

2016-6

Note, Cut and Strike Detection for Traditional Irish Flute Recordings

Islah Ali-MacLachlan

Birmingham City University, islah.ali-maclachlan@bcu.ac.uk

Maciej Tomczak

Birmingham City University, maciej.tomczak@bcu.ac.uk

Carl Southall

Birmingham City University, carl.southall@bcu.ac.uk

See next page for additional authors

Follow this and additional works at: <https://arrow.tudublin.ie/fema>



Part of the [Musicology Commons](#)

Recommended Citation

Ali-MacLachlan, I., Tomczak, M., Southall, C., Hockman, J. (2016). Note, cut and strike detection for traditional Irish flute recordings. *6th International Workshop on Folk Music Analysis*, Dublin, 15-17 June, 2016.

This Conference Paper is brought to you for free and open access by the 6th International Workshop on Folk Music Analysis, 15-17 June, 2016 at ARROW@TU Dublin. It has been accepted for inclusion in Papers by an authorized administrator of ARROW@TU Dublin. For more information, please contact arrow.admin@tudublin.ie, aisling.coyne@tudublin.ie.



This work is licensed under a [Creative Commons Attribution-NonCommercial-Share Alike 4.0 License](#)

Authors

Islah Ali-MacLachlan, Maciej Tomczak, Carl Southall, and Jason Hockman

NOTE, CUT AND STRIKE DETECTION FOR TRADITIONAL IRISH FLUTE RECORDINGS

Islah Ali-MacLachlan, Maciej Tomczak, Carl Southall, Jason Hockman

DMT Lab, Birmingham City University

islah.ali-maclachlan, maciej.tomczak, carl.southall, jason.hockman
@bcu.ac.uk

ABSTRACT

This paper addresses the topic of note, cut and strike detection in Irish traditional music (ITM). In order to do this we first evaluate state of the art onset detection methods for identifying note boundaries. Our method utilises the results from manually and automatically segmented flute recordings. We then demonstrate how this information may be utilised for the detection of notes and single note articulations idiomatic of this genre for the purposes of player style identification. Results for manually annotated onsets achieve 86%, 70% and 74% accuracies for note, cut and strike classification respectively. Results for automatically segmented recordings are considerably lower therefore we perform an analysis of the onset detection results per event class to establish which musical patterns contain the most errors.

1. INTRODUCTION

1.1 Background

Irish Traditional Music (ITM) is a form of dance music played on a variety of traditional instruments including the flute. Within the tradition of ITM, players from different backgrounds are individuated based on their use of techniques such as ornamentation, a key factor alongside melodic and rhythmic variation, phrasing and articulation in determining individual player style (McCullough, 1977; Hast & Scott, 2004).

To automatically detect a player's style in audio signals, a critical first step is to detect these notes and ornamentation types. In this paper we evaluate both notes and single note ornaments known as *cuts* and *strikes*. Both ornaments generate a pitch deviation: a cut is performed by quickly lifting a finger from a tonehole then replacing it; a strike involves momentarily covering the tonehole below the note being played. We also analyse the cut and strike elements of multi-note ornaments known as *short roll* and *long roll*.

Figure 1 shows the pitch deviation for cuts and strikes. Long and short rolls are also displayed, showing the inclusion of cut and strike figures. A long roll occupies the same duration as three eighth notes whereas a short roll is equivalent to two eighth notes. In practice, ITM follows a *swing rhythm*—while there is a regular beat, swing follows an irregular rhythm and therefore each eighth-note section may not be of equal duration in normal playing (Schuller, 1991).

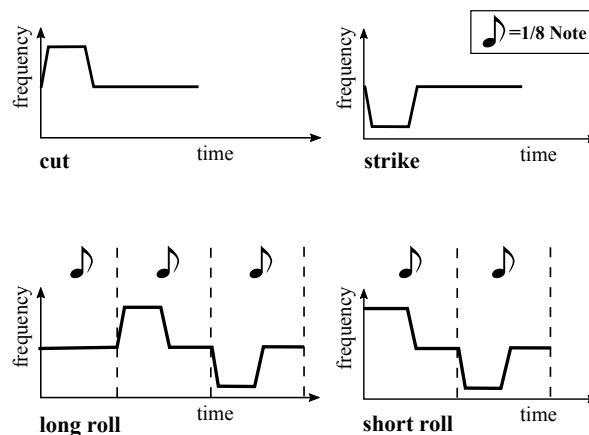


Figure 1: Frequency over time of *cut* and *strike* articulations showing change of pitch. *Long* and *short rolls* are also shown with pitch deviations (eighth note lengths shown for reference).

