

# Drum Groove Corpora

FRED HOSKEN[1]

*Georgetown University, United States of America*

TONI BECHTOLD

*Lucerne University of Applied Sciences and Art, Switzerland*

FLORIAN HOESL

*Lucerne University of Applied Sciences and Art, Switzerland*

LORENZ KILCHENMANN

*Lucerne University of Applied Sciences and Art, Switzerland*

OLIVIER SENN

*Lucerne University of Applied Sciences and Art, Switzerland*

**ABSTRACT:** Patterned microtiming deviations from metronomic regularity are ubiquitous in the performance of metered music. The relevance of microtiming to the perception of music has been studied since the 1980s. Most recently, microtiming has been investigated as a cause of groove (i.e., the pleasant urge to move in response to music). The study of microtiming relies on the availability of microtiming data. This report presents three large corpora of onset timings derived from drum kit performances in popular Anglo-American popular music styles. These data are made freely available (CC 4.0 license) to provide a resource for use by analysts and experimenters alike. They offer a common point of reference for future studies into the temporal facets of music performance. The datasets adhere to FAIR principles; they thus facilitate replication of analyses and experimental stimuli.

Submitted 2020 April 30; accepted 2020 September 16.

Published 2021 December 10; <https://doi.org/10.18061/emr.v16i1.7642>

**KEYWORDS:** *popular music, drum kit, drum pattern, microtiming, timing data, rhythm*

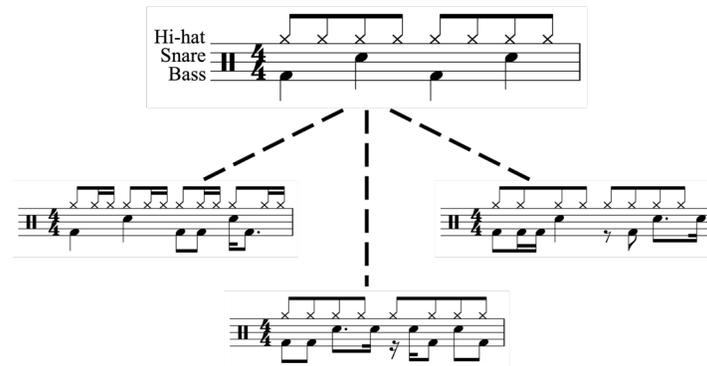
## BACKGROUND

HUMANS perform metered music with considerable temporal flexibility. Minute timing deviations from strict metronomic regularity of the order of tens of milliseconds are commonly referred to as microtiming. On a basic level, microtiming arises from the restricted accuracy of human motor performance (Rasch, 1988), yet motor jitter alone cannot explain the considerable size and consistent patterning of microtiming deviations in competently performed music. This indicates that microtiming serves aesthetic and expressive purposes relevant to style/genre (e.g., Johansson, 2010; Naveda et al., 2011) and individual performers (e.g., Prögler, 1995). Research has shown that *tempo rubato* in the performance of European art music clarifies the form of a piece (Clarke, 1985; Repp 1998, Juslin et al., 2001; Senn et al., 2012). Microtiming in Western popular music has been linked to elusive perceptual qualities such as music's *feel* (e.g., laid-back feel, Kilchenmann & Senn, 2011) and *groove* (the pleasurable urge to move in response to the music, Janata et al., 2012) but empirical results have been inconsistent (for an overview, see Senn et al., 2017; Hosken, 2020).

This report presents three corpora of timing data drawn from drum set performances in Anglo-American popular music styles. It contributes to a growing corpus of datasets that allow for the comprehensive study of microtiming phenomena across a variety of musical styles and genres, such as Cuban son and salsa (Poole et al., 2018; Maia et al., 2019), classical string quartet (Clayton & Tarsitani, 2019), solo piano (Goebel, 2001), Malian jembe (Polak et al., 2018), Western popular music drumming (Gillet & Richard, 2006), and others.



