

AUTOMATED DETECTION OF SINGLE- AND MULTI-NOTE ORNAMENTS IN IRISH TRADITIONAL FLUTE PLAYING

Münevver Köküer^{1,2}, Peter Jančovič², Islah Ali-MacLachlan¹, Cham Athwal¹

¹DMT Lab, Birmingham City University, UK

²School of Electronic, Electrical & Systems Engineering, University of Birmingham, UK
{munevver.kokuer, islah.ali-maclachlan, cham.athwal}@bcu.ac.uk
p.jancovic@bham.ac.uk

ABSTRACT

This paper presents an automatic system for the detection of single- and multi-note ornaments in Irish traditional flute playing. This is a challenging problem because ornaments are notes of a very short duration. The presented ornament detection system is based on first detecting onsets and then exploiting the knowledge of musical ornamentation. We employed onset detection methods based on signal envelope and fundamental frequency and customised their parameters to the detection of soft onsets of possibly short duration. Single-note ornaments are detected based on the duration and pitch of segments, determined by adjacent onsets. Multi-note ornaments are detected based on analysing the sequence of segments. Experimental evaluations are performed on monophonic flute recordings from Grey Larsen's CD, which was manually annotated by an experienced flute player. The onset and single- and multi-note ornament detection performance is presented in terms of the precision, recall and F -measure.

1. INTRODUCTION

Within Irish traditional music, ornaments are used extensively by all melody instruments. They are central to the style of the music, adding to its liveliness and expression. Amongst traditional players, the melody is merely a framework [3, 4] – dynamics, ornamentation and context will be added in real time. This is often different from classical music where a standard notation for each piece of music usually includes ornaments as written by the composer.

Ornaments are notes of a very short duration. They can be categorised into single-note and multi-note ornaments. Single-note ornaments are amongst the most common in Irish traditional music. Multi-note ornaments consist of a specific sequence of note and single-note ornaments.

Methods for ornament detection are typically based on detection of note onsets. Note onsets may be categorised as hard or soft. A hard onset, typical in percussive instruments, is characterised by a sudden change in energy. A soft onset shows a more gradual change in energy and it occurs in wind instruments, like flute. A variety of methods have been proposed for the detection of note onsets in music recordings, e.g., [1, 8, 11, 13, 17]. The methods typically exploit the change in the energy of the signal, which may be estimated in temporal or spectral domain. The use of phase has also been investigated, e.g., [1, 11], and combined with the fundamental frequency in [11]. It has been reported that reliable note onset detection for non-percussive instruments is more difficult to obtain due to the soft nature of the onsets [11].

An automated detection of ornaments is a challenging problem. This is because ornaments are of very short durations, which may cause them being easily omitted or falsely detected. Unlike note onset detection, this research area has received relatively little attention. An automatic location of ornaments for flute recordings based on MPEG-7 features was investigated in [5]. Transcription of baroque ornaments in two piano recordings by analysing rhythmic groupings and expressive timing was studied in [2]. This work used onset values from manually edited time-tagged audio. Several works employed spectral-domain energy-based onset detection, e.g., [9, 10, 16]. The work in [16] analysed ornamentation from Bassoon recordings. The work of a group from Dublin Institute of Technology, summarised in [9], is the only study on the detection of ornaments in Irish traditional flute music. This provided only some initial results and on a considerably smaller dataset.

In this paper, we extend our recent work presented in [14] and investigate automatic detection of single- and multi-note ornaments in flute playing. The presented ornament detection system is based on first detecting onsets and then exploiting knowledge of musical ornamentation. We explore the use of several different methods for onset detection and customisation of their parameters to detection of soft onsets of notes which may be also of very short duration. The detected onsets provide segmentation of the signal, where a segment is defined by the adjacent detected onsets. This segmentation, together with the musical knowledge of ornamentation is then used for the detection of single- and multi-note ornaments. Experimental evalua-



© Münevver Köküer^{1,2}, Peter Jančovič², Islah Ali-MacLachlan¹, Cham Athwal¹.

Licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). **Attribution:** Münevver Köküer^{1,2}, Peter Jančovič², Islah Ali-MacLachlan¹, Cham Athwal¹. "AUTOMATED DETECTION OF SINGLE- AND MULTI-NOTE ORNAMENTS IN IRISH TRADITIONAL FLUTE PLAYING", 15th International Society for Music Information Retrieval Conference, 2014.

tions are performed using recordings of Irish traditional tunes played by flute from Grey Larsen’s CD [15]. Results of ornament detection are presented in terms of the precision, recall and F -measure. The average F -measure performance for single- and multi-note ornaments is over 76% and 67%, respectively.

2. SINGLE- AND MULTI-NOTE ORNAMENTS IN IRISH TRADITIONAL FLUTE PLAYING

Ornaments are used as embellishments in Irish traditional music [15]. They are notes of a very short duration, created through the use of special fingered articulations.

Single-note ornaments, namely ‘cut’ and ‘strike’, are pitch articulations. The ‘cut’ involves quickly lifting and replacing a finger from a tonehole, and corresponds to a higher note than the ornamented note. The ‘strike’ is performed by momentarily closing an open hole, and corresponds to a lower note than the ornamented note.

Multi-note ornaments are successive use of single-note ornaments. To simplify the description, we refer to the ornamented note as the base note throughout the rest of this paper. The ‘roll’ consists of the base note, a ‘cut’, base note, a ‘strike’ and then returning to the base note. A shorter version of the roll, referred to as short-roll, omits the starting base note. The ‘crann’ consists of the base note that is cut three times in rapid succession and then returning to the base note. The short-crann omits the starting base note. The ‘shake’ commences with a ‘cut’, followed by a base note and a second ‘cut’ and then returning to the base note.

A schematic visualisation of the single- and multi-note ornaments is given in Figure 1. In the multi-note ornaments figure, the proportions of the length of the individual parts aim to approximately indicate the typical duration proportions. For instance, in theory, a roll would be split equally into three parts by the cut and the strike but in reality different players will time this differently according to the ‘swing’ of the tune, their muscle control and a host of other attributes that make up their personal style.

3. AUTOMATIC DETECTION OF ORNAMENTS

This section presents the developed automatic ornament detection system. We first give a brief description of the onset detection methods we employed and then describe how the detected onsets are used for the detection of single- and multi-note ornaments.

3.1 Methods for detection of onsets

Here we briefly describe three onset detection methods we employed. Two of the methods exploit the change of the signal amplitude over time, with processing performed in the temporal and spectral domain [1, 8]. The third method is based on the fundamental frequency [6, 11]. Each of the method requires several parameters to be set and their values are explored during experimental evaluations and presented later in Section 4.3. The implementation of the

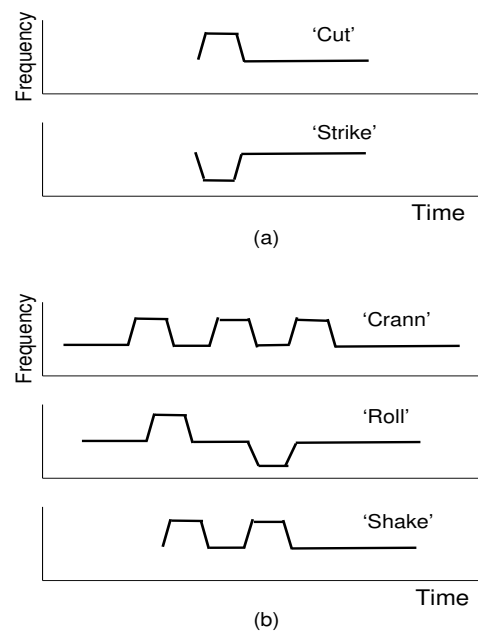


Figure 1. A schematic representation of single-note (a) and multi-note (b) ornaments.

temporal domain energy-based method used in parts some functions from the MIRtoolbox.

3.1.1 Signal energy: spectral domain

This method, also sometimes referred to as spectral-flux method, performs onset detection in the spectral domain. The signal is segmented into overlapping frames. Each signal frame is multiplied by Hamming window. The windowed frames are then zero padded and the Fourier transform is applied to provide the short-term Fourier spectrum. For each frequency bin, the differences between the short-term magnitude spectra of successive signal frames is computed. This is then half-wave rectified and the L_2 norm is calculated to provide the value of the detection function at the current frame. The peaks of the detection function, whose amplitude is above a threshold are used as the detected onsets. We explored the use of a fixed threshold value as well as computing the value adaptively based on the median of the detection function values around the current frame. Finally, if two consecutive peaks are found within a given time distance, only the first peak is used.

