PIANO TOUCH ANALYSIS: A MATLAB TOOLBOX FOR EXTRACTING PERFORMANCE DESCRIPTORS FROM HIGH-RESOLUTION KEYBOARD AND PEDALLING DATA

Michel Bernays
Université de Montréal
michel.bernays@umontreal.ca

Caroline Traube
Université de Montréal
caroline.traube@umontreal.ca

ABSTRACT

We hereby present the analytic tools developed as a MATLAB toolbox for exploring the most subtle features of piano touch and gesture. From high-sample-rate, high-precision precision key, hammer and pedal tracking data about a performance gathered thanks to the Bösendorfer grand piano-embedded CEUS digital recording system, the toolbox main functions can extract exhaustive features detailing the pianist's touch as a thorough account of nuances in articulation, timing, dynamics, attack and pedalling. Each performance and gestural control thereof is thus described over each of its notes and chords. By comparing several performances, it is possible to characterize the gestural control of expression in piano performance, as the correlations among piano touch features towards one examined expressive parameter. The analytic functions in the toolbox — with piano touch feature visualization and comparison, chords and notes selection tools, score-performance matching and advanced, automated statistical analyses and visualization thereof — allow for rigorous, quantitative exploration of expressive performance and its gestural control, which here is especially applied towards investigating the use of timbre as expressive device in piano performance. With these tools, we thus intend to build a gestural mapping of piano timbre.

1. INTRODUCTION

Musical performance is essential to the art and experience of music. Classical performers, in particular, aim to enlighten the composers' works, as reduced on scores, by their unique interpretation, thus expressing their creativity throughout this complex task. This holds undeniably true as regards piano performance, for which a vast repertoire has been composed in the few centuries since the inception of the instrument. An extensive, empiric knowledge of piano gesture and touch has thus developed within the pianistic world, to allow for the most vivid expressivity in piano performance. Many piano treatises, especially in the twentieth century [15] [18] [12], have provided guidelines for the most efficient and expressive gesture to improve the pianist sound. Through systematic studies and interview collections, researchers have highlighted the specificities of gesture, touch or fingering [3] strategies.

Moreover, along technological progress since the early twentieth century, new methods and tools have appeared that allow for a quantitative study of performance gesture.

The first datasets available on piano gesture in performances came as rolls recording key movements through mechanical design, such as Duo-art rolls which were used notably to study chord synchronization [27] and pedalling [11]. The Iowa Piano Camera, developed by Seashore et al. in 1936 for the "Objective Recording and Analysis of Musical Performance", could measure hammer motion and velocity with an elaborate, embedded slit-and-film system, and was used to analyze several dynamic, rhythmic/temporal and structural features [24]. Both systems could thus give out piano rolls indicating for each note the key depressed, its onset and offset (and thus duration), and in the last case, hammer velocity. In fact, this constituted a primitive version of the 1981 MIDI standard data. The subsequent MIDI recording pianos such as Yamaha Disklaviers, with the same but way improved pianoroll-like data acquisition abilities, vastly eased the quantitative recordings of some piano performance gestures. This way, systematic studies by Repp [22] [23], Parncutt [20] [21] and Goebl [5] [6] [7] among others, of such expressive features as timing, dynamics and articulation, have given us a better grasp the role of gesture in expressive piano performance.

Yet due to the limits inherent to the MIDI standard [17] with regard to the amount of information available on piano touch control, other technical methods of measurement have been called upon. In particular, and as early as 1929, Ortmann devised a mechanical recording system, involving springs, levers, a dynamograph and revolving drum, that could acquire continuous measurements of key motion and represent them as curves detailing the slightest fluctuations. It was thus used to study piano touch and tone [19] and determine the variations in touch, as key motion, for a single note with different tones intended. A few years later (1934), Hart, Fuller and Lusby [10] used an optical system to track hammer motion, and a mechanical key striker to identify the effect of various controlled key strokes on hammer motion and the acoustic signal.

Nowadays, equipment used to record expressive gesture in piano performance may include motion capture [26], sensors [9], or UV hand paint and lighting [14]. Such designs enable to acquire pianists' gesture as a whole, in order to characterize musical embodiment or physiological features. As regards the instrumental gesture itself, that is the efficient part of performance gesture that conveys energy directly to the keys, recent recording pianos have expanded over MIDI limitations and propose highly accurate tracking of key (and pedal) depression and hammer movement. Such detailed information is actually required for analyzing the subtle nuances involved in some aspects of piano performance, among which timbre and tone control.

2. AIMS

As our main research project indeed aims at understanding high-level pianists' ability to control timbre and tone in their performances according to their musical and emotional intentions, we investigate how the most fine-grained properties of pianists' touch are nuanced within musical performances, in order to obtain different timbres. Yet such nuances necessitate highly precise data from which the intricacies of key strokes can be thoroughly assessed. Such is the data that can be acquired by the Bösendorfer CEUS piano digital recording system. In order to analyze this fine-grained data, we have aimed to develop a dedicated MATLAB toolbox including a set of functions that we are hereby describing in this paper.

3. HIGH-RESOLUTION DATA ACQUISITION

The CEUS system we used is embedded in an Imperial Bösendorfer grand piano (Figure 1). [1] The system includes optical sensors behind the keys, hammers and pedals, microprocessors and electronic boards (Figure 2) that process sensor data and send it to an embedded computer into whose hard drive data is stored. CEUS is also a reproducing system, with solenoids attached to each key for replicating the exact motion stored from a human performance. The system tracks keys and pedals position (via their angle relative to rest level) at a standard rate of 500 Hz, and over 250 steps (8-bit encoding), which roughly means for keys a $40\mu m$ tracking accuracy for the point of the key (or about 30" for the key angle). As for hammers, two sensors monitoring the hammer ballistic travel towards the string give out its maximum velocity after it is launched by the key.

