

Annotation of soft onsets in string ensemble recordings

Tomczak, Maciej; Li, Min Susan; Bradbury, Adrian; Elliott, Mark; Stables, Ryan; Witek, Maria; Goodman, Tom; Abdikarim, Diar; Luca, Massimiliano Di; Wing, Alan; Hockman, Jason

DOI:

[10.48550/arXiv.2211.08848](https://doi.org/10.48550/arXiv.2211.08848)

License:

Creative Commons: Attribution-NonCommercial-ShareAlike (CC BY-NC-SA)

Document Version

Other version

Citation for published version (Harvard):

Tomczak, M, Li, MS, Bradbury, A, Elliott, M, Stables, R, Witek, M, Goodman, T, Abdikarim, D, Luca, MD, Wing, A & Hockman, J 2022 'Annotation of soft onsets in string ensemble recordings' arXiv.
<https://doi.org/10.48550/arXiv.2211.08848>

[Link to publication on Research at Birmingham portal](#)

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Annotation of Soft Onsets in String Ensemble Recordings

Maciej Tomczak^a Min Susan Li^b, Adrian Bradbury^b, Mark Elliott^c, Ryan Stables^a, Maria Witek^b, Tom Goodman^c, Diar Abdolkarim^b, Massimiliano Di Luca^b, Alan Wing^b, and Jason Hockman^a

^aBirmingham City University, Sound and Music Analysis (SoMA) Group in Digital Media Technology (DMT) Lab, Birmingham, UK;

^bUniversity of Birmingham, Sensory Motor Neuroscience (SyMon) Centre, Birmingham, UK;

^cWMG, University of Warwick, Coventry, UK

ARTICLE HISTORY

Compiled November 17, 2022

ABSTRACT

Onset detection is the process of identifying the start points of musical note events within an audio recording. While the detection of percussive onsets is often considered a solved problem, *soft* onsets—as found in string instrument recordings—still pose a significant challenge for state-of-the-art algorithms. The problem is further exacerbated by a paucity of data containing expert annotations and research related to best practices for curating soft onset annotations for string instruments. To this end, we investigate inter-annotator agreement between 24 participants, extend an algorithm for determining the most consistent annotator, and compare the performance of human annotators and state-of-the-art onset detection algorithms. Experimental results reveal a positive trend between musical experience and both inter-annotator agreement and performance in comparison with automated systems. Additionally, onsets produced by changes in fingering as well as those from the cello were found to be particularly challenging for both human annotators and automatic approaches. To promote research in best practices for annotation of soft onsets, we have made all experimental data associated with this study publicly available. In addition, we publish the ARME Virtuoso Strings dataset, consisting of over 144 recordings of professional performances of an excerpt from Haydn’s string quartet Op. 74 No. 1 Finale, each with corresponding individual instrumental onset annotations.

KEYWORDS

Music Information Retrieval; Onset Detection; Onset Annotation; String Ensembles

1. Introduction

Within the field of music information retrieval, note onset detection seeks to find the start points of musical events in an audio recording. Onset detection is used as a first step in a variety of high-level music processing applications in the creative, analytical and educational industries. The majority of onset detection approaches are based on subsampled versions of the audio signal (i.e., detection functions), which exhibit peaks at time points where some features of the signal change (e.g., energy, spectral content). Task performance is typically evaluated using standard metrics,

which assess onset presence within tolerance windows surrounding correct ground truth annotations. In comparison to symbolic representations of music (e.g., MIDI), audio waveforms present a challenge to onset annotation, obfuscating the precise location of true onset positions. Further, variability in strategies adopted by annotators can widen the discrepancy of ground truth labels associated with these events. Nonetheless onset detection for percussive sounds achieves near perfect performance in solo contexts; however, detection of soft onsets such as strings and woodwinds is a non-trivial task, due to a variety of playing styles and lack of phenomenological presence in the signal. This same dichotomy applies to the manual annotation process used to generate the ground truth used in evaluations—soft onset annotation is a more difficult task and generates lower levels of inter-annotator agreement.

While research has produced freely available datasets of soft onset annotations,¹ labels are typically generated from the responses of only a few participants. Further, relatively few studies have systematically investigated annotator strategies, quantified variability of annotations, or highlighted difficult contexts and instruments. To this end, we conduct three experiments to study annotation variability, propose a new dataset of professional string quartet recordings annotated by 24 participants, and extend a pre-existing strategy (Leveau & Daudet, 2004) for identifying the most reliable annotator among the set of participants towards generating reliable ground truth labels. The aim of the proposed experiments is to investigate best practices for consolidating multiple annotations created by musicians with wide-ranging annotation experience and state-of-the-art automatic onset detection approaches. Additionally, our experimental framework and the Virtuoso Strings dataset are well-suited for a variety of sound analysis tasks and music perception and cognition research.

1.1. Background

1.1.1. Onset Types

An important consideration when annotating note onsets of string instruments is distinguishing between the various methods of starting a note (e.g., *bowed*, *fingered*, *plucked*) and understanding the impact they have on how a listener might demarcate them as a beginning of an event. String instruments produce notes with relatively slow attack times, where the perceptual start of a sound differs from the moment our attention registers an onset. This phenomenon was investigated in the context of perceptual centres (P-centres), which refer to the moments at which humans perceive an event to occur (Morton et al., 1976, Danielsen et al. 2019). While sound onsets can be categorised based on their physical or perceptual characteristics (Collins, 2007), we are interested in associating ground truth annotations of physical onsets with musically motivated note types of string instruments. To that end, we first introduce a taxonomy based on physical characterisation of note onsets and then use it to define note types of string instruments investigated in this study.

One of the most fundamental methods for characterisation of music signals relies on envelopes. Coler and Lerch (2014) proposed the categorisation of string notes based on the standard attack, decay, sustain and release (ADSR) sections of temporal envelopes. These can be used to demarcate note, transition and rest regions of an audio signal. The bottom panel of Figure 1 shows idealised ADSR envelopes with labels for transition and rest regions during a cello performance. This model can be extended for

