A model of person perception from voices

Nadine Lavan and Carolyn McGettigan

Department of Experimental and Biological Psychology, Queen Mary University of London

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

Correspondence to:

Nadine Lavan, Department of Experimental and Biological Psychology, Queen Mary University of London, Mile End Rd, London E1 4NS, United Kingdom. E-mail: n.lavan@qmul.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 Sir Henry Wellcome Fellowship awarded to Nadine Lavan (220448/Z/20/Z) and a Research Leadership Award from the Leverhulme Trust (RL-2016-013) awarded to Carolyn McGettigan

Abstract

When hearing a voice, listeners can form a detailed impression of the person behind the voice. Existing models of voice processing focus primarily on one aspect of person perception - identity recognition from familiar voices - but do not account for the perception of other person characteristics (e.g., sex, age, personality traits). Here, we present a broader perspective, proposing that listeners have a common perceptual goal of perceiving who they are hearing, whether the voice is familiar or unfamiliar. We outline and discuss a model that achieves this goal via a common mechanism of recognising a familiar person, persona, or set of speaker characteristics.

Keywords: person perception, voice identity, unfamiliar voices, first impressions, familiarity

1 Introduction

When we hear a voice, we instantly form an impression of the person who is talking: This process can be referred to as *person perception from voices*. Person perception from voices can include the processing of any number of person characteristics: Is the person we are hearing likely to be male or female? Are they young or old? Do they sound like they are friendly, aggressive, or annoyed? Are they talking in their first language? If we know the specific person who is speaking, we may also be able to identify and name them.

Given that listeners may perceive a wealth of characteristics about the person they are hearing, it is surprising that existing cognitive and neural models of voice processing primarily focus on recognising the unique identity of a familiar person^{1,2}. Within these models, the perception of person characteristics from unfamiliar voices is considered only in terms of “becoming familiar” with a unique identity(e.g. ³) or via listeners being able to discriminate between two voice identities (e.g. ⁴). Indeed, most of the extant theoretical literature considers familiar and unfamiliar voice (identity) processing as being both qualitatively and mechanistically distinct processes⁴⁻⁷.

However, this framing of unfamiliar voice (identity) perception seems to be at odds with our subjective experience of familiar and unfamiliar voices alike as being rich in cues to person characteristics⁸. In this paper, we, therefore, put forward a new and more comprehensive account of person perception from both familiar and unfamiliar voices. We will first briefly review existing models of voice identity perception and discuss how familiar identity perception may indeed be a special case of person perception. We will then outline the evidence for the wide range of person characteristics that can be perceived from (unfamiliar) voices, beyond person identity. We then suggest that outside of experimental contexts, person perception from unfamiliar and familiar voices alike has a common goal of recognising *who* is being heard, but with crucially many more possible perceptual outcomes than recognising a unique individual. To achieve this common goal, we argue that instead of having distinct mechanisms for familiar vs unfamiliar identity perception, person perception from all voices is better characterised via a common mechanism involving the recognition of different person characteristics, be they identity-specific (for familiar voices) or not.

2 Existing models of voice (identity) perception

Over the last two decades, several models of voice (identity) processing have been proposed. These models were created based on evidence from different literatures and as a result, emphasise distinct aspects of voice processing. In the following section, we will outline the most prominent models, sketch out their remit, and show where these models differ from and/or complement each other.

A model of person perception from voices

2.1 A hierarchical model of (familiar) voice processing

Figure 1: Model of voice processing, adapted from Belin et al. (2011)

The arguably best-known and most general model of voice processing has been proposed by Belin and colleagues^{1,9,10} (see Figure 1), which in turn has inspired more recent and similar versions of this model^{3,11}. This model takes Bruce and Young's model of face perception¹² and directly applies it to voices: Following a low-level auditory analysis and a structural analysis of vocal sounds, processing continues along three functionally independent but partially interacting pathways. These pathways support the perception of speech, emotion, and identity information, respectively. Identity recognition is achieved in a "voice recognition unit", which is "activated by one of the voices known to the person" (p.131⁹). Voice recognition units interact with face recognition units (not shown in Figure 1) and once recognition in one or more modalities is achieved, a putative amodal "person identity node" is activated.

2.2 A mechanistic account of familiar voice identity recognition

Prototype-based coding is at the heart of a model of (familiar) voice identity perception by Lavner and colleagues², which provides a mechanistic account of familiar voice recognition. In this model, familiar voices are thought to be represented in terms of their (acoustic) deviation from a voice prototype. This prototype is described as being an average of all the voices that have been encountered by a listener. When a voice is heard, its deviant features in relation to the prototype are computed and the result of this deviant feature extraction is then matched to existing reference patterns or representations of known voices. If the distance between the deviant features of the perceived voice and the reference pattern is sufficiently small, the voice will be recognised as belonging to a particular familiar person. For a visualisation of the processing steps, see Figure 4a, which has been closely adapted from Lavner and colleagues' original model. How person perception from unfamiliar voices may work is not specified within Lavner and colleagues' model: Although their model includes a loop to indicate that perceived voices that do not fit an initial reference pattern are iteratively compared to other reference patterns – presumably until the matching voice is found and the familiar voice is recognised – there is no proposed mechanism for how truly unfamiliar voices are processed or are at least recognised as being unfamiliar.

2.3 Combining hierarchical and mechanistic models

While Belin and colleagues' hierarchical model of voice processing^{9,10} outlines a pathway to voice identity recognition, it does not specify the mechanism by which identity recognition is achieved within its proposed processing hierarchy. However, the framework can readily incorporate the mechanistic model proposed by Lavner and colleagues². These two models

have indeed recently been combined by Maguinness and colleagues³, who further extended this combined model to propose a mechanism for how unfamiliar voices become familiar, thus modelling voice identity learning. Maguinness and colleagues³ propose a hierarchical framework for voice identity perception, in which deviant features are extracted in relation to a prototype upon hearing familiar and unfamiliar voices alike. The model includes very similar hierarchical stages to Belin and colleagues' model^{9,10} and incorporates a prototype-based recognition mechanism as described by Lavner and colleagues². To model voice identity learning, Maguinness and colleagues³ propose that at first exposure to an unfamiliar voice, its pattern of deviant features does not fit any reference pattern for familiar voices and therefore cannot be recognised. Over repeated exposures, however, the pattern of deviant features will eventually become an established stored reference pattern or representation and can thus become recognisable as a familiar voice.

2.4 Drawing a distinction between familiar voice recognition and unfamiliar voice discrimination: A neuropsychological perspective

Kreiman and Sidtis⁵ present a view of voice identity perception that is strongly influenced by evidence from neuropsychological studies. The authors propose that unfamiliar and familiar voice identity perception are two mechanistically distinct processes. They cite neuropsychological studies, in which individuals with brain injuries completed two types of identity perception tasks: 1) Identity recognition, which traditionally requires the listener to link a single presented voice to a specific known identity, and 2) identity discrimination, in which the listener compares two voice samples and determines whether or not they were made by the same person. These studies showed a double dissociation, where individuals who had difficulties in voice identity processing (known as phonagnosia) following a brain injury were able to successfully complete identity discrimination tasks with unfamiliar voices but could not recognise familiar voices, and vice versa^{4,13,14} (see similar reports of dissociations in developmental phonagnosia¹⁵⁻¹⁷). This evidence is compelling and conclusively demonstrates that discrimination is not a prerequisite for recognition⁴. However, as other authors have noted^{3,18} manipulations of voice familiarity and experimental task overlap in these studies. Thus, although discrimination tasks were crucial in allowing neuropsychologists to establish identity perception impairments in phonagnosia, the conflation of task with voice type makes it unclear whether the dissociations reflect fundamental differences in familiar versus unfamiliar voice identity processing *per se*, differences in task demands, or aspects of both. Although the original studies are careful to discuss their results in light of both the experimental task and the voice familiarity, later research has often assumed that familiar and unfamiliar voice processing *per se* is entirely distinct processes without explicitly considering task demands (see Stevenage⁷ for a recent review in support of this assumption).