1.2 Related work

The approach undertaken in this paper utilises onset detection as a crucial first step in the identification of notes and ornaments. There are relatively few studies in the literature that deal specifically with onset detection within ITM, particularly with reference to the flute.

Onsets were found by Gainza et al. (2004) using band-specific thresholds in a technique similar to Scheirer (1998) and Klapuri (1999). A decision tree was used to determine note, cut or strike based on duration and pitch. Kelleher et al. (2005) used a similar system to analyse ornaments on the *fiddle* within Irish music, as bowed instruments also produce slow onsets.

Köküer et al. (2014) also analysed flute recordings through the incorporation of three kinds of information and a fundamental frequency estimation method using the YIN algorithm by De Cheveigné & Kawahara (2002). As in Gainza et al. (2004) a filterbank with fourteen bands optimised for the flute was used. More recently, Jančovič et al. (2015) presented a method for transcription of ITM flute recordings with ornamentation using hidden Markov models.

Unlike the above flute-specific methods, which rely on signal processing based onset detection algorithms, state of the art generalised onset detection methods use proba-

bilistic modelling. The number of onset detection methods using neural networks has substantially risen since Lacoste & Eck (2005). *OnsetDetector* by Eyben et al. (2010) uses bidirectional long short-term memory neural networks, and performs well in a range of onset detection tasks including solo wind instruments.

In this paper we perform an evaluation using several modern onset detection algorithms and a dataset comprised of 79 real flute performances of ITM. We then demonstrate how this information may be utilised towards the determination of notes and single note ornamentations.

The remainder of this paper is structured as follows: Section 2 details the method of segmentation, feature extraction and classification. In Section 3 we discuss evaluations of a range of onset detection methods and classification of notes, cuts and strikes. Results of the studies into onset detection and ornament classification are presented in Section 4 and finally conclusions and further work are discussed in Section 5.

2. METHOD

Figure 2 shows an overview of the proposed method. We extract features from audio segments representing events (notes, cuts, strikes) and propose an event type classification approach using the segmented event features.

For a fully automated method we use onset detection for segmentation. Event features are then extracted from inter-onset intervals (IOI). These features are used in a supervised learning algorithm to classify the segments as one of three distinct classes: notes, cuts and strikes. For onset detection, we attempt to use the top-performing algorithm from the evaluation presented in Section 3.2, and in the following discuss only the remaining feature extraction and classification stages.




Figure 2: Overview of the proposed classification method of notes, cuts and strikes in flute signals. The first phase shows feature extraction from segmented audio events and the second phase shows classification of the events.

2.1 Feature extraction

In order to capture the differences between each event type we extract features related to rhythm, timbre and pitch. An important distinction between notes, cuts and strikes is their duration, where notes are significantly longer than the two ornaments. To capture this we use the length ms of event segments. We then extract timbral features as

these are also important in class distinction. The change in timbre is caused by player’s fleeting finger motion as a tonehole is temporarily opened or closed. This results in a unique timbre that differs from notes. For that purpose we extract 13 Mel-frequency cepstral coefficients (MFCCs), excluding the first coefficient, and 12 chroma features features to accommodate for timbre and pitch changes in each of the articulations.

To extract features from the audio segments the input audio (mono WAV files) is down sampled to 11,025 Hz. Following the approach in Mauch & Dixon (2010) we calculate the MFCC and chroma features using a Hanning window of 1024 samples with 50% overlap. The extracted features are then normalised to the range [0,1] for every corresponding feature type. Each audio segment is assigned to its class Ω (e.g. note). An $n \times 26$ matrix F_{Ω} is created, where n represents the number of segments with 26 features (i.e., MFCCs, chroma, durations).

Each F_{Ω} segment appears in the context of musical patterns such as rolls, shakes or just consecutive notes in the recording. To account for the rhythmic, timbral and pitch changes of each event type in the context of these patterns, we concatenate the first derivatives of all features into every F_{Ω} segment.

2.2 Neural network classification

Audio segments are then classified into note, cut and strike classes using a feed-forward neural network.




Figure 3: Neural network architecture containing two hidden layers a^1 and a^2 , with weights w and biases b .

The proposed neural network, shown in Figure 3, consists of two hidden layers containing 20 neurons each. Back propagation is used to train the neural network, updating the weights and biases iteratively using scaled conjugate gradient of the output errors. A maximum iteration limit is set to 10,000 and the weights and biases are initialised with random non-zero values to ensure that training commenced correctly. A validation set is used to prevent over-fitting and cross entropy is used for the performance measure.

