MEASURING DISRUPTION IN SONG SIMILARITY NETWORKS

Felipe Falcão[1] **Nazareno Andrade**[1] **Flávio Figueiredo**[2]
Diego Silva[3] **Fabio Morais**[1]
[1] Universidade Federal de Campina Grande, Brazil
[2] Universidade Federal de Minas Gerais, Brazil
[3] Universidade Federal de São Carlos, Brazil
felipevf@copin.ufcg.edu.br, nazareno@computacao.ufcg.edu.br

ABSTRACT

Investigating music with a focus on the similarity relations between songs, albums, and artists plays an important role when trying to understand trends in the history of music genres. In particular, representing these relations as a similarity network allows us to investigate the innovation presented by these entities in a multitude of points-of-view, including disruption. A disruptive object is one that creates a new stream of events, changing the traditional way of how a context usually works. The proper measurement of disruption remains as a task with large room for improvement, and these gaps are even more evident in the music domain, where the topic has not received much attention so far. This work builds on preliminary studies focused on the analysis of music disruption derived from metadata-based similarity networks, demonstrating that the raw audio can augment similarity information. We developed a case study based on a collection of a Brazilian local music tradition called Forró, that emphasizes the analytical and musicological potential of the musical disruption metric to describe and explain a genre trajectory over time.

1. INTRODUCTION

Inflections on creative threads are prevalent events that can be observed multiple times throughout music history [1]. The emergence of punk rock in the early seventies, for example, changed the traditional rock and roll in many aspects to create a unique music expression [2]. The music branch of the punk culture brought heavy lyrics, aggressive looks, and even deep acoustic changes to the songs, which were more aroused and noisier than songs from previous decades. Such music aspects were replicated over time, as evidenced for example in the expert-curated influences credited in the AllMusic guide [3] to artists from the early stages of the punk rock (e.g. *Ramones*, *Bad Religion*, *Sex Pistols*). The guide attributes to these bands influence over more recent ones, such as *Green Day* and *The Offspring*.

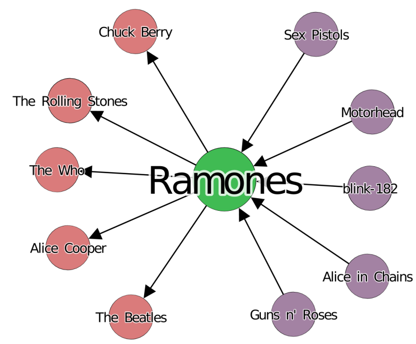

Figure 1. Network topology for a disruptive artist.

Regarding the different creative roles played by artists during the genre trajectories, the AllMusic guide defines the *Ramones* as *"inarguably the most relevant band in punk history, creating the stylistic prototype that would be followed by countless bands who emerged in their wake"* [1]. That points out to a particular innovative case where an artist had a significant and primary influence over the rupture of some well-established guidelines. Thus, an artist can be considered *disruptive* when your music contribution is developed in a self-sufficient way, abruptly shifting the present creative thread.

Conversely, AllMusic suggests a different nature of innovation when describes *Green Day* as *"influenced by the late-'70s punk predecessors, they went on to introduce a new, younger generation to the genre"* [2]. On the opposite of the disruptive movement by the *Ramones*, this excerpt allows describing *Green Day*'s creative potential as a consolidation of the previous practices, including their particular musical signature in the meantime.

Both musical creative natures can be represented by a network model, where the nodes represent artists and the edges symbolize the influence relation of one artist over another. Figure 1 shows a disruptive innovation by the *Ramones* based on influences metadata extracted from AllMusic, where edges indicate an "influenced by" relation. Predecessors (i.e., one that chronologically preceded another) of a focal node (in green) are represented by red

[1] https://www.allmusic.com/artist/ramones-mn0000490004/biography
[2] https://www.allmusic.com/artist/green-day-mn0000154544/biography

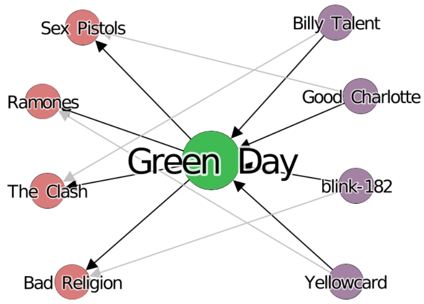

Figure 2. Network topology for a consolidator artist.

