

The trainable hearing aid: What will it do for clients and clinicians?

By Harvey Dillon, Justin A. Zakis, Hugh McDermott, Gitte Keidser,
Wouter Dreschler, and Elizabeth Convery

INTRODUCTION: THE PROBLEM

The extreme flexibility of modern hearing aids is presenting an increasing challenge to designers and clinicians: How should all the adjustable, or potentially adjustable, amplification parameters be set to make life best for the intended wearer?

Of course, all hearing aid fittings begin with a prescription of some sort. We aim to measure something about the client (often just the audiogram) and use this information to deduce how each amplification parameter (e.g., gain for low-frequency, high-intensity sounds) should be set. Prescriptive formulas are intended to make the prescription right for the average person with a particular set of measured characteristics (such as an audiogram).

Inevitably, however, some individuals will prefer settings that are different from average. Also, prescriptive formulas are derived from some combination of a theoretical rationale and experimental data, and possibly some formulas do not even achieve their goal of prescribing parameters that are right on average, let alone for any given individual.¹ Some parameters (e.g., speed of automatic noise reduction) are not individually prescribed at all. The designer of the hearing aid or of the fitting software sets them to the same value for all wearers, even though there may be little evidence to guide what that value should be.

For all these reasons, it is relatively common for a hearing aid to be better matched to the individual needs of its wearer if adjustments are made *away* from the prescribed response. That is, most clinicians should consider the prescribed response to be a reasonable starting point, not the end, of the adjustment journey.

This raises the difficult issue of how further adjustment should be performed. Of necessity, the clinician has made this adjustment, via the programming software, in the clinic. But there are many reasons why such an adjustment process may not lead to an optimal fitting:

- ❖ The clinic is usually a low-noise, low-reverberation environment, and often the only stimulus used is the clinician's voice. At best, such an adjustment can lead to an optimal fitting only for environments with these same ideal acoustic characteristics. It is well established that the optimal amplification characteristics vary with the acoustic environment.²
- ❖ Other stimuli can be used, but to simulate every situation in which people wear hearing aids would require a seemingly endless set of combinations. These include

noise at different levels, different degrees of reverberation, different types of non-speech sounds of interest, different directions of incoming sound (both sounds of interest and competing sounds), and different rates of change of any of the preceding items. While any individual aid wearer may wish to have the hearing aids optimized for operation in only a handful of places of particular importance, it would be a daunting task to try to approximate even these situations in the clinic. Even if it were possible to make the acoustic simulation realistic, the visual environment would not be, which might well affect the adjustment that is optimal.³

- ❖ If a small set of realistic environments could be created in the clinic, it is far from straightforward to structure the hearing aid adjustment process so that adjustments made when the client is listening to one sound do not undo those made while the person was listening to previous sounds. The result may well be an end adjustment that is worse than the prescribed starting point.
- ❖ Even if a converging adjustment occurs, the clinical time involved would be prohibitive. This will exacerbate the coming problem where, due to aging of the population in most developed countries, there will be more need than ever to streamline the clinical process if available clinical resources are to meet the enormous growth in demand for services that will occur over the next 25 years.

Currently, hearing aid wearers do not actually know if their hearing aids are optimally adjusted in any situation in which they use them. If the sound of the aid is so poor they find it unacceptable, they will return to the clinic, either to return the hearing aid(s) or to ask that they be adjusted. The clinician has to infer the acoustics of the situation, understand the nature of the client's dissatisfaction with the sound, and deduce (sometimes with software assistance) the parameter to be adjusted, the direction of the adjustment, and the extent of the adjustment.

Obviously, this process may go wrong or need to be repeated several times, especially if the client continues to experience new environments with different acoustic characteristics and perceive that the sound is not always quite right.

So, we are left with a problem: How can we optimally adjust the hearing aid? The remainder of this article describes a new concept, the trainable hearing aid, that can address




Figure 1. Hypothetical data showing the gain (averaged across frequency) preferred by a hearing aid user in six situations versus the overall sound pressure level in each situation

THE CONCEPT AND RANGE OF TRAINABILITY

With a trainable hearing aid, the wearer teaches the instrument how it should be adjusted. The wearer does so by using the aid in those situations in which he or she would like assistance with hearing. Therefore, the process takes place after the client leaves the clinic. Once the clinician explains the process to the wearer, the clinician need not spend further time on the process.

Trainable hearing aids will vary in their complexity and effectiveness. The first generation of device will likely be capable of learning only one thing: the volume control setting preferred by the client. A memory inside the instrument will take note of the setting last used by the client. When the hearing aid is next switched on, it will automatically set the volume to some combination of this most recent setting and the settings it remembers from the previous switch-on and/or from other times the hearing aid was used.

