ADAPTIVE MULTI-EXPOSURE REGULARIZATION FOR DYNAMIC HDR SCENE RECONSTRUCTION

Anonymous authors
Paper under double-blind review

ABSTRACT

Dynamic HDR reconstruction with Gaussian Splatting relies on neural tone mappers to model the nonlinear relationship between reconstructed HDR radiance and multi-exposure observations. Under varying exposure conditions, the pre-activation distribution within the tone mapper can shift substantially, altering its activation behavior and limiting the nonlinear responses effectively involved in reconstruction. To address this issue, we propose Adaptive Multi-Exposure Regularization (AME-REG), which calibrates the tone mapper's pre-activation distribution across exposure conditions using exposure-conditioned residual offsets. Extensive experiments across multiple dynamic HDR Gaussian Splatting pipelines demonstrate consistent improvements in HDR reconstruction quality, supporting the general applicability of our method.

1 INTRODUCTION

High dynamic range (HDR) reconstruction aims to recover scene radiance beyond the limited dynamic range of conventional image sensors. Recent Gaussian Splatting (GS)-based (Kerbl et al., 2023) HDR methods typically reconstruct an exposure-independent HDR scene representation and employ a learnable tone mapper to relate it to exposure-dependent LDR observations (Cai et al., 2024; Wu et al., 2024c; Zhang et al., 2025; Liu et al., 2026). This formulation has also been extended to dynamic scenes (Zhang et al., 2025; Liu et al., 2026), where temporal scene variation and multi-exposure observations need to be modeled jointly. During training, LDR observations acquired under different exposures jointly supervise the same underlying HDR scene representation through the tone mapper. Therefore, the tone mapper not only defines the forward radiometric transformation from HDR radiance to LDR intensity, but also affects how reconstruction supervision optimizes the HDR representation.

Existing methods have mainly improved the tone mapper through log-domain response modeling (Cai et al., 2024), adaptive or local tone mapping (Wu et al., 2024c; Liu et al., 2025), and temporal conditioning or additional radiometric constraints (Zhang et al., 2025; Liu et al., 2026). These efforts mainly focus on how the radiometric mapping is parameterized or conditioned. A comparatively less explored aspect is how exposure variation changes the operating region of its internal nonlinear stage. Under the commonly adopted log-domain formulation, changing the exposure introduces an additive shift to the mapper input, which further shifts the corresponding pre-activation distribution. As a result, different exposure brackets can place the same nonlinear mapping in different response regions and lead its units to exhibit substantially different activation states. Such exposure-dependent operating behavior can affect the nonlinear responses available during reconstruction, yet has received relatively limited attention in existing dynamic HDR methods.

To characterize this activation behavior, we analyze the tone-mapper architecture used in HDR-GS (Cai et al., 2024) as a representative case. As shown in Fig. 1, a large fraction of its hidden units exhibit activation frequencies close to either zero or one, indicating that they remain nearly always inactive or active over the observed samples. In the former case, a unit contributes little within the observed pre-activation range, whereas in the latter it behaves largely as a fixed affine component. Consequently, only a limited portion of the available nonlinear responses is effectively exercised by the data. We further observe that this concentrated activation behavior is accompanied by degraded HDR reconstruction across different scenes. These observations motivate explicit control of the pre-activation distribution to avoid excessively concentrated operating states.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

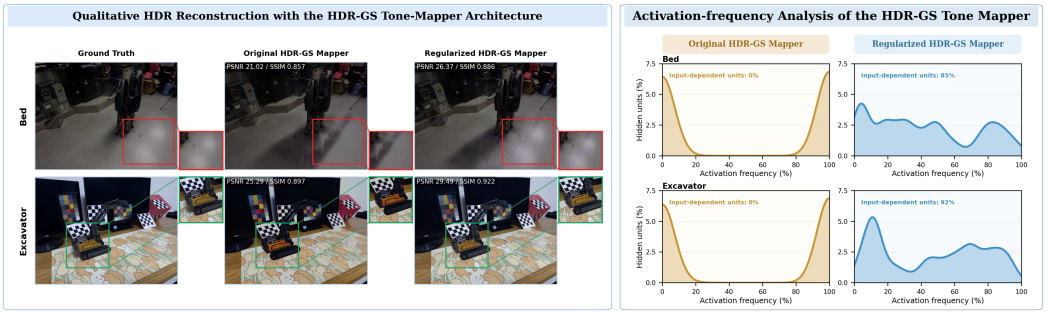

Figure 1: Qualitative HDR reconstruction and activation-frequency analysis of the HDR-GS tone mapper before and after regularization. Left: regularization consistently improves reconstruction quality across different dynamic scenes. Right: the original mapper exhibits activation frequencies strongly concentrated near the two extreme regimes, whereas regularization shifts substantially more units toward intermediate, input-dependent activation frequencies.

Motivated by these observations, we propose Adaptive Multi-Exposure Regularization (AME-REG), an exposure-aware regularization strategy for dynamic HDR reconstruction. AME-REG calibrates the pre-activation distribution of the nonlinear tone-mapping stage across different exposures. Specifically, we introduce exposure-conditioned residual offsets together with a quantile-based regularization loss to calibrate the pre-activation distribution of each exposure bracket within a suitable interval. This explicit calibration encourages the nonlinear stage to retain input-dependent responses across exposure conditions, allowing a broader set of nonlinear responses to participate in HDR reconstruction. As shown in Fig. 1, AME-REG produces a broader activation-frequency distribution and improves HDR reconstruction quality. The resulting regularization can be readily integrated into different dynamic HDR reconstruction frameworks.

Our contributions are summarized as follows:

- We revisit the role of the tone mapper in dynamic HDR reconstruction by analyzing its exposure-dependent behavior, and find that exposure variation shifts the pre-activation distribution of its nonlinear stage, altering its activation behavior and the nonlinear responses through which LDR supervision constrains the HDR scene representation.
- Based on this analysis, we propose AME-REG, an exposure-aware regularization strategy that calibrates the tone mapper's pre-activation distribution across exposure conditions through exposure-conditioned residual offsets and two simple L2 regularization terms.
- Extensive experiments on real and synthetic dynamic HDR benchmarks across multiple frameworks show consistent gains with AME-REG.

