Beyond correlation: acoustic transformation methods for the experimental study of emotional voice and speech

Pablo Arias, Laura Rachman, Marco Liuni, Jean-Julien Aucouturier

► To cite this version:

Pablo Arias, Laura Rachman, Marco Liuni, Jean-Julien Aucouturier. Beyond correlation: acoustic transformation methods for the experimental study of emotional voice and speech. *Emotion Review*, 2020, 10.1177/1754073920934544 . hal-02907502

HAL Id: hal-02907502

<https://hal.science/hal-02907502>

Submitted on 27 Jul 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Beyond correlation: acoustic transformation methods for the experimental study of emotional voice and speech.

Pablo Arias^{1,*}, Laura Rachman^{1,*}, Marco Liuni¹, Jean-Julien Aucouturier¹

Abstract

While acoustic analysis methods have become a commodity in voice emotion research, experiments that attempt, not only to describe, but to computationally manipulate expressive cues in emotional voice and speech have remained relatively rare. We give here a non-technical overview of voice-transformation techniques from the audio signal processing community that we believe are ripe for adoption in this context. We provide sound examples of what they can achieve, examples of experimental questions for which they can be used, and links to open-source implementations. We point at a number of methodological properties of these algorithms, such as being specific, parametric, exhaustive and real-time, and describe the new possibilities that these open for the experimental study of the emotional voice.

Keywords

voice, emotion, transformation, methods

Introduction

A typical approach to study how humans and other animals communicate emotions vocally uses acoustic analysis to quantify the physical features of vocalizations, such as their fundamental frequency (F0), intensity, or spectrum, and seek how they relate to the affect of the speaker (Bachorowski and Owren 1995) or listener (for a review, see e.g. Scherer 2003; Juslin and Laukka 2003). Several well-known tools exist in the community for this purpose, including PRAAT¹ for speech data (Boersma and Weenink 2002), OpenSMILE² (Eyben et al. 2013), MIRToolbox³ for song and musical data (Lartillot and Toivainen 2007), and Sound Analysis Pro⁴ (Tchernichovski and Mitra 2004) or Seawave⁵ (Sueur et al. 2008) for animal communication. The availability of dedicated software has an important impact on research: it gives access to audio signal processing techniques such as F0 extraction without needing a technical background; it helps standardize the definition of vocal features, by providing reference implementations (when one studies *jitter* while referencing PRAAT, others know what is meant and how to reproduce the work); and it provides an interdisciplinary interface between the research that creates these tools and the research that uses them, to share new techniques and new research needs.

One limitation of this methodology however is that it is intrinsically correlational. Analysing large corpora of speech or vocalizations to establish e.g. that happy voices have statistically higher F0s, faster rate and more animated intonations (Banse and Scherer 1996; for a recent review, see Kamiloglu et al. 2019) do not allow us to conclude that these features are biological signals that are causally involved in either the decoding of these emotions. For instance, it is now relatively well-described that smiling, the bilateral contraction of the zygomatics facial muscles, has perceivable acoustic consequences on the speaking voice that can be heard, e.g., on the phone (Tartter and Braun 1994). How listeners recognize smiles in speech, however, is

⁰¹STMS UMR9912 (IRCAM/CNRS/Sorbonne Universit), Paris (France)

* Equal contribution

0

Corresponding author:

Jean-Julien Aucouturier, STMS UMR9912, IRCAM, 1 Place Stravinsky, 75004 Paris (France)

⁰Email: aucouturier@gmail.com

¹<http://www.praat.org/>

²<https://www.audeering.com/opensmile/>

³<https://www.jyu.fi/hytk/fi/laitokset/mutku/en/research/materials/mirtoolbox>

⁴<http://soundanalysispro.com>

⁵<http://rug.mnhn.fr/seewave/>

Figure 1. The confusing case of fundamental frequency in smiled speech. Ohala (1980) demonstrated that the resonant frequencies (formants) of a cylindrical clay model of the vocal tract are lowered when the shape of a smile is carved into it (A, middle), and that these new frequencies are similar to that of a non-smiling, but shortened cylinder (A, bottom). He concluded that smiling only has mechanistic consequences on the resonances of the vocal tract filter (A). This is confirmed by the fact that listeners are able to recognize smiled speech in both pitched and non-pitched (whispered) vocalizations (B). However, it is typical that speakers simultaneously raise their pitch while smiling (C), and thus listeners use higher pitch as a cue to identify smiled speech (D), even though it is not causally implicated in smiling. The sole acoustic analysis of speech corpora therefore cannot fully elucidate what cognitive mechanisms are involved in how smile is communicated vocally. Figure adapted from Ohala 1980; Tartter and Braun 1994; Tartter 1980; Lasarcyck and Trouvain 2008.

remarkably complicated. On the one hand, physical models show that labial spreading, which reduces vocal tract length, has no mechanistic consequence on the F0 of the glottal source, but only on formant frequencies (Figure 1A; Ohala 1980; Drahota et al. 2008); and thus, listeners are able to identify smiled speech in the whisper register (Figure 1B; Tartter and Braun 1994), which has no audible F0. On the other hand, corpus analyses typically find strong associations between smiled speech and F0 (Figure 1C; Tartter 1980; Barthel and Quené 2015), which suggests that it is neurologically difficult for speakers to smile without simultaneously raising their pitch; and thus, in the normal register, listeners reliably use pitch as a cue to identify smiled speech, even though it is not causally implicated in its production (Figure 1D; Lasarcyck and Trouvain 2008; El Haddad et al. 2015). It is therefore clear, in this field of study like in any other (Casadevall and Fang 2008), that correlation does not imply causation and that the sole acoustic analysis of what is incidentally present in voice and

speech may obfuscate the mechanisms with which emotions are produced, or recognized (see Armstrong et al. 2019 for a similar argument on the signaling of body size by low F0 voices).

Rather than *describing* them, one would like the ability to *manipulate* the acoustic factors of interest in stimuli, in order to confirm experimentally that they causally lead to a change of behavior in the predicted direction when they are perceived. The manipulation of acoustic cues provides an approach complementary to corpus analyses, where the latter can establish a relation between two phenomena (e.g. shifted formants when people smile) and the former can be used to build a model and test for their involvement in perception⁶ (see also Goldstone and Lupyan 2016). Yet, while analysis tools are many, experiments that attempt to manipulate acoustic dimensions computationally in complex stimuli such as speech (Scherer 1972), music (Ilie and Thompson 2006), or animal vocalizations (Hienz et al. 2004) have been, until recently, remarkably rare. Perhaps because acoustic transformation tools are perceived to be too technical or of insufficient quality, a steady stream of research has even preferred less flexible, but more ethologically-valid ways to manipulate vocal characteristics, such as immersing animals in heliox (Nowicki et al. 1989; Rand and Dudley 1993).