If the hearing aid simply set the volume control to the setting last used, this would return volume control adjustment to how it was before digital hearing aids. Gain was determined by the physical position of the volume control and by the operation of any non-linear features (e.g., compression, noise reduction). Instead, the first generation of trainable hearing aids will give the user the average of the volume control settings the person has adopted previously. Thus, it will not be affected so much by whether the user was in a particularly quiet or particularly noisy place when he or she last adjusted the hear-

A greater level of sophistication will enable the trainable aid to match the volume setting the client selected to the acoustic environment at the time the adjustment was made. For example, the hearing aid could measure and record the overall level of sound in the environment for the few seconds prior to the volume control being adjusted. After being adjusted several times

in different environments, the hearing aid develops a picture of how much gain the client prefers in each environment.

Figure 1 shows what this might look like after just six adjustments of the hearing aid. As would be expected, the wearer generally prefers less gain (averaged across frequencies) as the environment gets louder. Armed with this information, how much gain would the ideal hearing aid automatically select if the user were then to enter a new environment with an overall level of, say, 80 dB SPL? The hearing aid could estimate this by fitting a curve (or in this case a two-segment straight line) to the data available to it. From the line in Figure 1, the hearing aid would deduce that the wearer is likely to select a gain of about 12 dB in this new situation, even though the user has never provided any training for this specific input level.

If the trainable aid does its job well enough, the wearer will not need to touch the volume control in this situation. Although the wearer has had to do only a single task—varying the volume control—the hearing aid has been able to deduce how the gain should vary as the input level changes. That is, the individ-

ual gain adjustments have trained the device to have an individually optimized low-level gain (in this case 24 dB), compression ratio (1.7), compression threshold (55 dB SPL), and, therefore, absolute gain for any input level.

Figure 2 shows a general block diagram for one form of a trainable hearing aid. For the relatively simple version of the trainable aid discussed above, the *acoustic measurement module* would consist of a simple sound-level meter that calculates overall sound pressure level averaged over a few seconds. The *programmable amplifier* is no different from the (digital processing) amplifier in any current non-linear hearing aid. It's just that its programming inputs remain permanently (and internally) connected to the *learning algorithm* module rather than being disconnected once the clinician has finished programming the hearing aid.

The simplest way to think about the learning algorithm module is that it maintains a record of user adjustments and the corresponding acoustic environments, along with the appropriate statistical or mathematical processes to deduce the set of amplification parameters that best match the user's preferences. In practice, there are ways to achieve this without requiring the learning algorithm to memorize every single adjustment and corresponding set of environment acoustics. The *user control* can be a simple rotary control, a toggle, or a pair of up-down buttons, and the control can be mounted on the hearing aid itself or on a remote control. The *microphone* and *receiver* are entirely conventional.

Note that, based on the data in Figure 1, the relationship between the gain preferred and the input level is imperfect in that the points deviate from the fitted curve. For example, on the two occasions when the overall level in the environment was 60 dB SPL, what caused the wearer

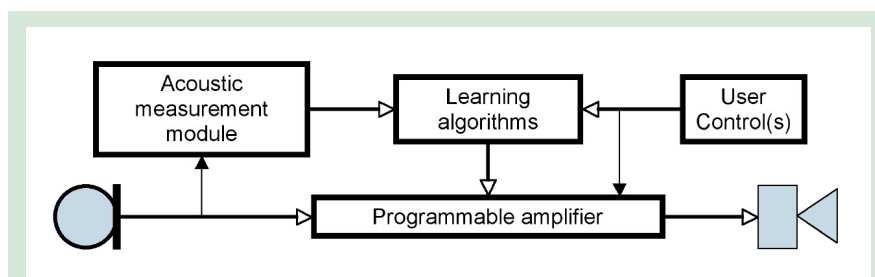


Figure 2. Block diagram for one general form of a trainable hearing aid.

to select a 17-dB gain on one occasion but 25 dB on the other? Perhaps on the first occasion the dominant sound was rain on the roof that the person did not want to listen to, while on the second it was speech that the client particularly wanted to understand. Consequently, another level of sophistication for trainable hearing aids will be to measure much more about the environment than just the overall level of sound.

For example, hearing aids could measure the spectral shape, the rate and extent to which the spectral shape varies, the apparent direction of sound, the modulation of sound within each frequency channel, and the way these modulations match each other across channels. These characteristics could be individually input to the learning algorithm, or they could be combined to classify environments in various ways, and the resulting classifications input to the learning algorithm.

