ETNA: EnTropy regularization for compressed Neural Avatar

Guillaume Loranchet[1,2], João Regateiro[1], Bharath Bhushan Damodaran[1], Franck Galpin[1], Adnane Boukhayma[2], Franck Multon[2], Pierre Hellier[2]

[1] InterDigital, Rennes, France [2] Univ Rennes, Inria, CNRS, IRISA, Rennes, France

BMVC 2026

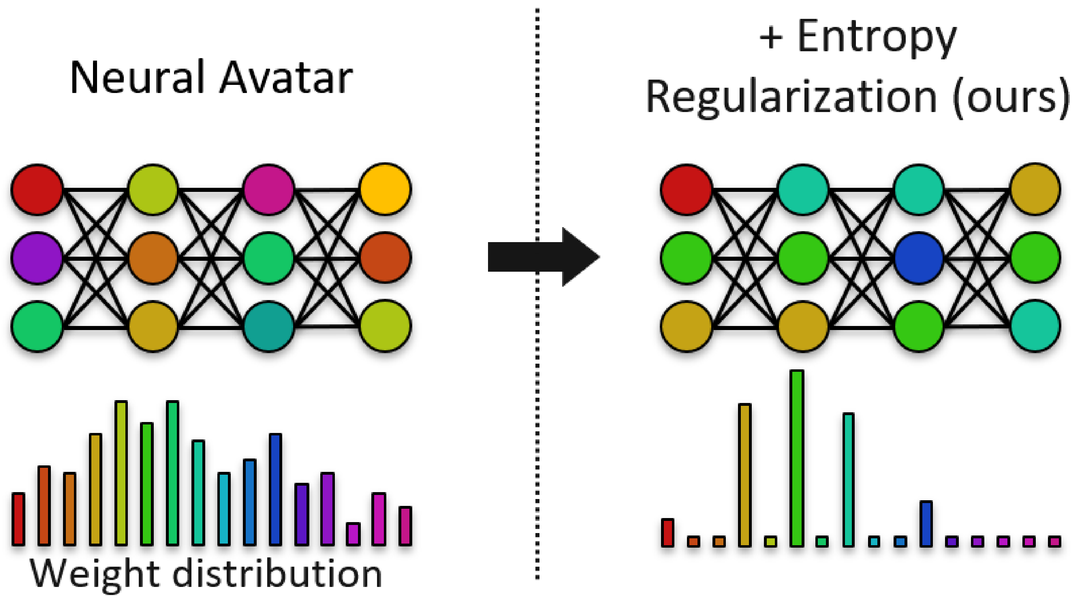
Introduction

ETNA trains compressible neural avatar weights while preserving reconstruction quality.

Standard quantization can introduce geometric artifacts at lower bitrates. ETNA shapes the weight distribution during training for efficient entropy coding.

Compact checkpoints reduce storage and transmission costs for shared virtual environments.

Regularization that leaves the forward pass unchanged.

