Amazing sound performance comes out of the Blu

Author
Leonard Cornelisse
Hearing Scientist & Manager, Hearing System Engineering
Favourite sound: Footsteps in snow

How do you make a great product like Discover Next even better?

Simple, you add the functionality that is most asked for – provide the client with the ability to modify the performance of the instrument in-the-moment. A secondary but equally important requirement is to make the client's preferred settings (after adjustment) available to hearing care professionals (HCPs) in fitting software, so they can guide and assist clients as they work to optimise the fitting. Finally and perhaps most important, you make it easier for clinicians to make meaningful adjustments to the sound performance in an automatic program which operates in a primarily mixed mode a large portion of the time. All of this comes together in Unitron's latest generation platform, Blu.

Although this required a rethink of the system architecture, it did not mean throwing away everything that has been built up over the years at Unitron. We have learned a lot from the development, usage and evolution of features like: Comfort/Clarity Balance, SmartFocus, Learn Now, Sound Conductor, SoundNav and its environmental classifier and SpeechPro.

Unitron's hearing instruments, much like instruments from other manufacturers, include a multitude of adaptive signal processing features (see Figure 1), such as:

- multi-channel, adaptive wide dynamic range compression (WDRC)
- impulse noise canceller (INC)
- frequency compression (Freq. Comp.)
- multiple microphone beamforming for directionality (Beamformer)
- spatial noise canceller (Spatial NC)
- speech enhancement (SE)
- noise reduction (NC)
- feedback phase canceller (Feedback Canceller)
- wind noise canceller (WNC)

Typical signal processing features in a hearing instrument are designed to adjust the amplified signal to achieve a variety of goals, such as:

- compensate for loss of audibility and loudness comfort,
- restoring acoustic cues that are disrupted by the physical presence of the hearing instrument, and
- improving the SNR by enhancing the speech and reducing noise through a variety of means.