This collection of timing data about drum set performance is made freely available in a format that is findable, accessible, interoperable, and reusable (FAIR Principles, Wilkinson et al., 2016) in the hope to offer a common resource to timing research that may be used in a consistent and replicable way across studies. It comprises of 62 drummers performing over 19,000 bars, nearly 700 minutes of drumming, in Anglo-American popular music styles. The timing data have been collected, processed, and filtered to offer machine- and human-readable information about how drummers perform. All of the drum patterns collated here have as their background structure the “archetypical rock beat” (Tamlyn, 1998, p. 11, see Figure 1). This archetype underlies a variety of drum patterns that feature in a range of musical styles such as pop, rock, funk, soul, disco, hip hop, drum and bass, and techno.



**Fig. 1.** The archetypical Anglo-American popular music drum groove (top) and, connected via dashed lines, examples of drum patterns that derive from it.

## METHODS

### Data Sources

This collection features one novel corpus of drum grooves (*The Loop Loft*) as well as a repackaging of two existing corpora (*Lucerne Groove Research Library* and *Google Magenta’s Groove MIDI Dataset*). Timing data about the three corpora have been obtained using different methods and each corpus’s drummers performed under different conditions.

The three sources are:

- 1) *The Loop Loft* (<https://www.thelooploft.com/>) is a commercial sample shop that provides short loops for DJs and producers to use in their creative work. The company invites performers into the studio to record short drum patterns while listening to a click track. Audio files for each microphone placed on each instrument within the drum kit are available allowing for clear identification of which drum was struck at what time. All audio is provided dry (no EQ, compression, reverb, etc.). Information about microphone types or proximity of microphones to the drumheads is unavailable. The audio of the click tracks is not available, so the exact timing of click track events is unknown. Each track title contains tempo information in bpm. Here, 1,467 tracks performed by four world-famous session musicians were purchased and analyzed using the `mirevents` function of the `MIRtoolbox` in `MATLAB` (Lartillot et al., 2008). Onsets below a threshold of 10% the maximum amplitude were discarded to remove bleed from other drums. Measurement precision is expected to be  $< 1$  ms.
- 2) The *Lucerne Groove Research Library* (<https://www.grooveresearch.ch/>) is a corpus of 251 drum grooves drawn from commercial recordings played by 50 highly acclaimed drummers in the fields of pop, rock, funk, soul, disco, R&B, and heavy metal. Two professional musicians transcribed the drum patterns by ear and manually identified each drum onset using spectrograms and oscillograms in `LARA` software (`www.hslu.ch/lara`, version 2.6.3). Onset measurement is estimated to be accurate to  $\pm 3$  ms for most of the music excerpts and, even in the most problematic cases, the timing measurement error is expected to rarely exceed  $\pm 10$  ms (see Senn et al., 2018 for full method). Drum patterns are provided in

MIDI and MP3 format. Since these drum performances are part of full-band recordings (i.e., not just the drums in isolation) drawn from 1956 to 2014, it is not knowable whether a click track was used in the performance, nor the precise location in time of a click track if one was used.

- 3) *Google Magenta's Groove MIDI Dataset* (<https://magenta.tensorflow.org/datasets/groove>) is a corpus of 503 drum patterns performed on a Roland TD-11 electronic drum kit by five professional drummers and four amateur players (Google employees). Drummers, who are anonymized in the set and referred to only by ID number, played on this MIDI drum kit to a click track. The TD-11 has a temporal resolution of 480 MIDI ticks per quarter note, so the lowest resolution (for a performance recorded at 50 BPM) is 2.5 ms and the mean resolution of all performances is 1.17 ms. The drummers performed drum patterns and solos for as long as they desired. This corpus was initially created as training data for a machine learning project into expressive drum performances (Gillick et al., 2019). The audio of the click tracks is not available, so the exact timing of click track events is unknown. The track titles inform about the tempo of each track (bpm).

### Data Filtering

The onsets of all corpora were filtered according to the exclusion criteria summarized in Figure 2.