2 RELATED WORK

HDR reconstruction and tone-mapping design. HDR reconstruction from multi-exposure LDR observations generally separates an exposure-independent HDR scene representation from the exposure-dependent radiometric mapping to observed LDR images (Debevec & Malik, 1997; Mitsunaga & Nayar, 1999; Hu et al., 2013; Yan et al., 2019; Sen et al., 2012). HDR-NeRF jointly models HDR radiance and the camera response function using a neural tone mapper, while HDR-Plenoxels further incorporates self-calibration of radiometric parameters such as exposure, white balance, and camera response (Huang et al., 2022; Kim et al., 2022; Mildenhall et al., 2021). With Gaussian Splatting, HDR-GS introduces HDR Gaussian primitives together with an MLP-based tone mapper, HDRGS adopts an asymmetric grid-based mapping, and GaussHDR further combines global radiometric mapping with local 3D and 2D tone mapping (Cai et al., 2024; Wu et al., 2024c; Liu et al., 2025). These methods establish the common HDR reconstruction paradigm in which a tone-mapping component bridges reconstructed HDR radiance and multi-exposure LDR supervision.

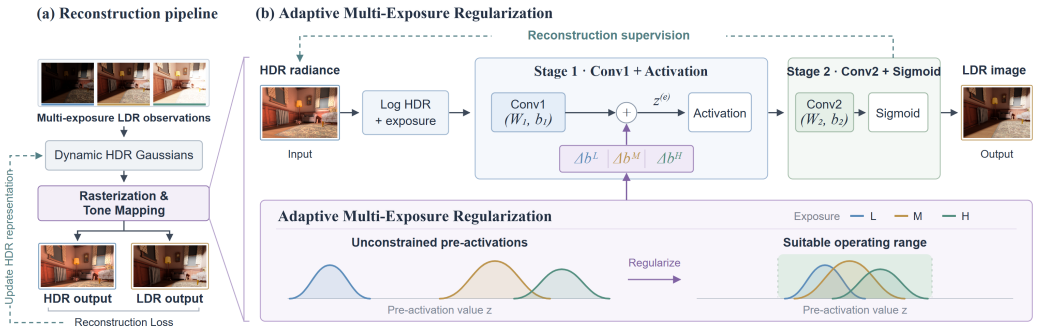

Figure 2: Overview of dynamic HDR reconstruction with the proposed AME-REG. (a) The overall reconstruction pipeline takes alternating-exposure LDR observations with fixed exposure metadata and optimizes a dynamic HDR Gaussian representation through joint HDR–LDR reconstruction. (b) We functionally decompose the adopted tone mapper into nonlinear response-basis construction and response aggregation and projection stages for analysis. Our regularization acts on the Stage-I pre-activations through bracket-conditioned calibration offsets and constrains them to a feasible interval derived from the pre-activation distribution, encouraging the nonlinear mapping to operate within an appropriate range across exposure brackets.

Dynamic Gaussian scene representations. Recent extensions of 3D Gaussian Splatting to dynamic scenes mainly differ in how temporal variation is represented (Pumarola et al., 2020; Li et al., 2021; Fridovich-Keil et al., 2023). Deformation-based methods, such as 4DGS and Deformable 3D Gaussians, maintain a canonical Gaussian representation and use learned deformation fields to predict time-dependent changes of Gaussian attributes (Wu et al., 2024b; Yang et al., 2024b). Trajectory-based methods instead explicitly parameterize Gaussian motion over time; SplineGS, for example, represents Gaussian motion with smooth spline trajectories (Park et al., 2025). A complementary line directly extends Gaussian primitives to the spatiotemporal domain (Li et al., 2024), as in 4D Gaussian Splatting, where temporal properties are incorporated into the Gaussian representation itself (Yang et al., 2024a). These dynamic Gaussian representations provide the scene-modeling foundations for recent dynamic HDR systems, including HDR-4DGS and Mono4DGS-HDR, which further incorporate multi- or alternating-exposure observations for HDR reconstruction (Zhang et al., 2025; Liu et al., 2026).

Tone-mapping and radiometric regularization. Beyond tone-mapper architecture, existing HDR reconstruction methods introduce different constraints to stabilize radiometric learning under varying exposures (Grossberg & Nayar, 2003; Martin-Brualla et al., 2021). HDR-HexPlane constrains radiometric estimation through learned exposure coefficients and a fixed camera response prior, whereas Mono4DGS-HDR introduces temporal luminance regularization to improve the temporal consistency of reconstructed HDR appearance (Wu et al., 2024a; Liu et al., 2026). Related dynamic HDR formulations also regularize the radiometric mapping itself; HDR-NSFF, for example, imposes a smoothness constraint on the camera response function (Shin et al., 2026). These approaches primarily constrain exposure estimation, response-function shape, or temporal output consistency. In contrast, the exposure-conditioned internal operating range of the nonlinear tone mapper has received comparatively less attention. Our work addresses this complementary aspect by regularizing its pre-activation distribution toward a suitable operating interval across exposure conditions.

3 METHOD

Our goal is to improve dynamic GS-based HDR reconstruction by regularizing the operating range of the exposure-conditioned tone mapper. The overall pipeline is illustrated in Fig. 2. Following prior HDR reconstruction methods, we employ an image-space exposure-conditioned tone mapper to map the rendered HDR radiance to LDR observations. We first introduce the commonly used structure and a functional decomposition of the tone mapper (Sec. 3.1). We then analyze how alternating ex-

posure brackets affect the pre-activation operating range of its nonlinear mapping (Sec. 3.2). Based on this analysis, we introduce Adaptive Multi-Exposure Regularization (AME-REG) that keeps the tone mapper within an appropriate operating range across exposure brackets (Sec. 3.3).

3.1 EXPOSURE-CONDITIONED TONE MAPPING FOR HDR RECONSTRUCTION

Existing HDR reconstruction methods employ a learnable tone mapper to model the camera response function (CRF), mapping reconstructed HDR radiance to the corresponding LDR observation under a given exposure (Huang et al., 2022; Cai et al., 2024; Liu et al., 2025). Following the standard log-domain formulation (Debevec & Malik, 1997), given a rendered linear HDR image $H_p \in \mathbb{R}^3_+$ and the recorded exposure E_k of exposure bracket k, the tone-mapper input is written as

$$x_{c,k,p} = \log(H_{c,p} + \epsilon) + \log(E_k + \epsilon), \tag{1}$$

where c, k, and p index the color channel, exposure bracket, and pixel, respectively. The tone mapper then transforms the exposure-conditioned signal $x_{c,k,p}$ into the observable LDR intensity.

