

A COMPUTATIONAL ANALYSIS OF REAL-WORLD DJ MIXES USING MIX-TO-TRACK SUBSEQUENCE ALIGNMENT

Taejun Kim¹ Minsuk Choi¹ Evan Sacks²
Yi-Hsuan Yang³ Juhan Nam¹

¹ Graduate School of Culture Technology, KAIST, South Korea

² 1001Tracklists, United States

³ Research Center for IT Innovation, Academia Sinica, Taiwan

{taejun, minsukchoi, juhan.nam}@kaist.ac.kr, evan@1001tracklists.com, yang@citi.sinica.edu.tw

ABSTRACT

A DJ mix is a sequence of music tracks concatenated seamlessly, typically rendered for audiences in a live setting by a DJ on stage. As a DJ mix is produced in a studio or the live version is recorded for music streaming services, computational methods to analyze DJ mixes, for example, extracting track information or understanding DJ techniques, have drawn research interests. Many of previous works are, however, limited to identifying individual tracks in a mix or segmenting it, and the sizes of the datasets are usually small. In this paper, we provide an in-depth analysis of DJ music by aligning a mix to its original music tracks. We set up the subsequence alignment such that the audio features are less sensitive to the tempo or key change of the original track in a mix. This approach provides temporally tight mix-to-track matching from which we can obtain cue-points, transition length, mix segmentation, and musical changes in DJ performance. Using 1,557 mixes from *1001Tracklists* including 13,728 tracks and 20,765 transitions, we conduct the proposed analysis and show a wide range of statistics, which may elucidate the creative process of DJ music making.

1. INTRODUCTION

A Disc Jockey (DJ) is a musician who plays a sequence of existing music tracks or sound sources seamlessly by manipulating the audio content based on musical elements. The outcomes can be medleys (mix), mash-ups, remixes, or even new tracks, depending on how much DJs edit the substance of the original music tracks. Among them, creating a mix is the most basic role of DJs. This involves curating music tracks and their sections to play, deciding the order, and modifying them to splice one section to another as a continuous stream. In each step, DJs consider various elements of the tracks such as tempo, key, beat,

chord, rhythm, structure, energy, mood and genre. These days, DJs create the mix not only for a live audience but also for listeners in music streaming services.

Recently, imitating the tasks of DJ using computational methods has drawn research interests [1–7]. On the other hand, efforts have been made to understand the creative process of DJ music making. In the perspective of reverse engineering, tasks extracting useful information from real-world DJ mixes can be useful in such a pursuit. In the literature, at least the following tasks have been studied. (1) *Track identification* [8–10]: identifying which tracks are played in DJ music which can be either a mix or a manipulated track. (2) *Mix segmentation* [11, 12]: finding boundaries between tracks in a DJ mix. (3) *Mix-to-track alignment* [13, 14]: aligning the original track to an audio segment in a DJ mix. (4) *Cue point extraction* [14]: finding when a track starts and ends in a DJ mix. (5) *Transition unmixing* [13, 14]: explaining how DJs apply audio effects to make a seamless transition from one track to another. However, the previous studies only focused on solving the tasks usually with a small dataset and did not provide further analysis using extracted information from the tasks. For example, Sonnleitner et al. [8] used 18 mixes for track identification. Glazyrin [11] and Scarfe et al. [12] respectively collected 103 and 339 mixes with boundary timestamps for mix segmentation. The majority of previous studies concentrated on identification and segmentation and few studies on the other three tasks used artificially generated datasets [13, 14].

To address the need of a large-scale study, we collected in a total of 1,557 real-world mixes and original tracks played in the mixes from *1001Tracklists*, a community-based DJ music service.¹ The mixes include 13,728 unique tracks and 20,765 transitions. However, tracks used in DJ mixes usually include various versions so-called “extended mix”, “remix”, or “edit”. Also, a few tracks in tracklists of the collected dataset are annotated incorrectly by users. Therefore, an alignment algorithm is required to ensure that the collected tracks are exactly the same versions as the ones used in the mixes. More importantly, the alignment will be a foundation for further computational analysis of DJ mixes. With these two motivations, we



© Taejun Kim, Minsuk Choi, Evan Sacks, Yi-Hsuan Yang, Juhan Nam. Licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). **Attribution:** Taejun Kim, Minsuk Choi, Evan Sacks, Yi-Hsuan Yang, Juhan Nam, “A Computational Analysis of Real-World DJ Mixes using Mix-To-Track Subsequence Alignment”, in *Proc. of the 21st Int. Society for Music Information Retrieval Conf.*, Montréal, Canada, 2020.