¹<https://www.ismir.net/resources/datasets/>

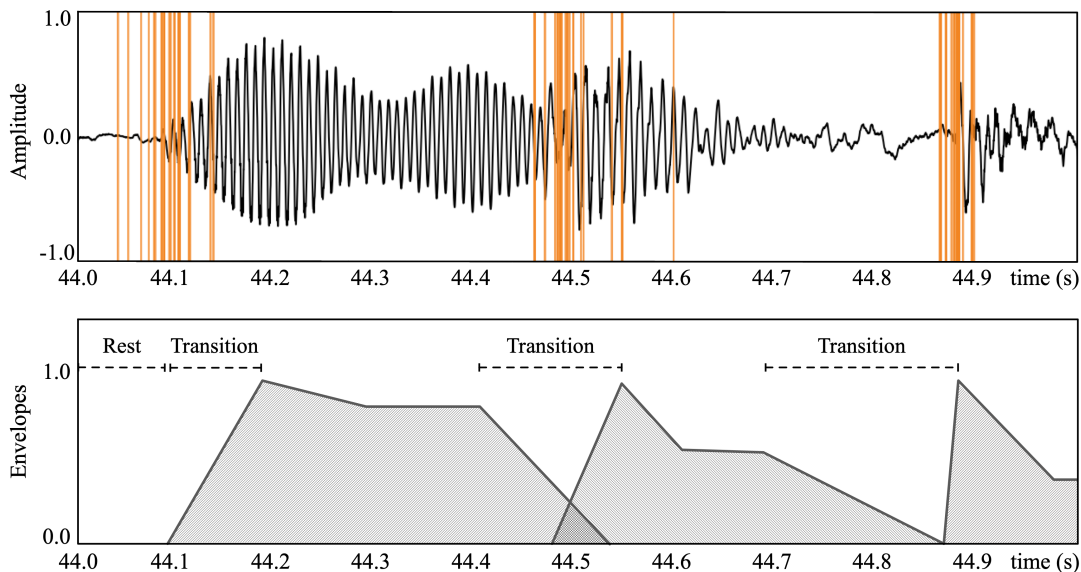


Figure 1. Audio waveform (top) with onset annotations from different annotators (orange) for a one second cello excerpt with three notes and envelopes (bottom) with labelled transition and rest regions. The idealised envelopes illustrate attack, decay, sustain and release regions containing annotated note onsets.

the characterisation of legato and detached transitions as well as rest-separated and decaying note types. Overlapping note envelopes are indicative of legato transitions between notes (e.g., first and second notes in Figure 1) and detached transitions (e.g., third note in Figure 1) are represented by pairs of notes, where the release section of the first envelope ends as the attack region of the following note begins. The rest-separated transitions can represent different types of notes which have musically-relevant silence between them, for example with decaying note types which fade gradually and do not exhibit a clear sustaining quality. In addition to descriptions of how sound changes it is important to understand the musical context of each change. To this end, different bowing techniques such as *détaché* (i.e., separate bow strokes for each note), *martelé* (i.e., hammered bow strokes) and *staccato* (i.e., note of shortened duration) have been studied in fields of acoustics and mechanics (Schoonderwaldt, 2009) as well as automatic bowing gesture detection (Dalmazzo & Ramírez, 2019).

We propose three main classes of note onsets for string instruments: *bowed onsets*, characterised by an attack of the bow on a string; *fingered onsets*, characterised by a change in pitch due to the engagement or release of a left-hand finger on a string; and *plucked onsets*, due to a hand digit plucking a string (e.g., pizzicato). Within each of these classes reside distinct types of onsets. For example, a *bowed onset* might be caused by the bow engaging with the string from above (e.g., spiccato, sautillé), a change of bowing direction (e.g., up to down bow), a change in bowing pressure or speed within a bow stroke (e.g., portato), or by a new string being engaged within a bow stroke (e.g., moving from the A to D string). The change in pitch resulting from a *fingered onset* might be due to a new finger engaging with the string (e.g., fourth finger following on from a first finger), or a new finger’s engagement being revealed by the release of another finger (e.g., a first finger following on from a third finger); either of these *fingered onset* types (i.e., engagement and release) can also be associated with changes to and from an *open* string on which no fingers are engaged. It should also be noted that these three classes of note onsets are not mutually exclusive; an

onset is often a combination of the two. *Fingered onsets*, especially those not involving changes to and from open strings, represent a more subtle class of onsets which are better described through pitch information changes rather than the discernible attack regions present in bowed onsets.

1.1.2. Onset Detection

The task of automatic onset detection typically involves three stages: *preprocessing*, *reduction*, and *peak-picking* (Bello et al., 2005). The *preprocessing* stage suppresses various aspects of the original audio signal (e.g., specific frequency bands) and enhances the target sounds of the studied musical domain (e.g., accentuating transients of *hard* onsets in drum recordings). The *reduction* stage generates an onset detection function referred to as an activation function in deep learning systems. The detection function outlines occurrences of transients in the form of a subsampled representation of the original signal. The final, *peak-picking* stage uses a detection or an activation function to select peaks which represent musical events in the original signal. The peak-picking process can be difficult as it relies heavily on clean detection or activation function representations, which can be easily affected when several instrument sounds overlap. A common problem linked to the preprocessing stage is the introduction of additional artifacts that can have a negative effect on the onset detection performance (Gillet & Richard, 2008). As such, current state-of-the-art deep learning onset detection approaches do not incorporate a preprocessing stage, instead reducing input representations such as logarithmic and Mel spectrograms to learned features during system optimisation, and perform peak picking on the output activation functions (Böck et al., 2012; Schlüter & Böck, 2014).

The MIR evaluation exchange (MIREX) *audio onset detection* evaluation challenge (MIREX, 2021) is used to test and benchmark new onset detection algorithms on a variety of musical audio signals. Ground truth information used within the challenge is achieved through cross-checking a single reference annotation by three experts. In cases where more than one reference annotation exists, the decision is not as clear. Leveau and Daudet (2004) have addressed the discrepancy for three annotation experiment participants using a small subset of the RWC music database (Goto et al., 2002),² including two solo bowed string instrument recordings. The authors compute the number of average consistent onsets (ACO) between annotators and report the mean timing differences between ACOs to measure the annotation difficulty for individual recordings.

1.2. Motivation

Onset annotations are typically generated by a single person or algorithm, and are cross-checked by one to three experts (MIREX, 2021; Fonseca et al., 2021). While economical, this is a relatively small set of views associated with the localisation of markers crucial for training and evaluation of automatic systems. The annotation of soft onsets is particularly challenging, as it is more difficult to localise start points within a time-domain signal without obvious and consistent signal characteristics for note events. It is therefore deemed important to consider multiple perspectives and strategies within the annotation process to identify problematic scenarios for both humans and algorithmic annotators, and to ensure the best annotations are chosen.

²<https://staff.aist.go.jp/m.goto/RWC-MDB/>

However, with the exception of Leveau and Daudet (2004), no studies have investigated the agreement between multiple onset annotations generated for the same musical material, and as such, there are no agreed-upon best practices for selecting the most appropriate markers among a set of annotators.