Prior HDR reconstruction methods employ MLP-based tone mappers, with separate mappings for the RGB channels (Huang et al., 2022; Cai et al., 2024; Liu et al., 2025). The tone mapper adopted in our experiments is the same as the existing structure: using one shallow MLP for each color channel, consisting of a hidden affine transformation followed by ReLU and a subsequent linear projection followed by Sigmoid. We use it as the representative tone mapper for the following analysis. Its computation is written as

$$a_{c,j,k,p} = w_{c,j} x_{c,k,p} + b_{c,j}, \qquad h_{c,j,k,p} = \phi(a_{c,j,k,p}), \tag{2}$$

followed by

$$y_{c,k,p} = \sum_j v_{c,j} h_{c,j,k,p} + \beta_c, \qquad \hat{I}_{c,k,p} = \sigma(y_{c,k,p}), \tag{3}$$

where $a_{c,j,k,p}$ denotes the pre-activation of hidden unit j, i.e., the input to the nonlinear activation $\phi(\cdot)$, and $h_{c,j,k,p} = \phi(a_{c,j,k,p})$ denotes the corresponding hidden response, and $v_{c,j}$ and β_c denote the output-layer weight and bias, respectively. For this mapper, $\phi(\cdot)$ and $\sigma(\cdot)$ are ReLU and Sigmoid, respectively.

The above computation naturally contains two functional stages. For the analysis below, we refer to the hidden transformation from $x_{c,k,p}$ to $h_{c,j,k,p}$ as *Stage I: nonlinear response-basis construction*, and the subsequent aggregation and output projection as *Stage II: response aggregation and projection*. This two-stage description is introduced to characterize the internal behavior of the existing tone mapper and does not modify its architecture. More elaborate HDR tone mappers may employ deeper nonlinear stacks, additional conditioning variables, or local tone-mapping components, but their computation can still be interpreted as first constructing exposure-dependent nonlinear responses and then aggregating or projecting them into the final LDR prediction. Our subsequent analysis focuses on the Stage-I pre-activation $a_{c,j,k,p}$ and how its operating range changes across exposure brackets. Architectural details of the adopted tone mapper are provided in Appendix A.1.

3.2 EXPOSURE-DEPENDENT OPERATING RANGE

Following the functional decomposition in Sec. 3.1, we next analyze how exposure variation affects the operating range of the nonlinear stage in the adopted tone mapper. Our analysis focuses on the Stage-I pre-activation $a_{c,j,k,p}$, whose distribution determines which region of the nonlinear mapping is covered by the observed samples. Although the tone-mapper parameters are shared across exposure brackets, the distribution of $a_{c,j,k,p}$ varies with exposure, leading the same nonlinear mapping to operate over different nonlinear operating regions.

Empirical observation. We first examine the operating behavior of the adopted tone mapper under the original alternating-exposure observations. As shown in Fig. 1, a large fraction of hidden units exhibit activation frequencies close to either zero or one before regularization. For the ReLU activation used in the representative mapper, these extreme frequencies indicate that the corresponding pre-activations remain predominantly on one side of the activation transition, such that the associated responses are either largely suppressed or dominated by a single linear regime. The activation

Table 1: Mean reconstruction quality over the four real dynamic HDR scenes. *Both HDR-HexPlane variants use the recorded exposure values instead of the original per-image exposure learning; †for GaussHDR, AME-REG replaces the original tone-mapping-specific supervision while the remaining reconstruction components are retained; ‡the HDR-4DGS baseline metrics on the real scenes are taken from the original paper (Zhang et al., 2025). See App. A for details.

Method	HDR			LDR		
	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓
HDR-GS	28.80	0.930	0.113	20.74	0.823	0.202
HDR-GS + AME-REG	**30.37**	**0.938**	0.147	**21.07**	**0.825**	0.238
GaussHDR	29.15	0.946	0.076	21.95	0.869	0.128
GaussHDR† + AME-REG	**31.09**	**0.947**	**0.075**	21.87	0.863	0.133
HDR-HexPlane	24.87	0.885	0.371	27.00	0.739	0.484
HDR-HexPlane* + AME-REG	**25.84**	**0.890**	**0.343**	**27.64**	**0.752**	**0.450**
HDR-4DGS	25.13	0.909	0.143	30.69	0.927	0.159
HDR-4DGS + AME-REG	**32.28**	**0.960**	**0.070**	**31.75**	0.900	**0.101**
HDR-SplineGS	24.23	0.947	0.065	22.98	0.856	0.115
HDR-SplineGS + AME-REG	**33.51**	**0.953**	0.065	**27.40**	**0.887**	**0.076**



Figure 3: Qualitative comparison on two representative scenes with the HDR-4DGS mapper.

statistics therefore provide an observable characterization of how the Stage-I pre-activation distribution occupies the nonlinear operating region. This observation motivates us to further examine the distribution itself and its dependence on exposure.

Exposure-dependent input shift. The dependence of the pre-activation distribution on exposure follows directly from the tone-mapper input formulation. Under Eq. 1, changing the exposure by a multiplicative factor $s > 0$ translates the tone-mapper input as

$$x' = x + \log s. \tag{4}$$

Since Stage I begins with an affine transformation, the corresponding pre-activation is shifted accordingly:

$$a'_{c,j,k,p} = a_{c,j,k,p} + w_{c,j} \log s. \tag{5}$$

Thus, even with unchanged tone-mapper parameters, different exposure brackets systematically shift the Stage-I pre-activation distribution. Consequently, the same nonlinear mapping can be evaluated over different pre-activation regions across exposure conditions.

Effect on nonlinear representation. The nonlinear responses produced by Stage I depend on the region of the response function covered by the pre-activation distribution. When the distribution is concentrated within a restricted region, the corresponding hidden responses exhibit limited input-dependent variation over the observed samples. Exposure-dependent shifts can therefore cause different brackets to utilize different portions of the shared nonlinear mapping, limiting the diversity of nonlinear responses involved in tone mapping.

