raw image features:

1206

Input Processing:

$$X_{LQ} = \text{PixelUnshuffle}_4(I_{LQ}) \in \mathbb{R}^{H/4 \times W/4 \times 48} \tag{18}$$

$$X_{GT} = \text{PixelUnshuffle}_4(I_{GT}) \in \mathbb{R}^{H/4 \times W/4 \times 48} \tag{19}$$

$$X_{concat} = [X_{LQ}; X_{GT}] \in \mathbb{R}^{H/4 \times W/4 \times 96} \tag{20}$$

1212

Encoder Architecture:

1214

$$\begin{aligned}
 1215 \quad E(X_{concat}) = & \text{AdaptiveAvgPool2d} \\
 1216 \quad & \text{LeakyReLU}_{0.1}(\text{Conv2d}_{128 \rightarrow 256}^{3 \times 3}) \\
 1217 \quad & \text{LeakyReLU}_{0.1}(\text{Conv2d}_{128 \rightarrow 128}^{3 \times 3}) \\
 1218 \quad & \text{LeakyReLU}_{0.1}(\text{Conv2d}_{64 \rightarrow 128}^{3 \times 3}) \\
 1219 \quad & \text{ResBlock}^6(\text{LeakyReLU}_{0.1}(\text{Conv2d}_{96 \rightarrow 64}^{3 \times 3}) \\
 1220 \quad & (X_{concat}))))))) \\
 1221 \quad & \\
 1222 \quad & \\
 \end{aligned} \tag{21}$$

1223

Output:

1224

$$\text{IPR}_{img} = \text{LayerNorm}(\text{MLP}(E(\text{output}))) \in \mathbb{R}^{256}$$

1226

1227

A.7.3 UNET TRANSFORMER ARCHITECTURE

1228

1229

Spatial Channel Layer Normalization (SCLN) SCLN captures global image statistics across spatial and channel dimensions:

1231

1232

$$\mu_{global} = \frac{1}{B \cdot C \cdot H \cdot W} \sum_{b,c,h,w} x_{b,c,h,w} \tag{22}$$

1235

1236

1237

$$\sigma_{global}^2 = \frac{1}{B \cdot C \cdot H \cdot W} \sum_{b,c,h,w} (x_{b,c,h,w} - \mu_{global})^2 \tag{23}$$

1238

1239

1240

$$\text{SCLN}(x) = \frac{x - \mu_{global}}{\sqrt{\sigma_{global}^2 + \epsilon}} \cdot \gamma \tag{24}$$

1241

where $\gamma \in \mathbb{R}^C$ is learnable channel-wise scaling.

1242 **Retinex Attention** The Retinex attention mechanism uses separate conditioning for reflectance
 1243 and illumination components:

1244 **Feature Conditioning:**

1245 $k_{v_r} = \text{Linear}(k_v[0 : 192]) \in \mathbb{R}^{3C/4 \times 1 \times 1}$ (25)

1246 $k_{v_i} = \text{Linear}(k_v[192 : 256]) \in \mathbb{R}^{C/4 \times 1 \times 1}$ (26)

1247 $x_r = x[:, 0 : 3C/4, :, :] \odot k_{v_r} + x[:, 0 : 3C/4, :, :]$ (27)

1248 $x_i = x[:, 3C/4 : C, :, :] \odot k_{v_i} + x[:, 3C/4 : C, :, :]$ (28)

1249 **Query-Key-Value Computation:**

1250 $Q = \text{DepthwiseConv}(\text{Conv}(x_r)) \in \mathbb{R}^{B \times C \times H \times W}$ (29)

1251 $KV = \text{DepthwiseConv}(\text{Conv}(x_i)) \in \mathbb{R}^{B \times 2C \times H \times W}$ (30)

1252 $K, V = \text{split}(KV, \text{dim} = 1)$ (31)

1253 **Attention with QK Normalization:**

1254 $Q_{norm} = \text{LayerNorm}(Q), \quad K_{norm} = \text{LayerNorm}(K)$ (32)

1255 $Q_{norm} = \frac{Q_{norm}}{\|Q_{norm}\|_2}, \quad K_{norm} = \frac{K_{norm}}{\|K_{norm}\|_2}$ (33)

1256 $\text{Attn} = \text{softmax} \left(\frac{Q_{norm} \cdot K_{norm}^T}{\sqrt{d_k}} \cdot \tau \right)$ (34)

1257 $\text{Output} = \text{Attn} \cdot V$ (35)

1258 where τ is a learnable temperature parameter.