Two lines of research have significantly advanced the quest for acoustic control and causal inference in voice and speech research: vocal morphing and speech synthesis. On the one hand, morphing —an algorithmic method to combine two voices by interpolating their spectral features (Kawahara and Matsui 2003) —, has allowed researchers to describe, e.g., how formants are processed to represent vocal identity (Latinus et al. 2013), whether averaged voices are perceived as more attractive (Bruckert et al. 2010) or whether vocal emotion are perceived categorically (Laukka 2005). However, morphing is generally performed between voices that differ in more than one acoustic dimension. For example, two morphed vocal identities that differ in their fundamental frequency, formant dispersion, and Harmonic-to-Noise Ratio (HNR), will inevitably generate experimental conditions where these acoustic features co-vary (Latinus et al. 2013). On the other hand, speech synthesis —a vast family of methods allowing to create artificial vocal stimuli from scratch by specifying part or all of their physical parameters (for a review, see e.g. Govind and Prasanna 2013; Malisz et al. 2019) —, has been used e.g. to reveal that vocal

⁶It should be noted that transformation do not help to establish the involvement of the cues in the production/expression of emotion; for this, one has to manipulate the actual emotional state of the signalers (e.g. Bachorowski and Owren 1995)

emotions can be recognised in isolated pitch contours (Scherer and Oshinsky 1977), to compare the emotional impact of various forms of non-linearities (Anikin 2019), or test the effect of formant frequencies on the recognition of smiled speech (Quené et al. 2012). While allowing theoretically unlimited control over the physical properties of the stimuli, synthesis methods have the caveat of decontextualising acoustic features (e.g. when replacing full speech with isolated pitch contours) and may also suffer from sonic artefacts (e.g. voices sounding robotic and artificial - although see this article's final prospective note on deep-learning techniques).

In our view, voice transformation —the technique to manipulate an original, natural vocal utterance in order to alter a specific acoustic dimension —provides a useful alternative⁷ to the morphing and synthesis approaches. By giving experimenters the ability to predict how behavioral, physiological or neural reactions vary depending on specific acoustic changes, while leaving all other features unchanged, transformations are well adapted to the nature of hypothesis-driven experimental research. Furthermore, because of recent improvements in vocal transformation algorithms, and their increasing availability in the computer-science communities (Stylianou 2009; Mohammadi and Kain 2017a), it may be time to consider these technologies part of the toolbox of vocal emotion researchers.

In this article, we review the neurological and acoustic bases of emotional voice production, and show that there are in fact well-established, high-quality algorithms that are able to parametrically transform a prerecorded vocal signal at all levels of the voice production pathway. We provide sound examples from our own recent work that illustrate these techniques, examples of experimental questions for which they can be used, and links to open-source software to replicate and extend these studies. Finally, we point at a number of methodological properties of these algorithms, such as being specific, parametric, exhaustive and real-time, and describe the new possibilities, and questions, that these entail for experimental research.

Voice transformations along the vocal production pathway

Voice is produced when the expiratory airflow from the lungs, generated by thoracic and abdominal muscles, sets the vocal folds of the larynx into oscillations. This sound wave, or glottal source, is resonated through the vocal tract, filtered by oral articulators such as the tongue and lips which amplify certain bands of energy (or formants) in its

frequency spectrum, and is finally radiated from the mouth and nose (Figure 2; Titze 1994). While the vibration of the vocal folds is a passive process, their oscillatory properties, the airflow that sets them into motion, and the resonance characteristics of the vocal tract are all controlled by over 100 respiratory, laryngeal, and oro-facial muscles (Simonyan and Horwitz 2011), whose motoneurons originate from the spinal cord and brainstem.

In emotional vocalizations, neural control over these muscles involve a hierarchy of cortical and subcortical systems including the periaqueductal gray (PAG), anterior cingulate cortex (ACC) and laryngeal motor cortex (LMC). These subcortical and cortical influences on muscle actuators at every stage of the vocal production pathway have different, complementary effects on the final acoustic properties of the vocal signals, and specific voice-transformation techniques exist to reproduce these changes in ecological voice and speech recordings.

Glottal source transformations

Changes in the subglottal pressure due to the contraction of thoracic and abdominal muscles, which are controlled from the anterior horn of the spinal cord, primarily lead to modulations of voice intensity. At moderate intensities, such as in e.g. happy, aroused voices compared to calm or sad voices, the effect of the modulation is carried linearly through the vocal pathway and can be simulated with a simple scalar multiplication of the recording's root mean square (RMS) intensity (see e.g. Ilie and Thompson 2006) or, for arbitrary intensity profiles, a piece-wise linear function as implemented e.g. in the reverse-correlation toolbox CLEESE⁸ (Burred et al. 2019).

Increased airflow, such as in pain cries or anger shouts, but also possibly altered neurological control over the laryngeal muscles, such as in stress or anxiety, may drive the vocal folds into non-linear/chaotic oscillatory regimes and, more generally, change the shape and periodicity of glottal pulses, which resulting in audible alterations of sound quality such as roughness, noisiness or breathiness (Figure 2C). In voice measurements, such non-linearities are often analysed in terms of jitter and shimmer (cycle-to-cycle variations in the period and amplitude of glottal pulses, respectively)

⁷Note that, strictly speaking, there is a technological overlap between voice transformation and certain forms of (re-)synthesis techniques, which analyse an original signal down to its generative parameters and then resynthesize a variant of the signal by altering these parameters. We review some of these techniques in the rest of the paper under the encompassing term of *transformation*.

⁸<http://forumnet.ircam.fr/product/cleese>

Figure 2. Neurological, anatomical and acoustic characteristics of the vocal production pathway. Voice is produced when the expiratory airflow from the lungs sets the vocal folds of the larynx into oscillations (right). This sound wave, or glottal source, is resonated through the vocal tract and filtered by oral articulators such as the tongue and lips (left) and is finally radiated from the mouth and nose. Neural control at different levels of this pathway involve a hierarchy of cortical and subcortical systems including premotor nuclei of the brainstem such as the nucleus ambiguus (Amb.), periaqueductal gray (PAG), anterior cingulate cortex (ACC) and laryngeal motor cortex (LMC). Voice-transformation techniques exist to parametrically manipulate a prerecorded vocal signal at all levels of the voice production pathway: vocal tract filter (A), pitch and intonation (B), and glottal source timbre (C).

and harmonic-to-noise ratio (HNR) Boersma and Weenink (2002). Such modulations of vocal source quality are important in emotional behaviors (Johnstone and Scherer 1999; Gobl and Chasaide 2010) and have been related, in listeners, to subcortical processing by the amygdala (Arnal et al. 2015).

Not all glottal source changes are easily simulated with voice transformations. In theory, analysis-resynthesis techniques can be used to, first, estimate the recording's series of glottal pulses (Degottex et al. 2013) and, then, resynthesize the vocal signal from a manipulated series of pulses with artificially-varied amplitude and period (Verma and Kumar 2005; Ruinskiy and Lavner 2008; Böhm et al. 2008). However, these techniques rely on an explicit model of pulse variability, which is typically learned from one or several target examples of naturally rough voices (Bonada and Blaauw 2013), and it is unclear how such predetermined patterns should be selected for arbitrary voices. Moreover, because of the computational complexity of the initial stage of glottal source estimation, these techniques cannot operate in real-time. Alternative approaches can also simulate variations in pulse periodicity by overlapping randomly time-shifted copies of the original recording (Loscós and

Bonada 2004) or modulating it at a divider of F0 to create subharmonics, as implemented e.g. in the ANGUS toolbox⁹ (Gentilucci et al. 2019). However, these only allow exploring a subset of all possible non-linearities (e.g. subharmonics, but not biphonation in general), and vocal source spectrum is one area of vocal production for which pure speech-synthesis approaches, in which variability in the glottal shape or periodicity can be specified explicitly (Brady 2005; Anikin 2018), may provide more experimental control than transformations.

Sound example S1: Female singing voice, first: original; second: manipulated with subharmonics with the ANGUS toolbox (Gentilucci et al. 2019).

