

LEARNING THE HELIX TOPOLOGY OF MUSICAL PITCH

Vincent Lostanlen^{*†} Sripathi Sridhar[†] Brian McFee[†] Andrew Farnsworth^{*} Juan Pablo Bello[†]

^{*} Cornell Lab of Ornithology, Cornell University, Ithaca, NY, USA

[†] Music and Audio Research Laboratory, New York University, New York, NY, USA

ABSTRACT

To explain the consonance of octaves, music psychologists represent pitch as a helix where azimuth and axial coordinate correspond to pitch class and pitch height respectively. This article addresses the problem of discovering this helical structure from unlabeled audio data. We measure Pearson correlations in the constant- Q transform (CQT) domain to build a K -nearest neighbor graph between frequency subbands. Then, we run the Isomap manifold learning algorithm to represent this graph in a three-dimensional space in which straight lines approximate graph geodesics. Experiments on isolated musical notes demonstrate that the resulting manifold resembles a helix which makes a full turn at every octave. A circular shape is also found in English speech, but not in urban noise. We discuss the impact of various design choices on the visualization: instrumentarium, loudness mapping function, and number of neighbors K .

Index Terms— Continuous wavelet transforms, distance learning, music, pitch control (audio), shortest path problem.

1. INTRODUCTION

Listening to a sequence of two pure tones elicits a sensation of pitch going “up” or “down”, correlating with changes in fundamental frequency (f_0). However, contrary to pure tones, natural pitched sounds contain a rich spectrum of components in addition to f_0 . Neglecting inharmonicity, these components are tuned to an ideal Fourier series whose modes resonate at integer multiples of the fundamental: $2f_0$, $3f_0$, and so forth. By the change of variable $f'_0 = 2f_0$, it appears that all even-numbered partials $2f_0 = f'_0$, $4f_0 = 2f'_0$, $6f_0 = 3f'_0$, and so forth make up the whole Fourier series of a periodic signal whose fundamental frequency is f'_0 . Figure 1 illustrates, in the case of a synthetic signal with perfect harmonicity, that filtering out all odd-numbered partials $(2p + 1)f_0$ for integer $p \geq 0$ results in a perceived pitch that morphs from f_0 to $2f_0$, i.e., up one octave [1].

Such ambivalence brings about well-known auditory paradoxes: tones that have a definite pitch class but lack a pitch register [2]; glissandi that seem to ascend or descend endlessly [3]; and tritone intervals whose pitch directionality depends on prior context [4]. To explain them, one may roll up the frequency axis onto a spiral or helix which makes a full turn at each octave, thereby aligning power-of-two harmonics onto the same radii.

The consonance of octaves has implications in several disciplines. In music theory, it allows for pitches to be grouped into pitch classes: e.g., in European solfège, *do re mi* etc. eventually “circle back” to *do*. In ethnomusicology, it transcends the boundaries between cultures, to the point of being held as a nearly universal attribute of music [5]. In neurophysiology, it explains the functional organization of the central auditory cortex [6]. In music information research, it motivates the

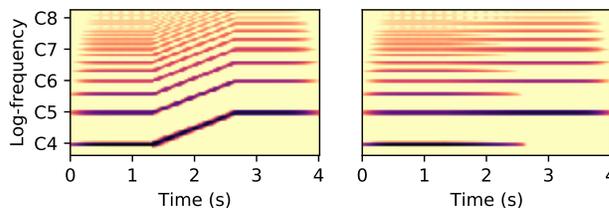


Fig. 1. Two continuous trajectories in pitch space, either by octave glissando (left) or by attenuation of odd-numbered partials (right). Darker shades indicate larger magnitudes of the constant- Q transform. Vertical ticks denote octaves. See Section 1 for details.

design of chroma features, i.e., a representation of harmonic content that is meant to be equivariant to parallel chord progressions but invariant to chord inversions [7, chapter 5].