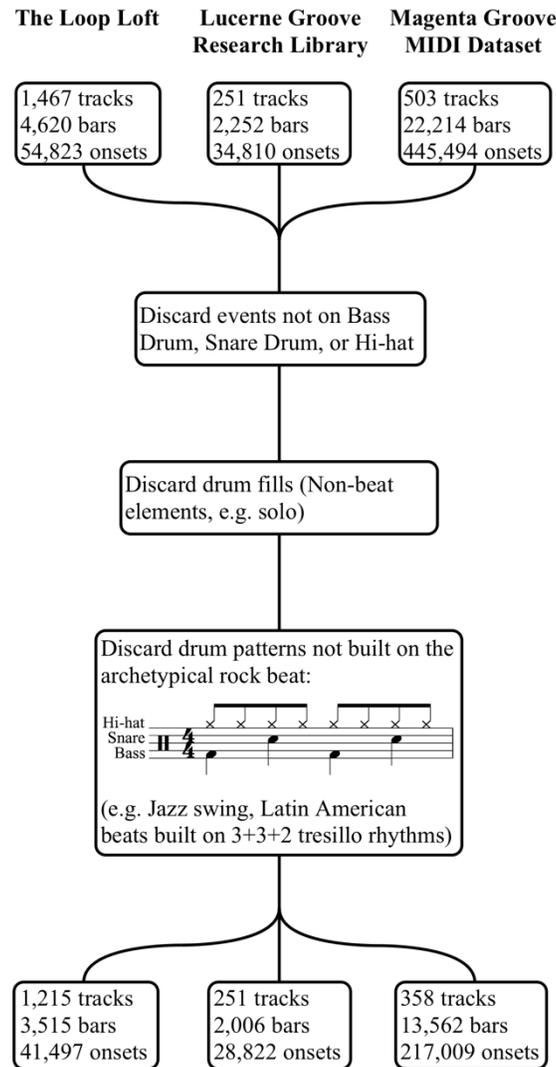
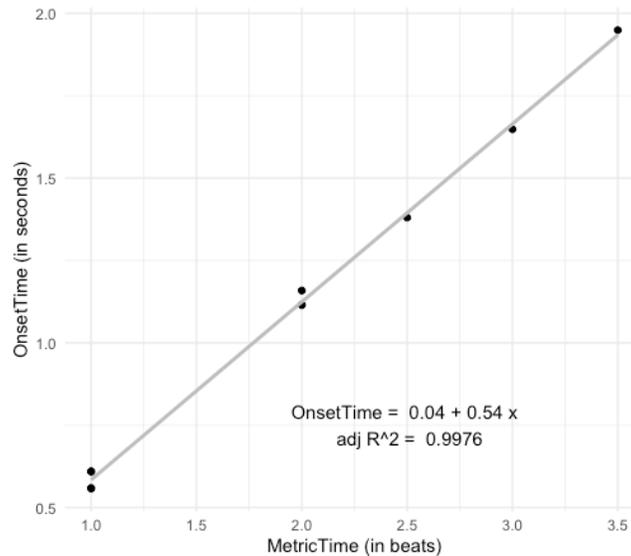


Fig. 2. Schematic representation of the data filtering process.

The differences in how the filtering process affected different corpora may be attributed to the unique nature of the performances captured by each set. For example, the *Lucerne* corpus is comprised entirely of the “core grooves” of famous songs in typical rock/pop styles and so no complete tracks were excluded from the set. The *Loop Loft* corpus, however, features numerous tracks explicitly labelled “Fill” and numerous tracks in Latin American styles, so several complete tracks were excluded. Likewise, the *Magenta* performances are in a range of styles, and often feature extended drum solos and groove patterns that are based around the tom-toms (these bars are filtered out). Unfiltered data sets are available upon request.

### Estimation of Microtiming Deviations

For each drum onset, the metric position within the track (*MetricTime*, in beats with divisions every 16<sup>th</sup> note) and an exact onset time (*OnsetTime*, in seconds) are available. The drummers in the *Loop* and *Magenta* corpora utilized click tracks throughout the recording process, but the onset times of the clicks are unavailable. Therefore, we used simple linear regression models to predict *OnsetTime* from *MetricTime* (see Figure 3).



**Fig. 3.** Example of the linear regression used to model the *MetronomicOnsetTime*. Points represent individual drum onsets, and the grey line is the fitted regression line.

