This is an accepted manuscript which is to be published in Leonardo Music Journal. The final version is available at https://www.mitpressjournals.org/doi/abs/10.1162/leon_a_01959#authorsTabList

How music AI is useful: Engagements with composers, performers, and audiences.

Oded Ben-Tal, Matthew Tobias Harris and Bob L. Sturm

Oded Ben-Tal (composer, researcher), Kingston University, Department of Performing Arts, Coombehurst House, KT27LA, UK. o.ben-tal@kingston.ac.uk. http://obental.wixsite.com/main

Matthew Tobias Harris (researcher), Queen Mary University of London, London E1 4NS, UK. operator@tobyz.net. http://tobyz.net

Bob L. Sturm (researcher), Royal Institute of Technology KTH, Tal, Musik och Hörsel (Speech, Music and Hearing), Lindstedtsvägen 24 SE-100 44 Stockholm, Sweden. bobs@kth.se. https://highnoongmt.wordpress.com

Abstract

Critical but often overlooked questions of research in artificial intelligence (AI) applied to music involve the impact of the resulting models for music. How and to what extent does such research contribute to the domain of music? How are the resulting models useful for music practitioners? In this article, we describe how we are addressing such questions by engaging composers, musicians, and audiences with our research. We describe two websites we have created that make our AI models accessible to a wide audience. We describe a professionally recorded album that we released to gauge the plausibility of material generated by our models and reviewers' comments on the music. Finally, we describe the use of our AI models as a tool for co-creation. Evaluating AI music models in these ways illuminate their impact on music making in a range of styles and practices.

1 Introduction

When applying artificial intelligence (more specifically Machine Learning techniques that underpin current research) in a creative field such as music, poetry or painting, one critical question to answer is, Why? Why should these technologies be applied to such an activity? Instead, much research and development in this area try to answer different questions, e.g., Can AI music system XYZ fool humans into believing its creations are by humans [1]? Or, how statistically similar are AI-generated outputs and the dataset used to train it [2]? Such measures may be useful in domains where success and failure are clearly defined (e.g. medical diagnosis), but when applied to *Art* these evaluation methods are insufficient. As Agres, et al [3] argue, evaluating creative systems requires looking beyond the generated outputs. The role of expertise and the perspectives of different target audiences are important aspects to consider.

Motivated by Wagstaff's key message in her position paper, "Machine Learning that Matters" [4], our research addresses the application of AI to the domain of musical practice: performance, composition, and improvisation. Our aim is twofold: to test how such AI systems can operate as part of a music ecosystem; and, to engage more actors in that ecosystem with the questions, problems, opportunities, and challenges that AI raises for music (and the other Arts by extension). We do this by engaging a range of users with the AI models we developed and critically examine the multiple ways in which these models can be used creatively within a diverse set of musical practices. This highlights the contribution such novel technology can make to the domain of music, and suggests where future developments could be most fruitful.

Machine learning (ML) essentially involves making a computer learn patterns by example, thereby sidestepping the codification of conventions that may not be so easy to express in computational language. This makes ML an attractive approach for modelling, generating and transforming music [5 – 10]. The majority of current work in music ML revolves around the same musical tasks that have been explored computationally almost as long as computers existed [3, 9], e.g., melody and harmony generation in a few known styles, such as jazz or JS Bach's chorales. This reflects both the availability of data needed to train the models and the extensive theorising which make it possible to interpret the outputs in musical context.

Our own research [11, 12, 14, 15] applies off-the-shelf deep ML to the specific domain of Western European folk music. The first data source for our models consisted of text-based transcriptions, in ABC notation [13], of traditional dance music mostly from Ireland and the UK. These transcriptions were crowd-sourced at thesession.org -- a community website for enthusiasts of that music. After data cleaning - removing incomplete transcriptions and unrelated examples such as Cage's 4'33" - we had over 23,000 transcriptions with which to train models. In brief, these models learn to predict a probability distribution over a vocabulary conditioned on the sequence generated up to that point. Iteratively sampling from that distribution leads to a generated transcription (for a more detailed discussion see [12]). We have also trained models on Scandinavian folk tunes collected from another data source [14]. With the aim of evaluating the potential contribution of such AI systems to music we have randomly subjected generated material to musical analysis (see section 3.2 of [11]). We have examined the performance of the system when prodded with un-idiomatic initial sequence (section 3.3 'Nefarious Testing'). We have solicited opinions from users of the same online forum where we obtained the data (section 3.5). And we have used the system for composition (section 3.4 and section 2.1 in

[15]). In this paper we discuss how we have extended our evaluation further by engaging wider audiences.