These datasets are recorded as binary files, wherein successive data chunks correspond to each timestamp — one every two milliseconds. Each chunk starts with a break code (255), followed by a 24-bit value indicating milliseconds since the recording start. It is followed by a series of 16-bit blocks, one for each key, pedal or hammer activated at this timestamp. Each block first contains the 8-bit number of the key depressed (as in MIDI, with the central C4 equal to 64, and up to 108), or pedal (109 to 111), or hammer (MIDI-key number + 128). The second 8-bit number gives out key/pedal position or hammer velocity. Datasets are stored under the `.boe` extension format.

[1] This piano is installed in a dedicated studio at BRAMS (International laboratory for Brain, Music and Sound research), *Université de Montréal.* (http://www.brams.org)

Figure 1. Imperial Bösendorfer grand piano in the BRAMS studio, with embedded CEUS system.

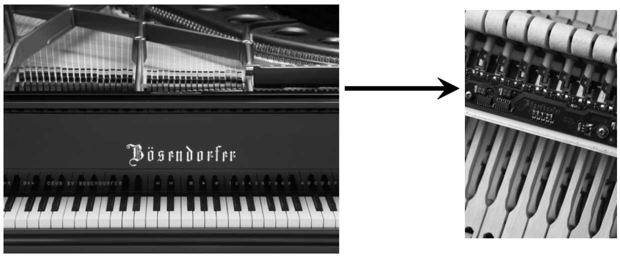

Figure 2. Details of the CEUS system: fallboard display/ interface and embedded electronics.

The Bösendorfer CEUS recording system thus constitutes an extremely precise tool to observe the finest subtleties in pianists' touch and measure the part of gesture actually efficient and transferred to the piano action. However, in order to get a clear understanding of piano touch peculiarities involved in one performance, the raw data must be processed in a more intelligible way. This is the role of the primary functions in our toolbox.

4. DATA PROCESSING

4.1. From streamed data files to piano rolls

The first requirement in the chain processing that would lead to a thorough *a posteriori* analysis of boe files, is to parse the raw files into MATLAB matrices. Binary-to-decimal conversion is handled by MATLAB, and our function `extract.m` deals with identifying each data chunk relative to one timestamp, essentially by finding the break tags "255", then storing timestamp-and-value (of key or pedal depression or hammer velocity) data couples on one line. The parser can also deal with the older boe format, wherein data is encoded in ASCII with hexadecimal numbers. We thus get a streamline of events that happened in the recorded performance.

Yet data is easier to use and interpret once restructured

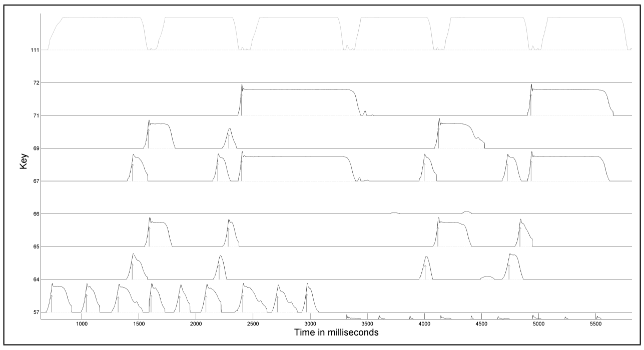

Figure 3. Pianoroll display of a performance (detail), wherein key motion is shown as blue lines, MHVs as red vertical arrows, and sustain pedal in light blue.

in a key-by-events matrix, wherein each line corresponds to one key/pedal/hammer, with its MIDI-like number as referent, and successive blocks indicate each event with its timestamp and value. From a timestamp-driven structure, the `notetime.m` function thus reorganises data into a key-by-key account of events. Special attention is given to correcting for possibly missing information, such as one key being depressed at two timestamps t and t+4 yet missing information at t+2. In all, missing or redundant timestamps are filtered, as well as singleton events. [2] This structured data is exported into a text file.

The data format is especially useful to get a visual representation of the recorded performance, in the form of a piano roll. To this aim, the `pianoroll.m` function can display keys and pedals motion, one line for each, along time. All recorded maximum hammer velocities are displayed for their relative key, with a red vertical arrow (Figure 3). This gives out the graphic equivalent of the well-known MIDI pianoroll display, but with the exact level of key depression instead of a fixed-velocity block per note.

However, within this data structure and pianoroll display, we can only observe the high-precision linear response from each key motion, and do not have direct information on the most basic musical structure, that is, the note. While visually obvious, this basic information we get immediately in MIDI here has to be retrieved through additional processing.

4.2. Retrieving notes

As events (key angle) are only registered when the key is out of its rest position, we can essentially retrieve note onsets by finding discontinuities in the timestamp sequence related to one key. Note onsets thus occur whenever two consecutive, timestamp-and-value blocks stored in "note-time" files possess non-consecutive timestamp values (*i.e.* separated by more than 2 ms). This process is run in the `notes.m` function, while improved with correction procedures for missing timestamps (up to 8 ms deep) and noisy information at note onset — *e.g.* a key somewhat being pushed down of less than a millimeter when the pianist sets his finger on top of it in preparation for the next note — which could shift onset time towards way earlier than actually performed. Retrieved notes that prove too feeble to launch the hammer and produce a sound are also filtered out. Moreover, the note-detection function can account for successive notes played on the same key, with the second note starting before the key is fully released from the first. Indeed, the double escapement action featured on grand pianos allows for faster note repetition after just a partial key release, down to a threshold called escapement point. As such two notes are not separated in the data by a discontinuity in timestamp sequence, we used a forward-loop exploration of local minima. The escapement threshold was assessed empirically through performance corpus analysis, and could be set at 140 (upon 250). Local minima lower than this threshold define the onset of a new note.