The fitted values of the regression models provided a metronomic reference (*MetronomicOnsetTime*, in seconds), and microtiming deviations (*MicrotimingSeconds*) were calculated as the difference between *OnsetTime* and *MetronomicOnsetTime* (see Equation 1).

$$\text{MicrotimingSeconds} = \text{OnsetTime} - \text{MetronomicOnsetTime}. \quad (1)$$

The slopes of the regression models were used to calculate the *Tempo* of the track. This, in turn, allowed to calculate *MicrotimingBeats* (see Equation 2).

$$\text{MicrotimingBeats} = \frac{\text{Tempo}}{60} \times \text{MicrotimingSeconds} \quad (2)$$

The performances of the *Lucerne* corpus were taken from commercial recordings that potentially exhibit tempo drift and may not have been recorded to a click track. Quadratic regression models were fitted

to the data to account for potential tempo drift. Higher polynomial orders were not found to significantly improve the fit of the model.

Regressions were chosen to model the metronomic beat locations over established beat-finding procedures that use, for example, recurrent neural networks in conjunction with autocorrelation and comb filters (e.g., madmom, Böck et al. 2016) for reasons of parsimony, replicability, and to minimize the addition of noise into the data. By relying only on the onset timing data and information about which location on the metric grid this onset corresponds to, we also avoid involving heuristics that may be sub-optimal, for example some beat tracking algorithms require assumptions to be made based on extra musical information (e.g., about style or genre – Böck et al. 2014). For our corpora, the variance unexplained by the simple beat locating model is minimal, with only two tracks (out of 1,824) having an adjusted  $R^2$  lower than .995. Regression model fit estimates for the three corpora can be found in Table 1.

## THE DRUM GROOVE CORPORA DATASET

### Data Structure

The *Drum Groove Corpora* dataset is made available online at <https://osf.io/3sejt/> in comma-separated values format (*DrumGrooveCorpora.csv*). For every onset (287,328 rows), the following data (13 columns) are provided:

- *Corpus*: the corpus to which the onset belongs (*Loop, Lucerne, Magenta*).
- *Drummer*: the name of the drummer (for the *Loop* and *Lucerne* corpora) or an uppercase letter where the name is unknown (*Magenta*).
- *Track*: a unique track name.
- *Year*: the recording year of the track (*Lucerne* and *Magenta* only; *Loop* recordings are all post-2010, though specific details are not available).
- *Strike*: indexes onsets of a track in the order of their *OnsetTime*.
- *Instrument*: the instrument of the drum kit on which the stroke is played (HH = Hi-Hat, SD = Snare Drum, or BD = Bass Drum).
- *MetricTime*: the metric time of the stroke (in beats), where 0.00 is bar 1, beat 1; 0.25 is one sixteenth note later; and 4.00 is bar 2, beat 1.
- *OnsetTime*: the onset time of the stroke (in seconds) measured from the start of the audio/MIDI recording.
- *MetronomicOnsetTime*: the metronomic onset time (in seconds), estimated by linear regression.
- *Tempo*: the tempo (in beats per minute), estimated by linear regression.
- *MicrotimingSeconds*: the difference between *OnsetTime* and *MetronomicOnsetTime* (in seconds).
- *MicrotimingBeats*: the microtiming deviation as a proportion of the beat.
- *TrackDuration*: the time difference between the first and last onset (in seconds).

## Data Description

Table 1 provides descriptive statistics for the three corpora. The *Magenta* corpus is by far the largest of the three (217,009 onsets), followed by the *Loop* (41,497 onsets) and *Lucerne* (28,822 onsets) corpora. *Magenta* also consists of relatively long tracks (37.88 bars on average) compared to the very short *Loop* drum performances (2.89), and the eight-bar period excerpts of the *Lucerne* corpus.

