ICML
International Conference
On Machine Learning

ICL
Intelligent Comput.
imaging Lab.

VERITAS LUX MEA

Triadic Dynamics Aware Diffusion Posterior Sampling for Inverse Problems: Optimizing Guidance and Stochasticity Schedules





Junseo Bang[1,*] Dong Ju Mun[1,*] Hoigi Seo[1] Seongmin Hong[2] Se Young Chun[1,2,3,†]

*Co-first authors †Corresponding author

1 Dept. of ECE, Republic of Korea, 2 INMC, 3 IPAI & AIIS, Seoul National University, Republic of Korea

Inverse Problems & Generative Priors

- Recover a clean signal x from a degraded observation y

$$y = \mathcal{A}x + n, \qquad n \sim \mathcal{N}(0, \sigma_n^2 \mathbf{I})$$

  - $\mathcal{A}$: forward operator (degradation)
  - n : additive Gaussian measurement noise
  - Ill-posed when dim(y) < dim(x) — infinitely many solutions → a prior is required

- Goal: sample from the posterior $p(x|y) \propto p(y|x) \cdot p(x)$
- Diffusion / flow matching models provide an expressive generative priors $p(x)$ for posterior sampling
- Flow matching → straighter trajectories → efficient few-step inference

Flow-based Posterior Sampling

Prior works schedule only subsets of these scales.
we schedule all three jointly and characterize their coupling.

CFG $\lambda(t)$	$v_t(\boldsymbol{x}_t) = v_\theta(\boldsymbol{x}_t, c_\emptyset) + \lambda(t)\big(v_\theta(\boldsymbol{x}_t, c) - v_\theta(\boldsymbol{x}_t, c_\emptyset)\big).$
Tweedie's formula	$\hat{\boldsymbol{x}}_{0\|t} = \boldsymbol{x}_t - \sigma_t v_t(\boldsymbol{x}_t), \quad \hat{\boldsymbol{x}}_{1\|t} = \boldsymbol{x}_t + (1-\sigma_t) v_t(\boldsymbol{x}_t),$
DC $\beta(t)$	$\tilde{\boldsymbol{x}}_{0\|t}(\boldsymbol{y}) = \hat{\boldsymbol{x}}_{0\|t} - \beta(t)\nabla_{\hat{\boldsymbol{x}}_{0\|t}}\mathcal{L}(\mathcal{A}\hat{\boldsymbol{x}}_{0\|t}, \boldsymbol{y}).$
Stoch $\eta(t)$	$\tilde{\boldsymbol{x}}_{1\|t} = \sqrt{1-\eta^2(t)}\,\hat{\boldsymbol{x}}_{1\|t} + \eta(t)\boldsymbol{\epsilon}.$
Euler update	$\boldsymbol{x}_{t+\Delta t} = (1-\sigma_{t+\Delta t})\tilde{\boldsymbol{x}}_{0\|t}(\boldsymbol{y}) + \sigma_{t+\Delta t}\tilde{\boldsymbol{x}}_{1\|t}.$

	Prior works	Ours (TriPS)
DC $\beta(t)$	partial	scheduled
CFG $\lambda(t)$	fixed	scheduled
Stoch $\eta(t)$	partial	scheduled
Coupling	—	**characterized**

Three key components

Data consistency
align with the measurement y

Classifier-free guidance
steer toward the text prompt

Stochasticity
inject random noise

FlowDPS: [Kim, ICCV 2025] FlowDPS : Flow-Driven Posterior Sampling for Inverse Problems

Scheduling the Scales Matters

Inpainting

PSNR: 18.69 dB
LPIPS: 0.124

PSNR: 23.23 dB
LPIPS: 0.089

Measurement | DC, CFG, Stoch. **Fixed** | DC, CFG, Stoch. **Scheduled** | Ground Truth