In this work, we conduct three experiments with 24 human annotators to investigate inter-annotator agreement, the effect of onset type (e.g., bow, finger), and the similarity between human and automated annotations. At present, onset detection datasets contain only a few string instrument examples; further, no datasets contain multiple recordings of the same piece with repeated musical instances and contexts. To mitigate this problem, we present the *Virtuoso Strings* dataset, consists of audio recordings of individual string instruments performed by a professional string quartet in a variety of musical and expressive conditions. The dataset is annotated by 24 human annotators with a range of musical training, and we extend the method by Leveau and Daudet (2004) for selecting the most consistent annotations towards the creation of ground truth data for training and evaluation of data-driven onset detection algorithms.

The remainder of this paper is structured as follows: Section 2 introduces the *Virtuoso Strings* dataset. Section 3 examines inter-annotator agreement and consistency, and Section 4 investigates annotation difficulty for four note onset types and presents our approach to identifying the ground truth annotations among many annotators. Section 5 compares the performance of five modern onset detection algorithms against that of the human annotators and the ground truth annotations. Finally, Section 6, provides concluding remarks and suggestions for future work.

2. Virtuoso Strings Dataset

In order to perform a detailed onset annotation analysis of string ensemble performance, we require a dataset of isolated recordings of individual instrument performers playing in multiple styles. Towards this goal, we present the *Virtuoso Strings* dataset, part of a larger collection of recordings and annotations captured through the Augmented Reality Music Ensemble (ARME) project for the study of musical performance synchronisation. This study presents the first freely available subset of Virtuoso Strings,³ which includes recordings of the first 49 bars from the Finale of Haydn’s string quartet Op. 74 No. 1 performed by the Coull Quartet. The scores are provided on the supporting website for this project.⁴

As seen from left to right in Figure 2, the first violin (VN1), second violin (VN2), viola (VA) and cello (VC) performers sat approximately one metre apart in a concert-style quartet arrangement. Recordings were realised through the use of instrument-mounted DPA 4099 series Lavalier microphones, which reduced the amount of bleed between instruments and ensured a clear quality recording of each separate sound source. All audio recordings were captured in a mono WAV format with a 44.1 kHz sample rate and 16-bit depth.

The selected Haydn excerpt was recorded 12 times (i.e., in 12 separate takes) each for 3 different musical conditions resulting in 48 recordings per condition, and 144 recordings in total. The conditions represent different interpretation instructions, and were chosen to span a wide range of performance types. It is thus possible to create more intricate contexts for the annotation of soft onset types which might have otherwise not been found if only standard concert performances were analysed. The normal con-

³<https://github.com/arme-project/virtuoso-strings>

⁴<https://github.com/arme-project/haydn-annotation-dataset>



Figure 2. The Coull Quartet during recording session.

dition (NR) represents a concert style performance, where musicians were asked to play as if on a concert stage. The speed condition (SP) includes spontaneous accelerando and decelerando initiated by a single musician (i.e., VN1, the leader) throughout each performance. The other musicians (i.e., VN2, VA, VC) were asked to follow the leader to the best of their ability. The leader was free to choose the magnitude of temporal variation of a single accelerando followed by a decelerando in each performance. In order to assess challenging contexts for onset annotation and automated detection, we include the deadpan condition (DP), for which musicians were asked to play with minimal expression in tempo and articulation. Both human and automated detection performance may then be compared against those of the NR and SP conditions towards understanding whether errors are inherent within the written score or manifest from musical expression adopted by performers.

3. Experiment 1: Inter-annotator Agreement

In this experiment, we analyse the agreement between different annotators, reporting on the consistency and asynchrony of annotations. We first introduce the annotation participants and experimental methodology, and subsequently outline two techniques utilised in the evaluation of human annotator performance: an existing visualisation method for identifying inter-annotator agreement (Ren et al., 2018), and a method for computation of average consistent onsets by Leveau and Daudet (2004), which we extend to be generalisable to many annotators with varied musical experience.

3.1. Participants

Twenty-four students and researchers from the University of Birmingham, Birmingham City University and University of Warwick took part in the experiment. Their musical experience, as defined by years of regular practice of a musical instrument and/or voice, ranged from no musical experience to professional musical career (minimum 0 years,

maximum 48 years, mean 28 years). All participants completed a questionnaire prior to undertaking the experiment, which requested information about their musical and annotation experience. Participants with musical experience then provided the number of years of experience, and listed their primary musical instrument. Participants were also asked if they had any known hearing impairments. Questionnaire queries and responses are available on the supporting website.⁴ Informed consent was obtained from all participants who contributed to this study.

3.2. Stimuli

The 24 participants annotated four recordings each (i.e., VA, VC, VN1, VN2) from the NR condition for take 12. The annotations from this experiment are denoted as NR12 and resulted in a total of 96 annotations, 24 overlapping per instrument. The final take (i.e., NR12) was chosen to ensure good quality performances, whereby each musician had had enough time to familiarise themselves with the piece and task instructions throughout the first 11 takes.

3.3. Onset Annotation

All participants were instructed to use the Sonic Visualiser (SV) software⁵ to annotate note onsets in each instrument’s recordings following a brief training session. The visual cues were presented to each participant as waveform representations (e.g., as in Figure 1), which provided visual orientation during event segmentation. Although annotation based only on visual cues can allow for semi-precise segmentation of some notes, every participant was instructed to confirm each annotation by listening to the recordings. Annotators were free to listen to each recording as many times as needed and to use the time-stretching functionality of SV for the more challenging sections of the performances; however, they were instructed not to use additional visual cues provided by SV (e.g., pitch tracks, spectral representations) to prevent overpowering auditory stimuli with too many visual cues.

The resulting annotations are represented by a sequence of time points demarcating note onsets, as annotated by participant *a*. The number of onsets per annotator may vary due to missed or mistakenly placed onsets. Short inter-onset-intervals (IOI) which could not reflect musical time intervals were determined to be mostly caused by participants mistakenly providing more than one keystroke for a single note event (i.e., double-clicking); a pre-processing stage was thus used to preserve the first onsets while removing subsequent erroneous onsets ($IOI < 30ms$). The same pre-processing was implemented in all experiments.

