Extended Abstract: Visualising Melodic Similarities in Folk Music

Chris Walshaw
Computing & Information Systems
University of Greenwich, London, UK
c.walshaw@gre.ac.uk

1. INTRODUCTION

The aim of this extended abstract is to outline a technique for visually exploring melodic relationships within folk tune collections. It stems from related work known as TuneGraph, [1], which allows users of abcnotation.com to explore melodic similarity. TuneGraph uses a similarity measure to derive a proximity graph representing similarities within the abc notation corpus. From this a local graph is extracted for each vertex, aimed at indicating close variants of the underlying tune. Finally an interactive user interface displays each local graph on that tune's webpage, allowing the user to explore melodic similarities.

As it stands TuneGraph only gives a localised view of the melodic relationships: this paper aims to look at exploring those relationships at a global (corpus-based) level.

2. METHODOLOGY

The essential idea is that, given a collection of tunes and a melodic similarity measure which can measure pairwise similarity between tunes (e.g. [2]), it is possible to construct a complete proximity graph of the corpus. Here the melodic similarity measure used is multilevel recursive sub-sequence alignment discussed in detail in [1] with some additional enhancements, tested in [3], also applied. However, the ideas are generic.

In the proximity graph each vertex represents a tune and edge weights represent similarities between tunes: the greater the similarity the larger the edge weight. If a similarity threshold, T, is applied so that an edge is only included if the two tunes it connects are sufficiently similar (if they match across at least some proportion T of their length, [3]) then a sparse proximity graph can be induced (the higher the threshold, the more sparse the graph). Subsequently, when the graphs are displayed, edge thickness is shown in proportion to the weight with similar vertices joined by thick edges and dissimilar ones by thin edges.

Following [1], the value for T used here is 1/6, a good compromise between restricting the number of edges (in order to make the visualisation tractable) but including enough to make the graph sufficiently rich. Further testing with other of values including 1/4 and 1/8 will be shown.

However, most reasonable values of the threshold typically generate a corpus graph with several *disconnected components* (subgraphs that are not connected by any edges) and often many *isolated vertices* (vertices with no incident edges – i.e. tunes that are not sufficiently similar to any other tune in the corpus to generate an edge).

This presents a problem for the investigation discussed here. An option is simply to visualise the largest connected component: however, this may only represent a small portion of the dataset. Accordingly, a straightforward scheme has been devised for connecting up the graph with a minimal number of zero weighted edges to help with the layout.

The setting of the edge weight to zero is important for the visualisation: since edge weights influence vertex placement, a zero weight edge will have minimal impact on the graph layout but the edge will mean that the two insufficiently similar vertices that it connects are positioned as close together as possible.

Once the corpus graphs are constructed they can be visualised using MultiLevel Force-Directed Placement (MLFDP) algorithms, e.g. [4], a standard technique for visualising large unstructured graphs.

3. RESULTS AND DISCUSSSION

3.1 Annotated datasets

The initial investigation explores two small datasets known to contain many related tunes and which have been annotated manually to indicate similar melodies, specifically those belonging to the same tune family.

The first of these datasets is the Annotated Corpus of the Meertens Tune Collection, version 2.0.1, [5]. This contains 360 Dutch folk melodies, each identified by experts as belonging to one of 26 tune families.

The second dataset contains 368 English morris dance tunes taken from the Morris Ring website[1]. Since morris music has several (approx. 35) traditions, each typically associated with a village, there are many tunes found in more than one tradition, but each tradition typically has a different variant of the tune. This dataset therefore contains 368 tunes which have been manually identified as belonging one of 113 tune families.

Because the datasets are annotated, the tune families induce a partition on the graph and different families can be visualised with different colours. Fig. 1 shows the results of the corpus graphs, visualised using MLFDP and overlaid with the partition induced by the tune families (zero-weight edges used to connect the graph are not displayed).

As can be seen, for both datasets the graph construction and visualisation is very complementary to the attribution of tunes to tune families: most edges go between tunes that are in the same family and even where they are not connected, most isolated vertices and small components (e.g. pairs) are close to other tunes in the same family.

Together these suggest that this kind visualisation can help to disambiguate tune families.

