

2016

A Corpus of Annotated Irish Traditional Dance Music Recordings: Design and Benchmark Evaluations

Pierre Beauguitte

Technological University Dublin, pierre.beauguitte@mydit.ie

Bryan Duggan

Technological University Dublin, bryan.duggan@tudublin.ie

John D. Kelleher

Technological University Dublin, john.d.kelleher@tudublin.ie

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



Part of the [Computer Sciences Commons](#)

Recommended Citation

Beauguitte, P., Duggan, B., & Kelleher, J. (2016). A Corpus of Annotated Irish Traditional Dance Music Recordings: Design and Benchmark Evaluations. *17th International Society for Music Information Retrieval Conference 2016*, New York City, August 7-11. doi:10.5072/zenodo.242990

This Conference Paper is brought to you for free and open access by the School of Computing at ARROW@TU Dublin. It has been accepted for inclusion in Conference papers by an authorized administrator of ARROW@TU Dublin. For more information, please contact yvonne.desmond@tudublin.ie, arrow.admin@tudublin.ie, brian.widdis@tudublin.ie, aisling.coyne@tudublin.ie, fiona.x.farrell@tudublin.ie.



This work is licensed under a [Creative Commons Attribution-Noncommercial-Share Alike 3.0 License](#)

A CORPUS OF ANNOTATED IRISH TRADITIONAL DANCE MUSIC RECORDINGS: DESIGN AND BENCHMARK EVALUATIONS

Pierre Beauguitte, Bryan Duggan and John Kelleher
School of Computing, Dublin Institute of Technology

pierre.beauguitte@mydit.ie, bryan.duggan@dit.ie, john.d.kelleher@dit.ie

ABSTRACT

An emerging trend in music information retrieval (MIR) is the use of supervised machine learning to train automatic music transcription models. A prerequisite of adopting a machine learning methodology is the availability of annotated corpora. However, different genres of music have different characteristics and modelling these characteristics is an important part of creating state of the art MIR systems. Consequently, although some music corpora are available the use of these corpora is tied to the specific music genre, instrument type and recording context the corpus covers. This paper introduces the first corpus of annotations of audio recordings of Irish traditional dance music that covers multiple instrument types and both solo studio and live session recordings. We first discuss the considerations that motivated our design choices in developing the corpus. We then benchmark a number of automatic music transcription algorithms against the corpus.

1. INTRODUCTION

Automatic music transcription is one of the main challenges in MIR. The focus of most recent research is on polyphonic music, since it is sometimes claimed that the problem of transcribing a melody from a monophonic signal is solved [4]. The standard method in evaluating transcription algorithms is to take a data driven approach. This requires a corpus of data with original audio and ground truth annotations. Several such corpora exist, for example those provided for the annual MIREX evaluations. However no corpus can be universal, and the difficulty of gathering such a dataset is often discussed [20].

We are interested particularly in Irish traditional dance music and applications of automatic and accurate transcription in this genre. This musical tradition has several characteristics that are not captured by the existing datasets. Furthermore, to the best of our knowledge, there is no corpus of annotated Irish traditional dance music available. Since this is of absolute necessity in order to evaluate any new transcription method, we have created

and released our own corpus. This paper presents its design and some baseline results for existing melody extraction algorithms.

The structure of the paper is as follows: Section 2 explains the rationale behind the creation of this corpus. We also detail the characteristics of the music we consider and the challenges it presents. Section 3 presents existing work related to our own. We then give the detailed design of our audio annotations dataset in Section 4. In Section 5, we evaluate four transcription algorithms on the dataset.

2. CHARACTERISTICS OF IRISH TRADITIONAL DANCE MUSIC

In common with similar, aurally transmitted musical traditions, Irish traditional dance music is subject to variations and ornamentation in its interpretation. The melodies, or *tunes*, are usually rather short and consist of two or sometimes more repeated parts.

The nature of the music is melodic and modal, as opposed to the more common harmonic and tonal aesthetics of classical and contemporary Western music. Most of the tunes are played on a D major or G major scale, but they can be in different keys including D ionian or mixolydian (major keys), E dorian or aeolian (minor keys) or G mixolydian (major key).