**Table 1.** Descriptive statistics for the three corpora.

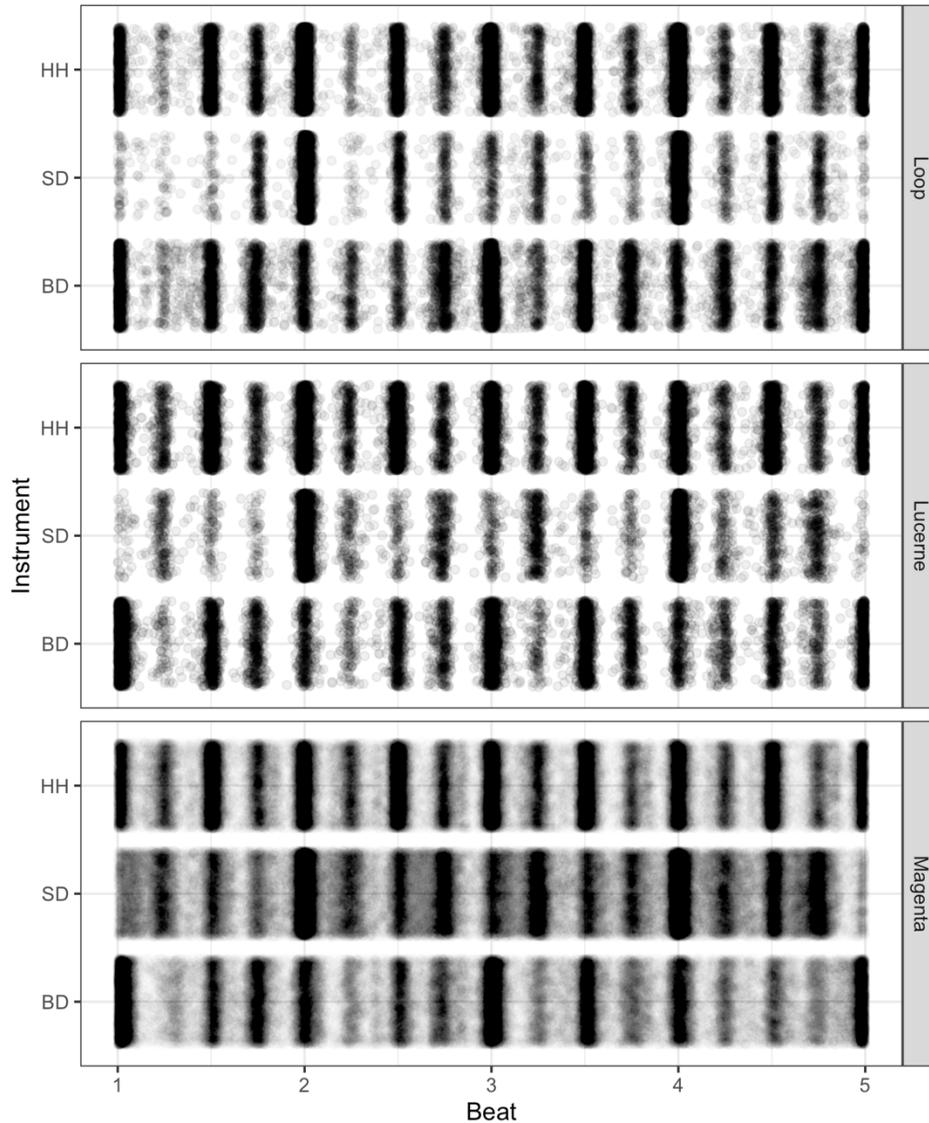
	Loop	Lucerne	Magenta
<b>Drummers<sup>a</sup></b>	4	50	9
<b>Tracks</b>	1,215	251	358
<b>Bars</b>	3,515	2,006	13,562
Mean bars per track	2.89	7.99	37.88
<b>Onsets</b>			
Hi-hat	19,102 (46%)	13,899 (48%)	72,250 (33%)
Snare drum	9,174 (22%)	6,317 (22%)	82,512 (38%)
Bass drum	13,221 (32%)	8,606 (30%)	62,247 (29%)
Total	41,497 (100%)	28,822 (100%)	217,009 (100%)
<b>Mean tempo (BPM)</b>	103.30	115.40	104.80
<b>Microtiming (in beats)</b>			
Mean	0.0000	0.0000	0.0000
Standard deviation	0.0277	0.0283	0.0509
Skewness	0.9732	-0.9481	-0.4165
Excess kurtosis	7.7995	7.7338	0.5998
Minimum	-0.1817	-0.2948	-0.1783
Maximum	0.1938	0.1820	0.1759
<b>Regression model fit (adj. <math>R^2</math>)</b>			
Mean	.9999	.9999	.9999
Minimum	.9859	.9998	.9974
Maximum	>.9999	>.9999	>.9999

*Note.* <sup>a</sup> The *Loop* and *Lucerne* corpora both feature performances by Omar Hakim, hence the total number of drummers in the data set = 62.

