

MACHINE ANNOTATION OF SETS OF TRADITIONAL IRISH DANCE TUNES

Bryan Duggan, Brendan O’ Shea

Mikel Gainza

Pádraig Cunningham

DIT School of Computing
Kevin St
Dublin 8
Ireland
{bryan.duggan,
brendan.oshea}@comp.dit.ie

Audio Research Group
DIT Kevin St
Dublin 8
Ireland
mikel.gainza@dit.ie

School of Informatics and
Computer Science
UCD
Dublin
Ireland
padraig.cunningham@ucd.ie

ABSTRACT

A *set* in traditional Irish music is a sequence of two or more dance tunes in the same time signature, where each tune is repeated an arbitrary number of times. A *turn* in a set represents the point at which either a tune repeats or a new tune is introduced. Tunes in sets are played in a *segue* (without a pause) and so detecting the turn is a significant challenge. This paper presents the MATS algorithm, a novel algorithm for identifying turns in sets of traditional Irish music. MATS works on digitised audio files of monophonic flute and tin-whistle music. Previous work on machine annotation of traditional music is summarised and experimental results validating the MATS algorithm are presented.

1. INTRODUCTION

Several papers address the necessity of developing MIR (Music Information Retrieval) systems that are adapted to the specific requirements of ethnic music and also to the needs of musicologists studying ethnic music [1-3]. While there are MIR systems that allow users to search for traditional Irish dance tunes using text based musical queries [4,5] and there are MIR systems that allow users to search for melodies using sung queries [6,7], there are no MIR systems that we are aware of that allow musicians to search for traditional Irish dance tunes using queries played on traditional instruments. Some examples of the above include the website thesession.org [4] which contains an extensive collection of over seven thousand traditional dance tunes in the ABC language; the system supports text queries by any of the metadata associated with a tune or melodic queries in the ABC language. Similarly, Melodyhound [6] a publicly accessible MIR system that supports sung queries and contains a large collection of traditional Irish dance tunes does not generate positive results when queries are presented in the form of melodies played on the tin-whistle or wooden flute.

Such a system would have many applications in the field of music archiving and retrieval, particularly given the many thousands of hours of archive music collected by organisations involved in the cataloguing of traditional music such as Na Píobairí Uilleann, Comhaltas Ceoltóirí Éireann and the Irish Traditional Music Archive. Similarly it is common at traditional music sessions, recitals and

even on commercial recordings for tunes to be named *gan ainm* (without name) when the tune in question does in fact have a name, composer and history. For a typical example see the CD recording [8].

Previous work proposes MATT2 (Machine Annotation of Traditional Tunes) as a system that can identify tunes played on either the flute or the tin whistle [9]. MATT2 takes advantage of a number of novel subsystems that significantly increase matching accuracy for traditional tunes played in a variety of regional styles by different musicians. These include an onset detection function developed for windblown instruments, an ornamentation compensation algorithm based on fuzzy histograms, a two thousand tune corpus of tunes in the ABC language (a natural fit for traditional music) and a melody normalisation algorithm that adapts tunes in the corpus to the way they might be played by a human musician. MATT2 is described in detail in [9] and we present an overview in section 3. The main purpose of this paper is to present our enhancements to the MATT2 system and specifically to present a new algorithm for annotating sets of traditional Irish dance tunes. Previous versions of MATT2 could only annotate single tunes, however in traditional music tunes are rarely played singly. More commonly tunes are played in groups of at least two tunes known as a *set* of tunes. A set typically consists of two three or four tunes played in succession without an interval [10,9]. Typically each tune in the set is played twice or three times before musicians advance to the subsequent tune in the set. A repetition or a change from one tune to the next in a set is known as a *turn*. As tunes in sets are always in the same time signature and often in the same key, the challenge therefore is in segmenting sets into tunes and repetitions. The approach presented in this paper tackles this problem by making use of melodic similarity calculated using a variant of the *edit distance* string matching algorithm described in section 3. The MATS algorithm described in this paper can identify the start and end of each repetition of a tune, can count the repetitions and can identify the title and associated metadata associated with each tune in a set.

