

Efficacy and effectiveness of a pattern-recognition algorithm

By Shilpi Banerjee, Laurel Olson, Karrie Recker, and Justyn Pisa

The MarkeTrak survey has shown that hearing aid users are least satisfied with their devices in noisy environments.¹ While speech is arguably the single most important signal in a listener's environment, amplification designed to maximize speech intelligibility may be ineffective in certain noisy conditions, which may decrease user satisfaction in those environments. Although hearing aids can often be adjusted manually with the push of a button or turn of a wheel, many people are unable or reluctant to "fuss with their hearing aids."

From the hearing aid user's perspective, the goals are simple: maximum intelligibility for speech with minimum interference from background noise, and optimum listening comfort and sound quality in all environments. Given that satisfaction with amplification is strongly correlated with the number of listening situations in which the user perceives a benefit,¹ it is in the best interests of clinicians and consumers alike to increase user satisfaction in a wide variety of environments.

Traditional approaches to noise management use information about level, modulation depth, or harmonicity to make decisions regarding the nature of the incoming sound.

Specifically, intense, unmodulated, or non-harmonic sounds are considered to be noise. The shortcoming of such methods is that they may erroneously classify a desired signal as noise. For example, music may be treated as noise because it is loud and has a steady temporal envelope; similarly, certain noises may be harmonic.

Pattern recognition is a scientific discipline whose goal is to classify objects into a number of categories or classes.² These objects can take various forms, including images, odors, and sounds. Voice-activated dialing and automatic transcription of speech to text are among current applications of pattern recognition.

Although efforts to classify acoustic signals are under way in many areas of audio processing, there are several challenges unique to the use of pattern recognition in hearing aids: (1) degradation of the input by noise and/or reverberation in the environment, (2) difficulty isolating the desired signal from the undesirable noise, (3) similarities in the characteristics of signal and noise, (4) exposure to an infinite variety of environments, (5) lack of opportunities for learning the characteristics of the signal and/or noise under ideal conditions, and (6) practical limitations on the

processing capabilities and battery drain of a hearing aid. (For comparison, note that the average cellular phone has roughly 10 times the processing capability and consumes about 10 times as much current as the most advanced hearing aids.)

This article describes the application in hearing aids of a pattern-recognition algorithm, known as Acoustic Signature, that is designed to differentiate among various types of acoustic stimuli in the environment and to apply this knowledge automatically.

ACOUSTIC SIGNATURE

As the capabilities of digital signal processing (DSP) in hearing aids have increased, so has the complexity of noise-management schemes. Acoustic Signature is the application of pattern recognition to acoustic signals arriving at the microphone(s) of the hearing aid. The algorithm uses several features of the input signal in making decisions, which dramatically decreases the error rate and improves the reliability and validity of the categorization.³

