MATT - A System for Modelling Creativity in Traditional Irish Flute Playing

Bryan Duggan

School of Computing,
Dublin Institute of Technology,
Kevin St., Dublin 8,
Ireland.

bryan.duggan@comp.dit.ie
http://www.comp.dit.ie/bduggan

Zheng Cui

School of Computing,
Dublin Institute of Technology,
Kevin St., Dublin 8,
Ireland.

zheng.cui@student.dit.ie

Pádraig Cunningham

Machine Learning Group,
Department of Computer Science,
Trinity College,
Dublin 2, Ireland.

padraig.cunningham@cs.tcd.ie

Abstract

This paper outlines ongoing work in the development of MATT (Machine Learning for Articulating Traditional Tunes). MATT uses a combination of case based reasoning and wave table synthesis to simulate the creative interpretation of traditional Irish tunes on the wooden flute. The paper presents a brief overview of the problem being addressed and describes the approach being undertaken. Some preliminary results are also presented.

1 Introduction

MATT (Machine learning for Articulating Traditional Tunes) is a system that seeks to model the creative interpretation of traditional Irish music by master flute players. As discussed in similar research, we do not propose that it is possible to entirely reproduce what musicians do when performing a piece of music. However, musical performance involves tacit knowledge about interpretation that humans acquire by observation and imitation [De Mantaras & Arcos, 2002]. This system examines the possibility of simulating the observation and imitation process in order to acquire the knowledge to interpret a piece of traditional Irish music in the style of master flute players. MATT uses a number of techniques to achieve this including case based reasoning (CBR) and wave table synthesis. CBR is a technique used in artificial intelligence to solve new problems by using or adapting solutions to existing problems [Smith & Medin, 1981]. A new problem is solved by retrieving one or more previously experienced cases, reusing the case, revising the solution, and retaining the new experience by incorporating it into the existing case-base. CBR is derived from theories of concept formation, problem solving and experiential learning in humans [Aamodt & Plaza, 1994]. The CBML Framework developed by Coyle *et al.* is used as a toolset for the CBR engine at the heart of MATT [Coyle, 2004].

In order to generate output that not only makes appropriate creative decisions, but also sounds realistic, a MIDI triggered wavetable VST (Virtual Studio Technology) is under development which models the sonic properties of a wooden flute [Steinberg, 2004]. The output of the CBR engine forms the input to the VST using an augmented MIDI protocol to render the final output of the system. The VST is built using approximately 1GB of samples generated from actual flute recordings. This paper mainly focuses on the work carried out in developing the CBR engine.

Section 2 of the paper explores ideas of creativity in traditional flute playing and elaborates on techniques used by flute players to add individual expression to the interpretation of tunes. It describes typical fingered articulations used by flute players. Section 3 outlines the approach being developed for MATT and describes corpus acquisition, case base building and target case retrieval algorithms. Section 4 presents some preliminary test results. Section 5 presents conclusions and future work.

2 Creativity in Traditional Irish Music

This project is primarily concerned with traditional dance music, as played on the wooden flute. 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* [Larson, 2003]. 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 (4 crochets in a bar, though usually transcribed as 8 quavers in a bar) while a waltz is generally played at slower pace and is in 3/4 time. The time signature, tempo and structure of a tune form are determined by the dance it accompanies. Most tunes consist of a common structure of two parts called either the *first* and *second* part or the *A* part and *B* part. Tunes are typically arranged into *sets*. A set consists of a number of tunes (commonly two or

three) played sequentially. Each tune in a set is usually repeated two or three times [Vallely, 1999].

The flute came into common use in traditional music in the 19th Centaury. The "Irish flute" is also known as the *concert flute* (because it is in concert pitch), the *timbre flute* (because it is made from wood), the *simple system flute* or the *fheadóg mhór* (big whistle). It has six holes tuned such that the lowest playable pitch (all holes closed) is the D above middle C, and the instrument will play a D scale (D, E, F#, G, A, B, C#) as the holes are uncovered sequentially to shorten the resonant length of the bore. The basic flute is often augmented with the addition of up to eight keys (typically made from silver, mounted on wooden blocks) used to play pitches which are impossible to produce on the basic flute. Figure 1 depicts a 6 keyed wooden flute made by Eamonn Cotter, an unkeyed flute made by Eamonn Cotter and an unkeyed bamboo flute made by Patrick Olwell in the key of F.