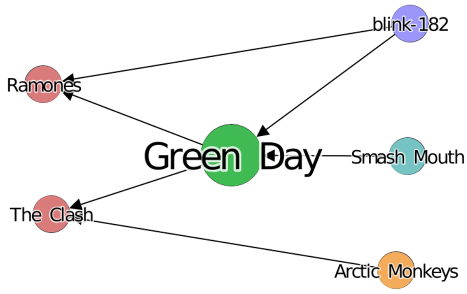

Figure 3. Demonstration for different types of influence.

nodes, whereas purple nodes represent the successors.

Due to its self-sufficient nature, one may expect that most of the artists (nodes) that succeed and are influenced by a disruptive artist connect to this artist, but not to its predecessors. A similar explanation can be given to describe the network topology for artists that consolidate the genres over time (Figure 2 for *Green Day*): these artists reaffirm a thread of influences, as their successors are usually influenced by both them and their predecessors.

Despite its simple semantics and noticeable potential in enriching musicological analyses about the history of genres, there is limited research that measures the disruption of songs over time. In particular, this work is based on the disruption quantification using a metric derived from audio similarity networks. Specifically, we model a network of similarity among songs and use their temporal precedence to explore patterns of similarity that reveal creative aspects and disruption over time.

Features extracted from raw audio are reportedly a rich source of similarity information [4], as they cover many music aspects, such as timbre [5, 6], harmony [7, 8], and rhythm [9]. Therefore, consider such types of data to construct similarity networks can be valuable in understanding how songs of the same genre are acoustically related. In this work, a musical disruption analysis is proposed over this similarity network, allowing to unveil some interesting findings of the disruption of songs over time. To promote a better interpretation of results, we collected a new audio dataset comprising songs of a definite style, called Forró.

In this analysis we represent songs as Mel-Frequency Cepstral Coefficients (MFCCs), using these representations to build a similarity network that connects songs with similar acoustics. Next, we process this network's topology to calculate the disruption metric for all songs, summarizing the most disruptive music pieces. This analysis allows us to validate the disruption metric in the music context. Both the dataset and the network representations generated during the experiments are made available for further studies.

2. BACKGROUND

2.1 Music Innovation & Corpora

Music innovation is not a popular main research topic, being usually mentioned in studies focused on modeling music influence [10–12]. In particular, Shalit et al. [13] proposed a dynamic topic model to represent music influence over time using metadata and audio. In their work a song is considered influential if its language gets replicated by subsequent works, while innovation is modeled as the extent of which the model accounts for a song when trained only with data from the past. Their findings leveraged the Million Song Dataset [14] to point to correlations between influence and innovation only during some short periods in the early 70's and mid-90's.

Associations between music influence and innovation were also investigated by Noyer and his collaborators [15]. Using a network to represent influences between artists, the authors tried to find topological differences between innovative and non-innovative artists. Their approach analysed artists data from 1951 to 2008, measuring innovation as the number of Grammy awards won by each artist. Conclusions identify that innovation in fact impacts network topology, showing that artists with more awards presented considerably more structural holes on their sub-networks.

Corpora plays a major role when representing the influence relations between artists, albums and songs. Both the Million Song Dataset [13,16] and the information available on the AllMusic music catalogue [15, 17, 18] have been used as audio and metadata source for many MIR tasks, including influence modeling.

2.2 Disruption Index

This present work builds on a network metric proposed by Funk & Owen-Smith [19] to measure destabilizing and consolidating influences of inventions over existing technology streams. Their CD_t index assumes that the degree of destabilizing influence (disruption) of an invention within an influence network should be measured in terms of "how future inventions make use of the technological predecessors cited by a focal patent". Given that notion, Figure 3 depicts an example of the three types of music influences for a focal node a (*Green Day* in the example) used for measuring disruption according with CD_t: Let n_i be the number of nodes i that reference only a and none of its predecessors (*e.g., Smash Mouth*), n_j be the number of nodes j that reference both a and at least one of its predecessors (*e.g., blink-182*), while n_k accounts for the number of nodes k that do not reference a but reference at least one of its predecessors (*e.g., Arctic Monkeys*). Disruption (henceforth referenced as D) is then measured as:

$$D = \frac{n_i - n_j}{n_i + n_j + n_k} \tag{1}$$

D ranges from -1 to 1, where values close to 1 indicate nodes with highly disruptive potential whereas measures around the negative extreme denote nodes that mostly consolidated influences over time and therefore were cited in parallel with their predecessors.

