Advanced Test Methods for Improving ANC Headphone Performance

Active noise-cancelling headphones are designed to electronically remove the sound coming from your surroundings. They work by using internal microphones that listen to what's happening in the world around you, inverting the noise, and sending it into the loudspeaker. The idea is both the output and the input will cancel out, leaving you with near-silence—or to whatever you choose to listen. This article attempts to gain better data to draw stronger conclusions about the ANC performance.

By
Hans W. Gierlich
(HEAD acoustics)

Photo 1: Testing an ANC headphone with the artificial head measurement system HMS II.3 and labCORE—the modular multi-channel front end platform

Active noise cancellation (ANC) in headphones has been available for several years, and has seen significant improvements in the amount of peak attenuation achievable. Today, customers use ANC in many different scenarios:

- Users want to shut themselves off from their environment. They'd like to avoid unwanted background noises (e.g., when traveling by plane, bus, or sitting in a noisy room). That is why they turn the ANC to maximum. In this scenario, it is not possible to play back any audio.
- In the second scenario, users still desire maximum noise cancellation but want to play back audio files such as music, news, or podcasts at the same time—or simply have speech in a phone call.
- The third use case scenario is about hearables and the idea of acoustic augmented reality. Users will want to hear the environment while they listen to audio or speech. The ANC is activated, but not at the maximum level. Thus, a pleasant and seamless interaction with the environment, and with other people around and on the other end of a phone call is still possible without having to increase the vocal effort and not use too much listening effort to understand the other person.

This article focuses mainly on the first and second scenario. The question we investigated is: How can we gain better data to draw stronger conclusions about the ANC performance? While testing ANC systems or headphones today users often have stationary noise playback and look at overall attenuation.

In some cases, ANC systems and headphones are also evaluated based on their ability in the bass (active), mid (cross-over) and treble (passive) regions, as well as how much they "leak" sound (see **Photo 1**). Those fundamental tests are good, but we identified four avenues that would allow us to improve our data.

First, instead of playing white or pink noise out of one or two speakers in a simplified noise field, we suggested playing back realistic background noise sources. There are standards existing in the sector of telecommunication testing such as European Telecommunications Standards Institute (ETSI) TS 103 224 that address this. This standard uses an eight-microphone array and an eight-loudspeaker configuration for recording sounds and playing them back in a laboratory.

The microphones are used to capture the spatial characteristics of the original sound field, so, during the equalization process and the playback process users are able to accurately playback the sound in terms of magnitude vs. frequency and also in terms of phase vs. frequency. This configuration is perfectly suitable for stressing beamforming devices or multi-microphone position devices, as well as providing a good impression on how the product behaves in the real world.

Second, when looking at the residual noise left behind by an ANC system, we suggested applying hearing models to gain a better understanding of the noise impact. The human brain is not a Fast Fourier Transform (FFT) analyzer, so it makes sense to evaluate ANC performance from the perspective of a hearing model (e.g., Zwicker Loudness) that is meant to approximate human hearing.

Third, we proposed looking at the temporal aspects of the ANC performance for two main reasons—one, our brains more closely resemble a time waveform analyzer and time domain events can, therefore, impact our sense of hearing, and two, the ANC systems and headphones may respond in different ways when they are subjected to transients in the background noise environment found in the real world. A simple level vs. time analysis can reveal much in this area. In contrast to the frequently used pink noise—which is very predictive, deterministic, and free of transients—the noise scenarios found in ETSI TS 103 224 (e.g., pub noise) include lots of activity in the background and lots of transients.

Finally, when activating an ANC system, the last thing we wanted was to compromise the integrity of the reproduced audio signal. That is why we chose to use sophisticated hearing models (e.g., the Relative Approach) to evaluate what the audio quality would be like when the ANC was switched on and when the ANC was switched off under various background noise scenarios.