Despite the wealth of evidence for the crucial role of octaves in music, there is, to this day, no data-driven criterion for assessing whether a given audio corpus exhibits a property of octave equivalence. Rather, the disentanglement of pitch chroma and pitch height in time–frequency representations relies on domain-specific knowledge about music [8]. More generally, although the induction of priors on the topology of pitch is widespread in symbolic music analysis [9], few of them directly apply to audio signal processing [10].

Yet, in recent years, the systematic use of machine learning methods has progressively reduced the need for domain-specific knowledge in several other aspects of auditory perception, including mel-frequency spectrum [11] and adaptive gain control [12]. It remains to be known whether octave equivalence can, in turn, be discovered by a machine learning algorithm, instead of being engineered ad hoc. Furthermore, it is unclear whether octave equivalence is an exclusive characteristic of music or whether it may extend to other kinds of sounds, such as speech or environmental soundscapes.

In this article, we conduct an unsupervised manifold learning experiment, inspired by the protocol of “learning the 2-D topology of images” [13], to visualize the helix topology of constant- Q spectra. Starting from a dataset of isolated notes from various musical instruments, we run the Isomap algorithm to represent each of these frequency subbands as a dot in a 3-D space, wherein spatial neighborhoods denote high correlation in loudness. Contrary to natural images, we find a mismatch between the physical dimensionality of natural acoustic spectra (i.e., 1-D) and their statistical dimensionality (i.e., 3-D or greater).

The companion website of this paper¹ contains a Python package to visualize octave equivalence in audio data.

This work is partially supported by NSF award 1633259 (BirdVox).

¹Companion website: <https://github.com/BirdVox/lostanlen2020icassp>

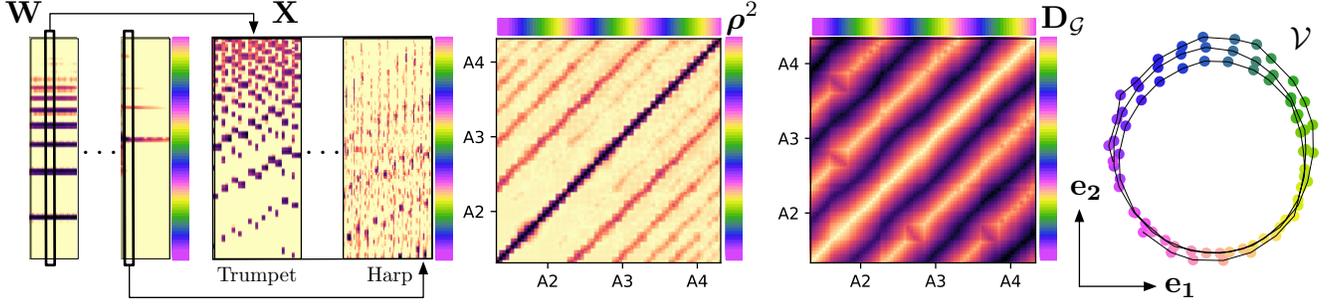


Fig. 2. Functional diagram of the proposed method, comprising: constant- Q transform \mathbf{W} , data matrix \mathbf{X} , extraction of Pearson correlations ρ^2 , shortest path distance matrix \mathbf{D}_G , and Isomap eigenbasis \mathbf{e} for the set of vertices \mathcal{V} . Darker shades in ρ^2 and \mathbf{D}_G indicate larger absolute values of the Pearson correlation and distance respectively. The hue of colored dots and the solid grey line respectively denote pitch chroma and pitch height. The first three dimensions in the Isomap embedding explain 36%, 35%, and 9% of the total variance in \mathbf{D}_G respectively. Note that our method is unsupervised: neither pitch chroma nor pitch height are directly supplied to the Isomap manifold learning algorithm. See Section 2 for details.

