



## Rethinking Hearing Aids as Recommender Systems

**Pasta, Alessandro; Petersen, Michael Kai; Jensen, Kasper Juul; Larsen, Jakob Eg**

*Published in:*  
CEUR Workshop Proceedings

*Publication date:*  
2019

*Document Version*  
Publisher's PDF, also known as Version of record

[Link back to DTU Orbit](#)

*Citation (APA):*  
Pasta, A., Petersen, M. K., Jensen, K. J., & Larsen, J. E. (2019). Rethinking Hearing Aids as Recommender Systems. *CEUR Workshop Proceedings*, 2439, 11-17.

---

### General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

# Rethinking Hearing Aids as Recommender Systems

Alessandro Pasta  
Technical University of Denmark  
Kongens Lyngby, Denmark  
alpas@dtu.dk

Kasper Juul Jensen  
Oticon A/S  
Smørum, Denmark  
kjen@oticon.com

Michael Kai Petersen  
Eriksholm Research Centre  
Snekkersten, Denmark  
mkpe@eriksholm.com

Jakob Eg Larsen  
Technical University of Denmark  
Kongens Lyngby, Denmark  
jaeg@dtu.dk

## ABSTRACT

The introduction of internet-connected hearing aids constitutes a paradigm shift in hearing healthcare, as the device can now potentially be complemented with smartphone apps that model the surrounding environment in order to recommend the optimal settings in a given context and situation. However, rethinking hearing aids as context-aware recommender systems poses some challenges. In this paper, we address them by gathering the preferences of seven participants in real-world listening environments. Exploring an audiological design space, the participants sequentially optimize three audiological parameters which are subsequently combined into a personalized device configuration. We blindly compare this configuration against settings personalized in a standard clinical workflow based on questions and pre-recorded sound samples, and we find that six out of seven participants prefer the device settings learned in real-world listening environments.

## CCS CONCEPTS

• **Information systems** → **Personalization; Recommender systems**; • **Human-centered computing** → *Ambient intelligence; User centered design.*

## KEYWORDS

Personalization, recommender systems, hearing healthcare, hearing aids

### ACM Reference Format:

Alessandro Pasta, Michael Kai Petersen, Kasper Juul Jensen, and Jakob Eg Larsen. 2019. Rethinking Hearing Aids as Recommender Systems. In *Proceedings of the 4th International Workshop on Health Recommender Systems co-located with 13th ACM Conference on Recommender Systems (HealthRecSys'19)*, Copenhagen, Denmark, September 20, 2019, 7 pages.

## 1 INTRODUCTION

Despite decades of research and development, hearing aids still fail to restore normal auditory perception as they mainly address the lack of amplification due to loss of hair cells in the cochlea [16], rather than compensating for the resulting distortion of neural activity patterns in the brain [22]. However, the full potential of

hearing aids is rarely utilized as devices are frequently dispensed with a “one size fits all” medium setting, which does not reflect the varying needs of users in real-world listening scenarios. The recent introduction of internet-connected hearing aids represents a paradigm shift in hearing healthcare, as the device might now be complemented with smartphone apps that model the surrounding environment in order to recommend the optimal settings in a given context.

Whereas a traditional recommender system is built based on data records of the form  $\langle user, item, rating \rangle$  and may apply collaborative filtering to suggest, for instance, new items based on items previously purchased and their features, recommending the optimal hearing aid settings in a given context remains highly complex. Rethinking hearing aids as recommender systems, different device configurations could be interpreted as items to be recommended to the user based on previously expressed preferences as well as preferences expressed by similar users in similar contexts. In this framework, information about the sound environment and user intents in different soundscapes could be treated as contextual information to be incorporated in the recommendation, building a context-aware recommender system based on data records of the form  $\langle user, item, context, rating \rangle$  [1]. However, addressing some challenges related to the four aforementioned data types is essential to make it possible to build an effective context-aware recommender system in the near future. In this paper, we discuss the main challenges posed when rethinking hearing aids as recommender systems and we address them in an experiment conducted with seven hearing aid users.