Super-resolution

PSNR: 28.99 dB
LPIPS: 0.157

PSNR: 29.35 dB
LPIPS: 0.103

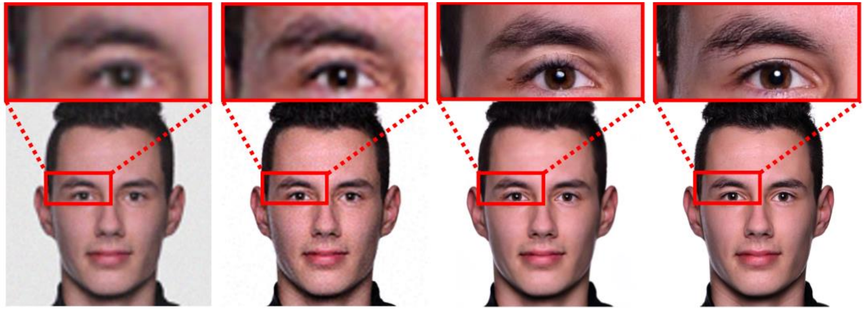

Measurement | DC, CFG, Stoch. **Fixed** | DC, CFG, Stoch. **Scheduled** | Ground Truth

- Same posterior sampler, only the schedules of $\beta(t), \lambda(t), \eta(t)$ differ, yet quality changes sharply.
- **How should the three scales be scheduled over time?**

Triadic Coupling Dynamics

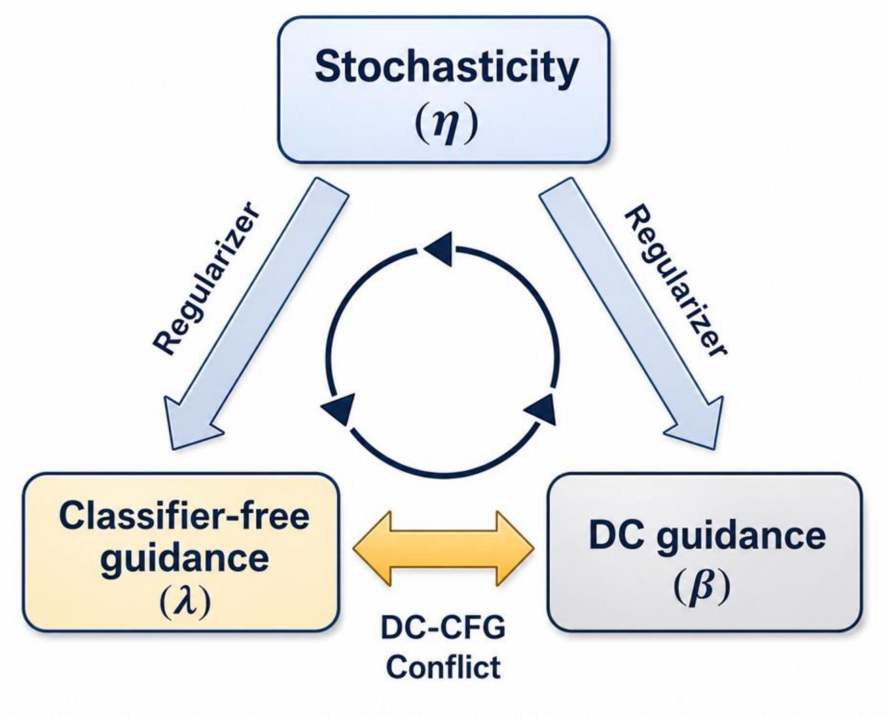

- Early-stage Guidance conflict
  - the two guidances oppose each other early in sampling
- Stochasticity as a Regularizer for DC guidance and CFG
  - Stochastic noise pulls trajectories back to high-probability regions