The original case study for D was developed over a database of patents granted in the US between 1977 and 2005. Their findings indicate that disruptive inventions are usually boosted by federal research funding initiatives, while commercial ties are more related to the consolidation of the *status quo*. Such validation was later expanded by Wu et al. [20], that enriched the dataset with scientific papers and software repositories, accounting now for a total of 65 million observations. The study concluded that disruptive products are associated with smaller teams, while larger groups mostly consolidated knowledge.

The disruption metric was recently experimented in the music context by Figueiredo and Andrade [17]. They leveraged influence metadata from AllMusic to create a network linking 32,568 artists according with their influence relations (i.e. a "link from artist a to b denotes that a has been influenced by b"). Disruptions are then extracted for all the network components, triggering discussions regarding disruptive and consolidator potential. In particular, they confirm the results of [13] about the lack of correlation between influence and disruption, also concluding that D translates structural insights that are not derived from any existing network metrics.

3. COLLECTED MUSIC DATA

Our musical disruption analysis uses audio data from a Brazilian cultural manifestation called Forró (composed by music, dance, and festivities), native from the northeast of Brazil during the second half of the 19^{th} century. The Forró music genre is composed of three preponderant instruments: accordion, triangle, and a percussive drum called *zabumba*. *Luiz Gonzaga* is the most prominent representative of this genre and is responsible for spreading his music to other regions of Brazil.

The audio data was obtained from the collaborative site *Forró em Vinil* [21], which organize and publish contents that register the history of Forró (e.g. albums, books, and pictures). The audio collection is maintained by media collectors that own long play records and CD's that are no longer produced by record labels. These collectors digitize their media and provide the audio files to the site's administrators, responsible for curating the collection. We built a dataset covering Forró songs ranging from the years 1945 to 2016, by scraping the site via a web crawler. Overall, 2,449 distinct albums were collected, grouping a total of 31,485 songs, each one annotated with artist, album, and release year.

To ensure that the collected data is only comprised of Forró songs, we excluded other genres found on a descriptive analysis phase. Moreover, to guarantee a chronological information required by the nature of this study, we filtered out albums without release year informed. These data corrections were necessary to satisfy genre-specific and time constraints requirements. Figures 4 and 5 show the song and album distribution over decades before and after dataset correction, respectively. Out of the original 2,449 albums, 2,293 satisfied the constraints, accounting for 27,352 songs which we considered.

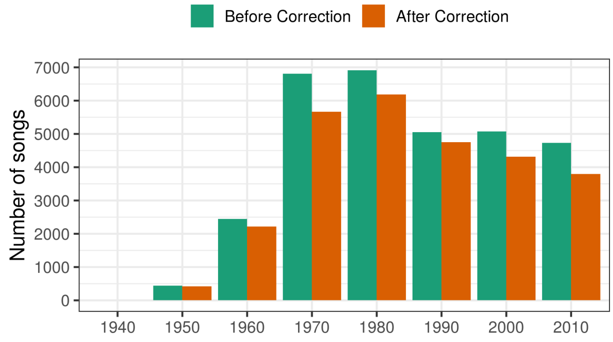

Figure 4. Histogram for number of songs over decades.

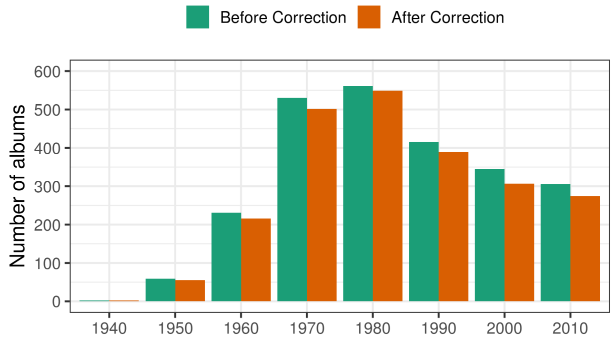

Figure 5. Histogram for number of albums over decades.