Figure 1

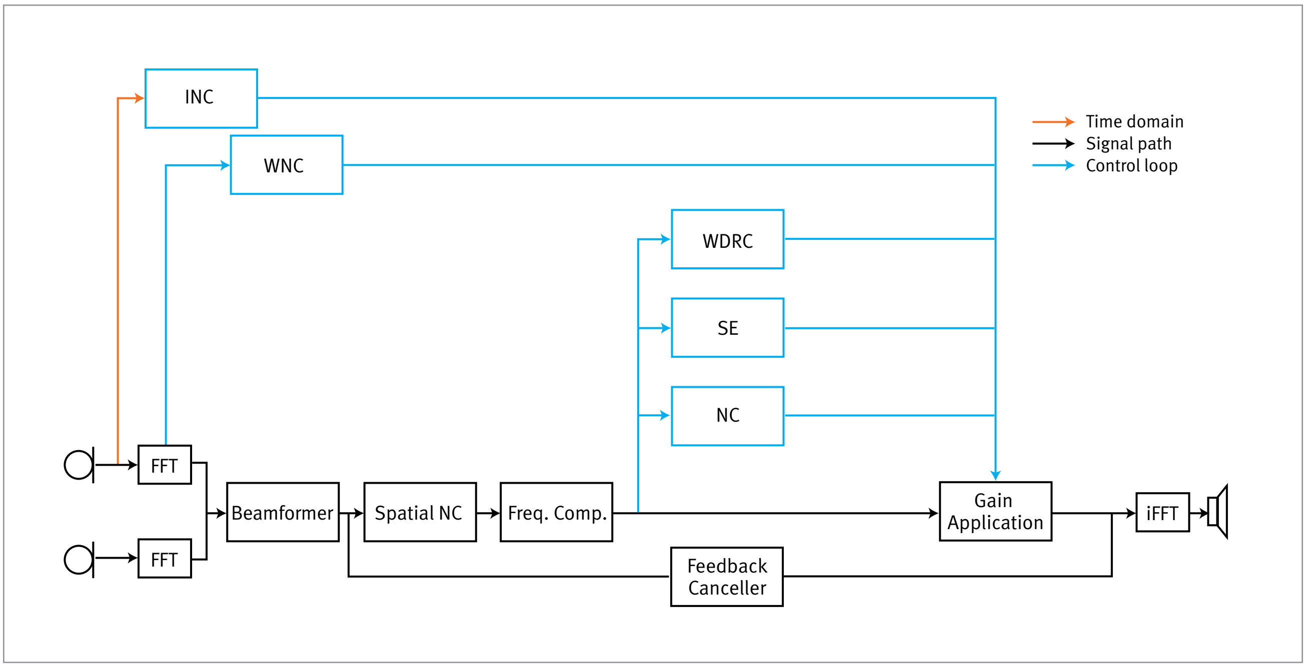

Figure 1. Simple block diagram of adaptive signal processing components in hearing instrument.

Why an automatic program?

Each of the adaptive features listed above can have a varying effect on the amplified signal depending upon the listening situation. In a hearing instrument system designed to operate all day every day, the performance of these adaptive signal processing features needs to be adjusted for the listening situation. Hence the need for an automatic program, which can select the correct combination of features and modify (steer) them, based on characterisation of the environment, using classification (such as SoundNav) and sophisticated steering (such as Sound Conductor) to select appropriate settings for each feature. The resulting steering, based on default settings, represents average preferences of listeners as determined during clinical trials.

What about personalisation?

However, we know that each listener and each listening situation is unique; and occasionally, the default average performance is not what the listener expects or prefers.

At Unitron, we believe that the highest satisfaction with a personal amplification system is achieved when the delivered sound performance matches the expected or desired sound performance for the hearing instrument user in the moment (see Figure 2 and Cornelisse, 2017). An optimal fitting finds the best match between delivered performance and expected performance with minimal intervention from the client.

Figure 2

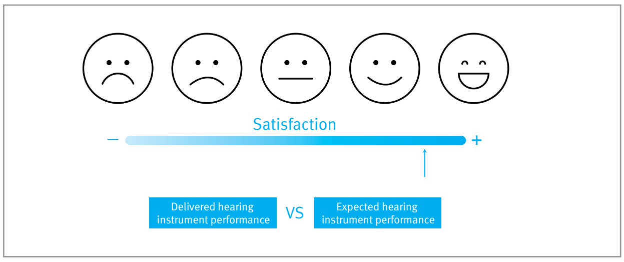

Figure 2. Satisfaction is related to how well the delivered sound performance of the hearing instrument matches the expected performance.

If wearers are not fully satisfied with the sound performance, then they should have the ability to make adjustments. Traditionally, this has been an arduous task that involved going back to the HCP, who would need to interpret the client's description of the problem (situation) and what was wrong (desired performance) and then make adjustments to one or more adaptive features in one or more environments in the automatic program. This process can be imprecise and may require several visits. It also runs the risk of making performance in other situations worse. The alternative for the hearing instrument wearer is to not make adjustments and get by with less than optimal performance in some scenarios.

Is there a better way?

When clients talk about making adjustments to sound performance there are typically three adjustment dimensions that they describe: focus, clarity of speech and minimisation noise. It should be noted that clients do not really know specifically which features to adjust, but rather, the user is describing sound performance effects that they wish to achieve. Unitron could have provided the client with controls that modify the performance of each individual adaptive feature directly, however that was not the direction chosen. Having control isn't about controlling everything, its about controlling what really matters or makes a difference – adjustments that align with these relevant perceptual dimensions are provided.

How to build the system?

In order to proceed, we needed to understand what we have today.

We recently collected over 55 hours of hearing instrument classification data (at a 1 second interval), in addition to other metrics including GPS location and EMA responses. A more complete description of the pilot investigation can be found in (Glista et all, 2020).

