Staining Doppler Audio

Matthew D. Hoffman, Perry R. Cook, and David H. Vilkomerson

***Abstract*—Carotid stenoses are responsible for many of the strokes occurring each year. We have developed a low-cost-of-use instrument for use in primary-care physicians' offices that can detect these stenoses before they produce a stroke. Non-specialists are to use the instrument with only limited training, so it is desirable to make the Doppler audio it produces as easy to interpret as possible. We present two techniques for doing so by processing or "staining" the original audio signal to make it easier for neophyte operators to recognize the information present in the signal. In a small user study, a technique based on adaptive filtering did not result in an improvement, but another technique based on adding a sine tone of varying frequency to the signal produced a significant improvement in accuracy over the unstained Doppler audio. The study also showed that untrained subjects were capable of performing quite well even on unstained audio.**

***Index Terms*—Doppler audio, Biomedical acoustics, Biomedical signal processing, Psychoacoustics**

I. Introduction

Stroke, which kills or disables more than half of the 700,000 Americans it strikes each year [1], occurs without warning symptoms 80% of the time [2]. Many of these strokes are caused by pieces of plaque breaking off from the region of the carotid bifurcation and being carried into small vessels in the brain where they block blood flow and kill the brain tissue. Surgically removing plaque in situations where plaque has reduced the diameter of the blood vessel by 60% or more reduces the stroke rate by two thirds [3].

To prevent strokes in this way, however, we need some way to determine what patients have significant amounts of arterial plaque. Using conventional ultrasound systems, with their skilled operators and physician reading, costs too much [4]. Further, screening procedures that require special appointments, as would be needed for ultrasound examination of the carotids, are typically underused: mammography, for example, is utilized by only about two thirds of the those who should use it, and colonoscopy, which greatly reduces the danger of colorectal cancer, by less than one third. By contrast, screening procedures carried out in physicians' offices, such as for hypertension or PSA (for prostate cancer) have 80% or higher utilization.

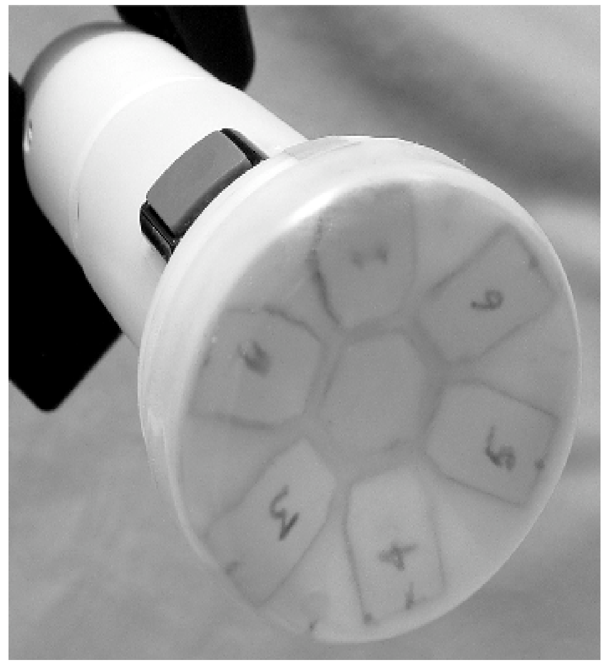

Fig. 1: The screening system scanhead.

To provide a way to detect carotid artery plaque, a special instrument was developed, designed for use in physicians' offices by nurses; it was described in detail in last year's IEEE Ultrasound Symposium [5]. The instrument detects the high blood velocity caused by plaque narrowing the carotid artery. By using a special ultrasound probe with 6 receiving transducers surrounding a large transmitter, angularly spaced at 60°, Doppler signals are received when the transmitter is over the carotid, independent of the orientation of the probe to the artery; the receivers are angled downward so that they receive signals from both deep and shallow carotid arteries. Using the presence of the Doppler signal to guide her, the operator moves the probe along the carotid from the lower neck to the jaw; when a high-frequency signal is heard, she pushes a button and the instrument automatically measures the velocity. No imaging is required.

As reported in [5], 34 carotids were blindly (i.e. before being examined by the conventional duplex Doppler colorflow ultrasound system) examined by this instrument. Its diagnoses of carotid stenoses compared favorably with that of the conventional ultrasound instrument operated by a skilled operator; 91% accuracy (using the conventional system as the gold standard) was achieved.