### 1.1 Rating

In order to be able to precisely and accurately recommend optimal device settings in every situation, gathering relevant user preferences (expressed as ratings) is essential. However, learning user preferences poses some challenges. Firstly, the device settings reflect a highly complex audiological design space involving multiple interacting parameters, such as beamforming, noise reduction, compression and frequency shaping of gain. It is important to explore the different parameters, in order not to disregard some parameters that might have relevant implications for the user listening experience, and to identify which parameters in an audiological design space [10] define user preferences in a given context. Secondly, the preferred device settings depend on the human perception of the

*HealthRecSys'19, September 20, 2019, Copenhagen, Denmark*

© 2019 Copyright for the individual papers remains with the authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0). This volume is published and copyrighted by its editors.

listening experience and it is therefore difficult to represent the perceptual objective using an equation solely calculated by computers [21]. Having to rely on user feedback, it is important to limit the complexity of the interface, to make the interaction as effective as possible. Thirdly, capturing user preferences in multiple real-world situations not only guarantees that the situations are relevant and representative of what the user will experience in the future, but it also allows the user to test the settings with a precise and real intent in mind. However, this increases the complexity of the task, since the real-world environment is constantly changing and a user might explore the design space while performing other actions (e.g. conversing).

A traditional approach to find the best parameter combination (i.e. the best device configuration) is *parameter tweaking*, which consists in acting on a set of (either continuous or discrete) parameters to optimize them. Similarly to enhancing a photograph by manipulating sliders defining brightness, saturation and contrast [21], the hearing aid user could control her listening experience by tweaking the parameters that define the design space and find the optimal settings in different listening scenarios. However, this method can be tedious when the user is moving in a complex design space defined by parameters that interact among each other [13]. One frequently used method to simplify the task of gathering preferences is *pairwise comparison*, which consists in making users select between two contrasting examples. A limitation of this approach is efficiency, given that a single choice between two examples provides limited information and many iterations are required to obtain the preferred configuration. Based on pairwise comparisons, an active learning algorithm may apply Bayesian optimization [2] to automatically reduce the number of examples needed to capture the preferences [3], assuming that the samples selected for comparison capture all parameters across the domain. Alternatively, one might decompose the entire problem into a sequence of unique *one-dimensional slider manipulation* tasks. As exemplified by Koyama et al. [13], the color of photographs can be enhanced by proposing users a sequence of tasks. At every step, the method determines the one-dimensional slider that can most efficiently lead to the best parameter set in a multi-dimensional design space defined by brightness, contrast and saturation. Compared to pairwise comparison tasks, the single-slider method makes it possible to obtain richer information at every iteration and accelerates the convergence of the optimization.

Inspired by the latter approach we likewise formulate the learning of audiological preferences in a given listening scenario as an optimization problem:

$$z = \arg \max_{x \in \mathcal{X}} f(x)$$

where  $x$  defines parameters related to beamforming, attenuation, noise reduction, compression, and frequency shaping of gain in an audiological design space  $\mathcal{X}$  [10] and the global optimum of the function  $f: \mathcal{X} \rightarrow \mathfrak{R}$  returns values defining the preferred hearing aid settings in a given listening scenario.

However, while it remains sensible to assume that individual adjustments would converge when crowdsourcing (i.e. asking crowd workers to complete the tasks independently) the task of enhancing an image [13], it is less likely that hearing impaired users would

have similar preferences due to individual differences in their sensorineural processing [16, 22]. Therefore, at least in the first phase, we need to ask the same user many times about her preferences, until her optimal configuration is found. Furthermore, in order to optimize the device in different listening scenarios, we need to ask the same user to move in the same design space multiple times. Altering the one-dimensional slider at every step of the evaluation procedure might make the task difficult, since the user would not know the trajectory defined by the new slider. We believe that decoupling the parameters and allowing users to manipulate one parameter at a time, moving in a one-dimensional space that is clearly understood, would allow them to better predict the effects of their actions and hence more effectively assess their preferences.