2.5 The scope and remit of existing models

What emerges from this overview is that existing theoretical models focus on identity perception, specifically *familiar* voice identity recognition^{2,9}. Where the perception of unfamiliar voices is considered^{3,4}, it is likewise always in the context of voice identity perception, proposing, either implicitly or explicitly, that familiar vs unfamiliar voice (identity) perception are distinct from one another. The models thus only consider one case in which listeners are trying to make sense of which specific person they are talking to (i.e., identity-specific perception) and entirely neglect to account for what listeners may do when identity-specific recognition is not possible (or not the focus of attention). That is, even though the perception of other person characteristics from (unfamiliar) voices, such as impressions of age, sex, regional accent, or perceived personality, is as much part of person perception as is identity-specific recognition, these core aspects of person perception are currently missing from theoretical models (but see Bruce and Young¹², who describe "visually derived semantic codes" for unfamiliar face perception).

3 Recognising familiar voices: A special case of person perception from voices

The perception of familiar voice identities has captured the imagination of researchers. An existing literature on face perception has fore-fronted the importance of identity-specific recognition (e.g. ¹²), which may in turn have inspired the particular focus of theoretical and empirical work on voice identity perception. Further, the recognition and perception of (personally) familiar voices have been viewed as being “special” whereas unfamiliar voice identities are thought to hold no such special status and lack (social) importance in perception^{5,6}. Below, we will briefly review some themes emerging from the (familiar) voice identity perception literature. We will also discuss why and how a “special” status has been conferred upon familiar voice perception in the literature.

Familiar voices can be defined as the voices of people whose identities are known to us, for example, because they are famous (e.g., actors, politicians), known to us in our personal lives (e.g. family, friends, colleagues), or because we have been introduced to them as part of an experimental paradigm (e.g. recognition training, passive exposure). Studies have shown that listeners can recognise familiar others from their voices, although the accuracy of identity recognition depends on the selection of voices/stimuli being heard and the characteristics of the listeners, as well as other factors^{5,19}: For example, longer voice recordings have been shown to increase identification accuracy for famous voices^{20,21}. On the other hand, acoustic manipulations (e.g. filtering or shifting of acoustic cues) or intentional voice disguise generally reduce the accuracy of familiar voice identity recognition^{22–26}, and tasks including within-person variability (e.g. a person laughing vs speaking vs shouting) pose perceptual challenges to familiar (and unfamiliar) voice identity perception^(e.g. 18,27). The type and degree of familiarity with a voice also matter: Personally-familiar voices (e.g., romantic partners) are recognised with very high accuracy even in perceptually-challenging tasks, while accuracy for lab-trained voices is substantially lower within those same tasks^{26,28}.

Identity-specific recognition is peculiar to familiar voices: Only a familiar person can be uniquely recognised – by name or otherwise – based on their voice, as “recognition” requires some stored representation of the specific thing that is to be recognised. This is not possible for unfamiliar voices, as a stored person-specific representation cannot exist for someone that we have never encountered before. This possibly sheds light on how the literature has come to see familiar and unfamiliar voice identity perception as distinct (see Section 2.4): It is impossible to run an identity recognition task with previously unheard voices – there is little point in asking a listener whether an unfamiliar voice is Tom or Dave – thus often necessitating the use of an alternative task, such as voice identity discrimination. While this is a practical solution, we note that explicit voice identity discrimination for unfamiliar voices is, in the end, quite an artificial task or listening situation. In this paper, we argue that identity-specific perception is unlikely to be a priority when listening to unfamiliar voices. Identity discrimination tasks, while valid in experimental contexts, thus bear little resemblance to how person-related information is processed from unfamiliar voices in everyday life.

Another unique aspect arising from familiar voice identity perception is that (personally) familiar voices at first glance appear to be much richer, personally-relevant signals than unfamiliar voices^{5,6,29}. In addition to the representation of the sound of a familiar voice, specific memories associated with the person will be accessible to a listener once a voice is recognised. The sound of a familiar voice may also activate representations of the person’s face^{30,31}, alongside biographical knowledge about them (even for people we “know” but have never met, such as celebrities). Similarly, we can access emotionally and socially salient information about, for example, whether we like this person or not, as well as specific memories of events and situations involving them^{5,6}. The emotional salience that may accompany hearing a familiar voice is readily expressed in the statement “it’s so good to hear your voice” and by – albeit anecdotal – accounts of emotional responses to voice recordings of personal significance (e.g. the story of a widow’s fondness for listening to the London Underground announcements recorded by her late husband²⁹). However, unfamiliar voices can still be emotionally evocative as a result of salient personal memories as they may be present for familiar voices: For

example, we may experience an emotional reaction when hearing someone speak in an accent that we have not heard since childhood, or hearing someone who *sounds like* a familiar individual. These situations are, however, likely to be exceptions, and not the rule, for unfamiliar voice perception.

Generally speaking, the experience of recognising a familiar voice is therefore more likely to trigger experiences that are more personally-relevant, specific, and richer compared to the experience of an unfamiliar voice. Thus, there are clearly aspects that are unique to familiar voice processing compared to the perception of unfamiliar voices. While these unique aspects can perhaps go some way to explaining the literature's focus on identity perception from voices, this however does not mean that the person perception from unfamiliar voices is uninteresting and can be neglected - after all, even the most treasured and personally-relevant familiar voices were once entirely unfamiliar to us.

4 Perceiving person characteristics from unfamiliar voices

If a voice is unfamiliar, what then is left for a listener to perceive about its owner in the absence of identity recognition? In this section, we will outline evidence showing that being unfamiliar with a specific talker in no way precludes *person perception* - in fact, listeners can and do routinely form rich and complex impressions of a person without being familiar with them or knowing their unique identity (see Figure 2).

Figure 2: Word cloud illustrating listeners' descriptions of six unfamiliar voices (see Lavan, 2023). The talkers were all female, speaking with a Standard Southern British accent and were reading out linguistically-neutral sentences for 2 minutes. Font size scales with the frequency with which individual descriptors were mentioned.

Empirical evidence from a range of fields, including evolutionary and social psychology, sociophonetics, and bioacoustics, demonstrates that human listeners will readily provide judgements of many person characteristics from an unfamiliar voice. Listeners perceive physical characteristics, such as the sex of a person^{32–34}, the person's age^{32–35}, aspects of body size such as height and weight^{36,37}, and physical strength^{38–40}. Listeners furthermore readily make judgements about other (perceived) social or psychological characteristics based on hearing a voice. These include a person's spoken accent⁴¹, native speaker status⁴², sexual orientation^{43,44}, social status⁴⁵, and occupation⁴⁶ (see also Kreiman & Sidtis⁵ for a general overview). A recent focus of interest has been on “first impressions” of personality traits (e.g.,

trustworthiness, competence, attractiveness). This work has, to some extent, shadowed prior research on face and social category perception^{47,48}. Research obtaining evaluations of speaker traits from voices - typically via rating scales - has found that first impressions are formed rapidly⁴⁹⁻⁵¹, are typically made with high inter-rater reliability for most trait categories, can be mapped onto a small number of underlying dimensions (e.g., warmth/trustworthiness vs dominance^{49,52}), and are consistent across different types of vocal stimuli (e.g. different languages and linguistic content⁵³⁻⁵⁵). While there is overall little empirical evidence that first impressions of a person are linked to their actual personality or character traits⁵⁶ (see also Box 1), these subjective impressions can nonetheless influence our daily interactions, behaviours and decisions (e.g.^{57,58}).

