“Please sort these sounds into 2 identities”: Effects of task instructions on performance in voice sorting studies

Nadine Lavan^{1,2}, Siobhan E. Merriman², Paayal Ladwa², Luke F.K. Burston², Sarah Knight¹ and Carolyn McGettigan¹

¹ *Department of Speech, Hearing and Phonetic Sciences, University College London*

² *Department of Psychology, Royal Holloway, University of London*

Correspondence to:

Nadine Lavan, Department of Speech, Hearing and Phonetic Sciences, University College London, 2 Wakefield Street, London WC1N 1PF, United Kingdom. E-mail:

n.lavan@ucl.ac.uk

or

Carolyn McGettigan, Department of Speech, Hearing and Phonetic Sciences, University College London, 2 Wakefield Street, London WC1N 1PF, United Kingdom. E-mail:

c.mcgettigan@ucl.ac.uk

Acknowledgements: This work was supported by a Research Leadership Award from the Leverhulme Trust (RL-2016-013) awarded to Carolyn McGettigan

THIS MANUSCRIPT HAS NOT BEEN PEER-REVIEWED

Abstract

In the current study we investigated the effects of two types of task instructions on performance of a voice sorting task by familiar and unfamiliar listeners. In Experiment 1, listeners were asked to sort 15 naturally-varying exemplars from two voice identities into as many identities as they perceived. Results replicate the findings of previous voice and face sorting studies: unfamiliar listeners form more clusters than familiar listeners, selectively failing to “tell people together”. That is, unfamiliar listeners erroneously perceive exemplars from the same identity as coming from different identities, while being highly accurate at telling apart the exemplars from different voices. In Experiment 2, we show that a change in task instructions - forcing participants to sort exemplars into two identities only - allows unfamiliar listeners to overcome this selective failure at “telling people together”. This improvement, however, comes at the cost of an increase in errors in telling people apart. For familiar listeners, similar non-significant trends are apparent. The current study thus shows that, even when being given the correct number of identities, listeners may fail to accurately perceive identity. These findings highlight voice identity perception in the context of natural within-person variability as a challenging task for both unfamiliar and familiar listeners.

Keywords: person perception, voice identity, voice sorting, within-person variability

Introduction

Voice identity perception in the context of variable stimuli has been shown to be a challenging task for listeners (see Lavan, Burton, Scott & McGettigan, 2018a for a review) because talkers constantly adjust what they are saying, and how they are saying it, for different surroundings and audiences. These adjustments introduce variability, such that the

sound of the voice can vary dramatically between, for example, what we sound like when we shout over background noise to when we mutter to ourselves while trying to solve a problem. In order to successfully perceive voice identity, listeners therefore need to not only be able to perceive different voices as indeed coming from different people: they also need to generalise across this within-person variability to link different-sounding samples of a single voice to the same person. This is especially challenging when dealing with voices that are unfamiliar to listeners. Studies have shown that listeners are less accurate in discriminating speakers when presented with unfamiliar voices containing within-person variability: Speaker discrimination accuracy has been shown to decrease when making judgements across disguised and undisguised voices (Reich and Duke, 1979), sung and spoken words (Peynircioğlu, Rabinovitz, & Repice, 2017), different vocalisations (Lavan, Scott, & McGettigan, 2016) and recordings produced in different languages (Wester, 2012). While familiar voice identity perception is generally more reliable, listeners can nonetheless be affected by within-person vocal variability. Lavan et al. (2016) have shown that, while familiar listeners systematically outperformed unfamiliar listeners on a speaker discrimination task, performance nonetheless decreased in the presence of extensive within-person variability (laughter vs. vowels sounds). Similarly, Wagner and Köster (1999) report that familiar voices are essentially unrecognisable when produced using falsetto voice.

