MixMash: A Visualisation System for Musical Mashup Creation

Catarina Maçãs*, Ana Rodrigues*, Gilberto Bernardes† and Penousal Machado*
*CISUC - Department of Informatics Engineering
University of Coimbra
Coimbra, Portugal
{cmacas, anatr, machado}@dei.uc.pt
†INESC TEC and University of Aveiro
Aveiro, Portugal
gba@inesctec.pt

***Abstract*—We present MixMash, an interactive tool to assist users in the creation of music mashups based on cross-modal associations between musical content analysis and information visualisation. Our point of departure is a harmonic mixing method for musical mashups by Bernardes et al. [1]. To surpass design limitations identified in the previous method, we propose a new interactive visualisation of multidimensional musical attributes—hierarchical harmonic compatibility, onset density, spectral region, and timbral similarity—extracted from a large collection of audio tracks. All tracks are represented as nodes whose distances and edge connections indicate their harmonic compatibility as a result of a force-directed graph. In addition, we provide a visual language that aims to enhance the tool usability and foster creative endeavour in the search for meaningful music mixes.**

***Keywords*-Music Mashup; Harmonic Mixing; Perceptual Relatedness; Information Visualisation; Force-Directed Layout; Cross-modal Associations;**

I. INTRODUCTION

Mashup creation is a music composition practice strongly linked to various sub-genres of Electronic Dance Music (EDM) and the role of the DJ [2]. It entails the recombination of existing (pre-recorded) musical audio as a means for creative endeavour. Existing mass preservation mechanisms allow users to access large collections of musical audio in digital format from which mashup practitioners can retrieve material for their mixes. To assist in the retrieval process, both industry and academia have devoted effort to enhance the experience of using digital tools for mashup creation by streamlining the time-consuming and complex search for compatible musical audio.

Early research on computational mashup creation, focused on rhythmic-only features, particularly those relevant to the temporal alignment of two or more musical tracks [3]. Recent research [1], [4], [5] has expanded the range of musical attributes under consideration, notably including harmonic-driven features to identify compatible musical audio, commonly referred to as "harmonic mixing".

In this paper, we extend Bernardes et al.'s [1] existing harmonic mixing method by addressing an identified limitation of its interface design. The method offers multidimensional information relevant to mashup creation in a navigable 2-dimensional interface, which intuitively guides users in the mixing process. The interface design has proven suitable for small music collections, yet the clutter resulting from the superposition of elements and the static representation of the tracks, limits usability when exploring larger collections. To address this limitation, we study the use of a force-directed graph to enhance the clarity and fluidity of a visualization. Through linked nodes on a canvas and a set of visual variables to characterise rhythmic, spectral and timbral track attributes, we express the hierarchical harmonic compatibility among tracks. Additionally, our interface allows users to easily highlight tracks with different attributes, further improving navigation through a musical audio collection.

The remainder of the paper is organised as follows. Section II, summarises the related work. Section III, describes the content-driven audio analysis that support our novel music visualization system. Section IV, presents a new approach for the visualization of compatible music tracks, as well as the description of the used graphic elements to characterise tracks. In Section V, we present the interface and interaction techniques, and in Section VI, we state the conclusions of our work and future directions.

II. RELATED WORK

A. Harmonic Mixing

We can identify four major harmonic mixing methods: key affinity, chroma vectors similarity, sensory dissonance minimisation and hybrid (hierarchical) models. The initial three approaches focus on single harmonic attributes only and the latter approach provides a hierarchical view over harmonic compatibility.

The earliest computational approaches to harmonic mixing were proposed by industry [6], [7] and adopted key distances (in the *Camelot Wheel* or circle of fifths representation [6]) as an indicator of harmonic compatibility. The large-scale harmonic nature of this indicator favours some degree of tonal stability of the mix by privileging the use of the same or a related diatonic key pitch set.

Two latter approaches have been proposed within the academic literature, which aim to find good matches between pitch-shifted audio tracks by maximizing the similarity between chroma vectors [4] and minimizing sensory dissonance [5]. These approaches privilege local alignments at the time scale of beat durations, over large-scale harmonic structure between audio slices with highly similar pitch class content.