Many tunes were originally dance tunes, with a fixed rhythm, tempo and structure, though nowadays the tunes are mostly performed for listening rather than dancing. The most common rhythms are reels, with a 4/4 time signature, and jigs, with a 6/8 time signature. Other types of tunes include polkas, hornpipes, slides, barndances and airs. An in-depth presentation of Irish music can be found in [22].

Although in the past Irish music was performed mostly by solo instrumentalists, and this type of performance is still highly valued today, in contemporary practice it is commonly played by ensembles of musicians. On commercial recordings, the most common configuration involves several melodic instruments with rhythmic and harmonic accompaniment.

The environment in which this music is most commonly played is that of *sessions*: gatherings of musicians, professional or amateurs, of a usually informal nature. When played in a session, Irish music can often be adequately qualified as heterophonic [23]. All players of melodic instruments (typically greater in number than rhythmic and



© Pierre Beauguitte, Bryan Duggan and John Kelleher. Licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). **Attribution:** Pierre Beauguitte, Bryan Duggan and John Kelleher. "A corpus of annotated Irish traditional dance music recordings: design and benchmark evaluations", 17th International Society for Music Information Retrieval Conference, 2016.

harmonic instruments) play the same tune together, but the result is often far from unison for several reasons. First of all, different instruments can play the same melody in different octaves (*e.g.* flute and tin whistle). Additionally, due to the acoustic limitations of certain instruments, or as an intended variation, some notes of a tune can be played in a different octave. The low B (B3) for example cannot be played on most traditional flutes. Consequently flute players often play a B4 instead, while a banjo player would play the “correct” note B3 and a whistle player would play an octave higher than the flute, B5. Yet, all would be considered as playing the same tune.

Another important aspect is the amount of variations present in Irish music. Because of personal or regional stylistic differences, the abundance of different sources (notations, archive or commercial recordings), and of the music being transmitted aurally (thus relying on the memory of the musician and therefore subject to interpretation), many different versions of a same tune may exist. Although musicians will often try to adapt to each other in order to play a common version during a session, it is not uncommon to hear some differences in the melodies. Finally, tunes are almost always ornamented differently by each individual musician depending on their style and personal preferences.

Our goal with this project is to create a representative corpus of annotated audio data. We will then be able to establish a baseline on the performance of existing transcription algorithms for this particular music style. This will facilitate the development and evaluation of new melody extraction algorithms for Irish traditional dance music. We believe this is of great interest for such applications as accurate transcription of archive recordings, improvement of popular digital tools for musicians [9], and monitoring of contemporary music practice [21].

3. RELATED WORK

Automatic melody transcription algorithms have been designed specifically for a *cappella* flamenco singing in [10], and evaluated against manual annotations made by experts musicians. Although the article cites previous work related to the creation of a representative audio dataset, the manual ground truths annotations were created specifically for that project.

Turkish makam music has also been the subject of research, in particular with the SymbTr project [12], whose goal was to offer a comprehensive database of symbolic music. Applications of automatic transcription algorithms have been studied in [5]. Ground truth annotations were obtained by manually aligning the symbolic transcription from the SymbTr database to existing audio recordings. The paper points out the predominance of monophonic and polyphonic Eurogenetic music in Music Information Retrieval evaluations, and the challenges presented by music falling out of this range, such as heterophonic music.

Applications of automatic transcription algorithms to large scale World & Traditional music archives are proposed in [1]. More than twenty-nine thousand pieces of au-

dio have been analysed. The obtained data have been made available through a Semantic Web server. No ground truth annotations were used to assess the quality of the transcriptions, that were obtained by the method presented in [3].

Manual and automatic annotations of commercial solo flute recordings of Irish traditional music have been conducted in [2], [11] and [13]. The focus in these papers is on style and automatic detection of ornaments in monophonic recordings. Consequently every ornament is finely transcribed in the annotations. An HMM-based transcription algorithm has been developed in [11] to recognise the different ornaments as well as the melody. An attempt was made at identifying players from these stylistic informations. The authors of [13] are planning on sharing the annotation corpus via the Semantic Web, but no data is publicly available yet.