Once the notes thus identified, we set on exploring the upper-level musical structure: chords.

4.3. Identifying chords

The function `chords.m` identifies groups of notes which have near-synchronous onsets. The task is performed in two rounds. First, a group of notes is formed when their onsets are less than 50 ms apart. A chord is thus defined, whose onset is defined as the earliest note onset. Another note can then be assigned to this chord if it provides the best-fitting, under-50-ms onset timing difference. This ensures the most synchronous notes are grouped together. In the second round, chords can be merged if any of the note onsets (instead of the earliest) within one chord falls within a looser interval, 100 ms, of any of the note onsets from other chords. The 50 and 100 ms onset synchronism intervals are consistent with mean and upper-bound values found by Repp [22] and Shaffer [25] and were tested empirically for matching between the chords identified in a corpus of performances and the designations in their respective scores.

Additionally, in the context of our target experiment which involved four different pieces, we were able to set thresholds for separating the range of each hand — that is, a note above which the left hand, and below which the right hand, never play. Chords could thus be separated by hand.

From the high-precision response tracking of key and pedal depressions and hammer velocities, we thus retrieve the fundamental musical structure of notes and chords. [3] This allows us to get back up to MIDI-like note identification, add the definition of chords, and deal with the fine-grained information through reduction to exhaustive and relevant features set to describe each note and chord.

[2] This procedure has essentially become a precautionary tale, as with the latest CEUS software updates such acquisition errors have been all but eliminated.

[3] The notion of "chord" is hereby used loosely, as it can account for a single note when not deemed synchronous to any other note played by the same hand.

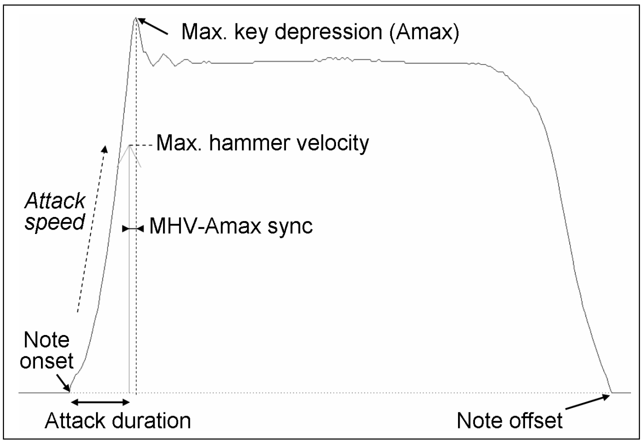

Figure 4. Illustration of some essential note features.

5. PIANO TOUCH DESCRIPTIVE FEATURES

Our toolbox also provides a thorough analysis of notes and chords, and extracts numerous features from the high-precision, 500 Hz tracking of keys and pedals depression and maximum hammer velocities record. The exhaustive information related to each musical structure of note or chord is reduced to a large set of features most relevant to understanding piano touch. The most relevant features to set alight were chosen on the basis of studies of various aspects of expressive piano performance that used MIDI data [22] [23] [2] [6], as well as studies focused on piano touch [19] [21] or keyboard action [16]. We have programmed, adapted, extended or added descriptive features which can be sorted in several broad categories: dynamic level, attack speed and type, duration and sustain, release, synchronism, intervals and overlaps, and pedals use.

5.1. Single note features

Each note is individually described by 46 features. First are its basic characteristics (Figure 4): key number, onset, offset and duration; maximum hammer velocity (MHV), maximum key depression angle (Amax) and their corresponding timestamps. From these are calculated several attack-descriptive features: attack duration (related to instants of both Amax and MHV), attack speed (as a ratio of Amax or MHV to its duration), and timing between Amax and MHV. Those features assess the attack as a description of dynamic level — the faster the attack and the more synchronous Amax and MHV, the higher dynamic level is.

Attack type can also be defined as percussive or not. Indeed, as Goebl et al. have shown [6] key trajectory differs whether the key is pressed, with the finger initially resting on the surface of the key, or struck percussively, with the finger starting above the key and reaching it with a non-zero velocity. The struck touch displays a fast initial key depression that slows down until reaching (or not) the keybed, while a pressed touch has the key depressed slowly first and gradually accelerating. We used two methods to elicit this behaviour: the ratio of key depression at half the attack duration to the maximum key depression, and the mean key depression during attack — akin to the area swept by the key depression curve during attack. Both features will have higher values with a percussive touch.

Two other ways of assessing the note profile were designed. First, critical points akin to the acoustic temporal envelope were retrieved, thus defining up to four zones in key depression: attack, decay (a short drop up in key depression certainly due to the reaction of the keybed felt, and found in many notes), sustain and release. Attack, sustain and release durations and ratio to the total note duration were thus assessed. And second, we defined empirically a threshold over which the key can be deemed deeply depressed. From there we could define three sections, and their durations, the first before the key reaches the threshold, then while the key is depressed over the threshold, and when it falls below it — akin in most cases to attack, sustain and release respectively.

