Phantom Curves

Scientific Discovery through Interactive Music Visualization

FABIAN C. MOSS, Universiteit van Amsterdam, The Netherlands
GIOVANNI AFFATATO, Politecnico di Milano, Italy
DANIEL HARASIM, École Polytechnique Fédérale de Lausanne, Switzerland

We introduce *phantom curves*, a novel music-theoretical concept based on the discrete Fourier transform (DFT), and document the creative process that led to their discovery. In particular, we emphasize the importance of interactive web applications for music visualization and analysis. This is demonstrated using the example of the application midiVERTO which affords interactions with the pitch-class content of musical pieces encoded in MIDI format without requiring in-depth understanding of the underlying mathematics. We illustrate the analytical value of studying families of phantom curves by applying the concept to music from a Broadway musical, a video game, and a Hollywood movie. This process of discovery thus testifies to the fact that digital tools can bridge disciplinary boundaries between music theory and mathematics, and this interaction can generate new scientific knowledge.

CCS Concepts: • **Applied computing → Sound and music computing**; **Interactive learning environments**; • **Information systems → Music retrieval**; • **Human-centered computing → Visualization toolkits**; **Scientific visualization**.

Additional Key Words and Phrases: music visualization, web application, scientific discovery, discrete Fourier transform

ACM Reference Format:
Fabian C. Moss, Giovanni Affatato, and Daniel Harasim. 2022. Phantom Curves: Scientific Discovery through Interactive Music Visualization. In *DLfM 2022: 9th International Conference on Digital Libraries for Musicology, 28–29 July, 2022, Prague*. ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3543882.3543886

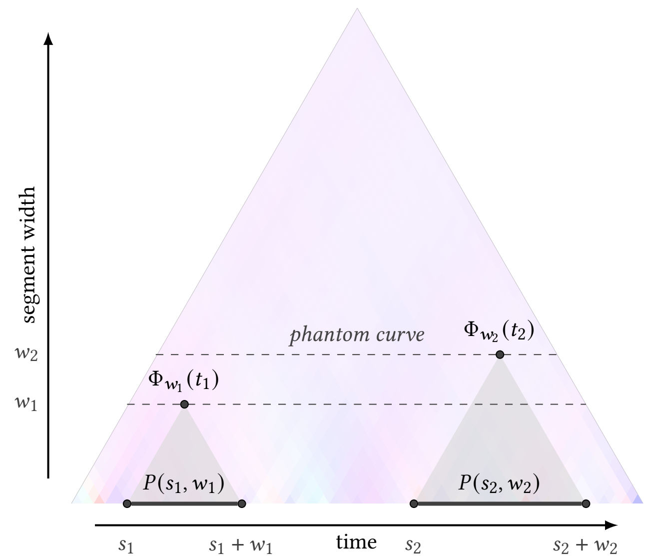

Fig. 1. Hierarchy of all possible timeline segments, viewed through a particular coefficient of the discrete Fourier transform (DFT), also called *wavescape* [16]. All segments with identical width w lie on a *phantom curve* Φ_w, shown as dashed lines. Each point $\Phi_w(t)$ on this curve corresponds to a particular segment $[s, s + w]$ with starting point s and pitch-class content $P(s, w)$; see Section 3 for details. Two example segments $[s_1, s_1 + w_1]$ and $[s_2, s_2 + w_2]$ with corresponding phantom curves Φ_{w_1} and Φ_{w_2} are shown.

1 INTRODUCTION

Scientific discovery can arise by many means. Results can be obtained from meticulous methodological rigor and carefully planned-through experiments, but oftentimes insights are also gained through spontaneous associations of previously unrelated facts or exploratory work. In particular for the latter mode of discovery, a stimulating environment is quintessential in order to allow creativity to flow [4, 9]. Collaborative inter- and transdisciplinary research for which knowledge and expertise from different fields is brought into interaction, is especially likely to provide such environments. Interdisciplinary work can, however, be impeded by specialized expert knowledge that is difficult to communicate to researchers with different backgrounds. In such situations, appropriately designed tools may help to build bridges between different research traditions and provide interfaces to otherwise restricted knowledge.

Here, we report an instance of scientific discovery made possible through interactive exploration using an online interface for music visualization. We document how the discovery of *phantom curves*, a new concept for music analysis, was enabled through the web application midiVERTO[1] for analyzing and visualizing tonal structures in pieces of music [5]. The dynamic nature of the application enables the exploration and comparison of different analyses by changing the parameters through a visual interface. This yields multiple views on tonal structures in a musical piece that can inform and augment further analysis. In our analyses below we show a selection of particularly insightful viewpoints.