3.1.2 Signal energy: temporal domain

Another method we employed performs the detection in temporal domain. The signal is passed through a bank of fourteen band pass filters, each tuned to a specific note on the flute in the range from D_4 to B_5 . The filters have non-overlapping bands, with the lower and the upper frequency being half way between the adjacent note frequencies. These fourteen notes are readily playable on an unkeyed concert flute. The signal in each band is full-wave rectified and then smoothed, resulting in amplitude envelope. The time derivative of the amplitude envelope is calculated in each band and this is smoothed by convolving

it with a half-Hanning window. We explored several ways of making decision about detected onsets. The information from all bands can be combined by summing together their smoothed derivative signals. Alternatively, a single band can be chosen as a representative at each time based on assessment of amplitudes of peaks around that time across all bands. Onsets are obtained by comparing the values of peaks to a threshold, which may be fixed or adaptive over time.

3.1.3 Fundamental frequency

In addition to methods exploiting the signal envelope, we also explore the use of the fundamental frequency (F_0). This has been reported to be beneficial for soft onset detection in [11]. Among a large variety of existing F_0 estimation algorithms, we employed the YIN algorithm [7] in this work. The F_0 estimation may result in so called doubling / halving errors. To help dealing with these errors, the F_0 estimates are postprocessed using a median filter. The length of this filter needs to be set sensitively – a longer filter may be preferable to deal with the F_0 estimation errors but this may also cause filtering out ornaments, which are characterised by their short duration.

The detection function at the frame time n , denoted as R_n , is based on calculating the change of F_0 over time. This can be performed by taking the difference between the F_0 estimate at the frame $(n + \Theta)$ and $(n - \Theta)$. The onset is detected as the first frame for which $abs(R_n) > \alpha_{F_0}$, where the value of the threshold α_{F_0} relates to the difference between frequencies of the closest possible notes.

3.2 Ornament detection

The detected onsets, as obtained using the methods described in Section 3.1, provide a segmentation of the signal, where each segment is formed based on the adjacent detected onsets.

We characterise each detected segment by some features, specifically, here we use the duration of the segment and its segmental fundamental frequency. For a given segment, its duration, denoted by D^{seg} , is obtained based on the detected onsets and its fundamental frequency, denoted by F_0^{seg} , is calculated as the median value of the F_0 s corresponding to all signal frames assigned to that segment. Finally, these segment features are used to determine whether the detected segment corresponds to a note or a single-note ornament and whether the sequence of segments corresponds to a multi-note ornament, and if single- or multi-note ornament is detected, then to determine its type.

3.2.1 Single-note ornament detection

As single-note ornaments are expected to be of a shorter duration than notes, we examined whether the duration of the detected segments can be used to discriminate these ornaments from notes. We conducted statistical analysis of the duration of notes and single-note ornaments in our recordings. This was performed using the manual onset annotations. The obtained distributions of the durations are depicted in Figure 2 – these indicate that the duration

can indeed provide a good discrimination between notes and ornaments. Based on these results, we consider that a segment is classified as a single-note ornament when its duration is below 90 ms, otherwise it is classified as a note.




Figure 2. The distribution of the duration of single-note ornaments (a) and notes (b) obtained using the development set.

The decision whether the detected single-note ornament is a ‘cut’ or ‘strike’ can be made based on comparing the values of the F_0^{seg} of the current and the following segment. This reflects the musical knowledge of ornamentation. If F_0^{seg} of the segment detected as ornament is higher than F_0^{seg} of the following segment, the ornament is classified as ‘cut’ and as ‘strike’ otherwise.

3.2.2 Multi-note ornament detection

The detection of multi-note ornaments, namely ‘crann’, ‘roll’ and ‘shake’, is based on analysing the features of a sequence of detected consecutive segments. We used a set of rules to determine whether the sequence corresponds to one of the multi-note ornament types or not. These rules reflect the definition of the multi-note ornaments as presented in Section 2 and for each ornament type are described below. Let us consider that r denotes the index of the first segment in the sequence of detected segments we are currently analysing. Let us denote by $\Delta F_0^{seg}(j, j+2)$ the difference between the F_0^{seg} for the segment $(r+j)$ and F_0^{seg} for the segment $(r+j+2)$, where j is an index to be set.