Photo 2: 3PASS setup in a laboratory

Test Setup with HEAD acoustics Equipment

When testing ANC headphones with HEAD acoustics equipment, the setup consists of an ACQUA system—the advanced communication quality analysis software that is connected to a labCORE front end. This modular multi-channel is equipped with coreBT—the Bluetooth module for labCORE. The front end is connected to a pair of Bluetooth headphones and sends the stimulating source signal from ACQUA to the headphones and into the ears. Then, the signal from the ears is sent back to the labCORE and further to the ACQUA analysis.

In the meantime, the eight loudspeakers are driven by the background noise simulation system 3PASS (see **Photo 2**). A predefined background

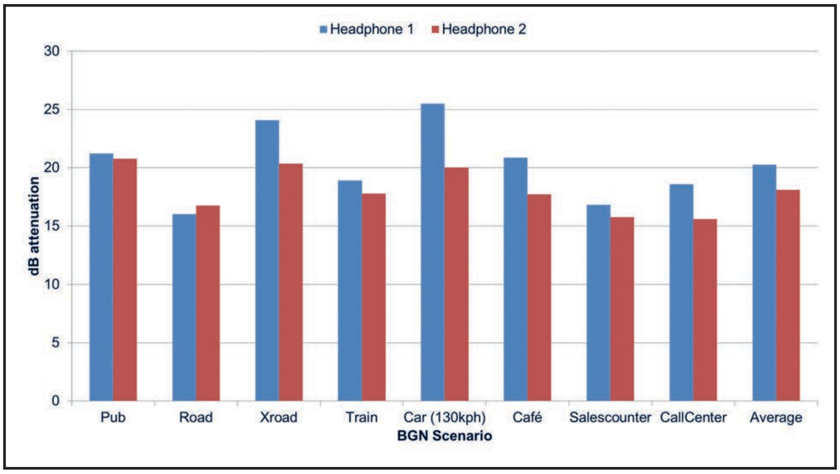

Figure 1: Overall active attenuation in decibels (dB)