¹ <https://www.1001tracklists.com>

Summary statistic	All	Matched
The number of mixes	1,564	1,557
The number of unique tracks	15,068	13,728
The number of played tracks	26,776	24,202
The number of transitions	24,344	20,765
Total length of mixes (in hours)	1,577	1,570
Total length of unique tracks (in hours)	1,038	913
Average length of mixes (in minutes)	60.5	60.5
Average length of unique tracks (in minutes)	4.1	4.0
Average number of played tracks in a mix	17.1	15.5
Average number of transitions in a mix	14.5	12.9

Table 1. Statistics of the *1001Tracklists* dataset. The original dataset size is denoted as ‘All’ and the size after filtering as ‘Matched’.

set up the mix-to-track subsequence dynamic time warping (DTW) [15] such that the mix can be aligned with the original tracks in presence of possible tempo or key changes. The warping paths from the DTW provide temporally tight mix-to-track matching from which we can obtain cue points, transition lengths, and key/tempo changes in DJ performances in a quantitative way. To evaluate the performances of the alignment and the cue point extraction methods simultaneously, we evaluate mix segmentation performances regarding the extracted cue points as boundaries dividing two adjacent tracks in mixes, comparing them to human-annotated boundaries. Furthermore, by observing the performance changes depending on the three different types of cue points, we analyze the human annotating policy of track boundaries.

Although DJ techniques are complicated and different depending on the characteristics of tracks, there has been common knowledge for making seamless DJ mixes. However, to the best of our knowledge, the domain knowledge has never been addressed in the literature with statistical evidence obtained by computational analysis. In this study, we analyze the DJ mixes using the results from the subsequence DTW mentioned above for the following hypotheses: 1) DJs tend not to change tempo and/or key of tracks much to avoid distorting the original essence of the tracks. 2) DJs make seamless transitions from one track to another considering the musical structures of tracks. 3) DJs tend to select cue points at similar positions in a single track.

The analysis is performed based on the results obtained from the subsequence alignment and provides insights statistically for tempo adjustment, key transposition, track-to-track transition lengths, and agreements of the cue points among DJs. We hope that the proposed analysis and various statistics may elucidate the creative process of DJ music making. The source code for the mix-to-track subsequence DTW, the cue point analysis and the mix segmentation is available at the link.²

2. THE DATASET

Our study is based on DJ music from *1001Tracklists*. We obtained a collection of DJ mix metadata via direct personal communication with *1001Tracklists*. Each entry of

²<https://github.com/mir-aidj/djmix-analysis/>

mixes contains a list of track, boundary timestamps and genre. It also contains web links to the audio files of the mixes and tracks. We downloaded them separately from the linked media service websites on our own. We found a small number of web links to tracks are not correct and so filtered them out by a mix-to-track alignment method automatically (see Section 3.3). The boundary timestamps of tracks in a mix are annotated by the users of *1001Tracklists*.

Table 1 summarizes statistics of the dataset. The original size of the dataset is denoted as ‘All’ and the size after filtering as ‘Matched’ in Table 1. Note that the number of played tracks is greater than the number of unique tracks as a track can be played in multiple mixes. The dataset includes a variety of genres but mostly focuses on House and Trance music. More detailed statistics of the dataset are available on the companion website.³

3. MIX-TO-TRACK SUBSEQUENCE ALIGNMENT

The objective of mix-to-track subsequence alignment is to find an optimal alignment path between a subsequence of a mix and a track used in the mix. This alignment result will be the basis of diverse DJ mix analysis concerning the cue point, track boundary, key/tempo changes and transition length. We also use it for removing non-matching tracks. This section describes the detail of computational process.

