The Society for the Preservation and Promotion of Machine Folk Music (v1.1)

The Official Tunebook with Interesting Information about a Fascinating Hobby

October 12, 2020

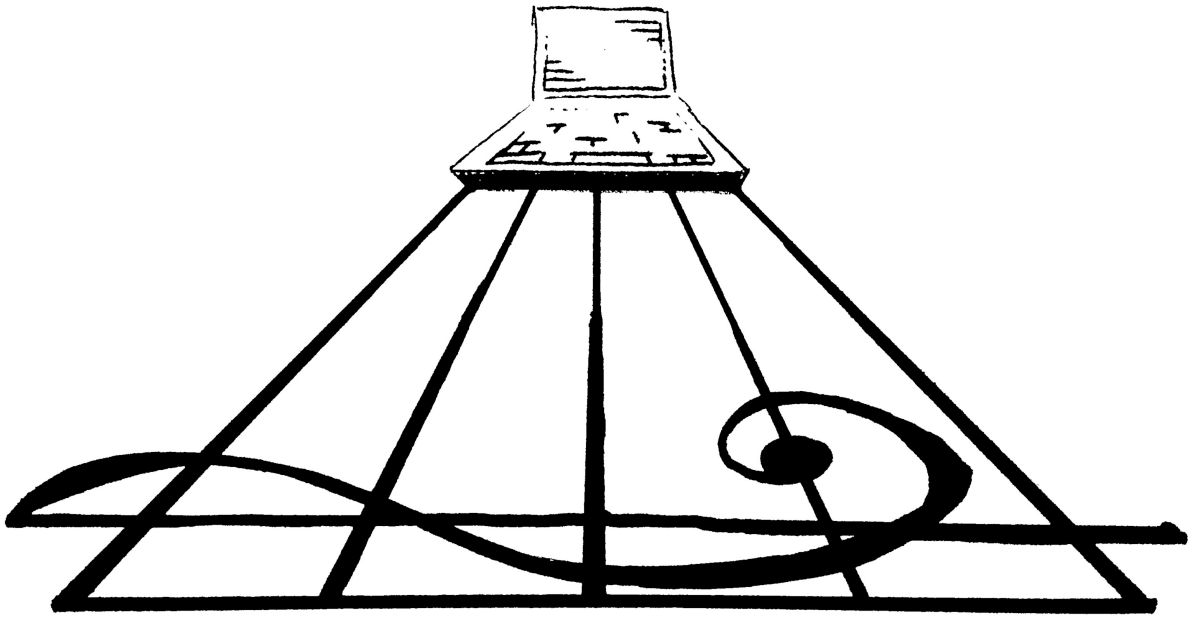

The Society for the Preservation and Promotion of Machine Folk Music (v1.1) was established in 2020, with support of the project, *Music at the Frontiers of Artificial Creativity and Criticism (MUSAiC)* (European Research Council Consolidator grant ERC-2019-COG, No. 864189), Bob L. T. Sturm presiding.