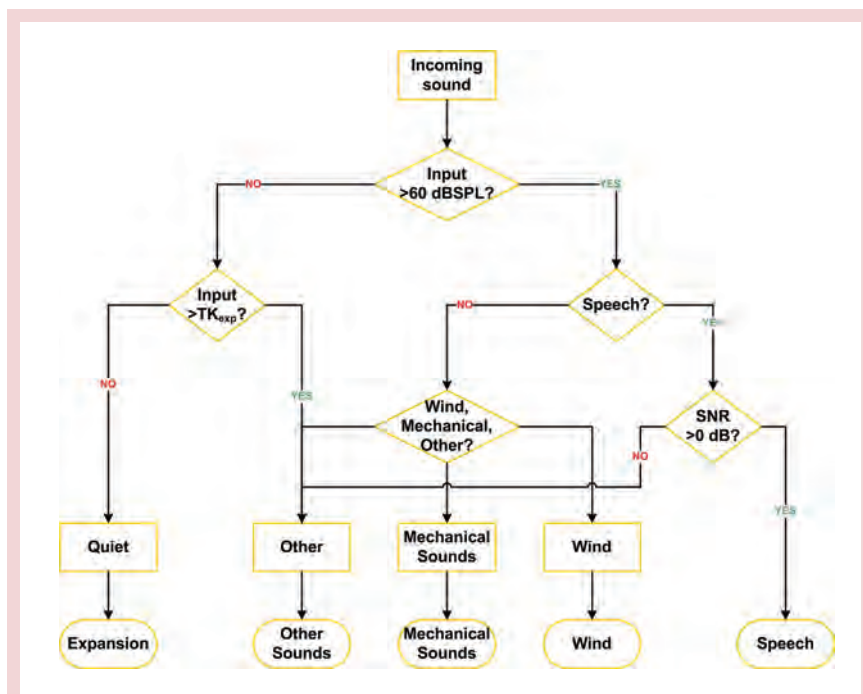


Figure 1. Classification scheme underlying Acoustic Signature. *TKexp* = expansion threshold, *SNR* = signal-to-noise ratio.

Acoustic Signature is based on the classification scheme shown in Figure 1. Sounds entering the microphone of the hearing aid are placed in one of five categories: *quiet*, *speech*, *wind*, *mechanical sounds*, and *other sounds*. These categories represent environments in which hearing aid users are often dissatisfied.

Classification of *quiet* is based on the input level, that is, inputs below the expansion threshold are considered *quiet*. The remaining categories rely on the spectral and/or temporal characteristics of the input. *Mechanical sounds* refer to sounds from machines, such as a vacuum cleaner, blender, or car. *Other sounds* are those not specifically identified in one of the other categories. Running water is an example. Speech that is unusable, due to a poor signal-to-noise ratio (SNR), is also classified as *other sounds*.

The primary purpose of Acoustic Signature is to preserve listening comfort in noisy environments without adversely affecting speech understanding. Thus, gain adaptation occurs for sounds classified as *quiet*, *wind*, *mechanical sounds*, or *other sounds*. In contrast, gain adaptation is not applied when *speech* is detected. It should be noted that decisions regarding the application of directionality to optimize speech understanding, i.e., Directional Speech Detector, are made independently of this pattern-recognition algorithm.

The following discussion describes efficacy and effectiveness of the Acoustic Signature algorithm. Because the rules for detection of quiet environments are simple and the clinical application of expansion

is well documented,⁴ this article focuses on the accuracy of classification and the appropriateness of adaptation for the remaining categories of sounds.

CLASSIFICATION ACCURACY

The first step in the clinical application of Acoustic Signature is to verify the efficacy of the algorithm itself. In other words, does it accurately classify inputs to the hearing aid as *speech*, *wind*, *mechanical sounds*, or *other sounds*?

Methods

To answer that question, 18 subjects, fitted bilaterally with BTE, ITE, ITC, or CIC Starkey Destiny 1200 hearing aids, participated in an experiment. While directional microphones were available on all styles except the CICs, the devices were tested in the omnidirectional mode. All other adaptive DSP features, including expansion, feedback cancellation, and dynamic directionality, were turned off during the evaluation.

Figure 2 shows the laboratory setup. The subject sat in the center of an eight-speaker array, marked by the "X" in the figure. The acoustic stimuli included average speech in quiet, loud speech in quiet, loud speech in noise at +5 dB SNR, a vacuum cleaner, a blender, and a treadmill. Each stimulus was presented randomly through speakers 1, 4, or 6 (0°, 135°, and 225° azimuth, respectively). Wind was generated by a quiet table fan that replaced speaker 1 at 0° azimuth. The medium and high settings were used to represent different wind speeds. Further, off-axis wind

directions were simulated by having the participant face speaker 2 (45° azimuth) and speaker 8 (315° azimuth). This was done at the high setting only.

The hearing aids were connected to the programming software throughout the evaluation so their state could be monitored at all times. The programming software contains buttons that are highlighted each time a certain category of sounds is detected, one each for *speech*, *wind*, *mechanical sounds*, and *other sounds*.

Results

There are two, equally important, aspects to analyzing the classification outcome: (1) appropriate categorization of the stimuli, and (2) the tendency for stimuli to be misclassified in a certain direction.