3.4. Method

3.4.1. Inter-annotator Agreement

Onset locations are typically evaluated with a single set of ground truth annotations created by expert annotators; however, as this ground truth does not exist, in this experiment we aim to identify annotator performance through the assessment of the inter-annotator agreement. We follow a similar approach to that of Ren et al. (2018) and Tomašević et al. (2021), who represented pairwise agreement between annotators

⁵<https://www.sonicvisualiser.org/>

using the standard F-measure metric. In an automatic detection scenario, if an event is detected in an acceptable range it is labelled as a true positive (TP). If a detected onset is not within this range it is labelled as a false positive (FP). Alternatively, if an annotated ground truth event does not coincide with any detected onset, it is labelled as a false negative (FN). Using these three event counts, precision (P) and recall (R) metrics can be calculated. Precision is the fraction of true positives among all detected events and recall is the fraction of true positives among all missed ground truth events. The two metrics are calculated as follows:

$$P = \frac{TP}{TP + FP}, \quad (1)$$

$$R = \frac{TP}{TP + FN}. \quad (2)$$

The precision and recall can be expressed as a single metric using the F-measure (F) calculation:

$$F = 2 \frac{PR}{P + R}. \quad (3)$$

A tolerance window is used to determine whether a detected onset is within an acceptable range either side of the onset in the ground truth. The window is represented as a duration either side of the detected onset (e.g., ± 25 ms). Tolerance windows are commonly set to the shortest played note within a metrical grid (e.g., 16th note, 32nd note). While automatic onset detection evaluations and studies of microtiming (i.e., intentional deviations from strict metronomic timing) for musicological purposes use shorter tolerance windows as greater precision is required, the typical tolerances in the literature range from 20 ms to 50 ms (Böck et al., 2012; Wu et al., 2018; Fonseca et al., 2021).

Once a detection or an activation function θ has been generated for each instrument recording under observation, the onsets must be located in the temporal domain from θ . Following Eyben et al. (2010), a threshold T can be determined using the mean of all frames in θ and a constant λ :

$$T = \text{mean}(\theta) * \lambda. \quad (4)$$

If the current frame n has its peak above T then the onset function Γ is given by:

$$\Gamma(n) = \begin{cases} 1, & \theta(n-1) < \theta(n) \geq \theta(n+1) \ \& \ T < \theta(n) \\ 0, & \text{otherwise.} \end{cases} \quad (5)$$

Inter-annotator agreement is evaluated using a similarity matrix of F-measure scores which represent higher agreements for comparisons where a larger proportion of onsets were within the set tolerance window ω ($\omega=25$ ms) which was set to a size commonly used in other work (Eyben et al., 2010; Böck et al., 2012; Ali-MacLachlan et al., 2016).

3.4.2. Consistent Onsets

To produce a generalisable and realistic measure of annotation variance and difficulty in annotation of each instrument, we extend the approach from Leveau and Daudet (2004) for computation of average consistent onsets (ACO). ACOs represent the means of TP onsets found from comparisons between pairs of annotations. We can also calculate timing differences between the found consistent onsets and take their average. Hence, higher mean timing differences will represent larger disagreement in how participants annotated a particular instrument. In Leveau and Daudet (2004), three annotators were studied, resulting in the following comparisons: first, TPs between participants 1 and 2 were found; then they were used to check that they were consistent in comparisons between subjects 2 and 3; and finally between subjects 3 and 1. Unfortunately, these comparisons will produce different results depending on the starting order of the annotators. To mitigate this, we take averages of the ACOs calculated from k repetitions. In each k , consistent onsets are calculated from a randomly generated sequence of annotators. We stop the computation after mean ACOs exhibit changes <1 ms.

To investigate the inter-annotator agreement at different analysis scales we calculate ACOs using tolerance windows between 25–100 ms. Additionally, in order to prevent the method from finding no consistent onsets for the lower tolerance window sizes we use a subset of 16 annotators with five or more years of musical experience.

3.5. Results and Discussion

Figure 3 shows matrices with pairwise F-measure scores for different instruments in NR12 performances. The labels on the x and y axes show years of musical experience of each of the 24 annotators. The comparisons between the annotators are sorted according to the number of years of musical experience, from the least experienced to most experienced (i.e., top to bottom and left to right). The results highlight higher F-measures in comparisons involving annotators with longer musical experience (e.g., lower right corner of the matrices). Lower agreement can be particularly seen between the cello annotations, which indicates higher difficulty in annotation of cello recordings as compared with the viola and the violins. A similar relationship can be observed in results for tolerance windows larger than 25 ms which are presented on the supporting website.⁶

The top panel of Figure 4 shows ACOs from NR12 performances plotted for tolerance windows between 25 ms and 100 ms. The total number of note onsets (i.e., as defined in the score) in VA, VC, VN1 and VN2 recordings are 116, 100, 167 and 150, respectively. While most ACO results approach the total note counts at higher tolerance windows, relatively few are found at tolerances <50 ms (e.g., <20 ACOs for each instrument at 25 ms). The bottom panel of Figure 4 shows mean timing differences calculated from ACOs for different tolerance windows. Timing differences for cello annotations are larger than for viola and violins. While in the lower tolerance window sizes the differences remain relatively small (i.e., ≈ 1.5 ms) they increase to approximately 5 ms difference at 100 ms tolerance. Additionally, a slight decrease in ACOs can be observed at 100 ms tolerance window for VN1, which is due to multiple onsets overlapping as the minimum average IOI for the NR12 performance is approximately 70 ms.

These results, as well as the pairwise F-measure results, indicate higher difficulty in annotation of cello recordings by the experiment participants with 5–15 years of