In this study, we showcase how synergies can be obtained by combining knowledge from mathematics, and music theory, as well as web development and user interface design. We first briefly describe the capabilities of midiVERTO and document the creative process leading to the discovery of phantom curves. Then, we provide the mathematical details of phantom curves and demonstrate the usefulness of this concept for music analysis drawing on three examples from different music styles, namely the main theme of the musical *The Phantom of the Opera* by Andrew Lloyd Webber (1986),

[1]The application can be accessed at https://dcmlab.github.io/midiVERTO/. It is open source and its code is hosted at https://github.com/DCMLab/midiVERTO/.

DLfM 2022, July 28–29, 2022, Prague

the music of the *Phantom Forest* from the PlayStation video game *Final Fantasy VI* by Nobuo Uematsu (1994), and the film score *Duel of the Fates* from *Star Wars Episode I: The Phantom Menace* by John Williams (1999). We conclude our contribution by outlining which potential future avenues can be taken to further develop the concept of phantom curves, in particular in relation to Music Information Retrieval (MIR) techniques.

2 DOCUMENTING A CREATIVE PROCESS

We now describe the process as well as some favorable factors that led to the discovery of *phantom curves*. Reconstructing creative processes *a posteriori* is difficult since it relies mainly on introspection, but we believe that reflections on how knowledge is generated may provide valuable experiences for others. While the concept of phantom curves has merit in itself, we emphasize in this section three factors that facilitated its discovery. First, the interdisciplinary work in the intersection of music theory and mathematics provided the initial impetus. Second, we were able to draw on prior work, most importantly recent work in mathematical music theory on the application of the discrete Fourier transform to pitch-class sets [1] and the combination of this line of research with hierarchical representations of tonal structure [7, 12] that culminated in the development of a visualization method called *wavescapes* [16]. The colored triangle in Figure 1 is an example for such a wavescape. While wavescapes provide a novel, intuitive way to grasp large-scale tonal structures in compositions, the impact of this method might be somewhat constrained because, in order to use wavescapes in one's own work, certain programming skills are required. We thus decided to develop an interactive version where users, in particular musicologists and music theorists with little programming affinity, would be able to interactively engage with wavescapes. Thus, the affordances of this tool are the third and final factor leading to the discovery of phantom curves. A web application was the natural choice for realizing this project, because it grants widest-possible accessibility, while reducing hardware and software requirements on the user's end.

When implementing the wavescapes and the corresponding Fourier coefficients (see below), we quickly realized that wavescapes incorporate two different notions of time, namely temporal location (starting point of a segment) and temporal extent (segment width). These are given by the horizontal and vertical components of a point in a wavescape, respectively. Moreover, we realized that these two components are not sufficiently explored in a static manner, and that the dynamic rendering in midiVERTO could provide new visualization benefits. In interaction with this tool, we observed that increasing window sizes directly corresponds to different 'cuts' through a wavescape. Moreover, seeing that these cuts 'shrink' with increasing segment width and that the corresponding points on the Fourier coefficients moved closer to one another, formed smoother and smoother curves, and finally contracted into a single point made us realize that we could conceptualize this as a family of curves (see mathematical details in Section 3). Coincidentally, the example piece we were looking at was the main theme of the musical *The Phantom of the Opera*, and combined with the fact that the curves relate to the latent long-term tonal structure hidden in pieces of music, naming the curves phantom curves was obvious.

3 PHANTOM CURVES

The way how pitches of a piece of music are organized – the piece's harmonic structure – can be studied at multiple hierarchical levels by considering the set of pitch-class distributions for all possible segments (i.e., all possible, potentially overlapping time intervals) of the piece [7, 11, 16]. Let T denote the total duration of a piece, measured in beats, quarter notes, or seconds, for example. Any of the piece's segments is uniquely specified by its starting time point $s \in [0,T]$ and its segment width $w \in (0,T]$, the segment's time duration. The *pitch-class content* of any such segment is obtained by counting and summing up all its pitch classes into a vector $P(s,w) \in \mathbb{R}^{12}_{\geq 0}$, where the total count of pitch class j equals the j-th vector entry $P(s,w)[j]$. These counts can also be weighted by the duration of the respective notes, and the formalization with P is generic to this choice. Here, we consider *pitch-class distributions* and thus normalize the pitch-class vectors $P(s,w)$ to sum up to one, $\tilde{P}(s,a) = P(s,a)/\sum_j P(s,a)[j]$.

