



EnslR: An Ensemble Algorithm for Image Restoration via Gaussian Mixture Models

Shangquan Sun^{1,2}, Wenqi Ren³, Zikun Liu⁴, Hyunhee Park⁴, Rui Wang^{1,2}, Xiaochun Cao³

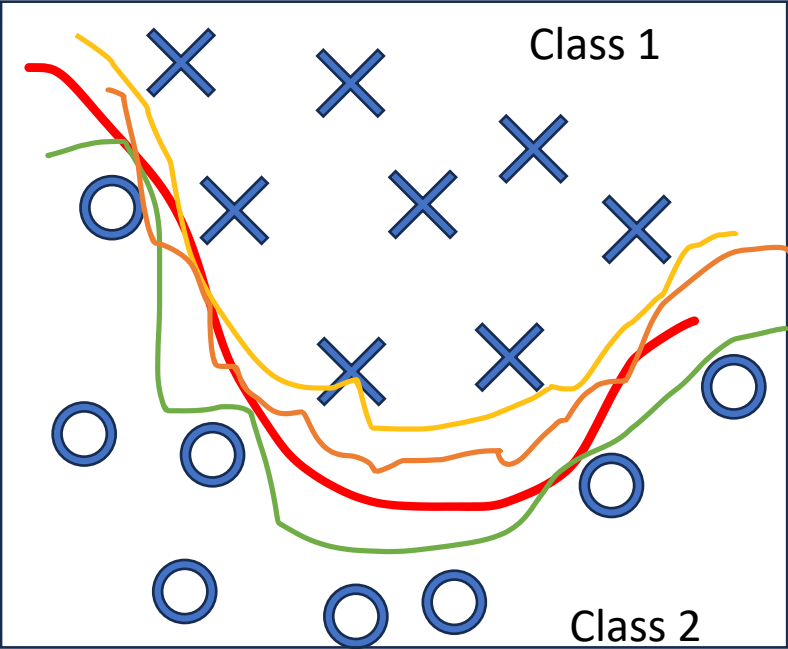
¹CAS, China ²UCAS, China ³Shenzhen Campus of SYSU, China ⁴Samsung

NeurIPS 2024



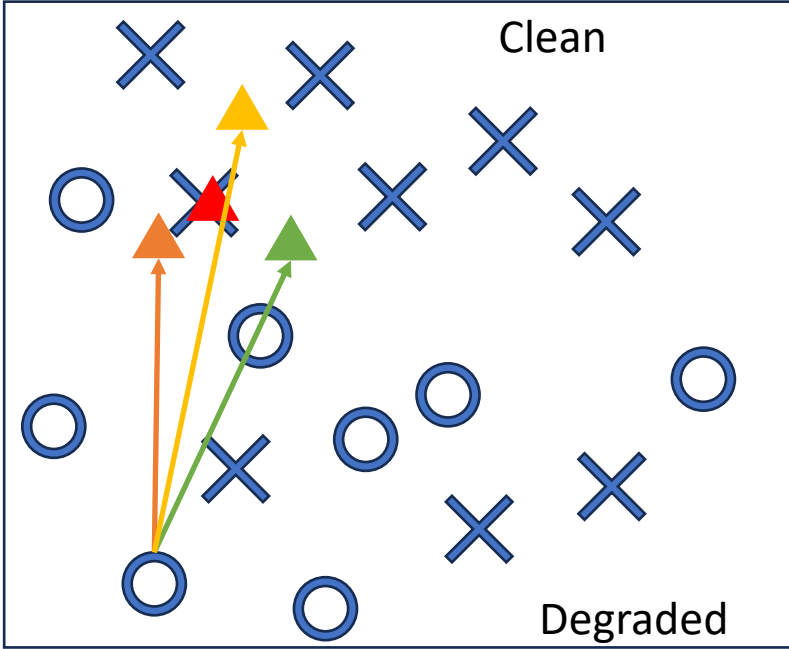
2024.11.2

Motivation

Classification: $\mathbb{R}^{3 \times H \times W} \rightarrow K$



Restoration: $\mathbb{R}^{3 \times H \times W} \rightarrow \mathbb{R}^{3 \times H \times W}$



Problem Formulation

- Ensemble for Image Restoration

$$\tilde{\mathbf{Y}}_n = \boldsymbol{\beta}_n^\top [\tilde{\mathbf{X}}_{1,n} \quad \cdots \quad \tilde{\mathbf{X}}_{M,n}], \quad \forall n \quad (1)$$

