

MATLAB Exercise – Mu-law Quantization for a Speech Signal

Program Directory: `speech_coding/mu_law_quantization`

Program Name: `mu_law_quantize_GUI25.m`

GUI data file: `mu_law_quantize.mat`

Callbacks file: `Callbacks_mu_law_quantize_GUI25.m`

TADSP: Section 11.4.2, pp. 691-676, Problem 11.27

This MATLAB exercise illustrates the impact of mu-law companding (compressing and expanding) and quantization of a speech signal.

Mu-law Quantization for a Speech Signal – Theory of Operation

This exercise quantizes a designated speech file using a mu-law quantizer with a bit rate of `nbits` per sample, where `nbits` is typically in the range of 2-10. For the chosen value of `nbits`, the program plots the quantized and the error signals along with an estimate of the signal and error power spectrum, and a histogram of the error signal values. The program also plays out the error signals in order to get a feeling for how much noise is introduced into the quantized signal at different bit rates.

Mu-law Quantization for a Speech Signal – GUI Design

The GUI for this exercise consists of two panels, 5 graphics panels, 1 title box and 9 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 mu-law quantized speech signal,
3. the mu-law quantization error signal,
4. the signal and error power spectrums,
5. the 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 9 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 popupmenu button that allows the user to select the speech file for analysis,
3. an editable button that specifies the value of `mu` for the mu-law quantizer; (default is `mu=40`),
4. an editable button that specifies the number of bits per sample, `nbits`, in the mu-law quantizer; (default is `nbits=10`),
5. a pushbutton to run the code and display the results of the mu-law quantization process on the five graphics panel displays,
6. a pushbutton to play out the original input speech signal,

7. a pushbutton to play out the mu-law quantized speech signal,
8. a pushbutton to play out the mu-law quantization error signal,
9. a pushbutton to close the GUI.

Mu-law Quantization for a Speech Signal – Scripted Run

A scripted run of the program 'mu_low_quantize_GUI25.m' is as follows:

1. run the program 'mu_low_quantize_GUI25.m' from the directory:


```
'matlab_gui\speech_coding\mu_low_quantization'
```
2. hit the pushbutton '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 'we were away a year ago_lrr.wav' for this example,
4. using the editable buttons, set the values of the mu-law compression to $\mu=40$, and the number of bits per sample for the mu-law quantizer to $n_{bits}=10$,
5. hit the 'Run Mu-Law Quantizer' button to compute and display the original speech signal in the upper graphics panel, the mu-law quantized signal in the second graphics panel, the mu-law quantization error signal in the third graphics panel, the power spectrums of the original signal and the mu-law quantization error signal in the fourth graphics panel, and the error signal histogram in the bottom graphics panel,
6. hit the 'Play Input Speech' button,
7. hit the 'Play Quantized Speech' button,
8. hit the 'Play Quantization Error' button,
9. experiment with different choices of speech file, and with different values for μ and n_{bits} ,
10. hit the 'Close GUI' button to terminate the run.

An example of the graphical output obtained from this exercise using the speech file 'we were away a year ago_lrr.wav' is shown in Figure 1. The value of μ was chosen as 40 for this example and the number of bits per sample, n_{bits} , was chosen as 10. The graphics panels show the original speech waveform (top graphics panel), the mu-law quantized signal (second graphics panel), the mu-law quantization error signal (third graphics panel), the power spectrums for the original signal and the mu-law error signal (fourth graphics panel), and the error signal histogram (bottom graphics panel).

Mu-law Quantization for a Speech Signal – Issues for Experimentation

1. Run the program for several different decreasing values of n_{bits} 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 n_{bits} do you begin to hear a difference? Also listen to the quantization noise. How is the noise different from the noise of a uniform quantizer? Do you think the noise could be considered to be white noise?
2. Compare the waveforms of the quantization noise and the original speech signal. Is what you observe consistent with a model for the quantized output signal of the form $\hat{x}[n] = x[n] + e[n]x[n]$, where $x[n]$ is the original signal and $e[n]$ is the quantization noise incurred in quantizing the output of the μ -law compressor? With such an approximation, the noise component at the output of the μ -law expander would be $f[n] = e[n]x[n]$. Determine σ_f^2 and $\phi_f[m]$ under the assumption of independence of $e[n]$ and $x[n]$ and use the results to help you interpret the spectrum and histogram plots.

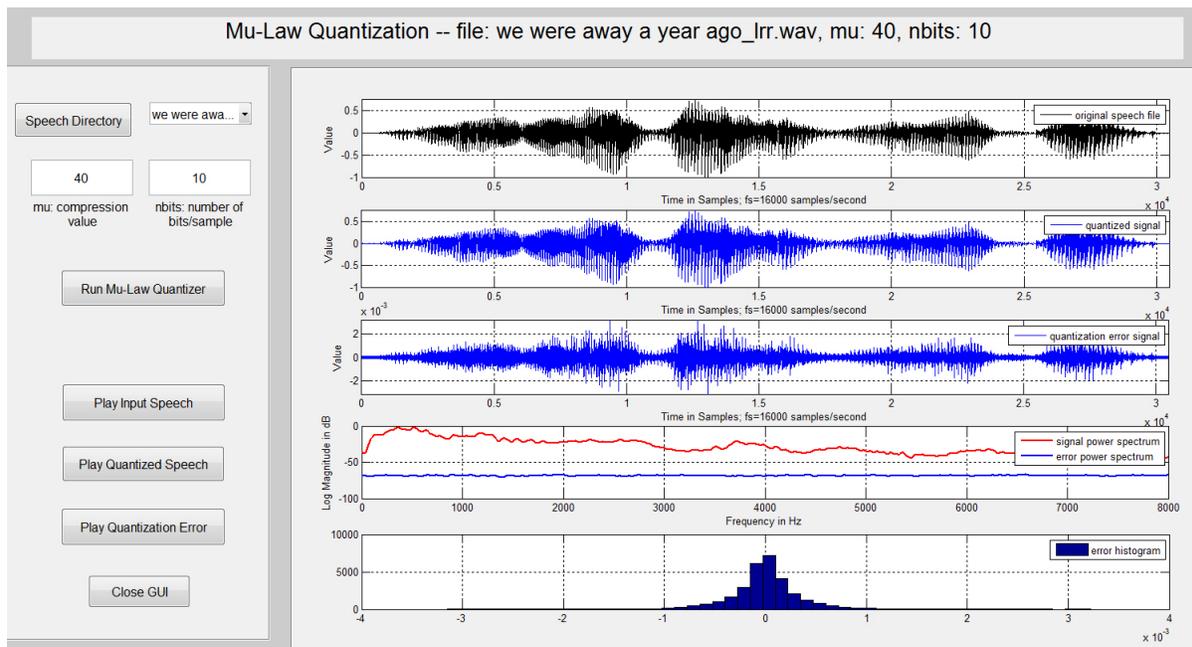


Figure 1: Graphical output from the mu-law quantization exercise. The five graphics panels show the original signal, the mu-law quantized signal, the mu-law quantization error signal, the power spectrums of the original signal and the quantization error signal, and the error signal histogram for a run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

3. With the mu-law quantization of the file "every salt breeze_suzanne.wav", run the exercise using different numbers of bits per sample. Is the spectrum flat for `nbits=10`? At what value of `nbits` does the white noise model appear to break down? When the spectrum is flat, is its level the same as with the uniform quantizer with the same number of bits?
4. Run the μ -law quantizer on some of the other signal signals in the `speech_files` folder. Are the results essentially the same?