Praise for Machine Folk and The Society

- "I like to think that music comes from somewhere that's living. Otherwise it's. . . something else. Algorithms, computers - they'll always be predictable. Not so much humans. ... My concern is that some people, somewhere and sometime, may consider one or more of these tunes - maybe all of them? - to be actual traditional tunes." *Ergo*
- "The most useless piece of shit ever. Fuck off with that crap." *fuckthat*
- "I don't know whether to applaud or cry." *John Kacur*
- "Machines aren't creative as this proves." *tz1*
- "Why. Machine art is not human therefore just a soulless imitation. If it imitates it cannot create anything new." *worddust, Woodbridge*
- "Isn't music robotic enough these days?" *rocksnoop1, Dover*
- "Jeez. A computer that noodles. That's all we need." *Mark M*
- "This sounds like evil devil work." *Jerone, LoM*
- "Snacka om själ-lösa låtar... MUSIK.. Speciellt folkmusik.. Ska komma fram via upplevelser, traditioners djupa prägling. Upplevelser osv. Där människan är fokus, där folkmusiken präglat det traditionella kulturella livet. ... När jag ser sådant här blir jag antingen förbannad eller skitskarp. *Stefan Johansson*
- "Let's make all humans redundant, brilliant! Has everybody really lost their soul?!" *pen, somewhere*
- "This takes away possibilities for real musicians to compose Music and earn a living!" *Per Selander*
- "Correct me if I'm wrong, but isn't music supposed to come from the soul. Isn't it an expression of our humanity aren't the sentiments expressed supposed to touch us and make us empathise with the song writer/singer. AI generated music does none of those things, therefore it's pointless. It's like comparing a machine made chair to a Chippendale. Yes you can sit on both of them, but the beauty of a handmade, beautifully crafted chair cannot be compared to a mass produced conveyor belt item." *babsg, Newcastle*
- "I think there is still some hope for humankind... Truthfully, if I learnt a tune by mistake that was written by a computer I would drop it. There are too many great tunes with a human story. A tune named "Johnny O'Leary's" will have at least some connection - how it got from there to here through people. A tune called "Macbook Pro's" just doesn't have the same allure." *bogman*
- "This computerized 'AI is just so non musically untalented lazy nerds can infiltrate the world of true musicians who love, created, and write the music from the joy, hurt, and life emanating from their hearts." *Radar Also, Hemet*
- "It's a niche interest within a niche interest that will go under the radar of all but those most interested in the area and won't have the slightest effect on the lives of the vast majority of us. It's probably less useful in practical terms than building algorithms to direct robot vacuum cleaners or self-driving cars, but why knock it as a field of study? Most career academics that I know are deeply buried in some very esoteric trench in their field." *Namloc*
- "Can we not technologically tamper with everything that is good and pure in this world? A computer farting out generated tunes in some academic lab somewhere is the beginning of the end. ... The sooner this experiment is confined to an anonymous university archive the better." *anonymous*
- "There are loads of crap tunes written by humans and there will continue to be (as long as there are people like me) in perpetuity. I'd rather play a great tune written by a computer than a crap one I wrote myself." *Conán McDonnell*

1 Introduction

This growing curated collection contains tunes arising from creative partnerships with music Ai.[1] Some Ai are trained on over 23,000 music transcriptions contributed by users of `thesession.org`, which is focused on Irish traditional dance music. These are named *folk-rnn* v1, v2 and v3. The different versions arise from different formats of the training material: v1 models character sequences in a single document, including tune titles; v2 models tokenized transcriptions that have been transposed to have a root of C; v3 uses a slightly different data representation from v2. Another Ai, *folk-rnn* (v. Swedish), was tuned using 4,000 transcriptions of Scandinavian folk music, collected from `folkwiki.se`.

Each transcription herein revealed itself to the real-world through a computational procedure involving on average one billion operations. They are not purely products of a cold and lifeless statistical algorithm attempting to imitate patterns it has learned from existing data — some human operator is needed to "flip the switch", and then to curate from the materials generated. Often the Ai generates poor results. Sometimes they are clearly derivative. In a few cases, they can be unusual and wonderful. The challenge is to find these "diamonds" in the raw materials. Then to bring them to life without ever heard them played before by a master musician. Even so, there remains the unsurpassable barrier of a tune existing without any trace whatsoever in a collective memory of a community of practitioners. Their context is digital vapor.

Some believe in a purity or divinity of Art, and the superiority of humans in making Art. Involving a machine in Art can be seen as a direct challenge to the purity of Art, not to mention a demotion of the human to machine. This argument elevates the tangible product over the intangible experience, and submits to a rather narrow and mystical notion of creating Art. Art is a human activity, full stop. It is an activity that occurs between our ears in our sensations, thoughts and memories. Art is not contained on the page or in the frame. It does not stop when we leave the concert hall or the museum. Art is entirely steeped in being human. The human decision to involve technology in that activity is only that: the suspension of pigment in egg yolk to make it stick to a surface; the suspension of pigment in slow-drying oil to make it blend and layer in ways superior to egg tempura; the suspension of pigment in plastic medium to make it sculptable and fast drying; the use of a coarse horse hair brush to paint several strands of hair at once; the use of a palette knife to make sharp straight edges; the use of the principles of geometry and perspective to create trompe-l'œil. How is involving Ai in creating music any different?