Finally, we gathered sustain and soft pedals use during the note: for each, their duration of use and amount of depression during the note, as well as their depression at note onset, offset and at the instant of MHV.

5.2. Chord features

Each chord is first described by basic features: number of notes within, its onset and offset (earliest and latest of its note onsets and offsets resp.), duration and maximum of its note Amax and MHV. Then, each of its notes is assigned, besides its 46 individual features, 10 additional characteristics of its synchronism towards the chord: note onset lag on chord onset, its ratio to the chord duration, and the onset lag amount (defined as the sum of all other keys depressions within the chord before the onset of the target note); note offset lead, ratio and amount with regard to chord offset (defined in the same way); its synchronism with the chord as ratio of note duration to chord duration; and its synchronism amount, defined as the ratio of the total amount from which the other notes within the chord are depressed for the duration of the target note, compared with their total depression amount. Each note within a chord is thus described by 56 features.

And each chord in itself is also described by 56 other features, with descriptions of internal synchronism and pedal use in addition to the primary features explained above. Indeed, the smallest, non-zero note onset lag defines the chord melody lead (how early the first note is compared with the next), and the smallest, non-zero note offset lead sets the one-note trail within the chord. As for soft and sustain pedals use, both are exhaustively featured as follow: duration and amount of depression during the chord; duration of deep-depression and mid-depression (when pedal depression level falls above or between certain thresholds, resp.); and assessment at chord onset and offset of pedal depression levels, activation (on or off) and timing (how long before or after the pedal was or will be activated).

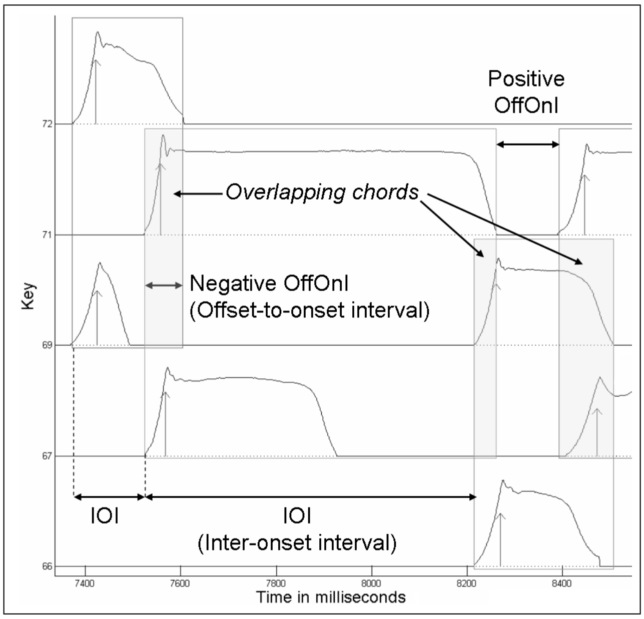

Figure 5. Pianoroll display of successive (framed) chords with their interval and overlap relations pointed out.

Finally, the following features assess the relations between chords. First, intervals are measured: inter-onset interval (IOI) from one chord onset to the next, interval from chord offset to the next chord onset (OffOnI), and its direction — negative indicates legato, positive staccato. Then, overlaps between chords are defined through their duration (how long one chord is overlapped by others), amount (of depression) and number (of chords in overlap with the target) (Figure 5). All those features of intervals and overlaps are calculated with regard to any chord (in onset temporal order), to same-hand chords only, and to chords played the other hand only. This is meant to elicit different articulation strategies.

In all, each multi-note chord is described by 56 chord-specific features, plus the mean and standard deviation of the 56 features describing each of its notes. This thus amounts to 168 features per chord. [4] Such an exhaustive account clearly states the depths at which the CEUS system lets us observe piano performance, and gather quantitative information about piano touch and its dynamics, percussiveness, articulation, depth, timing, pedalling,etc.

This features extraction is processed within the functions `notes.m` and `chords.m` (the latter calling the former), and is output in a three-dimensional, chord-by-note-by-feature matrix. For each chord of n notes, there are (n+3) vectors, the first being the features describing the chord itself, the second the mean of its note features and the third their standard deviation. Additionally, the mean and standard deviation of all chord features gives out a description of the performance, as a whole, through 322 characteristics. Moreover, each performance is similarly described with regard to left-hand chords only, right-hand chords only and left-vs-right hand differences. The results matrix can also be printed out as a formatted text file (restructured in 2D), or as two separate files accounting for chords alone and notes alone, resp.

[4] In the case of a single-note "chord", due to redundancy and no standard deviation, there remains 85 valid descriptors.

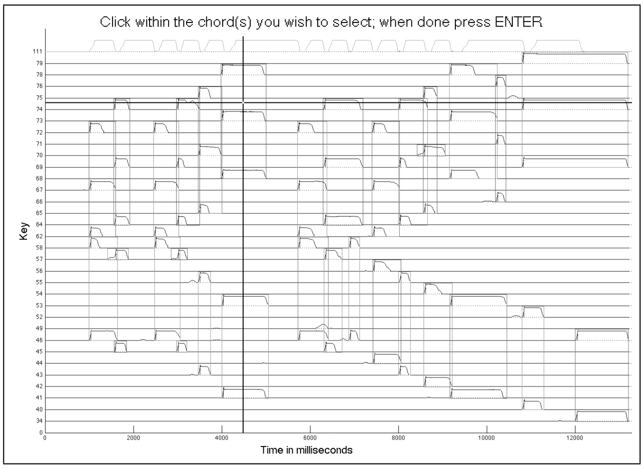

Figure 6. Chord selection interface, wherein a black target cursor (here centered around key 74 and timestamp 4500) indicates which chord a click within the frame will select.

