



2002

Best practices in the acquisition, processing, and analysis of acoustic speech signals.

Bartek Plichta

Follow this and additional works at: <https://repository.upenn.edu/pwpl>

Recommended Citation

Plichta, Bartek (2002) "Best practices in the acquisition, processing, and analysis of acoustic speech signals.," *University of Pennsylvania Working Papers in Linguistics*: Vol. 8 : Iss. 3 , Article 16.
Available at: <https://repository.upenn.edu/pwpl/vol8/iss3/16>

This paper is posted at ScholarlyCommons. <https://repository.upenn.edu/pwpl/vol8/iss3/16>
For more information, please contact repository@pobox.upenn.edu.

Best practices in the acquisition, processing, and analysis of acoustic speech signals.

Best Practices in the Acquisition, Processing, and Analysis of Acoustic Speech Signals

Bartek Plichta

1 Introduction

This article is intended to assist researchers in implementing audio technology in the field, studio, and research lab. It presents a set of best practices in the recording, processing, and analysis of acoustic signals. The author has no commercial interest in recommending particular pieces of hardware or software. Some brand names and manufacturers have been mentioned to help researchers find appropriate tools more easily.

2 Recording

2.1 Technical Basics

The purpose of recording is to capture the best possible signal. Regardless of the recording situation, one should consider at least three important acoustic parameters: frequency response, dynamic range, and signal-to-noise ratio (SNR). Frequency response describes the range of frequencies captured in the recording, measured in Hertz (Hz). Dynamic range is the ratio of the loudest to the softest part of the signal, measured in Decibels (dB). SNR is the ratio of signal amplitude to the amplitude of noise usually generated by the circuit, measured in dB. In the audio digitization process, SNR refers to the ratio of the maximum signal power to the quantization noise generated by the analog-to-digital converter. Table 1 summarizes the frequency response and dynamic range characteristics of speech and recording media.

| Medium | Frequency response: | Dynamic range: |
|------------------------------------------|---------------------|----------------|
| Speech | < 10000 Hz | 30-40 dB |
| Analog tape | < 10000-15000 Hz | 45 dB |
| DAT (Digital Audio Tape, 16 bit, 48 kHz) | < 24000 Hz | 96 dB |
| Digital telephone (ISDN) | < 4000 Hz | 48 dB |

Table 1. Frequency response and dynamic range characteristics of speech and recording media.

As we can see, the speech signal is not particularly demanding in terms of frequency response and dynamic range. In theory, a modern, high-quality analog or digital recording device should be able to capture the speech signal fairly accurately. The only exception is the phone line (analog or digital), which lacks detail in frequency and adds substantial amounts of noise to the signal. However, we can also see that digital recording devices such as DAT and hard disk recorders produce a considerably richer signal; one whose frequency spans the entire audible range and contains enough amplitude variation to allow for very detailed acoustic analysis.

2.1 Recording Hardware Setup

The recording process requires, minimally, the following set of hardware devices listed in Figure 1. Each of the devices used in the process may have a crucial impact on the overall quality of the acquired signal.

Typical analog recording setup

microphone– (phantom power)– pre-amplifier – analog recorder – analog tape

Typical digital recording setup

microphone– (phantom power)– pre-amplifier– A/D converter– HDD or DAT

Figure 1. Minimal recording hardware setups

2.2 Microphones

Two basic microphone types are most typically used for recording speech. The dynamic microphone has a diaphragm that consists of Mylar plastic that has a finely wrapped coil of wire (“voice coil”) attached to its inner face. This coil is suspended within a strong magnetic field. Whenever a sound wave hits the diaphragm, the coil is displaced in proportion to the amplitude of the wave, causing the coil to cut across the lines of magnetic flux supplied by the permanent magnet. Since the mass of the diaphragm and the coil is quite large, compared to the pressure changes in the sound wave, the dynamic microphone may not respond well to sharp transient sounds and it may fail to record minute changes in voice intensity. This does not mean that dynamic microphones should not be used for speech recording. On the contrary, there are several low-cost, rugged microphones, such as Shure SM58 or Shure SM48 (once recommended by Kay Elemetrics), that can be used successfully for many speech recording applications. Dynamic microphones do not require any external power supply, which makes them a good match to several field recorders (e.g., Marantz PMD 222).