4. MUSIC & SIMILARITY REPRESENTATIONS

We leverage Mel-Frequency Cepstral Coefficients (MFCCs) as feature for audio similarity estimation. MFCCs are robust music representations often used in many music information retrieval tasks [22], including genre classification [23], music recommendation [24, 25] and audio similarity [26, 27]. Moreover, given its reported ability to model timbre information [28], we expect that this feature will also capture relevant audio events when iterating over our data. In particular, we look for disruptive episodes when an artist included new instruments to the basic setup discussed on Section 3, which is something that actually happened during the history of Forró.

Similarly to what is proposed by Choi et al. [22] for their baseline feature, we employ as our audio feature the means and standard deviations for 20 MFFCs and their first order derivatives over the entire song. Audio processing techniques are aided by the Librosa package [29]. The result of this feature extraction methodology is a collection

Task #	# of classes (n in top-n)	Max. items per class	Sample size	Precision
1	20 classes	500 items	6572 items	**0.79**
2	50 classes	25 items	1125 items	**0.88**

Table 1. Sampling settings and reported precisions for artist (#1) and album (#2) classification tasks.

of 27,352 vectors (henceforth referenced as *feature vectors*), each one containing 80 elements that represent audio information.

An extra validation step is also conducted to confirm those feature vectors encode enough audio information to generate comprehensible music similarities. Two multi-class audio classification tasks are designed to measure the precision of machine learning classifiers when trained with feature vectors from the *Forró em Vinil* Dataset:

- Task 1: Artist classification: classification of artists among the *top-n* (those with more songs);
- Task 2: Album classification: classification of albums among the *top-n* (those with more songs);

SVM classifiers are used in both cases, given their efficiency in tasks with small training sets. Model training was done using scikit-learn [30] and experiments are run with 10-fold cross-validation using stratified splits. Models have their parameters optimized upon the use of grid-search on the validation phase and reported precision values are related to the best classifier after all splits are done. Table 1 summarizes both the dataset sampling strategies and scores for artist and album classifiers, indicating high precisions for both cases (79% and 88%, respectively).

These partial results present two interesting findings that support the next steps. First we can now fairly assume that our vector representations encode enough audio information to derive similarity measures. The second conclusion refers to the best performing kernel function considered by the grid-search routine for both classifiers: the Radial Basis Function (RBF) kernel. The RBF kernel models vector distance, and its mathematical definition [31] assigns to itself a similarity interpretation [32] (i.e., values ranging from 0 to 1, inversely proportional to the vector distance). Given its potential on providing similarity insights for sequential data, we opt for using the RBF kernel as similarity measure for pairs of feature vectors.

5. SIMILARITY NETWORK

To measure disruption we first need to construct a directed network connecting similar objects. When it comes to song similarity networks, the nodes are the songs and an edge between any pair of songs represent a similarity relationship. We now describe how our network was built.

Each song from the *Forró em Vinil* Dataset is a single node in this network. As for the edges, although the RBF kernel allows us to quantify the similarity between any pair of songs, the binary choice about whether or not we create

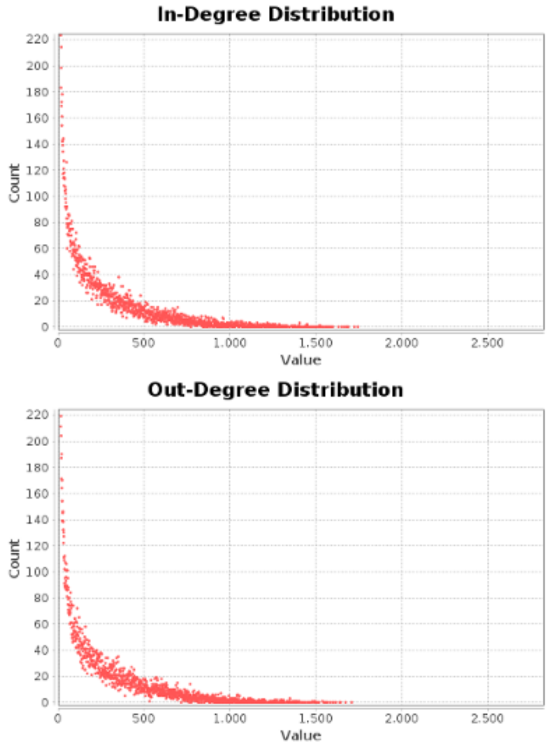

Figure 6. Distribution of in and out degrees.

an edge between two nodes depends on the definition of a similarity threshold above which we can safely ensure that a similarity edge exists. In order to empirically select this threshold, we leverage the fact that songs from the same album are arguably a fair ground truth for noticeable similarity (i.e., these songs usually share the same instrument and voice settings). Thus, we iterate over the whole dataset checking the average similarities between each song s from album a and every other song s' from a. This analysis informs an average similarity of $\approx$ 0.90, which from now on is used as threshold when creating edges.