2. ISOMAP EMBEDDING OF SUBBAND CORRELATIONS

2.1. Constant- Q transform and loudness mapping

Given a corpus of audio signals $\mathbf{x}_1 \dots \mathbf{x}_N \in \mathbb{L}^2(\mathbb{R})$ in arbitrary order, we define their constant- Q transforms (CQT) as the convolutions with wavelets $t \mapsto 2^\gamma \psi(2^\gamma t)$, where $2^{-\gamma}$ is a scale parameter:

$$\mathbf{W}\mathbf{x}_n : (t, \gamma) \mapsto 2^\gamma \int_{-\infty}^{+\infty} \mathbf{x}_n(t') \psi(2^\gamma(t - t')) dt. \quad (1)$$

The wavelet filterbank in the operator \mathbf{W} covers a range of J octaves. In the case of musical notes, we define a region of interest t_n in each \mathbf{x}_n as the frame of highest short-term energy at a scale of 93 ms. We compute their scalogram representation as the vector of CQT modulus responses at some $t = t_n$, for discretized values $u = 1 + Q\gamma$:

$$\mathbf{X}[n, u] = \left| \mathbf{W}\mathbf{x}_n \left(t_n, \frac{u-1}{Q} \right) \right|, \quad (2)$$

where the integer u ranges from 1 to QJ . Then, we apply a point-wise logarithmic compression Λ to map each magnitude coefficient $\mathbf{X}[n, u]$ onto a decibel-like scale, and clip it to -100 dB:

$$\Lambda(\mathbf{X})[n, u] = \max(-100, 10 \log_{10} \mathbf{X}[n, u]). \quad (3)$$

Unless stated otherwise, we set $|\psi|$ to a Hann window, the quality factor Q to 24, and the number of octaves J to 3 in the following. We compute constant- Q transforms with librosa v0.6.1 [14]. Section 3.4 will discuss the effect of alternative choices for Λ .

2.2. Pearson autocorrelation of log-magnitude spectra

In accordance with [13], we express the similarity between two features u and v in terms of their squared Pearson correlation $\rho^2[u, v]$. We begin by recentering each feature to null mean, yielding the matrix

$$\mathbf{Y}[n, u] = \Lambda(\mathbf{X})[n, u] - \frac{1}{N} \sum_{n=1}^N \Lambda(\mathbf{X})[n, u], \quad (4)$$

and then compute squared cosine similarities on all pairs (u, v) :

$$\rho^2[u, v] = \frac{(\sum_{n=1}^N \mathbf{Y}[n, u] \mathbf{Y}[n, v])^2}{(\sum_{n=1}^N \mathbf{Y}^2[n, u]) \times (\sum_{n=1}^N \mathbf{Y}^2[n, v])}. \quad (5)$$

Let \mathcal{V} be the set of all QJ features u in \mathbf{Y} . Adopting a manifold learning perspective, we may regard ρ^2 as the values of a standard Gaussian kernel $\kappa : \mathcal{V} \times \mathcal{V} \rightarrow \mathbb{R}$. Inverting the identity $\rho^2[u, v] = \kappa(u, v) = \exp(-2\mathbf{D}_{\rho^2}[u, v])$ leads to a pseudo-Euclidean distance

$$\mathbf{D}_{\rho^2}[u, v] = \sqrt{-\frac{1}{2} \log \rho^2[u, v]}. \quad (6)$$

2.3. Shortest path distance on the K -nearest neighbor graph

Following the Isomap manifold learning algorithm [15], we compute the K nearest neighbors of each feature $u \in \mathcal{V}$ as the set $\mathcal{N}_K(u)$ of K features $v \neq u$ that minimize the distance $\mathbf{D}_{\rho^2}[u, v]$, i.e., maximize the squared Pearson correlation $\rho^2[u, v]$. We construct a K -nearest neighbor graph \mathcal{G} whose vertices are \mathcal{V} and whose adjacency matrix $\mathbf{A}[u, v]$ is equal to $\mathbf{D}_{\rho^2}[u, v]$ if $v \in \mathcal{N}_K(u)$ or if $u \in \mathcal{N}_K(v)$, and infinity otherwise. We run Dijkstra's algorithm on \mathcal{G} to measure geodesics on the manifold induced by \mathbf{A} . These geodesics yield a shortest path distance function over \mathcal{V}^2 :

$$\mathbf{D}_G[u, v] = \begin{cases} \mathbf{D}_{\rho^2}[u, v] & \text{if } v \in \mathcal{N}_K(u) \text{ or } u \in \mathcal{N}_K(v) \\ \min_z (\mathbf{D}_G[u, z] + \mathbf{D}_G[z, v]) & \text{otherwise.} \end{cases} \quad (7)$$