Based on these results, a screening trial of 1400 patients

Manuscript received October 9, 2001. (Write the date on which you submitted your paper for review.) This work was supported in part by an NIH SBIR grant R44HL072534.

Matthew D. Hoffman is with the Department of Computer Science, Princeton University, Princeton, NJ 08540 USA phone: 609-258-8263; fax: 609-258-1771; (e-mail: mdhoffma at cs dot princeton dot edu).

Perry R. Cook is with the Department of Computer Science, Princeton University, Princeton, NJ 08540 USA (e-mail: prc at cs dot princeton dot edu).

David H. Vilkomerson is with DVX LLC, Kingston, NJ 08528 USA (e-mail: dv1dvx at comcast dot net).

using this instrument has been funded by the NIH. In this trial, the instrument will be operated by nurses in primary-care physicians' offices, rather than by the experienced technologist who used the instrument for the results reported in [5]. The success of Doppler diagnoses is known to be operator-dependent; this raises the question of whether the accuracy reported in [5] be achieved (or approached) by neophyte operators, i.e. can they recognize the highest pitched Doppler signal – indicating the worst stenosis – as they scan along the neck.

One way of helping such operators is to enhance the Doppler signal to make it more "informative". There has been much progress made in understanding how humans perceive sounds; this progress has led, for example, to coding for compression of music signals in MP3 players and Dolby noise reduction. We propose to improve the diagnostic clarity of Doppler signals, i.e. how easily an operator detects the higher frequency Doppler sound that indicates stenosis, by using psychoacoustic understanding and digital computing techniques. In analogy to the way pathologists stain the components of tissue different colors to enhance their ability to recognize diseased tissue, we call this process "staining Doppler audio."

II. Staining Technique

A skilled operator attempting to interpret the raw Doppler audio produced by our system is concerned almost entirely with picking out the highest frequency in the audio spectrum with a significant amount of energy, since this frequency corresponds to the highest velocity with which blood cells are moving through the artery. Although skilled operators are quite good at this task, it is possible that novice operators may be distracted by various other features of the audio, or may lack the ability to make fine distinctions between similar signals. Such factors could make it difficult for novice operators to successfully find the point of worst stenosis.

To address this potential problem, we process the raw Doppler audio signal to make the crucial piece of salient information encoded by the audio – instantaneous velocity – more perceptually obvious and unambiguous. We explored two techniques for accomplishing this goal: adaptive filtering and the addition of a guide tone. Both techniques make use of the 80% spectral rolloff statistic (a simple estimate of the psychoacoustic characteristic of brightness) as a control parameter.

A. Psychoacoustic Background

In psychoacoustics, "brightness" refers to the perception in a sound of high-frequency content relative to low-frequency content. For an operator trying to find the highest frequency present in a sound (such as Doppler audio), we might expect brightness to be the sort of perceptual feature he or she would cue into. In automatic audio analysis systems, acoustic features such as spectral centroid (the mean of all frequencies in the power spectrum, weighted by their respective powers) and spectral rolloff (the frequency below which some amount, e.g. 80%, of the power in the spectrum falls) are often used as proxies for brightness [6]. These automatically extracted statistics have been shown in psychoacoustic studies to correlate well with the perception of brightness [7].

Just-noticeable-difference (JND) psychoacoustic studies attempt to experimentally find the maximum amount that some perceptual feature can be changed without an average listener being able to notice the change. The smaller the JND, the more sensitive listeners are to small changes in that feature. One study found the JND for spectral centroid to be 0.15 centroid units, which equates to 15% of the fundamental frequency of the sound [8]. (It should be noted that this study was done on sounds that, unlike Doppler audio, are harmonic.) By contrast, the JND for pitch perception has been shown to typically be less than 1% [9]. The substantially higher sensitivity of the human ear to small changes in pitch than small changes in brightness suggests that presenting the information in the Doppler audio signal using a pitched tone may be fruitful.