The output for each layer of an L layered neural network can be calculated using:

$$a^{(l)} = f_l(a^{(l-1)}(t)W^l + b^l), \quad (1)$$

where, a^l is the output at layer l and W and b are the weight and bias matrices. The transfer function is determined by the layer, as shown in Eq. 2.

$$f_l(x) = \begin{cases} 2/(1 + e^{-2x}) - 1, & l \neq L \\ y = e^x / (\sum e^x), & l = L. \end{cases} \quad (2)$$

Classification is performed by finding the index of the maximum value within the output from the neural network.

3. EVALUATION

As the performance of the proposed method depends heavily on the accuracy of the chosen onset detection method, the aim of our first evaluation is to determine the best performing onset detection algorithm. We then perform an evaluation of our note and ornament classification.

3.1 Dataset

For both these evaluations, we require a dataset that is representative of a range of respected players with individual stylistic traits. The dataset is comprised of 99 solo flute recordings of between 16 and 81 seconds in length, spanning over 50 years. For the purpose of this study, 79 recordings were selected excluding the excerpts from Larsen (2003), as they contain tutorial recordings not representative of typical ITM performances.

The recordings are 16-bit/44.1kHz WAV files all recorded by professional ITM flute players. Annotations were made using either Sonic Visualiser by Cannam et al. (2010) or Tony by Mauch et al. (2015). The annotation was performed by an experienced flute player. Full details of the annotation methods may be found in Köküer et al. (2014) and Ali-MacLachlan et al. (2015).

Annotations associated with this dataset include the temporal location of onsets and the event type (e.g., note, ornament). Additional classes such as breaths were included in the note class as they contained pitch information from a previous note. The annotated event types are represented by 15,310 notes, 2,244 cuts, and 672 strikes.

3.2 Onset detection evaluation

In this evaluation we measured how well eleven onset detection algorithms were capable of identifying onsets related to notes, cuts and strikes within real-life flute recordings. We reviewed the wind instrument class results from MIREX and examined various studies that concerned detection of soft onsets within these instruments.

Specialised methods for soft onset detection have been proposed in the literature. *SuperFlux* by Böck & Widmer (2013b) calculates the difference between two near short-time spectral magnitudes and is optimised for music signals with soft onsets and vibrato effect in string instruments. *ComplexFlux* by Böck & Widmer (2013a) is based on the *SuperFlux* algorithm with the addition of a local group delay measure that makes this method more robust against loudness variations of steady tones. Similarly, *Log-FiltSpecFlux* introduced in Böck et al. (2012) was designed to deal with onsets of various volume levels but was optimised for real-time scenarios.

In addition, there are several other onset detection methods proposed in the literature that we tested. The *Onset-Detector* by Eyben et al. (2010) processes the input signal both in the forward and backward manner and outputs

peaks that represent the probability of an onset at the detected position. The *Energy* (Masri, 1996), *Spectral Difference* (Foote & Uchihashi, 2001), *Spectral Flux* (Dixon, 2006) and *Kullback-Leibler* (KL) (Hainsworth & Macleod, 2003) represent detection functions solely based in the spectral domain. Brossier (2006) presented a modification to the KL algorithm shown as *Modified Kullback-Leibler* in our evaluation. The *Phase-based* method by Bello & Sandler (2003) looks at phase deviation irregularities in the phase spectrum of the signal. Lastly, the *Complex Domain* approach by Duxbury et al. (2003) combines both the energy and phase information for the production of a complex domain onset detection function. Peak-picking for the evaluate approaches is performed with *Madmom*¹ and *Aubio*² MIR toolboxes.

The onset detection results were calculated using the standard precision, recall and F-measure scores that measure performance of each onset detection algorithm. Precision and recall are determined from the detected flute onsets if reported within 25 ms on either side of the ground truth onset times. The mean F-measure is calculated by averaging F-measures across recordings.

3.3 Note and ornament classification evaluation

To assess the performance of our presented note and ornament classification method, we perform two evaluations using the dataset from Section 3.1. In the first evaluation, we attempt to determine the worth of the chosen classification method and selected features alone. In this experiment, we rely on the manually annotated note onsets to segment the audio prior to the feature extraction and classification stages. In the second evaluation, we seek to determine the viability of a fully automated ornament detection approach that relies on onset detection for segmentation. In this evaluation we employ the top performing onset detection algorithm found in the onset detection evaluation detailed in Section 3.2. For the training of the automated method only the true positive onsets will be used to ensure that the neural network is trained with the features corresponding to their correct classes.