Obviously, a wide range of acoustic characteristics of sound might influence the type and degree of amplification the user would prefer. It is hard to conceptu-

alize the relationships in a simple graph, but inferential statistics can be used to derive the link between these acoustic characteristics and the amount of gain the client chose. Although the hearing aid can never actually read the client's mind (at least not in this decade), the hearing aid is able to "hear" everything that the client hears and take into account many of the factors that the client uses in deciding how far the hearing aid should be turned up or down.

The trick to doing this automatically is that the instrument must first learn which factors are actually important (in that they seem to affect what the client prefers) and what the link is between the value of each factor and the gain preferred by the client. The necessary statistical operations are well within the signal processing capabilities of the hearing aids that will be appearing over the next few years. Many of the more sophisticated acoustic parameters mentioned here are already being measured in current hearing aids, and the measurements are used to control amplification in a way determined by the hear-

ing aid designer or by the clinician. By contrast, in the trainable aid, the client determines the way these acoustic parameters affect the amplification provided.

ADJUSTING THE GAIN-FREQUENCY RESPONSE

So far we have talked only about the hearing aid automatically—but very intelligently—adjusting the gain in the same way as the user would adjust the volume control as situations changed. A hearing aid that almost always gave the loudness that the user preferred would be a great advance over existing hearing aids.⁴⁻⁶

However, as clinicians know only too well, there are many other things to adjust in a hearing aid. For example, the amount of amplification has to vary with frequency, and the amount at each frequency has to vary with input level, usually by different degrees. Therefore, the next level of sophistication for the trainable hearing aid is the individual adjustment of the shape of the gain-frequency response.

At first sight, it might seem necessary to present the user with something like

an audio mixing console, or at least a set of tone controls, to achieve this. However, the trainable aid has to be simple enough for people to use it with minimal instruction.

One solution, which we have successfully used in some of our experimental work with trainable aids, is to have a single control that takes on different functions from time to time.⁷ For example, it might be a volume control when first adjusted, then become a treble control, which turns into a bass control, before once again becoming a volume control.

If the user had to understand exactly what was occurring at any time, this would be a nightmare for them to use. However, if the instruction to the wearer is simply, "Turn this control to the position where the hearing aid sounds best, and then leave it there for 10 seconds or more," the client actually has a simple task to perform, even though the amplification parameter being trained varies from time to time.

A different solution is to give the user a small number of controls, preferably mounted on a remote control. The remote

control could be discarded once training was completed.

Our experimentation to date has indicated that users can reliably operate up to three controls that, between them, govern the broad shape of the gain-frequency response.⁸ Indeed, there are several viable ways that the controls can be linked to the shape of the gain-frequency response, although our experimental subjects preferred some arrangements to others.

By analogy with the simple volume control trainable aid, if the user can adjust the gain-frequency response at any one input level and if this is repeated in multiple environments, each with its own overall input level and spectral shape, then the learning algorithm is able to deduce, within each frequency region, the desired absolute gain at each input level, and thus the individually preferred compression ratios and compression thresholds.

In principle, there is no amplification parameter that cannot be trained by the user, provided he or she can hear something change in the sound quality as the underlying parameter is varied by moving

the control throughout its range. Quantities such as depth and speed of automatic noise suppression, selection and speed of adaptive directionality, compression speed, spectral enhancement, and frequency transposition can all be subject to training. In general, the more parameters we attempt to train, the more time the user will need to spend training the device.

Our experimental data with a trainable aid that controlled the non-linear gain-frequency response suggest that, not surprisingly, the trained hearing aid became more strongly preferred to non-trainable (but otherwise identical) hearing aids the longer the user spent training it.⁷ A few weeks of training is certainly long enough to accumulate the necessary number of aid adjustments to achieve a significant preference for the trained aid. This preference was marked and significant, even under double-blinded conditions.⁷

Clients do typically appear to be willing to spend a few weeks training their hearing aids, though some are willing to spend only a few days and others are willing to spend months.⁹ The effort may not

be any greater than the client would otherwise have spent adjusting a conventional volume control. And, unlike a conventional hearing aid, the need to adjust should decrease as the hearing aid learns.

Because trainability can be used in so many ways, there is unlikely to be a single “trainable hearing aid.” Rather, many companies will introduce trainability, and different devices will vary in their method of user control, the aspects of the acoustic environment that are monitored, and the amplification parameters that are controlled.