‘Crann’ is detected if the following is fulfilled: i) the sequence of F_0^{seg} follows the pattern ‘BHBHBHB’, where ‘B’ stands for a base note and ‘H’ for a note higher than the base note; ii) the segmental F_0^{seg} is similar for segments corresponding to the base note, i.e., the $\Delta F_0^{seg}(j, j+2)$ is within the given tolerance range β_{F_0} when j is individually set to 0, 2, and 4; and iii) the segment duration D^{seg} is below β_D for segments given by setting j from 1 to 5 and is above β_D for j set to 0 and 6. The ‘Short-Crann’ is using the same rules but taking into account that the starting base note is omitted.

‘Roll’ is detected if the following is fulfilled: i) the sequence of F_0^{seg} follows the pattern ‘BHBLB’, where ‘L’ stands for a note lower than the base note; ii) the value of $\Delta F_0^{seg}(j, j+2)$ is within the tolerance range β_{F_0} for j set

to 0 and 2; and iii) the segment duration D^{seg} is above β_D for j being 0 and 2 and is below β_D when j is 1 and 3. Again, the ‘Short-Roll’ is using the same rules but taking into account that the starting base note is omitted.

‘Shake’ is detected if the following is fulfilled: i) the sequence of notes follows the pattern ‘HBHB’; ii) the value of $\Delta F_0^{seg}(j, j+2)$ is within a given tolerance range β_{F_0} for j set to 1; and iii) the segment duration D^{seg} is below β_D when j is 1 and 2, and is above β_D when j is 3.

The parameters β_{F_0} and β_D were set to 20 Hz (except for ‘crann’ when 30 Hz was used) and 90 ms, respectively.

4. EXPERIMENTAL RESULTS

4.1 Data description

Evaluations are performed using recordings of Irish traditional tunes and training exercises played by flute from Grey Larsen’s CD which accompanied his book “Essential Guide to Irish Flute and Tin Whistle” [15]. The tunes are between 20 sec and 1 min 11 sec long. All recordings are monophonic and are sampled at 44.1 kHz sampling frequency. Manual annotation of the recordings to indicate the times of onsets and offsets and the identity of notes and ornaments was performed by the third author of this paper, who is a highly experienced musician with over 10 years of flute playing. The manual annotation is used as the ground truth in evaluations. The data was split into separate development and evaluation sets. The development set, consisting of 6 tunes (namely ‘Study5’, ‘Study6’, ‘Study17’, ‘Lady on the Island’, ‘The Lonesome Jig’, ‘The Drunken Landlady’), was used for finding the best parameter values of onset detection methods. The evaluation set, consisting of 13 tunes, was used to obtain the presented results. The list of the tunes from the evaluation set, with the number of notes and ornaments, is given in Table 1. In total, this set contains 3025 onsets, including notes and ornaments. Out of these there are 301 single-note ornaments, consisting of 257 cuts and 44 strikes, and 152 multi-note ornaments, consisting of 117 rolls (including short-rolls), 19 cranns (including short-cranns), and 16 shakes.

4.2 Evaluation measures

Performance of the onset and ornament detection is evaluated in terms of the precision (P), recall (R) and F -measure. The definition of these measures is the same as used in MIREX onset detection evaluations, specifically,

$$P = \frac{N_{tp}}{N_{tp} + N_{fp}}, R = \frac{N_{tp}}{N_{tp} + N_{fn}}, F = \frac{2PR}{P + R}$$

where N_{tp} is the number of correctly detected onsets / ornaments and N_{fp} and N_{fn} is the number of inserted and deleted onsets / ornaments, respectively. The onset detection is considered as correct when it is within ± 50 ms around the onset annotation.

The single-note and multi-note ornaments are considered to be detected correctly when the onsets, corresponding to the start and to the end of the ornament are within ± 50 ms and ± 100 ms range, respectively.

