
<http://folkrrnn.org> and <https://themachinefolksession.org/>

Matthew Tobias Harris
Queen Mary University of
London
London E1 4NS, UK
operator@tobyz.net

Bob L. Sturm
Royal Institute of Technology
KTH
Lindstedtsvägen 24, SE-100 44
Stockholm, Sweden
boblsturm@gmail.com

Oded Ben-Tal
Kingston University
Kingston Hill, Kingston upon
Thames, Surrey KT2 7LB, UK
o.ben-tal@kingston.ac.uk

Paste the appropriate copyright statement here. ACM now supports three different copyright statements:

- ACM copyright: ACM holds the copyright on the work. This is the historical approach.
- License: The author(s) retain copyright, but ACM receives an exclusive publication license.
- Open Access: The author(s) wish to pay for the work to be open access. The additional fee must be paid to ACM.

This text field is large enough to hold the appropriate release statement assuming it is single spaced in a sans-serif 7 point font.

Every submission will be assigned their own unique DOI string to be included here.

Abstract

We demonstrate 1) a web-based implementation of a generative machine learning model trained on transcriptions of folk music from Ireland and the UK (<http://folkrrnn.org>, live since March 2018); 2) an online repository of work created by machines (<https://themachinefolksession.org/>, live since June 2018). These two websites provides a way for the public to engage with some of the outcomes of our research investigating the application of machine learning to music practice [2, 3], as well as the evaluation of machine learning applied in such contexts. Our machine learning model is built around a text-based vocabulary, which provides a very compact but expressive representation of melody-focused music [4]. The specific kind of model we use consists of three hidden layers of long short-term memory (LSTM) units [1]. We trained this model on over 23,000 transcriptions crowd-sourced from an online community devoted to these kinds of folk music.¹ Several compositions created with our application have been performed so far, and recorded and posted online.² We are also organising a composition competition using our web-based implementation, the winning piece of which will be performed at the 2018 O'Reilly AI conference in London in October.³

¹<http://thesession.org>

²https://www.youtube.com/channel/UC7wzmG64y2lbTUeWji_qKhA

³<https://folkrrnn.org/competition/>

Technical specifications

Our demo requires only a computer (laptop provided), headphones (provided), a table and an internet connection. The space and time for setup are minimal.

Description of presentation

Our presentation will be an interactive demo of the web-based implementation and the online repository. Figure 1 shows the *folk-rnn* application after a user has generated a new transcription. This user has clicked on "Compose", and the model has generated the text seen at top. The implementation renders this in common practice notation, and provides a way to audition the tune. A user can also change the initialisation of the model, and can provide seed tokens to start, e.g., the beginning of an existing melody. Figure 2 shows the welcome screen of The Machine Folk Session. This website allows people to explore and submit tunes generated by or with computers.

Examples of use

Our model has found use in several new compositions that have been performed, e.g.:

1. "Bastard Tunes" by Oded Ben-Tal <https://www.youtube.com/watch?v=YZ2jb0ksOm4>
2. "Between the Lines" by Oded Ben-Tal <https://www.youtube.com/watch?v=GdvyIH-0Q1k>
3. "Safe Houses" by Úna Monaghan <https://www.youtube.com/watch?v=x6LS9MbQj7Y&t=115s>
4. "March to the Mainframe" by Bob L. Sturm <https://www.youtube.com/watch?v=TLzBcMvI15M>
5. "Dialogues with folk-rnn" by Luca Turchet <https://www.youtube.com/watch?v=pkf3VqPieoo>

Furthermore, several tunes generated by folk-rnn have also been interpreted by professional musicians, e.g.,

1. Sets #1, 2 & 3 by an Irish ensemble: <https://www.youtube.com/watch?v=IZKc363886Y>; https://www.youtube.com/watch?v=_qpHaSwiJ-o&t=35s; <https://www.youtube.com/watch?v=j7RpmahiZQ>
2. "Interpretations of Computer-Generated Traditional Music" by John Hughes (db) <https://www.youtube.com/watch?v=GmwYtNgHW4g>
3. "Transcriptions #1469, 1470 & 1472 from The folk-rnn (v2) Session Book, vol. 1 of 10" by Torbjörn Hultmark (sop. trom.): <https://www.youtube.com/watch?v=4kLxvJ-rXDs>

Outside of these recordings, we have used the model to create "The Endless folk-rnn Traditional Music Session"⁴ and 34 pdf volumes of 100,000 "machine folk" tunes.⁵

Acknowledgements

"Engaging three user communities with applications and outcomes of computational music creativity", AHRC AH/R004706/1.

REFERENCES

1. B. L. Sturm. 2018. What do these 5,599,881 parameters mean? An analysis of a specific LSTM music transcription model, starting with the 70,281 parameters of its softmax layer. In *Proc. Music Metacreation workshop of ICCG*.

⁴ <http://www.eecs.qmul.ac.uk/~sturm/research/RNNIrishTrad/index.html>

⁵ <https://highnoongmt.wordpress.com/2018/01/05/volumes-1-20-of-folk-rnn-v1-transcriptions/>