Of interest, was the performance of the Unitron environmental classifier. The classifier takes prototypical sound types, such as speech in noise or music and maps them using the sound type, plus overall level and SNR to create a mixed six listening environment structure (plus music as a seventh exclusive environment). The logged hearing instrument data lets us compare the overall signal level and SNR by environment at each point in time (~200,000 data points). The results are shown in Figure 3 and Table 1, which shows the distribution of SPL/SNR for each environment when the proportion was greater than 75%. As expected, the distribution of SPL/SNR varies by each sound environment and is clustered around the expected average SPL/SNR for each. That is, the cluster of data for the conversation in a small group is at favourable SNRs and moderate overall signal levels. Conversely, the no conversation noise data is clustered around poor SNRs and higher overall signal levels. It should be noted that an individual environment's proportion value only exceeded 75% roughly 46% of time. This means that for the other 54% of the time, the proportions were more mixed in nature, with no dominate single environment.

Reduced complexity in characterisation of the listening scenario

One challenge with classification in this way is that there are six mixed environments (aka dimensions), which adds an element of uncertainty when making adjustments. However, these can be collapsed onto a mixed space having two dimensions: communication and complexity. The communication dimension represents the probability that a conversation is taking place. The complexity dimension is a combination of elements representing the listening situation. These elements can include: a) the number of and variety in types of sound objects, b) number and location of competing sound sources, c) steadiness or variability

Figure 3

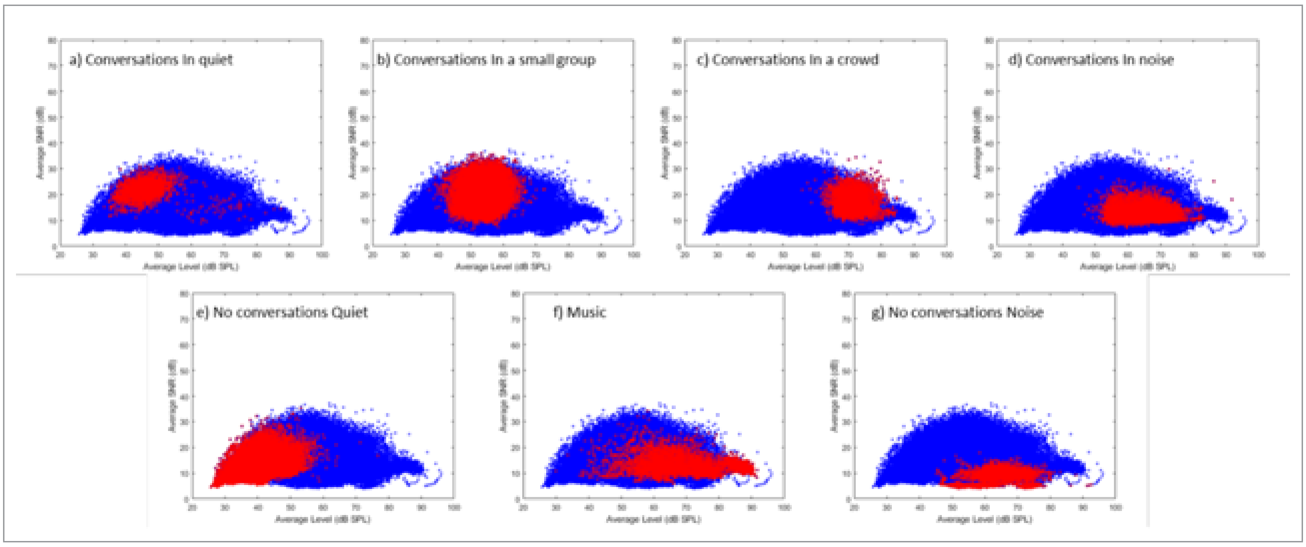

Figure 3. Scatterplot of overall signal level by estimated signal SNR. Each plot shows the distribution for all data (blue) and the distribution for the indicated classified environment (red). The data points for the selected environment are restricted to times when the class proportion was greater than 75%.