Theoretical accounts model familiar and unfamiliar voice identity perception as at least partially distinct processes. Maguiness, Roswadowitz & Von Kriegstein (2018) highlight how familiar and unfamiliar voice identity perception may dissociate on a neural and behavioural level. Kreiman and Sidtis (2011) propose that unfamiliar voices are processed based on their acoustic and perceptual features in a stimulus-driven way. In a situation that

would require listeners to explicitly discriminate between two stimuli, this may take the shape of close comparisons between perceptually salient features of the stimuli until a decision is reached as to whether the two stimuli are likely to be produced by the same person or two different people. In contrast to this, familiar voices are thought to be processed in a holistic way, based on the detection of perceptual features that are diagnostic for a familiar identity. These features may take the form of a particular quality of voice or characteristic speaking style that allow us to detect voices as being familiar – notably, the nature of these features may vary from voice to voice, and from listener to listener (Kreiman & Sidtis, 2011). Through the detection of such diagnostic features of a familiar voice, vocal signals produced by familiar speakers are thought to be matched to representations of the specific speaker's vocal inventory that are stored in long-term memory.

More empirical evidence for dissociations between familiar and unfamiliar voice identity perception has recently been reported for studies using voice sorting paradigms (Lavan, Burston & Garrido, 2018b; Lavan, Burston, Ladwa, Merriman, Knight & McGettigan, 2018c). In such voice sorting paradigms, listeners are exposed to naturally-varying recordings of two voice identities – specifically, recordings from dialogues extracted across different scenes and episodes of TV shows. Each voice identity is represented by a number of exemplars. Participants, who are either familiar with the TV show or not, are then asked to sort the recordings into voice identities. While familiar listeners sort the sounds into a relatively small number of voice identities, unfamiliar listeners tend to perceive a much larger number. An analysis of the composition of the perceived voice identities (i.e. examining the patterns of exemplars within them) reveals that both groups are highly accurate at “telling people apart” – that is, listeners are able to distinguish between different identities despite the within-person variability. This is evidenced by the very small number of mixing errors, which arise

when listeners sort exemplars from different identities into one perceived identity. Unfamiliar listeners, however, fail to “tell people together”, because they split naturally varying exemplars produced by one person into a larger number of perceived identities. While familiar listeners also make some “telling people together” errors, they tend to be able to perceive variable exemplars produced by the same voice as belonging to one identity only, despite obvious acoustic and perceptual discrepancies between the exemplars.

It has been proposed that this failure in “telling people together” is due to unfamiliar listeners having no person-specific representation of the voice that would allow them to know *how* a novel voice varies and thus recognise within-person variability as such. The subjective experience for listeners when exposed to naturally varying exemplars of unknown voices is thus mostly driven by the perceptual differences between exemplars, with within-person variability perceived as between-person variability. In voice sorting tasks, this then leads to the characteristic pattern of results, in which unfamiliar listeners form a greater number of clusters than identities present and thus fail to “tell people together”. Since this pattern of results seems to be specifically due to a failure to “tell people together”, we ran two voice sorting experiments, across which we manipulated the task instructions. First, in a free sorting task, we aimed to replicate findings from previous sorting tasks. We thus instructed participants to sort 30 brief recordings (2 identities x 15 exemplars each) into identities, without giving any guidance as to how many identities may be present. Then in a second experiment, and in line with a previous study using a face sorting paradigm (Andrews, Jenkins, Cursiter & Burton, 2015), we ran a restricted sorting task: Here, participants were told that only 2 identities were present and instructed to sort the recordings accordingly into 2 clusters. We predicted that, for the free sorting task, unfamiliar listeners would form more clusters than familiar listeners by failing to “tell people together”, that is, perceiving different

exemplars from the same voice identity as different voice identities due to the natural within-person variability (Lavan et al., 2018b, Lavan et al., 2018c, see also Andrews et al., 2015; Jenkins, White, Van Montfort & Burton, 2011; Redfern & Benton, 2017; Zhou & Mondloch, 2016 for faces). Based on previous voice sorting studies, we furthermore predicted that there would be only few errors for “telling people apart”. For restricted sorting, we predicted that errors in “telling people apart” and “telling people together” should be reduced for both groups.