In recent years it became increasingly popular among music theorists to study pitch-class sets and distributions via the discrete Fourier transform [1, 8, 15, 17]. Since the DFT can be applied to music in symbolic formats (e.g., MIDI) without prior interpretation of the musical material by a music theorist, it is well suited for distant-reading approaches in corpus studies, such as the comparison of different pieces' tonal organization at a high level of abstraction.

Definition 3.1 (Discrete Fourier Transform). The *discrete Fourier transform* (DFT) of any pitch-class vector x (i.e., any choice of $x = P(s,w)$) corresponds to the mapping

$$F\colon \mathbb{R}^{12}_{\geq 0} \to \mathbb{C}^{12}, \qquad F(x)[k] = \sum_{j=0}^{11} x[j] e^{i2\pi j \frac{k}{12}},$$

and $F(x)[k]$ is called the k-th *Fourier coefficient* of x.

The application of this transformation to pitch-class distributions relates to music theory in particular through the association of its coefficients to fundamental concepts: chromaticity (1st coefficient), dyadicity (2nd coefficient), triadicity (3rd coefficient), presence of seventh chords (4th coefficient), diatonicity (5th coefficient), and similarity to whole-tone scales (6th coefficient). Each one of the remaining coefficients is equivalent to one of the first six coefficients by symmetry [16]. The DFT has been combined with *keyscapes* [11] to define *wavescapes* [16], a method for visualizing pitch-class distributions hierarchically through the DFT. Figure 1 shows a particular wavescape, a triangle in which each cell shows the Fourier coefficient of a segment as a color value.

We call horizontal cuts through the wavescape of a particular Fourier coefficient *phantom curves*. The idea of phantom curves is that their shape indicates how a piece's tonality unfolds in time. Such curves can also be visualized on the unit disk in the complex plane which we call a Fourier coefficient space. Figure 2 shows both representations next to one another. Colors between these two visualizations are matched, a detailed definition of the color mapping can be found in [16].

Definition 3.2 (Phantom Curve). For a segment length $w \in (0, T]$, the *phantom curve* Φ_w is defined as the continuous mapping

$$\Phi_w\colon [0, T - w] \to \mathbb{C}^{12}, \qquad \Phi_w(t) = F\left(\tilde{P}(t, w)\right).$$

For an intuitive explanation of phantom curves, consider a sliding window of width $w \in (0, T]$ moving along the piece, from its start to its end. At each position, the pitch-class distribution of the segment covered by the sliding window is mapped to its Fourier coefficients, and the succession of these coefficients forms a curve in the coefficient space. Figure 1 shows two phantom curves for different segment widths, represented by dashed horizontal lines cutting horizontally through the wavescape. The sliding-window segments are shown as bold lines at the bottom of the triangle and the Fourier-coefficient values are plotted in the color domain. A phantom-curve plot in a coefficient space is shown to the right in Figure 2.

Phantom curves are based on the pitch-class content for segments of arbitrary but fixed width w, and this allows to analyze a piece of music from any level of detail. However, a single perspective is rarely sufficient for a comprehensive analysis, and multiple viewpoints may be more revealing. In our exploratory analyses, for example, we determined widths that correspond to appropriately detailed resolutions in dynamic interaction with the application midiVERTO. Hence, we searched for particular instances in the family of all a piece's phantom curves.

Definition 3.3 (Phantom Curve Family). The *family of phantom curves* for a given piece is

$$\Phi\colon (0, T] \to \left\{[0, T - w] \to \mathbb{C}^{12}\right\}, \qquad w \mapsto \Phi_w,$$

where Φ_w is defined as in Definition 3.2.

Note that each phantom-curve family can be understood as a contraction in the sense that for increasing values of w, the phantom curves Φ_w shrink.[2] This contraction can be conveniently visualized using the midiVERTO application. A special phantom curve is the one over the entire piece ($w = T$) for which τ_T is constant ($\tau_T(t) = 0$) and the phantom curve Φ_w contracts into a single point. This curve can be considered the *phantom singularity*, which corresponds to the point at the tip of the triangle in Figure 1.