Table 1

	%	SPL	SNR
All data	100	55.8	14.4
Conversations			
In quiet	2.8	46.1	21.6
In a small group	10.1	54.3	19.7
In a crowd	3.2	71.1	16.9
In noise	5.1	64.1	12.7
No conversations			
Quiet	19.7	39.9	12.6
Noise	3.3	63.8	9.0
Music			
Music	2.2	69.8	13.2

Table 1. The percentage of time that each individual environment proportion exceed 75% and the corresponding average SPL and SNR. Note that the total time that the class proportions exceeded 75% was 46.4% of the time. The remaining 53.6% was a mixed scenario with no class proportion exceeding 75%.

in background noise (do sound objects come and go or are the sound objects more persistent) and d) listening challenge/difficulty, (eg. speech to noise ratio).

The classifier environments can be mapped onto this two dimensional, mixed space of communication/complexity (see Figure 4). For example, conversation in quiet and conversation in a small group map to low complexity/high conversation (top left), whereas the noise (without speech) environment maps to high complexity/low communication (bottom right). Interestingly, there is a high correlation between the location of the classifier environments in the communication/complexity space and the overall level/SNR space. However, these two representations are not exactly the same, since the communication/complexity has the added benefit of being built on top of the prototypical sound types. Since the mapping to the communication/complexity space, at each point in time, includes the prototypical sound types there is a higher degree of confidence in the classification of the situation. For example, there can be a higher degree of certainty that a conversation is occurring when the sound type is speech in noise, than if just SPL and SNR was used to classify the signal. The addition of the sound types allows the mapping to go from a space that covers overall level/SNR to a space with the perceptual dimensions of communication and complexity.

Figure 4

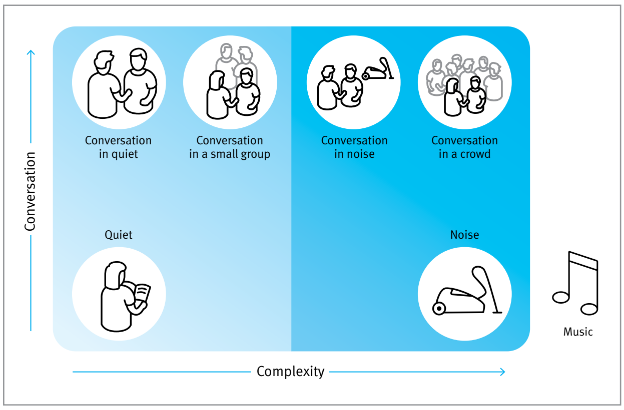

Figure 4. SoundNav classifier environments shown on two dimensional communication/complexity space.

An additional advantage of the communication/complexity space is that it can be divided into 'listening zones' and the appropriate signal processing in each zone becomes readily apparent. The communication/complexity space can be subdivided into four quadrants (see Figure 5) :

1. easy conversations
2. easy listening
3. challenging conversations
4. challenging listening

Figure 5

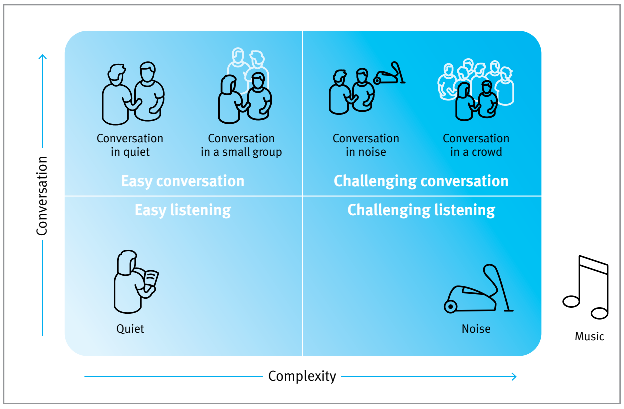

Figure 5. Communication/complexity space subdivided into listening zone quadrants.