Note that, because vocal folds or analog anatomical structures are present in a large number of species, source non-linearities are not unique to human vocalizations but are used as a signal of threat and alarm by a wide range of animals, including primates (Fitch et al. 2002), but also rodents (Blumstein and Recapet 2009), canids (Wilden et al. 1998), whales (Tyson et al. 2007), and birds (Fee et al. 1998).

⁹<http://forumnet.ircam.fr/product/angus>

The use of glottal source transformations can therefore be extended to the study of animal behavior.

Changes in the oscillatory properties of the vocal folds, linked e.g. to their length and opening, are mostly controlled by the intrinsic laryngeal muscles, innervated from the vagal nerve originating in the nucleus ambiguus (Amb.) of the medulla, and lead to modulations of vocal source timbre (as above) but also, and perhaps most importantly, of vocal F0 (Figure 2B). A range of techniques exist to manipulate F0. Simple algorithms, as used for example in altered auditory feedback research (Hain et al. 2001) and implemented in the DAVID toolbox¹⁰ (Rachman et al. 2018), are based on resampling or multiple delay lines (a technique that introduces a small delay to an audio signal in order to play it faster/slower, thus raising/lowering its pitch, Dattorro 1997) and may alter vocal tract filtering or formants unrealistically beyond small parametric changes. State-of-the-art techniques that allow separating source and filter information to avoid such artifacts are based on reconstructions of the signal's short-time Fourier transform (STFT) at non-uniform rates, such as the Pitch Synchronous Overlap and Add (PSOLA) method as implemented e.g. in PRAAT (Boersma and Weenink 2002), the Phase-Vocoder method (Moulines and Laroche 1995) as implemented e.g. in CLEESE (Burred et al. 2019), or pitch-adaptive analyses techniques such as the Adaptive Interpolation of Weighted spectrum method as implemented in STRAIGHT¹¹ (Kawahara 1997). These transformation methods not only allow raising or lowering the mean pitch of a recording, which may correspond to a baseline change of valence (Figure 3A, see e.g. Ilie and Thompson 2006), but can also manipulate the difference between the instantaneous and mean F0 to exaggerate or lessen variations, as seen for example in fearful vs sad vocalizations (Figure 3B, see e.g. Pell and Kotz 2011), create parametric F0 contours such as vibrato in anxious voices (Figure 3C, see e.g. Bachorowski and Owren 1995), or local intonations at the start or end of an utterance, as in surprised or assertive speech (Figure 3D, see e.g. Jiang and Pell 2017). In addition, most pitch-shifting methods can also be used to manipulate the duration or speech rate of utterances (a process known as time-stretching), producing faster or slower speech in positive or negative emotional states (Scherer and Oshinsky 1977).

Sound example S2: Male speech, first: original; second: pitch increased by 50 cents; third: pitch decreased by 50 cents; fourth: pitch modulated with a 8Hz vibrato. All transformations made with delay lines, using the DAVID toolbox (Rachman et al. 2018).

Sound example S3: Female speech, first: original; second-six: random pitch intonations, generated on 6 successive time-windows, using gaussian

distributions centered at ± 0 cents, $SD = \pm 200$ cents. All transformations made with a phase-vocoder, using the CLEESE toolbox (Burred et al. 2019).

Sound example S4: Female speech, first: original; second-sixth: random variations of speed rate, generated on 6 successive time-windows, using gaussian distributions centered at $\pm 0\%$, $SD = \pm 10\%$ original duration. All transformations made with a phase-vocoder, using the CLEESE toolbox (Burred et al. 2019).

Vocal tract transformations

The shape and resonating characteristics of the vocal tract, and thus, the spectral properties of the sound, are modulated by the articulators of the supraglottal region (e.g. lips, tongue, jaw), which are controlled by the oro-facial motoneurons of the facial and trigeminal nuclei of the pons (Figure 2A), but also by the extrinsic laryngeal muscles, whose motoneurons originate from the hypoglossal nucleus of the medulla and which have the ability to raise or lower the position of the larynx within the neck (Titze 1994). As for the glottal source, a number of techniques exist to transform vocal tract characteristics, without altering other aspects of vocal production. Simple techniques, such as the PRAAT “Change Gender...” method (Boersma and Weenink 2002), exploit the side-effects of the resampling pitch-shifting method to reduce or increase formant dispersion, thus simulating changes in vocal tract length and physical dominance (Boidron et al. 2016).

Other techniques enabling the manipulation of individual formant frequencies include formant resynthesis methods (Quené et al. 2012), spectral envelope manipulations by frequency warping (Arias et al. 2018b) and neural networks (Narendranath et al. 1995). While vocal tract characteristics have been mainly studied in the context of speaker identity (Mohammadi and Kain 2017b), recent work has suggested they are also actively manipulated by emotional oro-facial gestures such as those involved in the expression of disgust (Chong et al. 2018) or smiling (Arias et al. 2018a).

Sound example S5: Female speech, first: original; second-sixth: random variations of vocal tract filter, generated with spectral envelope frequency warping using the CLEESE toolbox (Burred et al. 2019).

Sound example S6: Female speech, first & third: original; second & fourth: formant shifted upwards to simulate smiling, generated with frequency warping (Arias et al. 2018b).

¹⁰<http://forumnet.ircam.fr/product/david>

¹¹http://www.wakayama-u.ac.jp/~kawahara/STRAIGHTadv/index_e.html

Figure 3. Four examples of pitch transformations on a single recording of the sentence “I would like a new alarm clock”, produced by a male English speaker. The original pitch values are presented in black and the transformed pitch values in magenta. The speech waveform is shown on the x-axis and the dotted line indicates the mean F0 of the recording. (A) 100-cent upward pitch shift, applied uniformly over time. (B) Exaggerated pitch dynamics by a 15% increase of pitch values with respect to the mean pitch. (C) Vibrato applied with a 100 cent depth and a rate of 8.5 Hz. (D) Increasing pitch at the end of the utterance.

As for glottal source above, vocal tract transformations are also relevant beyond human voice and speech, for the study of animal communication. Formant dispersion correlates with body size in rhesus macaque monkeys (Fitch 1997), and dominant frequency (that of the formant with highest amplitude) inversely correlates with body size across 91 mammalian species (Bowling et al. 2017). Researchers have used formant manipulations of animal vocalizations in playback experiments to show e.g. that red deer stags are more attentive and reply more to calls from bigger conspecifics (Reby et al. 2005), or that female koalas spend

more time attending to vocalizations simulating larger males (Charlton et al. 2012).

Properties of voice transformations for experimental research

Voice transformation techniques have a number of methodological properties that make them well suited for experimental research, and allow experimental paradigms that would not be otherwise possible with traditional analysis methods or using actor-recorded vocalizations. We list and comment here on a number of these properties, namely that of specificity, parametricity, exhaustivity, and real-time behavior. Despite their advantages for research, assessing whether a given acoustic transformation possesses each of these properties also brings a number of methodological and theoretical difficulties, which we discuss below.

Specificity

Vocal behavior is highly multidimensional and, either for anatomical or neural reasons, it is nearly impossible for speakers to produce for instance only the pitch aspects of a sad expression without simultaneously varying timbre or duration, or timbre aspects without simultaneously varying pitch (Figure 1). Co-variation of features in naturalistic recordings means that one cannot conclusively establish what acoustic property drive listeners’ emotional evaluations. The ability to transform voice along specific signal dimensions, such as pitch or timbre, while preserving all other aspects of the original, opens the possibility to create pairs of stimuli (original and transformed) that differ only in one experimental factor. Comparing behavioral or physiological measures within pairs thus allows controlling for co-varying factors and identify causal relations that would otherwise be missed in noise or, worse, spuriously attributed to the wrong factor.