4 A GALLERY OF UNCANNY EXAMPLES

We now illustrate the usefulness of phantom curves for music analysis by applying the concept to three pieces of music from 20th-century popular music culture: the main theme of a Broadway musical, a video game, and a cinema movie.These specific examples were chosen because each of them alludes to the supernatural, in other words: to phantoms. There is ample music-theoretical evidence that composers writing in the late-Romantic idiom (and 20th-century composers taking their inspirations from them), often draw on remarkably similar techniques when it comes to expressing the eerie and magical [2, 3, 6, 14]. These techniques are sometimes subsumed under the notion of extended tonality [10, 13]. According to these accounts, supernatural and eerie associations are often expressed by drawing on major-third relations that can not be fully accounted for in a diatonic context (representing the natural world), and thus perfectly embody the realm of the otherworldly.

[2] Mathematically, the line integral of Φ_w converges to 0 as w goes to T.

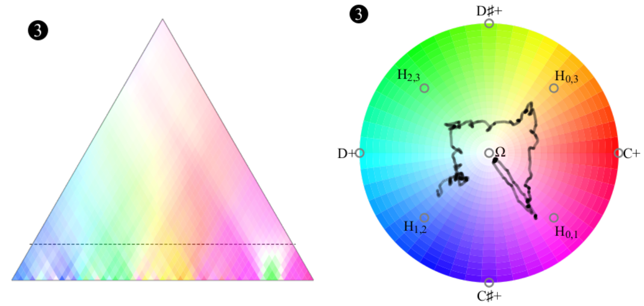

Fig. 2. Phantom curve on the wavescape (left, dashed) and 3rd Fourier coefficient space (right, dotted) for the main theme of *The Phantom of the Opera* (time resolution: ♪; segment width: 150 × ♪ = 18.75 × 𝅝). The labels shown to the right denote augmented triads (e.g., C+) and hexatonic scales that result from joining adjacent augmented triads (e.g., $H_{0,1}$ = C+ ∪ C♯+; the numbers 0 and 1 correspond to the pitch classes of C and C♯).

It should be noted, however, that midiVERTO can be used to analyze any piece in MIDI format and is not restricted to any particular style of music. Non-pitched instruments can, however, not be represented because the Fourier transform as presented here only takes the pitches into account.

The Phantom of the Opera – Main Theme (1986). The main theme of the musical *The Phantom of the Opera* by composer Andrew Lloyd Webber expresses the musical's overall narrative of a tortured soul seeking redemption through love. The character development of the main protagonist, the phantom, is foreshadowed by the essential harmonic progressions of the main theme. Figure 2 shows the third Fourier coefficient space. The four large color regions in blue, green, yellow, and pink correspond to the overall harmonic trajectory D minor – G minor – E minor – F minor. The regularity of this sequence becomes apparent when viewed through the lens of the third Fourier coefficient, revealing the square-like shape of the phantom curve. Because each of these key areas is locally diatonic, the corners of this square are closer to hexatonic scales (i.e., the union of two augmented triads; e.g., $H_{0,1}$ = C+ ∪ C♯+), than to a single augmented triad (e.g., C+) [16].

Final Fantasy VI – Phantom Forest (1994). In *Final Fantasy VI*, a cataclysmic event occurs midway through the game, and transforms the peaceful World of Balance into the World of Ruins. The music for the Phantom Forest by composer Nobuo Uematsu belongs to the first part and reveals interesting aspect of the music's tonal structure through the third Fourier coefficient. Here, the phantom curve exhibits a trident-like shape that results from the rondo-like form, with the overarching key of A minor (pink-violet) and key changes to B♭ major (blueish), F major (violet), C major (orange), and finally E♭ major (greenish). Locally, the melody contains many chromatic alterations and thus the 'fingers' of the trident seem to target the augmented triads. It can also be seen that this basic

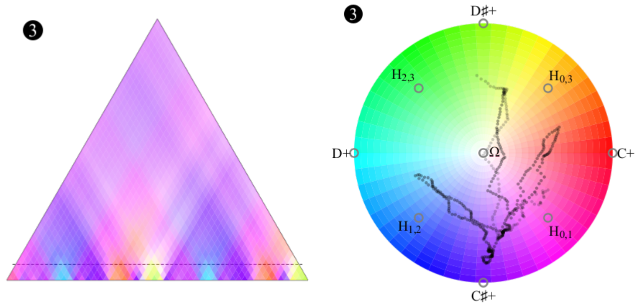

Fig. 3. Phantom curve on the wavescape (left, dashed) and 3rd Fourier coefficient space (right, dotted) for *Final Fantasy VI – Phantom Forest* (time resolution: ♪; segment width: 48×♪ = 6×𝅝). See Figure 2 for the explanation of the labels.