The studies above examine the perception of a wide range of person characteristics. Yet, there is considerable heterogeneity in both the theoretical motivations and the methodological approaches across disciplines. For example, while researchers in bioacoustics are often concerned with how actual physical variation in bodies gives rise to differences in voice acoustics, and evolutionary psychologists are interested in how these might signal qualities such as physical or reproductive fitness to potential adversaries or mates, social and experimental psychologists have been more focused in how perceived personality traits in the voice might relate to each other and a range of social outcomes for the speaker (e.g. employment, election, financial investment). These differences in research questions mean that different types of person characteristics - for example, height versus trustworthiness - are often studied completely separately, using different tasks and judgements (e.g., various types of scales vs categorical judgements vs judgements relative to the listener's own characteristics), and with varying focus on the "ground truth" of perception (see Box 1). So, although this body of work illustrates the potential richness of person percepts from voices, it is difficult to gain a full appreciation of how listeners form (complex) impressions of unfamiliar others from voices based on any individual study or field. Similarly, it remains difficult to gauge the relative salience and importance of these different perceived characteristics in relation to one another.

Recent studies in experimental psychology have taken a broader view on person perception from (unfamiliar) voices, attempting to examine which characteristics are spontaneously perceived from unfamiliar voices, and how multiple person percepts might be organised in relation to one another. Lavan⁸ asked participants to listen to a recording of a voice and provide a list of words that described the person they heard, thus allowing the listeners themselves to generate the characteristics rather than making judgements on characteristics predetermined by the experimenter (see also Pear⁴⁶). Listeners provided a wide range of descriptions based on the voices, including descriptions of physical, psychological, and social characteristics (see Figure 2 for a word cloud showing a summary of these free descriptions). While likely reflecting post-perceptual decisions and biases, in addition to perceptual processes, these data nonetheless provide evidence that listeners do form complex impressions of other people based on their voices. Additionally, the study reports that there was a structure to listeners' responses - listeners tended to first mention physical characteristics, like sex and age. Further recent evidence for such structured person perception has come from perceptual gating-type studies, in which listeners were exposed to brief voice clips of increasing exposure duration (e.g., from 25ms to 800ms) and asked at each duration to provide ratings of a range of physical, social, and psychological characteristics. Inter-rater agreement at each exposure duration was used as an index of impression formation, where high agreement indicates an impression has been formed, and no/low agreement indicates that no impression has been formed yet (i.e., responses are random). Despite using a very different approach to a free text description task (cf. Lavan⁸), these results also indicate a temporal hierarchy of person percepts, where impressions of physical characteristics (age, sex) are formed more quickly than impressions of social and trait characteristics^{50,51} (see also ⁵⁹ for faces).

The review above clearly shows that listeners can perceive a wealth of person-related information from unfamiliar voices. Indeed, the main impression listeners take away from

encountering an unfamiliar voice is unlikely to be limited to identity-specific information (e.g., “I did not know this person”; “Person X sounds different from Person Y in this conversation”) as current models might perhaps suggest. What is much more likely is that the listener comes away with a (possibly hierarchically-organised) combination of first impressions covering aspects of physicality, social characteristics and other traits (e.g. “This woman sounded really friendly and clever”). Taken together, the findings from unfamiliar voice perception thus strongly motivate the need for existing models of voice perception to be updated to include person perception *beyond identity recognition*.

Box 1: The role of ground truths in person perception

Much of the research on person perception from voices has addressed how accurately listeners can perceive identity- and person-related information. For identity-specific judgements, such as speaker discrimination or recognition, accuracy is an intuitive measure - the listener instinctively understands what it means to be “right”, and the researcher can easily verify their responses based on knowing which identities have been presented. Similarly, a range of other person characteristics - for example, height in cm or age in years - can be objectively compared with listener responses. However, often there are considerations around what constitutes a “correct” response. For example, when estimating a speaker’s height, is an “incorrect” estimate of 5cm taller/smaller than the ground truth considered as good accuracy or not? Thus, even when objective measures of characteristics are available, their interpretation is often not straightforward (see also discussions around the accuracy of perception for age or physical strength^{36,40}).

In other cases, a “ground truth” is much more difficult to access or obtain. For example, in the trait perception literature, measures of the talkers’ personality traits are in many cases not readily available. For this reason, research questions on the perception of trait characteristics are mostly focused on (changes in) mean ratings, independently of accuracy, as well as on the inter- and sometimes intra-rater agreement of trait percepts across listeners (e.g.^{50,51,53}). In the absence of an accuracy measure, agreement here is typically used to indicate 1) that responses are not random and 2) that shared representations underpin listener judgements. There are other situations in which accuracy can be assessed but it is low or at chance. Here, the inspection of the nature and patterning of listener responses can still provide valuable additional insights (e.g.⁶⁰).

Overall, while accuracy can be an interesting and informative measure, it is often only one way of looking at the data. Arguably, however, one of the primary outcomes of person perception from a voice is the individual listener’s impression of the characteristics of that person, be that impressions accurate or not. We would therefore argue that there is scope to ask new questions about person perception that have to date been largely unexplored, by going beyond assessments of the accuracy of perceiving certain person characteristics to examine the frequency and content of listener responses - for example, by (additionally) deriving measures of agreement or response bias that can be interpreted independent of their relationship to a ground truth.

5 A model for person perception from voices (PPV)

Models of voice perception need to account for the goal of *person perception* from familiar and unfamiliar voices alike, instead of framing the success or failure of the perception of specific identities as the primary and perhaps only objective. Outside of experimental tasks, listeners’ shared aim is to perceive *who a person is* at a holistic level, which is not restricted to identity-specific perception, but is instead based on whatever information about a person is available or meaningful to them. Assuming this common goal of person perception from all voices, we argue that person perception from familiar and unfamiliar voices alike can be achieved via a

common mechanism. We propose that a mechanism of *recognition* underpins all aspects of person perception from voices. This includes the recognition of specific known identities, but also extends to recognising any number of broader person characteristics.

As it stands, the voice identity perception literature mainly operationalises “recognition” as a process during which a vocal signal is attributed to a specific person, thus primarily considering recognition in relation to familiar voices and identity-specific perception (cf. Section 2). When listeners recognise a familiar person, they can find themselves being able to name the person (“It’s Aisha”) or identify them in the absence of naming (“It’s the woman with brown hair who works in the café across the road”, “It’s the actor from those superhero movies”). In contrast, (accurate) identity recognition from an unfamiliar voice is impossible. However, upon hearing an unfamiliar voice that is fairly high-pitched, likely youthful, and likely female with a Californian accent and a lot of vocal fry, listeners can readily *recognise* that they are listening to a young woman as opposed to a child. Depending on prior experience, listeners may even recognise that the voice described above fits the *persona* or social stereotype of “Valley Girl”⁶¹. The impression of an unfamiliar person from a voice - be that based on a few broad demographic characteristics or a well-formed persona - provides meaningful information about that person and arises in the absence of listeners being familiar with the person’s unique individual identity. Crucially, this impression can be entirely subjective and does not have to be accurate to exist (see Box 1).

Here, we reframe person perception as a process that has a common perceptual goal and which is achieved via a common mechanism of recognition of person-specific, persona-specific, and/or person-general (e.g., sex, age, etc.) characteristics. Informed by this view, we outline a model that for the first time provides an account of how meaningful perceptual outcomes for person perception from unfamiliar voices can arise, beyond identity-specific perception. Importantly, this model also eliminates the need for two mechanistically entirely distinct processes for perceiving information about people and their specific identity from familiar vs. unfamiliar voices (e.g., recognition vs discrimination^{4,5,7}). We will first describe below how these multiple possible recognition processes may unfold and interact within a processing hierarchy, followed by a description of potential common mechanisms that could underpin the recognition processes.