Variable application of signal processing within the communication/complexity space.

The listening needs for each quadrant can be described based on the communication/complexity dimensions. For example, the left hand side of the space has low complexity and it is likely that the listener will require little additional signal processing, beyond what is required for hearing loss compensation. Conversely if the listener is in the challenging conversation quadrant – a listening situation (such as conversation in a crowd) with high probability of communication and high complexity, then it is likely that additional signal processing to facilitate communication in competing signals (eg. noise) will be required. Finally, if the listener is in the low communication/high complexity quadrant (aka background noise), then their preference for signal processing will likely be towards listening comfort in noise.

Based on the listening zones, it is possible to predict what type of signal processing will be preferred in each quadrant (see Figure 6). For example;

- More directionality (aka wide fixed directional) should be applied in the challenging side as compared to the easy side (spatial awareness). Or in the case of SpeechPro, target dependent directionality should be applied on the challenging side.

- Speech enhancement should only be applied when the value of the communication dimension is high and more SE should be applied for more challenging communication scenarios.
- Similarly increasing noise reduction should be applied for more challenging communication scenarios. In this case, it is advantageous to apply slightly more NC for non-communication scenarios than for communication scenarios.

These descriptions of signal processing strength match the performance that was delivered with the combination of SoundNav and Sound Conductor at default settings. What is not shown in Figure 6, is an additional frequency response offset that could also contribute to the goal of enhancing speech or reducing noise or other more advanced signal processing techniques.

Fitting and client controls in Blu

From a technical implementation perspective, it is important that client adjustments applied in the field and traditional fitting, in an office or clinical setting, are treated the same way. The goal of these fitting procedures is to modify the sound performance of the instrument for a given situation. Adjusting the sound performance of a hearing instrument is relatively simple when the device has a traditional single listening environment/manual program. However, it becomes substantially more complex when the hearing instrument has a mixed environment automatic program. Figure 7 shows a conceptual map of the elements to consider. The hearing instrument automatic program will be comprised of: a) a classification system

