Linear predictive coding, example

Problem

Suppose the signal is a stationary Gaussian source X_n with zero mean and auto correlation function $R_{XX}(k) = E\{X_n X_{n+k}\}$. We know the following values of R_{XX}:

$$R_{XX}(0) = 5.98,\ R_{XX}(1) = 5.39,\ R_{XX}(2) = 4.55,\ R_{XX}(3) = 3.95$$

The variance of the signal is $\sigma_X^2 = R_{XX}(0)$.

We want to code the source using linear prediction followed by uniform quantization and perfect, memoryless source coding. The step size of the quantizer is chosen such that the resulting rate after the source coder is 5 bits/sample.

No predictor

For reference, we first see what result we get when not using any prediction. The rate $R = 5$ bits/sample is high enough that we can use the fine quantization approximation. This gives us

$$D \approx \frac{\pi e}{6} \cdot \sigma_X^2 \cdot 2^{-2R} \approx 0.008312$$

Expressed as a signal-to-noise ration, we get

$$\text{SNR} = 10 \cdot \log_{10} \frac{\sigma_X^2}{D} \approx 28.57 \text{ [dB]}$$

One-step predictor

We are now going to use a one-step linear predictor:

$$p_n = a_1 \hat{X}_{n-1} \approx a_1 X_{n-1}$$

The prediction error variance is

$$\sigma_d^2 = E\{(X_n - p_n)^2\} \approx E\{(X_n - a_1 X_{n-1})^2\} = (1 + a_1^2) R_{XX}(0) - 2a_1 R_{XX}(1)$$

Differentiate with respect to a_1 and set equal to zero

$$\frac{\partial}{\partial a_1} \sigma_d^2 = 2a_1 R_{XX}(0) - 2R_{XX}(1) = 0$$

which gives us

$$a_1 = \frac{R_{XX}(1)}{R_{XX}(0)} \approx 0.9013 \quad \Rightarrow \quad \sigma_d^2 \approx 1.1218$$

The resulting distortion after quantization is then

$$D \approx \frac{\pi e}{6} \cdot \sigma_d^2 \cdot 2^{-2R} \approx 0.001559$$

Expressed as a signal-to-noise ration, we get

$$\text{SNR} = 10 \cdot \log_{10} \frac{\sigma_X^2}{D} \approx 35.84 \text{ [dB]}$$

Two-step predictor

We are now going to use a two-step linear predictor:

$$p_n = a_1 \hat{X}_{n-1} + a_2 \hat{X}_{n-2} \approx a_1 X_{n-1} + a_2 X_{n-2}$$

The prediction error variance is

$$\begin{aligned} \sigma_d^2 &= E\{(X_n - p_n)^2\} \approx E\{(X_n - a_1 X_{n-1} - a_2 X_{n-2})^2\} \\ &= (1 + a_1^2 + a_2^2) R_{XX}(0) - 2a_1 R_{XX}(1) - 2a_2 R_{XX}(2) + 2a_1 a_2 R_{XX}(1) \end{aligned}$$

Differentiate with respect to a_1 and a_2 and set equal to zero

$$\frac{\partial}{\partial a_1} \sigma_d^2 = 2a_1 R_{XX}(0) - 2R_{XX}(1) + 2a_2 R_{XX}(1) = 0$$

$$\frac{\partial}{\partial a_2} \sigma_d^2 = 2a_2 R_{XX}(0) - 2R_{XX}(2) + 2a_1 R_{XX}(1) = 0$$

which can be written as the matrix equation

$$\mathbf{RA} = \mathbf{P}$$

where

$$\mathbf{A} = \begin{pmatrix} a_1 \\ a_2 \end{pmatrix}, \ \mathbf{R} = \begin{pmatrix} R_{XX}(0) & R_{XX}(1) \\ R_{XX}(1) & R_{XX}(0) \end{pmatrix}, \ \mathbf{P} = \begin{pmatrix} R_{XX}(1) \\ R_{XX}(2) \end{pmatrix}$$

This gives us

$$\mathbf{A} = \mathbf{R}^{-1}\mathbf{P} \approx \begin{pmatrix} 1.1490 \\ -0.2747 \end{pmatrix} \quad \Rightarrow \quad \sigma_d^2 \approx 1.0371$$

The resulting distortion after quantization is then

$$D \approx \frac{\pi e}{6} \cdot \sigma_d^2 \cdot 2^{-2R} \approx 0.001442$$

Expressed as a signal-to-noise ration, we get

$$\text{SNR} = 10 \cdot \log_{10} \frac{\sigma_X^2}{D} \approx 36.18 \text{ [dB]}$$

Three-step predictor

We are now going to use a three-step linear predictor:

$$p_n = a_1\hat{X}_{n-1} + a_2\hat{X}_{n-2} + a_3\hat{X}_{n-3} \approx a_1X_{n-1} + a_2X_{n-2} + a_3X_{n-3}$$

Similarly to the previous predictors, we can express the prediction error variance using the auto correlation function and then differentiate it with respect to each of the predictor coefficients and set equal to zero to get a linear equation system. Since we know the structure of this equation system we can immediately express it as a matrix equation

$$\mathbf{RA} = \mathbf{P}$$

where

$$\mathbf{A} = \begin{pmatrix} a_1 \\ a_2 \\ a_3 \end{pmatrix}, \ \mathbf{R} = \begin{pmatrix} R_{XX}(0) & R_{XX}(1) & R_{XX}(2) \\ R_{XX}(1) & R_{XX}(0) & R_{XX}(1) \\ R_{XX}(2) & R_{XX}(1) & R_{XX}(0) \end{pmatrix}, \ \mathbf{P} = \begin{pmatrix} R_{XX}(1) \\ R_{XX}(2) \\ R_{XX}(3) \end{pmatrix}$$

This gives us

$$\mathbf{A} = \mathbf{R}^{-1}\mathbf{P} \approx \begin{pmatrix} 1.2028 \\ -0.4997 \\ 0.1958 \end{pmatrix} \quad \Rightarrow \quad \sigma_d^2 \approx 0.9974$$

The resulting distortion after quantization is then

$$D \approx \frac{\pi e}{6} \cdot \sigma_d^2 \cdot 2^{-2R} \approx 0.001386$$

Expressed as a signal-to-noise ration, we get

$$\text{SNR} = 10 \cdot \log_{10} \frac{\sigma_X^2}{D} \approx 36.35 \text{ [dB]}$$

Analysis

As can be seen, the results gets better the larger the predictor is. This will always be true when we ignore the effect of the quantization on the prediction using our approximation. Note however that in the real world, having a long predictor might actually give a worse result than a shorter predictor, since we are feeding back the quantization noise into the predictor loop. See the lecture slides for examples of this effect.