Bernardes et al. [1] recently proposed a hierarchical harmonic mixing method, which combines previous approaches for harmonic compatibility in a single framework. One of the main novelties of the method is its interface design, which offers a global view over the harmonic compatibility of an entire music collection (many-to-many) for both small and large time scales of music, beyond the existing one-to-many relationships between a user-defined track and an audio collection [1] in related software.

B. Music Visualisation

Historically, numerous artists have created audiovisual associations, later referred to as graphic notation and visual music. Pioneering works by Kandinsky, Pfenninger, Cage, Fischinger, and Whitney explore combinations of visual principles (like colour) to emphasise the audiovisual experience [8]. In the present work, the visual translation of a collection of musical tracks was approached from two different points of view: music visualisation and the visual representation of a large collection of music tracks.

Regarding music visualisation, while some authors have solved this problem by focusing on the geometry of musical structure [9], others have found a visual solution by focusing on a specific set of musical characteristics [10]–[12]. For example, Wattenberg [10] uses arc diagrams to connect sequences containing the same pitches, revealing the structure of musical compositions. Sapp [11] on the other hand, presented a multi-time scale visualisation to compare differences between key's tonic pitch and view the harmonic structure and relationships in a musical composition. An experimental visualisation that aims to provide a perceptually relevant experience for the user was presented by Rodrigues et al. [13]. Overall, the aforementioned visualisations do not allow great modularity of data, often binding the visual clarity of each element and limiting the comparison of other musical characteristics.

Although visualisations of large collections of music tracks have been addressed [14]–[18], it still is an area in need of great developments. Grill and Flexer [14], similar to the work of Bernardes et al. [1], developed a visualisation strategy capable of representing perceptual qualities from a large collection of sounds. Although their musical focus was in sound textures, they also aimed at finding an intuitive and meaningful interface. For this purpose, they constructed an audiovisual language based on the cross-modal mechanisms of human perception. In the end, subjects were able to successfully associate sounds with the corresponding graphical representations. On the other hand, Hamasaki and Goto, Gulik et al. and Schwartz and Schnell [15], [17], [18] developed interactive visualizations of several layers of information, but do not visually represent the perceptual relevance of the musical characteristics.

C. Force-directed Graphs

Graph Visualisation is an important visualisation technique to handle the representation of relational structures in data, aiding in the analysis, modelling, and interpretation of complex network systems [19]. These models are characterised by the existence of two main elements: (i) nodes, representing an entity (e.g., person, cell, machine); and (ii) edges, representing the relation between nodes. This visualisation model is often applied to large and complex datasets, which often leads to a set of problems related to performance and clutter. As the graph grows in size, the required visual space and computational resources also increases [20]. To solve this, clustering techniques can be applied, grouping similar nodes by their semantics and/or position on the graph. By doing so, it is possible to reduce visual clutter, enhance clarity and performance, and to create a simplified overview of the network's structure, contextualising it and reducing complexity [20], [21].

The positioning of nodes in space can be defined through a different set of graph layouts such as, arc diagrams, tree maps, circular, and force-directed layouts; in this paper, we focus on the latter. Force-directed graphs are based on a physical system that organises the network through forces of repulsion and attraction applied continuously to each node. This technique facilitates the visual interpretation of the inherent structure of the data, improving the analysis and comprehension of complex networks [22].

Force-directed graphs can be used in a variety of areas, such as biology, medicine, literature, sociology [23]–[28], but also in the visualization of music collections [29], [30]. In the work of Muelder et al. [31], a force-directed layout is applied to visualise music libraries. Each piece of music is represented by a node, positioned depending on its similarities with other musics (e.g. same artist). Songrium [15], is a music browsing service that visualises relations among original songs and their derivative works. Each node represents a video-clip and each edge connects an original work to its derivatives. In addition, songs with similar moods get stronger forces, thus are positioned closer together.

In this work, we aim to expand the state-of-the-art by creating a visualisation that enables a user to analyse and relate a large collection of music tracks for mashup creation, with the long term goal of having a positive impact on user creativity [32].