The condenser microphone works on a different principle. A thin plastic diaphragm coated on one side with gold or nickel is placed at a close distance from a stationary backplate. Once a polarizing voltage (from a 48V phantom power supply) is applied to these plates, the two surfaces create capacitance that varies as the diaphragm moves in response to a sound wave. Since the diaphragm is very light, the response of the condenser microphone is very accurate, often producing a recording extremely rich in both frequency and dynamic response. Condenser microphones are usually more fragile than dynamic microphones and require a 48V phantom power supply. Some can use battery packs, but others rely on an external power source. This makes the condenser microphone cumbersome for field use, despite the on-board phantom power supply found in some DAT and HDD recorders.

2.3 Frequency Response

Each microphone type has a unique frequency response. It is important to remember that many manufacturers tailor a microphone's frequency response to accentuate particular parts of the spectrum to function optimally within a specific recording application. For the purposes of acoustic analysis, always choose microphones with a wide and flat frequency response curve.




Figure 2. Frequency response of a Shure Beta 87a microphone. Note the proximity effect at 10 mm.

2.3.1 Polar Patterns

A microphone's polar pattern should play a crucial role in its selection for a specific purpose. The polar pattern is a plot of the sensitivity of a microphone as a function of the angle around it (Figure 3). Several common polar patterns are used in microphones today. Omnidirectional microphones record sound equally from all directions. They are most commonly found in built-in or lavalier types, and are good for recording interviews, but their 360° pickup range introduces too much noise to be used reliably in acoustic analysis.




Figure 3. Cardioid polar patten of a Shure Beta 87a microphone

The cardioid (heart-shaped) pattern is most sensitive to sounds coming from the front. It is 6 dB less sensitive to sounds from 90 degrees to the side, and, in theory, is completely insensitive to sounds coming from the rear. The most important attribute of a cardioid (directional) microphone is its ability to discriminate between direct sounds (coming from the direction in which it is pointed) and reverberant, unwanted sounds from all other directions. This type of polar pattern usually produces signals that are substantially less noisy than those captured with an omnidirectional microphone (Figure 4).




Fig. 4a




Fig. 4b

Figure 3. Waterfall plot of a recording of, "Bob was positive that his wife".

Figure 4a. Recorded with a built-in, omnidirectional microphone.

Note the amount of room noise and lack of detail in spectral prominences.

Figure 4b. Recorded with a head-worm, directional microphone.

Note negligible amounts of noise and superb detail in the speech spectrum.

2.3.2 Proximity Effect

Usually, high-quality speech recordings require the sound source to be fairly close to the microphone's diaphragm. This may trigger a so-called proximity effect (Figure 2). Proximity effect is the increase in the low-frequency sensi-

tivity of a microphone when the sound source is close to it. This is particularly true of cardioid, directional microphones. To counter that, most high-end directional microphones use a low-frequency roll-off filter to restore the response to its flat, natural balance. Some microphones have a user-selectable switch to control the filter. The proximity effect may be responsible for speech spectra showing emphasis in the low-frequency range, around the first and second harmonics.

2.3.3 Cabling and Phantom Power

Cabling is a common cause of recording problems. To avoid noise (60 Hz hum) and phase problems, use professional quality balanced XLR (two conductors for the signal with neither connected to the shield) cables to connect the microphone to the pre-amplifier. If the pre-amplifier does not have balanced XLR inputs, one should use a balanced to un-balanced transformer (Ebttech Line shifter, CP8201 In-Line Transformer). The transformer's primary side matches the impedance of the microphone and is balanced, while the secondary is unbalanced and has high impedance that matches most unbalanced pre-amplifier inputs. Condenser microphones require a 48V phantom power. While some microphones, such as AKG C1000, can work from an internal battery source, others require an external phantom power supply, which many good pre-amplifiers and mixing consoles come equipped with.