To create our network, for each song s we query the similarity matrix among all songs and create an edge from s to a predecessor s' if their similarity is greater than or equal the threshold. Additionally, to limit our analysis to a timeframe where stylistic movements are observable, we enforce a time window within which two songs must fall to in order to enable connections between them. Here the size of this time window is 10 years, as it seems reasonable in this context and was also used in [20] when deriving influences between scientific papers.

Out of the 27,352 original songs from the dataset, 26,452 are included in the resulting graph, connected by 5,728,466 directed edges. This minimal difference from the original songs count is explained by the removal of disconnected nodes (i.e. songs that do not sound similar to any other). 98% of all nodes are densely connected to the same giant component and in and out-degrees distribution can be observed on Figure 6. Moreover, Figure 7 illustrates an ego-network extracted from the original structure.

Song	Artist	Album (year)	Disruption Index & Comments
Padrinho Cícero do Juazeiro	*Trio Juazeiro*	*Pedaço de fulô (1982)*	$D = 1$ ($n_i = 37, n_j = 0, n_k = 0$). A fast song (140bpm) with a clear and complex accordion arrangement. The sub-network focused around its node evidences connections with multiple songs from same albums, what might indicate the emergence of a new (disruptive) acoustic setting that was adopted by following artists, like *Clemilda* and *Roberto do Acordeon*.
Namorada de João	*Coroné Narcisinho*	*Forró do Serrado (1969)*	$D = 1$ ($n_i = 28, n_j = 0, n_k = 0$). The song brings a very audible triangle as part of its percussive setup, what can't be perceived in most of the songs from the same epoch. Dominant triangles can also be heard in many of the songs that succeded *Namorada de João*, as in *Esse forró eu danço* (*Abdias* - 1977).
Sem vergonheira	*Marinês*	*Canção da fé (1972)*	$D = 1$ ($n_i = 24, n_j = 0, n_k = 0$). Marinês is one of the first female Forró singers. *Sem vergonheira*, as many of her songs, presents a combination of a strong lead singing voice and effective backing vocals, an unusual practice back then. Similar strategy is used by some of its succeeding songs, like *Quebra-cabeça* (*Trio Nordestino - 1981*).
Derramaro o gai	*Luiz Gonzaga*	*O nordeste na voz de Luiz Gonzaga (1962)*	$D = 1$ ($n_i = 22, n_j = 0, n_k = 0$). The refined accordion melodies are undoubtedly the strongest aspect of *Luiz Gonzaga*'s work, and this song is proof of that. *Derramaro o gai* has multiple disruptive connections with other songs from its very same album, as well as similarities with songs from *Severino Januário*, his brother.
Lembranças	*Flávio José*	*Só confio em tu (1977)*	$D = 1$ ($n_i = 19, n_j = 0, n_k = 0$). The song empowers the acoustic guitar among the original Forró instrumentation, what was rare back in the late seventies. Similar songs by *Flávio José* solidify this new creative branch, imitated by artists like *Marinês* and *Genival Lacerda*.

Table 2. Top-5 of disruptive songs according to the D measure.

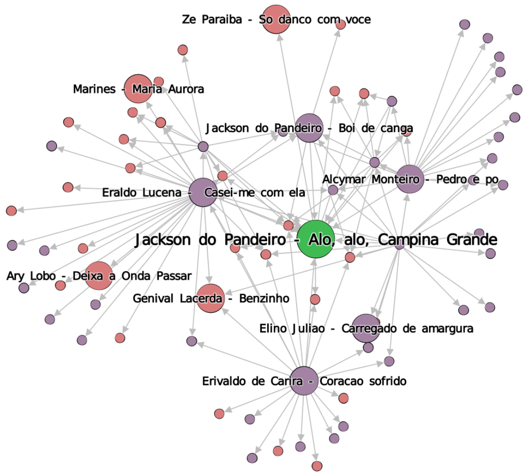

Figure 7. Ego-network for *Alô, alô, minha Campina Grande* by *Jackson do Pandeiro* ($D = -0.07$). Red nodes preceded and purple succeeded the focal, green, node.