Methods

Participants

84 participants completed the study. Sample size was determined to match previous face and voice sorting studies: there were at least 20 participants per sub-sample (Jenkins et al., 2011; Lavan et al., 2018b; Lavan et al., 2018b). Participants were recruited via Prolific.ac and survey exchange services, such as surveytandem.com. Those recruited via Prolific were compensated for their time. The study was approved by the local ethics committee. We recruited familiar and unfamiliar listeners: if participants reported having watched at least one season of *Breaking Bad*, they were assigned to the familiar group. Participants who reported not having seen any episodes of the TV show were assigned to the unfamiliar group. A number of participants were excluded based on the following criteria: 4 participants who had not seen a full season of *Breaking Bad* were excluded; 2 unfamiliar listeners were excluded since they reported having recognised one of the voices (Bryan Cranston) from other TV shows; 7 further listeners were excluded as they created more clusters in the restricted sorting condition than instructed; 4 listeners failed the attention check (see *Materials*; these listeners did not sort 2 exemplars of a computer-generated voice into a

separate cluster); finally, 1 listener was excluded since they moved less than 80% of the sounds on the slide (thus did not appear to engage with the task).

In total, 80 participants are thus included in the current experiments: 20 familiar (mean age: 23.1 years, SD: 3.6 years, 14 female) and 20 unfamiliar listeners (mean age: 23.0 years, SD: 4.4 years, 13 female) for the free sorting experiment, and a further 20 familiar (mean age: 26.3 years, SD: 5.0 years, 7 female) and 20 unfamiliar (mean age: 27.6 years, SD: 6.3 years, 14 female) for the restricted sorting experiment. The data from 13 participants for each listener group in the free sorting experiment was part of the data set reported in Lavan et al. (2018c): these are the participants who were first presented with the low-expressiveness condition (which is identical to the task in the current study). To reach the target of 20 participants, 7 participants for each listener group were additionally tested for the purposes of the current study.

Materials

The materials used in this experiment are the same as the one used for the “low expressiveness” condition in Lavan et al. (2018c). Audio clips (duration between 1.2 and 4 seconds) from two of the prominent characters of the TV show *Breaking Bad* (Hank Schrader and Walter White) were extracted. The clips contained meaningful utterances with only minimal background noise and there was no interference from other voices. No iconic catchphrases or otherwise diagnostic linguistic information (e.g. referring to a character’s job, etc.) were present in the stimuli. Other than featuring emotionally neutral speech, the sound clips thus varied naturally in all features (e.g. speaking style, verbal content, speaking environment and conversation partners). Exemplars were normalized for peak amplitude (to

0.400 Pa), long silences were cut and all sounds low-pass filtered at 10kHz (using a Hann pass-band filter with upper and lower edges 0Hz and 10000Hz, smoothing 20Hz) using Praat (Boersma & Weenink, 2018) to account for systematic differences in the audio quality. The 30 exemplars (2 identities [Hank, Walter] x 15 exemplars) were embedded into a Microsoft Powerpoint slide. Two identical exemplars spoken by a synthetic female voice, saying “Hello. My name is Sarah” were also included as attention checks to verify that participants were completing the task correctly (i.e. by forming a single identity cluster for the 2 female voice exemplars on each slide; see exclusion criteria). Each embedded sound was represented by a number on the screen. These numbers were evenly distributed across the slide, with no clusters being obvious from the outset.

Procedure

For the free sorting task, participants were instructed to sort the exemplars into clusters, so that each cluster included the exemplars produced by a single speaker, and thus represented a perceived speaker identity. Therefore, participants could form any number of clusters. For the restricted sorting task, participants were told that there were three speakers (i.e. 2 voices from *Breaking Bad* plus the female voice) present in the task and that they should therefore sort the 32 exemplars into 3 clusters. There was no limit on how many times participants could play the sounds. Participants completed the task online via Qualtrics.

Results

Confirmatory analysis: Number of clusters in the free sorting task

Figure 1 Number of perceived identities for familiar and unfamiliar listeners by instruction type. Bars show the means across participants, and each dot shows one participant. Boxes show the 95% confidence intervals for the means (only shown for the free sorting task, since all listeners followed the instructions for the restricted sorting task and formed two clusters only).