- For a reference set, formulate all data with Gaussian prior

$$\mathbf{y}_{1:N} | f_m, \mathbf{x}_{1:N} \sim \mathcal{N}(\mathbf{x}_{m,1:N}, \text{diag}(\boldsymbol{\Sigma}_{m,1}, \dots, \boldsymbol{\Sigma}_{m,N})), \quad (2)$$

- Partition data into bins:

$$\mathbf{y}_{r,1:N} = \mathbf{R}_r \cdot \mathbf{y}_{1:N}, \quad \mathbf{x}_{r,m,1:N} = \mathbf{R}_r \cdot \mathbf{x}_{m,1:N}, \quad (3)$$

- Within each bin, Gaussian prior still holds

$$\mathbf{y}_{r,1:N}^{(i)} | f_m, \mathbf{x}_{r,m,1:N} \stackrel{i.i.d.}{\sim} \mathcal{N}(\boldsymbol{\mu}_{r,m,1:N}, \boldsymbol{\sigma}_{r,m,1:N}), \quad \forall i \in [1, \dots, N_r], \quad (4)$$

- Then ensemble problem is split for estimating weights within bins

$$\mathbf{y}_{r,1:N}^{(i)} = \mathbb{E}_z \left[\mathbf{x}_{r,m,1:N}^{(i)} \right] = \sum_{m=1}^M \alpha_{r,m} \cdot \mathbf{x}_{r,m,1:N}^{(i)}; \quad P(\mathbf{y}_{r,1:N}^{(i)}) = \sum_{m=1}^M \alpha_{r,m} P(\mathbf{y}_{r,1:N}^{(i)} | z = m), \quad (5)$$

$$\{\alpha_{r,m}\}_{r,m} \in \arg \max P(\mathbf{y}_{1:N}). \quad (6)$$

Gaussian Mixture Models and EM Algorithm

- Gaussian Mixture Models:

$$\alpha_{r,m} \in \arg \max_{\alpha_{r,m}} P(\mathbf{y}_{r,1:N}) = \arg \max_{\alpha_{r,m}} \prod_{i=1}^{N_r} P(\mathbf{y}_{r,1:N}^{(i)}). \quad (7)$$

- Estimate ensemble weights by maximizing the log likelihood as

$$\begin{aligned} \log P(\mathbf{y}_{r,1:N}) &= \log \prod_{i=1}^{N_r} \sum_{m=1}^M \alpha_{r,m} \phi(\mathbf{y}_{r,1:N}^{(i)}; \boldsymbol{\mu}_{r,m,1:N}, \boldsymbol{\sigma}_{r,m,1:N}) \\ &= \sum_{i=1}^{N_r} \log \sum_{m=1}^M \alpha_{r,m} \phi(\mathbf{y}_{r,1:N}^{(i)}; \boldsymbol{\mu}_{r,m,1:N}, \boldsymbol{\sigma}_{r,m,1:N}) \\ &\geq \sum_{m=1}^M P(z = m | \mathbf{y}_{r,1:N}^{(i)}) \log \frac{\alpha_{r,m} \phi(\mathbf{y}_{r,1:N}^{(i)}; \boldsymbol{\mu}_{r,m,1:N}, \boldsymbol{\sigma}_{r,m,1:N})}{P(z = m | \mathbf{y}_{r,1:N}^{(i)})}, \end{aligned} \quad (8)$$

Gaussian Mixture Models and EM Algorithm

- Expectation step:

$$\gamma_{r,m,1:N} \leftarrow P(z = m | \mathbf{y}_{r,1:N}^{(i)}) = \frac{\alpha_{r,m} \phi(\mathbf{y}_{r,1:N}^{(i)}; \boldsymbol{\mu}_{r,m,1:N}, \boldsymbol{\sigma}_{r,m,1:N})}{\sum_{m=1}^M \alpha_{r,m} \phi(\mathbf{y}_{r,1:N}^{(i)}; \boldsymbol{\mu}_{r,m,1:N}, \boldsymbol{\sigma}_{r,m,1:N})}. \quad (9)$$

- Maximization step:

$$\alpha_{r,m} \leftarrow \frac{1}{N_r} \sum_{i=1}^{N_r} \gamma_{r,m,1:N} \quad (10)$$

$$\boldsymbol{\sigma}_{r,m,1:N} \leftarrow \frac{\sum_{i=1}^{N_r} \gamma_{r,m,1:N} (\mathbf{y}_{r,m,1:N}^{(i)} - \boldsymbol{\mu}_{r,m,1:N})^2}{\sum_{n=1}^{N_r} \gamma_{r,m,1:N}}. \quad (11)$$