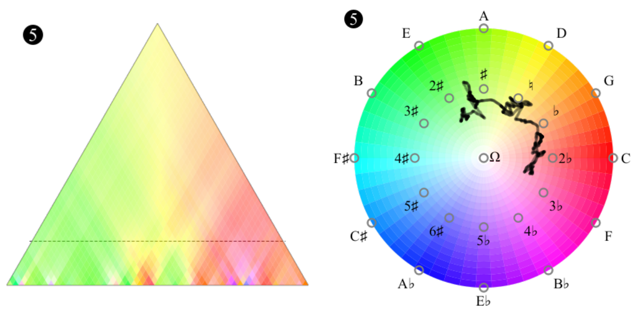

Fig. 4. Phantom curve on the wavescape (left, dashed) and 5th Fourier coefficient space (right, dotted) for *Duel of the Fates* from *Star Wars Episode I – The Phantom Menace* (time resolution: ♪; segment width: 240 × ♪ = 30 × 𝅝). The labels shown to the right denote the circle of fifths for single pitch classes (C, G, D, ...) and diatonic collections (♮, ♯, 2♯, ...).

structure is then repeated. Composing circular forms that allow potentially infinite repetition is often found in game music.

Star Wars Episode I – The Phantom Menace (1999). In one of the movie's most epic battle scenes, accompanied with a film score entitled *Duel of the Fates* by composer John Williams, the binary opposition between the main protagonists Jedi Qui-Gon Jinn and Obi-Wan Kenobi on one side and Sith apprentice Darth Maul on the other is represented by the juxtaposition of two diatonically unreconcilable keys: E minor and G minor. The minor-third distance of these keys leads to very different coloring, green and red, in the fifth Fourier coefficient (Figure 4). Interestingly, the dramatic arc of the conflict between the protagonists seems to be directly reflected in how these keys are used: the phantom curve shows a transition from the green E-minor area to the red G-minor area, with only a brief interlude of balance between the two (shown in yellow).

The analytical sketches of the three pieces above are rather succinct and in no sense meant to represent comprehensive analyses of their tonal and formal structure. Our goal is not to render complete analyses of the pieces, but primarily to demonstrate how phantom curves and midiVERTO can be used exploratively to generate and to interact with such analyses. Readers are invited to take up our analytical outlines as starting points for deeper explorations of the tonal structures of these pieces as well as others.

5 CONCLUSIONS

In this paper, we have introduced a novel concept for computational music analysis: *phantom curves*. We moreover documented the process of how we conceived of this concept in interaction with a freely-available interactive online tool, the web app midiVERTO, thereby demonstrating the scientific and pedagogical value of developing and disseminating such tools. The new idea of phantom curves was then applied to analyze three pieces of 20th-century music from a varied range of genres: a musical theme, video-game music, and the score of a central scene of a cinema blockbuster.

Several avenues present themselves to build upon our present work in the future. In particular, more rigorous characterizations of families of phantom curves seem to be a promising direction. For instance, studying the smoothness and curvature of these families might reveal interesting insights. Applying the concept not only to singular pieces but entire corpora of compositions might moreover reveal clusters of pieces with similar overall tonal plans and hence similar families of phantom curves (up to transposition), a related direction was explored in [7]. A naturally follow-up question would be whether such similarities are effected by factors such as genre, historical period, composer, or instrumentation. It would be moreover interesting to investigate whether and to what degree different analysts agree in their choice of phantom curves for a given piece and arrive at similar conclusions. We would like to encourage readers to explore the potential of midiVERTO, and to utilize it in order to learn more about music, as well as to discover new ideas through exploration and experimentation.

ACKNOWLEDGMENTS

This research project has received funding from the UNIL-EPFL dhCenter, an interdisciplinary research platform set up between Université de Lausanne (UNIL) and École Polytechnique Fédérale de Lausanne (EPFL), Switzerland. It was also supported in part by the University of Amsterdam Data Science Centre. The authors thank Martin Rohrmeier for his support and guidance as well as Jason Yust and Cédric Viaccoz for their valuable comments during the development process of the web application.

