Goal Conditioned Reinforcement Learning for Photo Finishing Tuning

Jiarui Wu[1,2], Yujin Wang[1], Lingen Li[1,2], Zhang Fan[1], Tianfan Xue[2,1]

[1]Shanghai AI Laboratory, [2]The Chinese University of Hong Kong

Project website

What is photo finishing tuning?

Photo Tuning: Given the *input image*, *photo finishing pipeline*, and the *tuning target*, automatic find the best set of parameters that produce the target rendering style.

Main idea: reinforcement learning to reduce tuning step

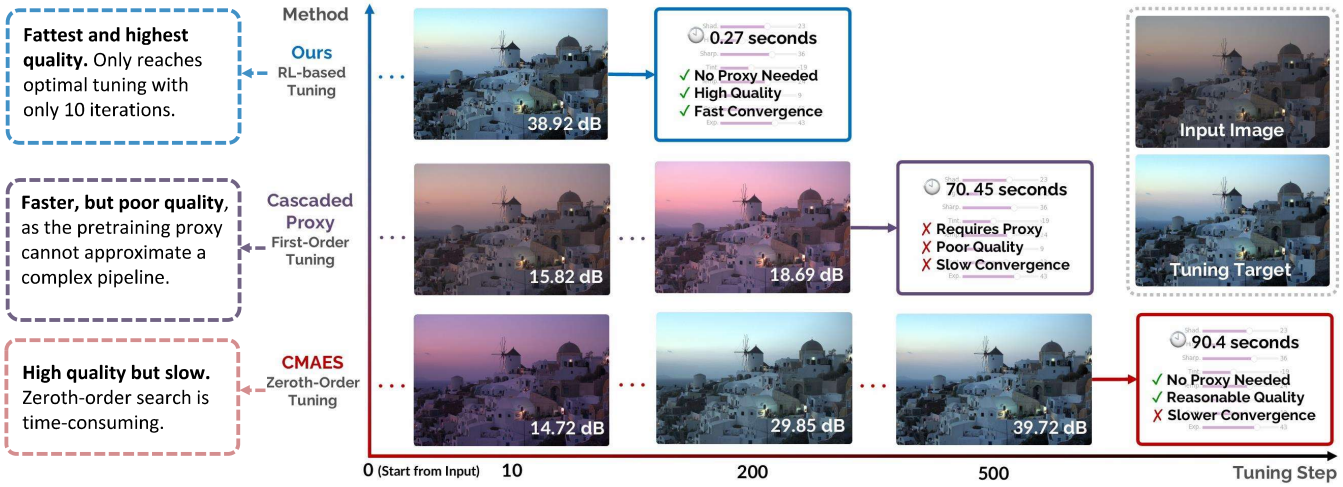

Method

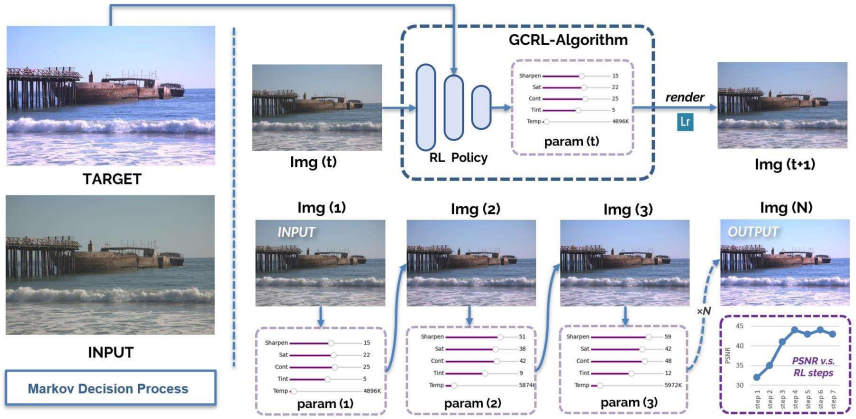

For more details regarding the our RL framework, please check out our paper.

Goal Conditioned Reinforcement Learning based Photo Tuning:

- By forming the problem into a Markov Decision Process and training the RL policy, we develop a smart searching algorithm that brings results closer to the target at each step.
- With a learned policy, our algorithm predicts the search direction more accurately than zeroth-order methods, without relying on a proxy.
- We also design a state representation to better model the relationship between the photo editing space and our policy network. See our paper for details.

Experimental results

Qualitative Results. Ours generalize well to unseen dataset.

	FiveK Test Set			HDR+ *(cross-dataset generalization)*		
Method	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS
CMAES	28.53	0.9586	0.0968	28.08	0.9539	0.1307
Monolithic Proxy	21.71	0.9104	0.2144	17.80	0.8940	0.3044
Cascaded Proxy	22.31	0.9115	0.1939	18.90	0.8982	0.2797
Ours	**35.89**	**0.9764**	**0.0305**	**31.54**	**0.9652**	**0.0563**

User study results (No. of votes of each methods).

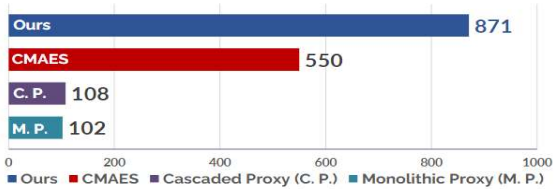

Qualitative Comparison of Photo Finishing Tuning on HDR+ dataset

Qualitative Comparison of Photo Stylization Tuning task