- Priors of Mean and Variance:

$$\boldsymbol{\mu}_{r,m,1:N} \leftarrow \frac{1}{N_r} \sum_{i=1}^{N_r} \mathbf{x}_{r,m,1:N}^{(i)} \quad (12)$$

$$\boldsymbol{\sigma}_{r,m,1:N} \leftarrow \frac{1}{N_r} \|\mathbf{x}_{r,m,1:N} - \boldsymbol{\mu}_{r,m,1:N}\|_2. \quad (13)$$

Look-up Table for inference and Overall Algorithm

- LUT to save the weights for inference

$$\tilde{\mathbf{X}}_{r,m,n} = \mathbf{R}_r \cdot \tilde{\mathbf{X}}_{m,n}, \text{ where } \mathbf{R}_r = \prod_{m=1}^M \mathbf{I}_{B_m}(\tilde{\mathbf{X}}_{m,n}). \quad (14)$$

$$\tilde{\mathbf{Y}}_n = \sum_{r=1}^{T^M} \tilde{\mathbf{Y}}_{r,n} = \sum_{r=1}^{T^M} \sum_{m=1}^M \alpha_{r,m} \tilde{\mathbf{X}}_{r,m,n}. \quad (15)$$

Algorithm 2: MPEM: EM algorithm with known Mean Prior

Input: $\mathbf{y}_{r,1:N}, \mathbf{x}_{r,1,1:N}, \dots, \mathbf{x}_{r,M,1:N}$
Output: $(\alpha_{r,1}, \dots, \alpha_{r,M})$

- 1 $N_r \leftarrow$ number of nonzero pixels in $\mathbf{y}_{r,1:N}$;
- 2 Initialize $\mu_{r,m,1:N}$ by Eq. 12;
- 3 Initialize $\sigma_{r,m,1:N}$ by Eq. 13;
- 4 **while not converge do**
- 5 **for** $i \in [1, N_r]$ **do**
- 6 **for** $m \in [1, M]$ **do**
- 7 Update $\gamma_{r,m,1:N}$ by Eq. 9;
- 8 **end**
- 9 **end**
- 10 **for** $m \in [1, M]$ **do**
- 11 Update $\alpha_{r,m}$ by Eq. 10;
- 12 Update $\sigma_{r,m,1:N}$ by Eq. 11;
- 13 **end**
- 14 **end**

Algorithm 1: EnsIR: an ensemble algorithm for image restoration

Input: A small reference dataset $\{\mathbf{x}_n, \mathbf{y}_n\}_{n=1}^N$ for ensemble weight estimation, test set $\{\hat{\mathbf{X}}_n\}$, M pre-trained models f_1, \dots, f_M , bin width b , Empty lookup table LUT

Output: Ensemble result $\{\tilde{\mathbf{Y}}_n\}$

Estimation Stage:


- 1 Obtain restoration predictions by $\mathbf{x}_{m,n} = \text{flatten}(f_m(\mathbf{X}_n))$, $\forall m \in \{1, \dots, M\}$;
- 2 Append restoration predictions and ground-truths into $\mathbf{y}_{1:N}$ and $\mathbf{x}_{m,1:N}$ based on Eq. 2;
- 3 Define bin set space $\mathbb{B} = \{[0, b), [b, 2b), \dots, [(T-1)b, 255]\}$;
- 4 **for each bin set** $(B_1, \dots, B_M) \in \mathbb{B}^M$ **do**
- 5 Compute the partition map $\mathbf{R}_r = \prod_{m=1}^M \mathbf{I}_{B_m}(f_m(\mathbf{x}_n))$;
- 6 Partition images and obtain range-wise patches $(\mathbf{y}_{r,1:N}, \mathbf{x}_{r,1,1:N}, \dots, \mathbf{x}_{r,M,1:N})$ by Eq. 3;
- 7 $(\alpha_{r,1}, \dots, \alpha_{r,M}) \leftarrow$ **MPEM** $(\mathbf{y}_{r,1:N}, \mathbf{x}_{r,1,1:N}, \dots, \mathbf{x}_{r,M,1:N})$;
- 8 Store $\text{LUT}[(B_1, \dots, B_M)] \leftarrow (\alpha_{r,1}, \dots, \alpha_{r,M})$;
- 9 **end**