5.1 The PPV (person perception from voices) model: Broadening existing models

Figure 3: A hierarchical model of Person Perception from Voices (PPV, see ^{10,12}). The figure shows an illustration of how different percepts may interact and inform each other. Boxes with a solid outline describe a recognised percept, the box with a dashed outline describes a computational process. Brackets at the bottom of the figure loosely map the processing stages

of the PPV to Belin et al.'s hierarchical model of voice identity perception¹⁰. Not all possible interactions between percepts are shown. For a detailed explanation of the model, please see Section 5 in the main text.

The PPV model (Figure 3) aims to describe how listeners may recognise one or more aspects of a person or a persona to provide a more comprehensive account of listening situations that result in person perception from vocal cues. The PPV model builds on previous model (See Section 2) and is therefore at its core a hierarchical framework (see^{10,12}), within which a common recognition mechanism operates (see Section 5.2) to account for the common goal of person perception from voices. In our model, following the extraction of person-relevant features from the acoustic voice signal, listeners recognise any number of person characteristics based on these features. In line with previous models, the PPV model indicates a possible route to account for recognition of familiar persons or personae (see below) and adds to this the capacity for recognition of one or more individual person characteristics (sex, age, health, etc.). Although visually depicted as two routes, these processes are not mutually exclusive to one another: While person-specific recognition may be prioritised depending on the listening situation, individual person characteristics may be recognised in parallel with recognising a familiar person/persona (see below for further considerations).

We first consider a direct route for the identification of a familiar person or persona. This part of the model is closely aligned with previous theoretical accounts of identity perception^{3,5,9,11}. Specifically, for this direct route, following the extraction of person-relevant features from a heard voice, this voice is recognised in a person-specific 'recognition unit' as a known person without there being any intermediate processing steps (see^{3,5,9,11}. Here, we extend this mechanism beyond identity-specific recognition to additionally support the recognition of known *personae* - these include familiar social stereotypes, such as a Valley Girl or City Banker, where the vocal signal yields a coherent and detailed percept of a known type of person without the identification of a unique individual.

We next consider the situation in which the listener hears an unfamiliar voice, from which a specific person or a persona cannot be recognised. When perceiving an unfamiliar voice, listeners recognise one or more physical, psychological, or/and social attributes of a voice. As indicated by the grey box in Figure 3, the recognition of several person characteristics may give rise to a percept of an unfamiliar person comprising the recognised characteristics (e.g., male, middle-aged, tall, friendly, German). However, even when the listening situation makes only one or two characteristics available to perception - for example, due to brevity or poor signal quality - the PPV model still allows for these to be recognised as individual percepts.

In existing models for the recognition of specific identities, there is only one possible match between a given voice and a stored representation of a specific person's identity, which represented by the direct route to the perception of a known person or persona in the PPV model. The PPV model proposes additionally the recognition of multiple person characteristics at the same time or closely staggered in time. This may happen on a regular basis when we encounter unfamiliar voices and derive complex, multi-attribute impressions from these voices. Emerging evidence from voice research suggests that this perception of different person characteristics from unfamiliar voices may indeed be structured. For example, Lavan⁵⁰ reports that impressions of physical characteristics, such as age, sex, and health, are formed more quickly than trait or social characteristics, such as level of education and trustworthiness. Based on this evidence, we tentatively place characteristics like sex and age "earlier" in the processing chain compared to trait and social percepts in Figure 3. We furthermore propose possible interactions between characteristics, which can reflect existing patterns of co-occurrence in the population (e.g., sex and height recognition: female talkers are usually smaller than male talkers). Therefore, our model allows for (some) bidirectional effects of facilitation (or inhibition) of any of these characteristics on one other (e.g., a percept of "female" may facilitate a percept of "smaller than average height"). Further investigation will need to test

whether such interaction of representations primarily happens within the recognition process, for example through direct mutual co-activation of stored representations, or via yet another route implicating additional systems (e.g., top-down effects of attention or memory).

Box 2: Structuring a perceptual hierarchy: Bottom-up, top-down, and post-perceptual influences

The PPV model's proposal for how people are perceived from voices differs from existing accounts in allowing for recognition and interaction of multiple person-related characteristics as opposed to the recognition (or not) of a unique familiar identity. The multi-attribute nature of the model warrants some further consideration of how multiple perceptual representations might be organised and prioritised in perception, both via bottom-up, top-down influences and post-perceptual.

A hierarchy of the recognition of person characteristics could be shaped by bottom-up perception of acoustic cues that carry person-relevant information. How quickly and how clearly these percepts might arise depends on the accessibility of the cues over time. For example, pitch cues to an adult's sex can be perceived within a few glottal cycles⁶², while information about a vocal tremor indexing old age or poor health requires longer exposure. This example would predict the faster recognition of sex over age/health perception. Additionally, some characteristics may be more readily discriminable than others based on their distinguishing acoustic cues - for example, voice pitch in humans is a highly salient and reliable cue to differences between adult male and female voices, while it is less reliable for children^{63,64}. Given these considerations, it may be that recognition of certain person characteristics can happen earlier than others that are less clearly marked (e.g., sex perception from adult voices being quicker than sex perception from child voices⁶³). A question that arises from such considerations is whether the bottom-up availability of certain cues makes recognition of some characteristics obligatory *en route* to less accessible or more specific percepts. For example, must a listener recognise that a voice is female in order to recognise that it belongs to their sister? Can rapid recognition of a highly familiar identity override the recognition of some of their physical, psychological, or social characteristics via bottom-up mechanisms?

Top-down effects originating from the wider cognitive system could also have impacts on the recognition process. For example, task demands requiring evaluation of a personality trait might lead to down-weighting of some characteristics in favour of others: e.g. participants do not reliably recognise their own voice within a larger set of voices when being asked to make attractiveness judgements⁶⁵. Similarly, general priors and context-specific expectations may affect which representations are prioritised (e.g. failing to recognise a familiar work colleague's voice when heard unexpectedly in a holiday setting). There is also potential for higher-order, post-perceptual effects to influence or override the outcomes of the recognition process. For example, behavioural research has shown that bias in the association of lower voice pitch with larger body size mediates the apparent relationship between pitch and rated dominance for male voices, suggesting that listeners use perceived body size as a heuristic for judgements of formidability⁶⁶. Thus, while a basic structure of person perception from voices may exist, determined by bottom-up availability and accessibility of acoustic cues, there must be scope within a processing hierarchy to dynamically re-structure the processing chain in response to highly variable listening situations and perceptual priorities.

Our model also includes the potential for bidirectional interactions between the recognition of a person or persona and the recognition of person characteristics. The potential for the recognition of person characteristics to influence the recognition of a known person or persona (upward arrow toward "Familiar Person(a)" in Figure 3) can describe situations in which an objectively familiar voice is not immediately recognised as such, but where recognition of the

specific identity emerges after additional information about the person has been accumulated from the voice signal in a largely bottom-up fashion (e.g., when hearing a low/moderate-familiarity voice). On the other hand, potential effects of a recognised familiar person or persona on the recognition of person characteristics (downward arrow from Familiar Person(a) in Figure 3) are perhaps less intuitive. These can however account for previous reports that familiarity with a specific person might affect the neural decoding of person sex and age from their face/voice⁵⁹, and evidence that familiarity yields faster judgements of sex from voices⁶⁷.

It is possible that an “indirect route” to (familiar) person or persona recognition via the prior recognition of other person characteristics may in fact be the typical means of recognising a familiar person or persona. We highlighted above that there may be a temporal or perceptual hierarchy to how different person characteristics are recognised: By extension, it may also be the case that, for example, sex recognition is faster than identity-specific recognition^{50,59,62}, such that listeners may routinely recognise some person-general characteristics before recognising the specific identity or persona. Despite the possible interactions, the PPV model still allows for direct person/persona recognition via matching to a single stored identity representation as outlined above (adapted from Belin and colleagues⁹, Maguinness and colleagues³), where prior or parallel processing of other person characteristics is not obligatory. Future work will therefore need to determine to what extent such a hierarchy may be obligatory, what factors determine how specific people or personae are recognised, as well as what the possible functional benefits of an indirect route to identity recognition may be (see Box 2 for further considerations regarding the structure of a perceptual hierarchy).