6. DISRUPTION ANALYSIS

We can now combine the disruption metric D with the similarity network proposed in the previous section to trigger discussions regarding the disruptive potential of songs over the history of Forró. Since disruption as modeled by Equation 1 depends on preceding data (i.e., n_j and n_k nodes), we decide to use songs prior to 1960 only as comparison data for the following decades, hence no disruptions for these are reported. In other words, the songs from the forties and fifties are a part of the graph (they impact the disruption of future songs), we just do not report their disruption. All the other songs have their disruption indexes derived according to the i, j and k as described on Section 2.2.

Table 2 depicts data from the disruption ranking and summarizes the five most disruptive songs of the *Forró em Vinil* Dataset, trying to support these findings with specifics related to the songs acoustics and their similarity relations. Although artist influence is not a mandatory requirement when determining disruption, it's meaningful to evidence that music pieces from representative artists such as *Luiz Gonzaga*, *Marinês* and *Flávio José* are considered disruptive according to our analysis.

We draw special attention to songs from *Sivuca* that are included among the most disruptive ones (eight songs with $D \geq 0.5$). This musician, widely acclaimed for his work both in Brazil and the United States, was a multi-instrumentalist with strong accordion and acoustic guitar skills. Many of his songs with high disruption in the network combine elements from a variety of genres other than Forró, like *Choro*, *Frevo*, Jazz and Blues. The acoustic richness assigned to his work as well as the uniqueness of the music performed by *Sivuca* generate a lot of internal similarity relations between his own songs, causing the high disruptions. To put it another way, when it comes to Forró, *Sivuca* was disruptive in the sense that his work was mostly influenced by himself, and himself only. This peculiar finding is a representative example of how the disruption metric can actually help to identify meaningful events hidden inside the history of genres.

Next, we leverage the disruption information to model how the creative thread for Forró was developed during the past seven decades. With this analysis we aim at finding exactly when the genre presented creative inflections and how often these events happen during its history. Due

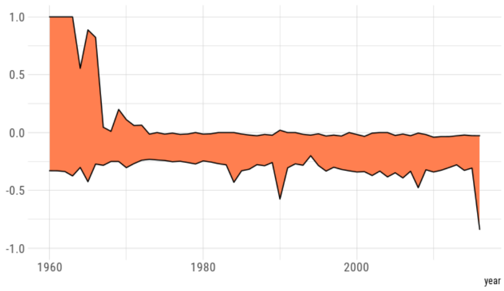

Figure 8. 5 to 95 percentile range of D over the years.

to the large number of songs with D around zero, caused by the dense network, we opt to summarize this disruption distribution in Figure 8 using 5 to 95 percentile ranges over the years. In overall, the higher disruptions of Forró are mostly concentrated on its first years, specially in the interval between 1960 and 1970. While at a first glance this may look like a natural consequence of these being the first songs in the dataset, recall that we omit an entire decade from Figure 8 (i.e., to filter out biases due to first mover advantage, songs from the 1950's impact the disruption of future songs but are not present in our analysis).

We further queried the ranking to understand what happened in the 1960's. Firstly, we see that this high creative load is guided by multiple disruptive songs from pioneers of the history of Forró, like *Luiz Gonzaga*, *Jackson do Pandeiro* and *Marinês*. When we investigate the biographies of these artists (from the AllMusic Guide), we point out facts such as: *Jackson do Pandeiro* [3] is considered: *"one of the most inventive and influential Brazilian musicians"*, *Luiz Gonzaga* [4] is cited as *''one of the most influential figures of Brazilian popular music in the twentieth century"*. Finally, *Marinês* was the first woman to have a Forró group [5]. Biographies were last accessed in August 2020.