In section 2 we discuss a pair of websites we have developed that allow online users to work with our models, generate and archive tunes and engage with others. In section 3 we discuss the recording and dissemination of an album with material generated by our models. In section 4 we discuss musicians, working outside the tradition of the training data we have used, interacting with our models.

2 Accessible Online Implementations

We have created two websites around our models making the model much more accessible. One is an interactive interface to our models, the other is a growing repository of music generated by or with folkrnn. Implementation details can be found at https://github.com/tobyspark/folk-rnn-webapp

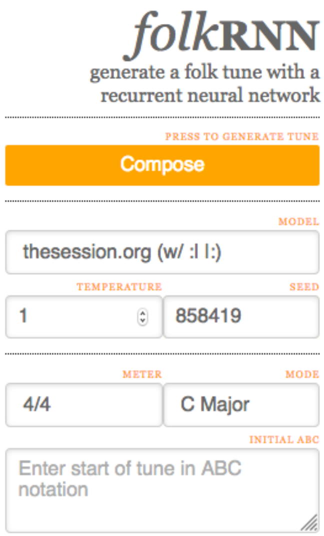

Figure 1 - music generation controls panel of https://folkrnn.org

2.1 Two web resources: https://folkrnn.org and https://themachinefolksession.org

The website https://folkrnn.org comprises an optimised, server-based implementation of our models, and a user interface that exposes functionality in a straightforward and appealing manner then the previous command line interface. The interface comprises a left-hand panel, shown in Fig. 1, that presents the music generation controls, with a main section that scrolls to hold each tune as it is generated by a single user. On the initial page load, this section shows information about the site, including the motivating ideas and a walk-through video showing the functionality in use.

The “Compose” button is the most prominent control on the page, clicking it results in a new tune appearing note-by-note as it is generated. Further controls are provided for iterative or deliberate use. A particular model can be selected, each differentiated by the data we used to train it. The temperature parameter can be raised or lowered, determining how “adventurous” the model acts. The seed parameter controls the internal pseudorandom state, to produce new transcriptions for the same (often default) parameters. It will change for each tune generated unless ‘pinned’ by manual input. The meter can be selected from a set of options, e.g. 4/4, 6/8, 9/8. The mode also can be selected from a set of options, i.e., C major, C minor, C dorian, C mixolydian. The “initial ABC” text box allows the beginning of a tune to be specified, which the model then completes. In addition to the textual ABC[13] representation output, staff notation and audio playback are provided; playback animation links all three representations. A user can download the result in MIDI format, or archive the result at The Machine Folk Session website.