4. PRESENTATION OF THE DATASET

We now describe the design of our corpus. First we present the set of audio recordings included, that was chosen to make the corpus representative of Irish traditional dance music. We then detail the annotation format used to deal with the characteristics of this genre.

4.1 Source of Audio Recordings

Three sources of recordings are included in the corpus:

- session recordings accompanying the Foinn Seisiún books published by the Comhaltas Ceoltóirí Éireann organisation, available with Creative Commons licence. These offer good quality, homogeneous examples of the heterophony inherent to an Irish traditional dance music session.
- Grey Larsen’s *MP3s for 300 Gems of Irish Music for All Instruments*, commercially available. These consist of studio quality recordings of tunes played on Irish flute, tin whistle and anglo concertina.
- personal recordings of the second author, a renowned musician, on the Irish flute. These are available together with the annotations.

The corpus comprises thirty tunes in total, which add up to more than thirty minutes of audio. We chose to include solo recordings as a way of comparing the performance of transcription algorithms on monophonic and heterophonic music. This set of tunes has been chosen to be representative of the corpus of Irish music in terms of type and key signature. Table 1 categorises the tunes in our corpus by tune type, key, and performance type.

The complete list of tunes with all relevant metadata is included with the dataset.

4.2 Format of the Annotations

We annotated each audio recording with note events, consisting of three values: onset time, duration and pitch. For the larger goal of obtaining a symbolic transcription, this format of annotation is more useful as well as easier to

	Reel	Jig	Hornpipe	Polka	Slide	Air
Dmaj	2	3	1	1	0	0
Gmaj	2	2	1	2	1	0
Amin	2	2	1	1	0	0
Emin	2	2	1	0	1	1
Bmin	1	0	0	0	1	0
Session	5	5	1	1	1	0
Solo	4	4	3	3	2	1

Table 1: Classification of tunes in our corpus by tune type, key, and performance type

obtain than a continuous pitch track labelling every audio frame. Despite the heterophonic nature of the session performances, there is always a single underlying monophonic melody. It is this melody we are interested in. For this reason there is no overlap between the notes, and the resulting transcription we present is monophonic.

Due to the heterophonic nature of Irish traditional music as played during a session, and to the slight tuning differences between the instruments, a single fundamental frequency cannot appropriately describe a note. Therefore we decided to report only MIDI note references instead of frequencies.

In session performances, the playing of ornamentation such as *rolls* and *cuts* [22] often results in several successive onsets for a single long note. Figure 1 shows an example of such a situation, where three different instruments interpret differently the same note (the short notes played by the flute are part of a *roll* and are not melodically significant, therefore they are not transcribed). This makes it difficult, even for experienced musicians and listeners, to know when repeated notes are to be considered as a single ornamented note or distinct notes. Because of this inherent ambiguity, it is not correct to associate one onset with one note in the transcription. For this reason, we decided to merge consecutive notes of identical pitch into one single note. A note in our transcriptions then corresponds to a change of pitch in the melody. For solo performances, there are some clear silences between notes (typically when the flute or whistle player takes a breath). Whenever such a silence occurs, we annotated two distinct notes even if they are of the same pitch. In the solo recordings present in the corpus, a breath typically lasts around 200ms. Notes that are repeated without pauses or cut by an ornament were still reported as a single note, in order to be consistent with the annotations of session recordings.

Manual annotations were made by the first author with the aid of the Tony software [14]. After importing an audio file, Tony offers estimates using the pYIN algorithm (note-level) presented in the next section. These were then manually corrected, by adding new notes, merging repeated notes and adjusting frequencies. The annotations were finally post-processed to convert these frequencies to the closest MIDI note references. With this annotation format, the dataset comprises in total more than 8600 notes.