Early-stage Guidance conflict

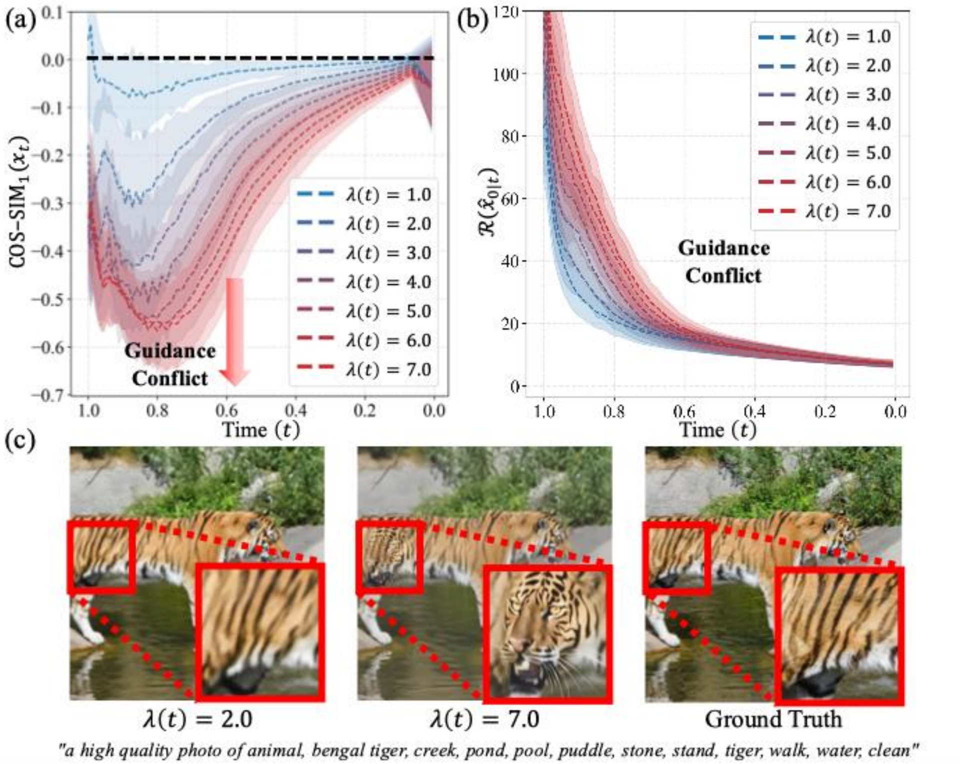

$$\text{COS-SIM}_1(\boldsymbol{x}_t) \triangleq \frac{\left\langle \tilde{b}_{\text{dc}}(\boldsymbol{x}_t; y),\, \tilde{b}_{\text{cfg}}(\boldsymbol{x}_t; c) \right\rangle}{\left\| \tilde{b}_{\text{dc}}(\boldsymbol{x}_t; y) \right\|_2 \left\| \tilde{b}_{\text{cfg}}(\boldsymbol{x}_t; c) \right\|_2},$$

$$\mathcal{R}(\hat{x}_{0|t}) \triangleq \left\| y - \mathcal{A}\hat{x}_{0|t} \right\|_2^2$$

- ❖ Early ($t \approx 1$): CFG opposes DC, $\text{COS-SIM}_1 < 0$.
- ❖ Prop. 1: larger $\lambda(t)$ slows residual-norm descent when $\langle \tilde{b}_{dc}, \tilde{b}_{cfg} \rangle < 0$.
- ❖ High CFG scale ⇒ semantic hallucinations that violate the measurement.

High-noise regime
⇒ keep $\lambda(t)$ low, $\beta(t)$ high.

Stochasticity as a Regularizer

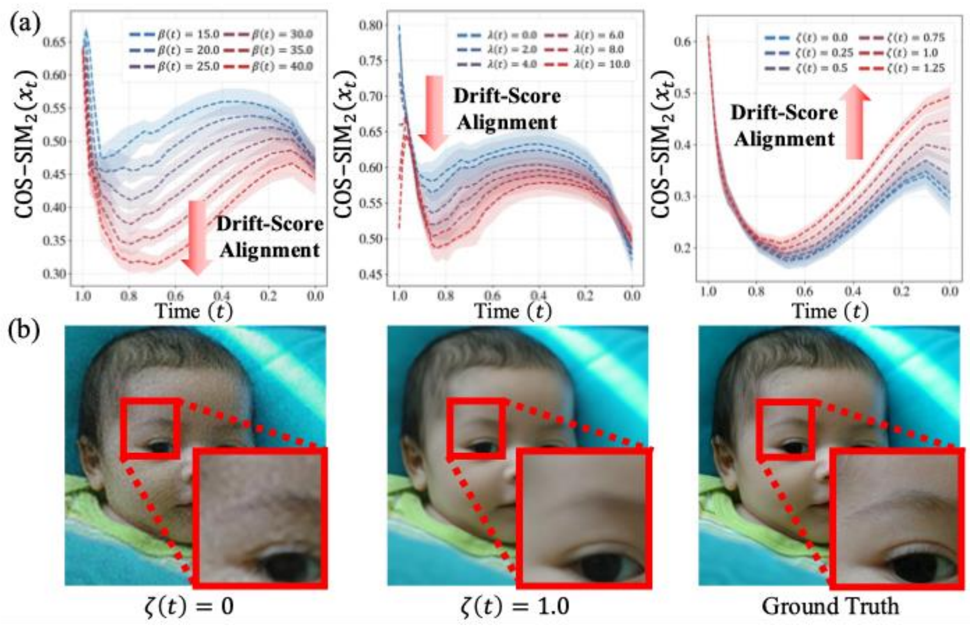

$$\mathrm{COS\text{-}SIM}_2(\boldsymbol{x}_t) \triangleq \frac{\left\langle b_{\det}(\boldsymbol{x}_t),\ \nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t) \right\rangle}{\left\| b_{\det}(\boldsymbol{x}_t) \right\|_2 \left\| \nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t) \right\|_2},$$