Section 2 of this paper briefly explains the domain of traditional Irish dance music. In Section 3 existing work on the MATT2 system is presented. Section 4 presents MATS (Machine Annotation of Traditional Sets), a novel

annotation algorithm which annotates sets of traditional tunes. Section 5 presents experimental results which establish the effectiveness of this new algorithm and section 6 presents conclusions and future work.

2. TRADITIONAL IRISH DANCE MUSIC

The most common forms of dance music are *reels*, *double jigs* and *hornpipes*. Other tune types include *marches*, *set dances*, *polkas*, *mazurkas*, *slip jigs*, *single jigs and reels*, *flings*, *highlands*, *scottisches*, *barn dances*, *strathspeys* and *waltzes* [11]. These forms differ in time signature, tempo and structure. For example a reel is generally played at a lively tempo and is in 4/4 time (four crochets in a bar, though usually transcribed as eight quavers in a bar), while a waltz is generally played at slower pace and is in 3/4 time. Most tunes consist of a common structure of two parts called the *A* part and *B* part. Tunes are typically played as *sets*. Certain common sets were originally put together to accompany set dances [10], while other sets have become popular as a result of being recorded by emigrant Irish musicians in America in the early part of the twentieth century.

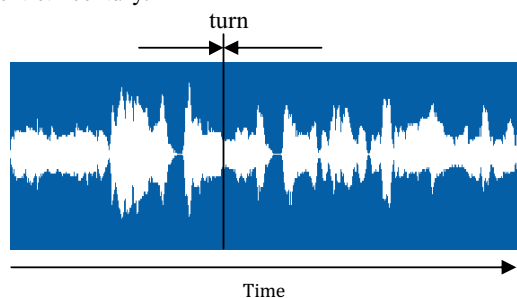


Figure 1: Waveform of the last phrase from the tune "Jim Coleman's" and the first phrase from the tune "George Whites Favourite" played in a set

The origin of many sets of tunes is unknown and musicians often compile new sets "on the fly" in traditional music sessions. Figure 1 shows a waveform plot from two tunes played in a set. The tunes were played on a wooden flute and as can be seen in the plot, there is no interval between the end of the first tune and the start of the second tune. Maddage *et al.* and other segmentation approaches generally look for repetitive patterns in a music recording [12]. This is not the case in our approach, where each tune in the set can be played once or many times.

When a traditional musician plays a tune, it is rarely played exactly as transcribed. In fact an experienced musician never plays the same tune twice identically, employing the subtleties of *ornamentation* and *variation* to interpret the tune [11]. For a discussion on the use of ornamentation in traditional music we refer to [11,13,14].

Ornamentation plays a key role in the individual interpretation of traditional Irish music [10]. The usage of

ornamentation is highly personal and large variations exist in the employment of ornamentation from region to region, instrument to instrument and from musician to musician. Tansey colourfully describes ornamentation in the following way:

"I put it to you therefore that it had to come from the throats of birds, the wild animals, the ancient chants of our forefathers, the hum of the bees and the mighty rhythms of the galloping hooves of wild horses all moulded together..." [15]

Ornamentation is difficult to detect correctly and state of the art ornamentation detection algorithms report a success rate of just 40% for multi-note ornaments [16,17]. Similarly, related work in classical music suggests that the playing of ornamentation (grace notes) requires adaptation of melodic similarity measures [18].

It is clear from this brief introduction that an MIR system for traditional dance music must therefore deal with many special problems, such as stylistic variation even within the same instance of a tune, the use of ornamentation which can skew melodic similarity measures and the collection of tunes into sets creating segmentation problems. Transposition invariance is not a requirement for MIR in traditional music as it is uncommon for tunes to be transposed into different keys [16].

3. MACHINE ANNOTATION OF TRADITIONAL TUNES (MATT2)

MATT2 works on mono, digital audio files in the WAV format recorded at 44KHz. A high level diagram of the subsystems that make up MATT2 are presented in Figure 2. MATT2 is described in detail in [9] and so a brief description is presented here.