## 1.2 Item

In order to enhance the hearing aid user experience, it is important to appropriately select the parameters that define the hearing aid configurations evaluated by users. Indeed, not only should the parameters have a relevant impact on the user listening experience, but the different levels of the parameters should also be discernible by untrained users. Three parameters have been demonstrated to be particularly important for the experience of hearing impaired users:

- (1) *Noise reduction and directionality*. Noise reduction reduces the effort associated with speech recognition, as indicated by pupil dilation measurements, an index of processing effort [23]. By allowing speedier word identification, noise reduction also facilitates cognitive processing and thereby frees up working memory capacity in the brain [18]. Moreover, fast-acting noise reduction proved to increase recognition performances and reduce peak pupil dilation compared to slow-acting noise reduction [23]. Given that the ability of users to understand speech in noisy environments may vary by up to 15 dB [4], it is essential to be able to individualize the threshold levels for the activation of noise reduction.
- (2) *Brightness*. While a lot of research has been focused on adapting the frequency-specific amplification which compensates for a hearing loss based on optimized rationales like VAC+ [5], rationales still reflect average preferences across a population rather than individual ones. Several studies indicate that some users may benefit from increasing high-frequency gain in order to enhance speech intelligibility [11, 12].
- (3) *Soft gain*. The perception of soft sounds varies largely among individuals. Hearing aid users with similar hearing losses can perceive sounds close to the hearing threshold as being soft or relatively loud. Thus, proposing a medium setting for amplification of soft sounds may seem right when averaging across a population, but would not be representative of the large differences in loudness perception found among individual users [17]. For this reason, modern hearing aids provide the opportunity to fine-tune the soft gain by acting on a compression threshold trimmer [14].

Taking a naive approach, treating each parameter independently, the preferences could subsequently be summed up in a general hearing aid setting, by simply applying the most frequently preferred values along each audiological parameter.

### 1.3 User

Hearing aids are often fitted based on a pure tone audiometry, a test used to identify the hearing threshold of users. However, as mentioned above, users perceive the sounds differently and might benefit from a fully personalized hearing aid configuration. For this reason, it is essential to fully understand what drives user preferences and which is the relative importance of users' characteristics and context. It is interesting to analyse whether users exhibit similar preferences when optimizing the hearing aids in several real-world environments and whether they result into similar configurations.

### 1.4 Context

Users often prefer to switch between highly contrasting settings depending on the context [11]. It has been shown that a context-aware hearing aid needs to combine different contextual parameters, such as location, motion, and soundscape information inferred by auditory measures (e.g. sound pressure level, noise floor, modulation envelope, modulation index, signal-to-noise ratio) [12]. However, these contextual parameters might fail to capture the audiological intent of the user, which depends not only on the characteristics of the sound environment but also on the situation the user is in. For this reason, in addition to retrieving the characteristics of the sound environment and the preferred device settings, it is also important to capture the contextual intents of users in the varying listening scenarios. Contextual information, in this exploratory phase, can be explicitly obtained by directly asking the user to define the situation she is in. However, in the future, to enable an automatic adaptation to the needs of users in real-world environments, relevant contextual information will need to be inferred using a predictive model that classifies the surrounding environment.

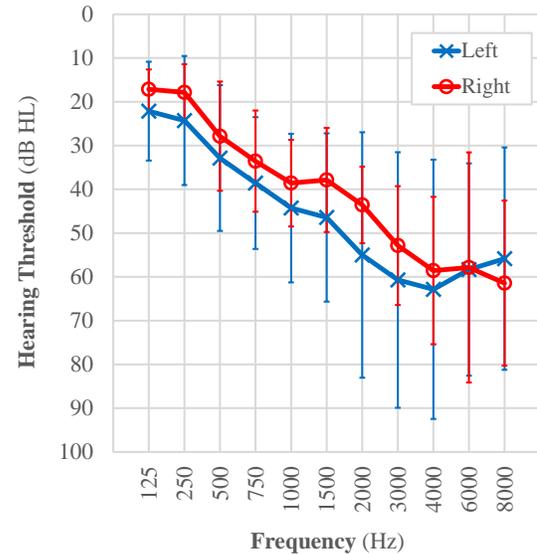
## 2 METHOD

### 2.1 Participants

Seven participants (6 men and 1 woman), from a screened population provided by Eriksholm Research Centre, participated in the study. Their average age was 58.3 years (std. 12 years). Five of them were working, while two were retired. They were suffering from a binaural hearing loss ranging from mild to moderately severe, as classified by the American Speech-Language-Hearing Association [6]. The average hearing threshold levels are shown in Figure 1. They were all experienced hearing aid users, ranging from 5 to 20 years of experience with hearing aids. All test subjects received information about the study and signed an informed consent before the beginning of the experiment.