Finally, at the top of the PPV hierarchy is the “wider cognitive system”, indicated by the cloud in Figure 3. We borrow this component from Bruce and Young’s model of face processing (in which the authors note: “The ‘cognitive system’, by convention, is somewhat cloudy.”, p. 311¹²). We model communication from the person/persona recognition machinery to the wider cognitive system, to support the storage and/or retrieval of person-relevant information such as biographical knowledge about a person, episodic memories of social interactions with them, and emotional reactions to them (see also Kreiman and Sidtis’ discussion of wider systems engaged when listening to familiar voices⁵). We furthermore also allow for effects of the wider cognitive system upon the recognition of person characteristics. These could, for example, include top-down influences of expectations or attentional modulations on person perception (see Box 2).

5.2 A common recognition mechanism for person perception from voices

In the sections above, we propose that all voice signals are processed using a common mechanism to achieve the recognition of a person, persona, and/or any person characteristic(s). There are various ways in which such a common recognition mechanism could be structured. Based on previous proposals from the voice (and face) perception literature that were heavily inspired by prototype-based accounts (as opposed to episodic-/exemplar-based accounts^{68–71}), Figure 4 illustrates two possible versions of this common mechanism.

The mechanism outlined in Figure 4a closely follows the existing prototype-based account of voice (identity) recognition by Lavner and colleagues² and Maguinness and colleagues³ (see Sections 2.2 and 2.3). In this account, voices are first encoded and perceptually analysed based on their acoustic features (see also ‘voice structural analysis’¹⁰). Following this analysis, voices are then compared to a stored voice prototype, which is based on an average(d) voice derived from a person’s accumulated experience with human voices. The outcome of this process is a pattern of deviant features, which describes how the heard voice differs from the voice prototype. This pattern of deviant features is then compared to stored reference patterns (representations): According to Lavner and colleagues² and Maguinness and colleagues³, like the pattern of deviant features computed for heard voices, these stored reference patterns (representations) also consist of deviant feature patterns of how known identities differ from

A model of person perception from voices

the voice prototype. While previous prototype models of voice identity recognition only consider comparisons to stored identity-specific reference patterns (representations) in order to model familiar voice identity recognition, we additionally propose that the same mechanism can be applied to stored reference patterns for specific known persons, personae (e.g., “Valley Girl”), and/or person characteristics (e.g., physical, social, psychological, see Figure 4a). In this comparison process, a distance (d) between the deviant feature pattern (of a heard voice) and a stored reference pattern (representing a person, persona or a person characteristic) is computed. If this distance (d) is small enough and thus falls within a recognition threshold, recognition of that person, persona or person characteristic takes place. If the distance (d) exceeds the threshold, the pattern of deviant features for the heard voice is iteratively compared to other relevant stored reference patterns (e.g., another identity or another sex/gender).

a) Recognition mechanism based on a prototype model of person perception (adapted from Lavner et al. 2001)

b) Alternative recognition mechanisms without prototype-based comparisons

Figure 4: Illustration of how recognition of a person, persona or person characteristic may be mechanistically achieved. a) A mechanistic account of the common recognition mechanism adapted from Lavner and colleagues². b) An alternative, simplified account of recognition without compulsory prototype-based coding of deviance. Across panels a) and b), each box with a solid outline describes a recognised percept, each box with a dashed outline describes a computational process, and each box with a dotted line describes a type of mental representation supporting recognition. The box “recognised person characteristic or person(a)” directly maps onto each of the boxes labelled as percepts (solid outline) in Figure 3.

The section above highlights that the extant prototype-based model of voice identity perception by Lavner and colleagues² can be readily extended to account for person perception more broadly. However, Figure 4a depicts only one possible instantiation of a prototype-based mechanism, based on comparisons of incoming signals with a *single* voice prototype. Existing prototype-based accounts for identity recognition propose a variation of the model illustrated above where there could be multiple, functionally-relevant voice prototypes (e.g., sex-specific prototypes^{3,72}). For a model to incorporate multiple *prototypes*, we would assume that each prototype would code for an average of a characteristic (identity, age, height, etc.), with a relevant prototype being selected ad-hoc depending on the listening situation. We note, however, that the assumption of multiple prototypes – already within the original model that only considers identity perception – brings with it a substantial increase of complexity within the mechanism. Similarly, there are questions about how, and perhaps more importantly, why any prototype would qualitatively and/or functionally differ from any representation, an issue that is exacerbated when considering multiple, potentially co-dependent recognition processes as the PPV model does, therefore requiring a large number of prototypes. We expand on these existing issues within prototype-based model of voice in Box 3.

Box 3: The origins of a functional role for prototype(s) in models of voice recognition

Many variations of prototype-based models (as opposed to e.g. exemplar-based models) have been proposed in different areas of psychology to describe the architecture of mental representations^{73–76}. Prototype-based models of voice (and face) identity recognition were originally invoked to account for empirical findings showing that highly distinctive voices and faces are recognised more quickly and with higher accuracy^{2,74,77}. By basing identity recognition on deviant features relative to a prototype, the magnitude or uniqueness of the deviance pattern increases with increasing distinctiveness. This could in turn explain better identity recognition performance for more deviant identities, thus ushering in a functional role for prototypes during identity recognition.

While this is a plausible account of the empirical findings in relation to distinctive identity recognition, there may be alternative explanations. For example, voices with properties far from the prototype may also lie in less populated parts of a notional voice space⁷⁸ where there are fewer competitors^{70,79}. Interestingly, and contrary to the finding for identity recognition, several classic studies of category abstraction had also found that distinctive category members were *less* well recognised in classification and recall tasks (see ⁷⁷ for details). This alternative prediction also rings true for the recognition of other person characteristics: A voice that is very deviant from a prototypical “female voice” would likely be less well recognised as a woman (but instead might be recognised as child or an adult of another gender). That is, distance from a prototype *per se* does not intuitively account for better recognition of any characteristic (or category) beyond perhaps a unique identity. There are thus two possibilities: On the one hand, the mechanisms for recognising specific identities could truly be qualitatively different from the perception of other types of characteristics (and categories), warranting prototype-based coding for one but not the other. If this were the case, what would be the reasons for the unique status of identity recognition? Alternatively, prototype-based coding of deviance may be functionally not obligatory for either identity- or person perception (see Figure 4b), and empirical findings underpinning prototype models for identity recognition can be explained otherwise.

If there is a – perhaps limited – functional role for prototype-based coding, prototypes and representations are proposed to be functionally and qualitatively distinct: The prototype is an average voice, while representations are patterns of deviance in relation to this prototype. When outlining mechanisms of identity recognition, some authors, however, also mention the use of sex-specific female and male prototypes – others even propose “multiple” prototypes the nature of which remains otherwise unspecified^{2,3,72}. We note there that inserting additional prototypes into the deviance-based recognition mechanism raises questions about how many, and in what order, different prototypes might be engaged in calculating deviance patterns (e.g., does speaker sex need to be first recognised in order to select the “female voice” prototype to

support recognition of other characteristics)? It is furthermore unclear what the outputs of multiple prototype comparisons might be (e.g. sequential prototype comparisons creating deviances of deviances, vs. individual deviance patterns for all invoked prototypes?), as well as what these would imply for the number and nature of required reference patterns used for recognition.

Given that there are questions around whether prototype-based coding of deviance is required in the first place^{70,79} and what the nature and number of functionally-relevant prototypes may be, it is worth considering whether recognition could be achieved without obligatory prototype-based coding. Please see Figure 4b and Section 5.2 for such a simplified mechanism.