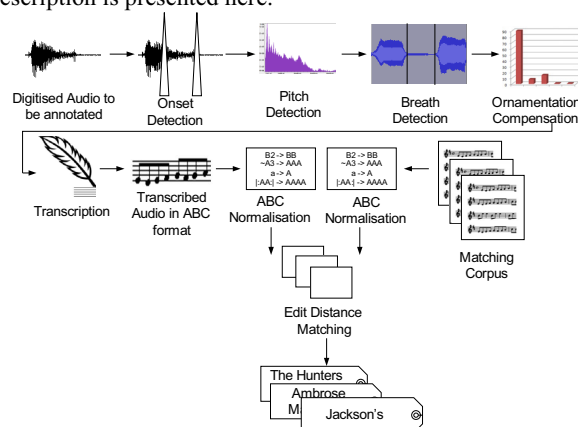


Figure 2: High level diagram of the MATT2 tune annotator

The audio file to be annotated is first segmented into candidate note onsets using an onset detection function

adapted from Gainza [16,17]. The onset detection function ODCF is based on time domain FIR (Finite Impulse Response) comb filters. ODCF discovers harmonic characteristics of the input signal and is therefore useful for detecting onsets in *legato* playing typical of windblown traditional instruments such as the flute and the tin whistle.

In order to detect the perceived pitch of a frame, the pitch detection sub-system performs a STFT (Short Term Fast Fourier Transform) on segments bounded by onsets detected by the onset detection system. The algorithm then calculates the pitch as being the interval between the two most prominent peaks in the FFT graph. This simple approach works well for the harmonics of the wooden flute and the tin whistle.

MATT2 incorporates a breath detector subsystem to transcribe a breath in the signal. A breath is marked if either the pitch detected by the pitch detector is less than 100Hz or the average amplitude of a candidate note cn is less than a 10% threshold th of the average amplitude of the entire signal s . Breaths detected before the transcription of the first pitched note and at the end of the transcription are ignored by the system.

MATT2 uses a heuristic to determine if the input signal was generated by a tin whistle or a wooden flute. A tin whistle in the key of D is pitched exactly one octave above a flute in the key of D, so if the algorithm counts more notes with a pitch above G5 (783.99hz) than below G5, then the algorithm concludes that the input signal contains a tin whistle and the pitches in the pitch spelling algorithm are shifted up accordingly.

Both the wooden flute and the tin whistle have a range of two octaves, though this can be extended by cross fingering techniques [11,19,13]. To tag each candidate note cn with a pitch spelling $pS(cn)$, each calculated note frequency is compared with the frequencies of the notes in the key of D4 Major and D5 Major $k_1... k_{16}$ the two octaves playable on a wooden flute.

The system eliminates notes whose durations are close to zero by merging their durations with subsequent notes. This has the effect of eliminating consecutive onsets (false positives in the ODF caused by noisy onsets) and also eliminating ornamentation notes such as those found in *rolls*, *cuts* *taps* and *crans* typical of traditional Irish music [11,20,19,13,15]. To achieve this, the quantisation subsystem first generates a histogram of all the note durations. The histogram bin with the highest value is considered to be the length of a quaver note. The histogram counts notes within +/-30% of the bin width. The algorithm also updates the bin width each time a candidate is counted, so that the bin widths represent the cumulative average lengths of notes counted. A transcription t is then generated in the ABC language of the input signal from the features extracted by the subsystems in MATT2.

MATT2 has a corpus Z of two thousand known tunes (and variations) in the ABC language drawn from the

collections of Norbeck [21]. To identify a tune, MATT2 firstly normalises both the transcription t and each string $c \in Z$. This process is described in detail in [9]. Normalisation minimises the effect of transcription errors and stylistic variation on the calculation of melodic similarity. The *edit distance* is then calculated for t in every $c \in Z$ and the tune with the lowest edit distance is returned as a match.