### 2.2 Apparatus

The participants were fitted according to their individual hearing loss with a pair of Oticon Opn S 1 miniRITE [8]. All had iPhones with iOS 12 installed and additionally downloaded a custom smartphone app connected to the hearing aids via Bluetooth. The app enabled collecting data about the audiological preferences and the corresponding context.



**Figure 1: Average hearing threshold (i.e. the sound level below which a person's ear is unable to detect any sound [7]) levels for the 7 participants. The participants had a hearing loss ranging from mild to moderately severe. Error bars indicate  $\pm 1$  standard deviation of the hearing thresholds.**

### 2.3 Procedure

The experiment was divided into four weeks. As shown in Table 1, the first three weeks were devoted to optimizing the three audiological parameters, one at a time. Each of the first three weeks, the participants were fitted with four levels of the respective parameter, while the other two parameters were kept neutral at a default level. For instance, in week 1, each participant could select between four levels of noise reduction and directionality. The participants were instructed to compare, using a smartphone app, the four levels of the parameter in different situations during their daily life and to report their *preference*. To ensure that the participants would evaluate the different levels in relevant listening situations and when motivated to optimize their device, they were instructed to perform the task on a voluntary basis. Moreover, every time they reported their preference, the participants were asked to specify:

- The *environment* they were in (e.g. office, restaurant, public space outdoor). Different environments are characterised by different soundscapes and pose disparate challenges for hearing aid users.
- Their *motion* state (e.g. stationary, walking, driving). Motion tells more about the activity conducted by the person, but may also mark the transition to a different activity or environment [9].
- Their audiological *intent* (e.g. conversation, work meeting, watching TV, listening to music, ignoring speech). Complementing the contextual information by gathering the intent of the participants in the specific situation might provide a

**Table 1: Study timeline**

Week	Activity
W. 1	Optimization of <i>noise reduction and directionality</i>
W. 2	Optimization of <i>brightness</i> (amplification of high-frequency sounds)
W. 3	Optimization of <i>soft gain</i> (amplification of soft sounds)
W. 4	Final <i>test of preference</i>

deeper insight into how the different audiological parameters help them in coping with different sounds.

- The *usefulness of the parameter* in the specific situation (on a scale ranging from 1 to 5). This evaluation is important not only to understand the relative importance of each preference, but also to assess the perceived benefit of the parameter in diverse situations.

The fourth week each participant compared two different device configurations in a blind test:

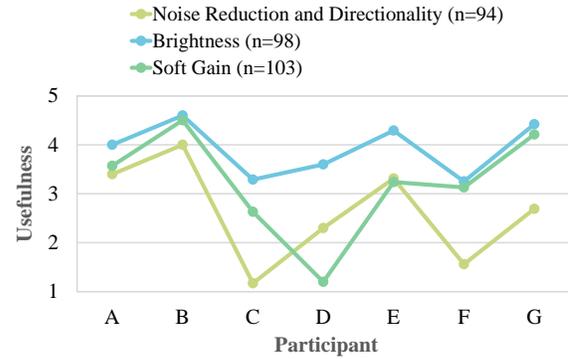
- An individually personalized configuration combining the most frequently selected preferences of the three audiological parameters gathered in *real-world listening environments* during the previous three weeks.
- A configuration personalized in a *standard clinical workflow* based on questions and on pairwise comparisons of pre-recorded sound samples capturing different listening scenarios including, for instance, speech with varying levels of background noise.

The participants were instructed to compare the two personalized configurations in different listening situations throughout the day and report their preference, while also labeling the context. At the end of the week, the participants were asked to select the configuration they preferred.

### 3 RESULTS

During the four weeks of test, the participants actively interacted with their devices, changing the hearing aid settings, overall, 4328 times (i.e. the level of the parameter during the first three weeks or the final configuration during the last week) and submitting 406 preferences. On average, the participants tried the different hearing aid settings 11 times before submitting a preference. Although one parameter affects the perception of the others, isolating them allows to analyse their perceived impact on the listening experience. As illustrated in Figure 2, the brightness parameter was on average rated higher in perceived usefulness. This result is consistent among the seven participants. Conversely, the noise reduction and directionality parameter resulted to have the lowest perceived usefulness for five participants out of seven. The soft gain parameter resulted to have an average perceived usefulness between those of the other two parameters.

