Principal Bit Analysis: Autoencoding with Schur-Concave Loss

Sourbh Bhadane, Aaron B. Wagner, Jayadev Acharya,

Cornell University

Autoencoders and their Limitations

- Given a centered dataset $x_1, x_2 \cdots x_n \in \mathbb{R}^d$ with an empirical covariance matrix K, an **autoencoder** consists of an encoder $f : \mathbb{R}^d \mapsto \mathbb{R}^k$ and a decoder $g : \mathbb{R}^k \mapsto \mathbb{R}^d$. In a **linear autoencoder (LAE)**, f and g are linear maps.
- Conventional LAEs do not identify principal directions of the dataset.
- Autoencoders are sometimes described as "compressing" the data.

LAE with Schur-Concave constraint

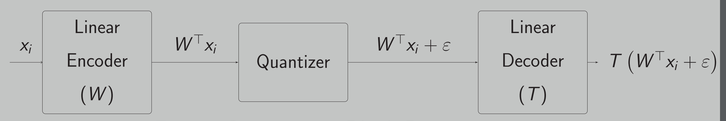

Figure: Compression Block Diagram

- Quantization assumptions:
  - Dither: If $Q()$ maps a real number to its nearest integer, then
    $$Q(z+\varepsilon) - \varepsilon \overset{d}{\sim} z + \varepsilon \quad \text{for } \varepsilon \sim \text{Unif}\left[-0.5, 0.5\right]$$
    .
  - To constrain number of bits, clip $w_j^\top x$ to the interval
    $$-\frac{\sqrt{(2a)^2 w_j^\top K w_j + 1}}{2} \qquad 0 \qquad \frac{\sqrt{(2a)^2 w_j^\top K w_j + 1}}{2}$$
  - Number of bits to represent $w_j^\top x$ is $\frac{1}{2}\log\left(4a^2 w_j^\top K w_j + 1\right)$.
- **Goal:** Minimize mean squared error (MSE) when $W^\top x$ is quantized, subject to a rate constraint on the number of bits required to represent quantized $W^\top x$.
- Observe $\left\{w_j^\top K w_j\right\}_{j=1}^d \mapsto \frac{1}{2}\log\left(4a^2 w_j^\top K w_j + 1\right)$ is **Schur-concave**.

Solve a more general nonconvex problem with any Schur-concave constraint ρ.

$$\begin{aligned} \inf_{W,T} \quad & \frac{1}{n}\sum_{i=1}^{n} \mathbb{E}_\varepsilon\left[\left\|x_i - T\left(W^\top x_i + \varepsilon\right)\right\|_2^2\right] \\ \text{subject to} \quad & R \geq \rho\left(\left\{w_j^\top K w_j\right\}_{j=1}^d\right). \end{aligned} \tag{1}$$

Theorem (Optimal LAE with Schur-concave constraint)

For Schur-concave $\rho : \mathbb{R}_{\geq 0}^d \to \mathbb{R}_{\geq 0}$ and $R > 0$, the set of matrices whose nonzero columns are eigenvectors of the covariance matrix K is optimal. If ρ is strictly Schur-concave and K contains distinct eigenvalues, this set contains all optimal solutions.

Remarks

- $W = US$, where U is eigenvector matrix and S is an unknown diagonal matrix.
- Let $\rho(x) = \sum_{i=1}^{d} \rho_{sl}(x_i)$, where $\rho_{sl} : \mathbb{R}_{\geq 0} \mapsto \mathbb{R}_{\geq 0}$.

Conventional LAE

- Conventional LAE: $W, T \in \mathbb{R}^{d \times k}$ where k is a parameter.
- Principal component analysis (PCA): $W = T = U_k$, where U_k is a matrix with top k eigenvectors as columns.
- PCA is a global optimal solution, but **not unique**.

We recover conventional LAEs by penalizing dimension,

$$\rho_{sl}(x) = 1\left[x > 0\right].$$

The optimal encoder is given by

$$S = \begin{bmatrix} \infty & & & & \\ & \ddots & & & \\ & & \infty & & \\ & & & 0 & \\ & & & & \ddots \\ & & & & & 0 \end{bmatrix}, \tag{2}$$

where top $\min\left(\lfloor R \rfloor, d\right)$ entries are ∞. Since latent variables are quantized in our formulation, PCA with parameter k is equivalent to S with top k diagonal entries equal to ∞.