If \mathcal{G} has more than one connected component, then \mathbf{D}_G is infinite over pairs of mutually unreachable vertices. On some datasets, ρ may lead to a disconnected K -nearest neighbor graph \mathcal{G} , especially for small K , thereby causing numerical aberrations in Isomap. However, we make sure that the effective bandwidth of the wavelet $\hat{\psi}$ is large enough in comparison with Q to yield strong correlations $\rho[u, u+1]$ for all u , and thus $u \in \mathcal{N}_K(u+1)$. With this caveat in mind, we postulate that \mathbf{D}_G is finite in all of the following. We set $K = 3$ in the following unless stated otherwise.

2.4. Classical multidimensional scaling and 3-D embedding

Let $\mathbf{S}_G[u, v] = \mathbf{D}_G[u, v]^2$. Classical multidimensional scaling (MDS) diagonalizes $\tau_G = -\frac{1}{2}(\mathbf{H}\mathbf{S}_G\mathbf{H})$ where $\mathbf{H}[u, v] = \delta(u - v) - \frac{1}{QJ}$ [16]. We denote by \mathbf{e}_m and λ_m the respective eigenvectors and eigenvalues of τ_G , satisfying $\tau_G \mathbf{e}_m = \lambda_m \mathbf{e}_m$. We rank eigenvalues in decreasing order, without loss of generality. Lastly, we display the Isomap embedding as a 3-D scatter plot with Cartesian coordinates $\mathbf{e}[u] = (\mathbf{e}_1[u], \mathbf{e}_2[u], \mathbf{e}_3[u])$ for every vertex $u \in \mathcal{V}$. We compute Isomap using scikit-learn v0.21.3 [17].

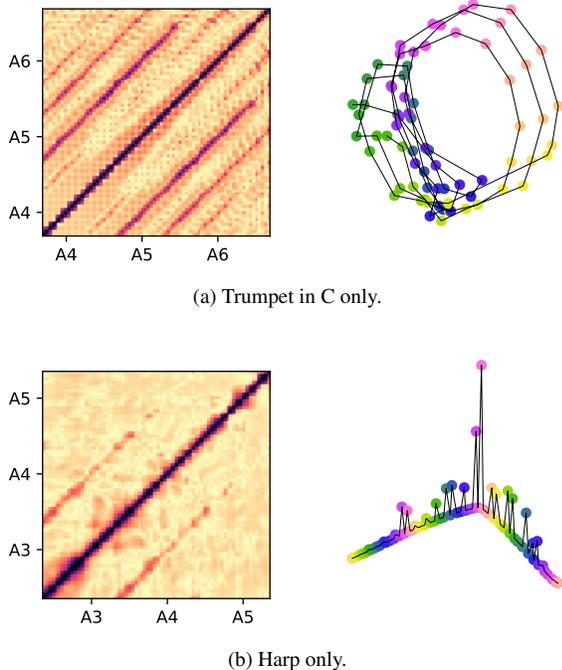


Fig. 3. Pearson correlation matrices ρ^2 (left) and Isomap embeddings (right) for two instruments in the SOL dataset: Trumpet in C (a) and Harp (b). In both cases, the number of neighbors is set to $K = 3$ and the loudness mapping Λ is logarithmic. See Section 3.2 for details.

3. EXPERIMENTS WITH MUSICAL SOUNDS

3.1. Main protocol

We extract \mathbf{X} from the isolated musical notes played by eight instruments in the SOL dataset [18]: accordion, alto saxophone, bassoon, flute, harp, trumpet in C, and cello. For each of these instruments, we include three levels of intensity dynamics: *pp*, *mf*, and *ff*. We include all available pitches in the tessitura (min = B_0 , median = E_4 , max = $C\sharp_8$). We exclude extended playing techniques, because some of them may lack a discernible pitch class [19]; thus resulting in a total of $N = 1212$ audio recordings for the *ordinario* technique. We hypothesize that cross-correlations across octaves are weaker than cross-correlations along the log-frequency axis. Because the helix topology of musical pitch relies on both kinds of correlation, we set the number of neighbors to $K = 3$.