Given the conceptual complexities that arise from obligatory prototype-based coding (be there one or multiple prototypes), we propose in Figure 4b a simplified common mechanism for recognising a familiar person, persona and/or different person characteristics. Specifically, the mechanism in Figure 4b removes voice prototypes as a functional component from the recognition process, such that the extracted person-relevant acoustic features of perceived voices can be directly compared to relevant stored representations. This alternative mechanism thus fully preserves one of the core ideas behind the original recognition mechanism (matching of relevant features of a heard voice to stored representations), but removes the original two-step mechanism introduced by deviance coding (deviant feature selection + comparison of deviant features to stored representation). By removing the functional role of a prototype, the distinction between (average-based) prototypes and (deviance-based) representations becomes unnecessary. Further, in the absence of a functionally-relevant prototype, the content of stored representations must change from being stored patterns of deviance (in relation to a prototype) to retaining and encoding all the person-relevant features. Whether this simplified approach is more reflective of the cognitive processes for perceiving people from voices than the prototype-based model remains to be tested.

We also stress at this point that we do not wish to claim that prototypes (or at least representations of what a central tendency of a category would be) do not exist: Listeners can have notions of what a prototypical human voice (vs. non-human voices), a prototypical female voice (vs a male voice) or even a notion of what sounds prototypical for a specific voice identity (vs another identity) sound like and prototypes may structure our experience of different types of voices. However, like some previous work⁷³, we raise questions about whether prototypes should 1) be included as functionally-obligatory components of recognition mechanisms and whether this inclusion should as a result 2) be able to define the nature and content of representations by mandating purely deviance-based coding and storage of information for representations^{2,3}.

Beyond what the different functional components of a common recognition mechanisms may be, there are some additional assumptions about the nature and content of representations within an expanded model of person perception. For example, due to originally only considering identity-specific recognition, existing models (whether prototype-based or not) assume that all representations (or reference patterns) are categorical in nature (Person X vs Person Y). This presents a challenge for the recognition of person characteristics that may appear continuous in nature, such as age or height. We discuss this and related questions in Box 4.

Box 4: The nature of voice-based representations – recognising categorical vs continuous characteristics

Representations are often central to mechanistic accounts of identity recognition. Such accounts implicitly assume that representations are essentially categorical - the stored reference pattern for Person X is distinct from the reference pattern for Person Y (although they can be mapped onto a common acoustic voice space^{72,78}). When thinking of recognition beyond identity, some questions arise about the nature of representations: How can categorical representations underpin the perception of dimensional characteristics like height, age, and perceived personality traits? For example, we could assume that categorical representations of continuous characteristics would need to encapsulate the dimensional nature of the characteristics in question. Listeners may, for example, have a representation of a person of average height, where the representation also encodes how height varies in the population. However, such an account cannot readily be incorporated into the mechanism outlined by prototype-based models of person recognition. Alternatively, listeners might need to access separate representations for a “tall” voice and a “small” voice. In a height estimation task, the listener’s response would then reflect a match with “tall”, a match with “small”, or no match (i.e. response around the known average height). Outside of experimental tasks, listeners may indeed only recognise cues to height if a voice sounds particularly tall or small to them. Cues for people of average height are not salient and may thus not be recognised (but perhaps inferred as ‘normal’).

Other considerations arise when trying to think about representations of characteristics that include possible sub-categories: For identity perception, the content of the representation may initially appear (reasonably) clear - it will encode some information about what a particular voice sounds like, be that based on deviance patterns or not. But what is the nature of representations for, e.g., accent perception? Accents can be perceived at a coarse level (e.g. ‘American’ vs ‘Australian’) or at a much more fine-grained level (‘American’ → ‘New Yorker’ → “person from Queens”). Do we have entirely independent representations for all the available levels of specificity, or are representations partially co-dependent? The same issues can of course rise for an identity: Do we have separate representations for what a familiar identity sounds like when they are very tired compared to when they are not tired (see also Lavan and colleagues⁸⁰ showing how notions of prototypicality - without necessarily having a functional role - may apply to identity recognition in the context of within-person vocal variability)?

These are just two examples of how the nature of mental representations is still relatively poorly understood for voice perception. Recent studies have shown that neuroimaging approaches may be able to shed light on where and how representations of individual identities (for voices⁸¹; for faces⁸²) and person characteristics (for faces⁸³) may be stored, as well as the time courses of processing these representations (for faces⁵⁹). Reverse correlation approaches have also been used to describe what representations of smiled speech⁸⁴ and trustworthy or dominant intonation patterns may sound like⁸⁵ (though see Knight and colleagues⁸⁶). It is a challenge for future research to uncover more of the nature and architecture of the representations that underpin person perception from voices and beyond.

6 Conclusions

The PPV model aims to provide a more comprehensive account of how listeners make sense of the person they are listening to, using an approach that incorporates and builds on aspects of the hierarchical frameworks and prototype-based mechanisms proposed within existing models of voice identity recognition. While the PPV model is more comprehensive than existing accounts of voice (identity) perception, its remit nonetheless has clear limits. For example, the current model has been proposed strictly within the framework of person perception from voices. However, we recognise that human person perception includes other sources and modalities of information, notably from faces and bodies. Other models of voice perception have explicitly depicted parallel/mirrored processing streams for face perception and allowed

for interaction between these face and voice processing systems^{9,11}. We too acknowledge this possibility and have mainly for ease of visualisation and interpretation restricted the present depiction and discussion of our model to voices only. With face processing models also being primarily focussed on familiar face identity recognition^{11,12}, we note that a version of the PPV model could also be applied to faces, extending current models to person perception in a broader, multimodal sense.

Similarly, other types of information, related to speech and affective states are encoded in voices and their perception has already been modelled in other accounts (e.g.^{9,11}). As acknowledged in those models, these parallel sources of information are likely to intersect with person perception from voices: A person might have a voice quality that does not immediately suggest a dominant personality, but the words they use might quickly compel a listener to think otherwise. These potential influences on person perception are currently not accounted for in the PPV model. This is not to suggest that we do not consider these intersections and interactions to be unimportant. Indeed, establishing the extent to which these processes are mutually dependent - in function, time, space (e.g., neurobiological location), and mechanism - will be crucial⁸⁷⁻⁹⁰.

With a view on future work, we hope that the PPV model presents the opportunity for greater inter-disciplinary synergy in testing its proposals, by taking a broader perspective on person perception from voices and thus capturing research questions from, to date, disparate empirical traditions. There is much work to be done to, for example, explore which and how different person characteristics are represented, which representations are processed and prioritised under which circumstances, how the perception of different person characteristics may interact and inform each other, and how precisely recognition is, in the end, mechanistically achieved for a range of different person characteristics. To answer these and similar research questions, studies will need to examine how the perception of multiple characteristics unfolds and is affected by experimental manipulations and situational contexts. We intend that the PPV model will provide a conceptual or theoretical backdrop to inform such studies.

References

1. Campanella, S. & Belin, P. Integrating face and voice in person perception. *Trends Cogn Sci* **11**, 535–543 (2007).
2. Lavner, Y., Rosenhouse, J. & Gath, I. The Prototype Model in Speaker Identification by Human Listeners. *International Journal of Speech Technology* **4**, 63–74 (2001).
3. Maguinness, C., Roswadowitz, C. & von Kriegstein, K. Understanding the mechanisms of familiar voice-identity recognition in the human brain. *Neuropsychologia* **116**, 179–193 (2018).
4. Van Lancker, D. & Kreiman, J. Voice discrimination and recognition are separate abilities. *Neuropsychologia* **25**, 829–834 (1987).
5. Kreiman, J. & Sidtis, D. *Foundations of Voice Studies*. (John Wiley & Sons, Ltd, 2011). doi:10.1002/9781444395068.
6. Sidtis, D. & Kreiman, J. In the beginning was the familiar voice: Personally familiar voices in the evolutionary and contemporary biology of communication. *Integrative Psychological & Behavioral Science* **46**, 146–159 (2012).
7. Stevenage, S. V. Drawing a distinction between familiar and unfamiliar voice processing: A review of neuropsychological, clinical and empirical findings. *Neuropsychologia* **116**, 162–178 (2018).
8. Lavan, N. How do we describe other people from voices and faces? *Cognition* **230**, 105253 (2023).
9. Belin, P., Bestelmeyer, P. E. G., Latinus, M. & Watson, R. Understanding voice perception. *Br J Psychol* **102**, 711–725 (2011).