To ensure an approximately equal proportion of training examples per class, we reduced the number of notes per recording to 6%, cuts to 30% and left in all strikes due to the proportion of these classes in the dataset. The classification evaluation is then performed using 5-fold cross validation.

4. RESULTS

4.1 Onset detection results

The results obtained from our experiment are shown in Table 1. The *OnsetDetector* method by Eyben et al. (2010) achieves the highest precision of 83% and F-measure of 78%. The high performance of this approach is in agreement with the results in the literature for the wind instrument class (Böck & Widmer, 2013a,b). While *Spectral*

¹ <https://github.com/CPJKU/madmom>

² <http://aubio.org/>

	P	R	F
OnsetDetector2015 Eyben et al. (2010)	0.8306	0.7510	0.7875
ComplexFlux2015 Böck & Widmer (2013a)	0.7414	0.6639	0.6996
SuperFlux2015 Böck & Widmer (2013b)	0.7659	0.6714	0.7144
LogFiltSpecFlux2015 Böck et al. (2012)	0.7597	0.6494	0.6989
Energy Masri (1996)	0.6870	0.5888	0.6270
Complex Domain Duxbury et al. (2003)	0.7548	0.6561	0.6999
Phase-based Bello & Sandler (2003)	0.7206	0.5522	0.6177
Spectral Difference Foote & Uchihashi (2001)	0.7087	0.5928	0.6416
Kullback-Leibler Hainsworth & Macleod (2003)	0.7926	0.4025	0.5265
Modified Kullback-Leibler Brossier (2006)	0.7659	0.1868	0.2890
Spectral Flux Dixon (2006)	0.5854	0.7618	0.6580

Table 1: Precision (P), Recall (R) and F-measure (F) for eleven onset detection methods. Maximum values for Precision, Recall and F-measure shown in bold.

class	notes	cuts	strikes
notes	86.97	8.43	8.54
cuts	6.79	70.46	16.96
strikes	6.24	21.07	74.49

Table 2: Confusion matrix for classification of notes, cuts and strikes using manually annotated onsets.

class	notes	cuts	strikes
notes	81.57	83.46	82.61
cuts	6.97	6.85	6.48
strikes	11.47	9.70	10.91

Table 3: Confusion matrix for classification of notes, cuts and strikes using a fully automated segmentation.

Flux achieved the highest recall score of 76% this is likely due to its overestimation of the onset positions thus resulting in a lower precision value. Consequently, in our note, cut and strike detection we use the onsets detected using the *OnsetDetector* as it outperforms other tested methods.

4.2 Note and ornament classification results

Table 2 presents a confusion matrix for note, cut and strike classification using features extracted from the annotated onset boundaries. The results demonstrate the effectiveness of the classification method for all three classes with 86% note, 70% cut and 74% strike detection accuracies. Misclassified notes are equally distributed across the other two classes demonstrating large timbral, pitch and rhythmic differences between note and ornament event types. The cuts and strikes are mostly misclassified as each other, which reflects their similar duration. These findings confirm the importance of duration in identifying the differ-

ence between ornaments (Gainza et al., 2004).

The results for a fully automated system evaluation are presented in Table 3. Here cuts and strikes were overwhelmingly misclassified as notes. These poor results are likely due to the imbalance between the number of annotated onsets and detected onsets. The evaluation using annotated onsets used in 916 notes, 670 cuts and 672 strikes, while the fully automated method used only 691 notes, 503 cuts and 518 strikes.

Training the system with features extracted from annotated segments and testing on automatically found segments did not improve on these results. To investigate the possible reasons for the poor classification results in Table 3, we conducted additional analysis of the onset detection results per event type.

4.3 Note, cut and strike onset detection accuracy

Cuts and strikes are components in multi-note ornaments such as rolls and shakes. To determine where onset detection errors occur we evaluate detection accuracy in relation to events that occurred immediately before and after the detected events. This evaluation allows us to see which event classes are most difficult to detect, and provide insight in the limitations of the real-life application of the proposed method for note, cut and strike detection.

Table 4 presents the onset detection results for each class of musical pattern. The classes consist of three event types where the central event is identified in bold. For example, *note cut note* is a detected cut with a note before and note afterwards, which exists within the event context of short and long roll or a single cut. The number of correctly detected onsets (true positives) is found as a percentage of the overall number of annotated onsets of that pattern.