The *Loop* and *Lucerne* corpora show considerable similarities, while the *Magenta* dataset seems to stand apart. For example, the proportions of hi-hat, snare drum, and bass drum strokes are comparable across the *Loop* and *Lucerne* sets, with a large proportion of strokes played by the hi-hat, followed by the bass drum and the snare drum. In the *Magenta* set, the snare drum strokes are in a majority, followed by hi-hat, and bass drum events.

The archetypical rock beat is strongly articulated in all three datasets: the bass drum strokes on beats one and three, the snare drum hits on the backbeats (two and four), and the regular eighth-note pulse of the hi-hat show a robust presence throughout all three *Drum Groove Corpora* (Figure 4). It can be seen, however, that the *Magenta* snare drum pattern is quite versatile, whereas, in the *Loop* and *Lucerne* corpora, the main role of the snare drum is to articulate the backbeats.

The mean microtiming deviations are zero for all three corpora (Table 1). This is a direct consequence of the regression approach taken to fit the *MetronomicOnsetTime* to each track. But the standard deviation of the microtiming deviations (*MicrotimingBeats*) in the *Magenta* corpus (0.0509 beats) is substantially greater than the standard deviations in the *Loop* (0.0277 beats) and *Lucerne* corpora (0.0283 beats). These differences in “tightness” are visible in Figure 4: the markings on each metric position are temporally (i.e., horizontally) more spread out in the *Magenta* corpus compared to the other two corpora.



**Fig. 4.** Strip plots of all onsets for bass drum (BD), snare drum (SD), and hi-hat (HH) for each corpus. Vertical spread within each strip is random jitter.

## CONCLUSIONS

This data report describes the *Drum Groove Corpora* dataset, a large collection of data that adheres to the FAIR principles and focuses on rhythmic patterns and microtiming deviations in Anglo-American popular music drumming. The *Drum Groove Corpora* dataset will prove to be useful in a variety of research settings, for example:

### Modelling of microtiming deviations

The dataset provides a basis to investigate and model the temporal processes that generate microtiming deviation patterns in popular music drumming. The *Loop* and *Lucerne* corpora invite study of microtiming profiles for different drummers, the three instruments of the drum kit, and different metric positions within the pattern. The *Magenta* corpus allows timing processes across longer performances to be studied, such as the elastic relationship between the performed tempo and the regular metronome click (akin to *tempo rubato*).

## Study of rhythmic patterns

The dataset allows for the study and classification of drum patterns that are derived from the archetypical rock beat. The *Lucerne* corpus offers a historic perspective in the sense that it represents a broad cross-section from more than half a century of popular music drumming. The *Loop* corpus focuses on recent performances by only four drummers and thus allows concentration on drummers' personal preferences and idiosyncrasies.

## Selection of experimental stimuli

The dataset may serve as metadata to help researchers screen and find experimental stimuli for listening experiments focused on microtiming deviations or rhythmic patterns. The audio files themselves can then be obtained from the respective sources:

- *Loop* – <https://www.thelooploft.com/>
- *Lucerne* – <https://www.grooveresearch.ch/index.php?downloads>
- *Magenta* – <https://magenta.tensorflow.org/datasets/groove>. See also the “Expanded Groove MIDI Dataset” that recreates all Magenta performances on 43 different virtual drum kits (Callender et al., 2020).

The data presented here was sourced from three corpora that show considerable differences with respect to the genre background of the music, popular music era, recording circumstances, performance duration, and the expertise of the drummers. These different conditions are reflected in the rhythmic and microtiming patterns enclosed in the three corpora and should be considered when the data is used for analysis.