2.4 Recording Techniques

Much of the success of a recording depends on the recording environment and microphone placement. Ideally, speech recordings should take place in soundproof studios or labs. If unavailable, find a relatively quiet room with as little low-frequency noise as possible. Typical sources of low-frequency noise include 60-Hz hum from electrical equipment, heating and air-conditioning ducts, elevators, doors, water pipes, computer fans, and other mechanical systems. If possible, switch off those devices during recording. Figure 4 illustrates the spectrum of typical ambient room noise. The low-frequency prominence of this kind of noise may interfere with acoustic analysis of F_0 and F_1 . It is quite difficult to filter out extraneous noise without removing information from the speech signal itself. High frequency noise should also be avoided, though it can be more easily filtered out of the signal before analysis.

The placement of the microphone directly affects the intensity of the recorded signal as well as the signal-to-noise ratio. The inverse square law guarantees a loss of approximately 6dB per doubling of distance from the sound source. A typical hand-held microphone is usually placed at a distance of 30 cm or so from the speaker's lips. Relative to a close placement of 4 cm,




Figure 4. Spectrum of typical room noise. Note prominence around 600 Hz

this represents the loss of about 18dB and an increased possibility of noise leaking into the recording. For this reason, use a head-mounted condenser microphone (e.g. AKG C410 or AKG C420) to maintain a close, constant distance to the source. Speech signals acquired this way are characterized by a high SNR and a broad range of intensity. It may also be useful to use a linear phase high-pass rumble filter (60 Hz cutoff and 24dB/octave attenuation), unless low-frequency components are expected in the signal.

2.5 Pre-Amplifiers

The main function of a pre-amplifier is to accept a very low-level signal (such as that from a microphone) and amplify it without adding noise. Good pre-amplifiers are not easy to build, as they have to be immune to all kinds of potential noise and signal distortion. It is, therefore, important to use the best possible pre-amplifier, particularly one that has a fairly high gain, broad dynamic range, high SNR, phantom power, and balanced XLR inputs. Good field pre-amplifiers are particularly hard to find. Sound Devices MP2, USB Pre, and Shure FP 23 are among some of the best, yet affordable field pre-amplifiers. The gain on the pre-amplifier should be set relatively high, but the speaker's voice amplitude range should be tested first to avoid signal overload. If intensity measurements are not intended, a soft compressor-limiter may be used to maximize amplitude and help prevent signal clipping.

2.6 Recording Devices

Several types of recording devices are used today. Some of the most common include analog cassette recorders, digital audio tape (DAT) recorders, minidisk recorders, and hard disk recorders. Some, especially field recorders, come with a built-in pre-amplifier and two independent input channels. When recording in the studio or lab, a hard disk recorder is highly recommended. In the field, however, other devices may also be considered.

2.6.1 Analog Recorders

As we saw in Table 1, analog tape is able to provide the necessary frequency response and dynamic range to capture a fair amount of detail in the speech signal. In fact, if a premium-grade microphone and pre-amplifier have been used, a professional cassette recorder, such as Marantz PMD 222, is capable of capturing fairly high-quality acoustic signals. Cassette recorders are usually quite durable and inexpensive. They operate on alkaline batteries and use standard audio cassettes, which are inexpensive and widely available. However, analog recorders typically have noisy transport mechanisms and offer no time code, which makes logging and analysis more difficult. In addition, analog tape is a fragile medium and should be avoided by researchers interested in long-term preservation of recordings. Perhaps the most serious disadvantage of analog recordings is that they must be digitized before any computer-assisted analysis can be performed. Although the audio digitization process can offer an extremely high quality of transfer, it is often done poorly, resulting in a considerable degradation of the original material.