⁶<https://arme-project.co.uk/demos/onset-annotation>

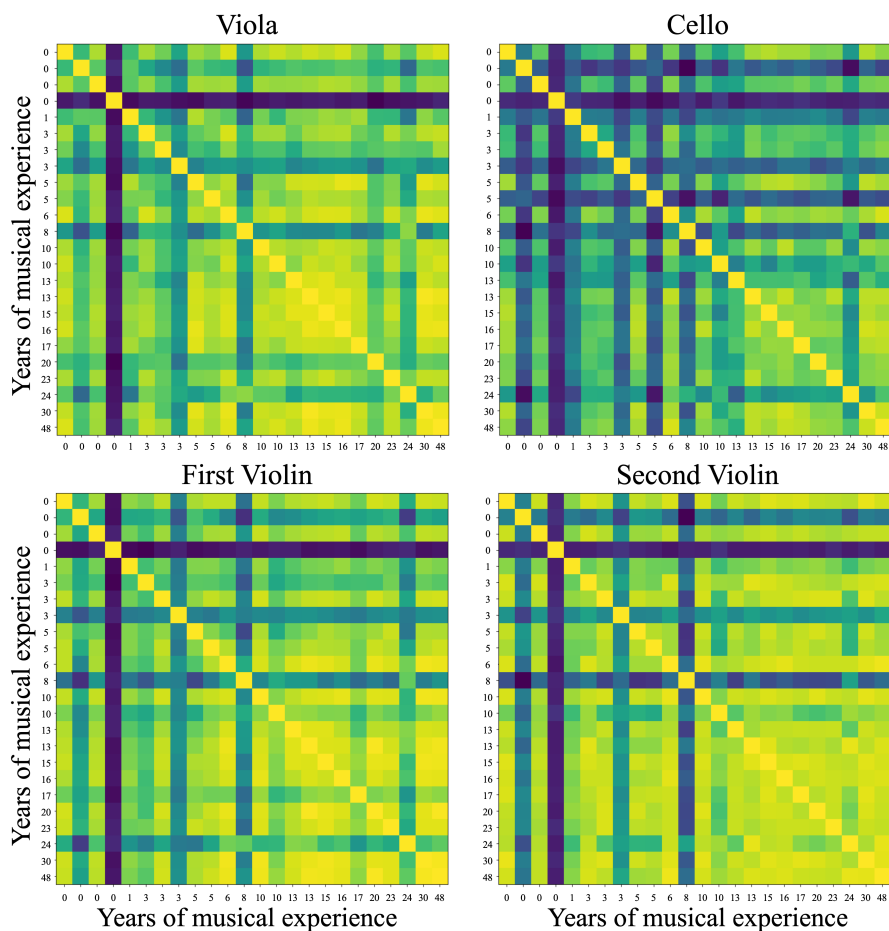


Figure 3. Pairwise F-measure scores calculated between 24 participants and sorted from least to most years of musical experience (i.e., top to bottom and left to right). High and low scores are indicated by bright and dark colours, respectively.

musical experience. While the timing differences between cello and other instruments are small (i.e., <2 ms at 25 ms tolerance), it is important to note that the number of ACOs found for the lower tolerance window sizes are very few when compared with the total numbers of notes in the performance (see Table 1). For example, at the 25 ms tolerance for VN1 there are approximately 20 consistently annotated onsets with an approximate 8 ms average timing difference between them out of the total of 167 note onsets in the performance. This result indicates that, on average, the experiment participants did not agree with the annotation of approximately 147 note onsets within the first violin recording in NR12. Overall, small numbers of ACOs for the smaller tolerance windows indicate there exist large disagreements between the annotators. Unfortunately, a general analysis of ACOs does not highlight which onsets manifest the highest annotation disagreements. Hence in the following experiment we conduct a more granular investigation which considers different types of onsets (see Section 1.1.1).




Figure 4. Average consistent onset counts (top) and average timing differences between consistent onsets (bottom) plotted for tolerance windows from 25 ms to 100 ms.

4. Experiment 2: Onset Type Analysis

This experiment investigates performances of all annotators on different note onset types in NR12 performances. Additionally, a new set of expert annotations is created in order to help evaluate automatic approaches to the detection of note onsets. These annotations include a high level of granularity with labels for each individual note played by each musician. In addition to note labels, every note onset time also includes a separate annotation denoting onset type. Two pairs of complementary types of onsets are presented: change of bow or finger within a bow, and fingered (stopped) or open string (unstopped).

4.1. Participants and Expert Annotator

The same twenty-four sets of annotations used within Experiment 1 are examined with respect to onset type. Following the method outlined in Section 3.4.2 we select the annotator whose onsets were closest to the average consistent onset (ACO) timings. Using a tolerance window of 25 ms, the annotations from participant 2 (a_2) are found to be the most consistent across the ACO calculations. Next, annotations made by a_2 for all NR12 instrument recordings are manually cross-checked to ensure no notes were missed when compared with the music score. Then, the checked instrument files are assigned to a new annotator a_0 , termed the expert annotator. The annotations of a_0 are available as part of the open dataset.⁴ The authors who modified and verified the expert annotations are all musically trained and have experience with string quartet

Inst.	All notes	Open string	Stopped note	Bow start	Finger change
VA	116	20	96	99	17
VC	100	6	94	89	11
VN1	167	5	162	112	55
VN2	150	18	132	115	35

Table 1. Number of notes per onset category in viola (VA), cello (VC), and first and second violins (VN1 and VN2) in NR12.

instrumentation (minimum 10 years).

4.2. Stimuli and Onset Type Annotation

A new set of annotations is added to those from expert a_0 describing four onset types. The four onset types are categorised into two pairs; each pair comprises complementary types of onset. The first pair defines them as being either an *open string* or *stopped note* (i.e., any fingered note). The second pair describes onsets as a *bow start* (i.e., first note within a bow direction) or *finger change* (i.e., any subsequent fingered notes played within a shared bow direction). Table 1 presents the numbers of onsets in each category.

4.3. Results and Discussion

Figure 5 shows the mean performances from all annotators in relation to the ground truth expert annotator a_0 (here considered ground truth) for each instrument and onset category calculated as a percentage of TP onsets. The reported results use tolerance window size of 25 ms and are presented as means from all annotation participants, while the per-annotator results are available on the supporting website.⁶

The highest mean performance from the annotators is observed in viola annotations for the open string onset category, indicating the highest agreement between all annotators for this onset type. The lowest level of agreement between all annotators is present in the cello annotations of the finger change onset type which can be seen to also be low for the other instruments. Mean accuracies across all participants for the open string and stopped notes as well as bow start and finger change onsets are 83.0%, 81.3%, 83.9% and 70.8%, respectively. Additionally, the mean accuracies across all participants for VA, VC, VN1 and VN2 are 82.0%, 72.4%, 82.8%, 81.7%. Overall, given the relatively small tolerance window there seems to be a stronger annotation agreement for all onset types except those of finger change onsets, which can be predominantly characterised through pitch change rather than modification of transient information.

5. Experiment 3: Onset Detection

In this experiment we evaluate annotations for NR, SP and DP conditions using several standard and state-of-the-art onset detection algorithms. The goal of this evaluation is to compare the onset detection performance of both signal processing and deep learning systems within a common evaluation framework. To highlight the performance of automatic approaches in relation to the proposed taxonomy of onset types, different systems are utilised in an in-depth discussion and evaluation of each onset type. Additionally, the highest-performing system is selected and used in comparative analysis