One pitfall important to avoid when it comes to discussing Ai in the Arts is this: thinking that the terms "intelligence" and "learning" mean what they commonly mean when it comes to people. These are "suitcase" terms with several meanings that can be confused. In the context of Ai, "intelligence" is a quite brittle thing that bears very little resemblance to human intelligence. And "learning" is merely the prosaic estimation of numbers in an algorithm that represent relationships like probabilities from data. Falling into this trap can give power to the Ai that it does not actually possess, leading to dystopian fantasies. To claim Ai will destroy music is the same as claiming, e.g., photography will destroy painting, perspective drawing will destroy narrative, and so on. Yes, the future is filled with composers and painters and photographers and writers ... The only real threat to the Arts are the lack of value for it in a society, and proscriptions of practice due to pressures economic, political, religious, etc.

[1] "Ai" is an acronym for "artificial intelligence". We do not capitalize the "i" because the "intelligence" is questionable.

2 Key references

1. The Society for the Preservation and Promotion of Machine Folk Music (v1.1) facebook group
2. The Google Drive of The Society
3. `folk-rnn.org`
4. `themachinefolksession.org`
5. The computer code and data behind folk-rnn.
6. Repository of The Endless folk-rnn Session website
7. The folk-rnn Session Books (34 volumes of 100,000 transcriptions)
8. "Let's Have Another Gan Ainm": An experimental album of Irish traditional music and computer-generated tunes (Associated Technical Report)
9. Sturm, Santos, Ben-Tal and Korshunova, "Music transcription modelling and composition using deep learning", in Proc. 1st Conf. Computer Simulation of Musical Creativity, July 2016.
10. Sturm and Ben-Tal, "Taking the Models back to Music Practice: Evaluating Generative Transcription Models built using Deep Learning," J. Creative Music Systems 2(1) Sep. 2017.
11. Sturm, "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 ICCC, 2018.
12. Sturm, Ben-Tal, Monaghan, Collins, Herremans, Chew, Hadjeres, Deruty and Pachet, "Machine learning research that matters for music creation: A case study," J. New Music Research 48(1):36-55, 2018.
13. Holzapfel, Sturm, and Coeckelbergh, "Ethical dimensions of music information retrieval technology," Trans. Int. Soc. Music Information Retrieval 1(1):44–55, 2018.
14. Sturm, Iglesias, Ben-Tal, Miron, and Gómez, "Artificial intelligence and music: Open questions of copyright law and engineering praxis," MDPI Arts 8(3), 2019.
15. Ben-Tal, Harris, and Sturm, "How music AI is useful: Engagements with composers, performers, and audiences," Leonardo Music Journal (accepted 2020).

3 Inspirations

1. Michael Snow, "The Last LP: Unique Last Recordings Of The Music Of Ancient Cultures"
2. Jennifer Walshe's The Avant Garde Archive of Ireland (`aisteach.org`)
3. Jennifer Walshe, "Ghosts of the Hidden Layer" seminar at Darmstädter Ferienkurse, July 25 2018
4. League of Automatic Music Composers

4 Appearances of Members of The Society

1. Feb 26 2020 at the Austrian Research Institute for Artificial Intelligence, Vienna
2. Mar 12 2020 at "AI for Humanity and Society 2020", Stockholm (cancelled due to COVID-19)
3. Sep 11 2020 at Ars Electronica 2020

5 Frequently Asked Questions

1. "What is *machine folk*?"
   Music that springs from human and machine (Ai) partnerships, played on folk instruments.

2. "Why *machine folk*?"
   Why not?

3. "I mean why not play *real* traditional music?"
   Let's! Playing *machine folk* is not an exclusionary practice. Plus there is something in *machine folk* that speaks of a people awash in computer technology, data harvesting, surveillance and automation.

4. "Do we really need to involve Ai in music?"
   No, in the same way that we don't really need to involve pianos.

5. "Why do you have a lower-case 'i' when you write 'Ai"'?
   To highlight the fact that the "intelligence" is questionable.

6. "Isn't music robotic enough these days?"
   "Robotic" is an aim for some styles, e.g., sequenced, electronic dance music. For *machine folk*, "robotic" need not be the aim, or an unavoidable outcome.