REFERENCES

[1] Emmanuel Amiot. 2016. *Music Through Fourier Space: Discrete Fourier Transform in Music Theory.* Springer.
[2] Richard Cohn. 2004. Uncanny Resemblances: Tonal Signification in the Freudian Age. *Journal of the American Musicological Society* 57, 2 (2004), 285–324.
[3] Richard Cohn. 2012. *Audacious Euphony: Chromatic Harmony and the Triad's Second Nature.* Oxford University Press, Oxford.
[4] Mihaly Csikszentmihalyi. 2013. *Creativity: The Psychology of Discovery and Invention.* Harper Perennial Modern Classics.
[5] Daniel Harasim, Giovanni Affatato, and Fabian C. Moss. 2022. midiVERTO: A Web Application to Visualize Tonality in Real Time. In *Mathematics and Computation in Music (Lecture Notes in Computer Science)*, Mariana Montiel, Octavio A. Agustín-Aquino, Francisco Gómez, Jeremy Kastine, Emilio Lluis-Puebla, and Brent Milam (Eds.). Springer International Publishing, Cham, 363–368. https://doi.org/10.1007/978-3-031-07015-0_31
[6] Frank Lehman. 2018. *Hollywood Harmony: Musical Wonder and the Sound of Cinema.* Oxford University Press, Oxford.

[7] Robert Lieck and Martin Rohrmeier. 2020. Modelling Hierarchical Key Structure with Pitch Scapes. In *Proceedings of the 21st International Society for Music Information Retrieval Conference (ISMIR 2020)*. Montreal, Canada, 811–818. https://doi.org/10.5281/zenodo.4245558

[8] Thomas Noll. 2019. Insiders' Choice: Studying Pitch Class Sets Through Their Discrete Fourier Transformations. In *Mathematics and Computation in Music (Lecture Notes in Computer Science)*, Mariana Montiel, Francisco Gomez-Martin, and Octavio A. Agustín-Aquino (Eds.). Springer International Publishing, Cham, 371–378. https://doi.org/10.1007/978-3-030-21392-3_32

[9] Martin Rohrmeier. 2022. On Creativity, Music's AI Completeness, and Four Challenges for Artificial Musical Creativity. *Transactions of the International Society for Music Information Retrieval* 5, 1 (2022), 50–66. https://doi.org/10.5334/tismir.104

[10] Martin Rohrmeier and Fabian C. Moss. 2021. A Formal Model of Extended Tonal Harmony. In *Proceedings of the 22nd International Society for Music Information Retrieval Conference, ISMIR 2021, Online, November 7-12, 2021*, Jin Ha Lee, Alexander Lerch, Zhiyao Duan, Juhan Nam, Preeti Rao, Peter van Kranenburg, and Ajay Srinivasamurthy (Eds.). 569–578.

[11] Craig Stuart Sapp. 2001. Harmonic Visualizations of Tonal Music. In *International Computer Music Conference Proceedings*, Vol. 2001. International Computer Music Association, San Francisco, 1–8. http://hdl.handle.net/2027/spo.bbp2372.2001.029

[12] Craig Stuart Sapp. 2005. Visual Hierarchical Key Analysis. *Computers in Entertainment* 3, 4 (2005), 1–19. https://doi.org/10.1145/1095534.1095544

[13] Arnold Schoenberg. 1969. *Structural Functions of Harmony*. Faber and Faber, London and Boston.

[14] Kenneth M. Smith. 2020. *Desire in Chromatic Harmony: A Psychodynamic Exploration of Fin de Siècle Tonality*. Oxford University Press, Oxford, New York.

[15] Dmitri Tymoczko. 2010. Set-Class Similarity, Voice Leading, and the Fourier Transform. *Journal of Music Theory* 52, 2 (March 2010), 251–272. https://doi.org/10.1215/00222909-2009-017

[16] Cédric Viaccoz, Daniel Harasim, Fabian C. Moss, and Martin Rohrmeier. 2022. Wavescapes: A Visual Hierarchical Analysis of Tonality Using the Discrete Fourier Transform. *Musicae Scientiae* (Jan. 2022), 10298649211034906. https://doi.org/10.1177/10298649211034906

[17] Jason Yust. 2019. Stylistic Information in Pitch-Class Distributions. *Journal of New Music Research* 48, 3 (2019), 217–231. https://doi.org/10.1080/09298215.2019.1606833