3.1 Feature Extraction

When DJs create a mix, they often adjust tempo and/or key of the tracks in the mix or add audio effects to them. Live mixes contain more changes in timbre and even other sound sources such as the voices from the DJ. In order to address the acoustic and musical variations between the original track and the matched subsequence in the mix, we use beat synchronous chroma and mel-frequency cepstral coefficients (MFCC). The beat synchronous feature representations enable tempo invariance and dramatically reduces the computational cost in the alignment. The aggregation of the features from the frame level to the beat level also smooths out local timbre variations. The chroma feature, on the other hand, facilitates key-invariance as circular shift of the 12-dimensional vector corresponds to key transposition. The MFCC feature captures general timbre characteristics. We used Librosa⁴ to extract the chroma and MFCC features with the default options except that the dimensionality of MFCC was set to 12 and the type of chroma was to chroma energy normalized statistics (CENS) [16].

3.2 Key-Invariant Subsequence DTW

We compute the alignment by applying subsequence DTW to the beat synchronous features [15]. We used an implementation from Librosa, adopting the transposition-invariant approach from [17]. Specifically, we calculated 12 versions of chroma features by performing all possible

³<https://mir-aidj.github.io/djmix-analysis/>

⁴<https://librosa.github.io/librosa/>

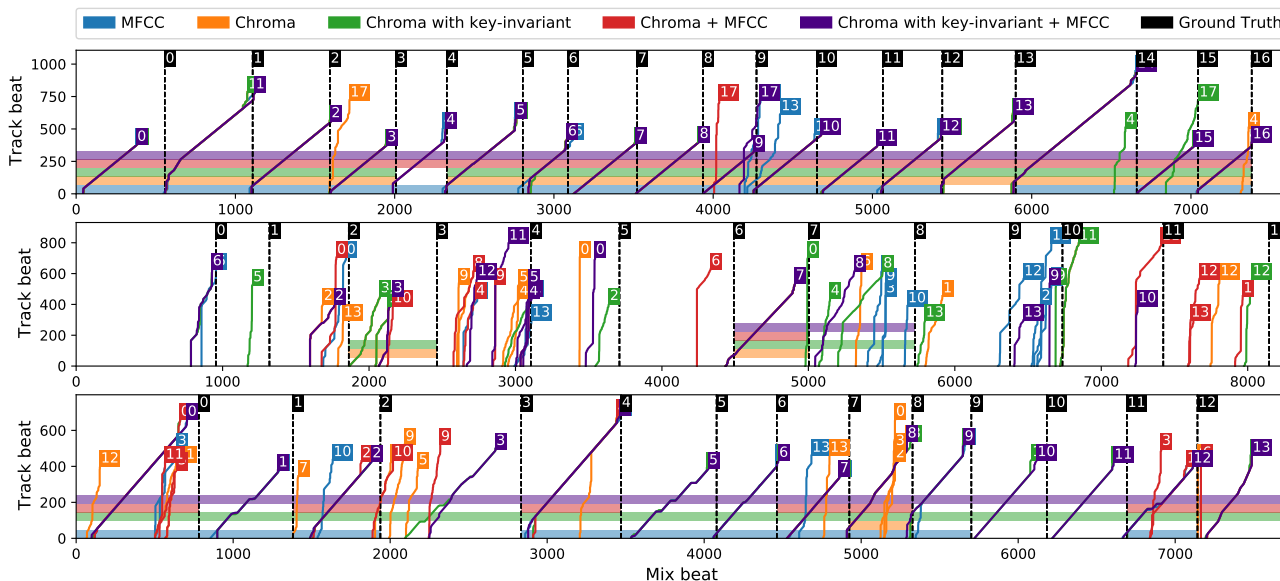


Figure 1. Visualizations of the result of a DTW-based mix-to-track subsequence alignment between a mix and the original tracks played in that mix. The colored solid lines show the warping paths of the alignment depending on the input feature, and whether or not applying the transposition-invariant method on the subsequence DTW. The tagged numbers on warping paths and ground truth boundaries indicate played and timestamped indices in the mix, respectively. A colored bar at the bottom of the figures is added if the alignment of the method is considered successful according to the match rate. (Top) A correctly matched example. (Middle) An unsuccessful example, due to the low sound quality of the mix. (Bottom) The alignment can be improved using the key-invariant chroma. Best viewed in color.

circular shifts on the original track side and select the one with the lowest matching cost in the subsequence DTW. This result returns not only the optimal alignment path but also the key transposition value of the original track.