Inference Stage:

- 10 **for each test data** $\hat{\mathbf{X}}_n$ **do**
- 11 **for each bin set** $(B_1, \dots, B_M) \in \mathbb{B}^M$ **do**
- 12 Retrieve $(\alpha_{r,1}, \dots, \alpha_{r,M}) \leftarrow \text{LUT}[(B_1, \dots, B_M)]$;
- 13 Partition input as $\tilde{\mathbf{X}}_{r,m,n} \leftarrow \mathbf{R}_r \cdot f_m(\hat{\mathbf{X}}_n)$, where $\mathbf{R}_r = \prod_{m=1}^M \mathbf{I}_{B_m}(f_m(\hat{\mathbf{X}}_n))$;
- 14 $\tilde{\mathbf{Y}}_{r,n} \leftarrow \sum_{m=1}^M \alpha_{r,m} \tilde{\mathbf{X}}_{r,m,n}$; /* Inner summation of Eq. 15 */
- 15 **end**
- 16 $\tilde{\mathbf{Y}}_n \leftarrow \sum_{r=1}^{T^M} \tilde{\mathbf{Y}}_{r,n}$; /* Outer summation of Eq. 15 */
- 17 **end**

A Simple Case with bin width=64

Restored \mathbf{x}_1



Pixels

0	0	1
0	0	0
0	0	67

Restored \mathbf{x}_2



Pixels

94	68	87
68	24	99
97	89	72

Restored \mathbf{x}_3



Pixels

98	99	97
88	91	89
92	94	93

Bin set 1:

$[0, 64)$ for \mathbf{x}_1
 $[0, 64)$ for \mathbf{x}_2
 $[0, 64)$ for \mathbf{x}_3

α_1

	0		24		91

Bin set 2:

$[0, 64)$ for \mathbf{x}_1
 $[0, 64)$ for \mathbf{x}_2
 $[64, 128)$ for \mathbf{x}_3

$\alpha_2 = [0.1, 0.2, 0.7] = \text{GMM}([0, 24, 91], 81)$

.....

Bin set 21:

$[64, 128)$ for \mathbf{x}_1
 $[64, 128)$ for \mathbf{x}_2
 $[64, 128)$ for \mathbf{x}_3

0	0	1	94	68	87	98	99	97
0		0	68		99	88		89
0	0	0	97	89	72	92	94	93

$\alpha_{21} = [0.1, 0.4, 0.5] = \text{GMM}([0.125, 84.25, 93.74], 91)$

.....

Bin set 4^3 :

$(192, 256]$ for \mathbf{x}_1
 $(192, 256]$ for \mathbf{x}_2
 $(192, 256]$ for \mathbf{x}_3

α_{64}

96	93	91
84	81	90
94	92	88

Result 2: Ablation Studies

Table 1: Ablation study of bin width b on Rain100H [74] with maximum step number 1000. “Runtime” is the average runtime [s].

b	16	32	64	96	128
Runtime	1.2460	0.1709	0.0265	0.0132	0.0059
PSNR	31.745	31.739	31.720	31.713	31.725
SSIM	0.9093	0.9095	0.9094	0.9093	0.9093

Table 2: Ablation study of maximum step number in the EM algorithm on Rain100H [74] with $b = 32$. “Time” is the time of EM algorithm [s].


#step	10	100	500	1000	10000
Time	12.108	28.516	30.409	30.518	30.524
PSNR	31.734	31.738	31.738	31.739	31.739
SSIM	0.9093	0.9094	0.9095	0.9095	0.9095

Table 6: The average runtime per image in seconds of the ensemble methods on Rain100H [74].

Method	Bagging [6]	AdaBoost [22]	RForest [7]	GBDT [23]	HGBT [34]	Average	ZZPM [92]	Ours
Runtime	1.0070	1.1827	9.8598	1.2781	0.1773	0.0003	0.0021	0.1709