Together, the empirical activation behavior and the exposure-induced shift reveal that the operating range of the Stage-I pre-activation distribution is closely related to the utilization of the nonlinear tone mapper. This observation motivates an exposure-aware calibration of the pre-activation distribution, where each exposure bracket is adaptively adjusted to maintain a suitable nonlinear operating range. Building on this insight, we introduce AME-REG, an adaptive multi-exposure regularization strategy that explicitly constrains the Stage-I pre-activation range across exposure brackets, as described in Sec. 3.3.

3.3 AME-REG: ADAPTIVE MULTI-EXPOSURE REGULARIZATION

Building on the operating-range analysis above, we introduce AME-REG to explicitly regulate the Stage-I pre-activation distribution of the tone mapper across exposure brackets. AME-REG combines bracket-conditioned calibration offsets with a quantile-derived feasible range, allowing each exposure-conditioned distribution to be adjusted according to its empirical operating range.

Bracket-conditioned range calibration. For each color channel c, exposure bracket k, and hidden unit j, AME-REG introduces a learnable calibration offset $\Delta b_{c,k,j}$ before the Stage-I nonlinear activation:

$$z_{c,j,k,p} = a_{c,j,k,p} + \Delta b_{c,k,j}, \tag{6}$$

with the calibrated hidden response given by

$$\tilde{h}_{c,j,k,p} = \phi\left(z_{c,j,k,p}\right). \tag{7}$$

The same offset is applied to all sampled pixels associated with a given (c, k, j), translating the corresponding pre-activation distribution while preserving its within-bracket structure. Separate offsets are maintained for the low-, middle-, and high-exposure brackets, enabling bracket-specific calibration of the nonlinear operating range. All calibration offsets are initialized to zero, preserving the original tone-mapper behavior at the beginning of training, and optimized by the range regularization during training.

Quantile-derived feasible range. The calibration offset is constrained to maintain the pre-activation distribution within a suitable nonlinear operating region. For the representative ReLU mapper, zero is the transition point of the hidden nonlinear activation and is therefore used as the reference point for defining the operating range.

Let $Q_\alpha^{(p)}(a_{c,j,k,p})$ denote the empirical α-quantile of the Stage-I pre-activations over the sampled pixels p. Given $q \in (0, 0.5)$, we require the calibrated distribution to retain at least a fraction q of samples on each side of the nonlinear transition. We use $q = 0.05$ by default and study its sensitivity in Sec. 4.3. The positive-side condition gives

$$\Delta b_{c,k,j} \geq -Q_{1-q}^{(p)}(a_{c,j,k,p}), \tag{8}$$

while retaining sufficient support on the negative side gives

$$\Delta b_{c,k,j} \leq -Q_q^{(p)}(a_{c,j,k,p}). \tag{9}$$

Combining the two conditions yields the feasible interval

$$\ell_{c,k,j} \leq \Delta b_{c,k,j} \leq u_{c,k,j}, \qquad \ell_{c,k,j} = -Q_{1-q}^{(p)}(a_{c,j,k,p}), \quad u_{c,k,j} = -Q_q^{(p)}(a_{c,j,k,p}). \tag{10}$$

Because the bounds are derived from the empirical pre-activation distribution of each exposure bracket, the feasible interval adapts to its bracket-specific operating range. Offsets within this interval incur no regularization penalty, while offsets outside the interval are calibrated by the range regularization.

Range regularization. Given the quantile-derived feasible interval, we implement the range constraint using two one-sided squared L2 penalties:

$$\mathcal{L}_{\text{range}} = \|[\ell - \Delta b]_+\|_2^2 + \|[\Delta b - u]_+\|_2^2, \tag{11}$$

where $[x]_+ = \max(x, 0)$. The first term penalizes calibration offsets below the lower bound, whereas the second penalizes offsets above the upper bound. The regularization is inactive when Δb lies inside the feasible interval and acts only when the calibration offset violates the prescribed range. The loss is averaged over the relevant exposure-bracket, color-channel, and hidden-unit dimensions.

Table 2: Mean reconstruction quality over the eight synthetic dynamic HDR scenes.

Method	HDR			LDR		
	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓
HDR-GS	26.91	0.889	0.157	27.17	0.897	0.157
HDR-GS + AME-REG	**27.21**	**0.894**	**0.139**	**27.67**	**0.905**	**0.139**
GaussHDR	30.02	0.949	0.046	30.62	0.955	0.055
GaussHDR† + AME-REG	**31.01**	**0.957**	0.046	**30.70**	0.952	0.058
HDR-HexPlane	23.24	0.752	0.343	28.39	0.836	0.254
HDR-HexPlane* + AME-REG	**24.28**	**0.756**	**0.334**	**29.54**	**0.856**	**0.230**
HDR-4DGS	25.89	0.859	0.103	30.96	0.914	0.098
HDR-4DGS + AME-REG	**28.80**	**0.919**	**0.076**	**33.01**	**0.957**	**0.070**
HDR-SplineGS	28.08	0.859	0.105	29.35	0.870	0.082
HDR-SplineGS + AME-REG	**29.05**	**0.884**	**0.080**	**30.33**	**0.894**	**0.062**

Training objective. The reconstruction and range objectives are optimized together as

$$\mathcal{L} = \mathcal{L}_{\text{recon}} + \lambda_{\text{range}} \mathcal{L}_{\text{range}}, \tag{12}$$

where λ_{range} controls the strength of AME-REG. We isolate the gradient flows of the two objectives: the statistics used to construct the feasible interval are detached when optimizing $\mathcal{L}_{\text{range}}$, while the calibration offsets are detached from $\mathcal{L}_{\text{recon}}$. Consequently, the offsets are updated only by the range regularization rather than being directly fitted by the reconstruction objective. Their calibrated values nevertheless determine the nonlinear operating range used in the reconstruction forward pass. Additional implementation details are provided in Appendix A.2.

4 EXPERIMENTS

We evaluate AME-REG on real and synthetic dynamic HDR scenes, analyze its effect on the exposure-conditioned pre-activation range, and assess its transferability across multiple HDR reconstruction backbones.

