

Perception: The Next Milestone in Learned Image Compression

Dr. Johannes Ballé

Google

Since its emergence roughly 5 years ago, the field of learned data compression has attracted considerable attention. Using machine learning in source coding promises faster innovation cycles, as well as better adaptation to novel data modalities and nonlinear distortion metrics. For example, image codecs can now be end-to-end optimized to perform best for specific types of images, by simply replacing the training set. They may be designed to minimize a given perceptual metric, or in fact any differentiable perceptual loss function, without the need to evaluate it during encoding. However, as has been demonstrated, many existing perceptual models are not useful for end-to-end training. They do not generalize well enough, no matter how well they may predict human judgements on image quality datasets. In this talk, I will first give an overview of the current state of learned image compression, and then focus on what I consider the next big milestone: finding new and better ways to model visual perception.