Figure 4A illustrates this situation with recent data from Arias et al. (2018a), in which we manipulated 20 original male and female recordings with a vocal tract transformation designed to simulate the acoustic effect of smiling, and asked 35 participants to evaluate their perceived happiness of the speaker. Each recording was transformed into two variants, smiled and unsmiled versions, respectively expanding or compressing formants by the same amount. These two versions are thus two opposite transformations relative to a neutral, non-manipulated stimulus. Because the transformation did not affect glottal source properties, the F0s of the original recordings were preserved by the two transformations, and ranged from 117 to 300 Hz. While there

Figure 4. Attributing ratings of speaker happiness to manipulated vocal tract cues in smiled speech. (A) Mean rating of happiness for 20 male and female recordings, as a function of the recording's mean pitch and sound manipulation : smile (dark green), neutral (black), unsmile (light green); each dot represents one separate recording; black arrows link unsmile with smile transformations of the same original recording. (B) Normalized ratings of happiness for male and female recordings manipulated with smile (light green) and unsmile (dark green) transformations; asterisks indicate significant statistical differences between the smile and unsmile distributions, computed with paired t-tests ($p < 0.01$); all ratings normalised by the corresponding neutral (non-manipulated) rating. (C) Difference of formant frequency (in cents) between smile (dark green) and unsmile (smile green) transformations and the corresponding non-manipulated sound (formant 1: F1, formant 2: F2) ; asterisks indicate significant statistical difference between the distributions, computed with paired t-tests ($p < 0.01$)

was a main effect of the smile transformation on ratings of perceived happiness within each triplet of recordings ($\chi^2(12)=55.2$, $p=1.0e-12$), variation of ratings across triplets was much greater because of the varied phrase content (positive vs negative semantics), prosody, speaker identity, etc. of the original recordings. These variations would likely mask the smaller yet remarkably consistent effect of vocal tract properties if investigated with another experimental design.

Caveat: Feature specificity is not a given in signal processing, though. For instance, pitch shifting algorithms based on harmonizers, as used in the altered vocal feedback literature, do not attempt to separate glottal and vocal tract characteristics and thus have side-effects on the signal's formants. Perhaps because of their technical nature, these effects are not always properly acknowledged in the psychological literature (e.g. *"The harmonizer shifts all frequencies, voice F0 as well as formants, and thus the shifted feedback signal sounds like a person's normal voice at a different F0 (details of the pitch-shifting algorithm are a trade secret of the manufacturer and are thus unavailable)"* - Hain et al. 2001), even though they may confound results attributed to the manipulation. Ensuring the right level of specificity, i.e. using traditional acoustic analysis methods such as PRAAT to validate the absence of effect of the transformation on vocal characteristics that could have a confounding impact on the measures of interest, should therefore be high on the agenda for researchers aiming to adopt these methodologies. This will be greatly helped by

the availability of open-source, validated software whose possibilities and limitations are well-documented.

Parametricity

A common research question aims to compare how emotional expressions are processed across stimulus conditions, testing e.g. whether listeners perceive emotional speech differently in their native or a foreign language (Scherer et al. 2001; Pell et al. 2009), on familiar or non-familiar speakers (Chen et al. 2016), male or female voices (Bonebright et al. 1996), or even across speech and music (Juslin and Laukka 2003). When using emotional stimuli produced by actors, differences in decoding performance across groups may arise either because of production or perception differences. For instance, if French listeners have difficulties processing emotional cues spoken by Japanese speakers, it may be because their auditory representations of the Japanese phonetic inventory are poor (a perception effect, Dupoux et al. 1999), because one does not use the same cues to express joy in Japanese and in French (a production effect, Kitayama et al. 2006), or both. **Balanced designs, in which stimuli produced by more than two culture groups are presented to all participants cross-culturally, is a useful strategy to disentangle such decoding and encoding differences: encoding differences should result in poor recognition across non-native groups, while decoding differences should result in impaired recognition only for a given group, but across stimuli from the other groups** (Matsumoto 2002; Yoshie

and Sauter 2019). **Once an in-group advantage has been established however, transformation tools that can parametrically manipulate stimuli then become useful to examine the causal mechanisms that underlie such cross-cultural differences.**

In programming language theory, parametricity is the property of a function that can handle input values identically without depending on their type (Pierce and Benjamin 2002). We use the term here to describe the property of voice transformations having a uniform acoustic effect regardless of the characteristics of the original signal, their speaker, content or culture. Parametricity opens the possibility to create emotional voice stimuli which utilize *exactly* the same prosodic cues in *exactly* the same manner (e.g. a 50-cent pitch increase in the first 2 seconds), and thus to separate the effect of production and perception in inter-group differences.

Data in Figure 4B illustrates this possibility by comparing the effect of the same smile transformation on a variety of recorded sentences by both male and female speakers. Even though the acoustics of male and female utterances widely differ, notably because of pitch dimorphism (Titze 1994), listeners' ratings of speaker positivity in both smile and unsmile transformations can be normalized with respect to original ratings, which allows us to compare the effect of identical cues between sexes. In the same experimental logics, audio transformations have been used to compare identical cues on speech, music and environmental sounds (Ilie and Thompson 2006; Ma and Thompson 2015).

Caveat: As for specificity, a number of technical considerations may reduce a given algorithm's parametricity. Some may result from algorithmic design choices. For instance, because F0/pitch is ill-defined on unvoiced portions of speech, pitch transformations, such as vibrato, are often designed to only operate on voiced sections of the signal, leaving transients untransformed. Transformations like vibrato therefore rely on the availability of relatively long voiced portions in phonemes (e.g., 250 ms for two cycles of an 8-Hz-vibrato to be perceived) and, even if specified with identical parameters, may not have as much physical effect on speakers with a faster rather than slower speech rate, or languages with a larger rather than smaller consonant/vowel ratio (see e.g. the failed generalization to Swedish in Rachman et al. 2018). Other breaches of parametricity may result from algorithmic limitations. Figure 4C illustrates this situation: despite the identical transformation parameters, the smile transformation as implemented in Arias et al. (2018a) unintentionally leads

to smaller physical changes of formant frequency F2 (as measured by PRAAT) in male voices than in female voices, making the higher ratings of perceived smiliness measured on transformed female voices difficult to attribute to a purely perceptual effect. It is therefore crucial for researchers to develop a technical understanding, as well as the methodological know-how, to critically assess a transformation's parametricity.

Exhaustivity

Although the human vocal apparatus can produce many different vocal sounds, languages generally use only a small subset of these sounds, and these are not uniformly distributed between and within languages (De Boer 2000). Many psychophysical procedures, such as adaptive staircase methods but also reverse correlation (Jack and Schyns 2017), require presenting participants with random distributions of stimuli in feature space, which is often impractical using naturalistic vocalizations (Belin et al. 2017). With voice transformations, experimenters can uniformly or adaptively sample a large space of prosodic variations, e.g. all vibrato frequencies between 1-10Hz, regardless of how common these may be in actual behavior, and are thus able to draw better inferences about how these features are processed.

For instance, in Ponsot et al. (2018), we used a phase-vocoder transformation to generate more than 70,000 random prosodic variations from a single recording of the word 'bonjour' (hello), and used reverse correlation to uncover participants' mental representations of a dominant or trustworthy way to pronounce that word. Similar paradigms can be applied to study how healthy participants and patients mentally represent emotional prosody, and address e.g. emotion perception deficits in pathologies like ASD or amusia.