6. ADDITIONAL FUNCTIONS

In addition to the main chord-and-note structuring and feature extracting functions, the toolbox also contains several analysis tools.

6.1. Selection

First, subsets of chords or notes can be selected with the `select_chords.m` or `select_notes.m` functions. There are two ways of selecting the subset: either by indicating notes or chords to select in a matrix or text file called as input argument, or graphically, by clicking on the chords/notes one wishes to select on the pianoroll display of the performance. With the first method, most useful for batch processing, notes to select are to be specified, each on one line, by their key number and timestamp falling between their onset and offset. The same is asked for chords to select, and here any key number falling within the range between the lowest and highest notes in the chord can be used as referent. In case such coordinates could refer to more than one note or chord (due to overlap), the closest fit in time range is selected. The graphical counterpart method works essentially the same way, with each click on the note or chord to select within the pianoroll figure sending back a key number and timestamp. Chord selection by click is eased by the framing of all chords in the pianoroll figure (Figure 6). The function then retrieves the features describing each selected note or chord, and calculates means and standard deviations over the subset. It outputs a 3D results matrix similar to the main performance-features matrix previously described, with the same printout options. This way, one can assess local characteristics within a performance.

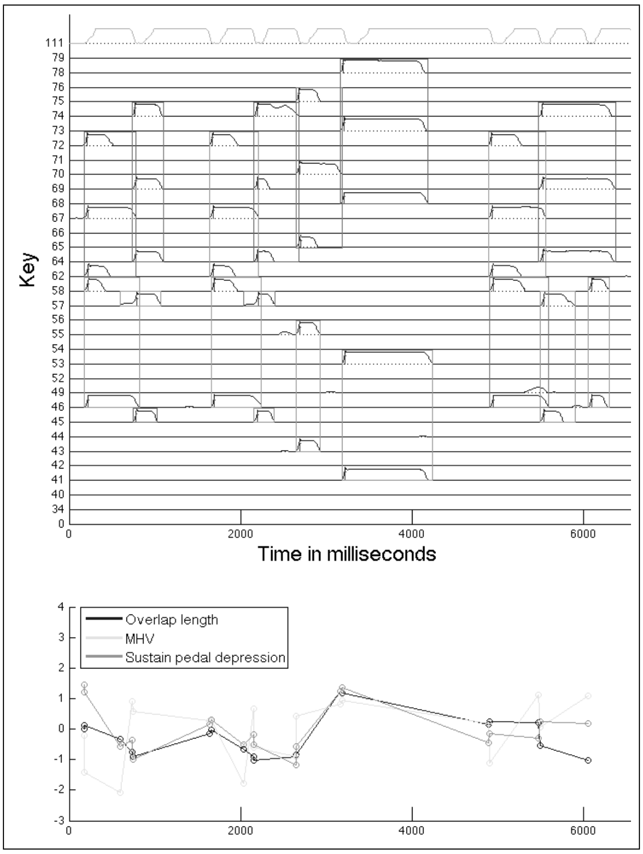

Figure 7. Graphical display of one performance piano roll (detail) and evolution over time of selected features.

6.2. Graphical feature representation

In order to visualize the evolution of features over the duration of a performance, the `g_pianostats.m` function plots the normalized feature values against the performance pianoroll display. First, a graphical user interface allows to select the features to plot and specify some options (hand separation, error bars for note features). The value of each selected feature, for each chord, is then plotted over time at the instant of chord onset (Figure 7). With the pianoroll as reference, one thus can see the evolution of the feature along the performance, and possibly identify its relation to the musical structure — *e.g.* phrasing — of the piece.

The second function, `g_compare.m`, is a graphical comparison tool, in time, of several performances. The same graphical user interface is used to select the features to display. For each selected feature, a separate plot tracks its evolution in time within each input performance, and compares its mean value within each performance. This allows to identify what differs between performances, and especially when differences occur (Figure 8).

6.3. Score matching

This function is aimed to assess the fit between two performances or between a performance and its score. It compares the notes identified in one performance to another, or

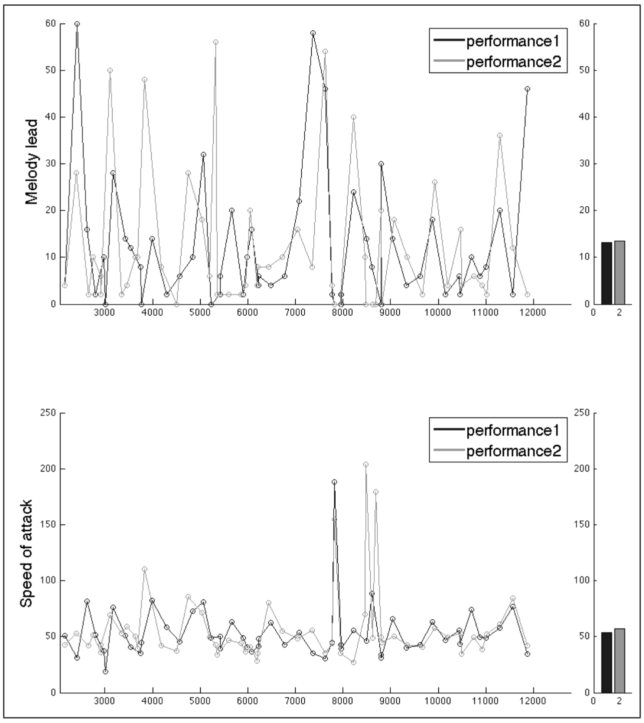

Figure 8. Comparison in time of 2 performances for 2 features, and mean thereof.

to the score (rendered as MIDI). The function only compares MIDI-like information of key number, onset, duration and velocity (MHV). It is an adaptation, with our custom, derived CEUS-data complying parser, of Ed Large's score matching functions [5] [13], essentially intended to assess performance errors (such as missing notes).