As can be seen in Table 4, low accuracies were found for notes following ornaments. The largest error was found in the *cut note note*. This pattern exists only in the context of single cuts and shakes and occurred 1579 times with only 574 correctly found instances.

Our proposed note, cut and strike detection method depends on the features extracted from the found inter-onset intervals. The events corresponding to the cut and strike classes are detected with 83% and 82% accuracies respectively. Detecting notes that exist directly after these ornaments in the onset detection stage augments the content of the features describing the ornament event types. This results in training data that does not represent the classes that we intended to capture.

5. CONCLUSIONS AND FUTURE WORK

In this paper we present a note, cut and strike detection method for traditional Irish flute recordings. Our chosen approach to this problem is that of inter-onset segment classification using feed-forward neural networks. To evaluate the effectiveness of this approach we first conducted an evaluation of various onset detection algorithms on our dataset with the hope of using this method as a first step in the feature extraction.

Musical pattern			Event context	Accuracy	True positives	Total onsets
note	note	note	<i>single notes</i>	83.36	8651	10378
note	cut	note	<i>short & long rolls & single cuts</i>	83.44	1870	2241
note	note	cut	<i>notes before a roll</i>	84.16	1637	1945
cut	note	note	<i>notes after single cuts & shakes</i>	36.35	574	1579
note	strike	note	<i>short & long rolls & single strikes</i>	82.39	552	670
cut	note	strike	<i>short & long rolls</i>	34.62	180	520
strike	note	note	<i>last notes in rolls</i>	16.22	84	518
cut	note	cut	<i>shakes</i>	31.03	45	145
strike	note	cut	<i>last notes in rolls</i>	20.69	30	145
note	note	strike	<i>notes before single strikes</i>	90.85	129	142

Table 4: Onset detection results for each event class (bold) in the context of events happening before and after the detected onset. Accuracy shown as percentage of the accurately detected onsets (true positives) from that pattern

When using ground truth onset annotations, we achieved 86%, 70% and 74% accuracies for note, cut and strike classification respectively. When using detected onsets to train the neural network we achieved poor classification results. We then performed an analysis of the detected onsets and the context in which they appear to establish both the degree of the errors and the musical patterns in which they occur.

In the future we intend to work on improving the automated detection of note events. We will also develop note and ornament classification methods with additional features and other neural network architectures (e.g., recurrent neural networks, networks with long short-term memory) in order to capture trends that appear in time-series data. We also plan to investigate how well the proposed system generalises to other instruments that are characterised by soft onsets such as the tin whistle and fiddle.

6. ACKNOWLEDGEMENTS

This work was partly supported by the project Characterising Stylistic Interpretations through Automated Analysis of Ornamentation in Irish Traditional Music Recordings under the Transforming Musicology programme funded by the Arts and Humanities Research Council (UK).

7. REFERENCES

- Ali-MacLachlan, I., Kökier, M., Athwal, C., & Jančovič, P. (2015). Towards the identification of Irish traditional flute players from commercial recordings. In *Proceedings of the 5th International Workshop on Folk Music Analysis*, Paris, France.
- Bello, J. P. & Sandler, M. (2003). Phase-based note onset detection for music signals. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 5, (pp. 49–52), Hong Kong.
- Böck, S., Krebs, F., & Schedl, M. (2012). Evaluating the Online Capabilities of Onset Detection Methods. In *Proceedings of the 13th International Society for Music Information Retrieval Conference (ISMIR)*, (pp. 49–54), Porto, Portugal.
- Böck, S. & Widmer, G. (2013a). Local Group Delay Based Vibrato and Tremolo Suppression for Onset Detection. In *Proceedings of the 14th International Society for Music Information Retrieval Conference (ISMIR)*, (pp. 589–594), Curitiba, Brazil.
- Böck, S. & Widmer, G. (2013b). Maximum filter vibrato suppression for onset detection. In *Proceedings of the 16th International Conference on Digital Audio Effects (DAFx)*, (pp. 55–61), Maynooth, Ireland.
- Brossier, P. M. (2006). *Automatic annotation of musical audio for interactive applications*. PhD thesis, Queen Mary, University of London.
- Cannam, C., Landone, C., & Sandler, M. (2010). Sonic visualiser: An open source application for viewing, analysing, and annotating music audio files. In *Proceedings of the 18th international conference on Multimedia*, (pp. 1467–1468), Firenze, Italy. ACM.
- De Cheveigné, A. & Kawahara, H. (2002). YIN, a fundamental frequency estimator for speech and music. *The Journal of the Acoustical Society of America*, 111(4), 1917–1930.
- Dixon, S. (2006). Onset detection revisited. In *Proceedings of the 9th International Conference on Digital Audio Effects (DAFx)*, volume 120, (pp. 133–137), Montreal, Canada.
- Duxbury, C., Bello, J. P., Davies, M., & Sandler, M. (2003). Complex domain onset detection for musical signals. In *Proceedings of the 6th International Conference on Digital Audio Effects (DAFx)*, London, UK.
- Eyben, F., Böck, S., Schuller, B., & Graves, A. (2010). Universal Onset Detection with Bidirectional Long Short-Term Memory Neural Networks. In *Proceedings of the 11th International Society for Music Information Retrieval Conference (ISMIR)*, (pp. 589–594).
- Foote, J. & Uchihashi, S. (2001). The beat spectrum: A new approach to rhythm analysis. In *Proceedings of the 2nd IEEE International Conference on Multimedia and Expo (ICME 2001)*, (pp. 881–884), Tokyo, Japan. IEEE.
- Gainza, M., Coyle, E., & Lawlor, B. (2004). Single-note ornaments transcription for the Irish tin whistle based on onset detection. In *Proceedings of the 7th International Conference on Digital Audio Effects (DAFx)*, Naples, Italy.
- Hainsworth, S. & Macleod, M. (2003). Onset detection in musical audio signals. In *Proceedings of the 29th Inter-*