III. MIXMASH

MixMash is a software application which aims to assist users in finding musical mixes in the context of mashup creation. It relies on metrics and methods presented in Bernardes et al. [1], which expand state-of-the-art harmonic mixing by providing a greater amount of relevant information to guide the mixing process. Its main novelty lies in a hierarchical harmonic mixing method, which includes metrics for both small-scale and large-scale structural levels, i.e. local (e.g., beats or phrases) and global (e.g, large sections or overall musical mashup) harmonic alignments between musical tracks, respectively. Moreover, this method considers three additional dimensions that can help users defining compositional goals in terms of rhythmic (onset density), spectral (region) and timbral qualities.

To promote an intuitive search for compatible samples across all tracks in a collection, a many-to-many mapping strategy was previously introduced in the interface design. This design opposes the ranked track list to a user-defined track, adopted in previous systems harmonic mixing software [6], [7], which is (i) inefficient computationally, as it recomputes intensive audio signal analysis every time a different audio track is selected as target, and (ii) limited in promoting creative endeavour and serendipity [1].

The back end of MixMash includes signal processing methods that automatically describe harmonic, rhythmic, spectral and timbral attributes from musical audio. The algorithmic strategies adopted to extract these attributes are extensively detailed in Bernardes et al. [1].

Both small-scale and large-scale harmonic compatibility metrics rely on indicators from the perceptually-motivated Tonal Interval Space [33], which represents musical tracks as points in a 12-dimensional space. Small-scale harmonic compatibility, S, combines two indicators: perceptual relatedness and dissonance, computed as distance metrics. The smaller the perceptual relatedness distances, the greater the affinity between two given tracks. To this indicator, we combine the degree of dissonance of the combination of two given audio tracks, which we aim to minimise following previous evidence from Gebhardt et al. [5]. Large-scale harmonic compatibility, L, is expressed as the distance from the 24 major and minor key locations, $k = 24$, and can be interpreted as the degree of association of a given sample to a musical key.

Three additional descriptions of rhythmic, spectral and timbral audio track attributes are computed. The rhythmic content is described by each track note onset density as the ratio between the number of onsets and the entire duration of the track (in seconds). The spectral content is described by the region a sample occupies in the 24-band perceptual Bark frequency scale. The timbral content similarity between tracks is computed as the cosine distance between Mel-frequency spectrum coefficients (MFCC).

Given a user-defined collection of n audio tracks, we compute (i) an $n + k$ square matrix, expressing the distance metrics for small, S, and large, L, scale harmonic compatibility metrics, (ii) two n-dimensional rhythmic and spectral track attribute vectors, and (iii) a $n{\cdot}n$ square timbral similarity matrix. All this is done using the Pure Data code from the original contribution [1].

IV. VISUALISATION MODEL

With the aim of improving the interaction and visualisation of the interface presented in Bernardes et al. [1], we implement a force-directed graph layout based on the *ForceAtlas2* algorithm [22] to create an emergent, organic and appealing environment to the user. To improve readability, we also define a graphical language to characterise and distinguish tracks. We implement a graphic interface (Fig. 1) which filters the track collection according to user preferences on harmonic, rhythmic, spectral and timbral attributes. We allow the user to manipulate the forces of attraction and repulsion between nodes, easing the reading of more cluttered zones (which represent groups of tracks with a strong harmonic compatibility). Moreover, to further avoid undesired clutter, we group compatible nodes into clusters. These clusters are represented in the graph as a group of nodes that can be expanded or withdrawn through interaction. In addition, we adopt a user-defined threshold value to constrain the number of connections between nodes and uniquely represent the tracks with strong harmonic compatibility.

A. *Force-directed system*

We apply a force-directed graph based on the work of Jacomy et al. [22] to represent the small- and large-scale harmonic compatibility distance metrics. In this layout, nodes represent both musical audio tracks and each one of the 24 major and minor musical keys. To distinguish the two types of nodes, we use different visual representations (discussed in IV-B). We have access to the computed harmonic compatibilities (see III), which resulted in a matrix of distances among tracks and keys. Through this matrix, we define the weight for each edge, and therefore the force of attraction between nodes. The more similar the nodes, the higher the force of attraction, and consequently the closer the nodes. These forces are applied to the nodes independently of their type, facilitating the interpretation of the graph and interaction when searching for harmonically compatible tracks. Additionally, by representing the musical keys through nodes, and consequently their relation to the tracks, we also represent the most probable key for each music track.