Caveat: While their parameters can be explored exhaustively, most transformation techniques do not implement boundaries on what is or isn't physiologically possible and thus, beyond a certain parameter range, may not not sound like authentic human speech. In addition, even within realistic parameter ranges, transformations may generate artefacts, such as unnatural timbres, doubling of F0 or smearing of transients, and give the transformed sound a robotic, artificial tone. Even if the impact of artificiality on emotional judgements remains unclear (Burleigh et al. 2013), one can rightly question the ecological validity of behaviors measured in response to such stimuli. While most artefacts can be avoided with simple heuristics (e.g. basing parameters on the range of variation measured in natural voices, or clipping sampling

distributions at $\pm 2SD$), quantifying how natural, or easily detectable, a given transformation sounds to participants can be remarkably complicated. First, the acceptance of transformed voice as authentic is heavily dependent on context. For instance, transformations may be more easily detected in situations where variants can be compared with the original sound. Second, judgements of naturalness are multifaceted, and likely incorporate evaluations of biological plausibility (*this doesn't sound human*), agentivity (*this sound tampered with*), vocal or social typicality (*no one in their right mind would do this*) or even semantics (*it does not make sense to say this with a happy voice*). Future research will benefit from more principled ways to measure transformation naturalness and its impact on participants.

Real-time behavior

One typical way to study the role of emotional expressions in social interactions is either to explicitly instruct social partners to display a certain emotion (e.g. Tice 1992), or to indirectly lead them to express it using a cover story (e.g. Van Doorn et al. 2012). Because many voice-transformation techniques allow real-time processing (typically, 50-100ms latency for transformations based on phase-vocoder, Lee et al. 2007), they open the possibility to control emotional expression in continuous, real-time interactions (e.g. on the phone) with no experimental demand, and possibly even without participants' awareness of the manipulation. For instance, we have used a vocal tract transformation to manipulate the perceived body size of mock patients calling a medical call center simulator (Boidron et al. 2016), and found that callers whose voice were perceived as indicative of physical dominance obtained a higher grade of response, a higher evaluation of medical emergency and longer attention from physicians than callers with strictly identical medical needs whose voice signaled lower dominance. Similar paradigms can be used e.g. to study how congruent or incongruent emotional expressions influence the outcome of group decision-making, or group creativity.

In addition, some voice-transformation techniques, such as those operating in the time-domain (Juillerat et al. 2008), can be so fast that they can not only manipulate a social partner's voice without disrupting the flow of interaction, but also manipulate the participant's own voice without disrupting their sensorimotor feedback (e.g. with less than 20ms latency between the original input voice and the manipulated output). This opens the possibility to build altered auditory feedback paradigms and test how e.g. hearing one's voice with a happier or sadder tone influences one's emotional experiences, judgments or decisions. For

instance, we used a time-domain transformation of pitch to modify participants' voices in a happy or sad direction as they read out an emotional neutral text, and found that participants who heard themselves with emotionally manipulated voices reported significantly different moods, as well as elevated levels of skin conductance (Aucouturier et al. 2016). Similar paradigms can be used to study e.g. whether personal emotional memories can be re-encoded with different valence when healthy participants or patients hear themselves narrate them with a transformed tone of voice.

Caveat: Once given the possibility to transform continuous speech in real-time interactions, the immediate next question concerns the contextual adaptation of transformations to speech content. While transformation parametricity is methodologically useful, applying the same pitch increase on spoken sentences regardless of, say, their original prosody or stress words (which are also marked with pitch - Pell 2001) may create unnatural or misadaptive expressions, and it is computationally unclear how to adapt transformations to speech content, especially in a real-time context. In that respect, voice transformations should not be considered experimental materials, but experimental methods. They do not replace actor-produced stimuli nor provide ready-made expressions with which to study the impact of emotions on behavior (in truth, transformed vocalizations may be less intense, less natural and less well-recognized than natural expressions, see e.g. Rachman et al. 2018). Rather, by manipulating signal properties at all stages of the vocal production pathway, they provide control over the physical properties of the stimuli, bringing unprecedented precision on the neurological, anatomical and acoustic components of what makes voice and speech emotional, but leaving it to the experimenter to construe how speakers and listeners integrate these components with other aspects of affective and cognitive processing.

A prospective note on deep learning techniques

In recent years, artificial voices have become an integral part of consumer electronic appliances (e.g. smart assistants, car navigators, augmentative and alternative communication) and the amount of funding private companies such as Amazon, Google or Apple have injected into the speech synthesis community has transformed the field into a fast-paced and competitive domain. In just a few years, neural machine learning (deep neural network, DNNs) operating on waveform samples, as opposed to e.g. spectrogram features,

has become the de-facto standard in voice synthesis (Oord et al. 2016).

The DNN architecture, which directly models waveforms using a neural network method trained with recordings of real speech, provides superior acoustic quality for several high-level vocal applications such as Text-to-Speech (Wang et al. 2017) and speech enhancement (Pascual et al. 2017), but also to generate expressive voices (Lee et al. 2017; Akuzawa et al. 2018) or convert their emotions (Luo et al. 2017). However, although highly flexible, these systems have so far failed to exhibit the level of parametric control that we argue here is needed for experimental applications. First, machine-learning speech generally emulates vocal patterns learned from large sets of recordings where all vocal features co-vary, and therefore typically lacks feature specificity and exhaustivity. Second, deep-learning architectures do not allow easy introspection into how information is represented in the network, making it difficult to know what exact vocal features are being manipulated and lacking interpretability. As far as we know, these limitations have so far prevented the application of DNNs for the kind of experimental work reviewed here (although see Sun et al. 2019).

However, progress in the field is fast. One promising line of research uses post-hoc, data-driven methods to reveal how stimulus information is encoded into network layers (Hsu et al. 2017). In the visual sciences, these methods have been used to compare what visual features human and machine use to achieve face classification (Xu et al. 2018), and similar approaches could be used for speech. Another relevant line of research aims to create DNN-based vocoders (Wu et al. 2019), in which, like in traditional vocoders reviewed above, speech synthesis can be controlled with specific pitch or duration parameters, while conserving the acoustic performance of deep learning models. Finally, Generative Adversarial Networks (GANs), a special class of DNN architecture capable of learning a deterministic mapping from one style of stimulus to another (Goodfellow et al. 2014), are increasingly used to create visual transformations (e.g. smiles Wang et al. 2018) and have also started to be applied to speech transformations. For example, GANs were recently used to transform a voice into its Lombard counterpart (a particular type of vocal effort which makes the voice more intelligible in background noise; Seshadri et al. 2019). All such advances open exciting new possibilities to create emotional voice and speech transformations, which will certainly find their way to the community in the upcoming years.

Conclusion

So far, the experimental study of emotional voice and speech has largely relied on acoustic analyses of datasets of natural emotional vocalisations, or the use of these recordings as stimuli in experimental research. While natural vocalizations have the advantage of being realistic and ecological, they often vary in several acoustic dimensions, making mechanistic conclusions difficult to establish. Yet, much is known about the essential anatomic building blocks involved in emotional vocal production (e.g. laryngeal muscle tension, vocal fold oscillatory regime, oro-facial gestures). These anatomic mechanisms are controlled by increasingly well-identified neural structures and have specific, and to some extent independent, acoustic consequences that can be modeled computationally —because they have a physical basis.