Quantitative Comparison: Super-Resolution

Datasets	Set5 [5]	Set14 [83]	BSDS100 [47]	Urban100 [30]	Manga109 [48]						
Metrics	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	
Base	SwinIR [39]	32.916	0.9044	29.087	0.7950	27.919	0.7487	27.453	0.8254	32.024	0.9260
	SRFormer [93]	32.922	0.9043	29.090	0.7942	27.914	0.7489	27.535	0.8261	32.203	0.9271
	MambaIR [27]	33.045	0.9051	29.159	0.7958	27.967	0.7510	27.775	0.8321	32.308	0.9283
Regr.	Bagging [6]	33.006	0.9050	29.119	0.7950	27.946	0.7498	27.546	0.8273	32.154	0.9270
	AdaBoost [22]	33.072	0.9049	29.175	0.7959	27.975	0.7503	27.786	0.8302	32.457	0.9286
	RForest [7]	33.032	0.9052	29.158	0.7954	27.964	0.7500	27.640	0.8287	32.287	0.9279
	GBDT [23]	33.085	0.9050	29.196	0.7956	27.980	0.7500	27.792	0.8311	32.467	0.9285
	HGBT [34]	33.078	0.9051	29.201	0.7959	27.984	0.7502	27.783	0.8310	32.444	0.9282
IR.	Average	33.097	0.9057	29.202	0.7964	27.983	0.7506	27.785	0.8313	32.466	0.9290
	RefESR [31]	33.091	0.9052	29.172	0.7960	27.972	0.7504	27.785	0.8312	32.447	0.9288
	ZZPM [92]	33.094	0.9057	29.203	0.7963	27.981	0.7506	27.786	0.8313	32.467	0.9290
	EnsIR (Ours)	33.103	0.9058	29.205	0.7964	27.984	0.7507	27.795	0.8315	32.468	0.9291




Quantitative Comparison: Deblurring

Datasets		GoPro [51]		HIDE [60]		RealBlur-R [58]		RealBlur-J [58]	
Metrics		PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
Base	MPRNet [82]	32.658	0.9362	30.962	0.9188	33.914	0.9425	26.515	0.8240
	Restormer [81]	32.918	0.9398	31.221	0.9226	33.984	0.9463	26.626	0.8274
	DGUNet [50]	33.173	0.9423	31.404	0.9257	33.990	0.9430	26.583	0.8261
Regr.	Bagging [6]	33.194	0.9418	31.437	0.9250	34.033	0.9456	26.641	0.8277
	AdaBoost [22]	33.205	0.9412	31.449	0.9251	34.035	0.9455	26.652	0.8280
	RForest [7]	33.173	0.9416	31.439	0.9247	34.039	0.9457	26.647	0.8280
	GBDT [23]	33.311	0.9418	31.568	0.9256	34.052	0.9465	26.684	0.8285
	HGBT [34]	33.323	0.9427	31.583	0.9267	33.986	0.9436	26.684	0.8296
IR.	Average	33.330	0.9436	31.579	0.9277	34.090	0.9471	26.689	0.8309
	ZZPM [92]	33.332	0.9436	31.580	0.9277	34.057	0.9468	26.688	0.8308
	EnsIR (Ours)	33.345	0.9438	31.590	0.9278	34.089	0.9472	26.690	0.8309



Quantitative Comparison: Deraining

Datasets	Rain100H [74]	Rain100L [74]	Test100 [88]	Test1200 [86]	Test2800 [25]						
Metrics	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM					
Base	MPRNet [82]	30.428	0.8905	36.463	0.9657	30.292	0.8983	32.944	0.9175	33.667	0.9389
	MAXIM [67]	30.838	0.9043	38.152	0.9782	31.194	0.9239	32.401	0.9240	33.837	0.9438
	Restormer [81]	31.477	0.9054	39.080	0.9785	32.025	0.9237	33.219	0.9270	34.211	0.9449
Regr.	Bagging [6]	31.461	0.9001	39.060	0.9782	31.865	0.9107	33.115	0.9152	34.216	0.9446
	AdaBoost [22]	31.472	0.9006	39.067	0.9782	31.866	0.9112	33.117	0.9153	34.221	0.9443
	RForest [7]	31.492	0.9012	39.089	0.9784	31.900	0.9127	33.147	0.9169	34.224	0.9447
	GBDT [23]	31.581	0.9058	39.044	0.9778	32.001	0.9236	33.276	0.9274	34.211	0.9446
	HGBT [34]	31.698	0.9075	39.115	0.9784	31.988	0.9241	33.305	0.9282	34.229	0.9450
IR.	Average	31.681	0.9091	38.675	0.9770	31.626	0.9225	33.427	0.9286	34.214	0.9449
	ZZPM [92]	31.689	0.9091	38.725	0.9771	31.642	0.9227	33.434	0.9286	34.231	0.9450
	EnsIR (Ours)	31.739	0.9095	39.216	0.9792	32.064	0.9258	33.445	0.9289	34.245	0.9451