Figure 1: Wooden flutes

Music is a creative art form and "individual expression" (*p-creativity*) is a *defining component* of traditional Irish music [Breathnach, 1977, Boden, 1996]. Individual creativity in traditional music takes three forms:

1. The composition of new tunes.
2. The arrangement of tunes.
3. The individual creativity of a musician in interpreting a tune.

This project concentrates on individual, interpretive creativity. When a traditional musician plays a tune, it is rarely played as transcribed, though unlike with jazz for example, traditional musicians never deviate from the structure or framework of the tune. In fact, experienced musicians rarely play the same tune twice, identically. Instead, a musician will employ the subtleties of *ornamentation* and *variation* to interpret the tune [Larson 2003].

Larson [2003] defines ornamentation as "*ways of altering or embellishing small pieces or cells of a melody that are between one and three eight-note beats long. These alterations and embellishments are created mainly through the use of special fingered articulations.*" Fingered articulations are a defining characteristic of traditional Irish music. The sound of most articulations is very brief. Although generated by inserting additional notes, the notes are played at such speed that they are not perceived as having a discernible pitch or duration. Fingered articulation techniques used in traditional music are summarised in Table 1.

Breath techniques also serve to articulate the tune. Additionally musicians will often play variations on the written melody, in particular if playing a tune several times in a set. Fingered articulation techniques combine with breath techniques and variations to give a player a variety of tools for musical expression.

Two musicians are involved in this project. Eamonn Cotter describes his flute playing as "intricate and detailed" [Cotter, 2006]. His playing is extensively ornamented and he makes full use of the keyed flute as he plays tunes in unusual keys and often inserts *accidentals* for effect. His playing has been compared to the great East Galway flute playing of Paddy Carty (who also used keys considerably) [Hurley, 2005]. He has recorded extensively, is a well known flute maker and has been a master flute teacher for over 20 years.

Articulation	Description
Cut	Used to separate two notes. A cut is articulated by playing a middle note momentarily at a higher pitch than the second note.
Tap	An articulation used to separate two notes. A tap is articulated by playing a middle note momentarily at a lower pitch than the second note.
Roll	An articulation used to separate three notes. The second note in the sequence is cut and the third note is tapped. Rolls can be played long (duration of three notes) or short (duration of two notes).
Cran	Similar to a roll, but the tap is replaced with a second cut. The second cut uses a different intermediate higher note than the first one.
Trill	A rapid alternation between the principal note and the note above it
Triplet	A stepwise rising or falling sequence of three notes played in quick succession in the rhythm of two notes.

Figure 2: Fingered flute articulation techniques

Catherine McEvoy is described as playing in the *Sligo-Roscommon style*. This style is very rhythmical, with dynamic breath control used to punctuate and phrase tunes. She uses less fingered ornamentation than Eamonn Cotter, but instead incorporates subtle variations on tunes in her playing. Catherine has recorded extensively and has been a member of the band Macalla since 1984. She is also an experienced and sought after teacher.

4 Overview of MATT

MATT consists of several sub-systems as illustrated in Figure 3.

The transcriber transcribes monophonic flute tunes into ABC format for use as a corpus. The ABC format supports such features as notes, rolls, cuts, taps, crans, trills, breath marks, bar divisions, sharps, flats, naturals, repeated sections, key changes, guitar chords, lyrics and variations [Mansfield, 2004]. As the ABC format does not have any way of representing dynamics, this will is not considered initially. Automatic transcription of traditional music is non-trivial, due to the use of the high speed articulations presented in Figure 2. This problem is currently being addressed by [Gainza *et al.* 2004] and so to provide an initial corpus of tunes for the system to learn from, transcriptions were done manually with the aid of software for sound source separation [Barry, 2004] and time scale modification [Roni, 2005]. Using this technique, three tunes were transcribed as played by Eamonn Cotter. The tunes selected were "The Golden Keyboard, Ambrose Maloney's and Jackson's". These tunes are reels (in 4/4 time) and are commonly played in traditional music *sessions*. Each tune contains the standard AABB repetitions encountered in traditional tunes and was played twice through, so that the transcriptions contain many *variations* on the same phrases.