Nevertheless, we do point out that the following decades were also presented with disruptive songs. In particular, we propose an artist by year investigation to unveil some insights regarding artists who have unsettled the creative structure of Forró. Figure 9 uses the same percentile approach as Figure 8 to summarize the disruption information for the six artists with higher averaged D for aggregated data (i.e., all songs from the artist in the network). Again we see *Luiz Gonzaga* and *Marinês* figuring as very disruptive artists, with a high creative production specially until 1980, when their careers came to an end (*Luiz Gonzaga* died in 1989 and *Marinês* reduced her production after late 1980). Their creative legacy seems to have been inherited by *Genival Lacerda* and *Flávio José*, other disruptive artists that have been active since the seventies and which often perform disruptive songs since then. These other artists provide further evidence that our ranking is not entirely explained by first mover advantage.

[3] `https://www.allmusic.com/artist/mn0000109367`
[4] `https://www.allmusic.com/artist/mn0000316340`
[5] `https://www.allmusic.com/artist/mn0000371916`

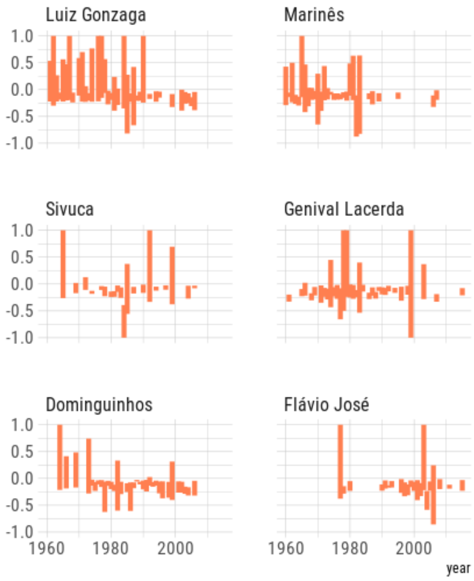

Figure 9. 5 to 95 percentile range of D for disruptive artists over the years.

7. FUTURE WORK & CONCLUSIONS

The present study proposed an audio-based approach to extend the experimentation of a disruption metric in the music context. A new dataset comprised of songs from a Brazilian music tradition was collected to allow for a specific case study. The data supported the generation of an audio similarity network that models the creative flow of songs over time. Results derived from the disruption index underline the semantic potential attached to it, by triggering discussions about specific times when the genre had creative inflections and even which artists were responsible for these events. We argue in favor of applying similar approaches to different contexts, as this may unveil interesting findings about the history of many music genres.

A complementary research direction encourages some enhancements on Equation 1. In particular, we advocate that this formula should also account for the nodes similarities encoded on the edges, instead of simply dealing with creative relations in a binary fashion. That would prevent future studies from having to define a similarity threshold to choose whether or not similarity edges are included in the network, as suggested by this work.

Reproducibility: Both the MFCCs for the *Forró em Vinil* Dataset and the generated similarity network (Graph Exchange XML Format) [6], as well as the code used during the experiments [7] are publicly available.

Acknowledgements: We thank the anonymous reviewers for their feedback. We also acknowledge CNPQ's Universal 2018 Grant: 421884/2018-5.

[6] https://zenodo.org/record/3820920
[7] https://github.com/nazareno/forro-disruption

8. REFERENCES

[1] P. Tschmuck, *Creativity and innovation in the music industry*. Springer, 2006.

[2] D. Simonelli, "Anarchy, pop and violence: Punk rock subculture and the rhetoric of class, 1976-78," *Contemporary British History*, vol. 16, no. 2, 2002.

[3] AllMusic, "Allmusic - record reviews, streaming songs, genres and bands," (accessed August, 2020), "https://www.allmusic.com".

[4] M. Müller, *Information retrieval for music and motion*. Springer, 2007, vol. 2.

[5] M. Muller and S. Ewert, "Towards timbre-invariant audio features for harmony-based music," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 18, no. 3, 2010.

[6] J. Pons, O. Slizovskaia, R. Gong, E. Gómez, and X. Serra, "Timbre analysis of music audio signals with convolutional neural networks," in *EUSIPCO*. IEEE, 2017.

[7] M. Müller and S. Ewert, "Chroma toolbox: Matlab implementations for extracting variants of chroma-based audio features," in *ISMIR*, 2011.

[8] F. Korzeniowski and G. Widmer, "Feature learning for chord recognition: The deep chroma extractor," in *ISMIR*, 2016.

