

MATLAB Exercise – Uniformly Quantize a Speech Signal

Program Directory: `speech_coding/uniform_quantization`

Program Name: `uniform_quantization_GUI25.m`

GUI data file: `uniform_quantization.mat`

Callbacks file: `Callbacks_uniform_quantization_GUI25.m`

TADSP: Section 11.4.1, pp. 680-691, Problem 11.25

This MATLAB exercise illustrates the process of uniform quantization of a speech sample using a range of B -bits per sample, and displays the unquantized signal, the quantized signal and the error signal at the selected bit rate along with plots of signal and error power density spectrum, and the histogram of the quantization error signal.

Uniformly Quantize a Speech Signal – Theory of Operation

This exercise quantizes a designated speech file using a mid-riser uniform quantizer with a bit rate of `nbits` per sample, where `nbits` is in the range of 2-10. For the designated value of `nbits`, the program plots the original signal, the quantized signal, and the error signal along with estimates of the signal and error power spectrum, and a histogram of the quantization error signal values. The program also enables playback of the original speech signal, the quantized speech signal, and the error signal in order to get a feeling for how much quantization noise is introduced into the quantized signal at different quantization bit rates.

Uniformly Quantize a Speech Signal – GUI Design

The GUI for this exercise consists of two panels, 5 graphics panels, 1 title box and 8 buttons. The functionality of the two panels is:

1. one panel for the graphics display,
2. one panel for parameters related to the the signal processing for signal quantization, and for running the program.

The set of five graphics panels is used to display the following:

1. the original (unquantized) speech signal,
2. the quantized speech signal,
3. the quantization error signal,
4. the signal and quantization error power spectrums,
5. the quantization error signal histogram.

The title box displays the information about the selected file along with the set of signal processing parameters. The functionality of the 8 buttons is:

1. a pushbutton to select the directory with the speech file that is to be analyzed using short-time analysis methods; the default directory is 'speech_files',
2. a popumenu button that allows the user to select the speech file for analysis,
3. an editable button that specifies the number of bits per sample, `nbits`, in the mid-riser uniform quantizer,
4. a pushbutton to run the uniform quantization code and display the results of the uniform quantization process on the five graphics panel displays,
5. a pushbutton to play out the original speech signal,
6. a pushbutton to play out the quantized speech signal,
7. a pushbutton to play out the quantization error signal,
8. a pushbutton to close the GUI.

Uniformly Quantize a Speech Signal – Scripted Run

A scripted run of the program 'uniform_quantization_GUI25.m' is as follows:

1. run the program 'uniform_quantization_GUI25.m' from the directory:

```
'matlab_gui\speech_coding\uniform_quantization',
```

2. hit the pushbutton 'Speech Directory'; this will initiate a system call to locate and display the filesystem for the directory 'speech_files',
3. using the popupmenu button, select the speech file for short-time feature analysis; choose the file 'sp41e_8k.wav' for this example,
4. using the editable buttons, set the value for the number of bits per sample for the uniform quantizer to `nbits=4`,
5. hit the 'Run Uniform Quantizer' button to compute and display the original speech signal in the upper graphics panel, the quantized signal in the second graphics panel, the quantization error signal in the third graphics panel, the power spectrums of the original signal and the quantization error signal in the fourth graphics panel, and the error signal histogram in the bottom graphics panel,
6. experiment with different choices of speech file, and with different values for `nbits`,
7. hit the 'Close GUI' button to terminate the run.

An example of the graphical output obtained from this exercise using the speech file 'sp41e_8k.wav' is shown in Figure 1. The number of bits per sample, `nbits`, was chosen as 10 for this example. The graphics panels show the original speech waveform (top graphics panel), the quantized signal (second graphics panel), the quantization error signal (third graphics panel), the power spectrums for the original signal and the error signal (fourth graphics panel), and the error signal histogram (bottom graphics panel).