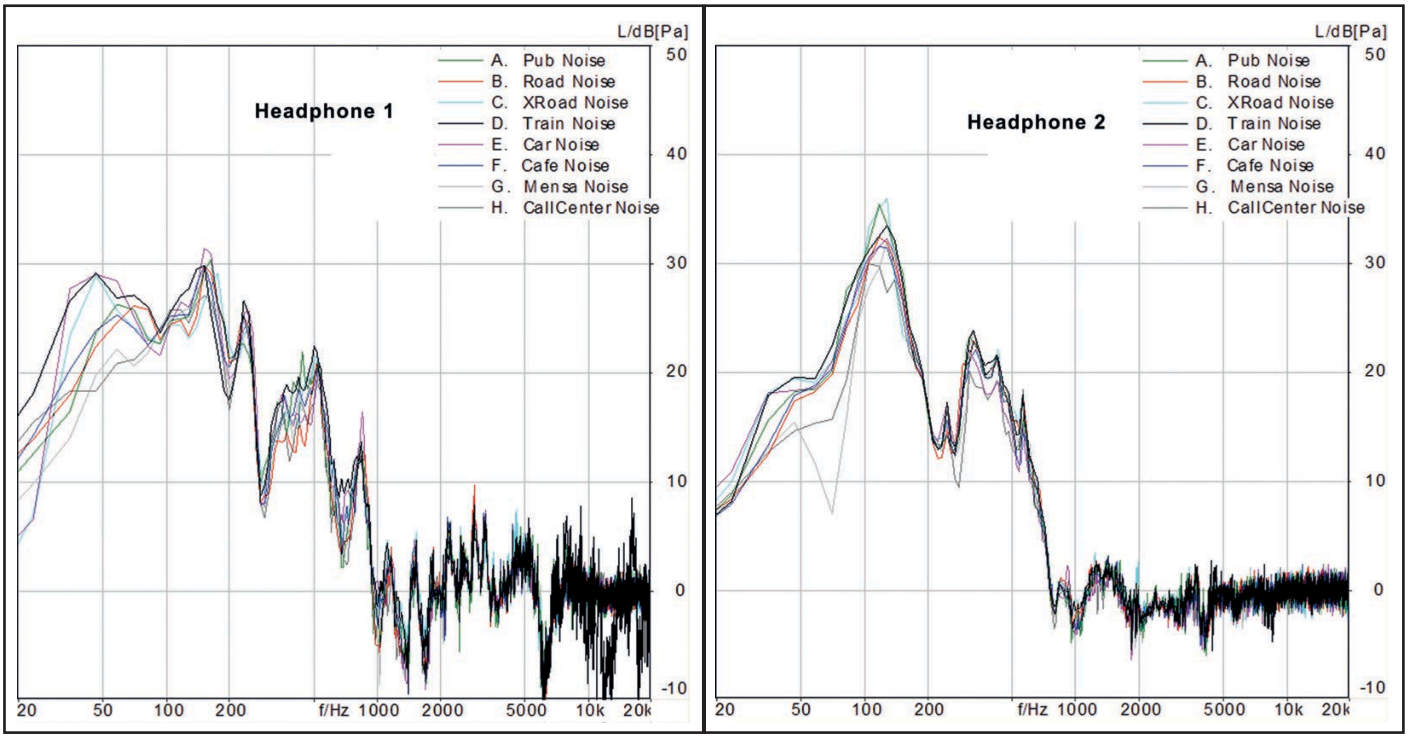

Figure 2: Overall active attenuation for Headphone 1 (a) and Headphone 2 (b)

noise was played back while measuring an ANC headphone. The ACQUA and the 3PASS computer are connected and synchronized such that any audio playback is going to occur simultaneously with the background noise playback. That made it possible to look at every time slice across headphone tests and compare them individually.

About the Author

Hans W. Gierlich started his professional career in 1983 at the Institute for Communication Engineering at RWTH, Aachen. In February 1988, he received a Ph.D. in electrical engineering. In 1989, Hans joined HEAD acoustics GmbH in Aachen as vice president. Since 1999, he has been head of the HEAD acoustics Telecom Division and in 2014, he was appointed to the board of directors. Hans is mainly involved in acoustics, speech signal processing and its perceptual effects, QOS and QOE topics, measurement technology, and speech transmission quality. He is active in various standardization bodies such as ITU-T, 3GPP, GCF, IEEE, TIA, CTIA, DKE, and VDA and chairman of the ETSI Technical Committee for "Speech and Multimedia Transmission Quality."

About HEAD acoustics—Telecom Division

HEAD acoustics was founded in 1986 and has been involved in noise and vibration, electroacoustic and voice quality testing since its inception. HEAD acoustics is based in Herzogenrath, Germany, with affiliates in China, France, Great Britain, Japan, South Korea, and the US, as well as a worldwide network of representatives. The Telecom Division of HEAD acoustics manufactures telecom test equipment and provides consulting services in the field of speech and audio quality. Moreover, HEAD acoustics closely co-operates with DECT Forum, ETSI, ITU-T, 3GPP, TIA, CTIA, GSMA, and other standardization bodies with regard to the development of quality standards for voice transmission and speech communication. In many partnership projects, HEAD acoustics has proven its competence and capabilities in conducting tests and optimizing communication products with respect to speech and audio quality under end-to-end as well as mouth-to-ear scenarios.

Comparison of Two ANC Headphones

Based on these advanced test methods and with the above described setup, we compared two commercially available ANC capable headphones. First, we measured the active attenuation achieved by these two headphones across different background noise scenarios and evaluated the performance according to traditional FFT-subtraction methods. We exposed the headphones to different noise scenarios (e.g., pub noise, road noise, call center noise, and more). The tests revealed that Headphone 1 performs generally better throughout the different scenarios (especially driving noise, which happens to resemble pink noise) than Headphone 2. The average active attenuation of Headphone 1 is about 20.2 dB, whereas Headphone 2 achieves an average score of 18.1 dB (see **Figure 1**).