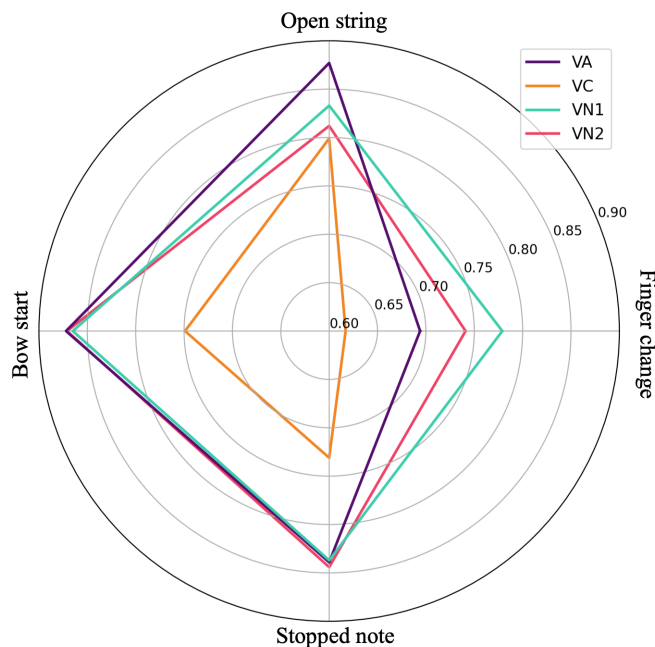


Figure 5. Mean true positive rates from all annotators calculated from comparisons against the ground truth expert annotations in NR12.

of onset annotation by individual annotators. For the sake of completeness and reproducibility, the extended evaluation results with all experimental data can be found on our complementary website.⁶

5.1. Participants

Four participants were selected from the previous experiments to take part in a third study due to their musical experience (minimum 5 years). The participants include a_{12} , a_{13} , a_{16} and a_{18} with 13, 5, 10 and 16 years of musical experience, respectively. In addition to the annotations from the four participants, we also use expert annotations a_0 for NR12.

5.2. Stimuli

Participants annotated 140 recordings, totalling the remaining repetitions 1–11 for NR and repetitions 1–12 for SP and DP conditions. A different number of recordings was assigned to each participant with the mean of 35 recordings per participant. As in the first experiment, the annotation experiment participants were not constrained by time and could use the time-stretching functionality of SV during the more challenging sections of each recording. All participants performed their annotations for this experiment after annotating the recordings used in Experiment 1.

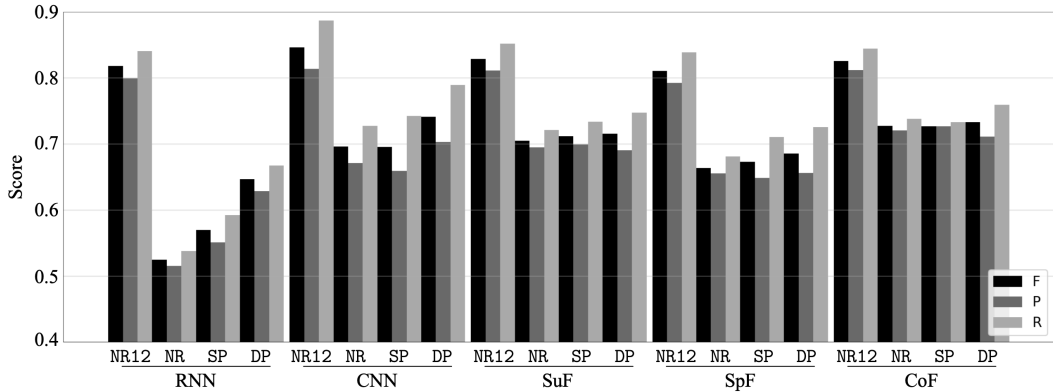


Figure 6. Mean F-measure (F), precision (P) and recall (R) calculated across VA, VC, VN1 and VN2. Different onset detection methods are compared against expert a_0 annotations in NR12 as well as annotations from a_{12} , a_{13} , a_{16} and a_{18} in NR, SP and DP conditions.

5.3. Method

5.3.1. Onset Detection Algorithms

To further compare and assess the performance of the compiled annotations we incorporate specialised methods for soft onset detection found in the literature. SuperFlux (SuF) by Böck and Widmer (2013b) computes the difference between the short-time magnitudes of nearby spectral frames and is optimised for music signals with soft onsets and vibrato effect in string instruments. ComplexFlux (CoF) by Böck and Widmer (2013a) is based on the SuF algorithm with the addition of local group delay measure which makes this method more robust against loudness variations of steady tones. We also use the standard spectral flux (SpF) method by Masri (1996). To compute detection functions from the three methods, we use log-filtered spectrograms as inputs with default parameters provided by madmom Python library (Böck et al., 2016).⁷ Additionally, we use two methods based on deep neural networks which achieve state-of-the-art onset detection results in various music contexts. To compute activation functions we use methods based on a recurrent neural network (RNN) by Böck et al. (2012) and a convolutional neural network (CNN) architecture by Schlüter and Böck (2014) both these methods use default parameters from the madmom library.

5.3.2. Evaluation Metrics

Standard precision (P), recall (R), and F-measure (F) scores (see Section 3.4.1) are used for evaluation of the onset detection performance. The mean F, P and R results are calculated as averages across annotations from all participants per instrument and onset detection method. The λ parameter is chosen using grid search across the NR12 dataset and is optimised for highest F-measure.

5.4. Results and Discussion

Figure 6 presents onset detection results as a correspondence between the onsets detected by the algorithms and compared against annotations from participants $a_{0,12,13,16,18}$. Results are reported for five algorithms using annotations from NR, SP and

⁷<https://github.com/CPJKU/madmom>

Inst.	Method	Open string	Stopped note	Bow start	Finger change
VA		20 onsets	96 onsets	99 onsets	17 onsets
VA	RNN	0.85	0.83	0.87	0.65
	CNN	0.95	0.90	0.94	0.71
	SuF	1.00	0.86	0.97	0.41
	SpF	0.95	0.85	0.89	0.76
	CoF	1.00	0.85	0.97	0.35
VC		6 onsets	94 onsets	89 onsets	11 onsets
VC	RNN	1.00	0.86	0.89	0.73
	CNN	1.00	0.81	0.83	0.73
	SuF	1.00	0.87	0.94	0.36
	SpF	1.00	0.85	0.90	0.55
	CoF	1.00	0.86	0.93	0.36
VN1		5 onsets	162 onsets	112 onsets	55 onsets
VN1	RNN	0.80	0.79	0.74	0.89
	CNN	1.00	0.91	0.94	0.85
	SuF	1.00	0.81	0.91	0.64
	SpF	1.00	0.79	0.78	0.84
	CoF	1.00	0.80	0.88	0.65
VN2		18 onsets	132 onsets	115 onsets	35 onsets
VN2	RNN	0.89	0.86	0.89	0.80
	CNN	0.89	0.92	0.97	0.74
	SuF	0.89	0.81	0.93	0.46
	SpF	0.94	0.81	0.90	0.57
	CoF	0.89	0.82	0.94	0.46
Means	RNN	0.88	0.84	0.85	0.77
	CNN	0.96	0.88	0.92	0.76
	SuF	0.97	0.84	0.94	0.47
	SpF	0.97	0.83	0.87	0.68
	CoF	0.97	0.83	0.93	0.46

Table 2. True positive rates for different onset types using expert annotations for NR12.

