Improving Hearing Aid Personalization Algorithm Efficiency with User Preference Correlations

Justin Jensen, Dhruv Vyas, Yumna Anwar, Erik Jorgensen, Octav Chipara, Yu-Hsiang Wu

Research Mentor and other collaborators

Yu-Hsiang Wu & Emily Andrews

Objectives

Determining the best hearing aid (HA) configuration among all possible configurations is challenging. Prior research has proposed using machine learning algorithms to evaluate the search space of available configurations by learning from user feedback to converge on an optimal configuration. Our previous study evaluated several HA personalization algorithms and compared their accuracy, consistency, and efficiency in converging to a user's most preferred option. The results indicated that although the best-performing algorithm—the Upper Confidence Bound (UCB) algorithm—could converge accurately and consistently, its efficiency (i.e. the total number of paired-comparisons necessary to identify the best configuration) was not good enough to be used in the real world. The purpose of the present study was to assess whether the efficiency of the UCB algorithm could be improved by incorporating correlations between patterns of preference from a group of users in addition to the feedback directly obtained from the user.

Design

Fifteen HA gain-frequency response configurations (i.e. fifteen presets) were developed using audiometric data from a national health database. Speech was recorded from the output of a HA programmed with each of the 15 preset configurations and presented to 32 older adults with hearing loss. The participants used a paired comparison paradigm to determine the order and strength of preference for all possible combinations of configurations (105 pairwise comparisons x 4 repetitions). Order of preference was determined using a Borda scoring method, where a configuration's Borda score is the ratio between the number of times it was preferred and the total number of pairwise comparisons for that configuration. Correlated pairs of configurations were identified by computing the Pearson correlation coefficient for all configuration pairs across all 32 users. A positively correlated pair has a correlation coefficient greater than 0.7, and the configurations have similarly high/low Borda scores. Alternatively, a negatively correlated pair has a correlation coefficient smaller than -0.7, and one configuration tends to be liked (high Borda score) while the other is disliked (low Borda score). Correlation results were used to determine how many paired comparisons could be removed from the search space and still yield accurate algorithm convergence. Root Mean Square (RMS) differences between correlated configurations were calculated.

Results

Positively correlated configurations had an average RMS difference of 4.52 dB, with differences driven by mid-to-high frequency bands (1 – 6 kHz). Negatively correlated configurations had an average RMS difference of 9.47 dB, with differences spread across all frequencies. Among the 105 possible pairwise comparisons, 33 correlated pairs of configurations were identified, meaning that the outcome of one trial could yield multiple conclusions about the outcomes of other pairwise trials. Accordingly, the correct order of a given user's preferences can be determined with only 72 comparison trials compared to the 105 total possible combinations, a reduction of 32%.

Conclusions

Incorporating correlations between user preferences can reduce the search space for an optimal HA configuration by as much as 32%. This can translate to a more efficient convergence of the selection algorithm, which would require less time and effort from the HA user to find their best HA configuration. Additional research is needed to assess to what degree this reduced search space improves efficiency.