## AUTHOR CONTRIBUTIONS

Fr.H. and O.S. prepared the manuscript; Fr.H. prepared the Loop Loft corpus; T.B., Fl.H. and L.K. prepared the Lucerne corpus and commented on the manuscript.

## SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://osf.io/3sejt/>. These materials are available under a CC 4.0 license: users are free to share, copy, redistribute, and adapt data as long as appropriate credit to this source is provided and any adaptations to the data are explicitly indicated.

## ACKNOWLEDGMENTS

This article was copyedited by Matthew Moore and layout edited by Diana Kayser.

## NOTES

[1] Correspondence can be addressed to: Fred Hosken, Georgetown University, 3700 O St NW, Washington, DC 20057. E-mail: [fred.hosken@georgetown.edu](mailto:fred.hosken@georgetown.edu).

## REFERENCES

- Böck, S., Korzeniowski, F., Schlüter, J., Krebs, F., & Widmer, G. (2016). madmom: A New Python Audio and Music Signal Processing Library. *Proceedings of the 24th ACM International Conference on Multimedia*, 1174–1178. <https://doi.org/10.1145/2964284.2973795>
- Böck, S., Krebs, F., & Widmer, G. (2014). A Multi-Model Approach to Beat Tracking Considering Heterogeneous Music Styles. *In Proceedings of the 15th International Society for Music Information Retrieval Conference (ISMIR)*.

Clarke, E. F. (1985). Structure and Expression in Rhythmic Performance. In R. West, I. Cross, & P. Howell (Eds.), *Musical Structure and Cognition* (pp. 209–236). London: Academic Press.

Callender, L., Hawthorne, C., & Engel, J. (2020) Improving Perceptual Quality of Drum Transcription with the Expanded Groove MIDI Dataset. <https://arxiv.org/pdf/2004.00188.pdf>

Clayton, M., & Tarsitani, S. (2019). *IEMP European String Quartet: A Collection of Audiovisual Recordings of String Quartet Performances, With Annotations*. Open Science Framework (OSF). <https://osf.io/usfx3/> [Data].

Gillet, O., & Richard, G. (2006). *ENST-Drums: An extensive audio-visual database for drum signals processing*. [Data].

Gillick, J., Roberts, A., Engel, J., Eck, D., & Bamman, D. (2019). Learning to Groove with Inverse Sequence Transformations. *Proceedings of the 36th International Conference on Machine Learning*. International Conference on Machine Learning (ICML), Long Beach, California. <http://arxiv.org/abs/1905.06118v2>

Goebel, W. (2001). Melody Lead in Piano Performance: Expressive Device or Artifact? *Journal of the Acoustical Society of America*, 110, 563–572. <https://doi.org/10.1121/1.1376133>

Hosken, F. (2020). The Subjective, Human Experience of Groove: A Phenomenological Investigation. *Psychology of Music*, 48(2), 182–198. <https://doi.org/10.1177/0305735618792440>

Janata, P., Tomic, S., & Haberman, J. (2012). Sensorimotor Coupling in Music and the Psychology of the Groove. *Journal of Experimental Psychology: General*, 141(1), 54–75. <https://doi.org/10.1037/a0024208>

Johansson, M. (2010). The Concept of Rhythmic Tolerance: Examining Flexible Grooves in Scandinavian Folk Fiddling. In A. Danielsen (Ed.), *Musical Rhythm in the Age of Digital Reproduction* (pp. 69–83). Ashgate. <https://doi.org/10.4324/9781315596983-5>

Juslin, P. N., Bresin, R., & Friberg, A. (2001). Toward a Computational Model of Expression in Performance: The GERM model. *Musicae Scientiae (Special Issue 2001/2002: Current Trends in the Study of Music and Emotion)*, 63–122. <https://doi.org/10.1177/10298649020050S104>

Kilchenmann, L., & Senn, O. (2011). Play in time, but don't play time: Analyzing Timing Profiles of Drum Performances. *Proceedings of the International Symposium on Performance Science 2011*, 593–598.