4.1 EXPERIMENTAL SETUP

Data. We evaluate on 12 dynamic HDR scenes: four real scenes (Bed, Tank, Excavator, and Toys) from HDR-4D-Real (Zhang et al., 2025), and eight synthetic scenes from HDR-4D-Syn, adapted from HDR-HexPlane (Wu et al., 2024a). Dataset details are provided in App. B.

Metrics. We report PSNR, SSIM (Wang et al., 2004), and LPIPS (Zhang et al., 2018) for HDR and LDR reconstruction. Following prior HDR protocols (Cai et al., 2024; Liu et al., 2025), HDR metrics are computed in the μ-law domain (Kalantari & Ramamoorthi, 2017; Huang et al., 2022; Cai et al., 2024), while qualitative results are visualized with Photomatix Pro (HDRsoft, 2026). HDR metrics are our primary target, whereas LDR metrics measure consistency with the exposure-conditioned observations.

Implementation details. For controlled analysis and ablations, we combine SplineGS (Park et al., 2025) with the exposure-conditioned tone mapper in Sec. 3.1, denoted HDR-SplineGS. We train on every other frame and evaluate on the remainder, yielding an approximately 1:1 train/test split. For AME-REG, we set $q = 0.05$, sample at most 4,096 pixels per iteration to estimate the Stage-I pre-activation distribution, and set $\lambda_{\text{range}} = 0.006$. All methods follow the HDR-4DGS image-downsampling setting (Zhang et al., 2025), with each baseline and its AME-REG variant evaluated under the same method-specific protocol. Additional implementation details are provided in App. A.

4.2 RECONSTRUCTION RESULTS

Quantitative comparison. Across the real and synthetic scenes (Tables 1 and 2), AME-REG improves HDR PSNR for every evaluated framework pair, with particularly large gains for HDR-4DGS



Figure 4: Qualitative comparison on two representative scenes with the HDR-GS mapper.

and HDR-SplineGS. The improvements are observed across different radiance representations and tone-mapping designs, including both GS-based and non-GS HDR reconstruction frameworks, indicating that the gain comes from regulating the nonlinear mapping behavior rather than modifying a specific reconstruction architecture. In contrast, LDR metrics remain broadly comparable, suggesting that the proposed regularization mainly improves the recovery of high-dynamic-range radiance information while preserving consistency with exposure-conditioned observations.

The gain is closely related to the operating-range constraint introduced by AME-REG. Dynamic HDR reconstruction requires a single tone mapper to explain observations from different exposure brackets, where large exposure-dependent shifts can cause pre-activations to accumulate in saturated regions and reduce the effective nonlinear response bases. By constraining the Stage-I pre-activation distribution within a feasible bracket-specific interval, AME-REG maintains more balanced responses across exposures and enables the tone mapper to better represent radiometric variations. This explains why frameworks with stronger exposure variation, such as HDR-4DGS and HDR-SplineGS, obtain larger improvements.

Qualitative comparison. We further compare HDR reconstructions with and without AME-REG on representative dynamic scenes (Figs. 3 and 4). The comparisons focus on challenging regions with large radiance variations, including highlights, shadows, and moving objects. AME-REG better preserves radiance structure and scene details in these regions, consistent with the quantitative improvements reported above. In particular, improvements are more evident in regions where different exposure brackets observe substantially different radiance levels, as these regions require a wider and more stable nonlinear operating range.

The activation-frequency distributions in Fig. 1 provide a complementary view of this behavior. After applying AME-REG, fewer tone-mapper units concentrate near extreme operating states, while more units contribute meaningful nonlinear responses. This supports our motivation that the improvement is not obtained by simply increasing the mapping capacity, but by maintaining effective utilization of the existing nonlinear response basis under exposure-dependent HDR reconstruction.

4.3 ABLATION STUDIES

We conduct controlled ablations on HDR-SplineGS to evaluate the regularization objective, alternative HDR objectives and nonlinear mappings, and the main design choices of AME-REG.

Regularization objective design. We compare AME-REG with fixed-50% L_1 and L_2 objectives that encourage an approximately balanced split of Stage-I pre-activations around the nonlinear transition. The 50% target provides a simple symmetric reference for testing whether a fixed operating target is sufficient. We also evaluate $q \in \{0.05, 0.10, 0.15\}$ under the same HDR-SplineGS training protocol.

The fixed-50% objectives already yield large gains over no regularization, confirming the importance of controlling the tone-mapper operating range. AME-REG further replaces the fixed target with a quantile-derived feasible interval. Across $q = 0.05$–0.15, SSIM remains unchanged while HDR PSNR increases slightly from 33.51 to 33.76 dB.

Comparison with existing HDR-specific objectives. We compare AME-REG with objectives adopted in existing HDR reconstruction methods, including HDR-NeRF, HDR-Plenoxels, and GaussHDR (Huang et al., 2022; Kim et al., 2022; Liu et al., 2025). These alternatives impose

Table 3: Ablations on HDR-SplineGS over the four real scenes. Top: comparisons with alternative objectives and nonlinear mappings. Bottom: design choices of AME-REG. Yellow highlights denote the adopted setting.

(a) Fixed 50%

Objective	PSNR ↑	SSIM ↑
No Reg.	24.23	0.947
Fixed-50% L_1	32.90	0.953
Fixed-50% L_2	33.32	**0.954**
AME-REG	**33.51**	0.953

(b) HDR objectives

Objective	PSNR ↑	SSIM ↑
HDR-NeRF	24.14	0.946
HDR-Plenoxels	24.38	0.947
GaussHDR	31.10	**0.956**
AME-REG	**33.51**	0.953

(c) Nonlinear mappings

Method	PSNR ↑	SSIM ↑
No Reg.	24.23	0.947
ReActNet	24.50	0.948
ACON-C	24.17	0.946
AME-REG	**33.51**	**0.953**

(d) Quantile q

q	PSNR ↑	SSIM ↑
0.05	33.51	**0.953**
0.10	33.69	**0.953**
0.15	**33.76**	**0.953**

(e) Placement

Placement	PSNR ↑	SSIM ↑
Stage I	**33.51**	**0.953**
Stage II	27.38	0.941
Both	27.71	0.913

(f) Exposure conditioning

Calibration	PSNR ↑	SSIM ↑
Global	23.96	0.951
Bracket	**33.51**	**0.953**

different forms of constraint or supervision on the radiometric mapping, providing complementary references for evaluating AME-REG under the same HDR-SplineGS setting.

