

Supplementary material: Slimmable Compressive Autoencoders for Practical Neural Image Compression

Fei Yang^{1,2}, Luis Herranz², Yongmei Cheng¹, Mikhail G. Mozerov²

¹ School of Automation, Northwestern Polytechnical University, Xi'an, China

² Computer Vision Center, Universitat Autònoma de Barcelona, Barcelona, Spain

{fyang, lherranz, mozerov}@cvc.uab.es, chengym@nwpu.edu.cn

A. Example of λ -scheduling

Fig. A1 (top) illustrates the evolution of the RD curve from the initial naive training (in black) until the end of λ -scheduling (in red). Fig. A1 (bottom) illustrates the schedule for the λ s of the different subCAEs.




Figure A1: Evolution of the RD curve (top) and λ during the λ -scheduling phase. Naive training shown in black (top).

B. Additional visualizations

Fig. A2a shows an illustrative example of the reconstructed images and (quantized) latent representations obtained by the different subCAEs. The slimmable structure of the last layer of the encoder results in a latent representation also structured in five groups of channels (i.e. 1-48, 49-72, 73-96, 97-144 and 145-192). One representative channel per group is shown. The corresponding bit allocation is shown in Fig. A2b. Note that smaller subCAEs also allocate fewer bits in the first group of channels.

Fig. A2c illustrates progressive decoding when (quality)

scalability is enabled, compared with the default slimmable decoding (non scalable bitstreams). In this case each group of channels of $w = 192$ in Fig. A2b correspond to the base stream (channels 1-48) and four enhancement streams, which can be progressively combined to improve quality. However, there is a noticeable increase in both rate (explained by the bit allocation shown in Fig. A2b) and distortion with the quality scalable bitstream.