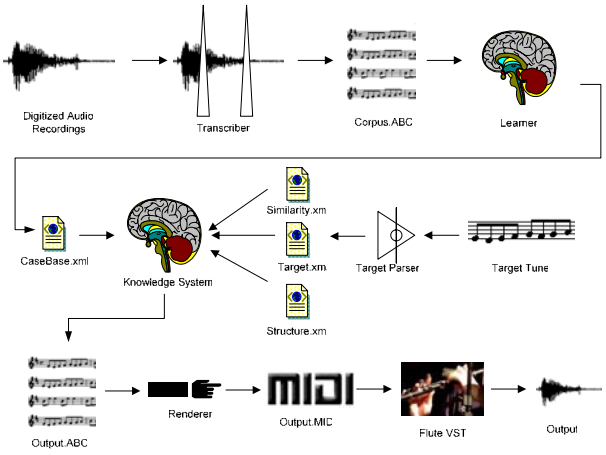

Figure 3: The main sub systems of MATT

Manual transcription of the same tunes recorded by Catherine McEvoy is underway, so it will be possible to compare the output of the system given corpora containing different playing styles.

Each transcribed tune represents approximately two minutes of playing and corresponds to 240 note events on average. Figure 4 is an extract from one of the tunes transcribed without expressive characteristics and the same extract as played by both musicians (some of the ABC headers have been removed for brevity).

The *learner* subsystem uses domain specific knowledge to compare transcriptions of tunes played by the musicians being modelled with transcriptions of the same tunes with no expressive characteristics. The learner divides each tune into half bars (4 quavers in length) and compares each half bar from the unornamented tune with the same half bar in the ornamented tune. The learner looks for instances of the articulations described in Figure 2 and also variations (in other words where the scores differ, but not in way that is an example of ornamentation). The learner builds a case base in XML format from the corpus in ABC format. *Cases* in a case base consist of vectors of features. The features selected for MATT are identified in Figure 5.

```
X:1
T:Ambrose Maloney's
H:No ornamentation
B3G ABGE|DGBG A3d|
BGGG A2ef|gedg eAAA|
B3G ABGE|DGBG A3d|
BGGG A2ef|ged=c BGGG|

X:2
T:Ambrose Maloney's
D:Eamonn Cotter: Track 1 - Traditional Irish Mu-
sic From County Clare
~B3G A{c}BGE|DGBG ~A3d|
B~G3 Az!Breath!ef|{a}g{ef}edg e~A3|
BdBG A{c}BGE|DGBG ~A3d|
B~G3 Az!Breath!ef|~g2d=c B~G3|

X:3
T:Ambrose Maloney's
H:Catherine McEvoy: Recorded live in the Cobble-
stone 01/02/2006
Bd{c}BG A{c}BGE|DG{c}BG ~A3d|
B~G3 Az!Breath!eg|{c'}gedg e~A3|
(3Bcd BD ABGE| Dz!Breath!BG ~A3d|
B~G3 ABea |{c'}ged=c B~G3|
```

Figure 4: Three transcriptions of the A part of the tune Ambrose Maloney's

Feature	Description
Notes	An n-gram of notes from the unornamented transcription
Duration	The duration in quavers of the n-gram
Key	The key the transcription is played in
TunePart	The tune part and repeat
Bar	Which bar the case occurs in
Position	The position within the bar
NewNote	An n-gram of notes from the ornamented transcription
Musician	The musician that played the piece transcribed
Tune Name	The name of the tune transcribed (not used for retrieval)

Figure 5: Feature vector of a case in MATT

Each tune results in on average 182 cases. We can conclude from this that in playing each tune, the musician being modelled has made 182 creative decisions. The corpus of three tunes used generates 547 cases. Examples of two cases identified are presented in Figure 6 (some features have been removed for brevity).