Consider, for example, the situation in which speech is classified as *speech* a high percentage of the time. This is a good outcome only if wind and mechanical sounds are not also frequently misclassified as *speech*. The misclassification of other categories of sounds as *speech* indicates an inherent bias within the algorithm toward that category. In other words, it is relatively easy, but not helpful, to obtain a high degree of accuracy for one category of sounds at the expense of accuracy for other categories.

Figure 3 shows the confusion matrix for the detection accuracy of Acoustic Signature. As expected, it classifies *speech*, *wind*, and *mechanical sounds* appropriately most of the time. However, it appears to have a significant tendency to classify speech and wind as *other sounds*.

Further investigation of the misclassified data revealed that in all but two cases, the overall power of the sound was at or below 60 dB SPL. As shown in Figure 1, the classification scheme is designed to classify inputs at 60 dB SPL or below as *quiet* or *other sounds*, depending on whether or not the expansion threshold is exceeded. All inputs at or below 60 dB SPL were classified as *other sounds* because expansion was disabled for this evaluation. Thus, the classification outcome was consistent with the design of the algorithm and evaluation.

Figure 4 summarizes the classification accuracy of Acoustic Signature. Accuracy is determined according to whether or not the classification was consistent with the design of the algorithm. Speech is arguably




Figure 2. Laboratory setup for assessment of detection accuracy. Left = setup for speech stimuli and mechanical sounds, right = setup for wind.

		CLASSIFICATION			
		SPEECH	WIND	MECHANICAL SOUNDS	OTHER SOUNDS
STIMULUS	SPEECH	75% (n=135)	1% (n=2)	<1% (n=1)	23% (n=41)
	WIND	<1% (n=1)	58% (n=84)	0% (n=0)	41% (n=59)
	MECHANICAL SOUNDS	1% (n=2)	0% (n=0)	87% (n=190)	12% (n=26)

Figure 3. Confusion matrix for detection accuracy of the algorithm.




Figure 4. Classification accuracy of Acoustic Signature.

the most important signal to which hearing aid users are exposed. The 95% classification accuracy for speech indicates that the transmission of usable speech signals to the ear is maximized. At 88%, mechanical sounds were also categorized with a high degree of accuracy. Lastly, the 92% accurate classification of wind is particularly noteworthy because it includes omnidirectional CICs.

The turbulent nature of wind makes it relatively easy to detect in directional hearing aids; the noise at the two microphones is uncorrelated and the resulting directional output is substantially higher than for the omnidirectional output. On the other hand, in an omnidirectional hearing aid, detection relies solely on the spectral and temporal characteristics of the wind. Overall, the evaluation verified excellent classification accuracy for all categories of sounds.

ADAPTATION

The next step in the clinical application of Acoustic Signature is validating its effectiveness, that is, ensuring that the hearing aid user perceives the value of the algorithm in realistic environments. This was evaluated in a field study of Destiny 1200 involving 12 external sites and 70 participants. For a full review of the study, see Olson and Pisa.⁵

Methods

Most participants had bilaterally symmetrical sensorineural hear-

ing loss, but there were four exceptions: three subjects with bilateral asymmetries greater than 15 dB and one with a bilaterally symmetrical mixed loss. Approximately 75% were experienced users of amplification. Participants ranged in age from 25 to 79 years, with a mean of 71. Four device styles—BTE, ITE, ITC and CIC—were evaluated and equally represented across participants. Figure 5 shows the mean audiometric configuration of all participants, grouped by device style.

Participants were fitted using the NAL-NL1* fitting formula, a proprietary modification of the NAL-NL1 prescription,⁶ which overestimates the required gain.⁷

The devices were programmed with two active memories, one of them exclusively for use on the telephone. Participants were instructed to use Memory 1, programmed as “Normal,” at all times. All adaptive DSP features, including Acoustic Signature, Directional Speech Detector, and Active Feedback Intercept, are active by default in the “Normal” memory.

During the first 2 weeks of the study, subjects were asked to evaluate the hearing aids under various listening conditions. The environments included blender, vacuum cleaner, wind, car and water. In each environment, participants rated their preference for gain adaptation on a 5-point scale. Specifically, they were asked, “Should the noise be: (1) much softer, (2) softer, (3) just right, (4) louder, or (5) much louder?”