Among these objectives, AME-REG achieves the highest HDR PSNR by directly regulating the exposure-conditioned Stage-I operating range (Table 3(b)).

Alternative nonlinear mappings. We further examine whether a more flexible nonlinear mapping can recover the reconstruction gain without explicit range regularization. Compared with standard ReLU, ReActNet and ACON-C (Liu et al., 2020; Ma et al., 2021) introduce more flexible activation behaviors that can potentially retain responses over a broader input range. We therefore replace the original ReLU with these nonlinear mappings while keeping the remaining HDR-SplineGS configuration unchanged.

ReActNet and ACON-C yield only minor changes over the unregularized setting, whereas AME-REG provides a substantially larger gain (Table 3(c)). This indicates that increasing the flexibility of the nonlinear mapping alone is insufficient to address the exposure-dependent operating-range issue.

Regularization placement. Applying the regularization to Stage I outperforms Stage II and both stages (Table 3(e)), supporting direct regulation of the Stage-I pre-activation distribution.

Exposure conditioning. We further compare a global calibration shared across all exposure brackets with the proposed bracket-conditioned calibration. The global variant uses one calibration offset per color channel and hidden unit, whereas AME-REG maintains separate offsets across exposure brackets. As shown in Table 3(f), bracket-conditioned calibration substantially improves HDR reconstruction, supporting our analysis that different exposure brackets induce different pre-activation shifts and benefit from bracket-specific calibration.

5 CONCLUSION

We identify inadequate control of the exposure-dependent operating range of the tone mapper as an important issue in dynamic GS-based HDR reconstruction. To address this issue, we introduce AME-REG, a pre-activation range regularization strategy that constrains the Stage-I pre-activation distribution within a quantile-derived feasible interval for each exposure bracket. This exposure-aware regularization promotes better utilization of the nonlinear operating range and improves the tone mapper's ability to model radiometric responses. Experiments across different dynamic HDR reconstruction frameworks and scenes demonstrate consistent improvements in HDR reconstruction quality, while controlled ablations validate that the proposed operating-range formulation is more effective than fixed targets or alternative nonlinear mappings.

AI USE STATEMENT

AI-assisted tools were used solely to improve the clarity, grammar, and presentation of the manuscript. They were not used to formulate the research ideas or methodology, conduct experiments, analyze data, or generate the scientific conclusions. The authors reviewed and take full responsibility for all content in the manuscript.

REFERENCES

Yuanhao Cai, Zihao Xiao, Yixun Liang, Minghan Qin, Yulun Zhang, Xiaokang Yang, Yaoyao Liu, and Alan Yuille. HDR-GS: Efficient high dynamic range novel view synthesis at 1000x speed via gaussian splatting. *In NeurIPS,2024*, 2024.

Paul E. Debevec and Jitendra Malik. Recovering high dynamic range radiance maps from photographs. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques*, pp. 369–378, 1997.

Sara Fridovich-Keil, Giacomo Meanti, Frederik Warburg, Benjamin Recht, and Angjoo Kanazawa. K-planes: Explicit radiance fields in space, time, and appearance, 2023. URL https://arxiv.org/abs/2301.10241.

M.D. Grossberg and S.K. Nayar. What is the space of camera response functions? In *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.*, volume 2, pp. II–602, 2003. doi: 10.1109/CVPR.2003.1211522.

HDRsoft. Photomatix Pro, 2026. URL https://www.hdrsoft.com/.

Jun Hu, Orazio Gallo, Kari Pulli, and Xiaobai Sun. Hdr deghosting: How to deal with saturation? In *2013 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1163–1170, 2013. doi: 10.1109/CVPR.2013.154.

Xin Huang, Qi Zhang, Ying Feng, Hongdong Li, Xuan Wang, and Qing Wang. HDR-NeRF: High dynamic range neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 18398–18408, 2022.

Nima Khademi Kalantari and Ravi Ramamoorthi. Deep high dynamic range imaging of dynamic scenes. *ACM Transactions on Graphics*, 36(4):144, 2017.

Nikita Karaev, Ignacio Rocco, Benjamin Graham, Natalia Neverova, Andrea Vedaldi, and Christian Rupprecht. CoTracker: It is better to track together. *arXiv preprint arXiv:2307.07635*, 2023.

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3D gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), 2023.

Jun-Seong Kim, Yu-Ji Kim, Ye-Bin Moon, and Tae-Hyun Oh. HDR-Plenoxels: Self-calibrating high dynamic range radiance fields. In *European Conference on Computer Vision*, pp. 105–121, 2022.

Zhan Li, Zhang Chen, Zhong Li, and Yi Xu. Spacetime gaussian feature splatting for real-time dynamic view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024.

Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes, 2021. URL https://arxiv.org/abs/2011.13084.

Jinfeng Liu, Lingtong Kong, Bo Li, and Dan Xu. GaussHDR: High dynamic range gaussian splatting via learning unified 3D and 2D local tone mapping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5991–6000, 2025.

Jinfeng Liu, Lingtong Kong, Mi Zhou, Jinwei Chen, and Dan Xu. Mono4DGS-HDR: High dynamic range 4D gaussian splatting from alternating-exposure monocular videos. In *International Conference on Learning Representations*, 2026.

Zechun Liu, Zhiqiang Shen, Marios Savvides, and Kwang-Ting Cheng. ReActNet: Towards precise binary neural network with generalized activation functions. In *European Conference on Computer Vision*, 2020.

Ningning Ma, Xiangyu Zhang, Ming Liu, and Jian Sun. Activate or not: Learning customized activation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021.

Ricardo Martin-Brualla, Noha Radwan, Mehdi S. M. Sajjadi, Jonathan T. Barron, Alexey Dosovitskiy, and Daniel Duckworth. Nerf in the wild: Neural radiance fields for unconstrained photo collections, 2021. URL https://arxiv.org/abs/2008.02268.