The function returns a graphical display of two corresponding MIDI piano rolls, wherein each note within one performance is linked to its corresponding note in the other performance/score. The cross-correlation matrix of the two performances/score in time is also displayed, and we added the calculation of a matching rate, which indicates the percentage of notes that match (in key number, timing, duration and velocity) between the two performances.

6.4. MIDI to boe conversion

Last but not least, we designed the `boe_gen.m` function to generate CEUS boe format files from MIDI input. First, the MIDI files are parsed into MATLAB variables with Ken Schutte and Tuomas Eerola `midi2nmat.m` function, from which output we conserve each note with its key number, onset, duration and velocity. MIDI velocity (7-bit) is linearly converted to CEUS MHV (8-bit). Our function can also directly take text files as input that use the same format as midi2nmat output. Instant of MHV can also be specified as a fifth parameter. If not (or MIDI input), instants of MHV are extrapolated as an empiric function of MHV and duration.

[5] Available online at https://www.jyu.fi/hum/laitokset/musiikki/en/... research/coe/materials/miditoolbox/matchingPerformancetoNotation.zip

Then, boe files can be generated with a straight rendering of the input information, that is, with MIDI-style notes of constant MHV/velocity along their duration, and a crenelated outlook. Yet the function can also generate more realistic boe files wherein key depression patterns are rendered as if key motion were precisely tracked. Depending on the input parameters, each note is assigned one of three fine-grained note prototypes each representing a typical key depression profile — as commonly identified in CEUS boe recordings. The note prototype is then warped by polynomial interpolation so as to fit the note input parameters, with longer durations accounted for by stretching the sustain phase (thus keeping attack and release in valid forms). The key-by-events "notetime" matrix is first created, and then transcribed in the output boe file.

This MIDI (or MIDI-like) conversion and/or augmentation into the CEUS system boe format has proven useful for testing the feature-extracting functions, and for comparing performances, feature-wise, to the "flat" reference of their MIDI-rendered scores.

7. APPLICATION: ANALYSIS OF PIANO TIMBRE GESTURAL CONTROL

7.1. Context

Amongst the various aspects of expressivity in piano performance, our research project focuses on timbre, and the control thereof through subtle nuances of touch and applied gesture. In order to investigate this timbre and tone dimension so widely acknowledged within the pianistic community as central to the art of piano playing [18] [4], we designed an experiment set in a musically relevant context. This has let us explore farther than the single-note studies which helped reinforce the long-held notion within the scientific world that mechanical constraints limited piano timbre control, for a single, isolated piano key, to the sole intensity of key stroke [10]. The study thus follows the groundbreaking, yet somehow overlooked, work of Ortmann [19], who could identify different patterns of key depression depending on tone intentions. He also posited that tone combinations and blends, when several notes are played simultaneously or successively and pedals are used — as is the case in musical context — add so many possibilities and factors of tone control [8]. We thus explored how features of articulation, timing, dynamics, attack, touch percussiveness, pedalling, and so on are nuanced in musical performances in order to produce different tones and timbres.

7.2. Method

To this aim, we first identified five timbre descriptors most salient and representative of the piano timbre range [1]: *Dry*, *Bright*, *Round*, *Velvety* and *Dark*. Four miniature pieces were composed so that each could be fittingly performed with each of the five timbres. And with the CEUS data acquisition system, we could record performances by four pianists with professional experience. They each played each piece, tree times with each of the five timbres. We thus compiled 240 boe recordings, from which we extracted the exhaustive characteristic features thanks to our Bosen Toolbox.

7.3. Analysis

We then set to compare all performances through their mean features, with regard to their respective timbre. To this aim, we developed several analysis functions as an extension of the Bosen Toolbox. These functions allow for automated processing of all features through several statistical tests from MATLAB Statistics Toolbox.

The main function performs several variance tests over each feature, with timbre as factor. The three tests used are one-way ANOVA, Welch robust test of equality of means, and Kruskal-Wallis non-parametric rank analysis of variance, with the assumptions required for test validity (normality of same-timbre groups distribution for ANOVA and Welch, and Levene's homogeneity of variance for ANOVA) tested as well. For each feature, with the p-value of the most powerful test whose assumptions are not violated as indicator of significance at the .05 level, we thus know whether this feature significantly varies depending on timbre, and therefore whether the feature can be useful in determining the timbre performed. If so, post-hoc tests (with Tukey's honestly significant difference criterion) are run to determine the timbre pair-wise significance for the feature, that is, which two timbres among all pair combinations the feature values can set apart from one another. All those results, for all features, are returned in a table.

In addition, the normalized means and standard deviations over all performances of same timbre are calculated, for each timbre and each feature. Those values can be graphically compared with a linear plotting function. Moreover, normalized feature values, regrouped by timbre, can be directly and more thoroughly displayed as box plots. Finally, the same information can be represented as a Kiviat graph — a.k.a. "radar" or "cobweb" chart.