Figure 2 illustrates our protocol and main finding. Isomap produces a quasi-perfect cylindrical manifold in dimension three. We color each dot $\mathbf{e}[u]$ according to a hue of $\theta(u) = \frac{2\pi}{Q}(u \bmod Q)$ radians. Furthermore, we draw a segment between each $\mathbf{e}[u]$ and its upper adjacent subband $\mathbf{e}[u + 1]$. Once these visual elements are included in the display of the scatter plot $(\mathbf{e}[u])_u$, the cylindrical manifold appears to coincide with the Drobisch-Shepard helix in music psychology [20, 21]. Indeed, hues $\theta[u]$ appear to align on the same radii, whereas the grey line γ grows monotonically with \mathbf{e}_3 . This result demonstrates that, even without prior knowledge about the perception of musical pitch, it is possible to discover octave equivalence in a purely data-driven fashion, by computing the graph of greatest cross-correlations between CQT magnitudes in a corpus of isolated musical notes.

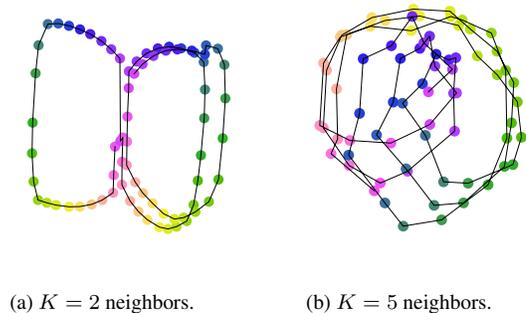


Fig. 4. Isomap embeddings for varying K -nearest neighbor graphs: $K = 2$ (left) and $K = 5$ (right). In both cases, the Pearson correlation matrix ρ^2 results from all instruments in the SOL dataset and the loudness mapping Λ is logarithmic. See Section 3.3 for details.

3.2. Varying the instrumentarium

We reproduce the main protocol on subsets of the SOL dataset, involving a single instrument at once. The bright timbre of brass instruments correlates with relatively loud high-order partials [22], resulting in large octave equivalence and a helical topology (see Figure 3 (a)). In contrast, harp tones carry little energy at twice the fundamental; yet, they induce sympathetic resonance along the soundboard, which affects nearby strings predominantly [23]. These two phenomena in combination favor semitone correlations over octave correlations, resulting in a rectilinear topology (see Figure 3 (a)).

3.3. Varying the number of neighbors K

We reproduce the main protocol with varying K -nearest neighbor graphs. Setting $K = 2$ results in multiple lobes, each corresponding to an octave, and connected at a single pitch class (see Figure 4). This confirms our hypothesis that large semitone correlations outnumber large octave correlations. Conversely, setting $K = 5$ results in a topology that is more intricate than the Drobisch-Shepard helix, involving correlations across perfect fourths and fifths..

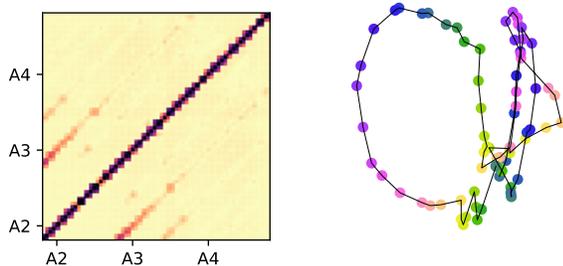
3.4. Varying the loudness mapping function Λ

We reproduce the main protocol with varying loudness mappings Λ (see Figure 5). On one hand, setting Λ to the identity function yields a K -nearest neighbor graph in which octave correlations are numerically negligible, except in the lower register. This results in an Isomap embedding which is circular in the bottommost octave and irregular in the topmost octaves. On the other hand, setting Λ to the cubic root function yields a helical topology. This experiment demonstrates the need for nonlinear loudness compression in the protocol; in contrast with [13], which relied on raw grayscale intensities in handwritten digits.