2.6.2 Digital Recorders

Digital recorders are becoming more and more common. They come in a variety of forms, such as DAT, minidisk, solid state PC card, CD, and hard disk. What distinguishes them from one another are the recording medium and format. For the purposes of acoustic analysis, use digital recorders that capture sound in an uncompressed, PCM (pulse code modulation) format.

At the heart of a digital recorder is the Analog-to-Digital converter (ADC). In digital systems, analog audio signal must be converted to digital form before it can be further processed. This entails sampling the signal at very short, successive, time intervals and converting the value of the sample to a binary number representing the amplitude of the waveform at that instant. The output of the ADC is a series of digital "words", typically at a rate of 44,100 words per second (sample rate). Each word contains a certain number (typically 16) of bits. Sample rate and bit depth are the most important factors determining the accuracy of the digital representation of the analog waveform. A typical digitization setting of a 44,100Hz sample rate and 16-bit word length ("CD quality") renders a recording with a broad frequency response (about 21,000 Hz) and an impressive dynamic range of 96 dB. Many modern ADCs are capable of sampling at the rate of 96 kHz and 24-bit resolution, which produces a very accurate signal with negligible quantization noise. Once the signal has been quantized, the values of each sample must be stored on the medium that the recording device uses.

2.6.3 DAT Recorders

DAT recorders use magnetic tape as a storage medium. Typical tapes store 120 minutes of uncompressed, high-quality mono recording at 48,000 Hz/16-bit. TASCAM DA-P1 offers sophisticated time coding, rugged construction, and good pre-amplifiers. It also comes with a built-in limiter and microphone/line inputs, and features the S/PIDF digital I/O interface. The transfer of digital audio from DAT to a computer hard drive for analysis is a lossless process, provided the correct digital transfer interface is used. DAT tape is as fragile as analog cassette tape, which again raises the issue of long-term preservation. Most digital-born recordings stored on DAT tapes will eventually have to be transferred to spinning disk or optical storage.

2.6.4 Minidisk Recorders

Minidisk recorders are popular among field workers due to their small size and ease of use. However, most minidisk recorders do not produce the quality required for detailed acoustic analysis. First, most portable minidisk recorders lack a quality pre-amplifier, making it difficult to interface them with premium-grade microphones. Minidisk field recordings are usually shipped with small, inexpensive lavalier condenser microphones whose frequency response and dynamic range are rather poor. In addition, there is virtually no control over the incoming signal. The A/D converter is inferior to that found on larger DAT or hard disk recorders. Many portable minidisk recorders lack the S/PDIF I/O interface, which makes it necessary to use a stand-alone minidisk deck for digital audio transfer. Finally, to achieve a small disc diameter, MD reduces data based on psychoacoustic principles. Prior to storage, the audio data rate (bit rate) of 1.41 Mbps (uncompressed PCM) is compressed using a perceptual coder (proprietary algorithm developed by SONY) to lower the bit rate to 292 Kbps (1/5 of the original). ATRAC preserves the sample rate but decreases the word length (resolution), resulting in a less faithful representation of the original and quantization noise, which is then masked by the algorithm. A lower word length results in decreased dynamic range, for which the algorithm cannot compensate. This produces a recording with a full 20-20000 Hz frequency response, but decreased resolution, increased quantization noise, and decreased dynamic range.

2.6.5 PC Card and CD-R Recorders

Marantz has developed a line of professional quality solid state PC card field recorders. They have impressive features, one of which is the ability to re-

cord uncompressed PCM 16-bit audio at the sample rate 44,100 Hz. These recorders are compatible with both compact flash and ATA-size PC cards. The data is stored as a Wav audio file, which can be easily downloaded to a PC with a PC card reader. Unfortunately, the cost of a PC card capable of containing an hour of uncompressed audio is still prohibitively high, and it is uncertain whether the electronic industry will continue to be interested in developing high-capacity, inexpensive memory cards. Recently, Marantz has released a portable CD recorder, PSD 300, which boasts the same professional features as other Marantz portables, but writes data directly to a CD-R or CD-RW. Since this is a new release, it remains to be seen how the PSD copes with quality issues as well as the demands of fieldwork.