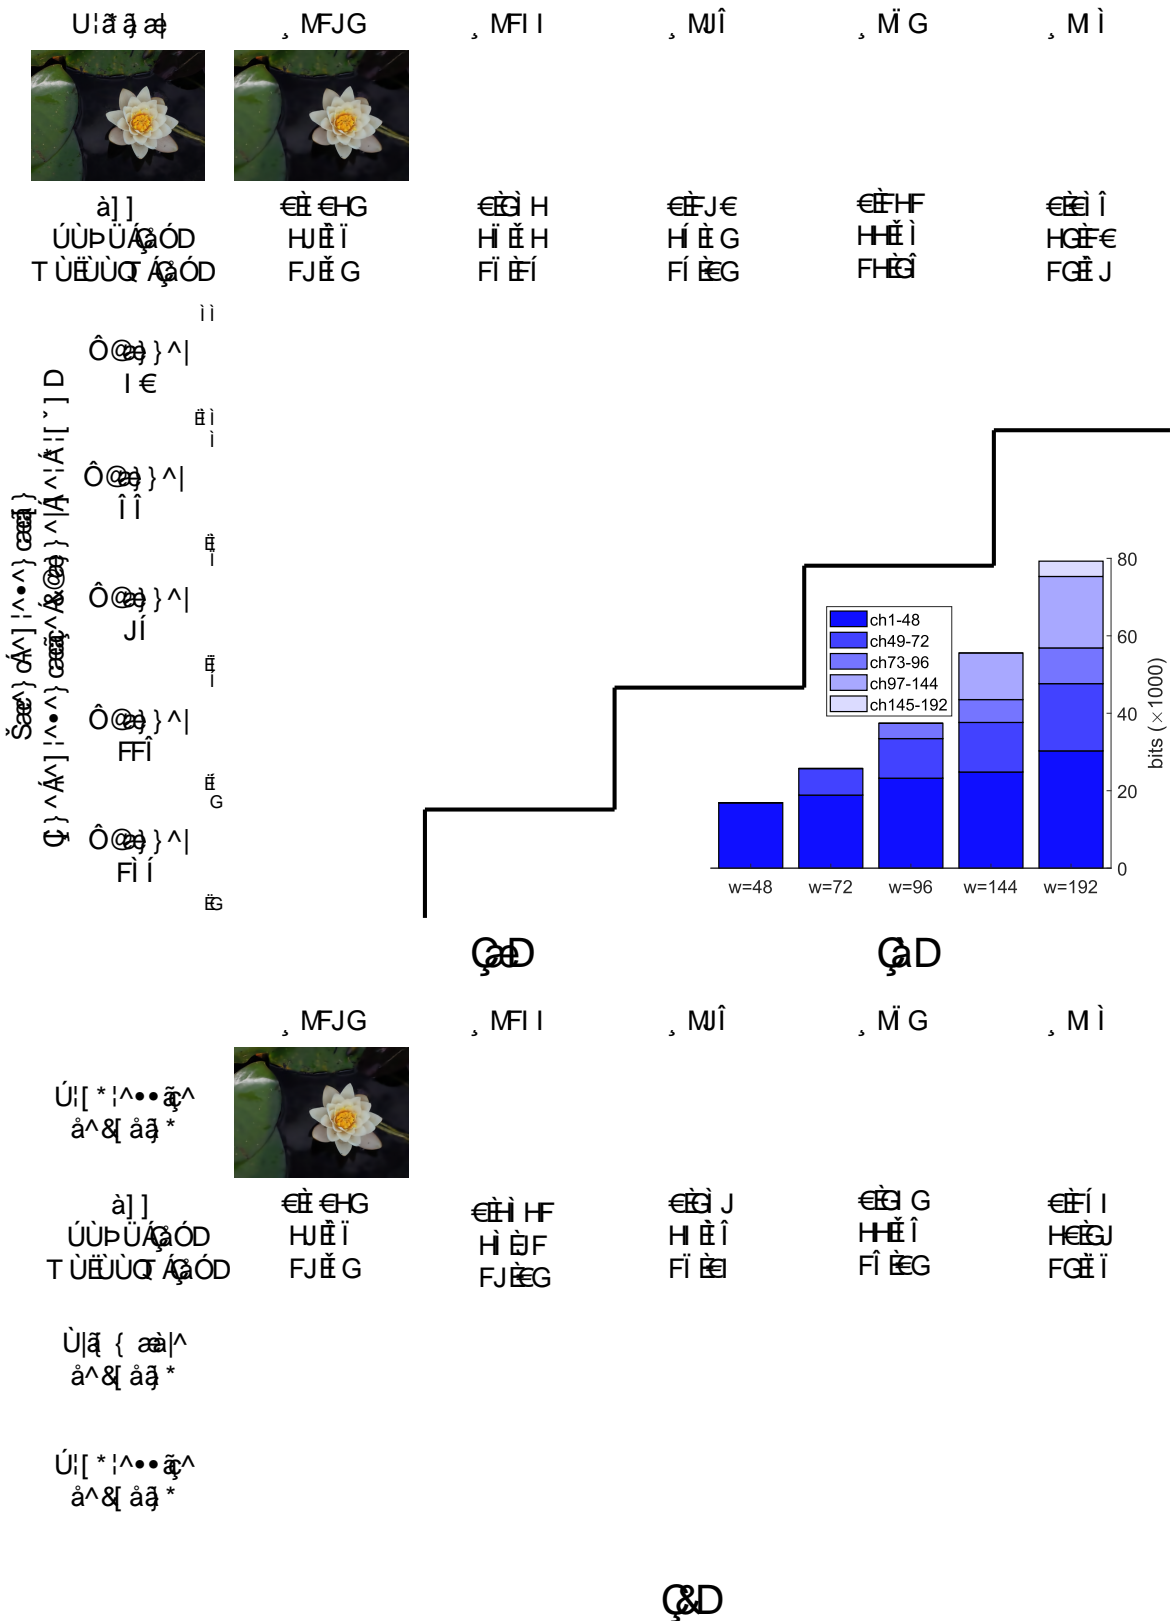


Figure A2: Illustrative example: (a, top) reconstructions using slimtable decoding, (a, bottom) selection of quantized latent maps (one per group of channels), (b) break down of bits spent in each group of channels, (c) scalable bitstream.