Recording, together with each preference, the perceived usefulness of the parameter in the specific situation also allows to understand how much each parameter contributes to the overall setting of the hearing aid. Figures 3, 4, 5 display the preferences of test participants for different levels of noise reduction and directionality, brightness, and soft gain, respectively. Only the preferences

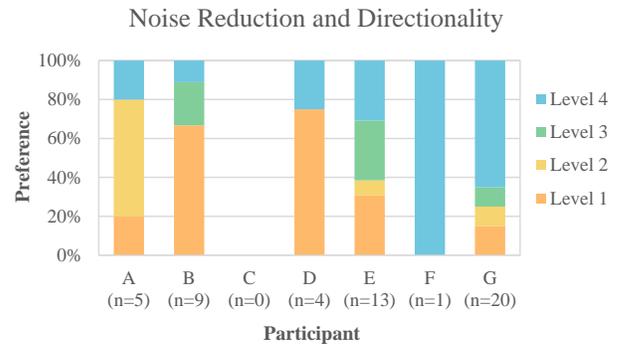


**Figure 2: Average perceived usefulness of three parameters (noise reduction and directionality, brightness, soft sounds). Brightness is perceived to be the most useful parameter. Noise reduction and directionality tends to be perceived as the least useful parameter.**

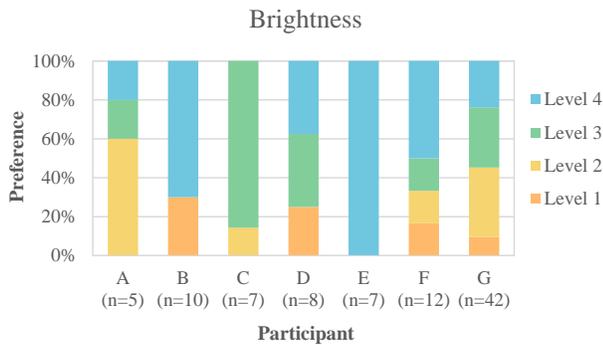
recorded in situations where the usefulness of the parameter is rated higher than two out of five are considered.

Firstly, the results indicate that the participants have widely different audiological preferences, rather than converging towards a shared optimal value. As the participants are ordered by age (A being the youngest), there seem, nevertheless, to be some common tendencies among younger or older participants across all parameters.

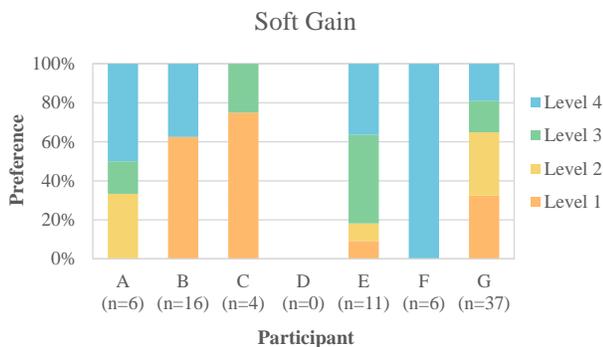
Secondly, most participants are not searching for a single optimum but select different values within each parameter. When adjusting the perceived brightness (Figure 4), six participants out of seven prefer, most of the time, the two highest levels along this parameter. Thirdly, the participants frequently prefer highly contrasting values within each parameter, depending on the context.