The *target parser* generates a target XML file from a given unornamented ABC file. It again divides the tune into

half bars. For every n-gram (where n > 2) of notes present in the half bar, the parser generates a target case. For example, the sequence of notes: A2EG generates three target cases: A2, EA and A2EA.

```
<case name="n7">
  <Notes>DAD</Notes>
  <Length>3</Length>
  <Key>EDor</Key>
  <TunePart>A1</TunePart>
  <Bar>4</Bar>
  <Position>7</Position>
  <NewNote>~D3</NewNote>
</case>
<case name="n531">
  <Notes>e2</Notes>
  <Length>2</Length>
  <Key>EDor</Key>
  <TunePart>B4</TunePart>
  <Bar>1</Bar>
  <Position>531</Position>
  <NewNote>ez!Breath!</NewNote>
</case>
```

Figure 6: Example cases generated by the learner

The *knowledge* subsystem takes as input the case base XML file, the target XML file generated by the two previously described subsystems and also a structure XML file and a similarity XML file. The structure XML file describes the hierarchy and cardinality of the features that can appear in a case [Coyle, 2002]. The similarity XML file describes the measure of the similarity between a target and a case [Coyle, 2004]. In MATT, the following features are used to calculate similarity: the duration of an n-gram (highly weighted so that only cases of identical duration are returned), the key of the tune, the tune part and repetition number (for example A1, A2, B1, B2), the bar number and the *edit distance* (again highly weighted) between the target n-gram of notes and the unornamented n-gram of notes from the case. Using the edit distance, allows phrases that are somewhat similar, but not necessarily identical to be returned.

Edit distance 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 n-gram to another [Grachten, 2004]. The retrieval system finds the highest scoring *k* cases for each half bar using the k-nearest neighbour algorithm. The case with the highest similarity score is chosen by the retrieval system and inserted into the output ABC file.

In order to make the output from the system sound realistic, a MIDI triggered wavetable VST instrument has been created using recordings of a real wooden flute playing each note staccato, legato and playing each of the articulations presented in Figure 2 for each note. Approximately 1GB of high quality samples are used in the VST instrument with the aim of creating an accurate rendering of the output of the retrieval engine.

The *renderer* takes the ABC file output from the retrieval subsystem and converts it to an augmented MIDI format which triggers the correct samples in the VST instrument. This subsystem uses software from the open source ABC2MIDI project [Shlien, 2006] and is still under development.

4 Testing

Two tests were carried out using the corpus of three tunes transcribed. A case base was built using three transcriptions of Eamonn Cotter's tunes to generate a case base of 547 cases. The unarticulated transcription of the tune "Jackson's" was used as a target for the knowledge system. As the tune "Jackon's" was used to build the case base, it was anticipated that the output produced would be broadly similar to the transcription of the tune used to build the case base. This was in fact the case and in most cases, the system chose the appropriate transformation. Interestingly the output contained several unusual artifacts. The artifacts were judged to be not entirely appropriate, though nonetheless they had a certain unique aesthetic novelty.

The second test was carried out using the corpus of three transcribed tunes to build the case base, but this time using an inexpressive transcription of the tune "The Copperplate" as a target. This tune is structurally similar to the tunes in the corpus (an AABB reel), but has an entirely different melody. The output from the knowledge system in the second test was a structurally correct tune (an AABB reel). It contained many appropriate examples of the articulations typically played in traditional music and listed in Figure 2. Additionally, the output contained some interesting variational artifacts, in other words instances where the same musical phrase was played differently from one repetition to the next. On the other hand in many cases the tune produced differed significantly from the target tune, beyond what would normally be considered "artistically appropriate".

5 Conclusions and Future Work

This paper outlined progress in developing MATT, a system which models creativity in traditional flute playing. An overview of creativity in Irish traditional music and the place of the flute in that tradition were provided. Techniques used by flute players were elaborated upon. The approach taken in developing MATT was elaborated and the results of testing carried out on the system in its present form were presented. From the tests carried out on the system we can conclude that MATT in its present form generates structurally correct, articulated traditional tunes. On the other hand, the fact that it also inappropriately alters the melody of the tune being articulated leads to the conclusion that the knowledge system at the core of MATT requires further work. An increased corpus will also enhance the quality of artefacts produced by the system.

Future work can be summarised as follows:

1. Enhancement of the knowledge system. The knowledge system at present uses a single case base and complex similarity to transform inexpressive to expressive tunes.

This area needs further work to produce accurate transformations. It would seem obvious that there should be two case bases, one for breath articulations and one for fingered articulations. Transformation would therefore become a two pass process. The calculation of melodic similarity also needs further work so as to avoid returning cases that are not melodically similar.

2. The VST sonic model of the wooden flute is under development at present. Currently although wavetables have been recorded for the full range of sounds a wooden flute can make, the VST just supports a subset of these. To evaluate the knowledge system, we propose using listening tests, which necessitate realist output from the system.
3. Work is also ongoing in transcribing tunes for the corpus. We anticipate much improved results with a larger corpus.