Results

Based on the design of the Acoustic Signature algorithm, we expected that blender and vacuum cleaner would be classified as *mechanical sounds*, wind would be identified as *wind*, car would be classified as *mechanical sounds* or *wind* (depending on whether or not the windows were rolled down), and water would be categorized as *other sounds*.

The participants’ responses were converted into numerical values such that *much softer* corresponded to a rating of 1, *just right* to a rating of 3, and *much louder* to a rating of 5. As shown

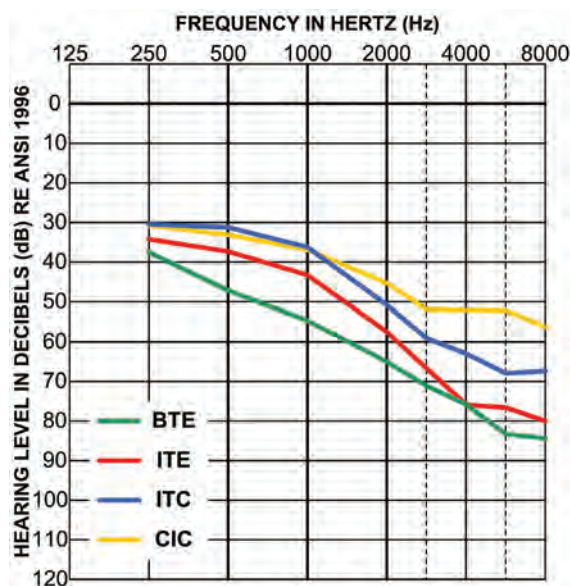


Figure 5. Mean audiograms (bilateral) of all participants grouped by device style.

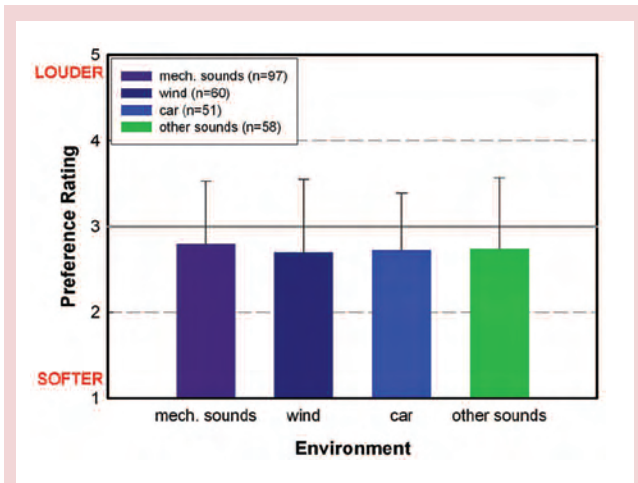


Figure 6. Mean preference ratings for gain adaptation of Acoustic Signature in various environments. The error bars represent one standard deviation from the mean. The reference line at 3 corresponds to a rating of “just right.”

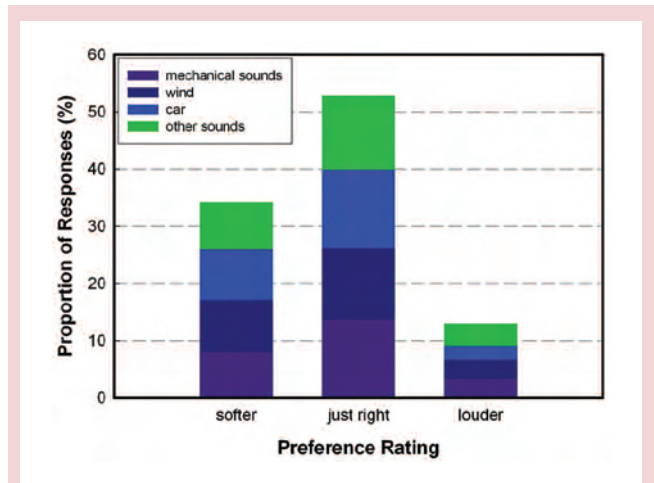


Figure 7. Preference ratings of participants who rated the amount of adaptation for each Acoustic Signature. Ratings of much softer and softer have been combined into a single category; the same is true for ratings of louder and much louder.

in Figure 6, a mean rating of approximately 2.8 was obtained for all classes of sounds. The fact that, on average, participants perceived the noise level to be just right indicates appropriate gain adaptation.