**Figure 3: Preferences for the 4 levels of noise reduction and directionality, which correspond (from level 1 to level 4) to increasing directionality settings, increasing levels of noise reduction in simple and complex environments and earlier activation of noise reduction [15]. The participants exhibited different noise reduction and directionality preferences and five of them preferred more than one level in different situations.**



**Figure 4: Preferences for the 4 levels of brightness, which correspond (from level 1 to level 4) to increasing amplification of high-frequency sounds. The participants exhibited different brightness preferences and six of them preferred more than one level in different situations.**



**Figure 5: Preferences for the 4 levels of soft gain, which correspond (from level 1 to level 4) to increasing amplification of soft sounds, thus increasing dynamic range compression [14]. The participants exhibited different soft gain preferences and five of them preferred more than one level in different situations.**

In order to combine the sequentially learned preferences, we summed up the most frequently chosen values along each parameter into a single hearing aid configuration. For each participant, we subsequently compared it against individually personalized settings configured in a standard clinical workflow based on questions and pre-recorded sound samples. After the fourth week, six out of seven participants responded they appreciated having more than one general hearing aid setting, as they used both configurations in different situations. They also wished to keep both personalized configurations after the end of the test. However, in a blind comparison of the two configurations, six out of seven participants preferred the hearing aid settings personalized by sequentially optimizing parameters in real-world listening scenarios.

## 4 DISCUSSION

Due to the aging population, the number of people affected by hearing loss will double by 2050 [20] and this will have large implications for hearing healthcare. Rethinking hearing aids as recommender systems might enable the implementation of devices that automatically learn the preferred settings by actively involving hearing impaired users in the loop. Not only would this enhance the experience of current hearing aid users, but it could also help overcome the growing lack of clinical resources. Personalizing hearing aids by integrating audiological domain-specific recommendations might even make it feasible to provide scalable solutions for the 80% of hearing impaired users who currently have no access to hearing healthcare worldwide [19]. The accuracy of the recommendation primarily depends on the ability of the system to gather user preferences, while the user explores a highly complex design space. In this study, we proposed an approach to effectively optimize the device settings by decoupling three audiological parameters and allowing the participants to manipulate one parameter at a time, comparing four discrete levels. The fact that the participants preferred the hearing aid configuration personalized in real-world environments suggests that the proposed optimization approach manages to capture the main individual parameter preferences.

Looking into the individual preferences learned when sequentially adjusting the three parameters, several aspects stand out. The results suggest that the *brightness* parameter has the highest perceived usefulness. This could be due to the fact that enhancing the gain of high frequencies may increase the contrasts between consonants and as a result improve speech intelligibility. Likewise, it may amplify spatial cues reflected from the walls and ceiling, improving the localization of sounds and thereby facilitating the separation of voices. The participants seemed to appreciate a brighter sound when listening to speech or when paying attention to specific sources in a quiet environment. Despite the advances in technology that reduce the risk of audio feedback and allow the new instruments to be fitted to target and deliver the optimal gain [8], in some situations most of the participants seemed to benefit from even more brightness. Conversely, users might prefer a more round sound in noisy situations or when they want to detach themselves.

Adjusting the *noise reduction and directionality* parameter is perceived as having the lowest usefulness. Essentially, this parameter defines how ambient sounds coming from the sides and from behind are attenuated, while still amplifying signals with speech characteristics. Although the benefits of directionality and noise reduction are proven, our results indicate that users find it more difficult to differentiate the levels of this parameter if the ambient noise level is not sufficiently challenging. The four levels of the parameter mainly affect the threshold for when the device should begin to attenuate ambient sounds. However, these elements of signal processing are partly triggered automatically based on how noisy the environment is. Therefore, in some situations, changing the attenuation thresholds (i.e. the parameter levels) might not make a difference. Thus, users may feel less empowered to adjust this parameter. On the other hand, the data also shows that participants actively select the lowest level of the parameter (level 1), which provides an immersive omnidirectional experience without attenuation of ambient sounds in simple listening scenarios. This suggests that, in some contexts,

users express a need for personalizing the directionality settings and the activation thresholds of noise reduction. Furthermore, previous studies have shown that the perception of *soft sounds* varies largely among individuals. Our results not only confirm that users have widely different audiological preferences, but also suggest they would benefit from a personalized dynamic adaptation of soft gain dependent on the context.