[9] G. T. Toussaint, L. Matthews, M. Campbell, and N. Brown, "Measuring musical rhythm similarity: Transformation versus feature-based methods." *Journal of Interdisciplinary Music Studies*, vol. 6, no. 1, 2012.

[10] J. Atherton and B. Kaneshiro, "I said it first: Topological analysis of lyrical influence networks." in *ISMIR*, 2016.

[11] N. J. Bryan and G. Wang, "Musical influence network analysis and rank of sample-based music." in *ISMIR*, 2011.

[12] N. Collins, "Influence in early electronic dance music: An audio content analysis investigation." in *ISMIR*, 2012.

[13] U. Shalit, D. Weinshall, and G. Chechik, "Modeling musical influence with topic models," in *ICML*, 2013.

[14] T. Bertin-Mahieux, D. P. Ellis, B. Whitman, and P. Lamere, "The million song dataset," 2011.

[15] E. Noyes, I. E. Allen, and S. Parise, "Artistic influences and innovation in the popular music industry," *Frontiers of Entrepreneurship Research*, vol. 30, no. 15, 2010.

[16] J. Serrà, Á. Corral, M. Boguñá, M. Haro, and J. L. Arcos, "Measuring the evolution of contemporary western popular music," *Scientific reports*, vol. 2, 2012.

[17] F. Figueiredo and N. Andrade, "Quantifying disruptive influence in the allmusic guide," in *ISMIR*, 2019.

[18] N. Collins, "Computational analysis of musical influence: A musicological case study using MIR tools." in *ISMIR*, 2010.

[19] R. J. Funk and J. Owen-Smith, "A dynamic network measure of technological change," *Management Science*, vol. 63, no. 3, 2017.

[20] L. Wu, D. Wang, and J. A. Evans, "Large teams develop and small teams disrupt science and technology," *Nature*, vol. 566, no. 7744, 2019.

[21] F. em Vinil, "Forró em vinil - um pequeno apanhado da música nordestina em vinil!" (accessed August 2nd, 2020), "https://www.forroemvinil.com".

[22] K. Choi, G. Fazekas, M. B. Sandler, and K. Cho, "Transfer learning for music classification and regression tasks," *CoRR*, vol. abs/1703.09179, 2017. [Online]. Available: http://arxiv.org/abs/1703.09179

[23] G. Kour and N. Mehan, "Music genre classification using MFCC, SVM and BPNN," *International Journal of Computer Applications*, vol. 112, no. 6, 2015.

[24] D. Bogdanov, M. Haro, F. Fuhrmann, E. Gómez, and P. Herrera, "Content-based music recommendation based on user preference examples," in *WOMRAD Workshop*, 2010.

[25] B. NIU, L.-z. KONG, S.-l. LUO, L.-m. PAN, and L. GUO, "Individuality music recommendation model based on MFCC and gmm [j]," *Transactions of Beijing Institute of Technology*, vol. 4, 2009.

[26] J. H. Jensen, M. G. Christensen, D. P. Ellis, and S. H. Jensen, "Quantitative analysis of a common audio similarity measure," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 17, no. 4, 2009.

[27] A. Flexer, D. Schnitzer, M. Gasser, and T. Pohle, "Combining features reduces hubness in audio similarity," *Children*, vol. 15, no. 15.95, 2010.

[28] F. De Leon and K. Martinez, "Enhancing timbre model using MFCC and its time derivatives for music similarity estimation," in *EUSIPCO*. IEEE, 2012.

[29] B. McFee, C. Raffel, D. Liang, D. P. Ellis, M. McVicar, E. Battenberg, and O. Nieto, "librosa: Audio and music signal analysis in python," in *SciPy Conf.*, 2015.

[30] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg *et al.*, "Scikit-learn: Machine learning in python," *The Journal of Machine Learning Research*, vol. 12, 2011.

[31] K.-M. Chung, W.-C. Kao, C.-L. Sun, L.-L. Wang, and C.-J. Lin, “Radius margin bounds for support vector machines with the rbf kernel,” *Neural computation*, vol. 15, no. 11, pp. 2643–2681, 2003.

[32] J.-P. Vert, K. Tsuda, and B. Schölkopf, “A primer on kernel methods,” *Kernel methods in computational biology*, vol. 47, 2004.