Figure 1 shows three examples of the alignment results when different combinations of features (MFCC, chroma, and key-invariant chroma) are used. When the alignment path of the subsequence satisfies a match rate (described in Section 3.3), we put a color strip corresponding to each feature in the bottom of the figure. Since we use beat synchronous representations for them, the warping paths become diagonal with a slope of one if a mix and a track are successfully aligned. The top panel in the figure shows an successfully aligned example for the most of tracks and features where all warping paths have straight diagonal paths.⁵ The middle panel shows a failing example because sounds from crowds are also recorded in the mix.⁶ The bottom panel shows a example where chroma with circular shift distinctively works better others as the DJ frequently uses key transposition on the mix.⁷

3.3 Filtering Using Match Rates

As stated above, we can measure the quality of the alignment from the warping path. Ideally, when every single move on the path is diagonal, that is, one beat at a time for both track and mix axis, we will obtain a perfect straight diagonal line. However, the acoustic and musical changes deform the path. We define the ratio of the diagonal moves

in a mix (one move per beat) as the *match rate* and use it for filtering out incorrectly annotated tracks. We experimentally chose 0.4 as a threshold. The size of the dataset after the filtering is denoted as “Matched” in Table 1. We only use the matched tracks for the analysis in this paper.

4. CUE POINT EXTRACTION

Cue points are timestamps in a track that indicate where to start and end the track in a mix. Determining the cue points of played tracks is an essential task of DJ mixing. This section describes extracting cue points using the warping paths obtained from the aforementioned mix-to-track subsequence alignment.

4.1 Term Definitions

We first define terms related to cue points. In the context of the track-to-track transition, a *cue-out* point is a timestamp that the previous track starts fading out and the next track starts fading in, and a *cue-in* point is when the previous track is fully faded out and only the next track is being played. The *transition* region is defined as the time interval from the cue-out point of the previous track to the cue-in point of the next track. Additionally, we define a *cue-mid* point as the middle of a transition, which can technically be considered as a boundary of the transition.

4.2 Methods

The mix-to-track alignment results naturally yield cue points of matched tracks. Figure 2 shows an example of extracted cue points (a zoomed-in view of the top figure in

⁵ <https://1001.tl/14jltncet>

⁶ <https://1001.tl/15fulzcl>

⁷ <https://1001.tl/bcx2z0t>

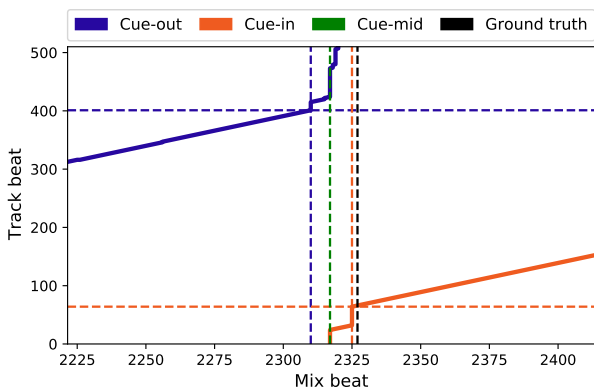


Figure 2. A zoomed-in view of a visualization of mix-to-track subsequence alignment explaining the three types of extracted cue points. The two solid lines indicate warping paths representing alignment between the mix and tracks. The vertical colored dotted lines represent the extracted cue points on the mix and the horizontal dotted lines represent the points on each track. The vertical black dotted line is a human-annotated ground truth boundary between the two tracks. The solid lines are the fourth and fifth warping paths from the top of Figure 1. Best viewed in color.

Figure 1). The two alignment paths drift from the diagonal lines in the transition region (between 2310 and 2324 in mix beat) because the two tracks cross-fades. Based on this observation, we detect the cue-out point of the previous track by finding the last beat where preceding 32 beats have diagonal moves in the alignment path. Likewise, we detect the cue-in point of the next track by finding the first beat where succeeding 32 beats have diagonal moves in the alignment path.

5. MIX SEGMENTATION

The goal of mix segmentation is to divide a continuous DJ mix into individual tracks, which can enhance the listening experience and can be a foundation of further analysis or learning of DJ mixes. Since DJs make seamless transitions, it is difficult to notice that a track is fading in or out. To quantitatively measure how difficult it is, a study analyzed how accurate humans are at creating the boundary timestamps and found that the standard deviation of the human disagreement for track boundaries in mixes is about 9 seconds, which implies it is difficult to find the optimal boundaries even for humans [12]. Furthermore, the ambiguous definition of the boundary and long lengths of transitions makes it difficult to annotate the boundary timestamps [8].