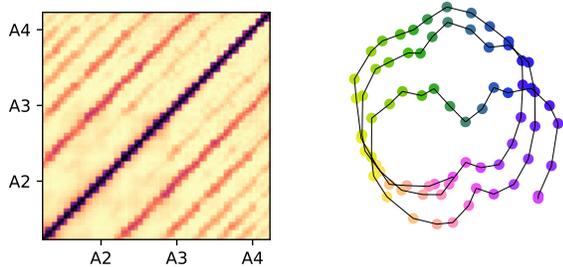
4. EXTENSION TO SPEECH AND URBAN SOUNDS

4.1. Experiment with speech data

We analyze the North Texas vowel database (NTVOW), which contains utterances of 12 English vowels from 50 American speakers, including children aged three to seven as well as male and female adults [24]; resulting in a total of $N = 9570$ audio recordings. As



(a) Linear loudness mapping function.



(b) Cubic root loudness mapping function.

Fig. 5. Pearson correlation matrices ρ^2 (left) and Isomap embeddings (right) for cubic root (a) and linear (b) loudness mappings. In both cases, ρ^2 results from all instruments in the SOL dataset and the number of neighbors is set to $K = 3$. See Section 3.4 for details.

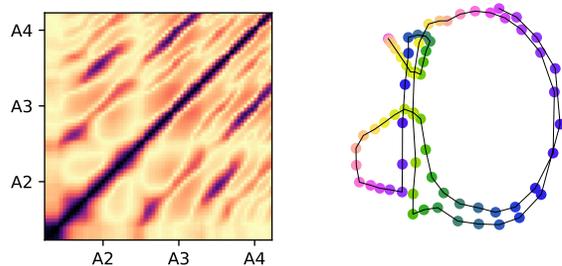
seen in Figure 6 (a), which includes the data from all age groups, there is some notion of pitch circularity seen in the mid-frequency range, but not so in the low and high-frequency ranges. This is because the distribution of f_0 in human speech is polarized around certain pitch classes, rather than uniform over the chromatic scale.

4.2. Experiment with environmental audio data

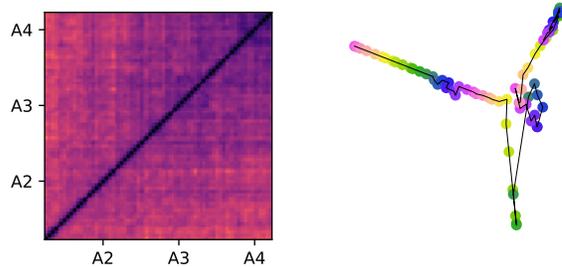
We analyze a portion of the SONYC Urban Sound Tagging dataset (SONYC-UST v0.4), which contains a collection of 3068 acoustic scenes from a network of autonomous sensors in various locations of New York City [25]. We restrict this collection to the acoustic scenes in which the consensus of expert annotators has confirmed the absence of both human speech and music. As a result of this preliminary curation step, we obtain $N = 233$ audio recordings from eight different sensor locations. Each of these scenes contains one or several sources of urban noise pollution, among which: engines, machinery and non-machinery impacts, powered saws, alert signals, and dog barks. Figure 6 (b) shows that no discernible correlations across octaves are observed in this dataset. This finding confirms the conclusions of a previous publication [26], which stated that “music audio signal processing techniques must be informed by a deep and thorough insight into the nature of music itself”.

5. CONCLUSION

The Isomap manifold learning algorithm offers an approximate visualization of nearest neighbor relationships in any non-Euclidean metric space. Thus, a previous publication [13] proposed to apply



(a) Spoken English vowels (NTVOW dataset).



(b) Urban sounds (SONYC-UST excluding speech and music).

Fig. 6. Extension of the proposed method to speech (a) and environmental soundscapes (b). In both cases, the number of neighbors is set to $K = 3$ and the loudness mapping Λ is logarithmic. See Section 4 for details.

Isomap onto cross-correlations between grayscale intensities in natural images. In this article, we have borrowed from the protocol of this publication to apply it on music data.