Initially, we counted the number of clusters formed by each participant in the free sorting task and the restricted sorting task (after removing the 'catch' exemplars). Figure 1 (right-hand plot) illustrates that all participants included in the restricted sorting task formed two clusters as per the instructions. For the free sorting task, data were not normally distributed in most cases, as indicated by Shapiro-Wilk tests. We therefore used non-parametric tests for the following analyses in the R environment using the *coin* package. Familiar listeners perceived significantly fewer clusters than unfamiliar listeners (Familiar: Mode = 2, Median = 4, Range = 2-9; Unfamiliar: Mode = 10, Median = 10, Range = 4-20; Wilcoxon's signed rank test: $Z = 4.17$, $p < .001$). This replicates the findings of other voice and face sorting studies: participants who are familiar with the faces or voices used in the studies perceive a lower number of identities than participants who are not familiar with the faces or voices.

Confirmatory analysis: The effects of familiarity and instruction type on performance for “telling people apart” and “telling people together”

Figure 2 a) Matrices of averaged listeners' responses for the voice sorting task for familiar and unfamiliar listeners. Within these 30 x 30 matrices (15 sounds files x 2 identities), each cell shows the probability that two exemplars were grouped within the same perceived identity: cells with a value of 1 indicate that the respective exemplars were always clustered together, cells with a value of 0 indicate that these sounds were never in the same clusters. b) Illustration of the different sections of the per-participant matrices that were analysed below.

To assess *how* listeners formed clusters, we created participant-wise response matrices. In these matrices, each cell codes for whether the corresponding pair of exemplars was placed within the same cluster (coded as 1) or placed in two separate clusters (coded as 0). The group-averaged response matrices are shown in Figure 2a. These matrices are symmetrical across the diagonal and can be conceptually divided into within-identity submatrices, indexing listeners' performance for “telling people together” and across-identity submatrices, indexing listeners' performance for “telling people apart” (see Figure 2b).

To explore the effects of familiarity (familiar vs. unfamiliar) and instruction type (free vs. restricted sorting) on listeners' performance for "telling people together" and "telling people apart", we computed the participant-wise averages of the within-identity and across-identity submatrices respectively (see Figure 2b). Perfect performance (i.e. forming two clusters of 15 exemplars, with correct assignment of all exemplars to their corresponding identity) would result in an average of 1 for the within-identity submatrices and an average of 0 for across-identity submatrix (for a detailed description of the analyses, see Lavan et al., 2018b).

Figure 3 a) Plots of the participant's accuracy for "telling people together" and "telling people apart" by group and task instruction. B) Plots of the cumulative error rates of "telling people apart" and "telling people together" plotted by group and task instructions. Boxes on all plots show the 95% confidence intervals for the means.

We first investigated how instruction type affected performance. For familiar listeners, Wilcoxon's rank-sum tests indicated that task performance for "telling together" and "telling apart" was not significantly affected by instruction type (telling together: $Z = 1.38$, $p = .168$; telling apart: $Z = 1.92$, $p = .055$). Numerical trends, however, indicate that, while performance improved for "telling people together", listeners made more errors for "telling people apart" during the restricted sorting task (Figure 3a). For unfamiliar listeners, performance for "telling people together" was dramatically better for the restricted sorting task ($Z = 5.20$, $p < .001$). Conversely, however, errors for "telling people apart" also significantly increased ($Z = 4.63$, $p < .001$).

We then compared task performance for the two listener groups. Unfamiliar listeners' performance was consistently worse compared to the performance of familiar listeners (all $Zs > 2.69$, all $ps < .008$), with the exception of "telling people apart" for the free sorting task ($Z = .62$, $p = .523$).

Exploratory analysis: Cumulative error rates

We conducted exploratory analyses to further explore how the instruction type affected the cumulative error rates across both "telling people together" and "telling people apart". To compute the cumulative error rates, we summed the individual error rates for "telling people together" (1 minus the mean probability of the lower triangle of each within-identity matrix for each participant) and the error rates for "telling people apart" (original telling apart measure) per participant, separately for the restricted and free sorting tasks (maximum cumulative error rate = 1). Wilcoxon's rank-sum tests showed that, while cumulative error

rates were dramatically lower for unfamiliar listeners for the restricted sorting task ($Z = 4.03$, $p < .001$, see Figure 3b), error rates remained stable for familiar listeners ($Z = .01$, $p = .989$).

