Compression with Bayesian Implicit Neural Representations

Zongyu Guo[*,1], Gergely Flamich[*,2], Jiajun He[2], Zhibo Chen[1], José Miguel Hernández Lobato[2]
[*]equal contribution, [1]University of Science and Technology of China, [2]University of Cambridge

Paper

TL;DR: We propose to compress data as variational Bayesian implicit neural representations, which supports joint rate-distortion optimization.

Motivation

Data compression with INRs:

1. Fit INR to data
2. Quantize weights
3. Encode quantized weights

Issues:

1. Overfitting ⇒ brittle weights
2. Quantization degrades performance
3. Cannot jointly optimize rate-distortion

Solution

Use variational Bayesian INRs!

Objective:

$$\mathcal{L}_\beta(\mathcal{D}, q_{\mathbf{w}}, p_{\mathbf{w}}) = \underbrace{\sum_{(\boldsymbol{x},\boldsymbol{y})\in\mathcal{D}} \mathbb{E}_{q_{\mathbf{w}}}[\Delta(\boldsymbol{y}, f(\boldsymbol{x} \mid \boldsymbol{w}))]}_{\text{distortion}} + \beta \underbrace{D_{\mathrm{KL}}[q_{\mathbf{w}} \| p_{\mathbf{w}}]}_{\text{rate}}$$

Coordinate Descent:

1. Optimize variational posteriors

$$q_{\mathbf{w}}^{(i)} = \arg\min_q \mathcal{L}_\beta(\mathcal{D}_i, q, p_{\mathbf{w};\boldsymbol{\theta}_p})$$

for $i = 1, \ldots, M$.

2. Update prior parameters

$$\boldsymbol{\mu}_p = \frac{1}{M}\sum_{i=1}^{M} \boldsymbol{\mu}_q^{(i)}$$

$$\boldsymbol{\sigma}_p = \frac{1}{M}\sum_{i=1}^{M} [\boldsymbol{\sigma}_q^{(i)} + (\boldsymbol{\mu}_q^{(i)} - \boldsymbol{\mu}_p)^2].$$