(a) Image (b) GT & LQ (c) MPRNet (d) HGBT [34] (e) Average (f) ZZPM [92] (g) Ours
 PSNR/SSIM 34.008/0.9359 37.052/0.9615 36.775/0.9603 36.777/0.9601 37.372/0.9639

Visualization

- Weight Map

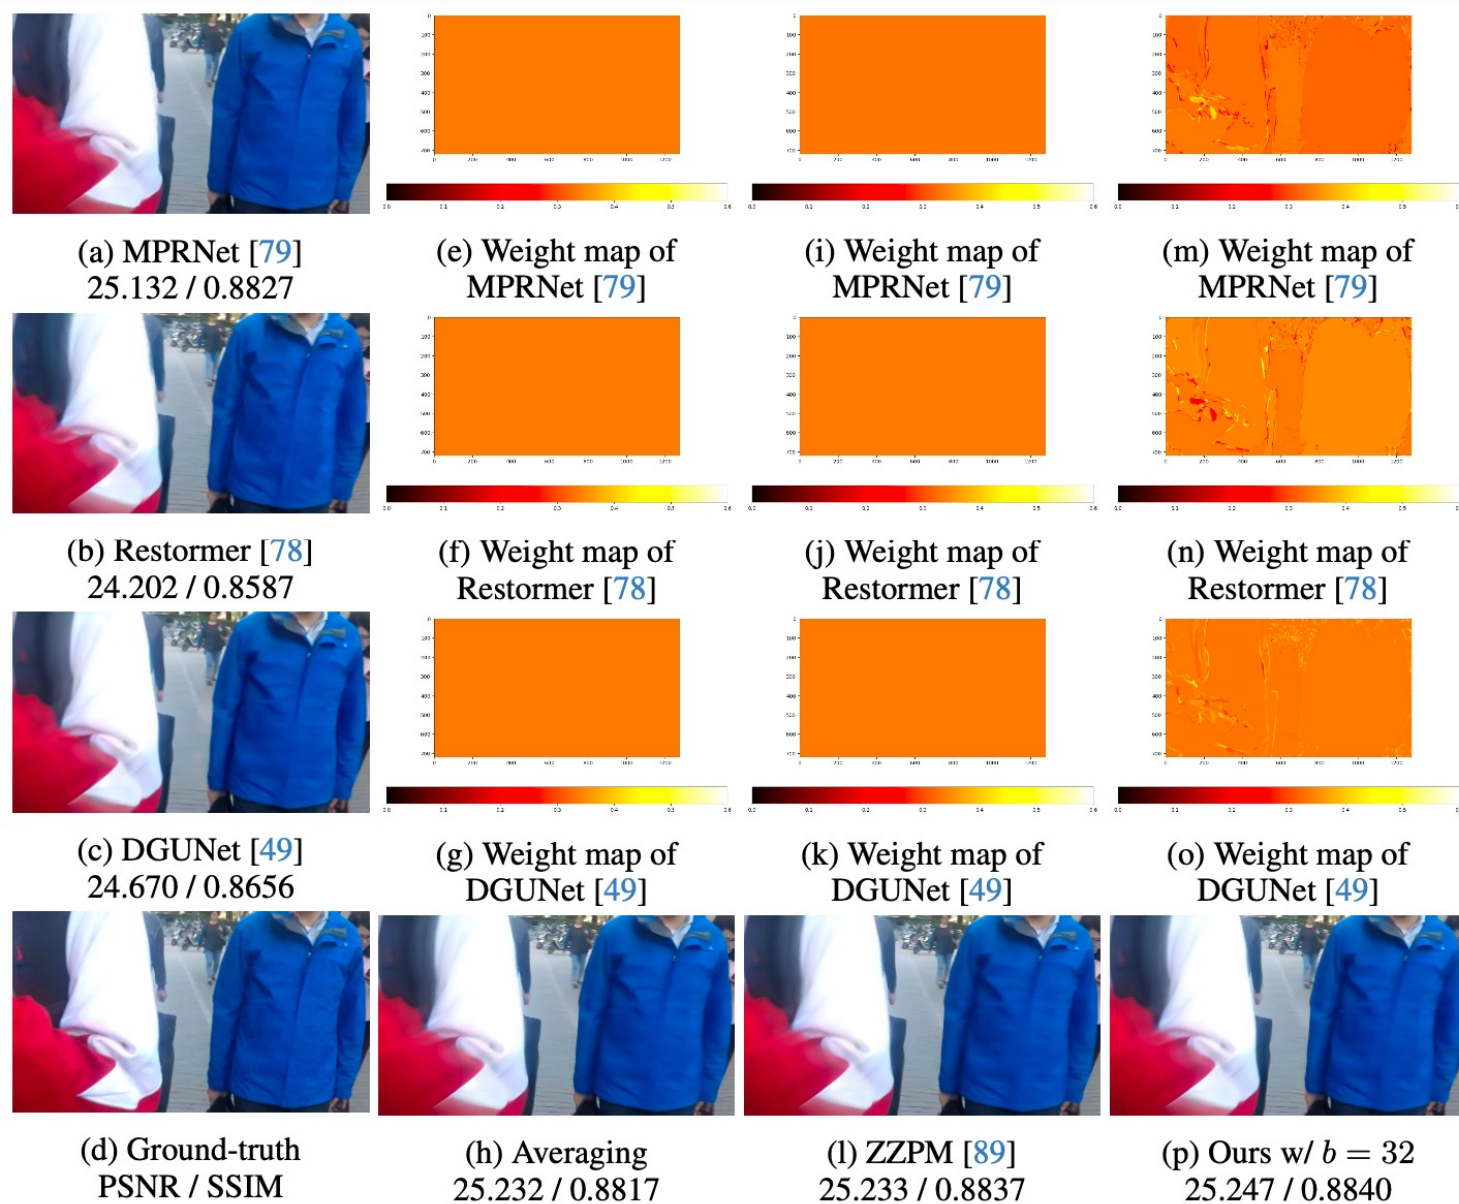


Figure 4: A sample of weight visualizations on HIDE [59]. Base models are DGUNet [49], MPRNet [79], and Restormer [78]. The first column shows the base model predictions and ground-truth. The second column shows the ensemble weights and result of the averaging strategy. The third column shows the ensemble weights and result of ZZPM [89]. The last column shows the ensemble weights and result of our method.

Visualization

- Feature by ResNet110 and PCA







Figure 5: A sampled group of feature visualizations on HIDE [59]. “Base” denotes the features of three base models, i.e., DGUNet [49], MPRNet [79], and Restormer [78].