Edit distance, also known as *Levenshtein distance* or *evolutionary distance* [22,23], is a concept from information retrieval and it describes the number of edits (insertions, deletions and substitutions) that have to be made in order to change one string to another. It is the most common measure to expose the similarity between 2 strings.

The edit distance $ed(x, y)$ between strings $x=x_1 \dots x_m$ and $y=y_1 \dots y_n$, where $x, y \in \Sigma^*$ is the minimum cost of a sequence of editing steps required to convert x into y . Σ is the alphabet of possible characters and Σ^* is the set of all possible sequences of $ch \in \Sigma$. Edit distance can be calculated using dynamic programming [23]. To compute the edit distance $ed(x,y)$ a matrix $M_{1..m+1,1..n+1}$ is constructed where $M_{i,j}$ is the minimum number of edit operations needed to match $x_{1..i}$ to $y_{1..j}$. Each matrix element $M_{i,j}$ is calculated as per (1). The minimum edit distance between x and y is given by the matrix entry at position $M_{m+1,n+1}$.

$$M_{i,1} \leftarrow i - 1, M_{1,j} \leftarrow j - 1 \quad \text{if } x_i=y_j$$

$$M_{i,j} \leftarrow \begin{cases} M_{i-1,j-1} & \text{else} \\ 1 + \min(M_{i-1,j-1}, M_{i-1,j}, M_{i,j-1}) \end{cases} \quad (1)$$

The algorithm can be adapted to find the lowest edit distances for x in substrings of y . This is achieved by setting $M_{1,j} = 0$ for all $j \in 1..n+1$. In contrast to the edit distance algorithm described above, the last row $M_{m+1,j}$ is then used to give a *sliding window* edit distance for x in substrings of y [23].

	D	G	G	G	D	G	B	D	E	F	G	A	B
B	0	0	0	0	0	0	0	0	0	0	0	0	0
D	1	1	1	1	1	1	0	1	1	1	1	1	0
E	2	1	2	2	2	1	2	1	0	1	2	2	1
E	3	2	2	3	3	2	2	2	1	0	1	2	3
E	4	3	3	3	4	3	3	3	2	1	1	2	3

Table 1: Edit distance for the string BDEE in DGGGDGBDEFGAB. This string represents the first 13 notes from the tune "Jim Coleman's" in normalised ABC format

An example of this variation on the edit distance applied to search for the pattern "BDEE" in "DGGGDGBDEFGAB" is given in Table 1. The

minimum edit distance positions are highlighted.

Variations on the edit distance algorithm have been applied in domains such as DNA analysis and automated spell checking and are commonly used in MIR systems [7,24].

With test input drawn from the playing of ten different musicians playing flute, whistle and fiddle, the system was able to correctly identify the tune in 86% of cases. In 96% of cases, the correct tune was identified within the top five closest matches [9].

4. MACHINE ANNOTATION OF TRADITIONAL SETS ALGORITHM (MATS)

In this section MATS is described. MATS is an enhancement to MATT2 described in the previous section. The purpose of MATS is to annotate tunes played in sets.

The shortest tune in the corpus Z used by MATT2 is a single jig. A single jig sj is a tune in 6/8 time with an A and B part played singly (48 quaver notes in duration). The shortest possible set therefore would contain two single jigs (96 notes) played with no repetitions. To annotate a set of tunes, MATS first uses a heuristic to determine if the string of transcribed notes t is longer than the length of the shortest set $length(sj) \times 2$.

When this is the case, the MATS algorithm is used instead of the minimum edit distance algorithm described in section 3. Pseudocode for the MATS algorithm is presented in Figure 3.

MATS first extracts a substring ss from t the transcription such that $length(ss) = length(sj)$ at position $p=1$ in t . MATS then searches the corpus Z using the edit distance algorithm described in section 3 to find a the closest match for ss . When a match is found MATS knows the name of the first tune and has c' , a transcription of the tune played with no repetitions from the corpus Z . MATS then generates an edit distance profile edp for c' , the matching tune, in t the transcription. edp is given as the last row of the edit distance matrix and can be understood as the positions where substrings in t match c' with the minimum edit distance.