Tune Title	Number of		Time (sec.)
	Notes	Ornaments (C-S-Ro-Cr-Sh)	
Study 11	76	20-0-0-0-0	26
Study 22	127	0-28-0-0-0	47
Maids of Ardagh	98	23-0-5-0-0	32
Hardiman the ..	112	12-0-7-1-0	28
The Whinny Hills ..	117	15-1-5-2-4	30
The Frost is All ..	151	27-2-12-0-0	41
The Humours of ..	289	59-7-12-14-0	82
The Rose in the ..	152	22-2-11-0-0	39
Scotsman over ..	153	18-0-9-2-0	38
A Fig for a Kiss	105	17-3-6-0-2	28
Roaring Mary	176	15-1-21-0-3	44
The Mountain Road	105	8-0-6-0-3	25
The Shaskeen	181	21-0-23-0-4	42

Table 1. The list of tunes contained in the evaluation set, with the number of onsets and ornaments and duration of each tune. The notation ‘C’, ‘S’, ‘Ro’, ‘Cr’ and ‘Sh’ stands for ‘cut’, ‘strike’, ‘roll’, ‘crann’ and ‘shake’, respectively.

4.3 Results of onset detection

We have performed extensive evaluations on the development set with different parameter values for each of the onset detection method. The best values of parameters for each of the method are given in Table 2. The achieved performance on the evaluation set using these parameters for each method is presented in Table 3. Note that these results include the onsets corresponding to both notes and ornaments. Performance difference of less than 1% was observed when the parameters were tuned specifically for the evaluation set. It can be seen that all methods provide good onset detection performance, with the F_0 -based method being slightly better than the energy-based methods. A method based on F_0 was shown to perform best for wind instruments also in [11], where it was also shown that its combination with other methods provided only slight improvement at similar P and R values. As such, in the following, we use only the F_0 -based method for evaluating the ornament detection performance. An example of a signal extract from one of the tune and the corresponding F_0 estimate and the detection function, with indicated true label and detected onsets, are depicted in Figure 3.

4.4 Results of single-note ornament detection

The results of single-note ornament detection are presented in Table 4 separately for ‘cut’ and ‘strike’. The achieved detection performance is significantly higher than that presented in previous flute studies using similar data [9]. The performance for ‘cut’ is close to the overall onset detection performance as presented in Table 3. The performance for ‘strike’ is considerably lower than for ‘cut’. This has also been observed in previous research and may be due to the nature the ‘strike’ is created. There was 5 substitutions of

Onset detection method with best values of the parameters	
sig-energy (spectral):	
–	frame length of 1024 samples (23.2 ms)
–	frame shift of 896 samples (20.3 ms)
–	threshold set as fixed at 2% of the maximum of the normalised detection function
–	minimum distance between peaks set to 10 ms
sig-energy (temporal):	
–	half-Hanning window of 35 ms
–	threshold set as fixed at 15% of the maximum of the normalised detection function
–	minimum distance between peaks set to 20 ms
F_0 :	
–	frame length of 1024 samples (23.2 ms)
–	frame shift of 128 samples (2.9 ms)
–	median filter of length 9 frames
–	parameter Θ set to 6 frames (17 ms)
–	parameter α_{F_0} set to 10 Hz

Table 2. Parameters of each onset detection method and their best values obtained based on the development set.

Algorithm	Evaluation performance (%)		
	Precision	Recall	F -measure
sig-energy (spectral)	94.9	85.0	89.7
sig-energy (temporal)	87.9	88.6	88.3
F_0	89.1	92.9	91.0

Table 3. Results of onset detection obtained by each of the employed method.

cut for strike and 1 substitution of strike for cut. These errors were contributed by slight inaccuracies in onset detection and F_0 misestimation.

	Single-note Ornament Detection		
	Precision (%)	Recall (%)	F -measure (%)
Cut	88.4	86.4	87.4
Strike	63.8	68.2	65.9

Table 4. Results of single-note ornament detection obtained by employing the F_0 -based onset detection method.