Figure 1: Example of ornamented notes on different instruments

4.3 Open Publication of the Dataset

The dataset is publicly available as a set of csv files.¹ Each file contains the annotation for an entire audio file. Each line represents a note (as onset time, duration, MIDI note). The annotations can be easily used in any evaluation framework as well as with the software Sonic Visualiser [6].

5. EVALUATION OF EXISTING TRANSCRIPTION ALGORITHMS

In order to establish some baselines for melody extraction in recordings of Irish traditional dance music, we evaluate transcriptions of our audio corpus returned by four different algorithms.

5.1 Presentation of the Algorithms

The melody extraction (or more generally automatic transcription) algorithms we use rely on different signal processing approaches, and return estimated melodies in two different formats. Some return frame-level estimates, or continuous pitch tracks, in which case some pre-processing detailed later in 5.2.1 is needed to conduct the evaluations. Others return note-level estimates of the same format as that used for our annotations, sometimes with a continuous pitch track as an intermediate step.

5.1.1 pYIN

pYIN [15] stands for probabilistic YIN, and is based on the standard frequency estimation algorithm YIN [7], used in conjunction with HMM-based pitch tracking. The initial algorithm returns frame-level estimates, but an additional segmentation step based on HMM modelling of note events was introduced in [14]. We evaluate the two versions of the algorithm, both open source and available as a Vamp plug-in.²

It is important to note that because we manually annotated our corpus with Tony, offering the note-level estimates from pYIN as first approximations, results of the evaluations might be biased in favour of this method.

5.1.2 Melodia

Melodia [19] first extracts a salience function by detecting peaks in the time/frequency representation of the au-

¹ <https://github.com/pierrebeauguitte/tuneset>

² <https://code.soundsoftware.ac.uk/projects/pyin>

dio signal and then extracts the best continuous pitch track possible. It is available as a Vamp plug-in.³

5.1.3 MATT

MATT [8] stands for Machine Annotation of Traditional Tunes. It returns note-level transcriptions, by estimating the fundamental frequency from the harmonicity of the signal, and segmenting the resulting continuous pitch track according to its continuity as well as the energy level. Although the method used is not at the state of the art of melody extraction algorithms (for an overview, see [18]), the fact that it was designed specifically for and fine-tuned to Irish traditional music makes it of great interest to us. A Java implementation is available online.⁴

5.1.4 Silvet

Silvet [3] is based on Principal Latent Component Analysis. Although it is designed for polyphonic music transcription, obtaining a single melody track is achievable by simply limiting the number of notes occurring at any time to one. It first generates a pitch track by factorising the spectrogram according to predefined templates. This is then post-processed with HMM smoothing, in a similar manner to the pYIN segmentation step. This approach has a much higher computational cost due to the complexity of spectrogram factorisation. It is available as an open source Vamp plug-in.⁵

5.2 Evaluations of the Transcriptions

In order to be consistent with our annotation policy (see 4.2), it is necessary to post-process the estimates of these algorithms in the following manner:

- align all frequencies to the closest MIDI note;
- merge consecutive notes of same pitch separated by less than 200ms (only for note-level estimates).

The second step is particularly critical for the note-level metrics of the MIREX Note Tracking task, but will also affect the frame-level metrics of the Melody Extraction tasks for the frames in the filled gaps.

All evaluations are performed with the `mir_eval` open source framework presented in [16].⁶ In order to assess the statistical significance of the differences in scores, we use the Wilcoxon signed rank test (to compare our samples with the performances reported in the original publications), and the Mann-Whitney U -test (to compare between our different samples).

5.2.1 Frame-Level Evaluation: Melody Extraction Task

The MIREX Melody Extraction task evaluates transcriptions at the frame-level. Pre-processing of both the ground truth and the estimates is necessary, and simply consists of aligning both on the same 10ms time grid. The pitch estimate for a frame is considered correct if it is distant from

the ground truth by less than a quarter of a tone (50 cents). The metrics also look for voicing errors: a voiced frame is one where a melody pitch is present. Five different metrics are computed for each tune. Results are shown, using box-plots, in Figure 2, (a) for the solo recordings, and (b) for the session recordings.