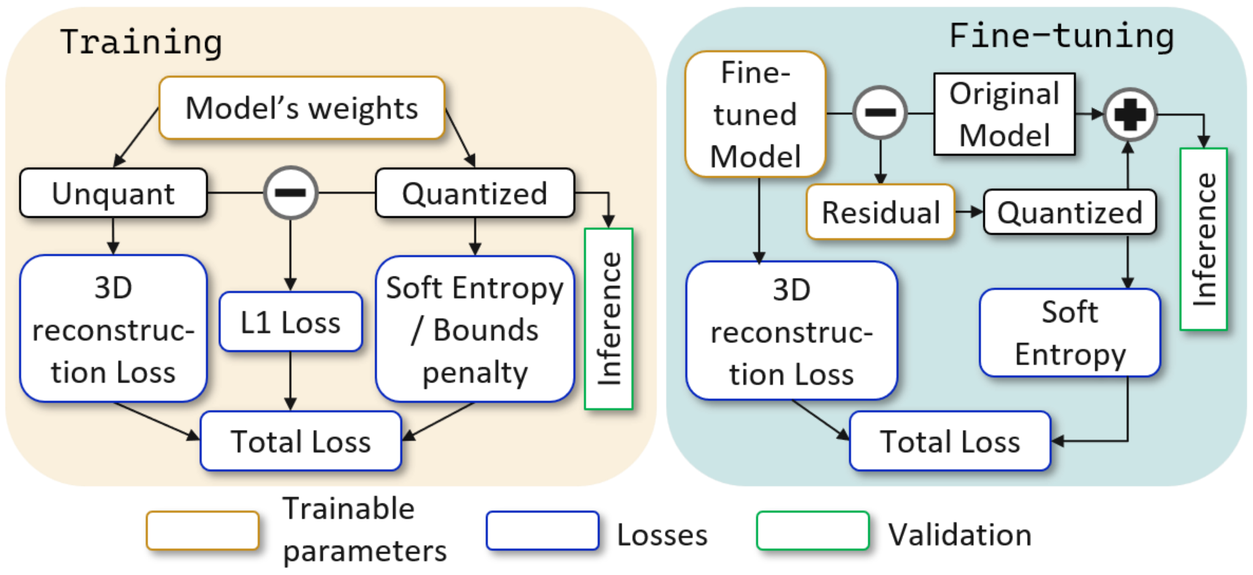


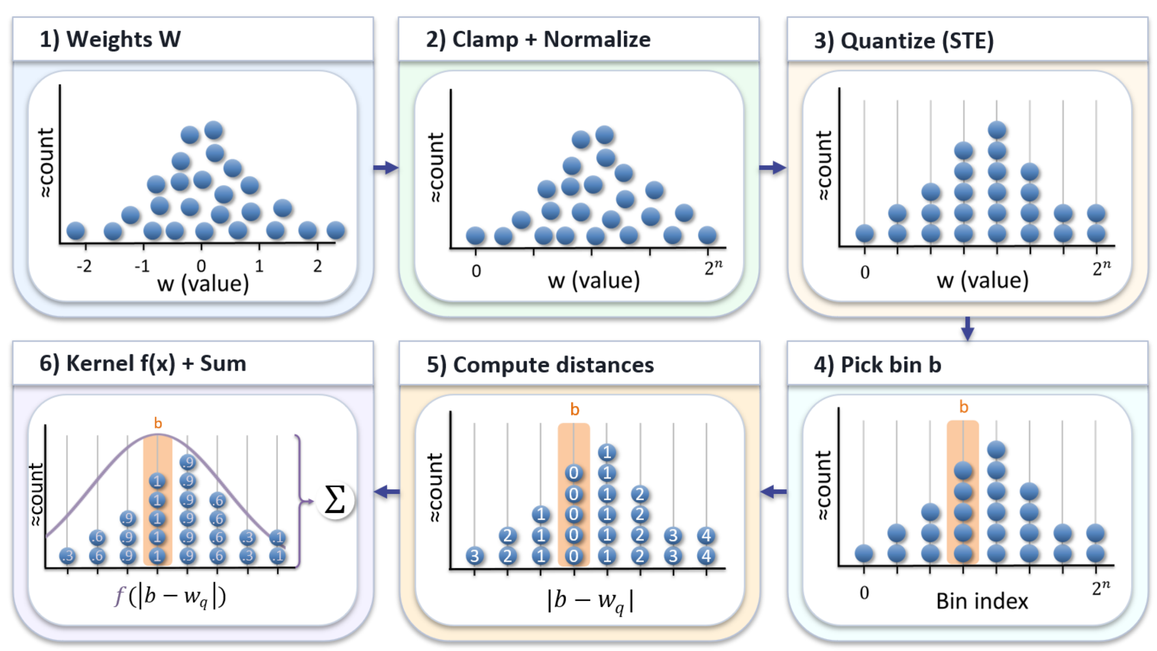
Method

1. Soft entropy A differentiable soft histogram encourages a low-entropy weight distribution.

2. Alignment and bounds L1 alignment and fixed bounds protect the quantized reconstruction.

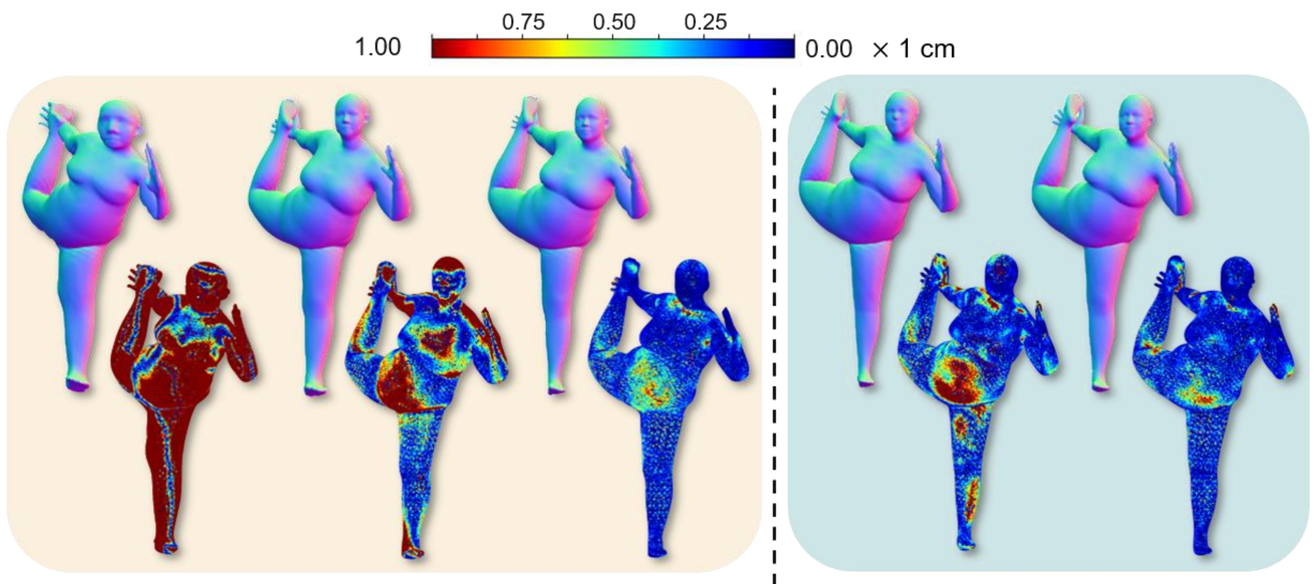