Figure 6

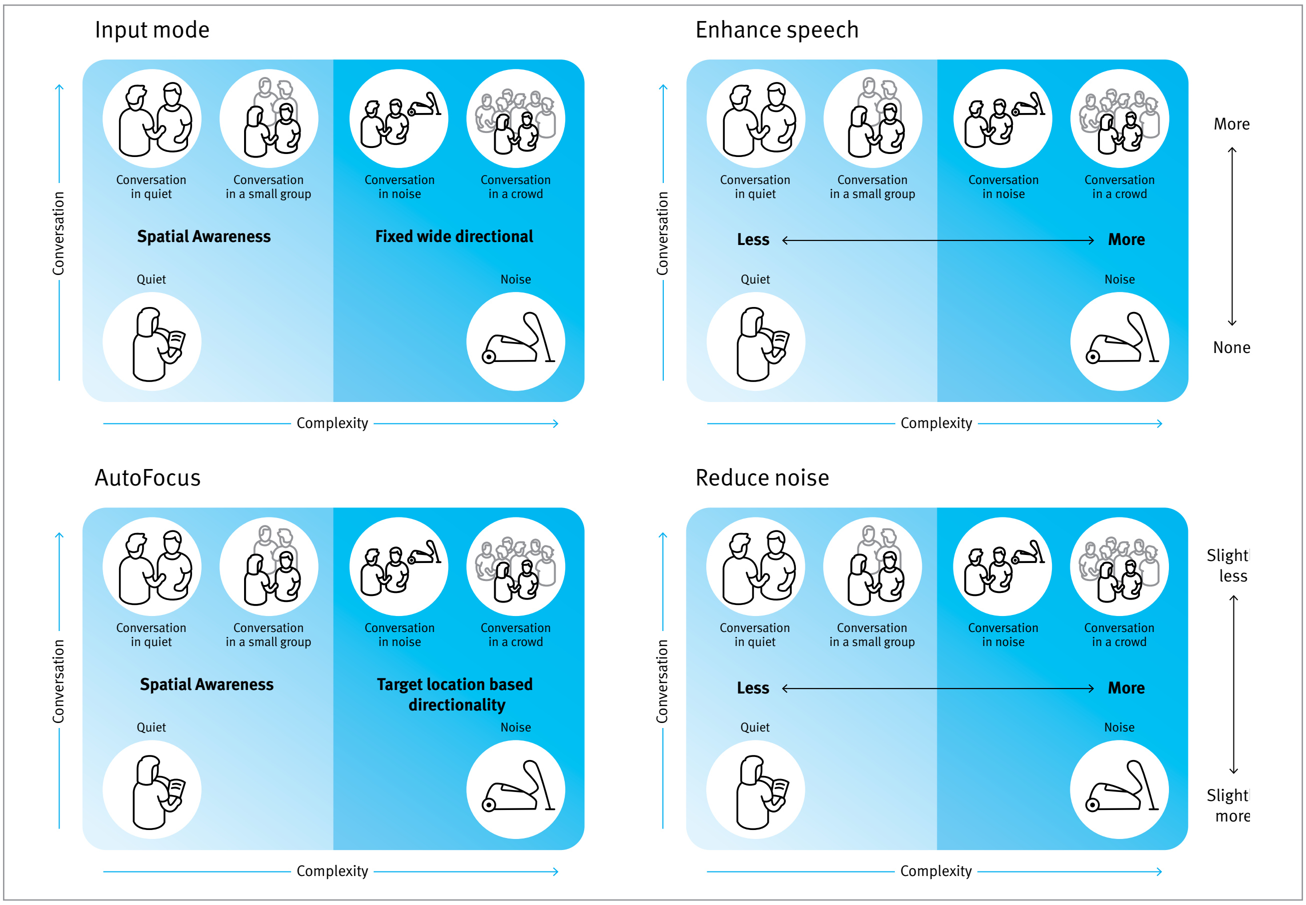

Figure 6. Conceptual application of signal processing to listening zones.

to characterise the listening scenario, b) a steering mechanism to adjust the adaptive features and c) the individual adaptive signal processing features. Our new automatic system, found in products on the Blu platform, Integra OS, brings together these aspects which contribute to the achieved sound performance. In addition, the HCP or client will want to make adjustments to the performance; and the listening scenario and the listener's intention will both influence where and when the adjustments are to be applied in the hearing instrument. These last two questions (where and when) are challenging to resolve.

1. Should an adjustment be applied to a specific situation (local), to all similar situations (generic local) or to all situations (global)?
2. For how long should the adjustment be applied? Is this a momentary change that will be removed after a period of time (volatile) or should the change be applied forever (persistent). A third mechanism could be to remember when an adjustment has been applied to the same situation numerous times and to then learn to apply the adjustment automatically.

Figure 7

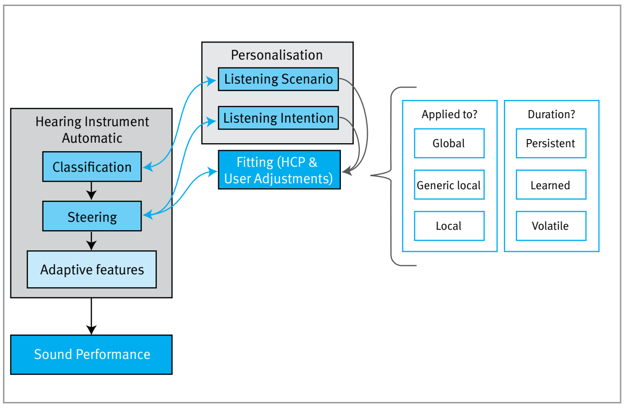

Figure 7. Mental model for 'fitting' automatic program.