In this article, we reviewed a wide range of recent (or not-so-recent) technologies that allow researchers to manipulate specific acoustic features along the voice production pathway. From the glottal source to the vocal tract, we presented the acoustic consequence of each of the building blocks of the vocal apparatus, as well as corresponding acoustic models and transformation algorithms from the signal-processing literature.

We suggest that using such transformations to control the content of vocal stimuli in experimental research is a promising line of work. This methodology allows researchers to formulate and test computational predictions about the behavioral, physiological and neural consequences of specific acoustic changes, enabling them to draw strong, causal links between the anatomic mechanisms involved in voice production and their subsequent reactions in listeners. Transformation technologies can easily be shared between research groups, made open source, and deployed across several types of studies (e.g. cross-linguistic, cross-species) and auditory modalities, such as non-verbal behaviour, speech and music. When possible, we gave here links to some of these tools that are available freely, as well as examples of studies that uses them, and hope that this list will only be growing.

However, in order for these technologies to be useful in an experimental context, they must deliver the proper type of acoustic control. We identified four of such constraints. First, in our view, transformations should be *specific*, i.e. transform sound in a single acoustic dimension mirroring an isolated anatomic mechanism (e.g. vocal fold saturation or zygomatic muscle contraction). Second, transformations should be *parametric*, i.e. have a uniform acoustic effect regardless

of the original signal (e.g. semantic content, age, identity, sex, species), thus allowing comparative studies. Third, transformations should be *exhaustive*, i.e. unconstrained by what speakers usually produce, but rather by what they *can* produce, in order to reduce sampling bias for psychophysical research. Finally, in an era where the study of social interactions is at the top of the cognitive-science research agenda, the community should favor transformations that can operate in real-time. We hope that these recommendations can be used as a ‘check-list’ for machine-learning and signal-processing researchers involved in creating new vocal transformations. If new tools, included those emerging from the recent trend of deep-learning research, follow these constraints, they will be more easily used for experimental research..

Acknowledgements

P.A., L.R. and M.L. produced the experimental data and software on which this review was based. J.J.A. wrote the manuscript, with contributions from all other authors. Data in Figure 4 collected at the Sorbonne-Universit INSEAD Center for Behavioural Sciences. Work funded by ERC StG 335536 (CREAM) and ANR REFLETS to J.J.A. We thank Gilles Degottex for feedback on Figure 1. Facial expressions in Figure 1 are from the extended Cohn-Kanade dataset (ck+), S106 and S113 (Lucey et al. 2010).

Funding

Work funded by ERC StG CREAM (335536) and ANR REFLETS (2017-2021) to J.J.A.

References

- Akuzawa K, Iwasawa Y and Matsuo Y (2018) Expressive speech synthesis via modeling expressions with variational autoencoder. *arXiv preprint arXiv:1804.02135*.
- Anikin A (2018) Soundgen: An open-source tool for synthesizing nonverbal vocalizations. *Behavior research methods* : 1–15.
- Anikin A (2019) The perceptual effects of manipulating nonlinear phenomena in synthetic nonverbal vocalizations. *Bioacoustics* : 1–22.
- Arias P, Belin P and Aucouturier JJ (2018a) Auditory smiles trigger unconscious facial imitation. *Current Biology* 28(14): R782–R783.
- Arias P, Soladie C, Bouafif O, Robel A, Segulier R and Aucouturier JJ (2018b) Realistic transformation of facial and vocal smiles in real-time audiovisual streams. *IEEE Transactions on Affective Computing*.
- Armstrong MM, Lee AJ and Feinberg DR (2019) 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.
- Arnal LH, Flinker A, Kleinschmidt A, Giraud AL and Poeppel D (2015) Human screams occupy a privileged niche in the communication soundscape. *Current Biology* 25(15): 2051–2056.
- Aucouturier JJ, Johansson P, Hall L, Segnini R, Mercadié L and Watanabe K (2016) Covert digital manipulation of vocal emotion alter speakers emotional states in a congruent direction. *Proceedings of the National Academy of Sciences* 113(4): 948–953.
- Bachorowski JA and Owren MJ (1995) Vocal expression of emotion: Acoustic properties of speech are associated with emotional intensity and context. *Psychological science* 6(4): 219–224.
- Banse R and Scherer KR (1996) Acoustic profiles in vocal emotion expression. *Journal of personality and social psychology* 70(3): 614.
- Barthel H and Quené H (2015) Acoustic-phonetic properties of smiling revised—measurements on a natural video corpus. In: *Proceedings of the 18th International Congress of Phonetic Sciences.—Glasgow, UK: The University of Glasgow*.
- Belin P, Boehme B and McAleer P (2017) The sound of trustworthiness: Acoustic-based modulation of perceived voice personality. *PLoS ONE* 12(10).
- Blumstein DT and Recapet C (2009) The sound of arousal: The addition of novel non-linearities increases responsiveness in marmot alarm calls. *Ethology* 115(11): 1074–1081.
- Boersma P and Weenink D (2002) Praat, a system for doing phonetics by computer. *Glott international* 5.
- Böhm T, Audibert N, Shattuck-Hufnagel S, Németh G and Aubergé V (2008) Transforming modal voice into irregular voice by amplitude scaling of individual glottal cycles. *Journal of the Acoustical Society of America* 123(5): 3886.
- Boidron L, Boudenia K, Avena C, Boucheix JM and Aucouturier JJ (2016) Emergency medical triage decisions are swayed by computer-manipulated cues of physical dominance in callers voice. *Scientific reports* 6: 30219.
- Bonada J and Blaauw M (2013) Generation of growl-type voice qualities by spectral morphing. In: *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, pp. 6910–6914.
- Bonebright TL, Thompson JL and Leger DW (1996) Gender stereotypes in the expression and perception of vocal affect. *Sex Roles* 34(5-6): 429–445.
- Bowling D, Garcia M, Dunn J, Ruprecht R, Stewart A, Frommolt KH and Fitch W (2017) Body size and vocalization in primates and carnivores. *Scientific reports* 7: 41070.