COMBINER: COMpression with Bayesian Implicit NEural Representations

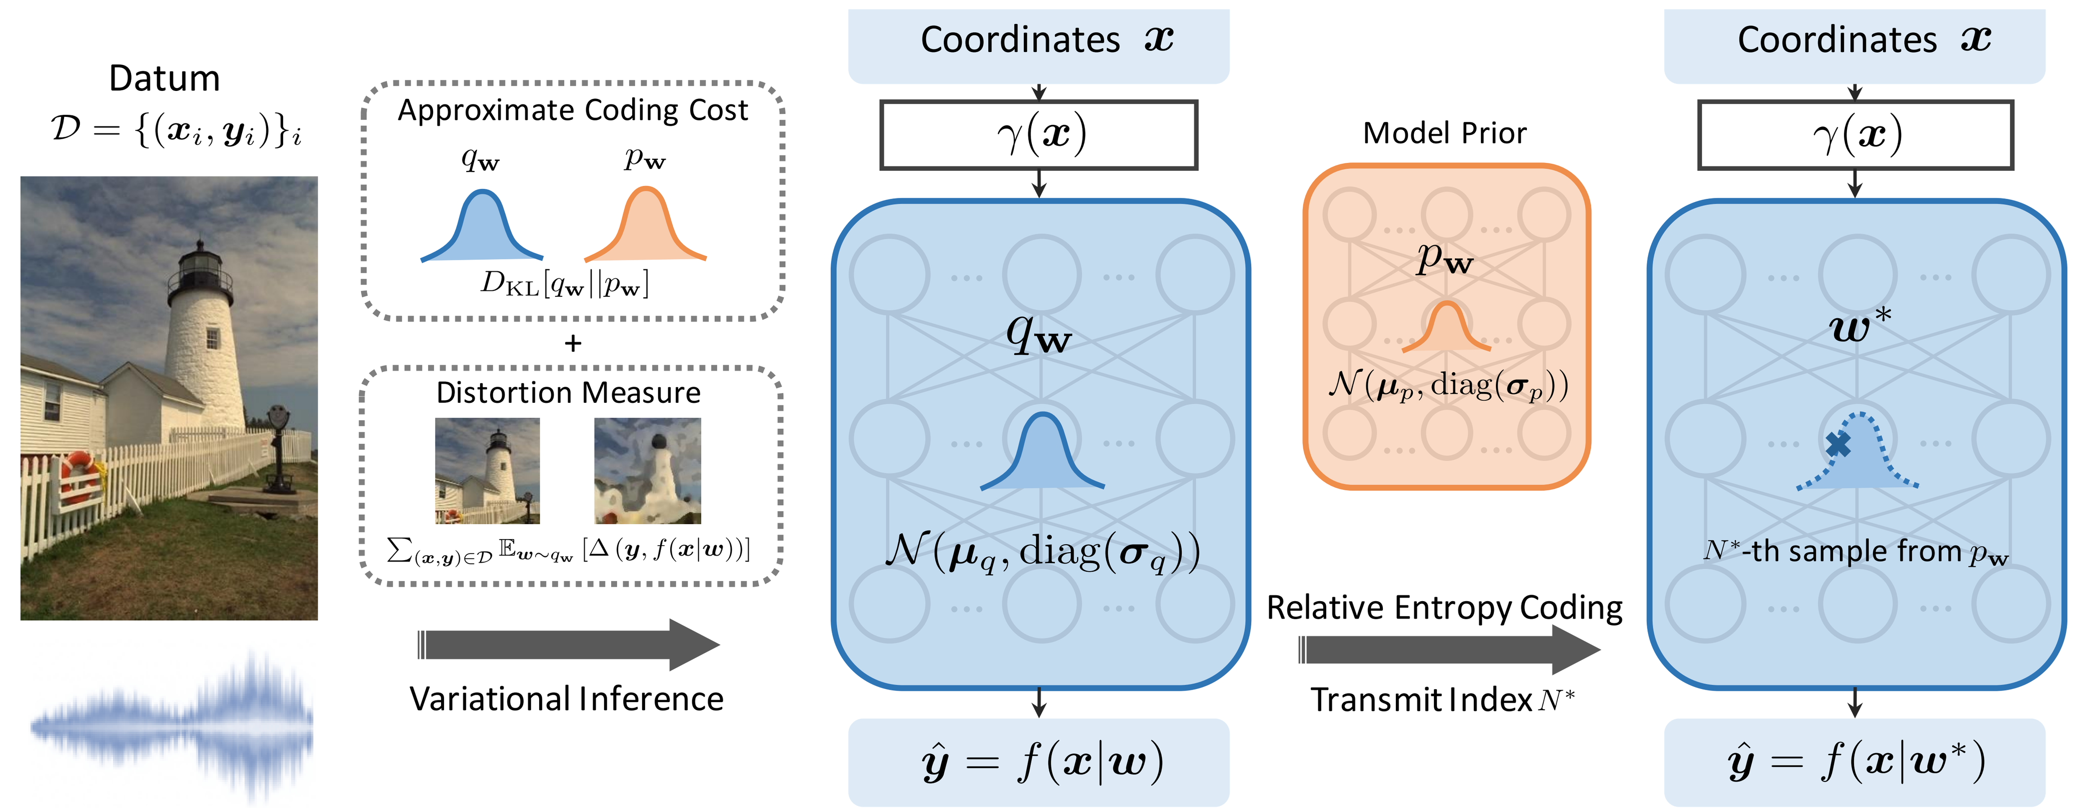

Image Compression

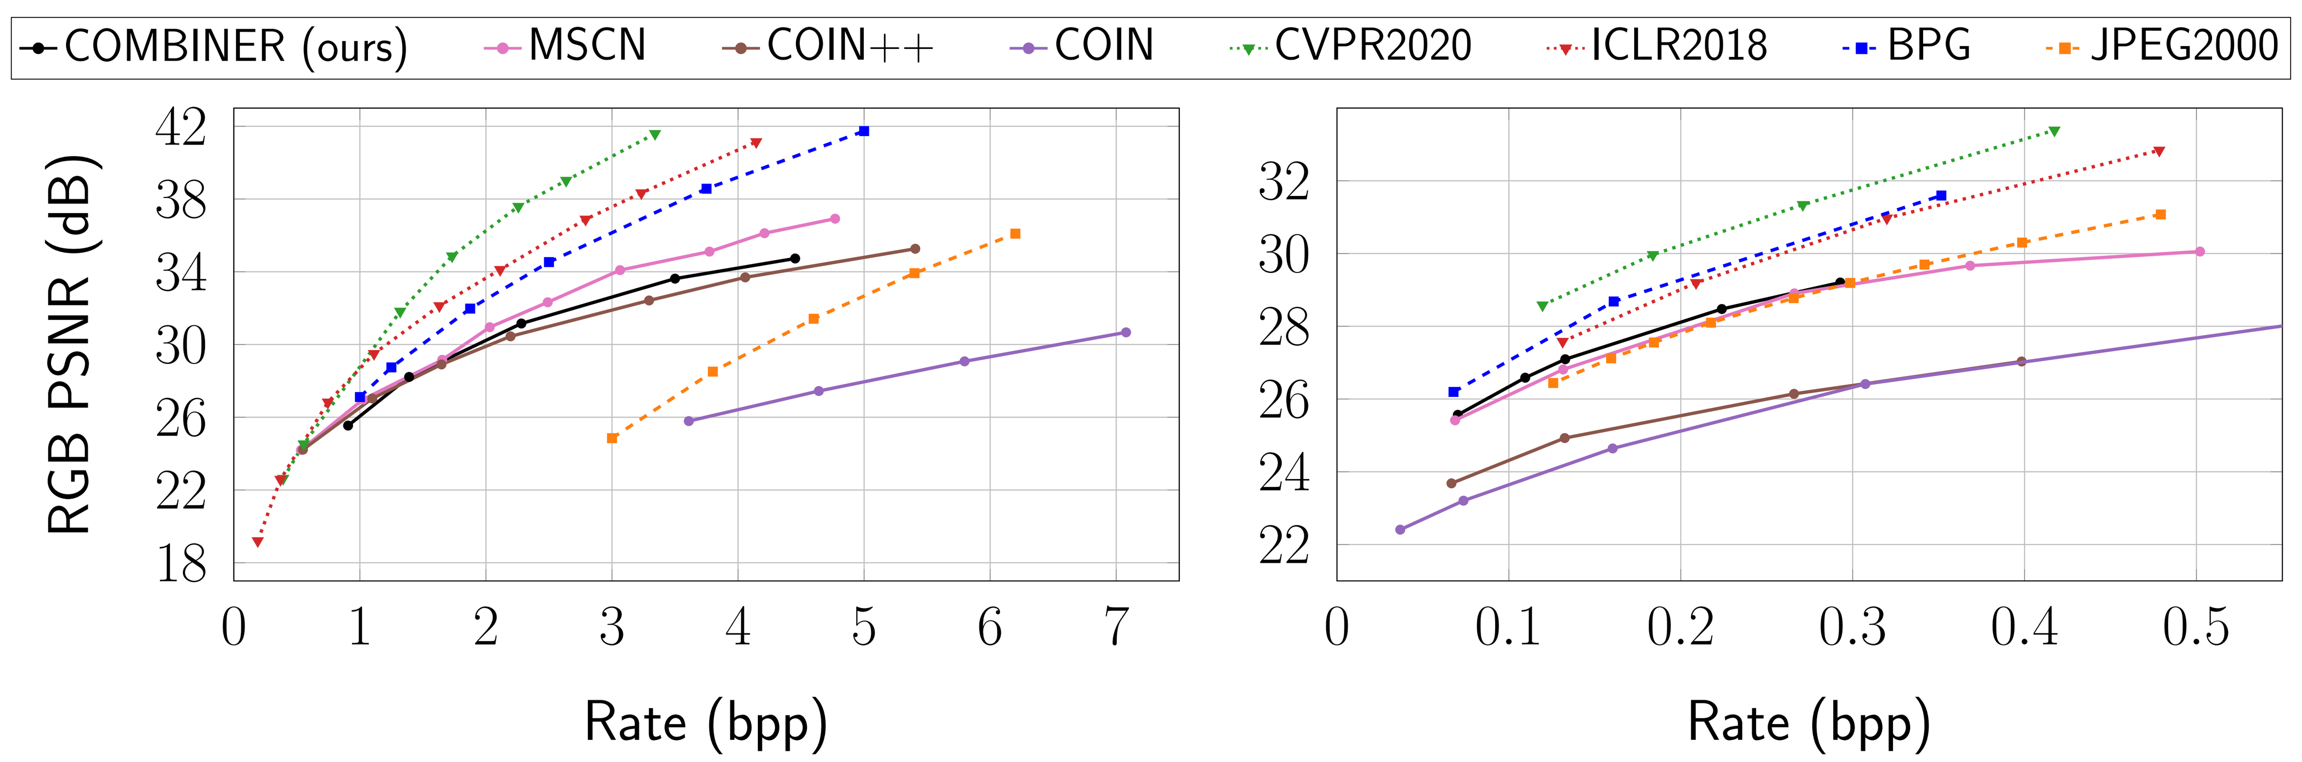

(a) CIFAR-10 dataset **(b) Kodak dataset**

bit-rate	Encoding (1 image, GPU A100 80G)			Decoding (1 image, CPU)
	Learning Posterior	REC + Fine-tuning	Total	
0.07 bpp		∼12 min 30 s	∼21 min 30 s	348.42 ms
0.11 bpp	∼9 min	∼18 mins	∼27 min	381.53 ms
0.13 bpp		∼22 min	∼31 min	405.38 ms
0.22 bpp	∼11 min	∼50 min	∼61 min	597.39 ms
0.29 bpp		∼68 min	∼79 min	602.32 ms

The encoding time and decoding time of COMBINER on Kodak dataset.