5.1 Cue Point based Estimation

Given the extracted cue point so far, we can estimate the track boundaries with three possible choices. The first is the position that the next track fully appears (cue-in point), the second is the position that previous track starts to disappear (cue-out point), and the last is the middle of the transition (cue-mid point). By comparing each of them

with human-annotated boundary timestamps, we can measure which type of cue point humans tend to consider as a boundary.

Figure 3 shows three histograms where each of them is computed from the differences between human-annotated boundary timestamps and one of the cue point types in beat unit. The overall trend shows that the distribution of cue-in point is mostly skewed towards zero. Interestingly, the distribution of cue-out point has more distinctive peaks around every 32 beat than the distribution of cue-in point. Considering the histogram of the transition length has peaks at every 32 beat as shown in Figure 6, this reflects that human annotators tend to label cue-in points as a boundary compared to cue-out (note that the transition length is computed by subtracting the cue-in point from the cue-out point). On the other hand, the distribution of cue-mid point has a gradually decreasing curve without peaks. While this distribution looks like having better estimates than the cue-out point, Table 3 shows an opposite result. That is, in terms of the number of cue points closest to the human annotations, the cue-out point is the second and the cue-mid point is the worst among the three types. These results indicate that the cue-mid point is a safe choice. That is, although the cue-mid point is least likely to be a boundary as shown in Table 3, the difference between the estimate and human annotation is relatively small because it is the middle of the transition region.

Table 2 shows the difference between human-annotated boundary timestamps and one of the cue point types in terms of median time (in seconds) on the left side. The overall trend confirms that the cue-in is the best estimate of track boundary and the cue-mid is a safer choice than the cue-out. The table also shows the result of “cue-best”. This is computed with the minimum difference among the three cue point types for each of the transition region. The result shows that the median time differences are dramatically decreased to 4-5 seconds. Table 2 also shows the difference between human-annotated boundary timestamps and the cue-in point in terms of hit rates on the right side. The hit rates are computed the ratio of correct estimates given a tolerance window. If the estimate is within the tolerance window on the human-annotated boundary timestamp, it is regarded as a correct estimate. We set three tolerance windows (15, 30, and 60 seconds) considering that the average tempo of tracks in the dataset is 127 beat per minute (BPM) and then the tolerance windows approximately correspond to 32, 64, 128 beats (multiples of a phrase unit). The result shows that the best hit rate with the 30 second window (about 64 beats) is above 80%. Given the long transition time as shown in Figure 6, the cue-in point may be considered as a reasonable choice.

5.2 Effect of Audio Features

Table 2 also compares the median time difference between human-annotated boundary timestamps and one of the cue point types for different audio features used in the subsequence DTW. In general, the chroma features are a better choices than MFCC (p-value of t-test < 0.001 for chroma

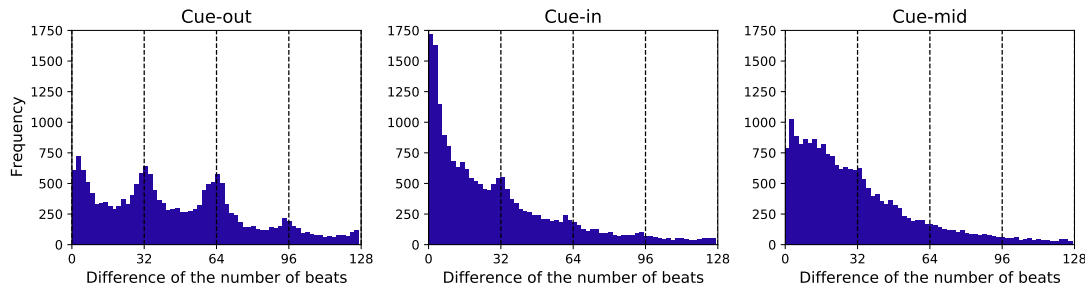


Figure 3. Histograms of distances to ground truth boundaries in the number of beats depending on the type of the cue point. The dotted lines are plotted at every 32 beats which is usually considered as a phrase in the context of dance music.