To provide different visualizations over the tracks' relations, we implement a set of mechanisms to define visible tracks in the graph. These mechanisms can be further explored by the user through the graphic interface. The main

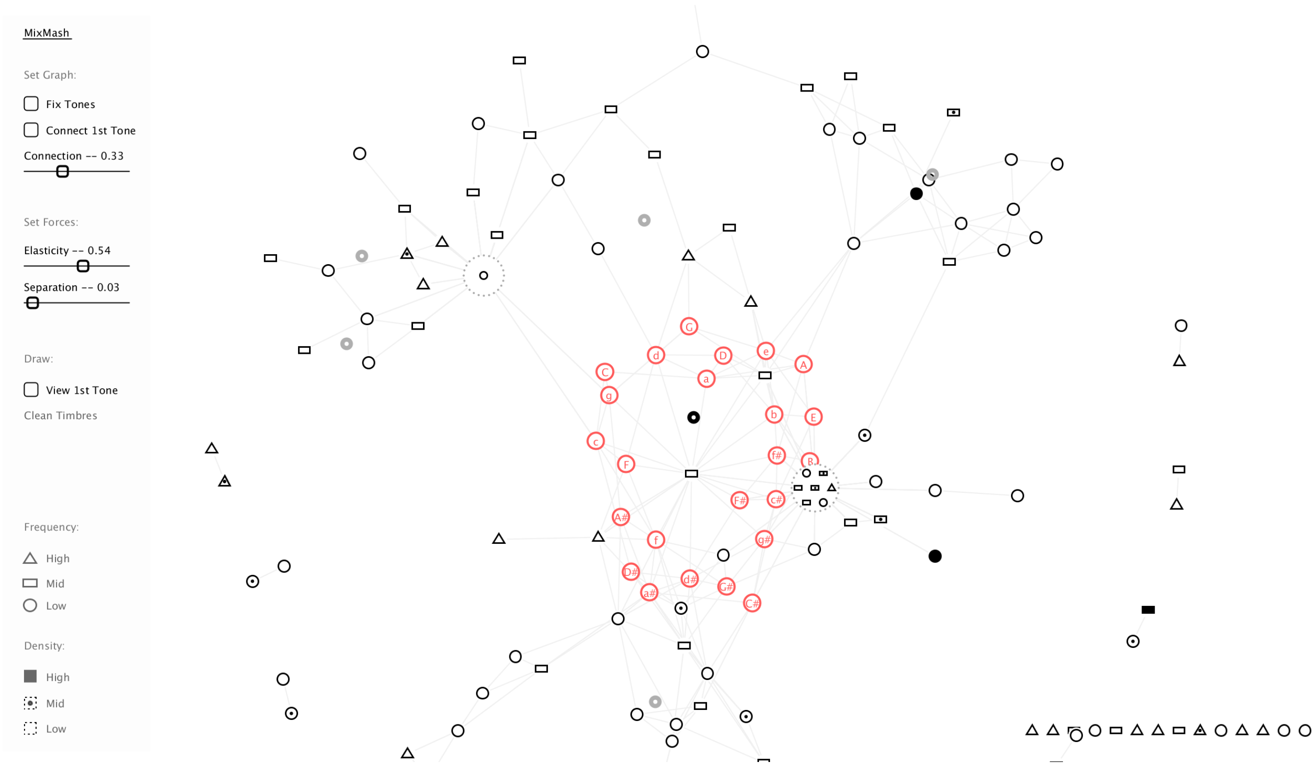

Figure 1. Screenshot of the visualization interface. On the left, the interaction panel; at the centre, the graph.

mechanism is to define a threshold value that determines whether two nodes are connected. For each node, we only search for the adjacent nodes whose harmonic compatibility value is lower than a predefined value (set in the interface). A second mechanism limits the connections between nodes and keys. If the latter mechanism is adopted by the user, tracks will connect only to the closest key, reducing clutter and enhancing the association between keys and musical audio tracks.

As the number of tracks in a collection can widely vary, we create another mechanism to prevent cluttered graphs. We used an agglomerative clustering algorithm and aggregate the nodes by their compatibility values. In addition, we define a minimum number of two nodes per cluster to prevent clusters being too small. We also connect these clusters to adjacent nodes. These adjacent nodes are retrieved from the list of adjacent nodes of each node present in the cluster. If one "outer" node is connected to an "inner" node, we connect the cluster to it. The attraction force between a cluster and an adjacent node is equal to the average force of all forces between "inner" nodes and that "outer" node.