DP conditions as well as expert annotations a_0 from NR12. The overall highest performing algorithms for all instruments are CNN and CoF. While all methods perform similarly on the NR12 recordings, a noticeable drop in performance can be observed for the NR, SP and DP. This may be due to the difference in annotation style between a_0 (see Section 4.1) and $a_{12,13,16,18}$ who were responsible for the remaining NR, SP and DP annotations. The largest performance drop can be seen by the RNN method for these three conditions. This may indicate that the RNN-based system is most sensitive to the variability in annotation strategies present in a_0 as its performance for NR12 is comparable with the other methods. Another reason for the lower onset detection performance in these conditions may be at least in part because they comprise a higher number of recordings as compared with the NR12. The mean results for NR12 include only four recordings but NR, SP and DP include 44, 48, and 48 recordings, respectively. While the mean precision scores are similar across all algorithms on recordings with NR12 expert annotations, the highest mean precision (0.814) and recall (0.887) are achieved by the CNN system. The mean recall results are higher than precision for the DP condition in all algorithms which is expected as these performances exhibit minimal expression in tempo and articulation making note onsets more regular than in the NR and SP conditions.

Onset detection results for different note onset types in the expert annotations NR12 are presented in Table 2. The accuracies are calculated as percentages of TP onsets per onset type, instrument and algorithm. Mean results for individual algorithms together

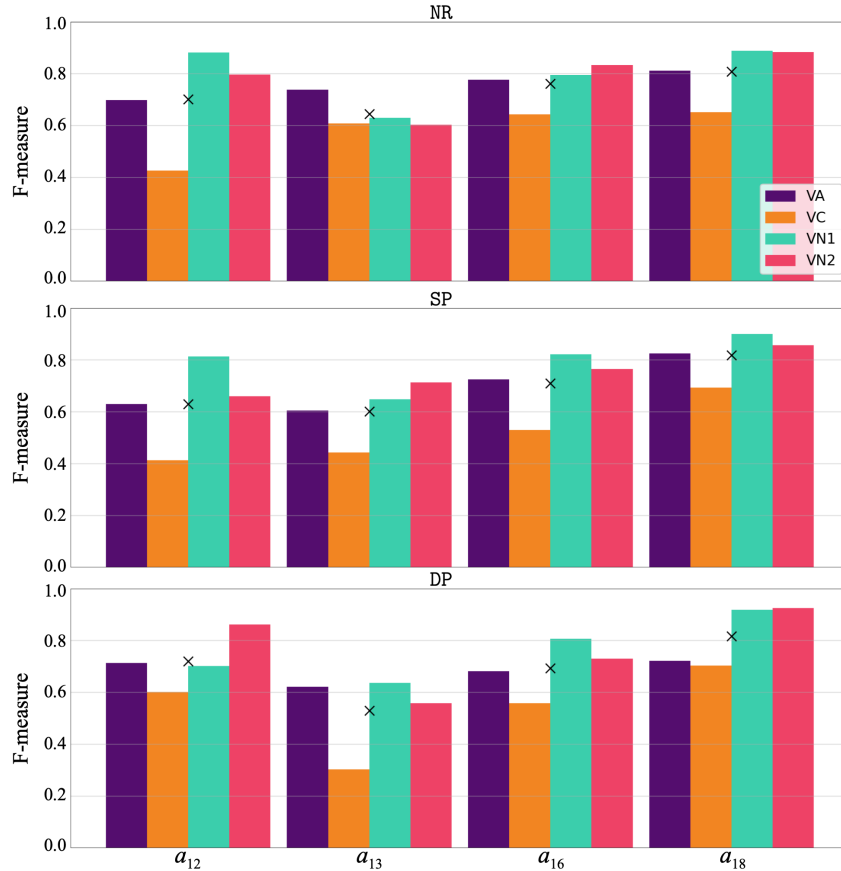


Figure 7. Mean F-measures acquired using the highest-performing CNN onset detection system for NR (top), SP (centre) and DP (bottom) conditions. The results are reported for annotators a_{12} , a_{13} , a_{16} and a_{18} with per annotator means (\times).

with a total mean accuracy calculated for each onset type are shown at the bottom of the table. The highest performing method is the CNN system with a mean 87.9% accuracy across all onset types. While the results indicate overall high detection rates for open string, stopped note and bow start onsets, the finger change onsets display much lower accuracies. This is a similar finding to the results presented in Experiment 2 (see Section 4.3), in which the majority of participants had the lowest average annotation accuracy for the finger change annotations. Interestingly, the overall high performing CNN and RNN systems show lower results for the viola and cello bow start onsets when compared to the SuF and CoF methods but outperform these methods in the detection of finger change onsets. The lower performance of SuF and CoF for the finger change onsets can also be seen in other instruments. This indicates that the deep learning systems demonstrate a more uniform high performance over all note onset types while the signal processing methods struggle with the more challenging finger change onsets. Additionally, the more easily detected open string onsets (compared with stopped notes) might be partially explained by the fact that the engagement of an open string often marks a change of string (e.g., from D to A string) and is thus associated with the fresh bow attack, even within a bow stroke, involved in such a move. A similar ease of identification of the open string note onsets can be observed in Experiment 2, where the annotators show high agreement (see Figure 5) in the

annotation of this type of onset across all instruments.

We also report on mean F-measures for each annotator who participated in this experiment. Figure 7 presents results obtained using the highest-performing CNN system, in which the detections are compared against onsets (i.e., ground truth) from each of the four participants. The CNN system achieved highest means (shown as \times) across instruments for annotator a_{18} in all conditions; 0.81 in NR, 0.82 in SP and 0.82 in DP. Additionally, the lowest mean detection performance per instrument can be observed in cello (0.55), and the highest in first violin (0.79). The lower F-measures for cello can be also observed in a_{18} who annotated the recordings with the highest similarity to the CNN onset detection system. The annotations from a_{18} for the cello in the DP condition are similar to that of the viola, which has a lower result than in the NR and SP conditions. The lower results for the viola in DP are also present in the a_{13} and a_{16} . While the cello and viola use the same tuning configuration (i.e., tuned to C-G-D-A), the cello is tuned an octave lower. The higher difficulty in exactly localising the cello note onsets could potentially be due to relatively thicker, low tension strings, which results in less discrete note attack characteristics when fingers or the bow engage and disengage with a string to effect a new note.