Feature	Median time difference (in seconds)				Cue-in hit rate		
	Cue-out	Cue-in	Cue-mid	Cue-best [†]	15 sec	30 sec	60 sec
MFCC	27.92	14.27	13.55	5.340	0.5187	0.7591	0.9023
Chroma	23.85	11.80	12.33	4.230	0.5837	0.7973	0.9286
Chroma with key-invariant	23.87	11.77	12.37	4.240	0.5843	0.7968	0.9282
Chroma + MFCC	23.41	11.48	12.16	4.380	0.5866	0.8035	0.9284
Chroma with key-invariant + MFCC	23.38	11.40	12.16	4.380	0.5881	0.8040	0.9288

Table 2. Mix segmentation performances depending on the type of cue point and the input feature used to obtain the warping paths. Median time differences between cue points and ground truths are shown on the left side and hit rates of cue-in points with thresholds in seconds are shown on the right side. “Key-invariant” indicates applying the key transposition-invariant method for the DTW. The best score of each criteria is shown in **bold**. † indicates the scores are computed using the best score among the three cue types.

Cue-out	Cue-in	Cue-mid
6,151 (30%)	10,844 (52%)	3,770 (18%)

Table 3. The number of ground truth boundary timestamps closest to the type of cue point.

with or without key-invariant). When both of chroma and MFCC are combined, the median time difference slightly reduces but it is statistically insignificant (p-value of t-test > 0.1). However, we observed that the subsequence DTW does not work well for some genres such as Techno which only contain drum and ambient sounds. This might can be improved by using MFCCs with a large number of bins or using mel-spectrograms. The use of key-invariant chroma generally does not make much difference because key transposition does not performed frequently as discussed in Section 6.2.

6. MUSICOLOGICAL ANALYSIS OF DJ MIXES

We hypothesize that DJs share common practices in the creative process in terms of tempo change, track-to-track transition, and cue point selection. In this section, we validate them using the results from the mix-to-track subsequence alignment and the cue point extraction.

6.1 Tempo Adjustment

We compare the estimated tempo of the original track to the tempo of each audio segment where the track is played in a mix. Figure 4 shows a histogram of percentage dif-

ferences of tempo between the original track and the audio segment in the mix. For example, a difference of 5% indicates the tempo of the original track is increased by 5% while played in the mix. As shown in the histogram, the adjusted tempo has an double exponential distribution, which means the adjusted tempo values are skewed towards zero. In detail, 86.1% of the tempo are adjusted less than 5%, 94.5% are less than 10%, and 98.6% are less than 20%. If one implements an track identification system for DJ mix that is robust to tempo adjustment, this distribution could be a reference.

6.2 Key Transposition

A function so-called “master tempo” or “key lock” that preserves pitch despite tempo adjustments is activated by default in modern DJ systems such as stand-alone DJ systems, DJ softwares, and even turntables for vinyl records. Therefore, key transposition is usually performed when a DJ intentionally wants to change the key of a track. As mentioned in Section 3.2, the transposition-invariant DTW can provide the number of transposed semitones as a by-product. We computed the statistics of key transposition using them (using DTW taking both MFCCs and key-invariant chroma). Figure 5 shows a histogram of key transposition between the original track and the audio segment in the mix. Only 2.5% among the total 24,202 tracks are transposed and, among those transposed tracks, 94.3% of them are only one semitone transposed. This result indicates that DJs generally do not perform key transposition much and leave the “master tempo” function turned on in most cases.

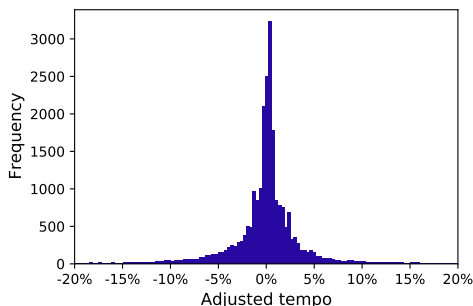


Figure 4. A histogram of adjusted tempo of tracks in mixes.

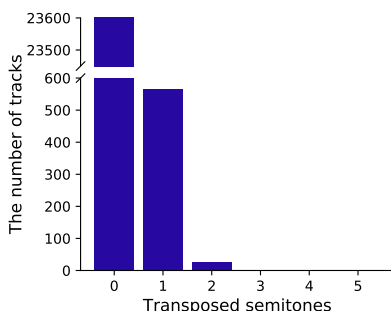


Figure 5. The number of tracks depending on the number of semitones in mixes.