Despite their methodological resemblance, the two studies lead to different insights. While [13] recovered a quasi-uniform raster, we do not recover a straight line from cross-correlations along the log-frequency axis. Instead, we obtain a cylindrical lattice in dimension three. Assigning pitch classes to subbands reveals that this lattice is akin to a Drobisch-Shepard helix [20], which makes a full turn at each octave. Thus, whereas [13] learned a 2-D topology from 2-D data, we learned a 3-D topology from 1-D data. Furthermore, after benchmarking several design choices, we deduce that the most regular helical shape results from: a diverse instrumentarium; a graph of $K = 3$ nearest neighbors; and a logarithmic loudness mapping. Lastly, we have discussed the limitations of our findings: although spoken vowels also exhibit a quasi-helical topology in subband neighborhoods, the same cannot be said of urban noise.

Beyond the realm of manifold learning, the present article motivates the development of weight sharing architectures that foster octave equivalence in deep representations of music data. Three examples of such architectures are: spiral scattering transform [27]; spiral convolutional networks [28]; and harmonic constant- Q transform [29]. Future work will extend this protocol to bioacoustic data, and comparing the influence of species-specific vocalizations onto the empirical topology of the frequency domain.

6. ACKNOWLEDGMENT

The authors wish to thank Stéphane Mallat for helpful discussions.