2.6.6 Hard Disk Recorders

Hard disk recording is usually associated with expensive recording studio multi-track recording devices connected to SCSI hard disk farms and racks full of processing hardware. However, recently, small USB devices have appeared on the market. USB Pre 1.5 is an impressive piece of hardware. It features two independent studio-quality microphone pre-amps, phantom power, 24-bit A/D converters, 106 dB dynamic range and a variety of inputs and outputs (tape, line, instrument, microphone, and S/PDIF). The device is powered through a computer USB port and can be used with most modern PC and Macintosh machines, including a variety of laptops. Recordings obtained with USB Pre are superb. Whenever the recording situation allows the use of a laptop, USB Pre (or similar device) should be used to acquire speech signals. At the time of writing this article, USB Pre is one of the best USB digital recording devices, though Digidesign has just released a similar product, the Mbox, which offers similar features to USB Pre, in addition to being fully compatible with Pro Tools 5.2 LE software.

3 Processing

3.1 Analog-to-Digital Conversion and Digital Data Transfer

In order to perform acoustic analysis on recorded speech data, the audio signal has to be converted into a digital audio file format, such as Wav or Aiff. Analog recordings have to be digitized and digital recordings need to be transferred to a personal computer via a digital audio file transfer interface. This is an important, yet often underestimated, stage in the process of preparing audio data for analysis.

The goal of A/D conversion (digitization) is to obtain the best possible digital representation of the original analog waveform. Choose a sample rate that will capture a broad range of frequencies and a bit-depth that will allow a wide dynamic range and a negligible amount of quantization noise. This requires a premium-quality, stand-alone A/D converter operating at the sample rate of at least 48,000 Hz and 24-bit resolution: It is crucial not to use a PCI multimedia sound card, which is built from inferior electronic components and, more importantly, allows electrostatic noise and distortion to leak into the captured acoustic signal (Figure 6)




Figure 5. Typical electrostatic noise generated by computer circuitry.

The A/D converter (e.g. Lucid AD 9624) should offer a variety of sample rates, oversampling, high quality anti-aliasing filters, and AES/EBU and S/PDIF digital outputs. Both AES/EBU (Audio Engineering Society/European Broadcasting Union) and S/PDIF (Sony/Philips Digital Interface) are fairly common on high-end digital audio devices. S/PDIF is used on a variety of consumer-level products, such as CD and minidisk players. It is also a common interface used on PCI digital I/O cards, which is why it is probably a better choice for most digital audio transfer applications.

| | AES/EBU | S/PDIF (IEC-958) |
|-----------------|-------------------|----------------------------|
| Cabling | 110 ohm shielded | TP 75 ohm coaxial or fiber |
| Connector | 3-pin XLR | RCA (or BNC) |
| Signal level | 3..10V | 0.5..1V |
| Modulation | biphase-mark-code | biphase-mark-code |
| Max. Resolution | 24 bits | 24 bits |

Table 2. Comparison of AES/EBU and S/PDIF interfaces

The analog playback device (e.g. TASCAM 122 mkIII) should be connected to the A/D converter. Make sure that the output levels on the tape deck match the input levels on the A/D converter. Use balanced XLR line level interface (+24 dBu min. gain, +7 dBu max. gain, 65k ohm impedance).

If the tape deck does not have this kind of output interface, use a signal level transformer (such as Ebtch Line shifter) and a pre-amplifier (Table 3).

| | Clip Level (1% THD) | Sensitivity (typical for 0 dB FS) | | Impedance (actual) |
|------|------------------------|-----------------------------------|-----------|-------------------------|
| | | min. gain | max. gain | |
| MIC | -12 dBu (0.2 V rms) | -10 dBu | -53 dBu | 2k ohm active-balanced |
| LINE | +24 dBu (12.3 V rms) | +24 dBu | +7 dBu | 65k ohm active balanced |
| DI | +9 dBu (2.2 V rms) | +8 dBu | -9 dBu | 10k ohm unbalanced |
| TAPE | +9 dBu (2.2 v rms) | +8 dBu | -9 dBu | 110k ohm unbalanced |

Table 3. Summary of typical signal level types.