B. Adaptive Filtering

Since the background noise and lowest-frequency content of the Doppler audio are for the most part not relevant to the operator's task of finding the segment of the carotid artery with the highest-velocity blood flow, our first staining technique involved filtering out these components to prevent them from distracting novice operators. We used a two-pole, two-zero digital bandpass filter to emphasize the frequency content around the 80% rolloff frequency, which was calculated over a 102 ms window every 51 ms. The filter was determined by placing two complex conjugate poles were placed at an angle corresponding to the 80% rolloff frequency and a radius of 0.9 from the origin. Two zeros were placed at $z = 1$ and $z = -1$ to produce a constant unity peak gain. Spectrograms comparing the original and filtered audio can be seen in figure 2.

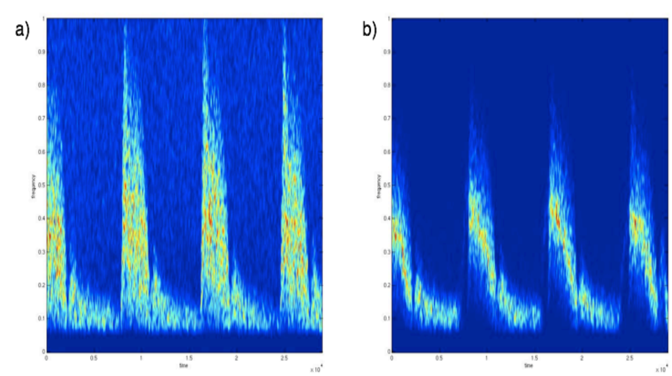

Fig. 2: Spectrograms of unstained (2.a) and filtered (2.b) Doppler audio. Time is shown in samples, frequency is shown as a fraction of the Nyquist frequency of 5000 Hz.

C. Guide Tone

Since (as discussed above) human beings are capable of finer pitch discrimination than brightness discrimination, for our second staining method we superimpose on the raw Doppler audio a sine tone whose frequency maps to the 80%

spectral rolloff of the Doppler audio (as calculated above). Since human pitch perception is roughly logarithmic, whereas the mapping of velocity to frequency in the raw audio is linear, we map rolloff values of 0-1 exponentially to pitch

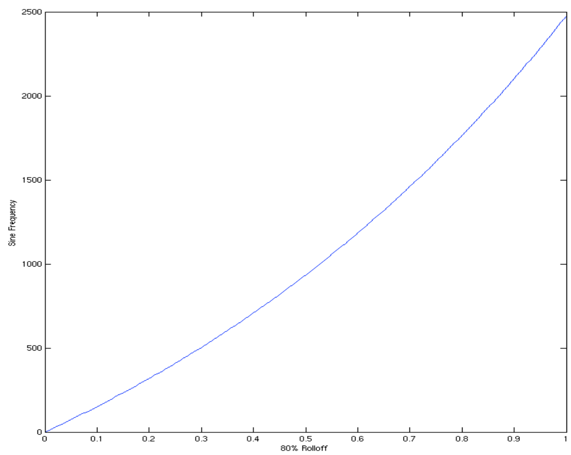

Fig. 3: Exponential mapping between calculated 80% rolloff value and frequency in Hz of synthesized sine tone.

values in a range between 50 and 2500 Hz. Figure 3 shows this relationship, and Figure 4 shows the spectrogram of an untreated recording of Doppler audio, a graph of the 80% rolloff values calculated (as above) every 51 ms, and the spectrogram of the original Doppler audio with a sine tone following the rolloff added.

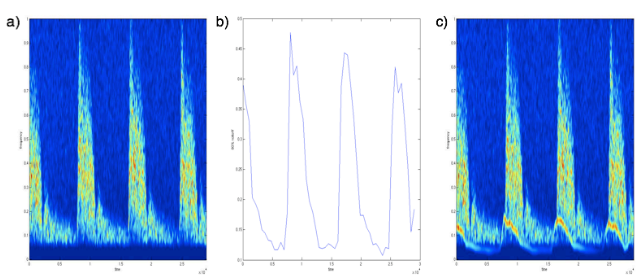

Fig. 4: Spectrogram of unstained Doppler audio recording (4.a), plot of the calculated 80% rolloff values over time for the same recording (4.b), and spectrogram of the original recording with a sine tone added as described above (4.c). Time is shown in samples, frequency and rolloff are shown as a fraction of the Nyquist frequency of 5000 Hz.

III. Experimental Setup

We conducted a user study to evaluate the abilities of inexperienced individuals to interpret both stained and unstained Doppler audio.