4.5 Results of multi-note ornament detection

Experiments for multi-note ornament detection were performed by analysing all the possible sequence patterns resulting from the detected segments – this consisted here of 3020 sequence pattern candidates. The results of multi-note ornament detection are presented in Table 5 separately for ‘roll’, ‘crann’ and ‘shake’. These results include also the short versions for ‘roll’ and ‘crann’. It can be seen that

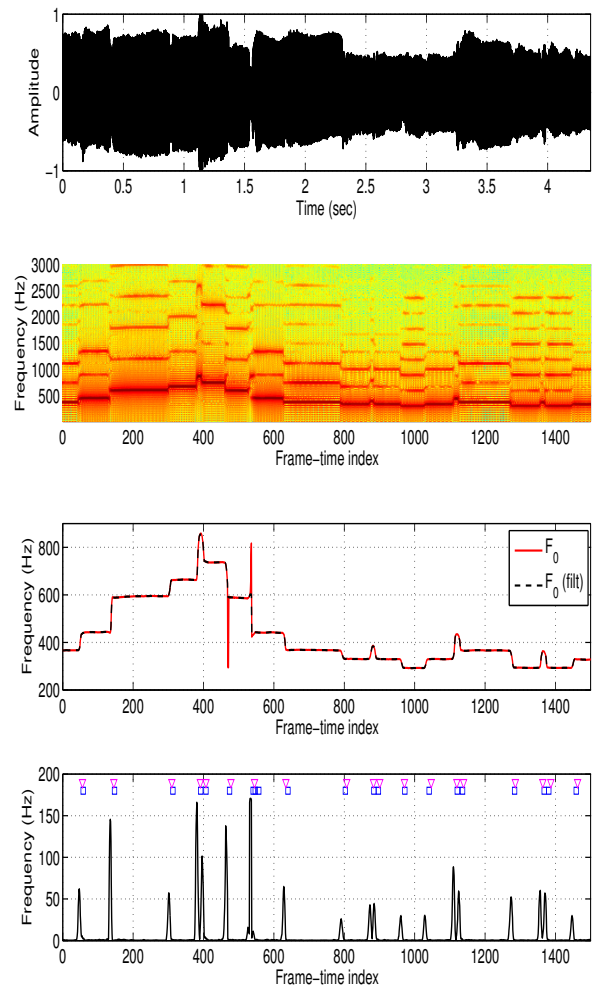


Figure 3. An extract from the tune ‘The Lonesome Jig’, depicting (from top to bottom) the waveform, spectrogram, F_0 estimation (unfiltered (red) and filtered (dashed black)) and the detection function with indicated detected onsets (blue \square) and true label (magenta ∇).

the performance for ‘shake’ is considerably lower than that for ‘roll’ and ‘crann’. This is due to the short sequence pattern of ‘shake’, consisting of only 4 parts, which makes it more likely to be accidentally match with other note sequence. We have also analysed the performance separately for the short and normal versions of the ‘roll’ and ‘crann’ ornaments. This showed that the F -measure performance for ‘roll’ was approximately 17% better than for ‘short-roll’. This trend was not observed for ‘short-crann’, which may be due its longer note sequence.

5. CONCLUSION AND FUTURE WORK

In this paper, we presented work on detection of single- and multi-note ornaments in Irish traditional flute music. We employed three different methods for onset detection and customised their parameter values to detecting soft onsets of possibly very short notes. The method based on the fundamental frequency (F_0) achieved around 91% onset detection performance in terms of the F -measure and

Multi-note Ornament Detection			
	Precision (%)	Recall (%)	F -measure (%)
Roll	87.5	67.0	75.9
Crann	86.7	68.4	76.5
Shake	50.0	50.0	50.0

Table 5. Results of multi-note ornament detection obtained by employing the F_0 -based onset detection method.

outperformed slightly the other two energy-based methods. The F_0 -based method was then used for evaluating the ornament detection performance. The discrimination between notes and single-note ornaments was based on the duration of segments defined by the adjacent detected onsets. The F_0 information of the current and the following segment was used to distinguish between ‘cut’ and ‘strike’ single-note ornaments. The achieved F -measure performance for ‘cut’ was over 87%, while for ‘strike’ over 65%. The multi-note ornament detection system was based on analysing the properties of a sequence of detected segments. This included the sequential pattern of segmental F_0 ’s, the duration of each segment, and the relationship of the segmental F_0 ’s among the segments. The average F -measure performance over all types of multi-note ornaments was over 67%.