7. "Isn't music supposed to come from the soul? Isn't it an expression of our humanity? Aren't the sentiments expressed supposed to touch us and make us empathize with the composer?"
   Many of us have sentiments for machines. Besides, it seems somewhat normal to care for an algorithm that one has watched grow and care for.

8. "Have you lost your soul?"
   Not that we are aware of; or, maybe yes? Let's ask the computer to generate a tune about it.

9. "Can we not technologically tamper with everything that is good and pure in this world?"
   We are definitely partial to the romantic idea of returning to the way things were before fire was discovered. But then we realize that most musical instruments wouldn't exist.

10. "Are you working to replace human composers and musicians?"
    No. Humans are an essential part of *machine folk*, from the programming and training of the Ai, to the curation of the generated music and its performance in the real world on real instruments.

11. "Won't this technology lead to replacing human composers and musicians?"
    If something can be automated, it will be automated. Regardless if an Ai can compose music, humans will continue to compose. Suffering is an essential part of the human condition.

12. "Who owns these tunes?"
    This is not yet clear, but the meaningfulness of this question is not clear in a society that has decided it's acceptable to not pay for music.[2]

13. "How will this help traditional music?"
    Does traditional music need help? We are more concerned with how this could harm traditional music.

14. "OK, How will this harm traditional music?"
    We are not at all sure yet, but are definitely keeping an eye out.

15. "By what criteria does The Society select computer generated tunes for this tunebook?"
    The criteria have yet to be formally defined, but certainly among them are things like musical sense, quirkiness, playability, and humor.

[2] Compliments to Holly Herndon.

The Boys Of Ballinaburre

folk-rnn (v2, with beamsearch n=2)

This tune was learned from folk-rnn (v2), which generated it under a particular sampling regimen. More information here. There's a claim that the turn resembles that of a traditional Irish tune named, "The Thatched Cabin". That particular tune does not appear in the training data of this folk-rnn model.

Why are you so hard to understand?

folk-rnn (v2)

This tune was learned from folk-rnn (v2). It is better played as a very slow hymn or air. The title refers to difficulties in the investigation of how the Ai model is working. This tune made its first appearance in Society Member Sturm's paper, "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 the Int. Conf. Computational Creativity, 2018.

Mickey Fitternaly's

folk-rnn (v2, with beamsearch n=2)

This reel was learned from folk-rnn (v2), which generated it under a particular sampling regimen. More information here. Society Member Sturm has slightly modified the last half of the sixth bar of the turn such that it repeats the 2nd bar of the turn.

The Liddle Shepe

folk-rnn (v2, with beamsearch n=2)

This lovely hornpipe was learned from folk-rnn (v2), which generated it under a particular sampling regimen. More information here.

Aloe Vera's Jig

folk-rnn (v2) + Sturm

When Society Member Näsström sent the photo at left to the Society, Member Sturm decided to ask folk-rnn (v2) to compose a jig about it. The output, seen in small notation, was then modified a bit by Sturm to create the piece. More information about this particular transcription is here.

Garrett Farmor's

folk-rnn (v2)

This tune was learned from folk-rnn (v2), and appears as #2857 in the folk-rnn (v2) Session Book, Vol. 1 of 10. More information on this tune can be found here. Society Member Sturm notates it here in 12/8 since it feels more like a slide than a jig.

Swing Swang Swung

folk-rnn (v2)

This tune was learned from folk-rnn (v2). The opening bears some similarity to "Swing Swang" played by Na Draíodóirí (a tune that is in the training data of this model). More information about this particular transcription is here.

Cleaper's Wheel

folk-rnn (v2) + Sturm

The Society members Larsson, Kmoch and Sturm found this tune together while browsing the Endless MF Session website (synthesized recording). It appears in the folk-rnn (v2) Session Book, Vol. 2 of 10. The title was taken from a tune generated by folk-rnn (v1).

Gallagher's Favourite

folk-rnn (v3) + Sturm

This tune appears in the folk-rnn (v3) Session Book, Vol. 3 of 4, as tune #5712.

Wolden Sailor

folk-rnn (v2) + Sturm

This tune appears in the folk-rnn (v2) Session Book, Vol. 7 of 10, as tune #18727.

Måndag på KTH

folk-rnn (v. Swedish)

This polska was learned from folk-rnn (v. Swedish). More information here.