Focusing on the optimization problem in the audiological design space, some indications can be inferred. The large differences among the participants suggest that, in a first phase, users' interaction is essential to gather individual preferences and thereby reach the optimum configuration for each single user. Simplifying the optimization task and offering a clear explanation of the one-dimensional slider made the process more transparent and increased users' empowerment. Once a recommender system is in place, this component might also prove useful in enhancing users' trust in the recommendations provided. Moreover, performing the optimization task in real-world environments ensured an accurate assessment and communication of users' preferences. In the short term, user preferences collected with this approach could flow into the standard clinical workflow and help hearing care professionals to fine-tune the hearing aids. However, a single static configuration, although personalized, might not fully satisfy the user. Our results indicate that such recommender systems should not simply model users as a sole set of optimized audiological parameters, because the preferred configuration varies depending on the context. It is therefore essential for these models to likewise classify the sound environment and motion state in order to infer the intents of the user. Being fully aware of the intent, by automatically labeling it, would add further value to the collected preferences and would allow to ask for user feedback in specific situations. That would make it feasible to verify hypotheses based on previous data, and progressively optimize several device configurations for different real-world listening scenarios. Once some configurations are learned, the hearing aids could automatically recommend them in specific situations and, by monitoring users' behavior, continuously calibrate to the preference of the user.

## 5 CONCLUSION

Internet-connected hearing aids open the opportunity for truly personalized hearing aids, which adapt to the needs of users in real-world listening scenarios. This study addressed the main challenges posed when rethinking hearing aids as recommender systems. It investigated how to effectively optimize the device settings by gathering user preferences in real-world environments. A complex audiological space was simplified by decoupling three audiological parameters and allowing the participants to manipulate one parameter at a time, comparing four discrete levels. The participants sequentially optimized the three audiological parameters, which were subsequently combined into a personalized device configuration. This configuration was blindly compared against a configuration personalized in a standard clinical workflow based on questions and pre-recorded sound samples, and six out of seven participants preferred the device settings learned in real-world listening environments. Thus, the approach seemed to effectively gather the main individual audiological preferences. The parameters resulted

to have a different perceived usefulness, differently contributing to the listening experience of hearing aid users. The seven participants exhibited widely different audiological preferences. Furthermore, our results indicate that hearing aid users do not simply explore the audiological design space in search of a global optimum. Instead, most of them select multiple highly contrasting values along each parameter, depending on the context.

## ACKNOWLEDGMENTS

We would like to thank Oticon A/S, Eriksholm Research Centre, and Research Clinician Rikke Rossing for providing hardware, access to test subjects, clinical approval and clinical resources.