Figure 4 shows the edit distance profiles for the set of tunes “Jim Coleman’s”, “George Whites Favourite” and “the Virginia” played in a set. The algorithm has identified the first tune as “Jim Coleman’s” and has subsequently generated an edit distance profile (the top plot in Figure 4) for the first tune in the transcription. The two troughs in this graph indicate the end of the two repetitions of the tune in the transcription. These can be considered as turns in the set.

The MATS algorithm then normalises the edit distance profile edp and passes the graph through a low pass filter that filters frequencies less than 10Hz. This has the effect of smoothing the graph. An example of a smoothed edit distance profile is given in Figure 5. This graph illustrates the top graph in Figure 4 after filtering has been applied.

The algorithm then detects troughs in the graph less than a threshold initially set to $t=0.3$. The algorithm varies this threshold dynamically by trying different values until the number of troughs in the graph is between one and five. It is rare in traditional music for a tune to be repeated more than five times in a set.

```

p ← 0
rem ← length(t) - p
while (rem >= sj)
begin
  ss ← substring(t, p, p + sj)
  foreach (c in Z)
begin
  ed_c ← min(ed(ss, c))
  if (ed_c < min_ed)
begin
  min_ed ← ed_c
  c' ← c
end
end
end
edp ← ed(c', t)
edp ← normalise(edp)
edp ← filter(edp, 10)
th ← 0.3
v ← troughs(edp, th)
foreach (tr in v)
begin
  convertToTime(tr)
end
r ← length(v)
p ← v[r]
print c', r
rem ← length(t) - p
end

```

Figure 3: Pseudocode for the MATS set annotation algorithm

The trough detection algorithm in MATS returns a vector of troughs \vec{v} , such that $length(\vec{v})$ is the number of troughs and the elements in \vec{v} are the positions of the bottom of the troughs. A trough in MATS need only have a descending wall as a trough can occur at the end of a tune and hence may not contain an ascending wall. An example of this is the third plot in Figure 4.

The algorithm repeats this process with a new p given as the last entry in the troughs vector to extract the second and subsequent tunes in the set until it is no longer possible to extract a substring ss of length $length(sj)$ starting at p because we have reached the end of t . The second tune in the set, “George Whites Favourite” was played once and there is a corresponding single trough in the graph of the edit distance function (the middle plot in Figure 4) for the tune from the corpus c' in the transcription t . The third tune “the Virginia” was repeated twice and so there are two troughs in the bottom plot in Figure 4.

5. RESULTS

In order to test the robustness of MATS we had a traditional musician record ten audio files of flute tunes played in sets. The recorded files are available at <http://www.comp.dit.ie/bduggan/mats>. The sets played in the input audio were taken from the Foinn Seisiún series of books published by Comhaltas Ceoltóirí Éireann [25].

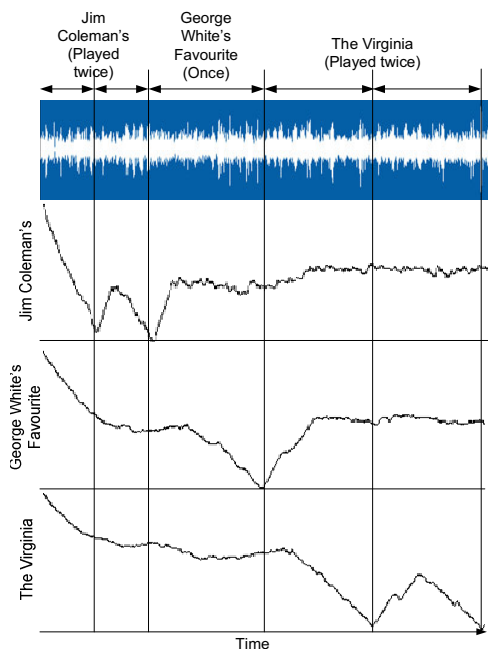


Figure 4: Edit distance profiles for three tunes played in a set

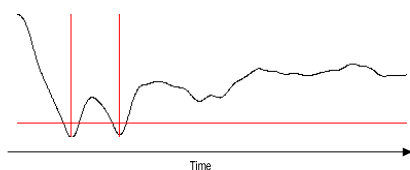


Figure 5: Filtered version of first graph in Figure 4. The dynamic threshold and detected troughs are marked