References

[Aamodt & Plaza, 1994] A. Aamodt and Enric Plaza. Case-based reasoning: Foundational issues, methodological variations, and system approaches, *AI-Communications*, 1994.

[Barry, 2004] Dan Barry. Sound Source Separation:Azimuth Discrimination and Resynthesis, Proc. of the 7th Int. Conference on Digital Audio Effects (DAFX-04), Naples, Italy, October 5-8, 2004

[Boden, 1996] Boden, M.A. *Dimensions of creativity*, Cambridge, Massachusetts: MIT Press, 1996

[Breathnach, 1977] Brendan Breathnach. Folk Music and Dances of Ireland, Mercier Press; Revised edition, June 1, 1977.

[Chambers, 2001] John Chambers. Frequently Asked Questions about ABC Music Notation, accessed March 2005.

[Cotter, 2002] Eamonn Cotter. The Golden Keyboard/Ambrose Moloney's/Jacksons, *Wooden Flute Obsession Vol. 1*, International Traditional Music Society, Inc., 2002

[Cotter, 2006] Eamonn Cotter, *Interview Notes*, 24 March, 2006

[Coyle, 2002] Lorcan Coyle, Conor Hayes, Pádraig Cunningham, Representing Cases for CBR in XML, Department of Computer Science,Trinity College Dublin

[Coyle, 2004] Lorcan Coyle, Dónal Doyle, Pádraig Cunningham, Representing Similarity for CBR in XML, Department of Computer Science, Trinity College Dublin

[De Mantaras & Arcos, 2002] Ramon Lopez de Mantaras and Josep Lluis Arcos. AI and Music: From Composition to Expressive Performance, *AI Magazine*, Fall 2002

[Driscoll, 2004] Sally Driscoll. A Trio of Internet Stars, ABC. *Fiddler Magazine*, Volume 11 Number 2, Summer 2004.

[Gainza et al. 2004] Mikel Gainza, Bob Lawlor and Eugene Coyle. Onset Detection and Music Transcription for the Irish Tin Whistle, ISSC 2004, Belfast, June 30 - July 2

[Grachten, 2004] Maarten Grachten, Josep Llu´ıs Arcos, and Ramon L´opez de M´antaras, TempoExpress, a CBR Approach to Musical Tempo Transformations, Advances in Case-Based Reasoning: 7th European Conference, ECCBR 2004, Madrid, Spain, August 30 - September 2, 2004.

[Heckroth, 1995] Jim Heckroth. *Tutorial on MIDI and Music Synthesis*, The MIDI Manufacturers Association, 1995.

[Hurley, 2005] Brad Hurley, An Interview with Eamonn Cotter, http://www.firescribble.net/flute/cotter.html, Accessed, April, 2006

[Larson, 2003] Grey Larsen. *The Essential Guide to Irish Flute and Tin Whistle.* Mel Bay Publications, Inc. 2003.

[Mansfield, 2004] Steve Mansfield. How to interpret abc music notation, http://www.lesession.co.uk/abc/abc_notation.htm, accessed March, 2005

[Ó hAllmhuráin, 1998] Gearoid Ó hAllmhuráin. *A Pocket History of Irish Traditional Music*, The O'Brien Press, Dublin, 1998.

[Roni, 2005] Roni Music. *The Amazing SlowDowner*, http://www.ronimusic.com/, accessed March 2005.

[Shlien, 2006] ABC2MIDI Home Page, http://abc.sourceforge.net/abcMIDI/, Accessed April, 2006.

[Smith & Medin, 1981] E. Smith, & D. Medin. *Categories and Concepts*, Harvard University Press, 1981.

[Steinberg, 2004] Steinberg Media Technologies GmbH, VST Module Architecture SDK Overview, January, 2004

[Vallely, 1999] Fintan Vallely. *The Companion to Irish Traditional Music*, New York University Press, 1999.

[Wallis & Wilson, 2001] Geoff Wallis and Sue Wilson. *The Rough Guide to Irish Music*. First Edition. Rough Guides, London, April 2001.

[Walshaw, 2005] Chris Walshaw. The abc home page, http://www.gre.ac.uk/~c.walshaw/abc/, accessed March 2005.