Furthermore, a function was designed to perform Principal Component Analysis over all significant features. PCA identifies, one by one, the linear combination of features (with individual weights assigned) which can explain as much of the total variance (or of the remaining variance) as possible. Within the space thus defined by those dimensions, each performance possesses its own coordinates. Our function selects all dimensions that explain a large enough chunk of the total variance (over 10%), and traces all the plans formed by the remaining dimensions, with the performances (color-tagged for timbre) set therein according to their coordinates. The description of those dimensions as linear combinations of the significant features is also stored. This method thus gives a rapid way to identify which combination of features is best able to differentiate between timbres.

7.4. Results

All these methods and processes were applied to the performances recorded for the experiment and their gestural descriptors extracted with the Bosen Toolbox. We separately analyzed the whole set of descriptors, and their means and deviations per performance, for full performances, for their left- or right-hand chords only, and for performances grouped piece by piece and pianist by pianist. While the complete interpretation of these exhaustive analyses is still currently under way, we have already identified significant correlations of several gesture descriptors to the timbre performed. The eight most representative among all significant timbre-discriminating gestural features are displayed as a Kiviat graph (Figure 9). Thus are elicited the different strategies in dynamics, pedalling, articulation, *rubato* and touch employed to perform a certain timbre. We can thus gather how piano touch and gesture features were used to control and vary timbre nuances. For instance, a counter-clockwise account of the Kiviat graph, starting to the right, indicates that playing with a *velvety* timbre required low yet quite varying dynamics, variation in attacks, heavy use of the soft pedal and to a lesser extent of the sustain pedal, a lot of *legato*, rather stable chord durations (*non rubato)* and a very soft, non-percussive touch. By comparison, playing *dark*, while most similar to playing *velvety* according to those eight gesture descriptors, involves a little more dynamic power and less variations, while speeds of attack still vary greatly. [6] The sustain pedal is used more, and the soft pedal less so. The articulation is even more *legato*, while the *dark* touch is much more percussive than the *velvety* touch.

The results thus consist in effect in the gestural mapping of piano timbre, to the extent of the gesture descriptors extracted and the timbres considered. The coarse timbre mapping with only these eight select features still allows for a unique description of the five timbres. More thorough analyses, involving more gesture descriptors for a finer account and interpretation of timbre gestural mapping, are in the works.

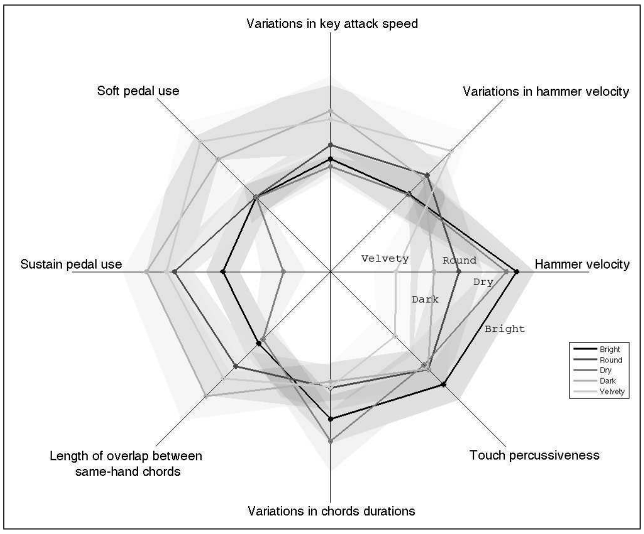

Figure 9. Kiviat plot of the eight most representative timbre-discriminating gesture descriptors (shaded zones correspond to $\pm$ 2 S.E.).

8. DISCUSSION

The Bosen Toolbox was developed to make the most use of the high-accuracy, high sample-rate Bösendorfer CEUS key/hammer/pedal tracking system data, and to offer an exhaustive and thorough account of piano touch and gesture, through meaningful features that can be interpreted in a musically relevant way.

Moreover, the toolbox statistical functions allow for analysis automation, and can be easily adapted to studying whichever performance factor, with a single configuration file to define all variables.

And, while the toolbox was especially designed for the CEUS data format, it could conceivably be used with any similar high-precision equipment, capturing the keyboard discretized per key (up to 108 discrete units) on one dimension, and high-accuracy key depression tracking on another. For instance, a high-frame-rate video tracking system could be positioned in such a way that its field of view would encompass a whole keyboard and key depressions would be visible (*e.g.* with the camera set at one end of the keyboard, or somewhere below it). The same would hold true for a motion capture system with reflectors installed on each key. More creative designs could also be conceived that do not even involve a keyboard. . .

To those ends, we plan on making the Bosen Toolbox, once finalized, available online — as a set of MATLAB m-files — under GNU licence. For now, the toolbox can be sent upon request.

In conclusion, the Bosen Toolbox is meant to help better understand the subtle nuances in expressive musical performance through which artists masterfully manage to convey emotion and feeling.

9. ACKNOWLEDGEMENTS

We wish to thank all the pianists who participated in the study. This research is supported by FQRSC (*Fonds de recherche du Québec - Société et culture*) through a doctoral research grant. Additional financial support was provided by CIRMMT (Centre for Interdisciplinary Research in Music Media and Technology) and OICRM (*Observatoire interdisciplinaire de création et de recherche en musique*). Research equipment, most crucially the Bösendorfer CEUS digital recording grand piano, is provided by BRAMS (International Laboratory for Brain, Music and Sound Research).

[6] Variations in hammer and key attack velocities resp. may not be linearly correlated, although they both concern the same broad-sense action, because they can be affected differently by differences in dynamic register and touch percussiveness.