Ben Mildenhall, Peter Hedman, Ricardo Martin-Brualla, Pratul Srinivasan, and Jonathan T. Barron. Nerf in the dark: High dynamic range view synthesis from noisy raw images, 2021. URL https://arxiv.org/abs/2111.13679.

T. Mitsunaga and S.K. Nayar. Radiometric self calibration. In *Proceedings. 1999 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (Cat. No PR00149)*, volume 1, pp. 374–380 Vol. 1, 1999. doi: 10.1109/CVPR.1999.786966.

Jongmin Park, Minh-Quan Viet Bui, Juan Luis Gonzalez Bello, Jaeho Moon, Jihyong Oh, and Munchurl Kim. SplineGS: Robust motion-adaptive spline for real-time dynamic 3D gaussians from monocular video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 26866–26875, 2025.

Luigi Piccinelli, Yung-Hsu Yang, Christos Sakaridis, Mattia Segu, Siyuan Li, Luc Van Gool, and Fisher Yu. UniDepth: Universal monocular metric depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024.

Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes, 2020. URL https://arxiv.org/abs/2011.13961.

Pradeep Sen, Nima Khademi Kalantari, Maziar Yaesoubi, Soheil Darabi, Dan B. Goldman, and Eli Shechtman. Robust patch-based hdr reconstruction of dynamic scenes. *ACM Trans. Graph.*, 31 (6), November 2012. ISSN 0730-0301. doi: 10.1145/2366145.2366222. URL https://doi.org/10.1145/2366145.2366222.

Dong-Yeon Shin, Jun-Seong Kim, Byung-Ki Kwon, and Tae-Hyun Oh. HDR-NSFF: High dynamic range neural scene flow fields. In *International Conference on Learning Representations*, 2026.

Qianqian Wang, Vickie Ye, Hang Gao, Weijia Zeng, Jake Austin, Zhengqi Li, and Angjoo Kanazawa. Shape of motion: 4D reconstruction from a single video. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2025.

Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4): 600–612, 2004.

Guanjun Wu, Taoran Yi, Jiemin Fang, Wenyu Liu, and Xinggang Wang. Fast high dynamic range radiance fields for dynamic scenes. In *International Conference on 3D Vision (3DV)*, 2024a.

Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Xinggang Wang. 4D gaussian splatting for real-time dynamic scene rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024b.

Jiahao Wu, Lu Xiao, Rui Peng, Kaiqiang Xiong, and Ronggang Wang. HDRGS: High dynamic range gaussian splatting. *arXiv preprint arXiv:2408.06543*, 2024c.

Qingsen Yan, Dong Gong, Javen Qinfeng Shi, Anton van den Hengel, Chunhua Shen, Ian D. Reid, and Yanning Zhang. Attention-guided network for ghost-free high dynamic range imaging. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1751–1760, 2019. URL https://api.semanticscholar.org/CorpusID:128359047.

Zeyu Yang, Hongye Yang, Zijie Pan, and Li Zhang. Real-time photorealistic dynamic scene representation and rendering with 4D gaussian splatting. In *International Conference on Learning Representations*, 2024a.

Ziyi Yang, Xinyu Gao, Wen Zhou, Shaohui Jiao, Yuqing Zhang, and Xiaogang Jin. Deformable 3D gaussians for high-fidelity monocular dynamic scene reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024b.

Kaixuan Zhang, Zhipeng Xiong, Minxian Li, Mingwu Ren, Jiankang Deng, and Xiatian Zhu. Dynamic novel view synthesis in high dynamic range. *arXiv preprint arXiv:2509.21853*, 2025.

Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018.

A Implementation Details

A.1 Tone-Mapper Architecture

The image-space tone mapper is implemented independently for the three color channels. For each channel, the exposure-conditioned scalar input is processed by a 1×1 convolution with $m = 128$ output channels, followed by ReLU, a 1×1 projection to one output channel, and Sigmoid. With $\epsilon = 10^{-6}$, the complete mapping is

$$x_{c,k,p} = \log(H_{c,p} + \epsilon) + \log(E_k + \epsilon), \tag{13}$$

$$a_{c,j,k,p} = w_{c,j} x_{c,k,p} + b_{c,j}, \tag{14}$$

$$h_{c,j,k,p} = \mathrm{ReLU}(a_{c,j,k,p}), \tag{15}$$

$$y_{c,k,p} = \sum_{j=1}^{m} v_{c,j} h_{c,j,k,p} + \beta_c, \tag{16}$$

$$\hat{I}_{c,k,p} = \sigma(y_{c,k,p}). \tag{17}$$

The affine transformation and ReLU form Stage I, which constructs exposure-dependent nonlinear response bases. The subsequent aggregation and Sigmoid projection form Stage II. This decomposition is functional and does not change the original mapper architecture. We apply AME-Reg to the Stage-I pre-activations $a_{c,j,k,p}$, before the ReLU transition.

Each color mapper contains 128 input weights, 128 hidden biases, 128 output weights, and one output bias, giving 385 parameters per channel and 1,155 parameters over RGB.

A.2 AME-Reg Implementation

Bracket-conditioned calibration offsets. The evaluated sequences contain three exposure brackets corresponding to low, middle, and high exposure. We maintain one calibration offset for each color channel, exposure bracket, and Stage-I hidden unit,

$$\Delta b \in \mathbb{R}^{3\times3\times128}, \qquad \Delta b^{(0)}_{c,k,j} = 0. \tag{18}$$

An exposure E is assigned to the nearest reference exposure in log space,

$$k(E) = \arg\min_i \left|\log E - \log E_i\right|. \tag{19}$$

For the alternating three-bracket sequences used in our experiments, this assignment directly recovers the low-, middle-, and high-exposure states. Zero initialization preserves the original tone-mapper response at the start of training.