1259 **Multi-Scale U-Net Architecture Encoder Path:**

1260 Level 1: $[B, 48, H, W] \xrightarrow{4 \times \text{TransformerBlock}} [B, 48, H, W]$ (36)

1261 $\downarrow \text{Downsample}$ (37)

1262 Level 2: $[B, 96, H/2, W/2] \xrightarrow{6 \times \text{TransformerBlock}} [B, 96, H/2, W/2]$ (38)

1263 $\downarrow \text{Downsample}$ (39)

1264 Level 3: $[B, 192, H/4, W/4] \xrightarrow{6 \times \text{TransformerBlock}} [B, 192, H/4, W/4]$ (40)

1265 $\downarrow \text{Downsample}$ (41)

1266 Level 4: $[B, 384, H/8, W/8] \xrightarrow{8 \times \text{TransformerBlock}} [B, 384, H/8, W/8]$ (42)

1267 **Decoder Path with Skip Connections:**

1268 Level 3: $\text{Upsample} + \text{Concat} + \text{ReduceChannel} \xrightarrow{6 \times \text{TransformerBlock}}$ (43)

1269 Level 2: $\text{Upsample} + \text{Concat} + \text{ReduceChannel} \xrightarrow{6 \times \text{TransformerBlock}}$ (44)

1270 Level 1: $\text{Upsample} + \text{Concat} \xrightarrow{4 \times \text{TransformerBlock}}$ (45)

1271 $\xrightarrow{4 \times \text{TransformerBlock}} \text{Conv2d}(96 \rightarrow 3) + \text{Residual}$ (46)

1272 **A.7.4 AUXILIARY CONSTRAINTS**

1273 **Anisotropic Diffusion** The anisotropic diffusion operator preserves edges while smoothing noise:

1274
$$\mathcal{A}(I) = \nabla \cdot (c(|\nabla I|) \nabla I)$$

1275 with diffusion coefficient:

1276
$$c(|\nabla I|) = \exp \left(-\frac{|\nabla I|^2}{s^2} \right)$$

1296 where $s \in [0.01, 1.0]$ is a learnable sensitivity parameter.
 1297

1298 **Texture Consistency Loss:**

1299
$$L_{tex} = \|\mathcal{A}(I_{input}) - \mathcal{A}(R_{pred})\|_1$$

 1300

1301 **Illumination Smoothness Loss:**

1302
$$L_{lum} = \sum_{i,j} w_{i,j} (|\nabla_x L_{i,j}|^2 + |\nabla_y L_{i,j}|^2)$$

 1303
 1304

1305 where $w_{i,j} = \exp(-|\nabla L_{i,j}|)$ provides gradient-aware weighting.
 1306

1307 **Polarized HVI Color Space** The polarized HVI transformation eliminates red discontinuity:
 1308

1309
$$H_{polar} = C_k \cdot S \cdot \cos(\pi H/3) \quad (47)$$

 1310

1311
$$V_{polar} = C_k \cdot S \cdot \sin(\pi H/3) \quad (48)$$

 1312

1313
$$I_{polar} = I_{max} = \max(R, G, B) \quad (49)$$

 1314

1315 where the adaptive intensity collapse factor is:
 1316

1317
$$C_k = k \cdot \sin(\pi I_{max}/2) + \epsilon$$

 1318

1319 with learnable parameter $k \in [0.1, 5.0]$.
 1320

Polarized Color Loss:

1321
$$L_{col} = \|H_{polar}^{pred} - H_{polar}^{gt}\|_1 + \|V_{polar}^{pred} - V_{polar}^{gt}\|_1 + \|I_{polar}^{pred} - I_{polar}^{gt}\|_1$$

 1322

1323 **A.7.5 TEACHER TRAINING OBJECTIVE**
 1324

1325 The complete teacher training loss combines:
 1326

1327
$$L_{teach} = L_{rec} + L_{vgg} + L_{sty} + \lambda_{tex} L_{tex} + \lambda_{col} L_{col} + \lambda_{lum} L_{lum}$$

 1328

1329 where:

1330
$$L_{rec} = \|I_{pred} - I_{gt}\|_1 \quad (\text{pixel loss}) \quad (50)$$

 1331

1332
$$L_{vgg} = \sum_l \lambda_l \|\phi_l(I_{pred}) - \phi_l(I_{gt})\|_2^2 \quad (\text{perceptual loss}) \quad (51)$$

 1333

1334
$$L_{sty} = \sum_l \|G_l(\phi_l(I_{pred})) - G_l(\phi_l(I_{gt}))\|_F^2 \quad (\text{style loss}) \quad (52)$$

 1335

1337 with $\lambda_{tex} = 0.05$, $\lambda_{col} = 0.05$, $\lambda_{lum} = 0.2$.
 1338

1339 **A.8 STAGE 2: STUDENT NETWORK ARCHITECTURE**
 1340

1341 **A.8.1 RECTIFIED FLOW FORMULATION**
 1342

1343 Rectified flow models feature synthesis through straight-line interpolation:
 1344

1345
$$\mathbf{x}_t = (1 - t)\mathbf{z} + t\mathbf{f}_{teach}, \quad t \in [0, 1]$$

 1346

1347 where $\mathbf{z} \sim \mathcal{N}(0, I)$ is noise and \mathbf{f}_{teach} are teacher features.
 1348

Velocity Field:

1349
$$\mathbf{v}(\mathbf{x}_t, t) = \frac{d\mathbf{x}_t}{dt} = \mathbf{f}_{teach} - \mathbf{z}$$

1350 A.8.2 VELOCITY PREDICTION NETWORKS
13511352 **Architecture for both ϵ_θ^{rex} and ϵ_θ^{img} :**

1353
$$\text{VelocityPredictor}(\mathbf{x}_t, t, \mathbf{c}) = \text{ResMLP}^5(\text{LeakyReLU}_{0.1}(\text{Linear}_{513 \rightarrow 256}([\mathbf{c}; t_{norm}; \mathbf{x}_t]))) \quad (53)$$

1354
1355
1356

1357 where the input is $[\mathbf{c}; t; \mathbf{x}_t] \in \mathbb{R}^{513}$ with time normalization $t_{norm} = t/t_{max}$.
13581359 **Velocity Matching Loss:**

1360
$$L_{vel} = \mathbb{E}_{t, \mathbf{z}, \mathbf{f}_{teach}} [\|\epsilon_\theta(\mathbf{x}_t, t, \mathbf{c}) - \mathbf{v}(\mathbf{x}_t, t)\|_2^2] \quad (54)$$

1361

1362 A.8.3 ODE INTEGRATION FOR INFERENCE
13631364 During inference, the ODE is solved using Euler's method:
1365

1366
$$\mathbf{x}_{t+\Delta t} = \mathbf{x}_t + \Delta t \cdot \epsilon_\theta(\mathbf{x}_t, t, \mathbf{c})$$

1367

1368 with adaptive step sizing $\Delta t = 1.0/N_{steps}$ for $N_{steps} \in [1, 5]$.
1369

1370 A.8.4 FLEX KNOWLEDGE DISTILLATION LOSS

1371 **Cross-Normalization** FLEX uses student statistics for normalizing both teacher and student fea-
1372 tures at each layer l :
1373

1374
$$\mu_{stud}^l = \frac{1}{H_l W_l} \sum_{h,w} \mathbf{f}_{stud}^{l,h,w} \quad (54)$$

1375

1376
$$\sigma_{stud}^l = \sqrt{\frac{1}{H_l W_l} \sum_{h,w} (\mathbf{f}_{stud}^{l,h,w} - \mu_{stud}^l)^2 + \epsilon} \quad (55)$$

1377

1378
$$\mathbf{f}_{teach}^{l,norm} = \frac{\mathbf{f}_{teach}^l - \mu_{stud}^l}{\sigma_{stud}^l} \quad (56)$$

1379

1380
$$\mathbf{f}_{stud}^{l,norm} = \frac{\mathbf{f}_{stud}^l - \mu_{stud}^l}{\sigma_{stud}^l} \quad (57)$$

1381

1382 **Percentile-Based Outlier Detection** For each layer l and channel c , we compute:
1383

1384
$$\tau_p^{l,c} = \text{Percentile}(|\mathbf{f}_{stud}^{l,c,norm}|, p) \quad (58)$$