The original publications introducing frame-level pYIN and MATT do not report these metrics. In [3], Silvet was only evaluated on corpora of polyphonic music, for which these metrics are not suitable. pYIN - notes was evaluated with the audio dataset from, and scores reported in [14] ranged from 0.83 to 0.85 for Overall Accuracy, and from 0.87 to 0.91 for Raw Pitch Accuracy. Evaluations conducted in [19] for Melodia on audio datasets from the MIREX evaluation resulted in Overall Accuracy of 0.77, Raw Chroma Accuracy of 0.83 and Raw Pitch Accuracy of 0.81. Scores obtained on our dataset are significantly lower for all these metrics and samples (p -values < 0.01), for both solo and session recordings. This seems to suggest that the genre of Irish traditional dance music does have some distinct characteristics, not limited to the heterophonic aspect of a session, that can present challenges for automatic transcription.

Comparing the Overall Accuracy on the solo and session subsets, it is interesting to see that MATT and both versions of pYIN have significantly lower scores on the session subset (p -values < 0.01), whereas Melodia and Silvet do not. We believe that this is because the Melodia and Silvet algorithms were specifically designed for polyphonic music analysis.

Raw chroma accuracy is, by definition, always higher than raw pitch accuracy. We observe that, on the solo recordings, this difference is only significant for Melodia (p -value < 0.05). On the session subset, it is very significant (p -values < 0.001) for all algorithms except Silvet. This suggests that Silvet is more robust for estimating fundamental frequencies.

5.2.2 Note-Level Evaluation: Note Tracking Task

We now present the results of the MIREX Note Tracking task. Although this task is primarily aiming at polyphonic music transcription systems, it also applies directly to monophonic music as long as both ground truth annotations and returned estimates are in a note-event format. In our case, this applies to pYIN - notes, MATT and Silvet.

Estimated notes are associated with their closest match from the ground truth, and a note is considered correctly transcribed if its onset is distant from the reference note by less than 50ms and its pitch by less than a quarter of a tone. Another way of evaluating the transcription is to also take the duration of notes into account. Commonly found instruments in Irish music have a wide range of acoustical characteristics: some (like the flute, the fiddle, the uilleann pipes) can be played *legato* or *staccato*, depending on personal or regional styles, or on the type of tunes performed; others (typically the banjo) can only be played *staccato*, with hard onsets and very little sustain. Consequently, the offset of the notes is of little interest for our evaluations,

