Generic Image Restoration with Flow Based Priors

Leonhard Helminger[1] Michael Bernasconi[1] Abdelaziz Djelouah[2] Markus Gross[1,2] Christopher Schroers[2]
[1]ETH Zürich / [2]DisneyResearch|Studios

Motivation

- Recent methods proposed alternatives and approximations to the true image prior
- With normalizing flows, we have an approach for a tractable *and* exact log-likelihood computation
- The goal is, to learn a distribution of target high quality content to serve as a prior in the MAP formulation.

Contributions

- Our work is the first that uses normalizing flows to learn a prior for generic image restoration
- We take advantage of the bijective mapping learned by our model to express the MAP problem of image reconstruction in latent space
- We propose to regularize the base distribution space with new loss terms, which yield a better behavior during the MAP inference

Overview

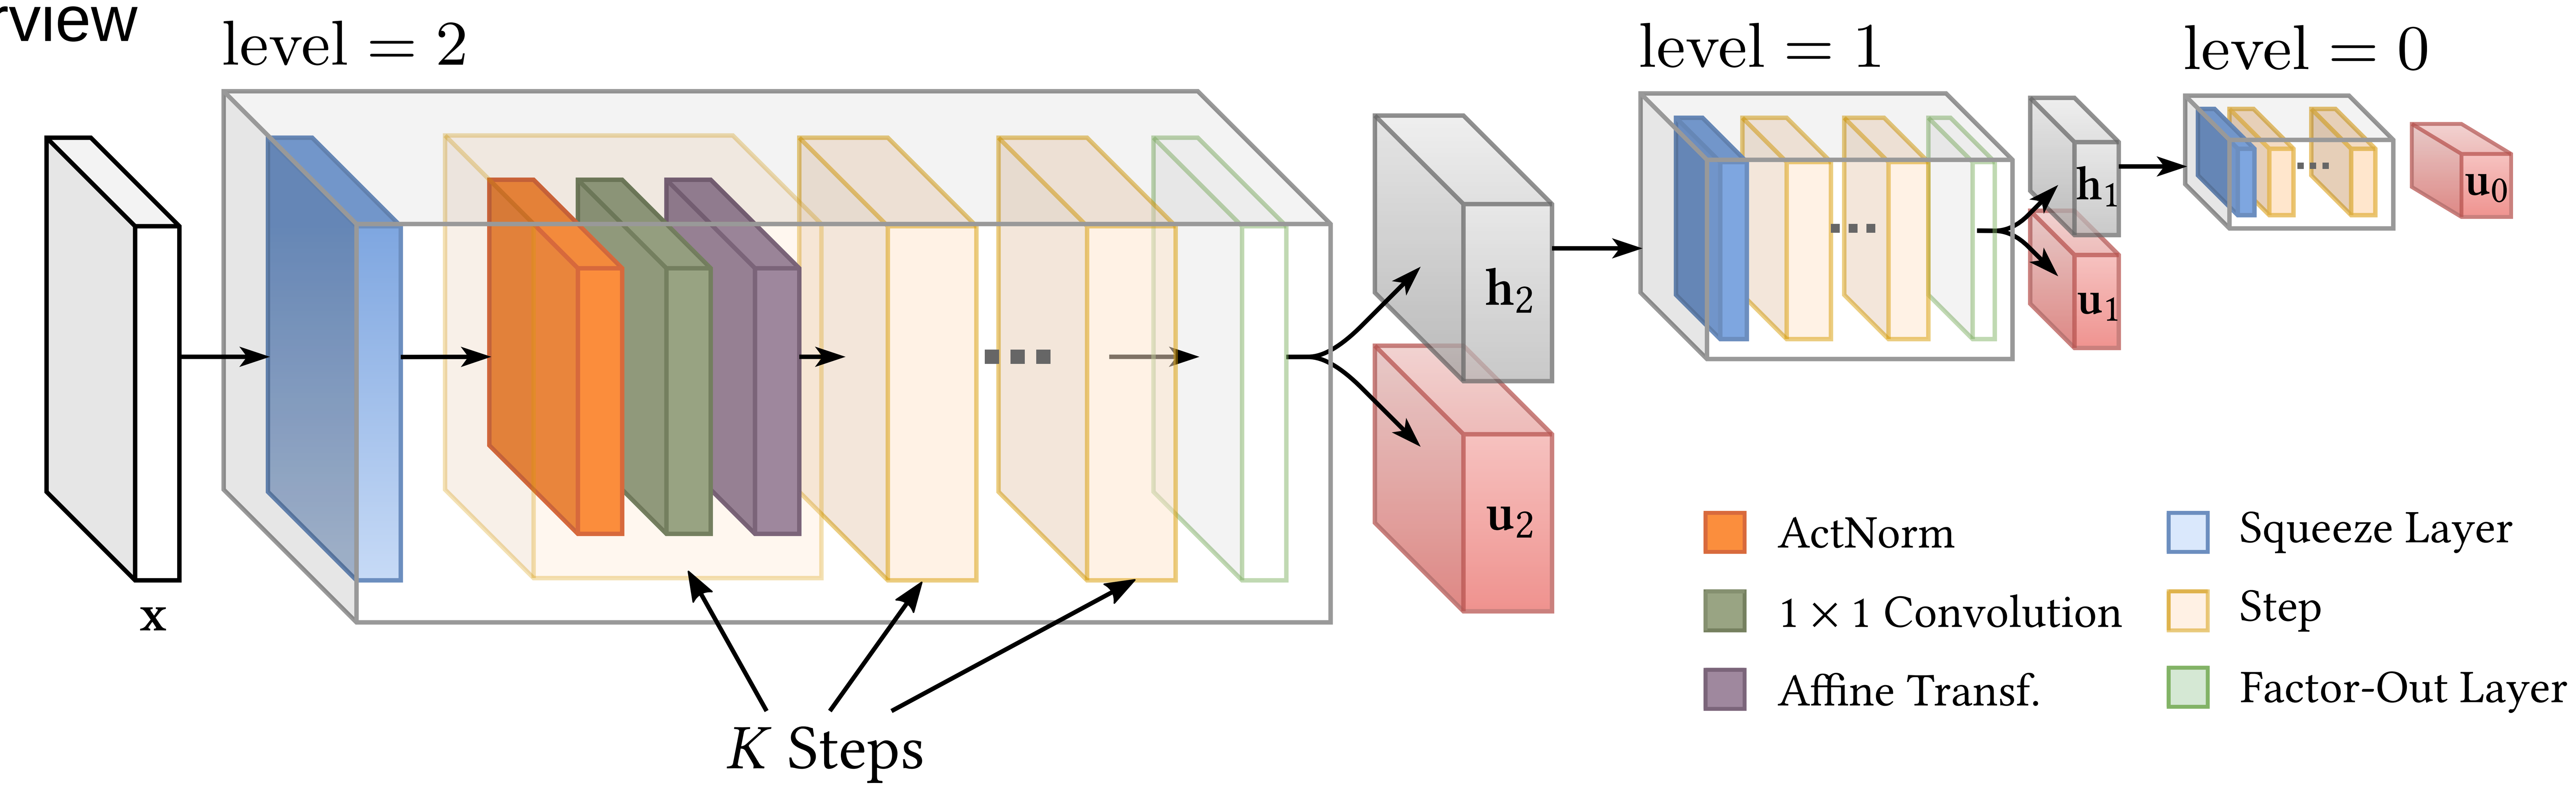

Phase 1: Training the flow based prior

$$\mathcal{L} = \mathcal{L}_{nll} + \beta_{ln}\mathcal{L}_{ln} + \beta_{ae}\mathcal{L}_{ae} + \underbrace{\beta_{in}\mathcal{L}_{in}}_{\text{optional}}$$

$$\text{where } \mathcal{L}_{nll} = -\log p_\theta(\mathbf{x} + \epsilon)$$

$$\mathcal{L}_{ln} = ||T_\theta\left(T_\theta^{-1}(\mathbf{x}) + \xi\right) - \mathbf{x}||_2^2$$

$$\mathcal{L}_{in} = ||T_\theta^{-1}(\mathbf{x}) - T_\theta^{-1}(\mathbf{x} + \eta)||_2^2$$

Phase 2: Image restoration using the prior

$$\mathbf{u}^\star = \arg\min_{\mathbf{u}} \Big[\underbrace{-\log p\left(\hat{\mathbf{x}} \,|\, T_\theta(\mathbf{u})\right)}_{\text{data}} \underbrace{-\log p_\theta\left(T_\theta(\mathbf{u})\right)}_{\text{prior}}\Big]$$

$$\text{with } \mathbf{x}^\star = T_\theta(\mathbf{u}^\star)$$

Figure 1: Our approach is able to remove the degradation and produces visually more pleasing results.