As in *ForceAtlas2*, we apply an attraction force to the centre of the canvas, making all nodes gravitate towards this point. This gravitational force is significantly weaker than the others, so it does not have much influence in the positioning of each node, preventing only the nodes from dispersing in the canvas.

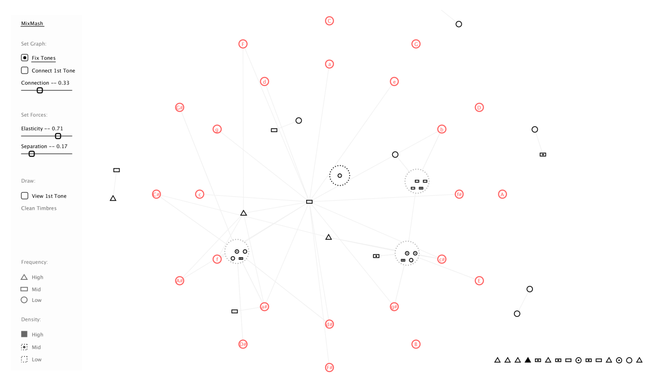

Figure 2. Fixing the keys in the circle of fifths. By clicking on the "Fix Tones" button, all nodes that represent a key are fixed on the canvas according to their positioning in the circle of fifths.

By default, the keys are positioned depending on their relations with the other tracks. This causes the nodes to have different positions at every run, and can create some fatigue in the user in the search of a specific key and its connected tracks. To facilitate this search, we implemented a mechanism that positions the keys according to the *Camelot Wheel* or circle of fifths representation (Fig. 2). As such, all connected keys have a fixed position, and only the musical audio tracks are influenced by attraction and repulsion forces.

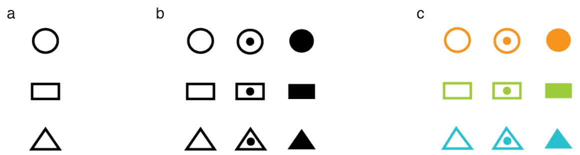

Figure 3. Audiovisual mappings: a) Shape - Spectrum Frequency from low frequencies (circle) to higher ones (triangle); b) Object Fill - Onset Density. Empty shape corresponds to low density, half-filled shape to normal/mid density, and full shape corresponds to high density; c) Colour - Spectrum Frequency from low frequencies (orange) to higher ones (blue).

B. Graphic Variables/ Audiovisual mappings

To ensure an efficient and interactive visualisation, the development of the visualisation model adheres to the following guidelines/challenges: (i) to each track there is a corresponding visual representation based on its musical attributes, creating a consistent visual feedback between similar tracks; (ii) perceptual foundations are used to guide the aesthetical choices concerning visual representation of tracks; (iii) a natural and intuitive interaction with the tool is promoted allowing the user to easily navigate among tracks to create his/her mashup.

For each track, the onset density, spectral region, and timbral attributes are mapped to a corresponding visual variable. The spectral region and onset density are subdivided into three levels of magnitude, allowing a more accurate interpretation and analysis of musical data.

The spectral attribute is split into high, medium and low regions. Its visual representation is primarily characterised by shape. Although there is a tendency to associate colour to pitch levels [8], [34], [35], it has been proposed that the use of shape or monochrome colour is more efficient than multicolour when the human eye has to distinguish several layers of information [36], [37]. As such, high frequency regions of the spectrum are associated with sharp contours (triangle), low frequencies were associated with rounded contours (circle) [38], and medium frequencies are associated with neutral contours achieved through the use of straight lines (rectangle) (see Fig. 3a). Colour is then used to reinforce these associations. Higher frequencies are related to a cold colour (blue) and low frequencies to a warm colour (orange) (see Fig. 3c).

Onset density is also split into high, medium and low density levels. Since onset density is related to the number of notes in a track, we convey this information into the visual parameter of shape-filling. Low density corresponds to an empty shape (only contours visible), medium density to a half-filled shape, and high density to a completely filled shape (see Fig. 3b).