³ <http://mtg.upf.edu/technologies/melodia>

⁴ <https://github.com/skooter500/matt2>

⁵ <https://code.soundsoftware.ac.uk/projects/silvet>

⁶ https://github.com/craffel/mir_eval

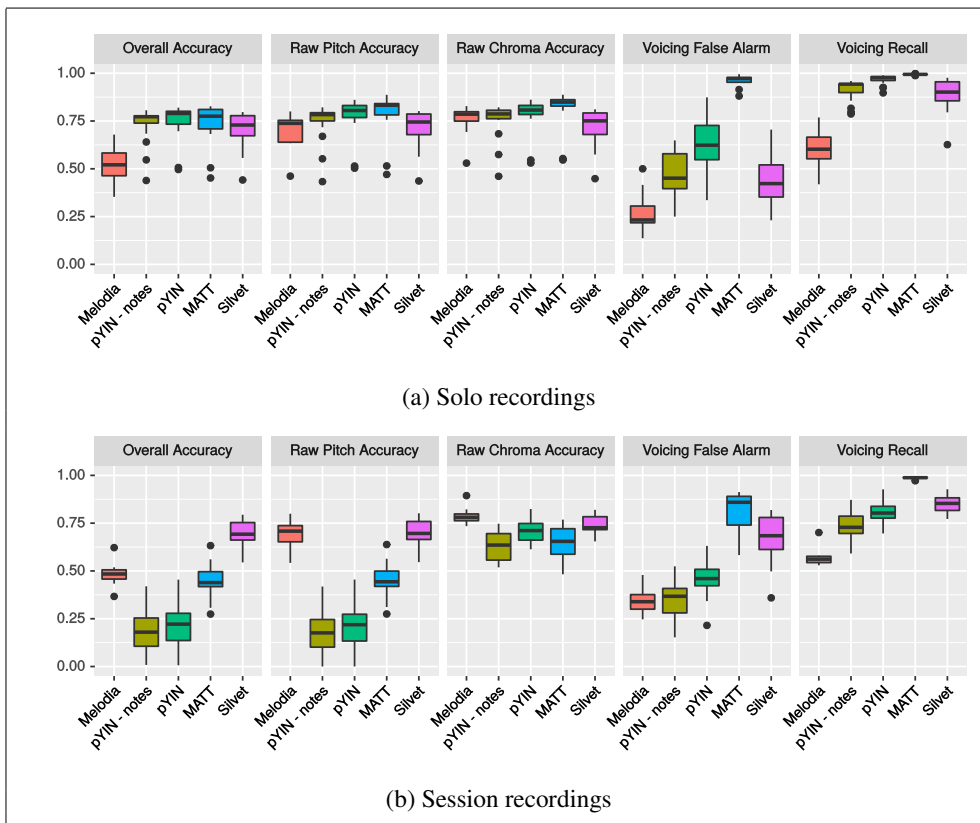


Figure 2: Scores of the MIREX Melody Extraction task evaluation metrics.

particularly in session recordings where all these instruments can play together. This is why we only report results for the first type of evaluation.

Precision, recall and F-measures are computed with the `mir_eval` framework [16], and plotted in Figure 3, (a) for the solo recordings and (b) for the session recordings. The figures also show the results for the Chroma correctness, where a note is correctly transcribed if its onset (± 50 ms) and its pitch class are correct. This is of interest to us because of the heterophonic nature of session performance described in Section 2.

It is interesting to observe that MATT achieves the highest F-measures on the solo recordings. However, only the difference with Silvet is statistically significant (p -values < 0.05). On the session subset, Silvet performs better than the other two algorithms, with high statistical significance (p -values < 0.01).

Precision scores are not significantly different on the solo recordings. On the session recordings, Silvet achieves significantly higher pitch precisions (p -values < 0.001). However, when looking at the chroma precision, the difference with pYIN is no longer significant. Scores for MATT remain significantly lower (p -values < 0.001).

Surprisingly, Silvet has higher F-measures on the session subset than on the solo subset, and this difference is statistically significant (p -value < 0.05). pYIN and MATT both score lower on the session subset. For the pitch F-measure, this difference is highly significant, with p -values

< 0.001 . For the chroma F-measure, only MATT scores significantly lower (p -value < 0.01).

6. CONCLUSION

In this paper, we introduced a new dataset of annotations of Irish traditional dance music. It is, to our knowledge, the first publicly available corpus of manually transcribed audio recordings in this genre. It covers a representative range of tune types, keys, and performance types. From the results obtained with state of the art automatic transcription algorithms, it appears that the heterophonic nature of this music presents challenges for MIR. It would have been of great interest to evaluate the HMM based algorithm used for flute ornaments recognition in [11], but unfortunately no implementation of it is publicly available.

These findings are good motivation to work towards the development of new methods for automatically transcribing Irish traditional dance music. We believe that this corpus will be of great use for this purpose, both for training data-driven models and to evaluate new algorithms.

We hope to make the corpus larger in the future, so that it includes a wider range of instruments and performance types (violin or *fiddle*, banjo, session recordings from other sources). We are also planning on making use of the Music Ontology [17] in later releases. Adopting the standards of the Semantic Web will hopefully allow more interaction with many resources such as, for example, the Europeana-Sounds database.