Exploratory analysis: Similar patterns of errors across listener groups and instruction types

To further investigate whether the geometries of responses are similar for different listener groups and instructions types, we computed Kendall's τ_a for average “telling together” and “telling apart” matrices across instruction type and across listener groups. If Kendall's τ_a was higher than 95% of the chance predictions ($p < 0.05$) obtained by shuffling the values within the comparisons of interest (5000 permutations), we rejected the null hypothesis. These analyses showed that matrices were significantly correlated across instruction types for both “telling apart” and “telling together”, for both familiar listeners (“telling together”: $\tau_a = .56$, $p < .001$; “telling apart”: $\tau_a = .30$, $p < .001$) and unfamiliar listeners (“telling together”: $\tau_a = .30$, $p < .001$; “telling apart”: $\tau_a = .35$, $p < .001$). Similarly, matrices for both “telling apart” and “telling together” were significantly correlated when comparing familiar listeners and unfamiliar listeners on the free sorting task (“telling together”: $\tau_a = .26$, $p < .001$; “telling apart”: $\tau_a = .12$, $p = .001$) and the restricted sorting task (“telling together”: $\tau_a = .18$, $p < .001$; “telling apart”: $\tau_a = .48$, $p < .001$). Thus, while manipulating the instruction changes overall performance, especially in the case of unfamiliar listeners, the patterns of responses and errors appear to remain partially stable across instruction type. Strikingly, these patterns also seem to be stable across listener groups, indicating that some of the underlying processes leading to errors may be shared despite the differences in familiarity.

Discussion

In the current study, we investigated how changes in the instructions for a voice sorting task affect familiar and unfamiliar listeners' sorting behaviour. Half of the participants received

task instructions previously used in other voice sorting tasks: here, listeners were asked to sort voice clips into as many identities as they perceived, without being given any information on how many voices were actually present. We replicated previous findings from the face and voice perception literature, showing that unfamiliar listeners perceived more identities to be present than familiar listeners (unfamiliar listeners median = 10 identities, mode = 10 identities; familiar listeners median: 4 identities, mode = 2 identities; see also Lavan et al., 2018b, Lavan et al., 2018c; Jenkins et al., 2011 for faces). Familiar and unfamiliar listeners were similarly accurate in telling people apart, with only relatively few mixing errors occurring for both groups. Unfamiliar listeners, however, selectively failed to “tell people together”; that is, they failed to generalise across the within-person variability and thus misperceived within-person variability as between-person variability. While performance for familiar listeners was not perfect (cf. Lavan et al., 2018b; Lavan et al., 2018c), errors were significantly less frequent for this listener group compared to unfamiliar listeners.

With a different set of task instructions, the pattern of results changed. Here, we implemented a restricted sorting task, where we asked listeners to sort the 30 exemplars into 2 identities (plus the identity included as an attention check) and thus forced listeners to “tell people together”. For these task instructions, the selective failure in “telling people together” largely disappeared in a separate group of unfamiliar listeners: the probability that unfamiliar listeners sorted exemplars from the same identity into the correct cluster rose sharply from 0.20 when given the free sorting instructions to 0.70 for the restricted sorting instructions. This drastic improvement in accuracy is, however, likely to be underpinned by very different perceptual processes compared to the ones underpinning good performance for familiar listeners: while familiar listeners can access voice identity representations that are likely to include information on how voices vary, thus allowing them to tell people