3. Stable training Start entropy regularization halfway through training. Stabilize with stochastic weight averaging.





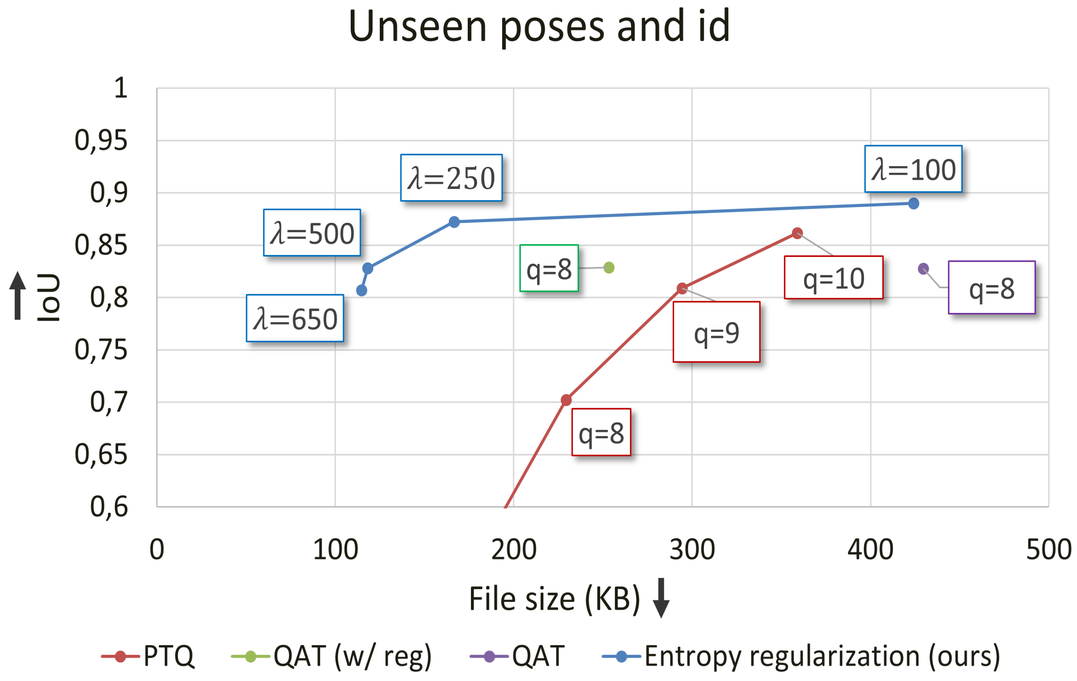
Training and fine-tuning (left). Differentiable soft-histogram construction (right).

Results



2.2× smaller checkpoints

Unseen poses and id



331 KB to 141 KB at comparable IoU ≈ 0.84