- Brady MC (2005) Synthesizing affect with an analog vocal tract: glottal source. In: *Toward Social Mechanisms of Android Science: A CogSci 2005 Workshop*. pp. 25–26.
- Bruckert L, Bestelmeyer P, Latinus M, Rouger J, Charest I, Rousselet GA, Kawahara H and Belin P (2010) Vocal attractiveness increases by averaging. *Current Biology* 20(2): 116–120.
- Burleigh T, Schoenherr J and Lacroix G (2013) Does the uncanny valley exist? an empirical test of the relationship between eeriness and the human likeness of digitally created faces. *Computers in Human Behavior* 29(3): 759–771.
- Burred JJ, Ponsot E, Goupil L, Liuni M and Aucouturier JJ (2019) Cleese: An open-source audio-transformation toolbox for data-driven experiments in speech and music cognition. *PloS one* 14(4): e0205943.
- Casadevall A and Fang FC (2008) Descriptive science. *Infection and immunity* 76(9): 3835–3836.
- Charlton BD, Ellis WA, Brumm J, Nilsson K and Fitch WT (2012) Female koalas prefer bellows in which lower formants indicate larger males. *Animal Behaviour* 84(6): 1565–1571.
- Chen B, Kitaoka N and Takeda K (2016) Impact of acoustic similarity on efficiency of verbal information transmission via subtle prosodic cues. *EURASIP Journal on Audio, Speech, and Music Processing* 2016(1): 19.
- Chong CS, Kim J and Davis C (2018) Disgust expressive speech: The acoustic consequences of the facial expression of emotion. *Speech Communication* 98: 68–72.
- Dattorro J (1997) Effect design, part 2: Delay line modulation and chorus. *Journal of the Audio engineering Society* 45(10): 764–788.
- De Boer B (2000) Self-organization in vowel systems. *Journal of phonetics* 28(4): 441–465.
- Degottex G, Lanchantin P, Roebel A and Rodet X (2013) Mixed source model and its adapted vocal tract filter estimate for voice transformation and synthesis. *Speech Communication* 55(2): 278–294.
- Drahota A, Costall A and Reddy V (2008) The vocal communication of different kinds of smile. *Speech Communication* 50(4): 278–287.
- Dupoux E, Kakehi K, Hirose Y, Pallier C and Mehler J (1999) Epenthetic vowels in Japanese: A perceptual illusion? *Journal of experimental psychology: human perception and performance* 25(6): 1568.
- El Haddad K, Dupont S, d’Alessandro N and Dutoit T (2015) An hmm-based speech-smile synthesis system: An approach for amusement synthesis. In: *Automatic Face and Gesture Recognition (FG), 2015 11th IEEE International Conference and Workshops on*, volume 5. IEEE, pp. 1–6.
- Eyben F, Weninger F, Gross F and Schuller B (2013) Recent developments in opensmile, the Munich open-source multimedia feature extractor. In: *Proceedings of the 21st ACM international conference on Multimedia*. ACM, pp. 835–838.
- Fee MS, Shraiman B, Pesaran B and Mitra PP (1998) The role of nonlinear dynamics of the syrinx in the vocalizations of a songbird. *Nature* 395(6697): 67.
- Fitch WT (1997) Vocal tract length and formant frequency dispersion correlate with body size in rhesus macaques. *The Journal of the Acoustical Society of America* 102(2): 1213–1222.
- Fitch WT, Neubauer J and Herzel H (2002) Calls out of chaos: the adaptive significance of nonlinear phenomena in mammalian vocal production. *Animal behaviour* 63(3): 407–418.
- Gentilucci M, Ardaillon L and Liuni M (2019) Composing vocal distortion: A tool for real-time generation of roughness. *Computer Music Journal* 42(4).
- Gobl C and Chasaide AN (2010) 11 voice source variation and its communicative functions. *The handbook of phonetic sciences* 1: 378.
- Goldstone RL and Lupyan G (2016) Discovering psychological principles by mining naturally occurring data sets. *Topics in cognitive science* 8(3): 548–568.
- Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A and Bengio Y (2014) Generative adversarial nets. In: *Advances in neural information processing systems*. pp. 2672–2680.
- Govind D and Prasanna SM (2013) Expressive speech synthesis: a review. *International Journal of Speech Technology* 16(2): 237–260.
- Hain TC, Burnett TA, Larson CR and Kiran S (2001) Effects of delayed auditory feedback (daf) on the pitch-shift reflex. *The Journal of the Acoustical Society of America* 109(5): 2146–2152.
- Hienz RD, Jones AM and Weerts EM (2004) The discrimination of baboon grunt calls and human vowel sounds by baboons. *J. Acoust. Soc. Am.* 116: 1692–1697.
- Hsu WN, Zhang Y and Glass J (2017) Learning latent representations for speech generation and transformation. *arXiv preprint arXiv:1704.04222*.
- Ilie G and Thompson WF (2006) A comparison of acoustic cues in music and speech for three dimensions of affect. *Music Perception* 23(4): 319–329.
- Jack RE and Schyns PG (2017) Toward a social psychophysics of face communication. *Annual review of psychology* 68: 269–297.
- Jiang X and Pell MD (2017) The sound of confidence and doubt. *Speech Communication* 88: 106–126.

- Johnstone T and Scherer KR (1999) The effects of emotions on voice quality. In: *Proceedings of the XIVth international congress of phonetic sciences*. Citeseer, pp. 2029–2032.
- Juillerat N, Schubiger-Banz S and Arisona SM (2008) Low latency audio pitch shifting in the time domain. In: *2008 International Conference on Audio, Language and Image Processing*. IEEE, pp. 29–35.
- Juslin PN and Laukka P (2003) Communication of emotions in vocal expression and music performance: Different channels, same code? *Psychological bulletin* 129(5): 770.
- Kamiloglu R, Fischer A and Sauter DA (2019) Good vibrations: A review of vocal expressions of positive emotions .
- Kawahara H (1997) Speech representation and transformation using adaptive interpolation of weighted spectrum: vocoder revisited. In: *Acoustics, Speech, and Signal Processing, 1997. ICASSP-97., 1997 IEEE International Conference on*, volume 2. IEEE, pp. 1303–1306.
- Kawahara H and Matsui H (2003) Auditory morphing based on an elastic perceptual distance metric in an interference-free time-frequency representation. In: *2003 IEEE International Conference on Acoustics, Speech, and Signal Processing, 2003. Proceedings.(ICASSP'03).*, volume 1. IEEE, pp. I–I.
- Kitayama S, Mesquita B and Karasawa M (2006) Cultural affordances and emotional experience: socially engaging and disengaging emotions in japan and the united states. *Journal of personality and social psychology* 91(5): 890.
- Lartillot O and Toiviainen P (2007) A matlab toolbox for musical feature extraction from audio. In: *International conference on digital audio effects*. Bordeaux, pp. 237–244.
- Lasarczyk E and Trouvain J (2008) Spread lips+ raised larynx+ higher f0= smiled speech?-an articulatory synthesis approach. *Proceedings of ISSP* : 43–48.
- Latinus M, McAleer P, Bestelmeyer PE and Belin P (2013) Norm-based coding of voice identity in human auditory cortex. *Current Biology* 23(12): 1075–1080.
- Laukka P (2005) Categorical perception of vocal emotion expressions. *Emotion* 5(3): 277.
- Lee E, Karrer T and Borchers JO (2007) An analysis of startup and dynamic latency in phase vocoder-based time-stretching algorithms. In: *ICMC*. Citeseer.
- Lee Y, Rabiee A and Lee SY (2017) Emotional end-to-end neural speech synthesizer. *arXiv preprint arXiv:1711.05447* .
- Loscos A and Bonada J (2004) Emulating rough and growl voice in spectral domain. In: *Proceedings of the International Conference on Digital Audio Effects*. Citeseer, pp. 49–52.
- Lucey P, Cohn JF, Kanade T, Saragih J, Ambadar Z and Matthews I (2010) The extended cohn-kanade dataset (ck+): A complete dataset for action unit and emotion-specified expression. In: *Computer Vision and Pattern Recognition Workshops (CVPRW), 2010 IEEE Computer Society Conference on*. IEEE, pp. 94–101.
- Luo Z, Chen J, Takiguchi T and Ariki Y (2017) Emotional voice conversion with adaptive scales f0 based on wavelet transform using limited amount of emotional data. In: *INTERSPEECH*. pp. 3399–3403.
- Ma W and Thompson WF (2015) Human emotions track changes in the acoustic environment. *Proceedings of the National Academy of Sciences* 112(47): 14563–14568.
- Malisz Z, Henter GE, Valentini-Botinhao C, Watts O, Beskow J and Gustafson J (2019) Modern speech synthesis for phonetic sciences: A discussion and an evaluation. In: *Proceedings of ICPhS*.
- Matsumoto D (2002) Methodological requirements to test a possible in-group advantage in judging emotions across cultures: Comment on elfenbein and ambady (2002) and evidence. *Psychological Bulletin* 128(2): 236–242.
- Mohammadi SH and Kain A (2017a) An overview of voice conversion systems. *Speech Communication* 88: 65–82.
- Mohammadi SH and Kain A (2017b) An overview of voice conversion systems. *Speech Communication* 88: 65–82.
- Moulines E and Laroche J (1995) Non-parametric techniques for pitch-scale and time-scale modification of speech. *Speech communication* 16(2): 175–205.
- Narendranath M, Murthy HA, Rajendran S and Yegnanarayana B (1995) Transformation of formants for voice conversion using artificial neural networks. *Speech communication* 16(2): 207–216.
- Nowicki S, Mitani JC, Nelson DA and Marler P (1989) The communicative significance of tonality in birdsong: responses to songs produced in helium. *Bioacoustics* 2(1): 35–46.
- Ohala JJ (1980) The acoustic origin of the smile. *The Journal of the Acoustical Society of America* 68(S1): S33–S33.
- Oord Avd, Dieleman S, Zen H, Simonyan K, Vinyals O, Graves A, Kalchbrenner N, Senior A and Kavukcuoglu K (2016) Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499* .
- Pascual S, Bonafonte A and Serra J (2017) Segan: Speech enhancement generative adversarial network. *arXiv preprint arXiv:1703.09452* .
- Pell MD (2001) Influence of emotion and focus location on prosody in matched statements and questions. *The Journal of the Acoustical Society of America* 109(4): 1668–1680.
- Pell MD and Kotz SA (2011) On the time course of vocal emotion recognition. *PLoS One* 6(11): e27256.
- Pell MD, Monetta L, Paulmann S and Kotz SA (2009) Recognizing emotions in a foreign language. *Journal of Nonverbal Behavior*