Lartillot, O., Toiviainen, P., & Eerola, T. (2008). A Matlab Toolbox for Music Information Retrieval. In C. Preisach, H. Burkhardt, L. Schmidt-Thieme, & R. Decker (Eds.), *Data Analysis, Machine Learning and Applications: Proceedings of the 31st Annual Conference of the Gesellschaft für Klassifikation e.V., Albert-Ludwigs-Universität Freiburg, March 7-9, 2007* (pp. 261–268). Springer-Verlag.

Maia, L. S., Fuentes, M., Biscainho, L. W. P., Rocamora, M., & Essid, S. (2019). *SAMBASET: A Dataset of Historical Samba de enredo Recording for Computational Music Analysis*. 20. <https://www.colibri.udelar.edu.uy/jspui/handle/20.500.12008/21751> [Data].

Naveda, L., Gouyon, F., Guedes, C., & Leman, M. (2011). Microtiming Patterns and Interactions with Musical Properties in Samba Music. *Journal of New Music Research*, 40(3), 225–238. <https://doi.org/10.1080/09298215.2011.603833>

Polak, R., Tarsitani, S., & Clayton, M. (2018). *IEMP Malian Jembe: A Collection of Audiovisual Recordings of Malian jembe Ensemble Performances, With Detailed Annotations*. Open Science Framework (OSF). <https://osf.io/m652x/> [Data].

- Poole, A., Tarsitani, S., & Clayton, M. (2018). *IEMP Cuban Son and Salsa: A Collection of Cuban Son and Salsa Recordings by the Group Asere, with Associated Annotations*. Open Science Framework (OSF). <https://osf.io/sfxa2/> [Data].
- Prögler, J. (1995). Searching for Swing: Participatory Discrepancies in the Jazz Rhythm Section. *Ethnomusicology*, 39(1), 21. <https://doi.org/10.2307/852199>
- Rasch, R. A. (1988). Timing and Synchronization in Ensemble Performance. In *Generative Processes in Music: The Psychology of Performance, Improvisation, and Composition* (pp. 70–90). Oxford: Clarendon Press. <https://doi.org/10.1093/acprof:oso/9780198508465.003.0004>
- Repp, B. H. (1998). Obligatory “Expectations” of Expressive Timing Induced by Perception of Musical Structure. *Psychological Research*, 61(1), 33–43. <https://doi.org/10.1007/s004260050011>
- Senn, O., Bullerjahn, C., Kilchenmann, L., & von Georgi, R. (2017). Rhythmic Density Affects Listeners’ Emotional Response to Microtiming. *Frontiers in Psychology*, 8. <https://doi.org/10.3389/fpsyg.2017.01709>
- Senn, O., Kilchenmann, L., Bechtold, T., & Hoesl, F. (2018). Groove in Drum Patterns as a Function of both Rhythmic Properties and Listeners’ Attitudes. *PLOS ONE*, 13(6), e0199604. <https://doi.org/10.1371/journal.pone.0199604>
- Senn, O., Kilchenmann, L., & Camp, M.-A. (2012). A Turbulent Acceleration Into the Stretto: Martha Argerich plays Chopin’s Prelude op. 28/4 in E minor. *Dissonance*, 120, 31–35.
- Tamlyn, G. N. (1998). *The Big Beat: Origins and Development of Snare Backbeat and Other Accompanimental Rhythms in Rock’n’Roll*. University of Liverpool. <http://tagg.org/xpdfs/TamlynPhD2.pdf>
- Wilkinson, M. D., Dumontier, M., Aalbersberg, Ij. J., Appleton, G., Axton, M., Baak, A., Blomberg, N., Boiten, J.-W., da Silva Santos, L. B., Bourne, P. E., Bouwman, J., Brookes, A. J., Clark, T., Crosas, M., Dillo, I., Dumon, O., Edmunds, S., Evelo, C. T., Finkers, R., ... Mons, B. (2016). The FAIR Guiding Principles for Scientific Data Management and Stewardship. *Scientific Data*, 3(1), 1–9. <https://doi.org/10.1038/sdata.2016.18>