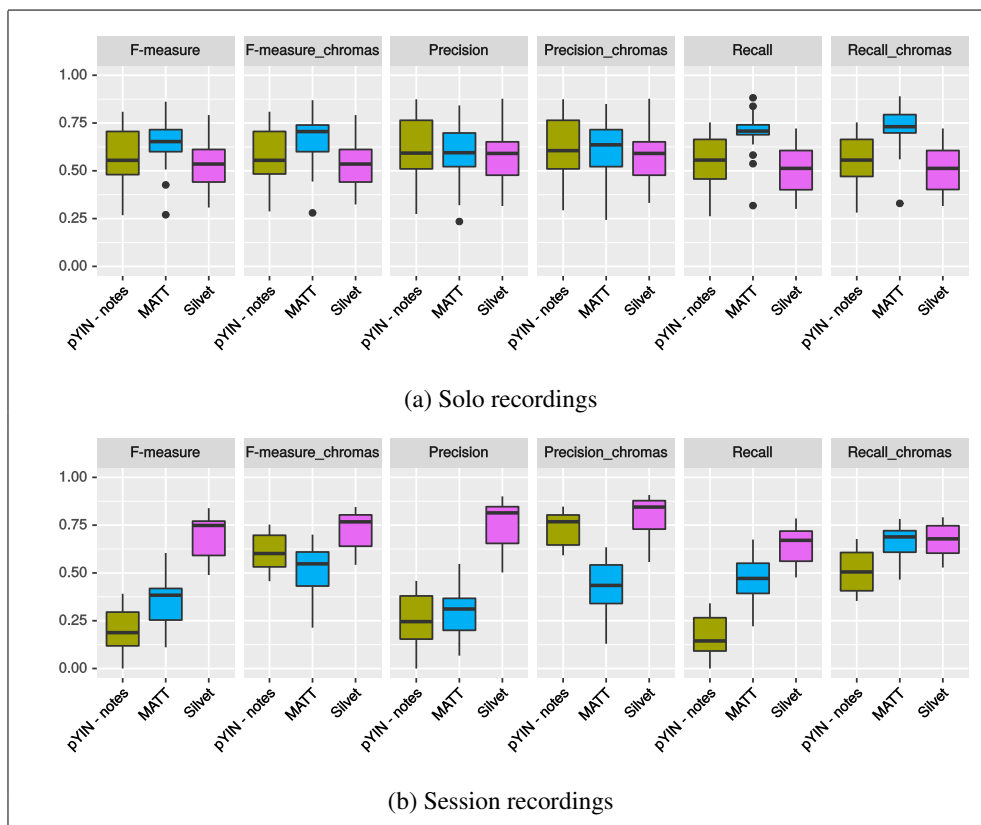


Figure 3: Scores of the MIREX Note Tracking task evaluation metrics.

7. REFERENCES

- [1] Samer Abdallah, Aquiles Alencar-Brayner, Emmanouil Benetos, Stephen Cottrell, Jason Dykes, Nicolas Gold, Alexander Kachkaev, Mahendra Mahey, Dan Tidhar, Adam Tovell, Tillman Weyde, and Daniel Wolff. Automatic Transcription and Pitch Analysis of the British Library World and Traditional Music Collections. In *Proceedings of the 5th International Workshop on Folk Music Analysis*, pages 10–12, 2015.
- [2] Islah Ali-MacLachlan, Münevver Köküer, Cham Athwal, and Peter Jančovič. Towards the Identification of Irish Traditional Flute Players from Commercial Recordings. In *Proceedings of the 5th International Workshop on Folk Music Analysis*, pages 13–17, 2015.
- [3] Emmanouil Benetos and Simon Dixon. A Shift-Invariant Latent Variable Model for Automatic Music Transcription. *Computer Music Journal*, 36(4):81–94, December 2012.
- [4] Emmanouil Benetos, Simon Dixon, Dimitrios Gianoulis, Holger Kirchhoff, and Anssi Klapuri. Automatic music transcription: challenges and future directions. *Journal of Intelligent Information Systems*, 41(3):407–434, July 2013.
- [5] Emmanouil Benetos and André Holzapfel. Automatic transcription of Turkish microtonal music. *The Journal of the Acoustical Society of America*, 138(4):2118–2130, 2015.
- [6] Chris Cannam, Christian Landone, and Mark Sandler. Sonic visualiser: An open source application for viewing, analysing, and annotating music audio files. In *Proceedings of the 18th ACM international conference on Multimedia*, pages 1467–1468. ACM, 2010.
- [7] Alain de Cheveigné and Hideki Kawahara. YIN, a fundamental frequency estimator for speech and music. *The Journal of the Acoustical Society of America*, 111(4):1917, 2002.
- [8] Bryan Duggan. *Machine annotation of traditional Irish dance music*. PhD thesis, Dublin Institute of Technology, June 2009.
- [9] Bryan Duggan and Brendan O’Shea. Tunepal: searching a digital library of traditional music scores. *OCLC Systems & Services: International digital library perspectives*, 27(4):284–297, October 2011.
- [10] Emilia Gómez and Jordi Bonada. Towards Computer-Assisted Flamenco Transcription: An Experimental Comparison of Automatic Transcription Algorithms As Applied to A Cappella Singing. *Computer Music Journal*, 37(2):73–90, 2013.
- [11] Peter Jančovič, Münevver Köküer, and Baptiste Wrena. Automatic Transcription of Ornamented Irish