Looking at the data from a frequency-based perspective (20 Hz to 20 kHz; ANC on minus ANC off), it once again showed that Headphone 1 slightly outperforms Headphone 2 (see **Figure 2**). While the peak attenuation at 150 Hz isn't quite as high as Headphone 2, Headphone 1 is clearly doing a better job attenuating the low frequencies and as well as offering a much higher upper frequency (1 kHz vs. ~750 Hz) than Headphone 2.

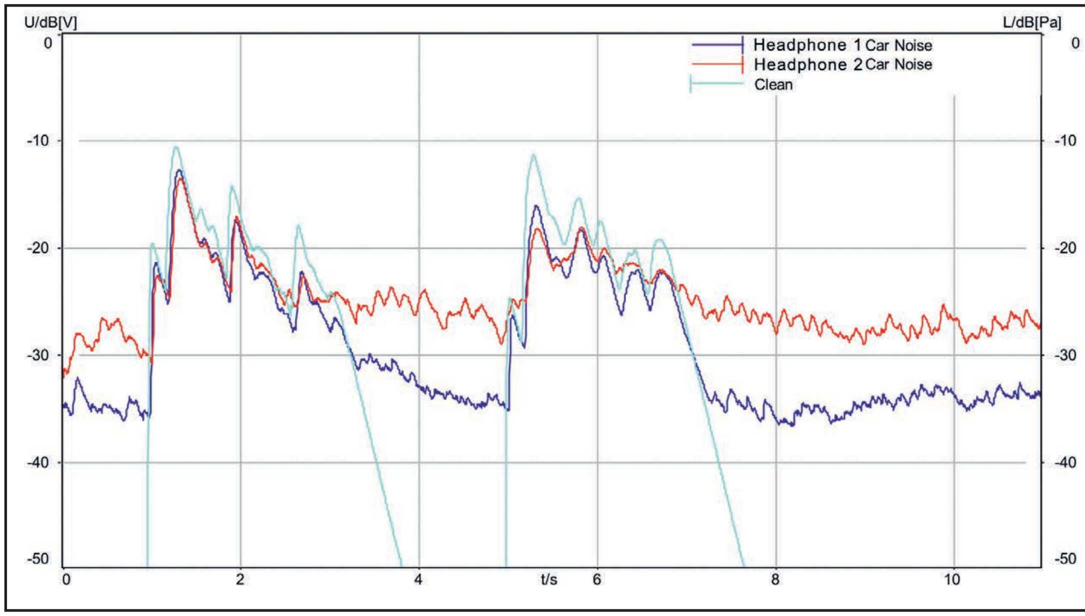

Figure 3: Level vs. Time analysis

The conclusion we drew is that the traditional objective FFT-based analysis methods would have crowned Headphone 1 the unequivocal king of ANC performance.

When we examined the temporal aspects in the level vs time analysis (see **Figure 3**), where we compared the headphones in noise to the level vs. time of the speech stimulus, it indicated that Headphone 1 provided better speech tracking. Furthermore, the noise level was generally lower and less prone to fluctuations.