- 33(2): 107–120.
- Pierce BC and Benjamin C (2002) *Types and programming languages*. MIT press.
- Ponsot E, Burred JJ, Belin P and Aucouturier JJ (2018) Cracking the social code of speech prosody using reverse correlation. *Proceedings of the National Academy of Sciences* 115(15): 3972–3977.
- Quené H, Semin GR and Foroni F (2012) Audible smiles and frowns affect speech comprehension. *Speech Communication* 54(7): 917–922.
- Rachman L, Liuni M, Arias P, Lind A, Johansson P, Hall L, Richardson D, Watanabe K, Dubal S and Aucouturier JJ (2018) David: An open-source platform for real-time transformation of infra-segmental emotional cues in running speech. *Behavior research methods* 50(1): 323–343.
- Rand AS and Dudley R (1993) Frogs in helium: the anuran vocal sac is not a cavity resonator. *Physiological Zoology* 66(5): 793–806.
- Reby D, McComb K, Cargnelutti B, Darwin C, Fitch WT and Clutton-Brock T (2005) Red deer stags use formants as assessment cues during intrasexual agonistic interactions. *Proceedings of the Royal Society B: Biological Sciences* 272(1566): 941–947.
- Ruinskiy D and Lavner Y (2008) Stochastic models of pitch jitter and amplitude shimmer for voice modification. In: *Electrical and Electronics Engineers in Israel, 2008. IEEEI 2008. IEEE 25th Convention of. IEEE*, pp. 489–493.
- Scherer KR (1972) Acoustic concomitants of emotional dimensions: Judging affect from synthesized tone sequences. In: *Eastern Psychological Association Meeting, Boston, MA*.
- Scherer KR (2003) Vocal communication of emotion: A review of research paradigms. *Speech communication* 40(1-2): 227–256.
- Scherer KR, Banse R and Wallbott HG (2001) Emotion inferences from vocal expression correlate across languages and cultures. *Journal of Cross-cultural psychology* 32(1): 76–92.
- Scherer KR and Oshinsky JS (1977) Cue utilization in emotion attribution from auditory stimuli. *Motivation and emotion* 1(4): 331–346.
- Seshadri S, Juvela L, Alku P and Räsänen O (2019) Augmented cyclegrams for continuous scale normal-to-lombard speaking style conversion. *Proc. Interspeech 2019* : 2838–2842.
- Simonyan K and Horwitz B (2011) Laryngeal motor cortex and control of speech in humans. *The Neuroscientist* 17(2): 197–208.
- Stylianou Y (2009) Voice transformation: a survey. In: *Acoustics, Speech and Signal Processing, 2009. ICASSP 2009. IEEE International Conference on. IEEE*, pp. 3585–3588.
- Sueur J, Aubin T and Simonis C (2008) Seewave, a free modular tool for sound analysis and synthesis. *Bioacoustics* 18(2): 213–226.
- Sun P, Anumanchipalli GK and Chang EF (2019) Brain2char: A deep architecture for decoding text from brain recordings. *arXiv preprint arXiv:1909.01401* .
- Tartter VC (1980) Happy talk: Perceptual and acoustic effects of smiling on speech. *Perception & psychophysics* 27(1): 24–27.
- Tartter VC and Braun D (1994) Hearing smiles and frowns in normal and whisper registers. *The Journal of the Acoustical Society of America* 96(4): 2101–2107.
- Tchernichovski O and Mitra P (2004) Sound analysis pro user manual. *CCNY, New York* .
- Tice DM (1992) Self-concept change and self-presentation: The looking glass self is also a magnifying glass. *Journal of personality and social psychology* 63(3): 435.
- Titze I (1994) *Principles of Voice Production*. Prentice Hall.
- Tyson RB, Nowacek DP and Miller PJ (2007) Nonlinear phenomena in the vocalizations of north atlantic right whales (*eubalaena glacialis*) and killer whales (*orcinus orca*). *The Journal of the Acoustical Society of America* 122(3): 1365–1373.
- Van Doorn EA, Heerdink MW and Van Kleef GA (2012) Emotion and the construal of social situations: Inferences of cooperation versus competition from expressions of anger, happiness, and disappointment. *Cognition & emotion* 26(3): 442–461.
- Verma A and Kumar A (2005) Introducing roughness in individuality transformation through jitter modeling and modification. In: *Proceedings.(ICASSP'05). IEEE International Conference on Acoustics, Speech, and Signal Processing, 2005.*, volume 1. IEEE, pp. I–5.
- Wang W, Alameda-Pineda X, Xu D, Fua P, Ricci E and Sebe N (2018) Every smile is unique: Landmark-guided diverse smile generation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 7083–7092.
- Wang Y, Skerry-Ryan R, Stanton D, Wu Y, Weiss RJ, Jaitly N, Yang Z, Xiao Y, Chen Z, Bengio S et al. (2017) Tacotron: Towards end-to-end speech synthesis. *arXiv preprint arXiv:1703.10135* .
- Wilden I, Herzel H, Peters G and Tembrock G (1998) Subharmonics, biphonation, and deterministic chaos in mammal vocalization. *Bioacoustics* 9(3): 171–196.
- Wu YC, Hayashi T, Tobing PL, Kobayashi K and Toda T (2019) Quasi-periodic wavenet vocoder: a pitch dependent dilated convolution model for parametric speech generation. *arXiv preprint arXiv:1907.00797* .
- Xu T, Zhan J, Garrod OG, Torr PH, Zhu SC, Ince RA and Schyns PG (2018) Deeper interpretability of deep networks. *arXiv*

preprint arXiv:1811.07807 .

Yoshie M and Sauter DA (2019) Cultural norms influence nonverbal emotion communication: Japanese vocalizations of socially disengaging emotions. *Emotion* .