1385

1386
$$M_{reliable}^{l,c,h,w} = \mathbb{I}[|\mathbf{f}_{stud}^{l,c,norm,h,w}| \leq \tau_p^{l,c}] \quad (59)$$

1387

1388 where $p = 95\%$ is the outlier percentile threshold.
13891390 **Resolution-Aware Weighting** Dynamic resolution weighting prevents high-resolution features
1391 from dominating:
1392

1393
$$w_l^{res} = \max \left(\left(\frac{H_{base} W_{base}}{H_l W_l} \right)^{0.25}, 0.1 \right) \quad (60)$$

1394

1395 where $(H_{base}, W_{base}) = (64, 64)$ and (H_l, W_l) is the spatial resolution at layer l .
13961397 **Complete FLEX Loss** The final FLEX loss combines masked feature matching with dual weight-
1398 ing:
1399

1400
$$L_{FLEX} = \sum_l w_l^{layer} \cdot w_l^{res} \cdot \frac{\sum_{c,h,w} M_{reliable}^{l,c,h,w} \cdot \|\mathbf{f}_{teach}^{l,c,norm,h,w} - \mathbf{f}_{stud}^{l,c,norm,h,w}\|^2}{\sum_{c,h,w} M_{reliable}^{l,c,h,w} + \epsilon} \quad (61)$$

1401

1402 where w_l^{layer} are predefined layer importance weights and the denominator normalizes by the num-
1403 ber of reliable (non-outlier) elements.
1404

1404 A.8.5 TRAJECTORY CONSISTENCY REGULARIZATION
14051406 **Smooth Transitions:**

1407
$$L_{trans} = \sum_{i=1}^{N-1} \|\mathbf{f}_{pred}^{i+1} - \mathbf{f}_{pred}^i\|_2^2$$

1408
1409

1410 **Target Alignment:**

1411
$$L_{target} = \|\mathbf{f}_{pred}^{final} - \mathbf{f}_{teach}\|_2^2$$

1412

1413 **Semantic Consistency:**

1414
$$L_{cons} = \sum_{i=1}^N \text{cos_dist}(\mathbf{f}_{pred}^i, \mathbf{f}_{teach})$$

1415
1416

1417 **Complete Trajectory Loss:**

1418
$$L_{traj} = \alpha_{trans} L_{trans} + \alpha_{target} L_{target} + \alpha_{cons} L_{cons}$$

1419

1420 with $\alpha_{trans} = 0.1$, $\alpha_{target} = 0.5$, $\alpha_{cons} = 0.2$.
14211422 A.8.6 TWO-PHASE TRAINING PROTOCOL
14231424 **Phase 1: Velocity Learning**

1425
$$L_{phase1} = L_{vel}^{rex} + L_{vel}^{img} + \lambda_{KD} L_{KD} + \lambda_{traj} L_{traj}$$

1426

1427 **Phase 2: Full Network Training**

1428
$$L_{phase2} = L_{rec} + \lambda_{FLEX} L_{FLEX} + \lambda_{vel} (L_{vel}^{rex} + L_{vel}^{img})$$

1429

1430 with $\lambda_{FLEX} = 0.15$, $\lambda_{vel} = 0.05$.
14311432 A.9 IMPLEMENTATION DETAILS
14331434 A.9.1 NETWORK DIMENSIONS AND PARAMETERS
14351436 **Stage 1 (Teacher):**1437

- RGFormer dimensions: $\text{dim} = 48$
- Multi-head attention heads: $[1, 2, 4, 8]$
- Transformer blocks per level: $[4, 6, 6, 8]$
- FFN expansion factor: 2.66

14381439 **Stage 2 (Student):**
14401441

- Velocity predictor features: 256
- Rectified flow timesteps: 4
- ODE integration steps: 1 – 5

14421443 A.9.2 TRAINING HYPERPARAMETERS
14441445 **Stage 1:**
14461447

- Learning rate: 2×10^{-4}
- Batch size: 16
- Training iterations: $500k$

14481449 **Stage 2:**
1450

1458 • Phase 1 learning rates: $lr_{rex} = lr_{img} = 2 \times 10^{-4}$
1459 • Phase 2 learning rate: 1×10^{-4}
1460 • Phase 1 iterations: $50k$
1461 • Phase 2 iterations: $200k$
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511