The sets consisted of single and double jigs and reels played multiple times in sets. In total, the sets contained 23 separate tunes with 48 turns we were interested in annotating. In carrying out this experiment, we were interested in establishing if MATT2 could correctly figure out the timings of turns and could identify the names of the tunes.

Correctly identified	96%
Incorrectly identified	4%

Table 2: Correctly and incorrectly identified tunes

MATT2 successfully identified 22 out of the 23 tunes, and recognised each input audio file as a set and so used the MATS set annotation algorithm (Table 2).

Table 3 shows a sample of the data collected in this experiment for the audio file used to generate Figure 4 and Figure 5. To establish a ground truth for the experiment, a human domain expert manually annotated the turns in the sets of tunes. In the human and machine columns are listed the onset time for turns in the set. Onset times for changes from one tune to the next are highlighted. From this table it can be seen that on average MATS was within .85 seconds of the human annotations.

Tune	Human	Machine	Difference
1	20.68	21.10	0.43
1	41.42	41.9	0.48
2	82.72	83.15	0.43
3	123.88	124.44	0.56
3	164.49	166.85	2.36
Average			0.85

Table 3: Human & machine annotated turns

The overall annotation accuracy is obtained by calculating two different measures *precision* and *recall*. The value of *precision* is calculated as per (2) where *TP* and *FP* are the true positives (correctly identified turns) and false positives (incorrectly identified turns). *recall* is calculated as per equation (3) where *FN* is the number of false negatives (turns in the input signal not detected by the algorithm).

$$precision = \frac{TP}{TP + FP} \tag{2}$$

$$recall = \frac{TP}{TP + FN} \tag{3}$$

TP	FN	FP	precision(%)	recall(%)
39	9	6	87%	81%

Table 4: Annotation accuracy

Table 4 shows the annotation accuracy. In can be seen from *precision* and *recall* that the algorithm provides a high degree of accuracy at detecting turns. Because the algorithm can successfully identify turns, it can also correctly extract a suitable prefix from the subsequent tune in the set and so can identify the tune. *FN*'s were caused by the algorithm failing to correctly identify the transitions between tunes in a set. When this happens the algorithm cannot extract a representative prefix from the next tune and so all subsequent turns are usually misidentified. In

some cases, *FP*'s were within a few seconds of the two second threshold we had set.

6. CONCLUSIONS & FUTURE WORK

This paper presented a novel algorithm that addresses a problem in the domain of Irish traditional dance music, that of annotating sets of tunes. As a set can contain an arbitrary number of tunes played segue without an interval and as tunes in sets are repeated an arbitrary number of times, are always in the same time signature and often in the same key, the significant challenge in this problem is in recognising where one tune ends and the next tune starts. The results presented prove that MATS is effective at segmenting sets, counting repetitions and at annotating individual tunes played in a set. To our knowledge this is the first time this problem has been addressed in an MIR system and we suggest that the proposed approach can be adapted to segmenting repeated tunes from other genres played in a segue.

The corpus used currently contains reels and jigs and in future work it will be augmented with the full complement of traditional tunes in different time signatures. One interesting feature not yet exploited is the metadata typically present in an ABC transcription. Effectively the time and key signature of an input audio file can be determined by *melodic similarity* with a known tune. This can be exploited in several interesting ways. Firstly, if the first tune in a set were to be identified as a reel, the search for subsequent tunes can be limited to reels, thus speeding up annotation. Conversely, if a number of reels were to be identified in a set and a single tune in a different time signature was to be identified this could be recognised as a potential error.