Visualization

- Pixel Distribution within bin set

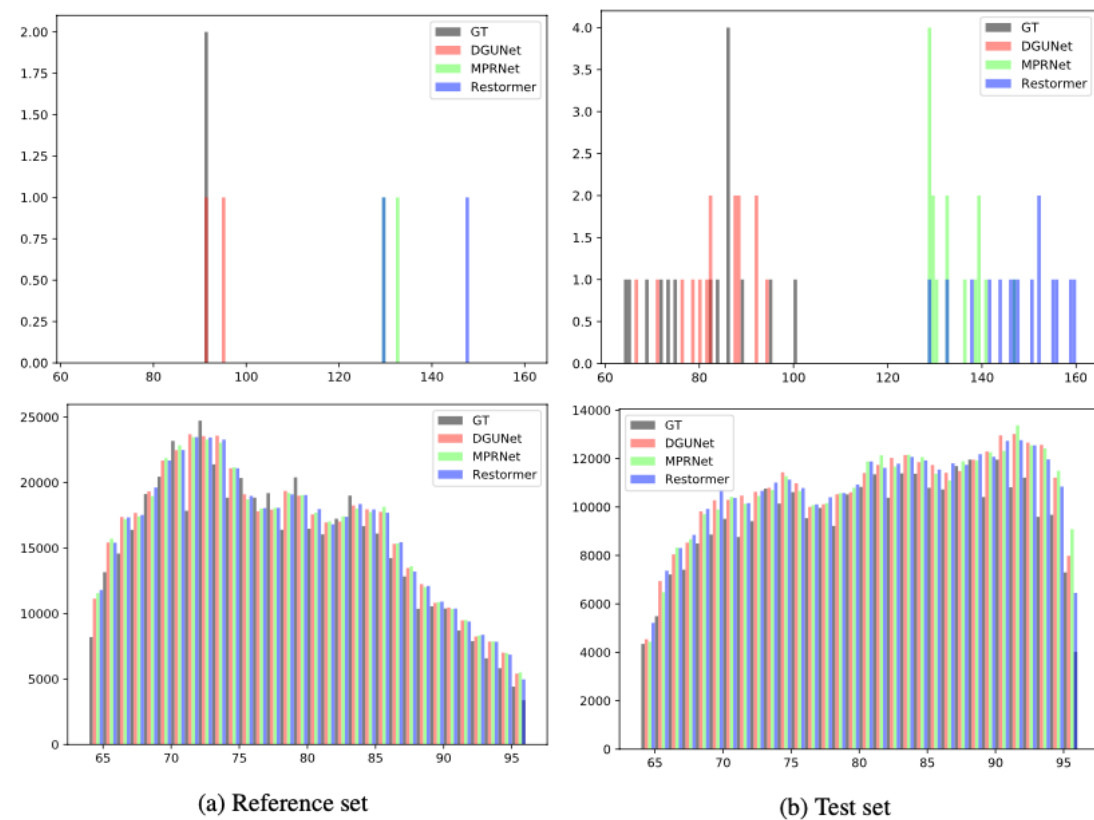
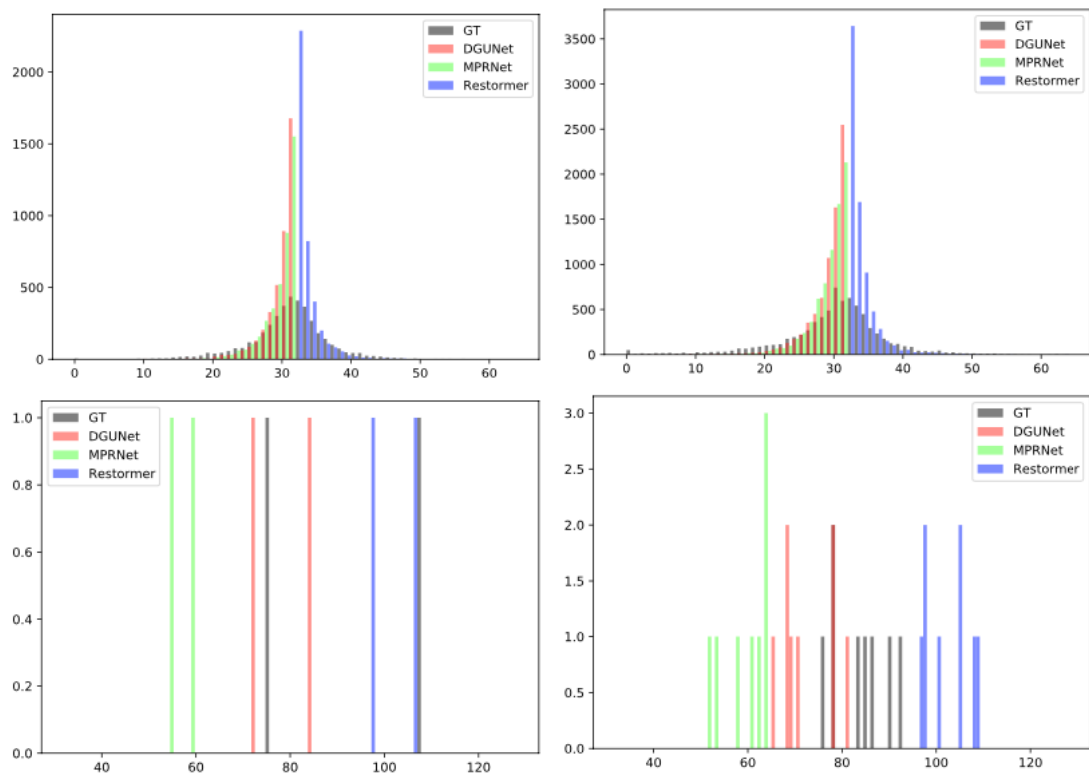


Figure 6: A group of distribution visualizations on HIDE [59]. The bin sets of the first row is ($B_1 = [0, 32)$, $B_2 = [0, 32)$, $B_3 = [32, 64)$). The bin sets of the second row is ($B_1 = [64, 96)$, $B_2 = [32, 64)$, $B_3 = [96, 128)$). The bin sets of the third row is ($B_1 = [64, 96)$, $B_2 = [128, 160)$, $B_3 = [128, 160)$). The bin sets of the last row is ($B_1 = [64, 96)$, $B_2 = [64, 96)$, $B_3 = [64, 96)$). Base models are DGUNet [49], MPRNet [79], and Restormer [78].

Thank You