Ablation Study

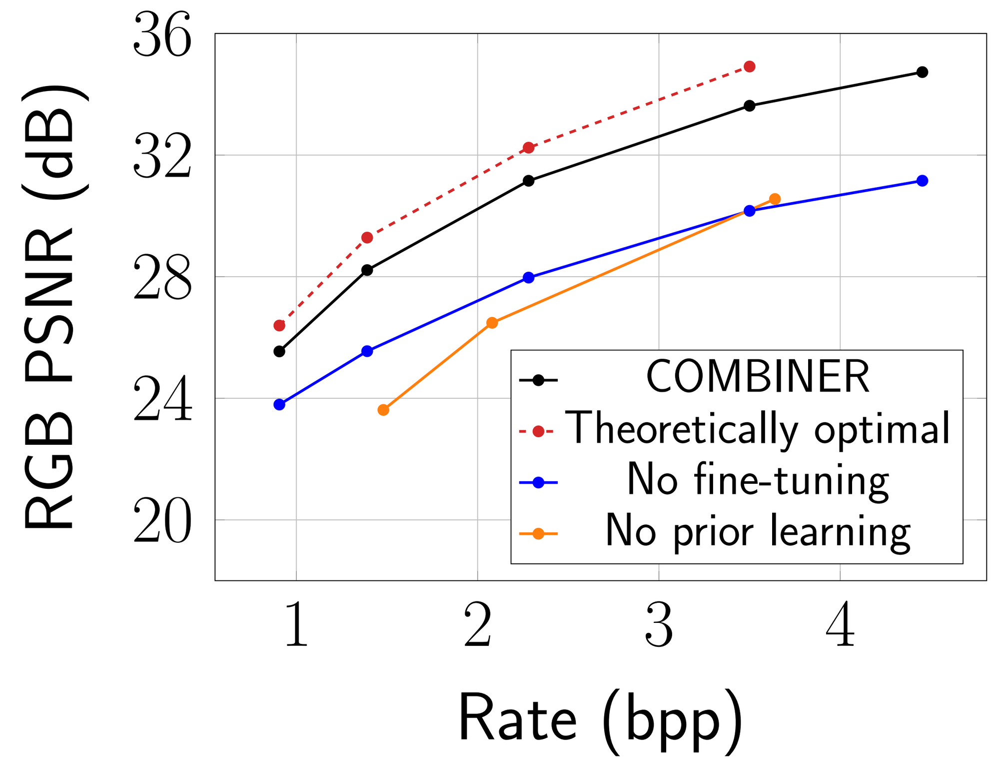

Audio Compression

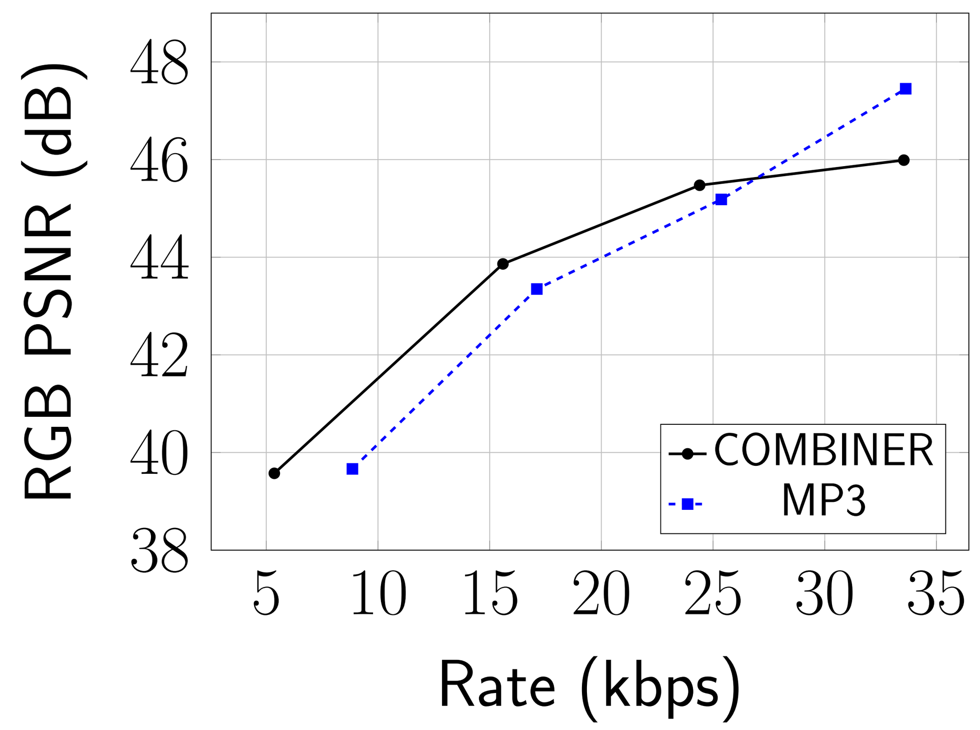

Fine-tuning