10. Belin, P., Fecteau, S. & Bédard, C. Thinking the voice: neural correlates of voice perception. *Trends Cogn Sci* **8**, 129–135 (2004).
11. Young, A. W., Frühholz, S. & Schweinberger, S. R. Face and Voice Perception: Understanding Commonalities and Differences. *Trends Cogn Sci* **24**, 398–410 (2020).
12. Bruce, V. & Young, A. Understanding face recognition. *British Journal of Psychology* **77**, 305–327 (1986).
13. Van Lancker, D. R., Kreiman, J. & Cummings, J. Voice perception deficits: neuroanatomical correlates of phonagnosia. *J Clin Exp Neuropsychol* **11**, 665–674 (1989).
14. Van Lancker, D. R., Cummings, J. L., Kreiman, J. & Dobkin, B. H. Phonagnosia: a dissociation between familiar and unfamiliar voices. *Cortex* **24**, 195–209 (1988).
15. Garrido, L. *et al.* Developmental phonagnosia: a selective deficit of vocal identity recognition. *Neuropsychologia* **47**, 123–131 (2009).
16. Roswadowitz, C. *et al.* Two Cases of Selective Developmental Voice-Recognition Impairments. *Current Biology* **24**, 2348–2353 (2014).
17. Roswadowitz, C., Schelinski, S. & von Kriegstein, K. Developmental phonagnosia: linking neural mechanisms with the behavioural phenotype. *NeuroImage* **155**, 97–112 (2017).
18. Lavan, N., Scott, S. K. & McGettigan, C. Impaired generalization of speaker identity in the perception of familiar and unfamiliar voices. *Journal of Experimental Psychology: General* **145**, 1604–1614 (2016).
19. Mathias, S. R. & von Kriegstein, K. How do we recognise who is speaking? *Front Biosci (Schol Ed)* **6**, 92–109 (2014).
20. Bricker, P. D. & Pruzansky, S. Effects of stimulus content and duration on talker identification. *J Acoust Soc Am* **40**, 1441–1449 (1966).
21. Schweinberger, S. R., Herholz, A. & Sommer, W. Recognizing famous voices: Influence of stimulus duration and different types of retrieval cues. *Journal of Speech, Language, and Hearing Research* **40**, 453–463 (1997).
22. Lavner, Y., Gath, I. & Rosenhouse, J. The effects of acoustic modifications on the identification of familiar voices speaking isolated vowels. *Speech Communication* **30**, 9–26 (2000).
23. Schmidt-Nielsen, A. & Stern, K. R. Identification of known voices as a function of familiarity and narrow-band coding. *The Journal of the Acoustical Society of America* **77**, 658–663 (1985).
24. Abberton, E. & Fourcin, A. J. Intonation and speaker identification. *Language and Speech* **21**, 305–318 (1978).
25. Reich, A. R. & Duke, J. E. Effects of selected vocal disguises upon speaker identification by listening. *J Acoust Soc Am* **66**, 1023–1028 (1979).
26. Hollien, H., Majewski, W. & Doherty, E. T. Perceptual identification of voices under normal, stress and disguise speaking conditions. *Journal of Phonetics* **10**, 139–148 (1982).
27. Lavan, N., Burston, L. F. K. & Garrido, L. How many voices did you hear? Natural variability disrupts identity perception from unfamiliar voices. *Br J Psychol* **110**, 576–593 (2019).
28. Kanber, E., Lavan, N. & McGettigan, C. Highly accurate and robust identity perception from personally familiar voices. *J Exp Psychol Gen* **151**, 897–911 (2022).
29. McGettigan, C. The social life of voices: studying the neural bases for the expression and perception of the self and others during spoken communication. *Frontiers in Human Neuroscience* **9**, (2015).
30. von Kriegstein, K., Kleinschmidt, A., Sterzer, P. & Giraud, A.-L. Interaction of face and voice areas during speaker recognition. *J Cogn Neurosci* **17**, 367–376 (2005).
31. von Kriegstein, K. *et al.* Simulation of talking faces in the human brain improves auditory speech recognition. *Proceedings of the National Academy of Sciences* **105**, 6747–6752 (2008).

32. Smith, D. R. R. & Patterson, R. D. The interaction of glottal-pulse rate and vocal-tract length in judgements of speaker size, sex, and age. *J Acoust Soc Am* **118**, 3177–3186 (2005).
33. Lavan, N. The effect of familiarity on within-person age judgements from voices. *British Journal of Psychology* **113**, 287–299 (2022).
34. Braun, A. Age estimation by different listener groups. *International Journal of Speech, Language and the Law* **3**, 65–73 (1996).
35. Moyse, E. Age Estimation from Faces and Voices: A Review. *Psychologica Belgica* **54**, 255–265 (2014).
36. Van Dommelen, W. A. Speaker height and weight identification: A re-evaluation of some old data. *Journal of Phonetics* **21**, 337–341 (1993).
37. Gonzalez, J. Estimation of speakers' weight and height from speech: A re-analysis of data from multiple studies by Lass and colleagues. *Perceptual and Motor Skills* **96**, 297–304 (2003).
38. Sell, A. *et al.* Adaptations in humans for assessing physical strength from the voice. *Proc Biol Sci* **277**, 3509–3518 (2010).
39. Raine, J., Pisanski, K., Oleszkiewicz, A., Simner, J. & Reby, D. Human Listeners Can Accurately Judge Strength and Height Relative to Self from Aggressive Roars and Speech. *iScience* **4**, 273–280 (2018).
40. Raine, J., Pisanski, K., Bond, R., Simner, J. & Reby, D. Human roars communicate upper-body strength more effectively than do screams or aggressive and distressed speech. *PLOS ONE* **14**, e0213034 (2019).
41. Bayard, D., Weatherall, A., Gallois, C. & Pittam, J. Pax Americana? Accent attitudinal evaluations in New Zealand, Australia and America. *Journal of Sociolinguistics* **5**, 22–49 (2001).
42. Piske, T., MacKay, I. R. A. & Flege, J. E. Factors affecting degree of foreign accent in an L2: a review. *Journal of Phonetics* **29**, 191–215 (2001).
43. Linville, S. E. Acoustic correlates of perceived versus actual sexual orientation in men's speech. *Folia Phoniatr Logop* **50**, 35–48 (1998).
44. Munson, B., McDonald, E. C., DeBoe, N. L. & White, A. R. The acoustic and perceptual bases of judgments of women and men's sexual orientation from read speech. *Journal of Phonetics* **34**, 202–240 (2006).
45. Giles, H. Evaluative Reactions to Accents. *Educational Review* **22**, 211–227 (1970).
46. Pear, T. H. *Voice and personality, as applied to radio broadcasting*. 247 (Wiley, 1931).
47. Oosterhof, N. N. & Todorov, A. The functional basis of face evaluation. *Proceedings of the National Academy of Sciences* **105**, 11087–11092 (2008).
48. Fiske, S. T., Cuddy, A. J. C. & Glick, P. Universal dimensions of social cognition: warmth and competence. *Trends in Cognitive Sciences* **11**, 77–83 (2007).
49. McAleer, P., Todorov, A. & Belin, P. How Do You Say 'Hello'? Personality Impressions from Brief Novel Voices. *PLOS ONE* **9**, e90779 (2014).
50. Lavan, N. The time course of person perception from voices. Preprint at <https://doi.org/10.31234/osf.io/7rzem> (2022).
51. Mileva, M. & Lavan, N. How quickly can we form a trait impression from voices? *Journal of Experimental Psychology. General* (2023).
52. Guldner, S. *et al.* Human talkers change their voices to elicit specific trait percepts. Preprint at <https://doi.org/10.31234/osf.io/afky7> (2022).
53. Groyecka-Bernard, A. *et al.* Do Voice-Based Judgments of Socially Relevant Speaker Traits Differ Across Speech Types? *J Speech Lang Hear Res* **65**, 3674–3694 (2022).
54. Baus, C., McAleer, P., Marcoux, K., Belin, P. & Costa, A. Forming social impressions from voices in native and foreign languages. *Sci Rep* **9**, 414 (2019).
55. Mahrholz, G., Belin, P. & McAleer, P. Judgements of a speaker's personality are correlated across differing content and stimulus type. *PLOS ONE* **13**, e0204991 (2018).