Having explained what we mean by a trainable hearing aid, let’s distinguish it from a data-logging aid. A trainable hearing aid may well require the hearing aid to log data—user inputs, at least, and acoustic environment characteristics for a more sophisticated device. Furthermore, the data logged by a trainable hearing aid could subsequently be perused by the clinician (in person or over the Internet), which is another similarity with data-logging aids. However, if the hearing aid only logs data, for later perusal and action by the clinician or the fitting software, and does not automatically change amplification on the basis of input from the user that it logs, it is not user-trainable.

IMPACT IN THE CLINIC

Trainability is likely to have a profound effect on how hearing aids are prescribed, measured, and adjusted in the clinic. It may also affect how a succession of appointments is structured, and significantly decrease the average number and/or length of appointments.

Most obviously, if the purpose of trainability is to enable the client to customize the broad shape of the gain-frequency response to his or her particular listening environments, then there is no point in spending valuable clinical time achieving a *close* match to a prescribed response. While a good prescription should still be used to ensure that the sound is safe and reasonable from the first time the hearing aid is worn, it does not seem worthwhile spending the time to do real-ear measurements and subsequent adjustments, provided that the manufacturers’ software is reasonably accurate at approximating the prescription from the outset.

Further time saving should occur

to complain about the sound quality. Suppose, for example, that the client finds that the hearing aid sounds unpleasantly loud and rumbly in traffic noise. Instead of scheduling the client for an adjustment appointment (or worse, advising the client that his or her brain will get used to it and accommodate!), the clinic will simply instruct the client to head for the nearest traffic and repeatedly adjust the control to whatever sounds best in that situation.

Assuming the trainable aid enables the client to control the gain-frequency response, it will progressively turn down the low-frequency gain for high-level sounds, if that modification will solve the problem. If the trainable aid has sophisticated environment-monitoring capabilities and if it is only in traffic noise that the client benefits from this adjustment, then the automatic adjustment will take place only in traffic noise and not affect other high-intensity, low-frequency sounds, such as the wearer’s own voice.

Pretty quickly, the client will learn that the permanent solution to unacceptable sound in any situation is to use, and for a limited time, adjust the hearing aid in that situation. In the process, some clients will likely begin to feel more “ownership” over the adjustment of the hearing aid than if they regard it as something beyond their control. Other clients will no doubt prefer to leave everything to “the expert.”

A survey of the clinical (adult, hearing aid wearing) population has demonstrated an overwhelmingly positive response to the concept of trainability.⁹ Our hope is that trainable hearing aids will lead to more effectively customized hearing aids, as well as to a greater sense of ownership of the fitting, and that this combination of improvements will reduce the incidence of unused and returned hearing aids. However, we will have no data on this until trainable hearing aids are available in a commercial and cosmetically acceptable form.

OTHER POTENTIAL BENEFITS

Trainable hearing aids could also be set to detect long-term changes in the gain selected by the user. As the most likely explanations for this would be deterioration of the user’s hearing or the hearing aid’s becoming faulty (e.g., progressively blocked with wax), it would even be pos-

sible to schedule a return visit to the hearing clinic. What clients think about their hearing aid talking to them (preferably at some quiet moment) has not yet been tested by our research team!

Clinicians have long suspected that hearing aid wearers’ amplification requirements change once they acquire experience with the hearing aids. It is commonly believed that new users prefer less overall gain and less high-frequency emphasis than experienced users.

Several manufacturers include acclimatization or adaptation managers in their software to modify the standard prescriptions. While research support for these views is minimal,^{10,11} the trainable aid provides an easy way to allow for these effects if the concern is real and causes no disadvantage if it is not. Users can start and stop training the trainable aid any time they wish. A reasonable clinical strategy would therefore be to instruct clients to train the aid intensively during the first few weeks, then forget about training and enjoy the benefits of the trained aid for the next few months. At some later stage, they could resume training for an additional week or two, and if the user’s preferences have changed, the trainable aid will accommodate to such changes without further involvement by the clinician.

Clinicians may wonder if they will still be needed at all if the trainable hearing aid becomes popular. The answer is an easy yes, though how they spend their time may well change. If they spend less time on hearing aid measurement, adjustment, and re-adjustment, then they will have more time for:

- ❖ Understanding clients’ needs, beliefs, and motivation or lack of motivation prior to a fitting. Extensive research indicates that a client’s motivation to use hearing aids is an important determinant of success with amplification.¹²⁻¹⁶ Additional counseling aimed at uncovering clients’ beliefs and providing additional perspectives for them to consider prior to fitting may be effective in increasing hearing aid use and satisfaction.
- ❖ Giving information after the hearing aid fitting about the use of assistive listening devices, communication or environmental strategies in difficult listening situations, and telephone usage.

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.