- national Computer Music Conference (ICMC)*, (pp. 163–166)., Singapore.
- Hast, D. E. & Scott, S. (2004). *Music in Ireland: Experiencing Music, Expressing Culture*. Oxford: Oxford University Press.
- Jančovič, P., Köküer, M., & Baptiste, W. (2015). Automatic transcription of ornamented irish traditional music using hidden markov models. In *Proceedings of the 16th ISMIR Conference*, (pp. 756–762)., Malaga, Spain.
- Kelleher, A., Fitzgerald, D., Gainza, M., Coyle, E., & Lawlor, B. (2005). Onset detection, music transcription and ornament detection for the traditional irish fiddle. In *Proceedings of the 118th Audio Engineering Society Convention (AES)*, Barcelona, Spain.
- Klapuri, A. (1999). Sound onset detection by applying psychoacoustic knowledge. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 6, (pp. 3089–3092)., Phoenix, Arizona.
- Köküer, M., Ali-MacLachlan, I., Jančovič, P., & Athwal, C. (2014). Automated Detection of Single-Note Ornaments in Irish Traditional flute Playing. In *Proceedings of the 4th International Workshop on Folk Music Analysis*, Istanbul, Turkey.
- Köküer, M., Kearney, D., Ali-MacLachlan, I., Jančovič, P., & Athwal, C. (2014). Towards the creation of digital library content to study aspects of style in Irish traditional music. In *Proceedings of the 1st International Workshop on Digital Libraries for Musicology*, London, U.K.
- Lacoste, A. & Eck, D. (2005). Onset detection with artificial neural networks for mirex 2005. *Extended abstract of the 1st Annual Music Information Retrieval Evaluation eXchange (MIREX)*, held in conjunction with ISMIR.
- Larsen, G. (2003). *The essential guide to Irish flute and tin whistle*. Pacific, Missouri, USA: Mel Bay Publications.
- Masri, P. (1996). *Computer modelling of sound for transformation and synthesis of musical signals*. PhD thesis, University of Bristol.
- Mauch, M., Cannam, C., Bittner, R., Fazekas, G., Salamon, J., Dai, J., Bello, J., & Dixon, S. (2015). Computer-aided Melody Note Transcription Using the Tony Software: Accuracy and Efficiency. In *Proceedings of the 1st International Conference on Technologies for Music Notation and Representation (TENOR)*, (pp. 23–30)., Paris, France.
- Mauch, M. & Dixon, S. (2010). Simultaneous estimation of chords and musical context from audio. *Audio, Speech, and Language Processing, IEEE Transactions on*, 18(6), 1280–1289.
- McCullough, L. E. (1977). Style in traditional Irish music. *Ethnomusicology*, 21(1), 85–97.
- Scheirer, E. D. (1998). Tempo and beat analysis of acoustic musical signals. *The Journal of the Acoustical Society of America*, 103, 588.
- Schuller, G. (1991). *The Swing Era: The Development of Jazz, 1930-1945*. Oxford Paperbacks.