6.3 Transition Length

Once we extract cue-in and cue-out points in the transition region, we can calculate the transition length. This can provide some basic hints on how DJ makes the track-to-track transition in a mix. Figure 6 shows a histogram of transition lengths in the number of beats. We annotated the dotted lines every 32 beat which is often considered as a phrase in the context of dance music. The histogram has peaks at every phrase. This indicates that DJs consider the repetitive structures in the dominant genres of music when they make transitions or set cue points.

6.4 Cue Point Agreement among DJs

Deciding cue points of played tracks is a creative choice in DJ mixing. Observing the agreement of cue points on a single track among DJs may elucidate the possibility of finding some common rules. To the end, we collected all extracted cue points for each track and computed the statistics of deviations in cue-in points and cue-out points among DJs. Specifically, we computed all possible pairs and their distances separately for cue-in points and cue-out points. Since the two distributions were almost equal, we combined them into a single distribution in Figure 7. From the results, 23.6% of the total cue point pairs have zero deviation. 40.4% of them were within one measure (4 beats), 73.6% were within 8 measures and 86.2% were within 16 measures. This indicates that there are some rules that DJs share in deciding the cue points. It would be interesting to perform detailed pattern analysis to estimate the cue points using this data in future work.

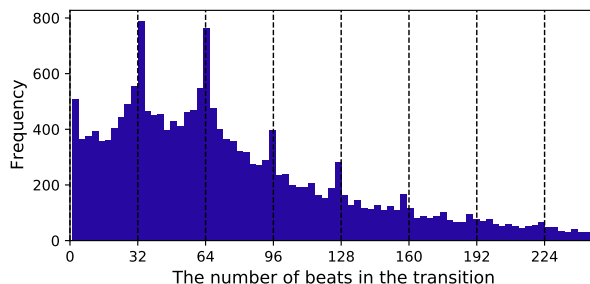


Figure 6. A histogram of the transition lengths in number of beats. The dotted lines are plotted at every 32 beats.

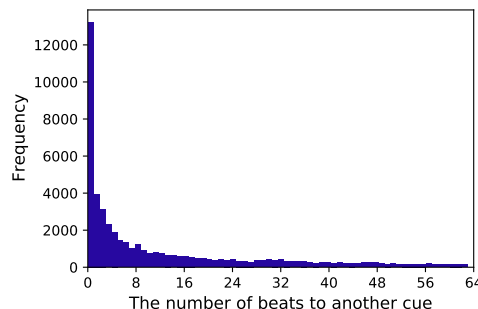


Figure 7. A histogram of distances between cue points of a single track in the number of beats.

7. CONCLUSIONS

We presented various statistics and analysis of 1,557 real-world DJ mixes from *1001Tracklists*. Based on the mix-to-track subsequence DTW, we conducted cue point analysis of individual tracks in the mixes and showed the possibility of common rules in the music making that DJs share. We also investigated mix segmentation by comparing the three types of cue point to human-annotated boundary timestamps and showed that humans tend to recognize cue-in points of the next tracks as boundaries. Finally, we showed the statistics of tempo and key changes of the original tracks in DJ performances. We believe this large-scale statistical analysis of DJ mixes can be beneficial for computer-based research on DJ music. The cue point analysis can be the ground for the precise definition of cue points and the tempo and key analysis can provide a guideline of the musical changes during the DJ mixing.

As a future work, we plan to estimate cue points within a track as a step towards automatically generating a mix [3, 4]. The cue point estimation has many application such as DJ software and playlist generation on music streaming services. This will require structure analysis or segmentation of a single music track, which is an important topic in MIR. Furthermore, we plan to analyze the transition region in a mix to investigate DJ mixing techniques. For example, it is possible to estimate the gain changes in the cross-faded region by comparing the two adjacent original tracks and the mix [13, 14]. The methods can be extended to the spectrum domain. Such detailed analysis of mixing techniques will allow us to understand how DJs seamlessly concatenate music tracks and provide a guide to develop automatic DJ systems.