10. REFERENCES

[1] Bernays, M. & Traube, C. "Verbal expression of piano timbre: Multidimensional semantic space of adjectival descriptors", in A. Williamon, D. Edwards & L. Bartel (Eds.), *Proceedings of the International Symposium on Performance Science (ISPS2011)*, Toronto, ON. European Association of Conservatoires (AEC), Utrecht, Netherlands, pp. 299–304, 2011.

[2] Bresin, R. & Battel, G.U. "Articulation strategies in expressive piano performance", *Journal of New Music Research*, 29(3): 211–224, 2000.

[3] Clarke, E.F., Parncutt, R., Raekallio, M. & Sloboda, J.A. "Talking fingers: an interview study of pianists' views on fingerings", *Musicae Scientiae*, I(1): 87–107, 1997.

[4] Dawei, Y. "Melody and timbre in piano performance", *Canadian Social Science*, 2(1): 69–72, 2006.

[5] Goebl, W. *The Role of Timing and Intensity in the Production and Perception of Melody in Expressive Piano Performance*. PhD thesis, *Institut für Musikwissenschaft, Karl-Franzens-Universität*, Graz, Austria, 2003.

[6] Goebl, W., Bresin, R. & Galembo, A. "Touch and temporal behavior of grand piano actions", *Journal of the Acoustical Society of America*, 118(2): 1154–1165, 2005.

[7] Goebl, W., Flossmann, S. & Widmer, G. "Investigations into between-hand synchronization in Magaloff's Chopin", *Computer Music Journal*, 34(3): 35–44, 2010.

[8] Gustafson, A.E. *Tone production on the piano: the research of Otto Rudolph Ortmann'*. DMus thesis, University of Texas at Austin, 2007.

[9] Hadjakos, A., Aitenbichler, E. & Mülhäuser, M. "Potential use of inertial measurement sensors for piano teaching systems: motion analysis of piano playing patterns", *Proceedings of the 4th i-Maestro Workshop on Technology-Enhanced Music Education*, Genova, Italy, pp. 61–68, 2008.

[10] Hart, H.C., Fuller, M.W. & Lusby, W.S. "A precision study of piano touch and tone", *Journal of the Acoustical Society of America*, VI: 80–94, 1934.

[11] Heinlein, C.P. "A discussion of the nature of pianoforte damper-pedalling together with an experimental study of some individual differences in pedal performance", *Journal of General Psychology*, 2: 489–508, 1929.

[12] Kochevitsky, G. *The Art of Piano Playing: A Scientific Approach*. Summy-Birchard Music, Secaucus NJ, 1967.

[13] Large, E.W. "Dynamic programming for the analysis of serial behaviors", *Behavior Research Methods*, 25(2): 238–241, 1993.

[14] MacRitchie, J. *Elucidating Musical Structure through Empirical Measurement of Performance Parameters*. PhD thesis, University of Glasgow, UK, 2011.

[15] Matthay, T. *The Visible and Invisible in Pianoforte Technique*. Oxford University Press, 1932.

[16] McPherson, A. & Kim, Y. "Multidimensional gesture sensing at the piano keyboard", *Proceedings of the 29th ACM Conference on Human Factors in Computing Systems (CHI 2011)*, Vancouver, B.C., pp. 2789–2798, 2011.

[17] Moore, F.R. "The dysfunctions of MIDI", *Computer Music Journal*, 12(1): 19–28, 1988.

[18] Neuhaus, H. *The Art of Piano Playing*. Trans. from original edition (1958) by K.A. Leibovitch, Barrie & Jenkins, London, 1973.

[19] Ortmann, O.R. *The Physiological Mechanics of Piano Technique: An Experimental Study of the Nature of Muscular Action as Used in Piano Playing and of the Effects Thereof Upon the Piano Key and the Piano Tone*. Reprint of original edition (1929) by E.P. Dutton, New York, 1962.

[20] Parncutt, R., Sloboda, J.A. & Clarke, E.F. "Interdependence of right and left hands in sight-read, written, and rehearsed fingerings of parallel melodic piano music", *Australian Journal of Psychology*, 51(3): 204–210, 1999.

[21] Parncutt, R. & Troup, M. "Piano", in R. Parncutt & G. McPherson (Eds.), *The Science and Psychology of Music Performance: Creative Strategies for Teaching and Learning*, Oxford University Press, pp. 285–302, 2002.

[22] Repp, B.H. "Patterns of note onset asynchronies in expressive piano performance", *Journal of the Acoustical Society of America*, 100(6): 3917–3932, 1996.

[23] Repp, B.H. "Acoustics, perception, and production of legato articulation on a computer-controlled grand piano", *Journal of the Acoustical Society of America*, 102(3): 1878–1890, 1997.

[24] Seashore, C.E. *Objective Analysis of Music Performance*. University of Iowa Press, 1936.

[25] Shaffer, L.H. "Performances of Chopin, Bach, and Bartók: Studies in motor programming", *Cognitive Psychology*, 13: 326–376, 1981.

[26] Thompson, M.R. & Luck, G. "Effect of pianists' expressive intention on amount and type of body movement", *Proceedings of the 10th International Conference on Music Perception and Cognition (ICMPC10)*, Sapporo, Japan, pp. 540–544, 2008.

[27] Vernon, L.N. "Synchronization of chords in artistic piano music", in C.E. Seachore (Ed.), *Objective Analysis of Music Performance*, University of Iowa Press, pp. 307–345, 1936.