7. REFERENCES

- [1] S. Doraisamy, H. Adnan, and N. Norowi, "Towards a MIR System for Malaysian Music," *7th International Conference on Music Information Retrieval, Victoria, Canada, 8 - 12 October 2006*.
- [2] K. Jensen, J. Xu, and M. Zachariassen, "Rhythm-based segmentation of popular chinese music," *Proceedings of 6th International Conference on Music Information Retrieval (ISMIR'05)*, pp. 374-380.
- [3] A. Nesbit, L. Hollenberg, and A. Senyard, "Towards Automatic Transcription of Australian Aboriginal Music," *5th International Conference on Musical Information Retrieval, Barcelona, Spain October 10-14, 2004*.
- [4] "The Session"; <http://www.thesession.org/>.
- [5] Bryan Duggan, "Learning Traditional Irish Music using a PDA," *IADIS Mobile Learning Conference, Trinity College Dublin*, Jul. 2006.
- [6] "Melodyhound"; http://www.melodyhound.com/query_by_humming.0.html.
- [7] W. Birmingham et al., "Musart: Music Retrieval Via Aural Queries," *Ann Arbor*, vol. 1001, pp. 48109-2110.
- [8] N. Sweeney, "The Whinny Hills of Leitrum," 2008.
- [9] B. Duggan, B. O'Shea, and P. Cunningham, "A System for Automatically Annotating Traditional Irish Music Field Recordings," *Sixth International Workshop on Content-Based Multimedia Indexing, Queen Mary University of London, UK, Jun. 2008*.
- [10] F. Valley, *The Companion to Irish Traditional Music*, New York University Press, 1999.
- [11] G. Larson, *The Essential Guide to Irish Flute and Tin Whistle*, Mel Bay Publications, Inc., 2003.
- [12] N. Maddage et al., "Content-based music structure analysis with applications to music semantics understanding," *Proceedings of the 12th annual ACM international conference on Multimedia*, 2004, pp. 112-119.
- [13] F. Valley, *Timbre: The Wooden Flute Tutor*.
- [14] B. Duggan, Z. Cui, and P. Cunningham, "MATT - A System for Modelling Creativity in Traditional Irish Flute Playing," *Third ECAI Workshop on Computational Creativity*, Riva Del Garda, Italy: 2006.
- [15] S. Tansey, *The Bardic Apostles of Innisfree*, Tanbar Publications, 1999.
- [16] M. Gainza, *Music Transcription within Irish Traditional Music*, PhD Thesis, Dublin Institute of Technology, Faculty of Engineering, 2006.
- [17] M. Gainza and E. Coyle, "Automating Ornamentation Transcription," *Acoustics, Speech and Signal Processing, 2007. ICASSP 2007. IEEE International Conference on*, vol. 1, 2007.
- [18] R. Typke, "Music Retrieval Based on Melodic Similarity," *Doctoral thesis, Utrecht University*, 2007.
- [19] C. Hamilton, *The Irish Flute Players Handbook*, Cork: Breac Publications, 1990.
- [20] N. Keegan, *The Words of Traditional Flute Style*, MPhil Thesis, University College Cork, Music Department, 1992.
- [21] H. Norbeck, "ABC Tunes," 2007; <http://www.norbeck.nu/abc/index.html>,
- [22] V. Levenshtein, "Binary Codes Capable of Correcting Deletions, Insertions and Reversals," *Soviet Physics Doklady*, vol. 10, 1966, p. 707.
- [23] G. Navarro and M. Raffinot, *Flexible Pattern Matching in Strings: Practical On-Line Search Algorithms for Texts and Biological Sequences*, Cambridge University Press, 2002.
- [24] K. Lemstrom and S. Perttu, "SEMEX-An Efficient Music Retrieval Prototype," *First International Symposium on Music Information Retrieval (ISMIR)*, 2000.
- [25] P. Brian, *Foinn Seisiún*, Comhaltas Ceoltóirí Éireann, 2004.