The A/D converter needs to be connected to a PCI (though USB and FireWire are becoming common) digital audio I/O card (such as Midiman Delta DiO 2496 via a S/PDIF interface. Select the digital I/O card as the recording interface in the audio recording software (e.g. Sonic Foundry Sound Forge 5.0 (PC) or BIAS Peak VST (Mac)). The digital audio signal should be captured with this software and saved either as Wav (PC) or Aiff (Mac) file at the sample rate and bit depth that the A/D converter was set to. It is also possible to capture digital audio signal directly into acoustic analysis software, such as CSL or Praat, though this is not recommended due to the fact that specialized recording and processing software offers considerably more control over the incoming signal. Also, USB Pre may be used as a high-quality, stand-alone A/D converter. In this case the digital audio signal is transferred to a PC via the USB interface, which eliminates the need to install a separate PCI digital I/O card and makes it possible to capture digital audio on a laptop. In addition, USB Pre has a pair of tape-level inputs, to which a cassette deck can be directly connected (see Table 3).

4 Analysis

4.1 Preparing Files for Analysis

48,000 Hz/24-bit PCM audio files are not suitable for acoustic analysis, though they can store master, preservation copies of the recordings. Such files should be described in simple metadata terms (such as Dublin Core or OLAC) and stored on a dependable optical storage medium, such as CD-R

or DVD . In order to prepare files for acoustic analysis, use several digital signal processing (DSP) techniques. First, down-sample the files to 11,025 Hz and change their resolution to 16 bits. The downsampling process should always include the use of anti-aliasing filters, and the resolution change should include dithering to minimize the effects of quantization noise. Both Sonic Foundry Sound Forge 5.0 and BIAS Peak VST offer adequate DSP tools to do that. The Nyquist theorem guarantees that the frequency response of a 11,025 Hz is exactly half that value. This is adequate for male speech. For female and child voices, use a higher sample rate (e.g., 22,000 Hz) to make sure that the spectrum contains higher frequency transient sounds.

4.2 Restoration

Often, linguists find themselves dealing with old, noisy recordings. Digitization does not improve the quality of the recording, nor does it remove any of its imperfections. However, a few simple DSP techniques can be used to clean up and enhance the recording. The example below shows how DSP techniques have effectively removed some of the unwanted noise from an old DARE recording. The original signal contained low frequency noise whose spectrum overlapped with that of speech, particularly around the F1 area. After restoration, the amplitude of noise has been decreased by over 20dB, thus making the file appropriate for reliable acoustic analysis.





Fig. 7a




Fig. 7b

Figure 6. "I remember" from DARE, DM 0735- S1. (1) original digitized at 48/24, (2) converted to 16 bit, (3) downsampled to 11,025 Hz, (4) 2:1 compressed starting at -15 dB, (5) volume adjusted. 7a. before, 7b. after.

4.3 LPC in Acoustic Analysis

Linear Predictive Coding (LPC) is often used by linguists as a formant extraction tool. There are a few important details about LPC that may help

avoid common analysis errors. LPC analysis assumes that a signal is the output of a causal linear system. It also assumes that the vocal-tract system is an all-pole filter and that the input to the system is an impulse train. Because of these assumptions, LPC analysis is usually most appropriate for modeling vowels which are periodic and for which the vocal-tract resonator does not usually include zeroes (e.g., in nasalized vowels). The order of an LPC model is the number of poles in the filter. Usually, two poles are included for each formant + 2-4 additional poles to represent the source characteristics. For adult speakers, average formant spacing is in the 1000 Hz range for males and in the 1150 Hz range for females. The LPC order is related to the sample rate of the audio file: 10000 Hz - LPC order = 12-14 (males) and 8-10 (females); 22050 Hz - LPC order = 24-26 (males) and 22-24 (females)