together, unfamiliar listeners do not have such representations. To arrive at this more accurate solution for the restricted sorting task, unfamiliar listeners are likely to perform an explicit matching process akin to an *ABX* task: if a listener identifies that a specific exemplar is less likely to belong to one of the clusters, they will have to sort it into the second cluster to follow the “restricted sorting” task instructions. While our data show that unfamiliar listeners are likely correct in these matching decisions, as reflected in dramatically better performance in the restricted sorting task, these listeners are unlikely to actually perceive all the exemplars within each cluster as coming from the same identity. A similar improvement, albeit statistically non-significant, was also apparent for familiar listeners. While familiar listeners were already performing well on this aspect of the task for free sorting, the probability of correctly “telling together” the different exemplars increased from 0.73 to 0.83 with the restricted sorting instruction¹. Crucially, however, these improvements for both groups came at a cost: unfamiliar listeners made significantly more mixing errors for the “restricted sorting” instructions compared to the “free sorting” instructions. The same pattern was also apparent for familiar listeners, although again it was not statistically significant. These findings align reasonably with what has been reported in the face perception literature. However, for face sorting, Andrews et al. (2015) do not report a significant increase in mixing errors (i.e. errors in “telling people apart”) between free sorting and restricted sorting groups. Finding a significant increase in “telling apart” errors for voices is thus an interesting example of how face and voice identity processing may differ from one another. These differences between modalities may arise due to different processes being

¹ While familiar listeners generally perform well for “telling people together” in voice (and face) sorting tasks, we note that due to the nature of the tasks, performance may be somewhat inflated for “free sorting”: in these tasks, listeners who arrived at the correct solution of forming 2 clusters with one identity in each ($N = 7$ in the current data set) may have indeed recognised every single exemplar as belonging to the relevant identity. Alternatively, they may have gleaned that only two identities are likely to be present, and thus assigned ambiguous sounds to the more likely cluster in the absence of genuine recognition. Thus, “perfect performance” in the task may reflect a combination of successful voice identity processing and task-specific decision-making strategies.

at work during face and voice identity processing, reflecting differences in the nature of the stimuli (visual vs. auditory; static vs. dynamic). Since voice identity processing appears to be generally less robust compared to face identity processing (Barsics, 2014; Stevenage, Hugill, Lewis, 2012), these differences could, however, also reflect a floor effect for mixing errors in the face sorting study.

In the current study, exploratory analyses highlighted that, while the “restricted sorting” instructions reduced the overall number of errors made by unfamiliar listeners due to the dramatic decrease in “telling people together” errors, cumulative error rates remained stable for familiar listeners across the two instruction types. Thus, for familiar listeners, the reduction of errors made in “telling people together” was balanced by the increase in errors for “telling people apart”. An analysis of the geometries of the response patterns comparing groups and instruction types shed further light on these errors. Here, significant correlations across listener groups and task instructions indicated that patterns of errors are shared for the voice stimuli used in this study. Thus, despite overall differences in accuracy between groups and task instructions, exemplars that were relatively harder to tell together/apart in the free sorting task were still the ones that were difficult to tell together/apart in the restricted sorting task – for unfamiliar and familiar listeners alike. The correlations of geometries therefore reflect the presence of exemplar-specific effects: within the set of naturally varying stimuli, some exemplars are likely to be a relatively bad likeness of a speaker’s voice due to specific acoustic and perceptual features. These exemplars are thus more difficult for familiar listeners to explicitly recognise (or, alternatively, these exemplars may be more susceptible to misrecognition; see Lavan et al., 2018c). In the absence of explicit recognition, familiar and unfamiliar listeners alike should thus process identity based on the acoustic and perceptual features of a signal (see Kreiman & Sidtis, 2011). These features

may make, for example, an exemplar of Walter White's voice stand out as sounding equally dissimilar to the remaining Walter and Hank exemplars. Listeners will perceive this exemplar as a single-exemplar identity in the free sorting task and, for the restricted sorting task, will sometimes assign the exemplar to "Walter"/Cluster A and sometimes to "Hank"/Cluster B. Thus, such exceptional exemplars will exhibit higher errors rates independent of task instructions and listeners groups, giving rise to similarities in the geometries of the response patterns.

The current study thus further highlights the challenges of voice identity perception in the context of within-person variability. Within-person variability has previously been shown to lead to failures in "telling people together", mostly for unfamiliar voices. In the current study, we additionally show that, even when provided with the correct number of identities, within-person variability can be sufficiently substantial for listeners to assign an exemplar into the wrong identity when forced to make a decision. This is most strongly exemplified by the significant increase in errors for "telling people apart" for unfamiliar listeners given the correct number of identities. This is particularly striking given that the within-person variability included in the current stimuli is unlikely to cover each speaker's full vocal inventory: all exemplars were extracted from relatively neutral utterances and did not, for example, include any whispered speech, laughter or screaming (see e.g. Lavan et al., 2016 for Lavan et al., 2018c for effects of such expressive voices on identity perception).