A. *Test Data Set*

We began with 100 3-second recordings of Doppler audio taken from 34 carotid arteries of actual patients at a sampling rate of 10 KHz. A test set of 27 pairs of recordings was chosen from the original 100. Each pair consisted of two recordings taken from the same patient and artery. Additionally, we required the measured velocities associated with the two recordings to differ by at least 15 cm^3/s to prevent possible measurement errors from corrupting the experimental data.

These 27 pairs were randomly divided into 3 subsets of 9 pairs. For each subject, one set was left unstained and the other two were processed as described above (i.e., the second set was adaptively filtered, and a guide tone was added to the third set). The staining method for each subset varied from subject to subject.

B. *Subjects*

We tested 11 subjects in total. Subjects were mostly science and engineering graduate students in their mid to late twenties, as well as two older computer scientists in their 40s and 50s. Two subjects were female, the other nine male. Two subjects suffered from some amount of hearing loss. Levels of musical training and critical listening experience varied. However, an effort was made to seek out subjects without such training. None of the demographic factors enumerated above had any measurable impact on subjects' performance. None of the subjects had any previous experience evaluating Doppler audio.

C. *Test Interface*

A simple testing interface was developed in OpenGL to facilitate the experiment. A screenshot of the interface can be found in figure 5. The screen shows a simple representation of the bifurcation of a carotid artery, with buttons on either side of the bifurcation. When the subject runs the mouse over one of these buttons, a looped recording of stained or unstained Doppler audio plays. The subject compares the two recordings associated with the two buttons, and clicks on the button associated with the recording that the subject estimates to be of higher velocity. The subject's answer is recorded, a message is printed to the console saying whether or not the subject answered correctly, and the next pair of recordings is automatically presented.

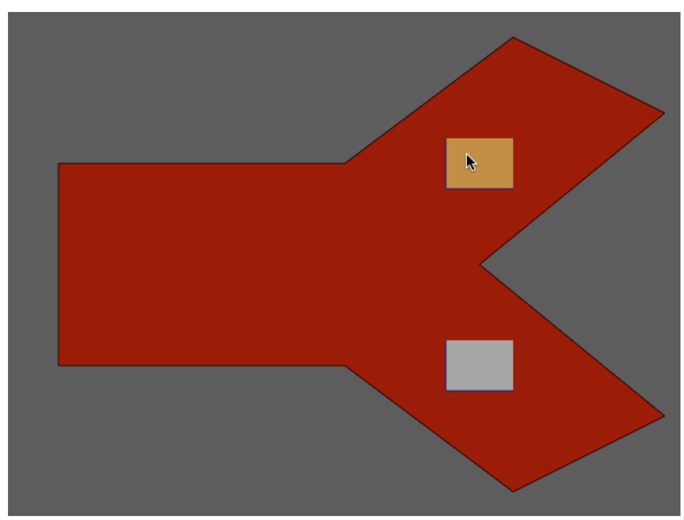

Fig. 5: Screen capture of the test interface.

D. *Testing Procedure*

Each subject was briefly given some background information about the project, and then presented with several recordings of Doppler audio not included in the test set and their measured velocities to give the subject a sense of how high- and low-velocity Doppler audio sounds. The subject was then shown how to use the test interface, and went through three practice trials of unstained recordings, filtered recordings, and recordings with guide tones added. These practice rounds used pairs drawn from the same set of demonstration recordings mentioned above, and the whole process of playing demonstration recordings and running practice rounds in general took less than 5 minutes.

Once the subject had been given a chance to practice comparing stained and unstained recordings and get used to the interface, the actual experiment began. Using the test interface, the subject first compared a subset of 9 pairs of unstained recordings, then a subset of 9 pairs of filtered recordings, then a final subset of 9 pairs of recordings with guide tones added. The subsets were varied and selected as described above. This portion of the experiment typically took between 5 and 10 minutes.

All audio was presented using an inexpensive pair of Sony MDR-CD60 headphones. Subjects were tested in an office environment with significant ambient machine noise, such as one might expect to be present in a physician's office.