LPC usually requires a very good speech sample. Many recordings done with omnidirectional microphones contain too little speech detail and too much noise to get reliable LPC readings. Figure 8 shows two spectrograms of the word "Shannon" uttered by the same speaker, one recorded with an omnidirectional microphone, the other with a head-worn cardoid one. The corresponding LPC results show quite different formant frequency and bandwidth values. Clearly, the sample taken with a head-worn microphone contains substantially more detail and is more appropriate for analysis. The additional difficulty comes from the fact that the vowel /æ/ in this sample is nasalized due to the following /n/. Only a very high quality recording makes it possible to use LPC in this case. By comparing the values obtained from spectrogram reading to those obtained by LPC, we can safely conclude that the LPC values in 8b are fairly accurate.




Figure 7. LPC analysis of two speech samples. Figure 8a. spectrogram and LPC graph of "Shannon" recorded with a built-in, omnidirectional microphone. Figure 8b. sample recorded with a directional head-worn microphone. LPC parameters (frame length = 20ms, filter order = 12, pre-emphasis = 0.9) were kept constant. LPC results:

| (formant & bandwidth Hz) | F1 | F2 |
|--------------------------|----------|-----------|
| built-in microphone | 871; 138 | 1670; 263 |
| head-set microphone | 521; 146 | 1770; 45 |

5 Conclusion

Computer-assisted analysis of acoustic speech signals can be dramatically improved if appropriate sound acquisition and processing tools and methodologies are used. It is often true that great results can be obtained with simple means and inexpensive hardware. It is our hope that this article will help researchers achieve that goal.

References

- AKG. <http://www.akg-acoustics.com/>
A Division Of Sound Enhancements, Inc. <http://www.swizzarmy.com/EbtechSite/products.htm>
- Harris, M.N. and R. Kelly, D.A.McLeod, M.J.Story. 1998. "Effects in High Sample Rate Audio Material". DCS Ltd.Hubber, David M. and Philip Willams. 1998. Modern Microphone Techniques. Vallejo, CA: Mix Books
- Johan, Lilkenkrants. 1997 Speech Signal Processing. In *The Handbook of Phonetic Sciences*. Eds. Hardcastle William and Laver John. Oxford: Blackwell.Karl, J.H. 1989. *An Introduction to Digital Signal Processing*. Academic Press.
- Kay Elemetrics. <http://www.kayelemetrics.com>
- Plichta, Bartek. 2001. Digitizing Speech Recordings for Archival Purposes. Working Paper. Matrix, Michigan State University.
- Pohlmann, Ken. 2000. *Principles of Digital Audio*. New York: McGraw-Hill.
- Rabiner, Lawrence R., and Ronald W. Schafer. 1978. *Digital Processing of Speech Signals*. Englewood Cliffs: Prentice-Hall.Shure. <http://www.shure.com>
- SONY Corporation. 1993. Minidisk Specifications. Digital Audio Disk Corporation.Sound Devices. <http://www.sounddevices.com>
- Stevens, Kenneth N. 1998. *Acoustic Phonetics*. Cambridge, Mass.: MIT Press. Story, Mike and R. Kelly, D.A.McLeod. 1998. Resolution, Bits, SNR and Linearity. DCS Ltd.Story, Mike. 1997. "A Suggested Explanation for (Some of) the Audible Differences between High Sample Rate and Conventional Sample Rate Audio Material". DCS Ltd.
- Titze, Ingo. 1994. *Workshop of Acoustic Voice Analysis*. National Center for Voice and Speech.
- Vanderkooy, J., and S.P. Lipshitz. 1984 "Dither in Digital Audio". *JAES*, vol.32, no. 11.Winholotz William S. and Ingo R. Titze. 1997. Conversion of a Head-Mounted Microphone Signal into Calibrated SPL Units. *Journal of Voice*. Vol. 11, No. 4, pp. 417-421.