At Unitron we have decided to take a two pronged approach to offer user adjustability. We expect performance in the automatic program to be well matched to the client's expectations. When the wearer wishes to modify the sound performance there are two options:

1. Stay in the automatic program and apply global volatile macro modifiers. These adjustments are applied to the automatic program, are applied in the moment and are temporary.
2. Switch to an appropriate manual program (or optional app program) and make additional 'generic local' persistent adjustments. Adjustments in the manual program are permanent, that is the next time the client selects the same manual program the previously modified settings will be applied. The assumption is that the client will select the same manual program for similar situations in which the same sound performance is desired. In this sense the manual programs are considered 'generic local', because the user can select the same program for multiple similar situations.

Since, in everyday use, the hearing instrument is in the automatic program, the client will most likely first make an adjustment to the automatic program. In this case, we offer a simple button to provide 'more' , either:

- along the dimension of Comfort – to reduce background noise.

The boost buttons in the Remote Plus app for the automatic program are macro controls that adjust a variety of adaptive signal processing features and allow the user to get a quick adjustment without a need to think about what individual settings to adjust. These controls are temporary and are meant to be applied in the moment, without leaving the automatic program. The original settings are restored when the client reboots the instruments.

If the macro controls in the automatic program do not provide satisfactory sound performance, then the client can select an optional program (or manual program) via the app that is suited for the situation. Unitron provides a range of pre-set programs to cover common situations that may pose a challenge for the listener. Each program offers a preset configuration suitable for the listening scenario and these programs can be personalised further. The user can adjust the program based on his or her individual preferences for the three dimensions of focus, enhancing speech and reducing noise, using sliders in the app. In this case, the changes are persistent and considered to be generic local, that is suitable for similar listening situations (eg. restaurant).

Summary

One of the primary goals in developing Blu was to facilitate the personalisation process for both the hearing care professional and the client. Our aims were to:

1. make it easier for the HCP to make meaningful adjustments to the sound performance in an automatic program which operates in a primarily mixed mode a large portion of the time.
2. make the client's preferred settings (after adjustment) available to the HCP in fitting software, so that he or she can guide and assist clients as they work to optimise the fitting.
3. provide the client with the ability to modify the performance of the instrument in-the-moment using an app.

Three key components that contribute to making the Blu platform a complete solution offering a high degree of personalisation, are:

What: Recognition that adjustment capability should offer changes in meaningful dimensions; focus, enhance speech and reduce noise.

How: Reduce the classification from a six destination mixed environment model to a two dimensional communication/complexity space thereby simplifying the conceptual model for application of adjustments.

Where and when: Utilise a simple mental model of applicability of user adjustments. Temporary changes are applied to the automatic program. In addition, the client has a variety of optional app programs which maintain changes to allow the client to further fine tune performance in specific situations when they are not fully satisfied with the sound performance of the automatic program.

Every day we navigate through a world of impromptu paths and unplanned detours. With the Moxi Blu family of devices and the latest generation of automatic sound optimisation, plus enhanced personalisation capabilities now available in the Remote Plus app, your clients will be prepared for wherever life takes them.

References

1. Danielle Glista, PhD, Robin O'Hagan, CDA, Leonard Cornelisse, MSc, Tayyab Shah, PhD, Donald Hayes, PhD, Sean Doherty, PhD, Jason Gilliland, PhD, and Susan Scollie, PhD, 2020, Combining passive and interactive techniques in tracking auditory ecology in hearing aid use, May 2020, Hearing Review online, https://www.hearingreview.com/inside-hearing/research/techniques-in-tracking-auditory-ecology-in-hearing-aid-use

2. Leonard Cornelisse, 2017, A conceptual framework to align sound performance with the listener's needs and preferences to achieve the highest level of satisfaction with amplification, Unitron Technical Paper (2017-09 027-6249-02)

027-6627-02/V1.02/2021-04/dr