Uniformly Quantize a Speech Signal – Issues for Experimentation

1. If, instead of using a designated speech file, you substituted a ramp signal (i.e., a signal that begins at -1 and ends at $+1$, with a few thousand linearly spaced samples in between these two signal values), how would each of the five plots shown in Figure 1 differ from those of speech signal plots? Hint: if you set `nbits` to 0, this ramp signal is generated and is used to replace the speech signal in the analysis. The results of such a change of signal are shown in Figure 2.
2. Run the program for several different decreasing values of `nbits` using the input sentence "every salt breeze_suzanne.wav". In each case, listen to both the original and the quantized signal. At what value of `nbits` do you begin to hear a difference? Also listen to the quantization error signal. At what value of `nbits` does it no longer sound like 'white noise'?
3. What is the maximum size of the quantization error in the case `nbits=10`? How does it compare to the theoretical value for the number of bits selected?
4. With the file "every salt breeze_suzanne.wav", use different numbers of bits. Looking at the noise power spectrum, at what value of `nbits` does the white noise model appear to break down? For values of `nbits` when the spectrum is flat, is its level consistent with the theoretical model for the number of bits in the quantizer?
5. Now consider the histogram of error samples. Is it approximately uniform? For what value of `nbits` would you say that the uniform model is no longer a good approximation?
6. Run the program on file '1A.waV'. Does the white noise model work for this signal? What might explain the nature of the error histogram for this signal?
7. Run the uniform quantizer on some of the other speech signals in the `speech_files` folder. Are the results the same across different speech files?

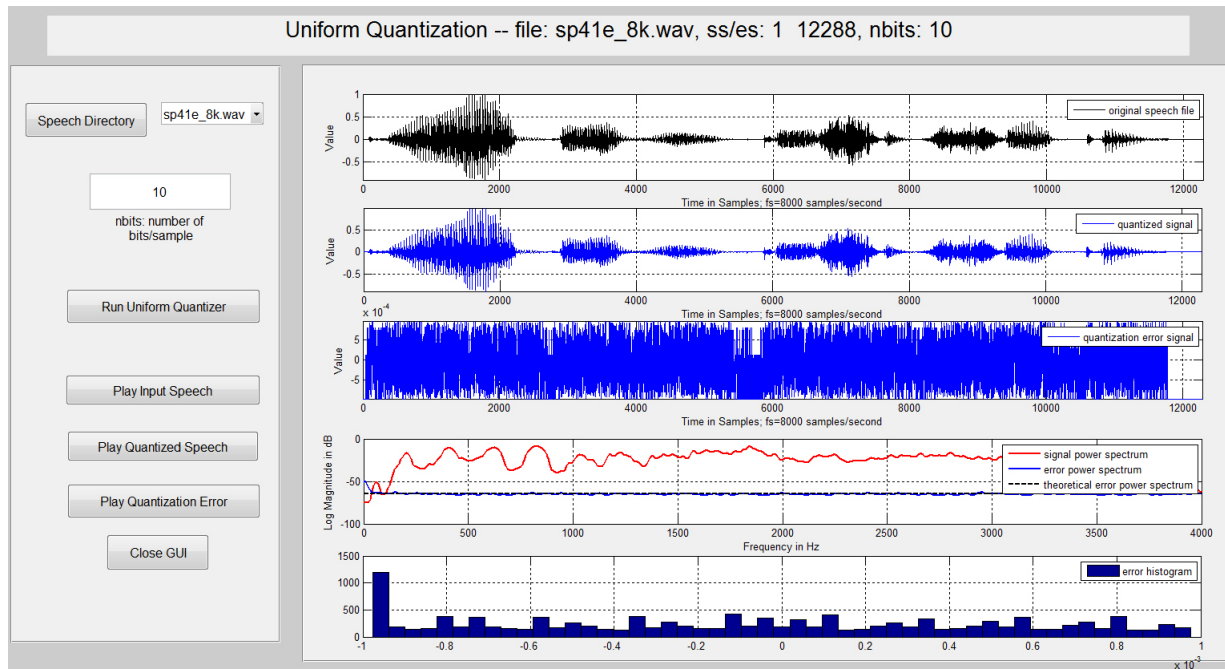


Figure 1: Graphical output from the uniform quantization exercise. The five graphics panels show the original signal, the quantized signal, the quantization error signal, the power spectrums of the original signal and the quantization error signal, and the error signal histogram for a run using `nbits=10` bits per sample.




Figure 2: Graphical output from the uniform quantization exercise with a ramp replacing the designated speech file. The five graphics panels show the original signal, the quantized signal, the quantization error signal, the power spectrums of the original signal and the quantization error signal, and the error signal histogram for a run using `nbits=4` bits per sample.