8. ACKNOWLEDGEMENT

We greatly appreciate *1001Tracklists* for offering us the mix metadata employed in this study. We note that the metadata used for this analysis was obtained with permission from *1001Tracklists*, and suggest that people who are interested in the data contact *1001Tracklists* directly. This research was supported by BK21 Plus Postgraduate Organization for Content Science (or BK21 Plus Program), Basic Science Research Program through the National Research Foundation of Korea (NRF-2019R1F1A1062908), and a grant from the Ministry of Science and Technology, Taiwan (MOST107-2221-E-001-013-MY2).

9. REFERENCES

- [1] Y.-T. Lin, C.-L. Lee, J.-S. Jang, and J.-L. Wu, "Bridging music via sound effects," in *2014 IEEE International Symposium on Multimedia*. IEEE, 2014, pp. 116–122.
- [2] R. M. Bittner, M. Gu, G. Hernandez, E. J. Humphrey, T. Jehan, H. McCurry, and N. Montecchio, "Automatic playlist sequencing and transitions," in *Proc. International Society for Music Information Retrieval Conference (ISMIR)*, 2017, pp. 442–448.
- [3] D. Schwarz, D. Schindler, and S. Spadavecchia, "A heuristic algorithm for DJ cue point estimation," in *Proc. Sound and Music Computing (SMC) Conference*, 2018.
- [4] L. V. Veire and T. De Bie, "From raw audio to a seamless mix: creating an automated DJ system for drum and bass," *EURASIP Journal on Audio, Speech, and Music Processing*, vol. 2018, no. 1, p. 13, 2018.
- [5] A. Kim, S. Park, J. Park, J.-W. Ha, T. Kwon, and J. Nam, "Automatic DJ mix generation using highlight detection," in *International Society for Music Information Retrieval Conference (ISMIR), Late-Breaking Paper*, 2017.
- [6] Y.-S. Huang, S.-Y. Chou, and Y.-H. Yang, "Generating music medleys via playing music puzzle games," in *Proc. AAAI Conference on Artificial Intelligence*, 2018.
- [7] —, "DJnet: A dream for making an automatic DJ," in *International Society for Music Information Retrieval Conference (ISMIR), Late-Breaking Paper*, 2017.
- [8] R. Sonnleitner, A. Arzt, and G. Widmer, "Landmark-based audio fingerprinting for DJ mix monitoring," in *Proc. International Society for Music Information Retrieval Conference (ISMIR)*, 2016, pp. 185–191.
- [9] P. S. Manzano, "Audio fingerprinting techniques for sample identification in electronic music," Master's thesis, Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU), Erlangen, Germany, 2016.
- [10] P. Lopez Serrano Erickson, "Analyzing sample-based electronic music using audio processing techniques," Ph.D. dissertation, Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU), Erlangen, Germany, 2019.
- [11] N. Glazyrin, "Towards automatic content-based separation of DJ mixes into single tracks," in *Proc. International Society for Music Information Retrieval Conference (ISMIR)*, 2014, pp. 149–154, source code available at [Online] <https://github.com/nglazyrin/MixSplitter>.
- [12] T. Scarfe, W. Koolen, and Y. Kalnishkan, "Segmentation of electronic dance music," *International Journal of Engineering Intelligent Systems for Electrical Engineering and Communications*, vol. 22, no. 3, p. 4, 2014, source code available at [Online] <https://github.com/ecsplendid/DanceMusicSegmentation>.
- [13] L. Werthen-Brabants, "Ground truth extraction & transition analysis of DJ mixes," Master's thesis, Ghent University, Ghent, Belgium, 2018.
- [14] D. Schwarz and D. Fourer, "Methods and datasets for DJ-mix reverse engineering," in *Proc. International Symp. on Computer Music Multidisciplinary Research (CMMR)*, 2019, pp. 426–437.
- [15] M. Müller, *Fundamentals of music processing: Audio, analysis, algorithms, applications*. Springer, 2015, ch. 7.2.3. Subsequence DTW.
- [16] M. Müller and S. Ewert, "Chroma toolbox: MATLAB implementations for extracting variants of chroma-based audio features," in *Proc. International Society for Music Information Retrieval Conference (ISMIR)*, 2011.
- [17] M. Müller and M. Clausen, "Transposition-invariant self-similarity matrices," in *Proc. International Society for Music Information Retrieval Conference (ISMIR)*, 2007, pp. 47–50.