Timbre is often referred in literature as the colour of instruments [34], [35]. Inspired by this definition, we represent this attribute by a coloured circle in the top side of the nodes, only visible it the user selects it (see Fig. 4b). Then, all tracks with similar timbres will also gain a circle coloured with the same colour code as the clicked node (see Fig. 4c). The choice for the colour code does not rely on any type of association. Our aim is only to have a clear distinction between colours so when multiple timbres are at displayed they can be easily distinguished.

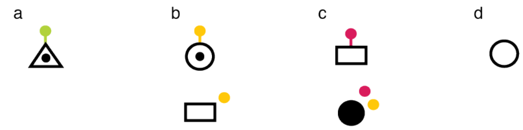

Figure 4. Timbre representations. a) One track is selected and there are no other tracks with similar timbres. b) One track is connected timbre wise to the selected track c) Track can have more than one timbre similarity d) The track is not related to any of the selected tracks.

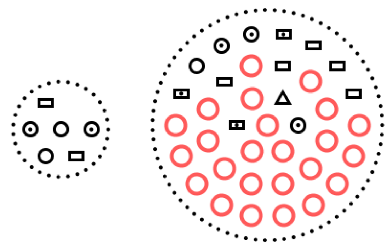

Figure 5. Representation of two clusters with different sizes.

To differentiate the nodes representing keys from audio tracks, we colour the outline of the key nodes with red and represent the key's tonic pitch with typography, allowing a direct reading of the key (see Fig.1).

For the clusters' representation, instead of adding more complexity to the canvas through the definition of a different shape, we use the elements of the cluster itself to represent it. As such, the cluster is represented by a sample from the "inner" nodes. This sample is delimited by a circle, whose size depends on the size of the cluster (Fig. 5). Thereafter, the user can differentiate the clusters from the nodes, and, at the same time, have an overview of their elements.

V. Interface and Interaction

Visual representation of data elements strongly impacts how we visualise large amounts of data. In this section, we discuss how our visual representation of musical concepts are displayed in light of a carefully designed interactive visualization.

Our system enables the user to explore the graph by modifying its organisation (e.g., fixing the keys according to the circle of fifths), changing the connection threshold between nodes, adjusting the forces of attraction and repulsion, and highlighting specific characteristics of the tracks. This is all accessible through an interactive panel on the left side of the interface (Fig. 1). We created a video (accessible through https://vimeo.com/270076175) to better demonstrate these interactions.

Once the system has been initialised, the user will see the nodes and clusters stabilising their position in the centre of the canvas. The user can then zoom and pan the graph to view more details. To listen to the tracks, the user has to move the mouse over each node. To select a node the user needs to left click. With this action, he/she will listen again to the track sound. If the user clicks in the 's' key from the keyboard, the music will stop playing. By continuously selecting nodes, the user is saving the tracks, which can be listened at the same time by clicking in the 'space' key. With this strategy, we aim to guide user through different mixing possibilities towards the creation of their own mashups. All these interaction techniques aim to enhance the understanding of the tracks harmonic compatibility and foster user creativity, by allowing the user to efficiently explore a large musical audio collection towards specific composition goals.

To view the compatible timbres of a certain track, the user has to click on the right button of the mouse. In doing so, the closest tracks (in terms of timbre) will be complemented with a circle as explained in IV-B.

To expand each cluster, the user has to click inside the cluster. By doing so, the nodes inside the cluster will be visible, as well as a donut shape. The latter, positioned in the centre of all corresponding nodes, functions as a button to close the cluster. By passing the mouse over this donut shape, all the nodes belonging to that cluster will be highlighted through a stroke coloured in magenta.

Overall, this set of choices are important to achieve our ultimate goal of an intuitive and meaningful interactive visualization tool in the context of musical mashup creation.