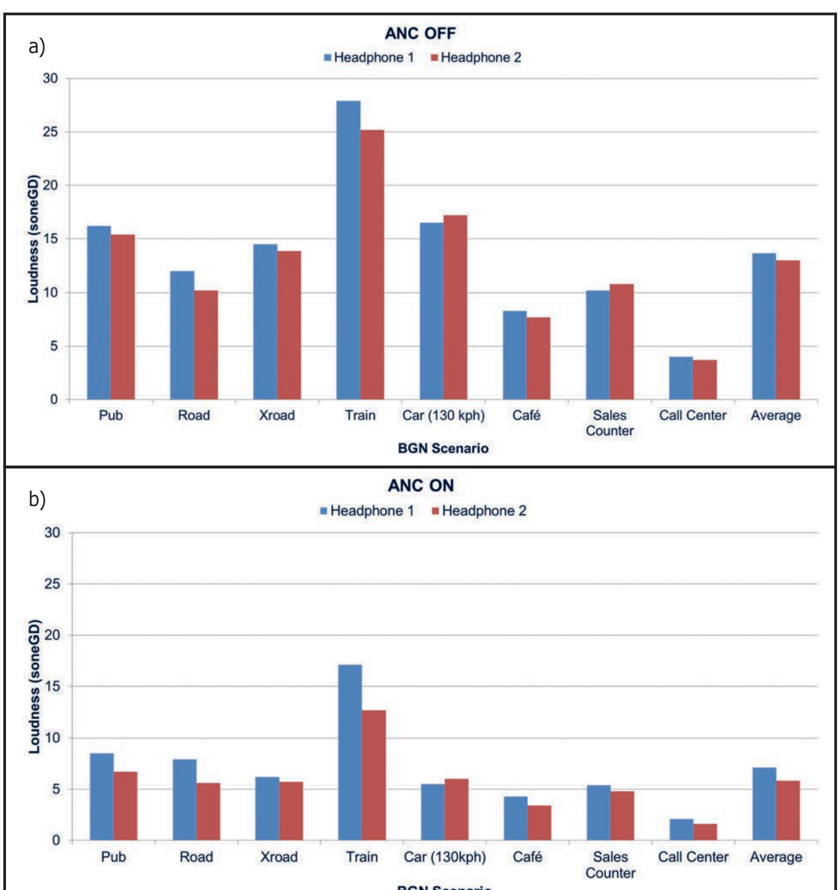

Figure 4: Zwicker Loudness in ANC Off (a) and ANC On mode (b)

However, when we looked at the data through the lenses of a frequency-based hearing model (Zwicker Loudness), which takes into account things such as frequency masking to get a better sense of how a human brain perceived a noise event, we got a vastly different picture. In ANC off mode (only passive attenuation), the overall noise through Headphone 1 is perceived as louder (see **Figure 4**). When switching ANC on, of course the loudness rating of both headphones 1 and 2 decreases, yet the gap between them widens. Headphone 2 reduced the overall noise perceived not significantly more compared to the noise perceived at Headphone 1. When we looked at Zwicker Loudness levels for the two headphones in individual noise scenarios, Headphone 2 outperformed Headphone 1, except in the driving noise scenario (most akin to pin noise).

Another advanced test method is the Speech Intelligibility Index (SII). Ironically, the SII metric does not measure speech, rather it measures a noise environment and then estimates how intrusive that noise environment is against idealized speech. It turned out again that across the different noise scenarios, Headphone 2 achieved better active and passive SII scores and that ANC improved SII more in Headphone 2 than in Headphone 1 (see **Figure 5**).

So while the traditional methods painted one picture, we were starting to see a different image of the two headphone ANC systems appear.

As an additional way to evaluate the ANC performance, we decided to use a new Mean Opinion Scores (MOS) metric from ETSI TS 103 281. It uses the sophisticated Relative Approach hearing model, which is both temporally and spectrally sensitive to approximate human hearing. It was originally designed for evaluating speech quality in the presence of noise for a communication device that is transmitting sound, so the application of the metric in the receiving direction is a little bit experimental. However, the metric actually gave us three scores in one. We got an S-MOS score (indicator of speech quality), an N-MOS score (indicator of noise suppression quality), and a G-MOS score (global score). We used the N-MOS portion of ETSI TS 103 281 metric to evaluate the noise suppression quality of both headphones—which revealed that both perform well, however, Headphone 2 once again performed a little bit better. When applying ANC, both headphones improved equally well and the overall difference between them was minor (N-MOS Headphone 1 = 3.2; N-MOS Headphone 2 = 3.3)—but Headphone 2 performed technically better.