IV. Results and Discussion

Table 1 presents the combined accuracy rates for the 11 subjects. On average, subjects were capable of correctly detecting the higher-velocity carotid based on the unstained audio 8 out of 9 times, which suggests that operators may be able to use the system effectively without extensive training. The results are clearly much better than chance for both unstained and stained audio, but adding a guide tone leads to a 45% decrease in error over the unstained audio ($p = 0.03$), and filtering the audio actually leads to a 27% increase in error ($p < 0.21$).

TABLE I
Accuracy Statistics for All Subjects

	No Stain	Adaptive Filter	Guide Tone
Subject 1	9 / 9	7 / 9	7 / 9
Subject 2	8 / 9	9 / 9	9 / 9
Subject 3	9 / 9	7 / 9	9 / 9
Subject 4	8 / 9	8 / 9	8 / 9
Subject 5	7 / 9	9 / 9	9 / 9
Subject 6	8 / 9	7 / 9	8 / 9
Subject 7	8 / 9	7 / 9	8 / 9
Subject 8	7 / 9	8 / 9	9 / 9
Subject 9	8 / 9	7 / 9	8 / 9
Subject 10	8 / 9	8 / 9	9 / 9
Subject 11	8 / 9	8 / 9	9 / 9
Average Subject	8 / 9 = 88.89%	7.73 / 9 = 85.86%	8.45 / 9 = 93.94%
Standard Deviation	0.633 / 9 = 7.03%	0.786 / 9 = 8.73%	0.688 / 9 = 7.64%

The poor performance of subjects listening to the adaptively filtered audio suggests that the filtering process removes important information from the signal. This could perhaps be addressed by increasing the bandwidth of the filter, replacing it with a lowpass filter that would leave the highest frequencies intact, or using a different algorithm to control the center frequency of the filter.

Subjects' reactions to the different versions of the Doppler audio varied substantially. Some, for example, found the guide tone variously distracting, annoying, or too quiet. One found all of the sounds irritating. Others (some of whom did perfectly on the unstained audio) said that the guide tone was helpful and made the task easier.

This experiment did not in any way measure or control for the speed or confidence with which subjects completed their task. This may be worth exploring in the future, since patients are likely to react better to an operator who can screen them with speed and confidence.

V. Conclusions

We implemented and tested two techniques for staining Doppler audio to make novice operators better able to make effective use of the carotid screening system described in [5]. Although the adaptive filtering technique did not produce an improvement in subject performance, the addition of a guide tone led to a significant reduction of subjects' already low error. Our results suggest that even novice operators may be able to use our system to effectively screen patients for stroke risk due to stenosis in their carotid arteries, especially if the audio is enhanced by leveraging the fine pitch resolution of the human auditory system.

References

[1] American Heart Association, "Heart Disease and Stroke Statistics – 2006 Update" [Online]. Available: http://circ.ahajournals.org/cgi/content/short/113/6/e85

[2] North American Symptomatic Carotid Endarterectomy Trial Collaborators, "Beneficial effect of carotid endarterectomy in symptomatic patients with high-grade stenosis," *New England Journal of Medicine*, 325:445-453, 1991.

[3] A. Halliday, A. Mansfield, J. Marro, C. Peto, R. Peto, J. Potter, et al., "Prevention of disabling and fatal strokes by successful carotid endarterectomy in patients without recent neurological symptoms: randomized controlled trial," *The Lancet*, 363(9420): 1491-502, May 2004.

[4] M. Benade, C. Warlow, "Costs and benefits of carotid endarterectomy and associated preoperative arterial imaging – A systematic review of health economic literature," *Stroke*, 23:629-38, 2002.

[5] D. Vilkomerson, T. Chilipka, R. Outcault, K. Goldman, "An instrument for screening for carotid stenoses," in *Proc. 2005 IEEE International Ultrasound Symposium*, pp. 397-402.

[6] G. Tzanetakis, P. Cook, "MARSYAS: a framework for audio analysis," *Organized Sound*, Vol. 4 Num. 3, pp. 169-175, Dec 1999.

[7] E. Boring and S. Stevens, "The nature of tonal brightness," *Proc. Of the National Academy of Science*, 22, pp. 514-521, 1936.

[8] R. A. Kendall and C. Carterette, "Difference threshold for timbre related to spectral centroid," in *Proc. Of the 4th International Conference on Music Perception and Cognition*.

[9] J.G. Roederer, *Introduction to the Physics and Psychophysics of Music*. New York, New York: Springer, 2001.