VI. Conclusions and Future Work

In this paper, we present a new approach to the work of Bernardes et al. [1] concerning visualisation and interaction techniques. Instead of a static visualisation of musical tracks, we implement a visualization system which relies on forces of attraction and repulsion to place the tracks depending on their harmonic compatibility. As future work, we intend to improve the clustering algorithm by giving the user the possibility to choose how to cluster nodes—by key, onset density, spectral region or timbral similarity. To improve the readability, we intend to change the nodes' size depending on their compatibility—nodes with higher compatibility, grow in size, emphasising highly compatible tracks. As the number of tracks can increase depending on the user, we aim to implement a fish-eye zoom technique, so the user can have detail in certain areas without losing context of the surrounding tracks. Finally, we intend to create a timeline in which the user can position the chosen tracks so they can play at different timings, enabling the composition of mashups with complex temporal structures.

Acknowledgments

The work is supported by Portuguese National Funds through the FCT-Foundation for Science and Technology, I.P., under the project IF/01566/2015 and the grant SFRH/BD/129481/2017. Project TEC4Growth - Pervasive Intelligence, Enhancers and Proofs of Concept with Industrial Impact/NORTE-01-0145-FEDER-000020, financed by the North Portugal Regional Operational Programme (NORTE 2020), under the PORTUGAL 2020 Partnership Agreement, and through the European Regional Development Fund (ERDF).

References

[1] G. Bernardes, M. E. P. Davies, and C. Guedes, "A hierarchical harmonic mixing method," *Lecture Notes in Computer Sciences: Post-proceedings of the CMMR'17 Conference*, In press.

[2] J. Shiga, "Copy-and-persist: The logic of mash-up culture," *Critical Studies in Media Communication*, vol. 24, no. 2, pp. 93–114, 2007.

[3] G. Griffin, Y. Kim, and D. Turnbull, "Beat-sync-mash-coder: A web application for real-time creation of beat-synchronous music mashups," in *Proceedings of ICASSP*, 2010, pp. 437–440.

[4] M. E. P. Davies, P. Hamel, K. Yoshii, and M. Goto, "Automashupper: Automatic creation of multi-song music mashups," *IEEE Trans. ASLP*, vol. 22, no. 12, pp. 1726–1737, 2014.

[5] R. B. Gebhardt, M. E. P. Davies, and B. U. Seeber, "Psychoacoustic approaches for harmonic music mixing," *Applied Sciences*, vol. 6, no. 5, 2016.

[6] Mixed in Key, "Mashup 2 [software]," last accessed on 28 March 2017. [Online]. Available: http://mashup.mixedinkey.com

[7] Native Instruments, "Traktor pro 2 [software]," last accessed on 1 September 2017. [Online]. Available: https://www.native-instruments.com/en/products/traktor/dj-software/traktor-pro-2/

[8] M. McDonnell, "Visual music," *Visual Music Marathon, Boston Cyberarts Festival Programme*, 2007.

[9] T. Bergstrom, K. Karahalios, and J. C. Hart, "Isochords: visualizing structure in music," in *Proceedings of Graphics Interface 2007*. ACM, 2007, pp. 297–304.

[10] M. Wattenberg, "Arc diagrams: Visualizing structure in strings," in *Information Visualization, 2002. INFOVIS 2002. IEEE Symposium on*. IEEE, 2002, pp. 110–116.

[11] C. S. Sapp, "Harmonic visualizations of tonal music." in *ICMC*, 2001, pp. 419–422.

[12] J. Snydal and M. Hearst, "Improviz: visual explorations of jazz improvisations," in *CHI'05 extended abstracts on Human factors in computing systems*. ACM, 2005, pp. 1805–1808.

[13] A. Rodrigues, A. Cardoso, and P. Machado, "Harmonic constellation: An audiovisual environment of living organisms," in *Fourth International Workshop on Musical Metacreation (MUME-2016)*, 2016.

[14] T. Grill and A. Flexer, "Visualization of perceptual qualities in textural sounds," in *ICMC*. Citeseer, 2012.

[15] M. Hamasaki and M. Goto, "Songrium: a music browsing assistance service based on visualization of massive open collaboration within music content creation community," in *Proceedings of the 9th International Symposium on open collaboration*. ACM, 2013.

[16] A. Rauber, E. Pampalk, and D. Merkl, "The som-enhanced jukebox: Organization and visualization of music collections based on perceptual models," *Journal of New Music Research*, vol. 32, no. 2, pp. 193–210, 2003.

[17] D. Schwarz and N. Schnell, "Sound search by content-based navigation in large databases," in *Sound and Music Computing (SMC)*, 2009.