56. McAleer, P. & Belin, P. The perception of personality traits from voices. in (eds. Fruholz, S. & Belin, P.) 585–606 (Oxford University Press, 2018).
57. Tigue, C. C., Borak, D. J., O'Connor, J. J. M., Schandl, C. & Feinberg, D. R. Voice pitch influences voting behavior. *Evolution and Human Behavior* **33**, 210–216 (2012).
58. Torre, I., Goslin, J., White, L. & Zanatto, D. Trust in artificial voices: A 'congruency effect' of first impressions and behavioural experience. in *Proceedings of the Technology, Mind, and Society* 1–6 (Association for Computing Machinery, 2018). doi:10.1145/3183654.3183691.
59. Dobs, K., Isik, L., Pantazis, D. & Kanwisher, N. How face perception unfolds over time. *Nat Commun* **10**, 1258 (2019).
60. Lavan, N., Smith, H., Jiang, L. & McGettigan, C. Explaining face-voice matching decisions: The contribution of mouth movements, stimulus effects and response biases. *Attention, Perception, & Psychophysics* **83**, 2205–2216 (2021).
61. D'Onofrio, A. Persona-based information shapes linguistic perception: Valley Girls and California vowels. *Journal of Sociolinguistics* **19**, 241–256 (2015).
62. Owren, M. J., Berkowitz, M. & Bachorowski, J.-A. Listeners judge talker sex more efficiently from male than from female vowels. *Perception & Psychophysics* **69**, 930–941 (2007).
63. Hazan, V. Speech Communication Across the Life Span. *Acoustics Today* **13**, 36–43 (2017).
64. Cartei, V. & Reby, D. Effect of Formant Frequency Spacing on Perceived Gender in Pre-Pubertal Children's Voices. *PLOS ONE* **8**, e81022 (2013).
65. Hughes, S. M. & Harrison, M. A. I like My Voice Better: Self-Enhancement Bias in Perceptions of Voice Attractiveness. *Perception* **42**, 941–949 (2013).
66. Armstrong, M. M., Lee, A. J. & Feinberg, D. R. A house of cards: Bias in perception of body size mediates the relationship between voice pitch and perceptions of dominance. *Animal Behaviour* **147**, 43–51 (2019).
67. Burton, A. M. & Bonner, L. Familiarity influences judgments of sex: The case of voice recognition. *Perception* **33**, 747–752 (2004).
68. Foulkes, P. & Docherty, G. The social life of phonetics and phonology. *Journal of Phonetics* **34**, 409–438 (2006).
69. Goldinger, S. D. Echoes of echoes? An episodic theory of lexical access. *Psychol Rev* **105**, 251–279 (1998).
70. Valentine, T. A Unified Account of the Effects of Distinctiveness, Inversion, and Race in Face Recognition. *The Quarterly Journal of Experimental Psychology Section A* **43**, 161–204 (1991).
71. Nosofsky, R. M. Exemplar-based accounts of relations between classification, recognition, and typicality. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **14**, 700–708 (1988).
72. Latinus, M., McAleer, P., Bestelmeyer, P. E. G. & Belin, P. Norm-Based Coding of Voice Identity in Human Auditory Cortex. *Curr Biol* **23**, 1075–1080 (2013).
73. Rosch, E. *Principles of categorization*. 270 (MIT Press, 2002).
74. Valentine, T. & Bruce, V. Recognizing familiar faces: The role of distinctiveness and familiarity. *Canadian Journal of Psychology / Revue canadienne de psychologie* **40**, 300–305 (1986).
75. Lakoff, G. *Women, Fire, and Dangerous Things: What Categories Reveal about the Mind*. (University of Chicago Press, 2008).
76. Posner, M. I. Abstraction and the process of recognition. in *The psychology of learning and motivation* (Academic Press, 1969).
77. Light, L. L., Kayra-Stuart, F. & Hollander, S. Recognition memory for typical and unusual faces. *J Exp Psychol Hum Learn* **5**, 212–228 (1979).
78. Baumann, O. & Belin, P. Perceptual scaling of voice identity: common dimensions for different vowels and speakers. *Psychol Res* **74**, 110–120 (2010).

79. Valentine, T., Lewis, M. B. & Hills, P. J. Face-Space: A Unifying Concept in Face Recognition Research. *Quarterly Journal of Experimental Psychology* **69**, 1996–2019 (2016).
80. Lavan, N., Burton, A. M., Scott, S. K. & McGettigan, C. Flexible voices: Identity perception from variable vocal signals. *Psychon Bull Rev* **26**, 90–102 (2019).
81. Tsantani, M., Kriegeskorte, N., McGettigan, C. & Garrido, L. Faces and voices in the brain: A modality-general person-identity representation in superior temporal sulcus. *NeuroImage* **201**, 116004 (2019).
82. Carlin, J. D. & Kriegeskorte, N. Adjudicating between face-coding models with individual-face fMRI responses. *PLOS Computational Biology* **13**, e1005604 (2017).
83. Tsantani, M. *et al.* FFA and OFA encode distinct types of face identity information. *J. Neurosci.* (2021) doi:10.1523/JNEUROSCI.1449-20.2020.
84. Ponsot, E., Arias, P. & Aucouturier, J.-J. Uncovering mental representations of smiled speech using reverse correlation. *J Acoust Soc Am* **143**, EL19 (2018).
85. Ponsot, E., Burred, J. J., Belin, P. & Aucouturier, J.-J. Cracking the social code of speech prosody using reverse correlation. *PNAS Proceedings of the National Academy of Sciences of the United States of America* **115**, 3972–3977 (2018).
86. Knight, S., Lavan, N., Kanber, E. & McGettigan, C. The social code of speech prosody must be specific and generalizable. *Proc Natl Acad Sci U S A* **115**, E6103 (2018).
87. Lee, J. J. & Perrachione, T. K. Implicit and explicit learning in talker identification. *Atten Percept Psychophys* **84**, 2002–2015 (2022).
88. Kapadia, A. M., Tin, J. A. A. & Perrachione, T. K. Multiple sources of acoustic variation affect speech processing efficiency. *The Journal of the Acoustical Society of America* **153**, 209–223 (2023).
89. Remez, R. E., Fellowes, J. M. & Rubin, P. E. Talker identification based on phonetic information. *Journal of Experimental Psychology: Human Perception and Performance* **23**, 651–666 (1997).
90. Sheffert, S. M., Pisoni, D. B., Fellowes, J. M. & Remez, R. E. Learning to recognize talkers from natural, sinewave, and reversed speech samples. *J Exp Psychol Hum Percept Perform* **28**, 1447–1469 (2002).