Principal Bit Analysis (PBA)

Solve original problem by choosing

$$\rho_{sl}(x) = \frac{1}{2}\log\left(\frac{\gamma}{\sigma^2}x + 1\right)$$

- $\gamma = 1$: classical waterfilling solution
- $\gamma \in (1, 2]$: convex optimization problem
- $\gamma > 2$: **nonconvex optimization problem**

PBA takes as input $\lambda > 0$ and outputs the optimal rate-distortion Lagrangian solution. By sweeping λ, we obtain the rate-distortion curve.

Variable-Rate Compression

- Ballé et al. [1] defined an autoencoder-based variable-length compressor objective called **nonlinear transform coding (NTC)**,
  $$\inf_{f,g} \mathbb{E}_{x,\varepsilon}\left[\left\|x - g\left(Q\left(f(x) + \varepsilon\right) - \varepsilon\right)\right\|_2^2\right] + \lambda H\left(Q\left(f(x) + \varepsilon\right) - \varepsilon|\varepsilon\right). \tag{3}$$
- Assuming a Gaussian source, linear f and g, and under independent encoding of each dimension, the NTC Lagrangian is
  $$\inf_{W,T} \mathbb{E}_{x,\varepsilon}\left[\left\|x - T\left(W^\top + \varepsilon\right)\right\|_2^2\right] + \lambda \cdot \sum_{i=1}^{d} h\left(w_i^\top x + [\varepsilon]_i\right). \tag{4}$$

Theorem: Under the above assumptions, any W that achieves the infimum has all its nonzero columns as eigenvectors of K.

Experiments

- We compare a PBA-based fixed-rate compressor with PCA. For images, we compare against JPEG, JPEG2000. For audio, we compare against AAC.
- We compare across three metrics. 1) SNR $= 10 \cdot \log_{10}(P/MSE)$, 2) Structural similarity index (SSIM/MS-SSIM) 3) Performance on downstream tasks, specifically classification accuracy.

SNR Performance

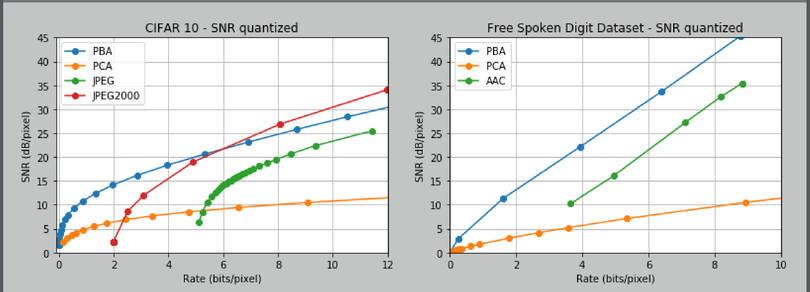

Figure: Left: CIFAR-10, Right: Free Spoken Digits Dataset

SSIM Performance

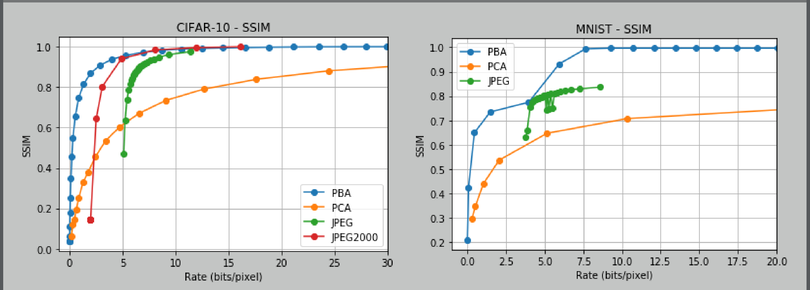

Figure: Left: CIFAR-10, Right: MNIST

Classification Accuracy Performance

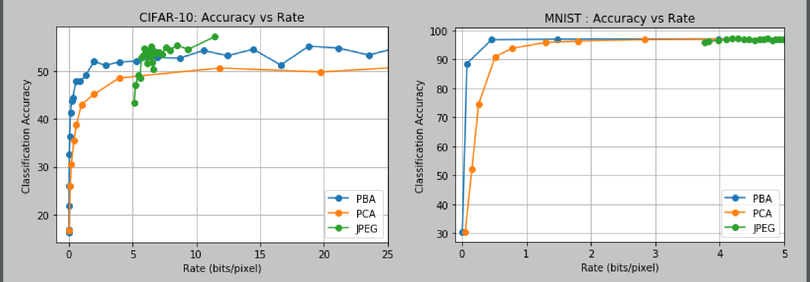

Figure: Left: CIFAR-10, Right: MNIST