The variability in preference ratings is of as much interest as the mean scores.

Figure 7 shows the distribution of preference ratings across the various categories of sounds. Three key observations can be made from these data:

First, more than half the subjects reported that the amount of gain adaptation provided by the algorithm in the

study was appropriate. Second, approximately 35% of participants would have preferred the noise to have been *softer*, i.e., more gain adaptation. Third, a much smaller proportion of participants (approximately 10%) reported a preference for *louder* noise, i.e., less gain adaptation.

The absence of a one-size-fits-all solution is not surprising. Some of these differences may be attributed to individual variability. In addition, previous experience with amplification may play a substantial role in this outcome.

At the upper end of the continuum might be someone experienced with linear amplification who is accustomed to intense sounds being loud. Such a person may reject gain adaptation because it makes intense noises, such as a vacuum cleaner, less loud. Similar reasoning may prompt this person to reject wide dynamic range compression (WDRC).

At the lower end of the continuum are new users of hearing aids and those less tolerant of noise and/or amplified sound in general. One important conclusion to draw from these data is the need for flexibility in gain adaptation. In other words, the hearing professional must be able to adjust the amount of gain adaptation to meet the individual client's preferences. While the programming software did provide options for the amount of gain reduction, this feature was not used in this study because its goal was to evaluate the appropriateness of the default setting.

Although it's not apparent from the quantitative data, the car proved to be the trickiest environment for the Acoustic Signature. Whether participants loved it or hated it depended on *what* they wanted to hear. Some were glad that the car was quieter while others were displeased that they could not always hear the radio or another passenger. One subject was happy that the devices became quieter with two sets of twins in the back seat!

A car is a complex listening environment in which numerous objective and subjective factors influence the outcome. A memory dedicated for use in the car may be warranted to address the specific needs of individual hearing aid users.

SUMMARY AND CONCLUSIONS

The basic conclusions of the study on the Acoustic Signature algorithm are as follows:

(1) Despite realistic constraints on hearing aids, the algorithm uses the principles of pattern recognition to categorize acoustic stimuli and automatically applies that knowledge.

(2) It classifies inputs to the hearing aid as *speech*, *wind*, *mechanical sounds* or *other sounds* with a high degree of accuracy.

(3) The amount of gain adaptation provided by Acoustic Signature is appropriate for most hearing aid users. The programming software, Inspire OS, offers additional options for fine-tuning the adaptation for individual needs.

Shilpi Banerjee, PhD, is Manager, Algorithm Research; **Laurel Olson**, MA, is Manager, Clinical Product Research; and **Karrie Recker**, AuD, and **Justyn Pisa**, AuD, are both Research Audiologists, all at Starkey Laboratories, Inc. Readers may contact Dr. Banerjee at shilpi@starkey.com.

REFERENCES

1. Kochkin S: MarkeTrak VI: Factors impacting consumer choice of dispenser and hearing aid brand; Use of ALDs and computers. *Hear Rev* 2002;9(12):14-23.
2. Theodoridis S, Koutroumbas K: *Pattern Recognition*. San Diego: Academic Press, 2003.
3. Duda RO, Hart PE, Stork DG: *Pattern Classification*. New York: John Wiley and Sons, Inc., 2001.
4. Plyler P, Hill A, Trine T: The effects of expansion on the objective and subjective performance of hearing instrument users. *JAAA* 2005;16:101-113.
5. Olson L, Pisa J: Clinical trial of the nFusion hearing system. *Hear Rev* 2006, in press.
6. Dillon H, Katsch R, Byrne D, et al.: The NAL-NL1 prescription procedure for non-linear hearing aids. *National Acoustics Laboratory Annual Report* 1997: 4-7.
7. Smeds K: Is normal or less than normal overall loudness preferred by first-time hearing aid users? *Ear Hear* 2004;25(2):159-172.