- ❖ Stronger $\beta(t)$ or $\lambda(t)$ ⇒ drift leaves the higher probability region ($\mathrm{COS\text{-}SIM}_2 \downarrow$).
- ❖ More $\eta(t)$ ⇒ realigns drift with the score ($\mathrm{COS\text{-}SIM}_2 \uparrow$)
- ❖ Stochasticity at early timestep helps high $\mathrm{COS\text{-}SIM}_2$ at later timesteps.

Monotonic triadic scheduling trend: $\beta(t) \downarrow,\ \lambda(t) \uparrow,\ \eta(t) \downarrow$

Triadic Schedule Optimization

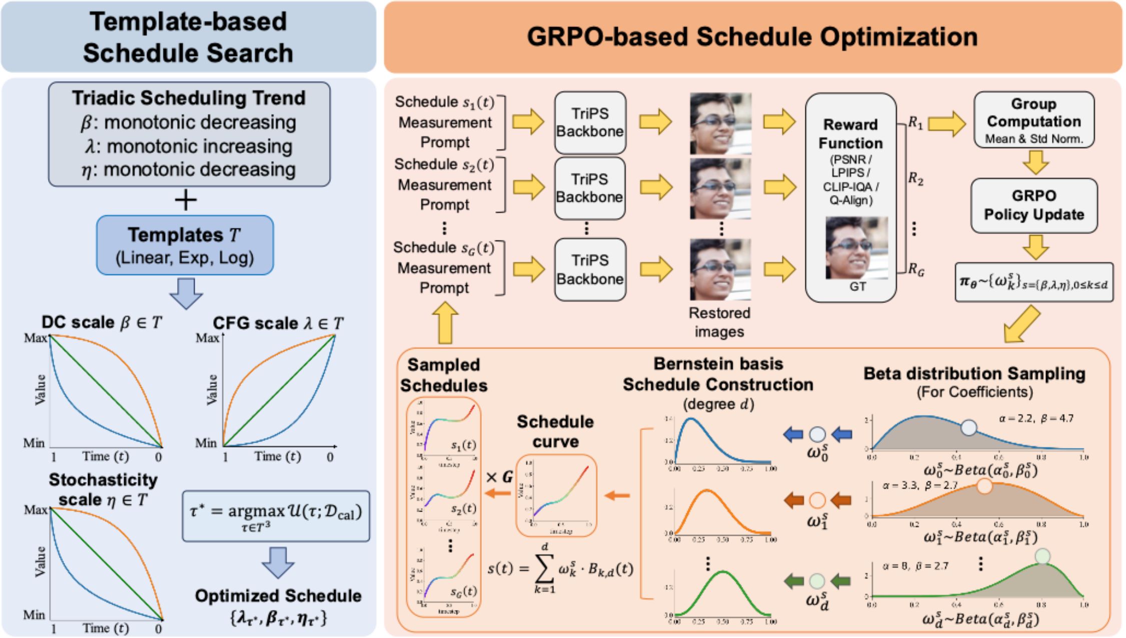

TriPS$_T$ · Template search

Coarse search over functional priors $\mathcal{T}$ = {linear, exp, log}.

$$\tau^\star = \arg\max_{\tau \in \mathcal{T}^3} \mathcal{U}(\tau; \mathcal{D}_{cal})$$

TriPS$_G$ · GRPO optimization

Fine curves beyond fixed forms via Bernstein-Beta parameterization:

$$\tilde{s}(t) = \sum_{k=0}^{d} w_k^{(s)} B_{k,d}(t),$$

$$w_k^{(s)} \sim \mathrm{Beta}(a_k^{(s)}, b_k^{(s)}), \quad s \in \{\lambda, \beta, \eta\}$$

Hybrid distortion + perceptual reward:

$$R = w_{dist} R_{dist} + w_{perc} R_{perc}$$

GRPO: [Shao, arXiv 2024] Deepseekmath: Pushing the limits of mathematical reasoning in open language models

Quantitative results — Linear Inverse Problems

Table 1. Quantitative comparison on linear inverse problems. Results are averaged over 1,000 FFHQ and 800 DIV2K samples using 28 NFEs with Gaussian noise ($\sigma_n = 0.03$). **Bold** and $\underline{\text{underline}}$ denote the best and second-best performance, respectively.