There are several points we are currently considering to extend this work. First, we plan to analyse the errors made by each of the onset detection methods and accordingly explore whether their combination could lead to detection performance improvements. This would also include exploration of the use of other onset detection methods, including other F_0 estimation algorithms and possible incorporation of the sinusoidal detection method we presented in [12]. Second, we will explore a compensation for variations in tempo across the recordings. Finally, we plan to employ probabilistic rules for detection of multi-note ornaments which should allow for better handling of the variations due to player’s style.

Acknowledgement

This work was supported by a project under the ‘Transforming Musicology’ programme funded by Arts and Humanities Research Council (UK).

6. REFERENCES

- [1] J. P. Bello, L. Daudet, S. Abdallah, Ch. Duxbury, M. Davies, and M. B. Sandler. A tutorial on onset detection in music signals. *IEEE Trans. on Speech and Audio Processing*, pages 1–13, 2005.
- [2] G. Boenn. Automated quantisation and transcription of musical ornaments from audio recordings. In *Proc. of the Int. Computer Music Conf. (ICMC)*, pages 236–239, Copenhagen, Denmark, Aug. 2007.
- [3] B. Breathnach. *Folk music and dances of Ireland*. Osian, London, 1996.
- [4] C. Carson. *Last night’s fun: in and out of time with Irish music*. North Point Press, New York, 1997.
- [5] M. Casey and T. Crawford. Automatic location and measurement of ornaments in audio recording. In *Proc. of the 5th Int. Conf. on Music Information Retrieval (ISMIR)*, pages 311–317, Barcelona, Spain, 2004.
- [6] N. Collins. Using a pitch detector for onset detection. In *Proc. of the 5th Int. Conf. on Music Information Retrieval (ISMIR)*, pages 100–106, Spain, 2005.
- [7] A. de Cheveigne and H. Kawahara. Yin, a fundamental frequency estimator for speech and music. *Journal of the Acoustical Society of America*, 111(4):1917–1930, April 2002.
- [8] S. Dixon. Onset detection revisited. In *Proc. of the 9th Int. Conf. on Digital Audio Effects (DAFx)*, pages 133–137, Montreal, Canada, Sep. 2006.
- [9] M. Gainza and E. Coyle. Automating ornamentation transcription. In *IEEE Int. Conf. on Acoustics, Speech and Signal Processing*, Honolulu, Hawaii, 2007.
- [10] M. Gainza, E. Coyle, and B. Lawlor. Single-note ornaments transcription for the Irish tin whistle based on onset detection. In *Proc. of the Digital Audio Effects (DAFX)*, Naples, Italy, 2004.
- [11] A. Holzapfel, Y. Stylianou, A. C. Gedik, and B. Bozkurt. Three dimensions of pitched instrument onset detection. *IEEE Trans. on Audio, Speech, and Language Processing*, 18(6):1517–1527, Aug. 2010.
- [12] P. Jančovič and M. Kökür. Detection of sinusoidal signals in noise by probabilistic modelling of the spectral magnitude shape and phase continuity. In *Proc. of the IEEE Int. Conf. on Acoustics, Speech, and Signal Processing (ICASSP)*, pages 517–520, Prague, Czech Republic, May 2011.
- [13] A. Klapuri. Sound onset detection by applying psychoacoustic knowledge. In *Proc. of the IEEE Int. Conf. on Acoustics, Speech, and Signal Processing (ICASSP)*, pages 3089–3092, March 1999.
- [14] M. Kökür, I. Ali-MacLachlan, P. Jančovič, and C. Athwal. Automated detection of single-note ornaments in Irish traditional flute playing. In *Int. Workshop on Folk Music Analysis*, Istanbul, Turkey, June 2014.
- [15] G. Larsen. *The Essential Guide to Irish Flute and Tin Whistle*. Mel Bay Publications, Pacific, Missouri, USA, 2003.
- [16] M. Puiggros, E. Gómez, R. Ramirez, X. Serra, and R. Bresin. Automatic characterization of ornamentation from bassoon recordings for expressive synthesis. In *9th Int. Conf. on Music Perception and Cognition*, Bologna, Italy, 2006.
- [17] E. D. Scheirer. Tempo and beat analysis of acoustic musical signals. *Journal of the Acoustical Society of America*, 103(1):588–601, 1998.