Finally, we wanted to evaluate the speech quality performance in the presence of background noise

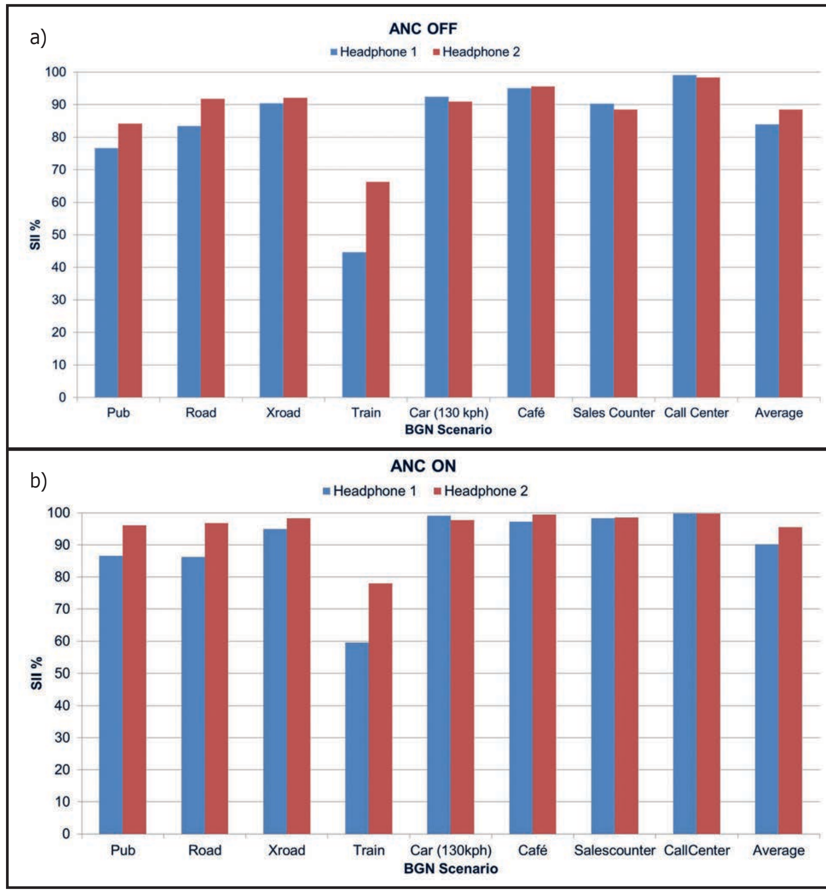

Figure 5: Speech Intelligibility Index (SII) with ANC Off (a) and ANC On (b)

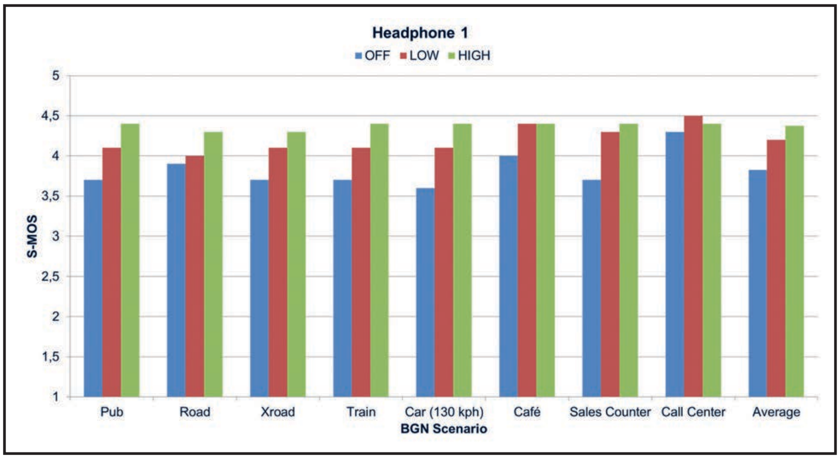

Figure 6: ANC effects on speech quality using Headphone 1

by using the S-MOS portion of ETSI TS 103 281. Headphone 1 can be used in three different ANC settings: ANC on, ANC low, and ANC high. In general, the measurements showed that we tend to get improved speech quality when increasing the level of active noise cancellation (see **Figure 6**).