Flow Matching Model (SD3.5-M)																				
	Super-Resolution ×8					Super-Resolution ×12					Motion Deblurring					Gaussian Deblurring				
Method	PSNR↑	SSIM↑	FID↓	LPIPS↓	MUSIQ↑	PSNR↑	SSIM↑	FID↓	LPIPS↓	MUSIQ↑	PSNR↑	SSIM↑	FID↓	LPIPS↓	MUSIQ↑	PSNR↑	SSIM↑	FID↓	LPIPS↓	MUSIQ↑
FFHQ (768 × 768)																				
ReSample	24.65	0.708	98.07	0.196	32.06	22.96	0.632	172.19	0.357	21.09	25.54	0.747	96.62	0.216	31.65	26.37	0.746	97.47	0.160	31.50
FlowChef	27.53	0.759	57.24	0.147	49.05	26.38	0.731	110.15	0.209	38.70	24.88	0.716	63.48	0.237	39.22	27.30	0.754	46.40	0.152	44.26
FlowDPS	27.92	$\underline{0.772}$	$\underline{23.80}$	0.120	$\underline{54.36}$	26.84	0.745	30.54	$\underline{0.156}$	48.33	25.15	0.721	43.18	0.222	48.51	26.02	0.731	45.00	0.204	47.11
FLAIR	$\underline{28.88}$	0.768	55.51	0.123	52.26	27.25	0.752	$\underline{29.31}$	0.158	46.63	28.80	0.695	21.57	0.095	52.48	28.60	0.729	**18.41**	0.090	$\underline{54.71}$
TriPS_T (Ours)	**29.03**	**0.789**	26.66	$\underline{0.113}$	53.65	**27.45**	$\underline{0.754}$	35.13	0.161	$\underline{52.63}$	**31.20**	$\underline{0.809}$	$\underline{17.28}$	$\underline{0.060}$	$\underline{63.52}$	**29.95**	**0.804**	25.02	$\underline{0.084}$	51.12
TriPS_G (Ours)	28.55	0.762	**22.18**	**0.107**	**61.69**	$\underline{27.38}$	**0.762**	**28.22**	**0.154**	**53.92**	**31.20**	**0.813**	**15.89**	**0.059**	**64.37**	$\underline{29.60}$	$\underline{0.782}$	$\underline{21.21}$	**0.074**	**61.92**
DIV2K (768 × 768)																				
ReSample	20.55	0.535	75.67	0.238	26.82	19.54	0.459	119.22	0.372	20.57	21.47	0.588	61.79	0.222	31.17	21.74	0.559	80.79	0.221	24.22
FlowChef	22.08	0.561	47.47	0.213	41.70	20.88	0.508	84.28	0.297	35.94	19.62	0.482	74.01	0.366	41.24	21.57	0.532	52.84	0.251	39.52
FlowDPS	22.14	0.545	35.18	0.175	39.87	21.07	0.469	48.36	**0.251**	33.47	19.88	0.473	52.23	0.322	40.01	20.46	0.473	58.56	0.307	36.80
FLAIR	$\underline{22.90}$	0.592	41.23	0.167	42.30	21.12	$\underline{0.520}$	$\underline{42.16}$	0.256	46.23	23.90	0.614	22.17	0.129	52.24	22.70	0.561	32.26	0.157	$\underline{44.96}$
TriPS_T (Ours)	**23.05**	**0.607**	$\underline{31.80}$	**0.158**	$\underline{45.24}$	**21.27**	0.518	43.74	$\underline{0.255}$	$\underline{50.14}$	**26.29**	**0.728**	$\underline{15.49}$	**0.066**	$\underline{59.16}$	$\underline{23.94}$	**0.646**	$\underline{28.57}$	$\underline{0.123}$	44.47
TriPS_G (Ours)	22.78	$\underline{0.594}$	**27.84**	$\underline{0.163}$	**50.14**	$\underline{21.13}$	**0.531**	**37.48**	0.257	**52.26**	$\underline{26.19}$	$\underline{0.715}$	**14.94**	**0.066**	**59.89**	**23.97**	$\underline{0.644}$	**26.88**	**0.121**	**45.19**