[18] R. van Gulik, F. Vignoli, and H. van de Wetering, "Mapping music in the palm of your hand, explore and discover your collection," in *Proceedings of the 5th International Conference on Music Information Retrieval*, 2004.

[19] I. Meirelles, *Design for information*. Rockport publishers, 2013.

[20] I. Herman, G. Melançon, and M. S. Marshall, "Graph visualization and navigation in information visualization: A survey," *IEEE Transactions on visualization and computer graphics*, vol. 6, no. 1, pp. 24–43, 2000.

[21] D. Kimelman, B. Leban, T. Roth, and D. Zernik, "Reduction of visual complexity in dynamic graphs," in *International Symposium on Graph Drawing*. Springer, 1994, pp. 218–225.

[22] M. Jacomy, T. Venturini, S. Heymann, and M. Bastian, "Forceatlas2, a continuous graph layout algorithm for handy network visualization designed for the gephi software," *PloS one*, vol. 9, no. 6, 2014.

[23] A. J. Enright and C. A. Ouzounis, "Biolayout–an automatic graph layout algorithm for similarity visualization," *Bioinformatics*, vol. 17, no. 9, pp. 853–854, 2001.

[24] K.-I. Goh, M. E. Cusick, D. Valle, B. Childs, M. Vidal, and A.-L. Barabási, "The human disease network," *Proceedings of the National Academy of Sciences*, vol. 104, no. 21, pp. 8685–8690, 2007.

[25] C. Chen, "Citespace ii: Detecting and visualizing emerging trends and transient patterns in scientific literature," *Journal of the Association for Information Science and Technology*, vol. 57, no. 3, pp. 359–377, 2006.

[26] J. Heer and D. Boyd, "Vizster: Visualizing online social networks," in *Information Visualization, 2005. INFOVIS 2005. IEEE Symposium on*. IEEE, 2005, pp. 32–39.

[27] F. Gilbert, P. Simonetto, F. Zaidi, F. Jourdan, and R. Bourqui, "Communities and hierarchical structures in dynamic social networks: analysis and visualization," *Social Network Analysis and Mining*, vol. 1, no. 2, pp. 83–95, 2011.

[28] X. Fu, S.-H. Hong, N. S. Nikolov, X. Shen, Y. Wu, and K. Xuk, "Visualization and analysis of email networks," in *Visualization, 2007. APVIS'07. 2007 6th International Asia-Pacific Symposium on*. IEEE, 2007, pp. 1–8.

[29] F. Vavrille. Liveplasma. [Online]. Available: http://www.liveplasma.com/artist-Hard-Fi.html

[30] M. Gibney. Music-map, the tourist map of music. [Online]. Available: https://www.music-map.com

[31] C. Muelder, T. Provan, and K.-L. Ma, "Content based graph visualization of audio data for music library navigation," in *Multimedia (ISM), 2010 IEEE International Symposium on*. IEEE, 2010, pp. 129–136.

[32] N. Henry, J.-D. Fekete, and M. J. McGuffin, "Nodetrix: A hybrid visualization of social networks," *IEEE transactions on visualization and computer graphics*, vol. 13, no. 6, pp. 1302–1309, 2007.

[33] G. Bernardes, D. Cocharro, M. Caetano, C. Guedes, and M. E. P. Davies, "A multi-level tonal interval space for modelling pitch relatedness and musical consonance," *Journal of New Music Research*, vol. 45, no. 4, pp. 281–294, 2016.

[34] E. L. Mudge, "The common synaesthesia of music," *Journal of Applied Psychology*, vol. 4, no. 4, pp. 342–345, 1920.

[35] J. Ward, B. Huckstep, and E. Tsakanikos, "Sound-colour synaesthesia: To what extent does it use cross-modal mechanisms common to us all?" *Cortex*, vol. 42, no. 2, pp. 264–280, 2006.

[36] R. Arnheim, *Art and visual perception*. Univ of California Press, 1974.

[37] A. Chatterjee, *The aesthetic brain: How we evolved to desire beauty and enjoy art*. Oxford University Press, 2013.

[38] C. Spence, "Crossmodal correspondences: A tutorial review," *Attention, Perception, & Psychophysics*, vol. 73, no. 4, pp. 971–995, 2011.