Since the ETSI TS 103 281 metric uses speech files as its stimulus, we did have to make some assumptions about audio quality: If the speech quality is good, the audio quality is good, too. Looking at Headphone 1's performance across noise environments, this holds true except in the case of call center noise. This is relatively low-level noise (people talking quietly in the background, keyboard typing, etc.). When applying too much ANC, it lead to a slight drop in speech quality. Perhaps this was due to the fact that call center was the quietest predefined noise scenario and the ANC system in Headphone 1 was so aggressive that it sought something to cancel—and, unfortunately, affected the speech quality—or that we are affected by the self-noise of the ANC system.

Headphone 2 can be used in four different ANC settings: ANC off, ambient normal, ambient voice, and ANC on (see **Figure 7**). Ambient normal and ambient voice settings are meant to apply gain at the higher frequencies in order to cancel out the passive attenuation and essentially give the impression that you aren't wearing any headphone. In addition, the ambient voice setting has some slight attenuation at the lower frequencies in an attempt to improve speech quality. In some sense, those two settings are an attempt at creating a "hearables" type product as we defined in scenario three. As with Headphone 1, increasing the amount of ANC in Headphone 2 led to the same positive trend of improving speech quality.

Comparing the S-MOS performance of Headphone 1 and Headphone 2 showed that Headphone 1 seemed to perform better both with ANC on and ANC off (ANC off/ Average S-MOS: Headphone 1 = 3.8; Headphone 2 = 3.8 // ANC on/Average S-MOS: Headphone 1 = 4.3; Headphone 2 = 4.2).

Conclusion

Comparing the performance of both headphones, using the different test methods across different background noise scenarios, revealed a more balanced view of the two.

- Headphone 1 scored well using traditional FFT-based analysis techniques, level vs. time and in speech quality reproduction with and without ANC.
- Headphone 2, however, tended to score better when using hearing models and psychoacoustic metrics for evaluating the residual noise.

The primary message is that judging the ANC performance solely on FFT-based analysis can be misleading.

For those who want to explore further, we have presented four areas where testing can be expanded in order to gain better insight into ANC headphone performance. First, more realistic noise source with spatial characteristics should be applied for better use case validation. Second, use loudness models and psychoacoustic analyses to judge impact of residual noise. Third, look at temporal aspects and behavior regarding performance of ANC systems and regarding the human brain's performance. Finally, assess the influence on audio quality in different ANC settings.

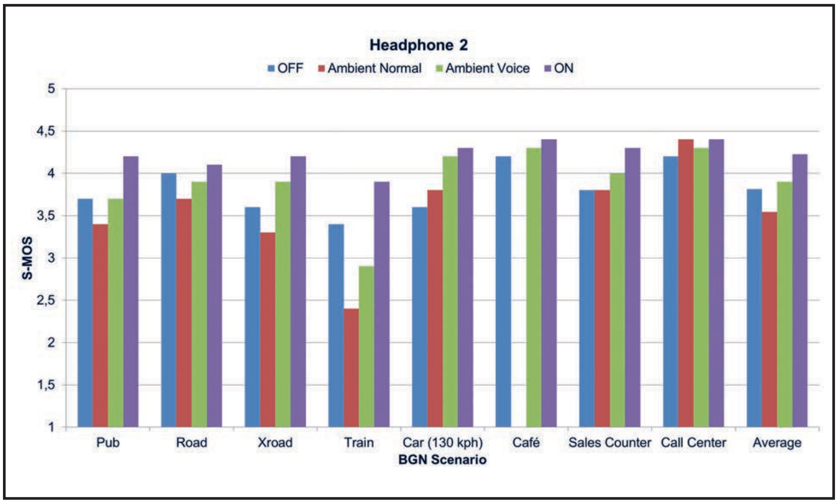

Figure 7: ANC effects on speech quality using Headphone 2