ReSample: [Song, ICLR 2024] Solving Inverse Problems with Latent Diffusion Models via Hard Data Consistency
FlowChef: [Patel, ICCV 2025] FlowChef: Steering of Rectified Flow Models for Controlled Generations
FlowDPS: [Kim, ICCV 2025] FlowDPS : Flow-Driven Posterior Sampling for Inverse Problems
FLAIR: [Erbach, NeurIPS 2025] Solving Inverse Problems with FLAIR

Qualitative results — Linear Inverse Problems

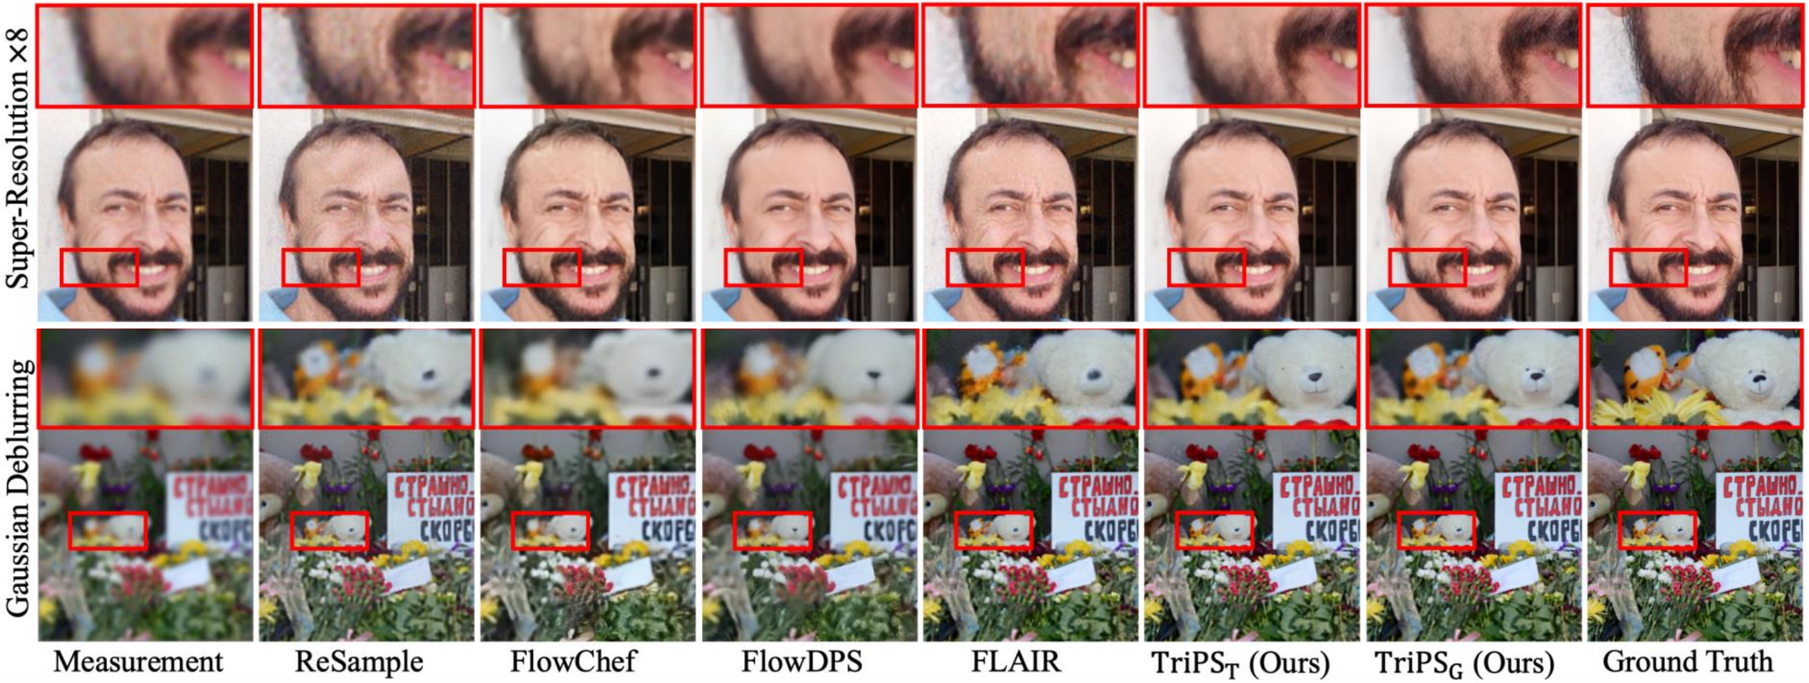

Reward-guided Perception–Distortion Control

$$R = w_{\text{dist}} R_{\text{dist}} + w_{\text{perc}} R_{\text{perc}}$$

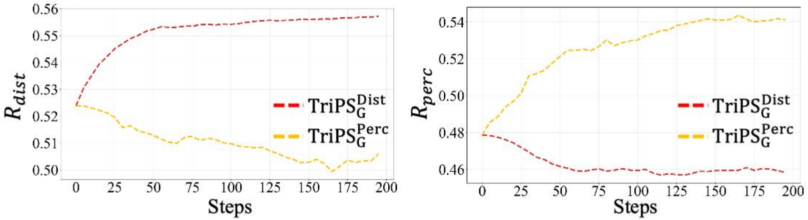

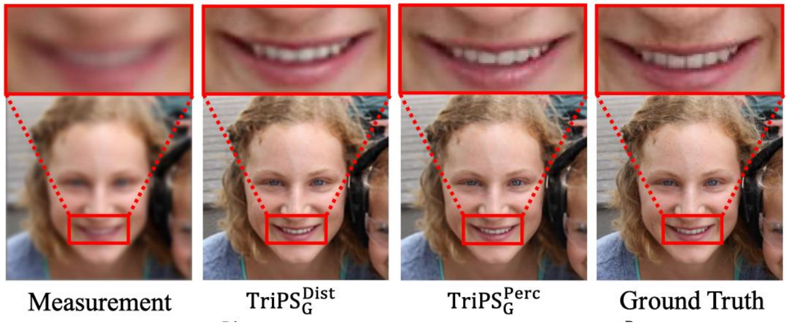

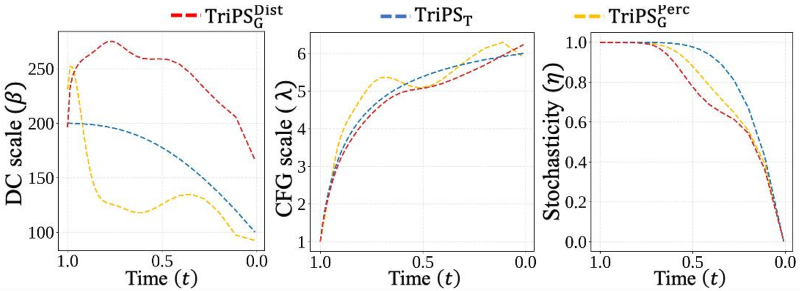

(a) Reward Guided Perception-Distortion Controller

Case	PSNR↑	SSIM↑	KID↓	LPIPS↓
$\text{TriPS}_G^{\text{Dist}}$	**30.62**	**0.840**	0.036	0.072
TriPS_T	30.33	0.822	0.016	0.070
$\text{TriPS}_G^{\text{Perc}}$	29.89	0.793	**0.011**	**0.067**

- Tuning the reward weights slides TriPS along the perception–distortion frontier

Conclusion & Limitation

Contributions

- Analysis of Triadic coupling dynamics
  - DC↔CFG conflict at early timesteps, stochasticity regularizes the trajectory.
- Triadic schedule optimization framework
  - TriPS_T template search + TriPS_G GRPO for fine temporal curves.

Limitations

- Additional computational overhead
  - TriPS_T and TriPS_G adds calibration / optimization cost over fixed schedules.

Thank You!

For the paper, code, and more results, scan the QR or visit:
https://mundongju.github.io/TriPS/

Acknowledgements: This work was supported in part by Institute of Information & Communications Technology Planning & Evaluation (IITP) grants funded by the Korea government(MSIT) [No. RS-2021-II211343, Artificial Intelligence Graduate School Program (Seoul National University) / No. RS-2025- 02314125, Effective Human-Machine Teaming With Multimodal Hazy Oracle Models], the National Research Foundation of Korea (NRF) grant funded by the Korea government(MSIT) (No. RS-2025-02263628), the BK21 FOUR program of the Education and Research Program for Future ICT Pioneers, Seoul National University, and Samsung Electronics Co., Ltd. (IO251217-14748-01).

ICML
International Conference
On Machine Learning