7. REFERENCES

- [1] Diana Deutsch, Kevin Dooley, and Trevor Henthorn, "Pitch circularity from tones comprising full harmonic series," *Journal of the Acoustical Society of America*, vol. 124, no. 1, pp. 589–597, 2008.
- [2] Diana Deutsch, "The paradox of pitch circularity," *Acoustics Today*, vol. 7, pp. 8–15, 2010.
- [3] Jean-Claude Risset, "Pitch control and pitch paradoxes demonstrated with computer-synthesized sounds," *Journal of the Acoustical Society of America*, vol. 46, no. 1A, pp. 88–88, 1969.
- [4] Claire Pelofi, Vincent de Gardelle, Paul Egré, and Daniel Pressnitzer, "Interindividual variability in auditory scene analysis revealed by confidence judgements," *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 372, no. 1714, pp. 20160107, 2017.
- [5] Edward M. Burns, "Intervals, scales, and tuning," in *The Psychology of Music*, pp. 215–264. Elsevier, 1999.
- [6] Jason D. Warren, Stefan Uppenkamp, Roy D. Patterson, and Timothy D. Griffiths, "Separating pitch chroma and pitch height in the human brain," *Proceedings of the National Academy of Sciences*, vol. 100, no. 17, pp. 10038–10042, 2003.
- [7] Meinard Müller, *Fundamentals of music processing: Audio, analysis, algorithms, applications*, Springer, 2015.
- [8] Dmitri Tymoczko, "The geometry of musical chords," *Science*, vol. 313, no. 5783, pp. 72–74, 2006.
- [9] Louis Bigo and Moreno Andreatta, "Topological structures in computer-aided music analysis," in *Computational Music Analysis*, David Meredith, Ed., pp. 57–80. Springer International Publishing, Cham, 2016.
- [10] Ching-Hua Chuan and Elaine Chew, "Polyphonic audio key finding using the spiral array ceg algorithm," in *2005 IEEE International Conference on Multimedia and Expo. IEEE*, 2005, pp. 21–24.
- [11] Neil Zeghidour, Nicolas Usunier, Iasonas Kokkinos, Thomas Schaiz, Gabriel Synnaeve, and Emmanuel Dupoux, "Learning filterbanks from raw speech for phone recognition," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 5509–5513.
- [12] Yuxuan Wang, Pascal Getreuer, Thad Hughes, Richard F. Lyon, and Rif A. Saurous, "Trainable frontend for robust and far-field keyword spotting," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 5670–5674.
- [13] Nicolas Le Roux, Yoshua Bengio, Pascal Lamblin, Marc Joliveau, and Balázs Kégl, "Learning the 2-D topology of images," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2008, pp. 841–848.
- [14] Brian McFee, Matt McVicar, Stefan Balke, Carl Thomé, Colin Raffel, Dana Lee, Oriol Nieto, Eric Battenberg, Dan Ellis, Ryuichi Yamamoto, Josh Moore, Rachel Bittner, Keunwoo Choi, Pius Friesch, Fabian-Robert Stöter, Vincent Lostanlen, Siddhartha Kumar, Simon Waloschek, Seth Kranzler, Rimvydas Naktinis, Douglas Repetto, Curtis "Fjord" Hawthorne, CJ Carr, Waldir Pimenta, Petr Viktorin, Paul Brossier, Joao Felipe Santos, Jackie Wu, Erik Peterson, and Adrian Holovaty, "librosa/librosa: 0.6.1," May 2018.
- [15] Joshua B. Tenenbaum, Vin De Silva, and John C. Langford, "A global geometric framework for nonlinear dimensionality reduction," *Science*, vol. 290, no. 5500, pp. 2319–2323, 2000.
- [16] Warren S. Torgerson, "Multidimensional scaling: I. theory and method," *Psychometrika*, vol. 17, no. 4, pp. 401–419, 1952.
- [17] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [18] Guillaume Ballet, Riccardo Borghesi, Peter Hoffmann, and Fabien Lévy, "Studio Online 3.0: An Internet "killer application" for remote access to Ircam sounds and processing tools," in *Actes des Journées d'Informatiques Musicales (JIM)*, 1999.
- [19] V. Lostanlen, J. Andén, and M. Lagrange, "Extended playing techniques: the next milestone in musical instrument recognition," in *Proceedings of the International Conference on Digital Libraries for Musicology (DLfM)*, 2018.
- [20] Roger N. Shepard, "Circularity in judgments of relative pitch," *Journal of the Acoustical Society of America*, vol. 36, no. 12, pp. 2346–2353, 1964.
- [21] Fred Lerdahl, *Tonal pitch space*, Oxford University Press, 2004.
- [22] Emilie Poirson, Jean-François Petiot, and Joël Gilbert, "Study of the brightness of trumpet tones," *Journal of the Acoustical Society of America*, vol. 118, no. 4, pp. 2656–2666, 2005.
- [23] Jean-Loïc Le Carrou, François Gautier, and Roland Badeau, "Sympathetic string modes in the concert harp," *Acta Acustica united with Acustica*, vol. 95, no. 4, pp. 744–752, 2009.
- [24] Peter F. Assmann and William F. Katz, "Time-varying spectral change in the vowels of children and adults," *Journal of the Acoustical Society of America*, vol. 108, no. 4, pp. 1856–1866, 2000.
- [25] Mark Cartwright, Ana Elisa Mendez Mendez, Graham Dove, Jason Cramer, Vincent Lostanlen, Ho-Hsiang Wu, Justin Salamon, Oded Nov, and Juan Pablo Bello, "SONYC Urban Sound Tagging (SONYC-UST): A multilabel dataset from an urban acoustic sensor network," in *Proceedings of the International Workshop on Detection and Classification of Acoustic Scenes and Events (DCASE)*, 2019.
- [26] Meinard Müller, Daniel P. W. Ellis, Anssi Klapuri, and Gaël Richard, "Signal processing for music analysis," *IEEE Journal of Selected Topics in Signal Processing*, vol. 5, no. 6, pp. 1088–1110, 2011.
- [27] Vincent Lostanlen and Stéphane Mallat, "Wavelet scattering on the pitch spiral," in *Proceedings of the International Conference on Digital Audio Effects (DAFX)*, 2016, pp. 429–432.
- [28] Vincent Lostanlen and Carmine Emanuele Cella, "Deep convolutional networks on the pitch spiral for musical instrument recognition," in *Proceedings of the International Society on Music Information Retrieval (ISMIR) Conference*, 2016.
- [29] Rachel Bittner, Brian McFee, Justin Salamon, Peter Li, and Juan Pablo Bello, "Deep saliency representations for f_0 estimation in polyphonic music," in *Proceedings of the International Society on Music Information Retrieval (ISMIR) Conference*, 2017.