- Traditional Flute Music Using Hidden Markov Models. In *Proceedings of the 16th International Society for Music Information Retrieval Conference*, 2015.
- [12] M. Kemal Karaosmanoğlu. A Turkish makam music symbolic database for music information retrieval: SymbTr. In *Proceedings of the 13th International Society for Music Information Retrieval Conference*, 2012.
- [13] Münevver Köküer, Daithí Kearney, Islah Ali-MacLachlan, Peter Jančovič, and Cham Athwal. 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, DLFM '14*, pages 1–3. ACM Press, 2014.
- [14] Matthias Mauch, Chris Cannam, Rachel Bittner, George Fazekas, Justin Salamon, Jiajie Dai, Juan Bello, and Simon Dixon. Computer-aided Melody Note Transcription Using the Tony Software: Accuracy and Efficiency. In *Proceedings of the First International Conference on Technologies for Music Notation and Representation*, 2015.
- [15] Matthias Mauch and Sam Dixon. pYIN: A fundamental frequency estimator using probabilistic threshold distributions. In *Proceedings of the 39th International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 659–663. IEEE, 2014.
- [16] Colin Raffel, Brian McFee, Eric J. Humphrey, Justin Salamon, Oriol Nieto, Dawen Liang, and Daniel PW Ellis. mir_eval: A Transparent Implementation Of Common MIR Metrics. In *Proceedings of the 15th International Society for Music Information Retrieval Conference*, 2014.
- [17] Yves Raimond, Samer A. Abdallah, Mark B. Sandler, and Frederick Giasson. The Music Ontology. In *Proceedings of the 8th International Society for Music Information Retrieval Conference*, pages 417–422, 2007.
- [18] Justin Salamon, Emilia Gomez, Daniel P.W. Ellis, and Guilhem Richard. Melody extraction from polyphonic music signals: Approaches, applications, and challenges. *Signal Processing Magazine, IEEE*, 31(2):118–134, 2014.
- [19] Justin Salamon and Emilia Gómez. Melody extraction from polyphonic music signals using pitch contour characteristics. *IEEE Transactions on Audio, Speech, and Language Processing*, 20(6):1759–1770, 2012.
- [20] Li Su and Yi-Hsuan Yang. Escaping from the Abyss of Manual Annotation: New Methodology of Building Polyphonic Datasets for Automatic Music Transcription. In *International Symposium on Computer Music Multidisciplinary Research*, 2015.
- [21] Norman Makoto Su and Bryan Duggan. TuneTracker: tensions in the surveillance of traditional music. In *ACM conference on Designing Interactive Systems*, pages 845–854. ACM Press, 2014.
- [22] Fintan Vallely, editor. *The Companion to Irish Traditional Music*. Cork University Press, Cork, second edition edition, 2011.
- [23] Sean Williams. *Focus: Irish Traditional Music*. Routledge, 2013.