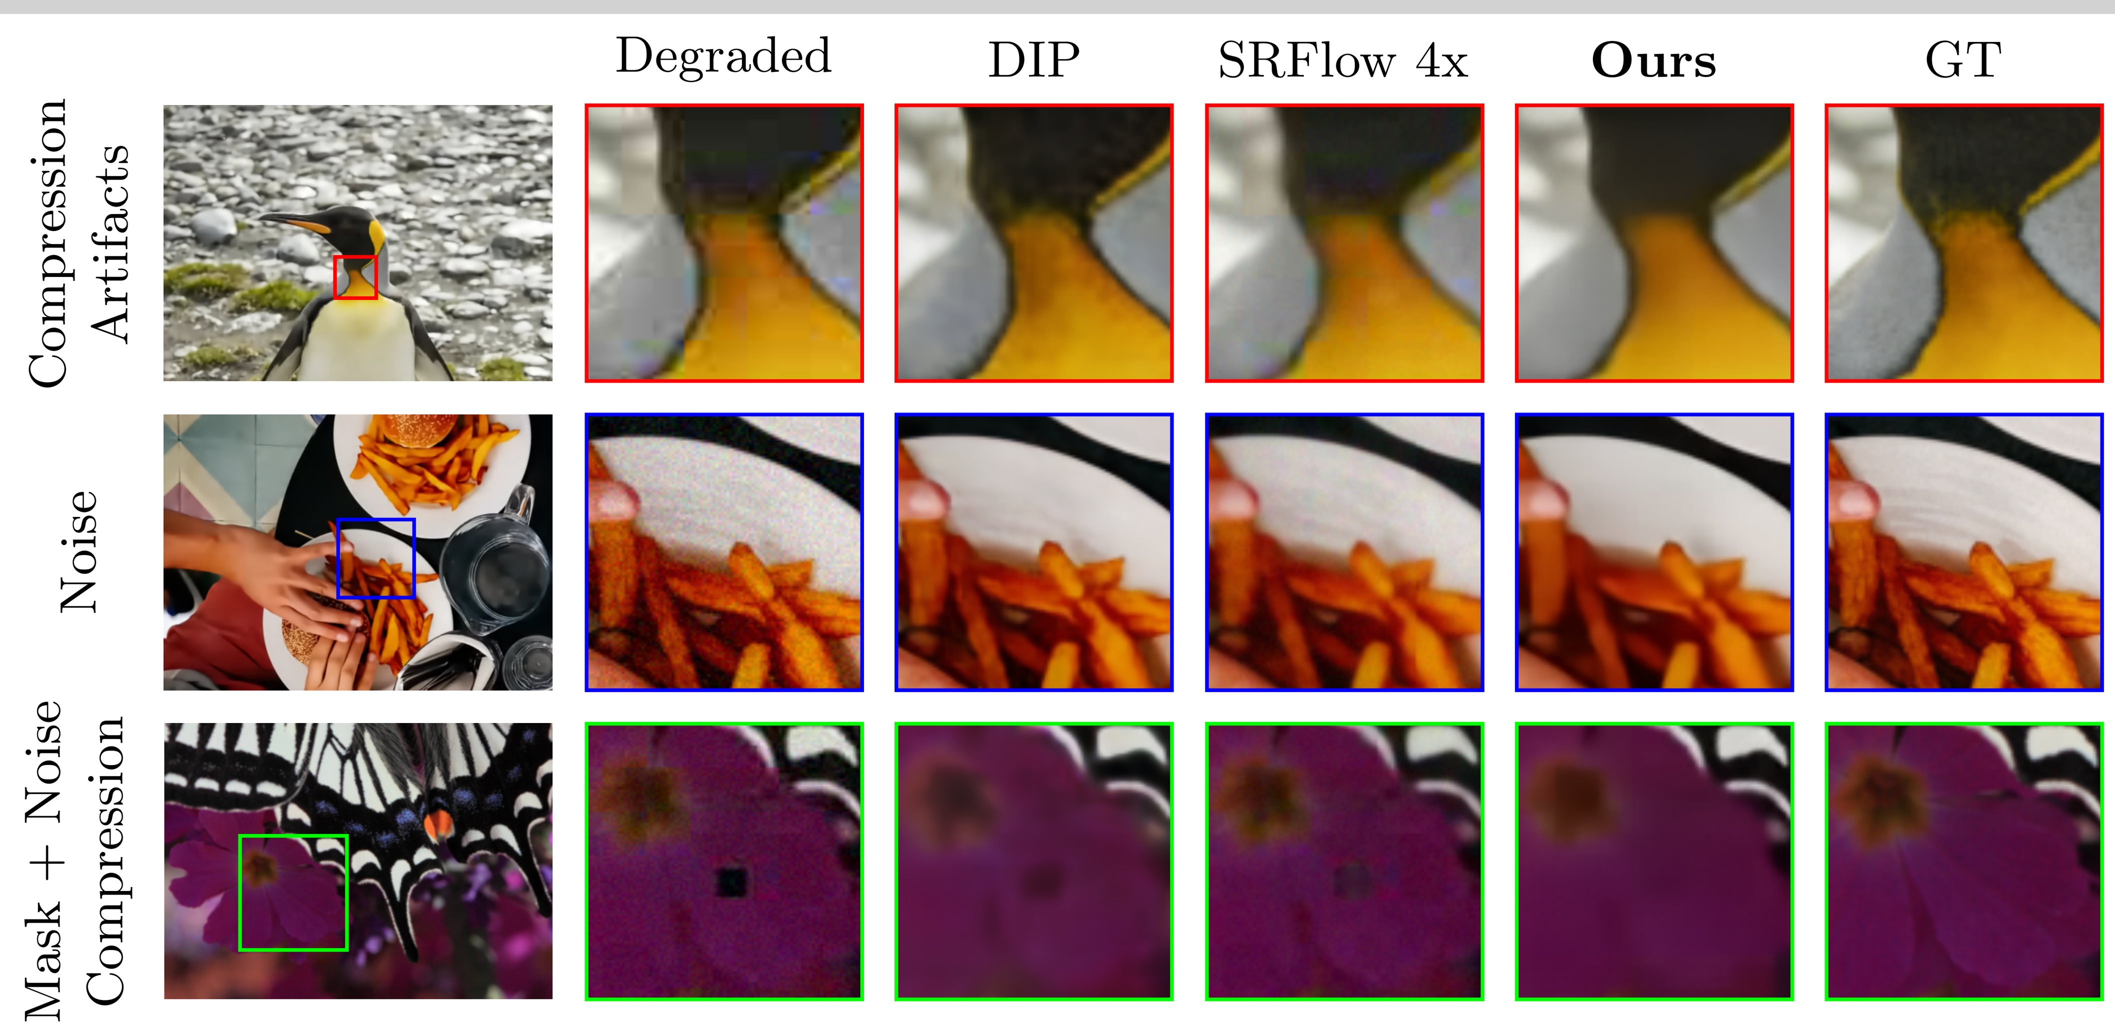

Figure 2: Results on DIV2K dataset. The proposed prior is used to restore images of arbitrary size.

		DIP	Double-DIP	SRFlow 4x	SRFlow 8x	Ours
JPEG	PSNR↑	28.16	-	27.75	25.27	**30.29**
	SSIM↑	0.85	-	0.80	0.71	**0.86**
	MSSSIM↑	**0.97**	-	0.95	0.91	0.96
	LPIPS↓	**0.17**	-	0.25	0.31	0.23
Noise	PSNR↑	**30.22**	-	27.32	24.73	28.99
	SSIM↑	**0.92**	-	0.81	0.69	0.87
	MSSSIM↑	**0.98**	-	0.96	0.92	0.96
	LPIPS↓	**0.07**	-	0.23	0.34	0.21
Multi Degr.	PSNR↑	26.41	27.62	27.57	25.21	**29.87**
	SSIM↑	0.78	0.80	0.78	0.69	**0.85**
	MSSSIM↑	0.92	0.94	0.94	0.91	**0.96**
	LPIPS↓	0.26	0.27	0.26	0.34	**0.23**

Table 1. Quantitative evaluation on DIV2K and comparison with state of the Art.