6. Conclusions

In this work we presented a new dataset of 24 human note onset annotations and a set of expert annotations for an excerpt from the Finale of Haydn’s string quartet Op. 74 No. 1 performed by a professional quartet in three distinct performance styles. Investigating the human annotations in the context of inter-annotator performance and automatic onset detection methods highlighted the challenging aspects of the annotation process and informed best annotation practices. We proposed a categorisation of note onset types of string instruments, which enabled us to evaluate which types caused the highest disagreement between human annotators and proved the most challenging to the state-of-the-art automatic methods. Within our evaluation we found cello to be the most difficult instrument to annotate by the majority of experiment participants, a finding to be taken into consideration during the note onset dataset curation process. Additionally, we show that finger change onsets are the most challenging type to identify by both human annotators and algorithms. Both the high-quality audio recordings and the experiment annotations, as well as the expert annotations, are published as open datasets. In the future, we plan to continue the study of stylistic differences across players and across different recording repetitions. Lastly, we plan to explore semi-automatic and instrument-specific annotation approaches as well as to continue investigation into the variability between automatic onset detection systems and human annotators, towards improving synchronisation in interactive music systems.

Acknowledgements

This project is kindly funded by Engineering and Physical Sciences Research Council (EPSRC) grant with reference EP/V034987/1. We would like to thank the musicians of the Coull Quartet for their dedication and professionalism as well as all experiment participants who helped with this study.

References

- Ali-MacLachlan, I., Tomczak, M., Southall, C., & Hockman, J. (2016). Note, cut and strike detection for traditional irish flute recordings. In *Proceedings of the International Workshop on Folk Music Analysis (FMA)*.
- Bello, J. P., Daudet, L., Abdallah, S., Duxbury, C., Davies, M. E., & Sandler, M. B. (2005). A tutorial on onset detection in music signals. In *IEEE Transactions on Speech and Audio Processing*, 13(5), 1035–1047.
- Böck, S., Arzt, A., Krebs, F., & Schedl, M. (2012). Online real-time onset detection with recurrent neural networks. In *Proceedings of the International Conference on Digital Audio Effects (DAFx), York, UK*.
- Böck, S., Korzeniowski, F., Schlüter, J., Krebs, F., & Widmer, G. (2016). madmom: A new python audio and music signal processing library. In *Proceedings of the ACM International Conference on Multimedia (ACM-MM), Amsterdam, Netherlands*, 1174–1178.
- Böck, S., Krebs, F., & Schedl, M. (2012). Evaluating the online capabilities of onset detection methods. In *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR), Porto, Portugal*, 49–54.
- Böck, S., & Widmer, G. (2013a). Local group delay based vibrato and tremolo suppression for onset detection. In *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR), Curitiba, Brazil*, 361–366.
- Böck, S., & Widmer, G. (2013b). Maximum filter vibrato suppression for onset detection. In *Proceedings of the International Conference on Digital Audio Effects (DAFx), Maynooth, Ireland*.
- Coler, H. v., & Lerch, A. (2014). CMMSD: A data set for note-level segmentation of monophonic music. In *Proceedings of the Audio Engineering Society (AES) Conference: Semantic Audio*, 1–10.
- Collins, N. M. (2007). Towards autonomous agents for live computer music: Realtime machine listening and interactive music systems. *PhD Thesis. University of Cambridge*.
- Dalmazzo, D., & Ramírez, R. (2019). Bowing gestures classification in violin performance: a machine learning approach. *Frontiers in Psychology*, 10(1), 1–14.
- Danielsen, A., Nymo, K., Anderson, E., Câmara, G. S., Langerød, M. T., Thompson, M. R., & London, J. (2019). Where is the beat in that note? Effects of attack, duration, and frequency on the perceived timing of musical and quasi-musical sounds. In *the Journal of Experimental Psychology: Human Perception and Performance*, 45(3), 402–418.
- Eyben, F., Böck, S., Schuller, B., & Graves, A. (2010). Universal onset detection with bidirectional long short-term memory neural networks. In *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR), Utrecht, Netherlands*, 589–594.
- Fonseca, J., Fuentes, M., Bonini, F. B., & Davies, M. E. P. (2021). On the use of automatic onset detection for the analysis of maracatu de baque solto. In *Perspectives on Music, Sound and Musicology*, 209–225.
- Gillet, O., & Richard, G. (2008). Transcription and separation of drum signals from polyphonic music. In *IEEE Transactions on Audio, Speech, and Language Processing (TASLP)*, 16(3), 529–540.
- Goto, M., Hashiguchi, H., Nishimura, T., & Oka, R. (2002). RWC music database: Popular, classical and jazz music databases. In *Proceedings of the International Conference on Music Information Retrieval (ISMIR), Baltimore, USA*, 287–288.
- Leveau, P., & Daudet, L. (2004). Methodology and tools for the evaluation of automatic onset detection algorithms in music. In *Proceedings of the International Conference on Music Information Retrieval (ISMIR), Barcelona, Spain*.
- Masri, P. (1996). Computer modelling of sound for transformation and synthesis of musical signals. *PhD Thesis. Electrical & Electronic Engineering, University of Bristol*.
- MIREX. (2021). Audio onset detection evaluation. Accessed on the 8th of November 2022: https://www.music-ir.org/mirex/wiki/2021:Audio_Onset_Detection.
- Morton, J., Marcus, S., & Frankish, C. (1976). Perceptual centers (P-centers). *Psychological*

- Review*, 83(5), 405–408.
- Ren, Y., Nieto, O., Koops, H. V., Volk, A., & Swierstra, W. S. (2018). Investigating musical pattern ambiguity in a human annotated dataset. *In Proceedings of the International Conference on Music Perception and Cognition (ICMPC)*.
- Schlüter, J., & Böck, S. (2014). Improved musical onset detection with convolutional neural networks. *In Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Florence, Italy*, 6979–6983.
- Schoonderwaldt, E. (2009). Mechanics and acoustics of violin bowing. (*Unpublished doctoral dissertation*). Stockholm: KTH Royal Institute of Technology.
- Tomašević, D., Wells, S., Ren, I., Volk, A., & Pesek, M. (2021). Exploring annotations for musical pattern discovery gathered with digital annotation tools. *In the Journal of Mathematics and Music*, 15(2), 194–207.
- Wu, C.-W., Dittmar, C., Southall, C., Vogl, R., Widmer, G., Hockman, J., ... Lerch, A. (2018). A review of automatic drum transcription. *In IEEE/ACM Transactions on Audio, Speech, and Language Processing (TASLP)*, 26(9), 1457–1483.