Pre-activation sampling and quantile estimation. At each iteration, we sample at most 4,096 pixels to estimate the empirical Stage-I pre-activation distribution. Samples are grouped by color channel, exposure bracket, and hidden unit. For a group (c, k, j), let a_p denote its sampled base

pre-activations and let $\delta = \Delta b_{c,k,j}$. We use $q = 0.05$ by default and require the calibrated samples to retain at least a fraction q on each side of the ReLU transition at zero:

$$q \leq \frac{1}{N} \sum_{p=1}^{N} \mathbf{1}[a_p + \delta > 0] \leq 1 - q. \tag{20}$$

The positive-side requirement gives $\delta \geq -Q_{1-q}(a)$, while the negative-side requirement gives $\delta \leq -Q_q(a)$. The resulting feasible interval is therefore

$$-Q_{1-q}(a) \leq \delta \leq -Q_q(a), \tag{21}$$

which is the scalar form of Eq. 10. Because the bounds are recomputed from the empirical distribution, the interval adapts to the current bracket-specific operating range rather than imposing a fixed target activation frequency.

Range loss and gradient flow. Offsets inside the feasible interval incur no range penalty. Offsets below or above the interval are penalized by the two one-sided squared terms in Eq. 11. We use $\lambda_{\text{range}} = 0.006$ in the total objective. When estimating the feasible interval, the rendered HDR samples and shared tone-mapper parameters are detached. Conversely, the calibration offset used in the reconstruction forward pass is detached from the reconstruction loss. Consequently, the range loss updates only the calibration offsets without directly modifying the shared scene representation or base tone mapper, while the reconstruction loss optimizes the original reconstruction parameters under the calibrated operating range without directly fitting the offsets.

Parameters and optimization. The calibration tensor introduces $3 \times 3 \times 128 = 1{,}152$ parameters. Its optimizer uses an initial learning rate of 5×10^{-4}, matching the initial learning rate of the red-channel tone mapper, while the base RGB tone-mapper learning rates retain their original exponential schedules. Quantile estimation uses the sampled pixels only during training; the learned calibration offsets are used directly during inference and do not require quantile computation.

A.3 Training Configuration

Train/test split and image resolution. We train on every other frame and evaluate on the remaining frames, producing an approximately 1:1 train/test split. All methods use the image-downsampling protocol of the HDR-4DGS benchmark. Each baseline and its AME-Reg variant use the same data split and the same method-specific training configuration.

Camera estimation. Alternating exposures can introduce large appearance differences between adjacent frames and weaken photometric camera optimization. For the camera-estimation branch only, we select a middle-exposure image I_{mid} as the photometric reference and histogram-match each low- or high-exposure observation to it. Denoting histogram matching by $\mathcal{H}(\cdot;\cdot)$, the camera-estimation input is

$$\tilde{I}_k = \begin{cases} \mathcal{H}(I_k; I_{\text{mid}}), & E_k \neq E_{\text{mid}}, \\ I_k, & E_k = E_{\text{mid}}. \end{cases} \tag{22}$$

This normalization is not used by the reconstruction branch: the original observation I_k and its recorded exposure E_k remain unchanged for HDR–LDR supervision.

Monocular backbone priors. For HDR-SplineGS, monocular metric depth and long-range point tracks are obtained with UniDepth (Piccinelli et al., 2024) and CoTracker (Karaev et al., 2023), respectively, and dynamic masks follow the Shape-of-Motion (Wang et al., 2025) preprocessing used by the SplineGS training pipeline.

A.4 Integration with Reconstruction Backbones

For each backbone, calibration offsets are inserted before the input-side nonlinear activation while the scene representation, data split, and method-specific optimization protocol remain unchanged. HDR-HexPlane uses the recorded exposures for both variants. For GaussHDR, AME-Reg replaces the original tone-mapping-specific supervision; the remaining reconstruction components are retained. HDR-SplineGS provides the common Stage-I interface used for controlled analysis and ablations.

B Benchmark and Evaluation Protocol

B.1 Benchmark Usage

All evaluated sequences are taken from the established HDR-4DGS benchmark: four real sequences (Bed, Tank, Excavator, and Toys) and eight synthetic sequences. We use the provided exposures, cameras, and HDR references, introduce no additional dataset, and apply the same split and preprocessing to every method.

B.2 Evaluation Protocol

We report PSNR, SSIM, and LPIPS for both HDR and LDR reconstruction. HDR metrics use the benchmark's μ-law mapping, whereas LDR metrics use the exposure-conditioned observations. Metrics are evaluated on held-out frames and averaged over the corresponding real or synthetic split. Photomatix Pro is used only to visualize HDR results and does not affect quantitative evaluation.

C Additional Method Analysis

C.1 Exposure-Dependent Operating Regions

For a multiplicative exposure change $s > 0$, the log-domain input is translated by $\log s$. The corresponding Stage-I pre-activation becomes

$$a'_{c,j,k,p} = a_{c,j,k,p} + w_{c,j} \log s. \tag{23}$$

Thus, fixed tone-mapper parameters can operate over different nonlinear regions across exposure brackets. For ReLU, activation frequencies near zero or one indicate that most samples occupy the inactive region or a single active linear regime, respectively. AME-Reg does not pull this frequency toward a fixed central target; it uses empirical quantiles to define a bracket-adaptive feasible interval.

C.2 Forward and Backward Interpretation

Before the output Sigmoid, the adopted tone mapper is a univariate piecewise-affine function. Wherever it is differentiable,

$$\frac{\partial y_{c,k,p}}{\partial x_{c,k,p}} = \sum_{j=1}^{m} v_{c,j} w_{c,j} \mathbf{1}[a_{c,j,k,p} > 0]. \tag{24}$$

A unit that is inactive over nearly all observed samples contributes little on the empirical support. A unit that is active over nearly all samples contributes an approximately fixed affine component. Units whose activation states vary over the observed range can introduce data-supported changes of slope and therefore provide input-dependent nonlinear responses.

The same activation mask also appears in the reconstruction gradients. Let $g_{c,k,p} = \partial \mathcal{L}_{\text{recon}} / \partial y_{c,k,p}$ absorb the downstream Sigmoid and reconstruction gradients. Then

$$\frac{\partial \mathcal{L}_{\text{recon}}}{\partial b_{c,j}} = \sum_{k,p} g_{c,k,p} v_{c,j} \mathbf{1}[a_{c,j,k,p} > 0], \tag{25}$$

$$\frac{\partial \mathcal{L}_{\text{recon}}}{\partial w_{c,j}} = \sum_{k,p} g_{c,k,p} v_{c,j} x_{c,k,p} \mathbf{1}[a_{c,j,k,p} > 0]. \tag{26}$$