[1] https://themorrisring.org/music

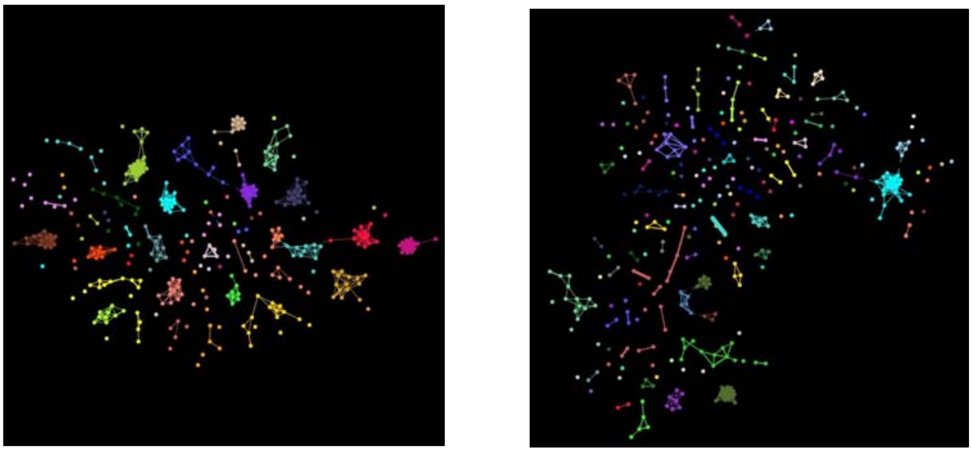

Figure 1. Visualisations of the Meertens (left) and Morris Ring (right) corpus graphs.

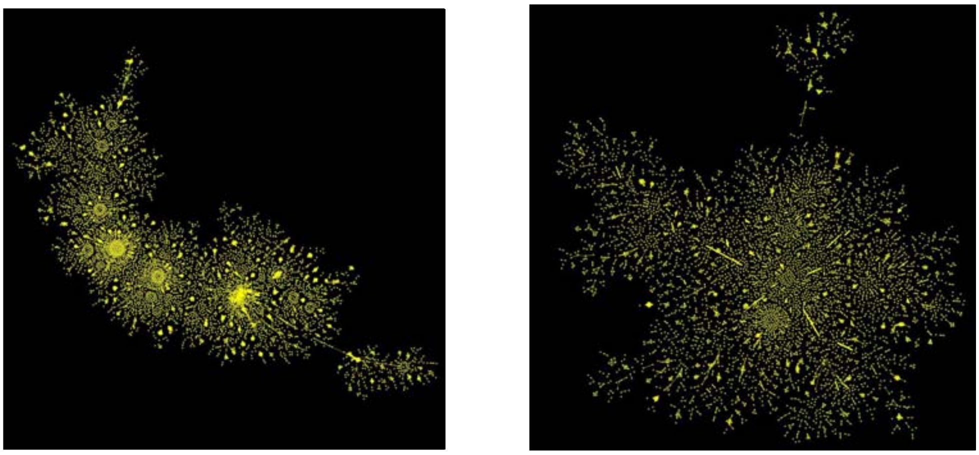

Figure 2. Visualisations of the Village Music Project (left) and TheSession (right) corpus graphs.

3.2 Collection-based datasets

The second set of results come from two tune collections which are not annotated (so not partition data is available). These are the Village Music Project[1] with around 5,600 English tunes transcribed from 18^{th} & 19^{th} century manuscripts and a subset of 5,000 of the ~30,000 tunes found at The Session[2], a community site which hosts a large collection of Irish traditional music.

Fig. 2 shows visualisations of the resulting corpus graphs. The structures are similar to their smaller counterparts in Fig. 1 although with many more disconnected vertices (these are much larger and much more disparate datasets). Nonetheless, tightly bound clusters of similar tunes are clearly visible, together with other structures, such as super-connectors (isolated vertices surrounded by sunflower like structures of other isolated vertices) and weak linkage (long, lightly-weighted edges which indicate loose connections between different subgraphs).

Although it is not easy to draw any conclusions from these two visualisations, it is encouraging to imagine an interactive exploration tool with user-driven zoom features and including score rendering and MIDI playing facilities.

4. REFERENCES

[1] C. Walshaw, "Constructing Proximity Graphs To Explore Similarities in Large-Scale Melodic Datasets," in *6th Intl Workshop on Folk Music Analysis*, 2016.

[2] B. Janssen, P. van Kranenburg, and A. Volk, "Finding occurrences of melodic segments in folk songs employing symbolic similarity measures," *J. New Music Res.*, p. (to appear), 2017.

[3] C. Walshaw, "Tune Classification using Multilevel Recursive Local Alignment Algorithms," in *7th Intl Workshop on Folk Music Analysis*, 2017.

[4] C. Walshaw, "A multilevel algorithm for force-directed graph-drawing," *J. Graph Algorithms Appl.*, vol. 7, no. 3, 2003.

[5] P. van Kranenburg, B. Janssen, and A. Volk, "The Meertens Tune Collections : The Annotated Corpus (MTC-ANN) Ver. 1.1 and 2.0.1," 2016.

[1] http://www.village-music-project.org.uk/

[2] https://thesession.org/