Voice sorting has proven an extremely useful tool for quantifying voice identity perception within a task that is readily understandable and intuitive to participants. In prior investigations of identity perception from voices, familiarity has been somewhat conflated with task design, such that – broadly – familiar listeners are tested on recognition and unfamiliar listeners on

voice matching or discrimination tasks. These traditional tasks load differentially on “telling together” and “telling apart”, while sorting tasks allow us to explore both simultaneously. In the current study, we have replicated previously-reported differences in the type and rate of errors in voice identity perception shown by familiar and unfamiliar listeners. To this, we have added new insights into the effects of task instruction on performance, showing that informing listeners of the number of underlying voice identities drastically improves the ability of unfamiliar listeners to “tell people together” – however, these improvements are not sufficient for familiar-like performance. The challenge for future work will be to chart how familiar levels of performance emerge as voices are learned, in particular the timecourse of errors and how this relates to the level of exposure and familiarity with voice identities.

References

- Andrews, S., Jenkins, R., Cursiter, H., & Burton, A. M. (2015). Telling faces together: Learning new faces through exposure to multiple instances. *The Quarterly Journal of Experimental Psychology*, 68(10), 2041-2050.
- Barsics, C. G. (2014). Person recognition is easier from faces than from voices. *Psychologica Belgica*, 54(3), 244-254.
- Jenkins, R., White, D., Van Montfort, X., & Burton, A. M. (2011). Variability in photos of the same face. *Cognition*, 121(3), 313-323.
- Kreiman, J., & Sidtis, D. (2011). *Foundations of voice studies: An interdisciplinary approach to voice production and perception*. John Wiley & Sons.
- Lavan, N., Burston, L. F., & Garrido, L. (2018). How many voices did you hear? Natural variability disrupts identity perception from unfamiliar voices. *British Journal of Psychology*.
- Lavan, N., Burston, L. F., Ladwa, P., Merriman, S. E., Knight, S., & McGettigan, C. (2018). Breaking voice identity perception: Expressive voices are more confusable for listeners. *PsyArXiv*. May, 23.
- Lavan, N., Burton, A. M., Scott, S. K., & McGettigan, C. (2017). Flexible voices: identity perception from variable vocal signals.
- Lavan, N., Scott, S. K., & McGettigan, C. (2016). Impaired generalization of speaker identity in the perception of familiar and unfamiliar voices. *Journal of Experimental Psychology: General*, 145(12), 1604-1614.

- Maguinness, C., Roswadowitz, C., & von Kriegstein, K. (2018). Understanding the mechanisms of familiar voice-identity recognition in the human brain. *Neuropsychologia*.
- Peynircioğlu, Z. F., Rabinovitz, B. E., & Repice, J. (2017). Matching Speaking to Singing Voices and the Influence of Content. *Journal of Voice*, 31(2), 256-e13.
- Redfern, A. S., & Benton, C. P. (2017). Expressive faces confuse identity. *i-Perception*, 8(5), 2041669517731115.
- Reich, A. R., & Duke, J. E. (1979). Effects of selected vocal disguises upon speaker identification by listening. *The Journal of the Acoustical Society of America*, 66(4), 1023-1028.
- Stevenage, S. V., Hugill, A. R., & Lewis, H. G. (2012). Integrating voice recognition into models of person perception. *Journal of Cognitive Psychology*, 24(4), 409-419.
- Wagner, I., & Köster, O. (1999). Perceptual recognition of familiar voices using falsetto as a type of voice disguise. In *Proceedings of the XIVth International Congress of Phonetic Sciences, San Francisco* (pp. 1381-1385).
- Wester, M. (2012). Talker discrimination across languages. *Speech Communication*, 54(6), 781-790.
- Zhou, X., & Mondloch, C. J. (2016). Recognizing “Bella Swan” and “Hermione Granger”: No own-race advantage in recognizing photos of famous faces. *Perception*, 45(12), 1426-1429.