## REFERENCES

- [1] Gediminas Adomavicius and Alexander Tuzhilin. 2011. Context-Aware Recommender Systems. In *Recommender Systems Handbook*, F. Ricci, L. Rokach, B. Shapira, and P. Kantor (Eds.). Springer, Boston, MA, USA, 217–253. [https://link.springer.com/chapter/10.1007/978-0-387-85820-3\\_7](https://link.springer.com/chapter/10.1007/978-0-387-85820-3_7)
- [2] Eric Brochu, Vlad M. Cora, and Nando de Freitas. 2010. A Tutorial on Bayesian Optimization of Expensive Cost Functions, with Application to Active User Modeling and Hierarchical Reinforcement Learning. *CoRR abs/1012.2599* (2010).
- [3] Eric Brochu, Nando de Freitas, and Abhijeet Ghosh. 2007. Active Preference Learning with Discrete Choice Data. In *Proceedings of the 20th International Conference on Neural Information Processing Systems (NIPS '07)*, 409–416.
- [4] Mead C. Killion. 2002. New Thinking on Hearing in Noise: A Generalized Articulation Index. *Seminars in Hearing* 23 (January 2002), 057–076. <https://doi.org/10.1055/s-2002-24976>
- [5] Susanna L. Callaway and Andreea Micula. 2016. *Client Target and Real-ear Measurements*. Technical Report. Oticon A/S, Smørum, Denmark.
- [6] John Clark. 1981. Uses and Abuses of Hearing Loss Classification. *ASHA: a journal of the American Speech-Language-Hearing Association* 23 (August 1981), 493–500.
- [7] European Commission. 2019. Glossary: Hearing threshold. Retrieved August 16, 2019 from [https://ec.europa.eu/health/scientific\\_committees/opinions\\_layman/en/hearing-loss-personal-music-player-mp3/glossary/ghi/hearing-threshold.htm](https://ec.europa.eu/health/scientific_committees/opinions_layman/en/hearing-loss-personal-music-player-mp3/glossary/ghi/hearing-threshold.htm)
- [8] Josefine J. Jensen. 2019. *Oticon Opn S Clinical Evidence*. Technical Report. Oticon A/S, Smørum, Denmark.
- [9] Benjamin Johansen, Maciej J. Korzepa, Michael K. Petersen, Niels H. Pontoppidan, and Jakob E. Larsen. 2018. Inferring User Intents from Motion in Hearing Healthcare. In *Proceedings of the 2018 ACM International Joint Conference and 2018 International Symposium*. <https://doi.org/10.1145/3267305.3267683>
- [10] Benjamin Johansen, Michael K. Petersen, Niels H. Pontoppidan, and Jakob E. Larsen. 2019. Modelling User Utterances as Intents in an Audiological Design Space. In *Workshop on Computational Modeling in Human-Computer Interaction (CHI 2019)*.
- [11] Benjamin Johansen, Michael K. Petersen, Maciej J. Korzepa, Jan Larsen, Niels H. Pontoppidan, and Jakob E. Larsen. 2018. Personalizing the Fitting of Hearing Aids by Learning Contextual Preferences From Internet of Things Data. *Computers* 7, 1 (2018). <https://doi.org/10.3390/computers7010001>
- [12] Maciej J. Korzepa, Benjamin Johansen, Michael K. Petersen, Jan Larsen, Niels H. Pontoppidan, and Jakob E. Larsen. 2018. Learning Preferences and Soundscapes for Augmented Hearing. In *IUI Workshops*.
- [13] Yuki Koyama, Issei Sato, Daisuke Sakamoto, and Takeo Igarashi. 2017. Sequential Line Search for Efficient Visual Design Optimization by Crowds. *ACM Trans. Graph.* 36, 4, Article 48 (July 2017), 11 pages. <https://doi.org/10.1145/3072959.3073598>
- [14] Nicolas Le Goff. 2015. *Amplifying Soft Sounds - a Personal Matter*. Technical Report. Oticon A/S, Smørum, Denmark.
- [15] Nicolas Le Goff, Jesper Jensen, Michael S. Pedersen, and Susanna L. Callaway. 2016. *An Introduction to OpenSound Navigator™*. Technical Report. Oticon A/S, Smørum, Denmark.
- [16] Nicholas A. Lesica. 2018. Why Do Hearing Aids Fail to Restore Normal Auditory Perception? *Trends in Neurosciences* 41, 4 (April 2018), 174–185. <https://doi.org/10.1016/j.tins.2018.01.008>
- [17] Jeremy Marozeau and Mary Florentine. 2007. Loudness Growth in Individual Listeners with Hearing Losses: A Review. *The Journal of the Acoustical Society of America* 122, 3 (2007), EL81–EL87. <https://doi.org/10.1121/1.2761924>
- [18] Elaine H. N. Ng, Mary Rudner, Thomas Lunner, Michael Syskind Pedersen, and Jerker Rönnberg. 2013. Effects of Noise and Working Memory Capacity on Memory Processing of Speech for Hearing-aid Users. *International Journal of Audiology* 52, 7 (2013), 433–441. <https://doi.org/10.3109/14992027.2013.776181>

- [19] World Health Organization. 2013. *Multi-Country Assessment of National Capacity to Provide Hearing Care*. Geneva, Switzerland.
- [20] World Health Organization. 2019. Deafness and Hearing Loss. Retrieved June 30, 2019 from <https://www.who.int/news-room/fact-sheets/detail/deafness-and-hearing-loss>
- [21] Antti Oulasvirta, Per O. Kristensson, Xiaojun Bi, and Andrew Howes. 2018. *Computational Interaction*. Oxford University Press.
- [22] Jonathan E. Peele and Arthur Wingfield. 2016. The Neural Consequences of Age Related Hearing Loss. *Trends in Neurosciences* 39, 7 (2016), 486–497. <https://doi.org/10.1016/j.tins.2016.05.001>
- [23] Dorothea Wendt, Renskje K. Hietkamp, and Thomas Lunner. 2017. Impact of Noise and Noise Reduction on Processing Effort. *Ear and Hearing* 38, 6 (2017), 690–700. <https://doi.org/10.1097/aud.0000000000000454>