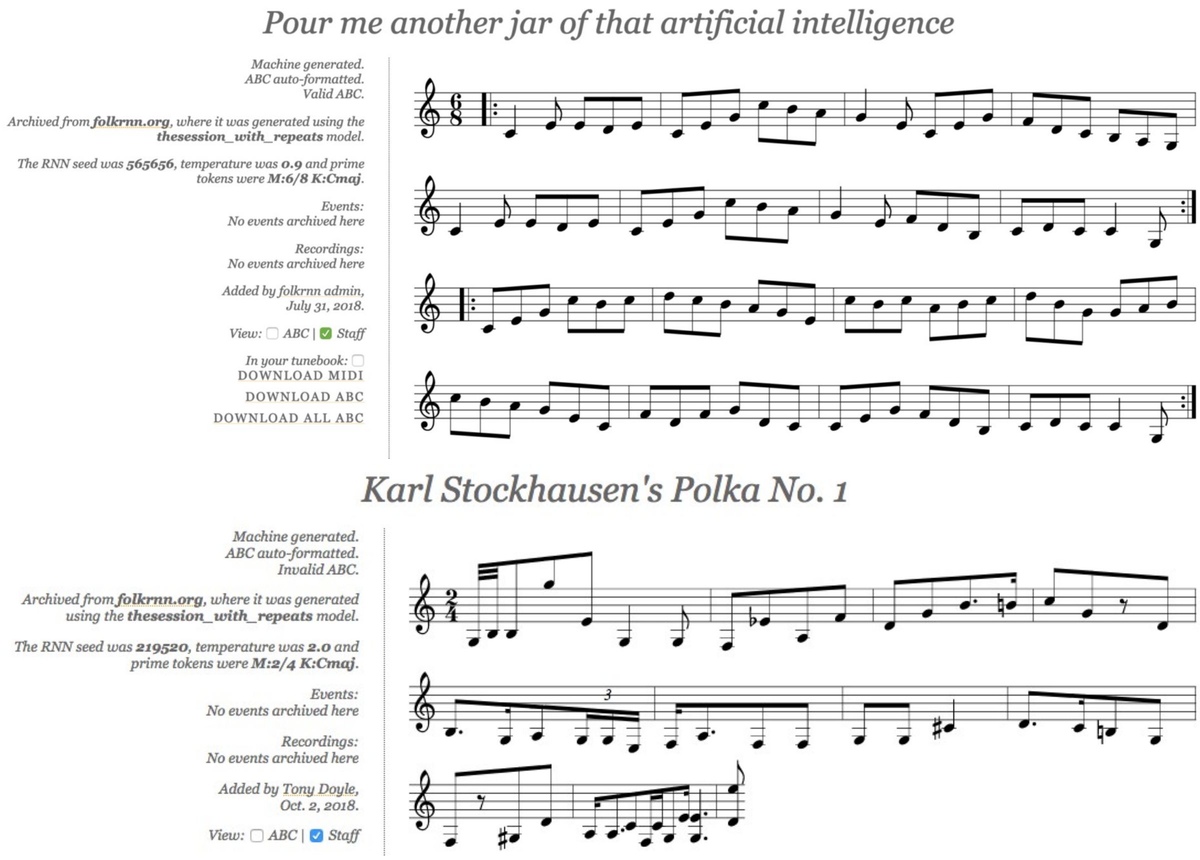

Figure 2 - Two tunes generated by folkrnn.org and archived at to themachinefolksession.org

The website https://themachinefolksession.org serves as a community-driven archive dedicated to music created by or with folkrnn. The site is primarily organised around tunes. On any given tune’s page the original submission can be seen along with any backstory, settings (i.e. an edit or variation of the original tune), performances (as video or audio recordings), comments, and links to any events that featured the tune. Users registered to the website can add tunes to their tunebooks. Inspired by folk sessions elsewhere, we are experimenting with features such as ‘tune of the month’ where the community selects a tune to all learn that month, and contribute their particular takes (though this has not been a success so far). Figure 2 shows two example tunes submitted by users to the website. One is clearly outside the idiom of Irish traditional music.

2.2 Usage

Our analysis of the server data of our websites shows its use, and the impact of media attention. During the first 235 days of activity at folkrnn.org, 24562 tunes were generated by approximately 5700 users. Activity for the first 18 weeks averages a median of 155 tunes weekly. Since then overall use increased with a median of 665 tunes generated weekly (as of August 2019). This period also features usage spikes. The largest, correlating to a mention in German media [16], shows an 18.4x increase in tunes generated per week. The 5700 people who have engaged with the online implementation in this period compares with around 250 people who have engaged with the command line tool in the three years of its existence [17].

The data provides evidence of human-machine co-composition using the folkrnn.org system. There are 4007 transcriptions where each tune has one or more generation parameter changed from the previous. We see an average of 6 (mean: 5.9, stddev: 8.7) iterations in such processes and they account for 57% of all transcriptions generated. *Temperature* is the most used parameter, at 40%. This has the simplest action of the generation parameters in the UI – since it is a simple numeric value that can be increased or decreased. Changing temperature also results in more obviously dramatic changes in the generated material; increasing the temperature from 1 to 2 will often yield tunes that do not sound traditional at all (as “Stockhausen’s Polka” in fig. 2 illustrates). The “Initial ABC” textbox is used 20% of the time, which is notable as this requires text manipulation on the part of the user.

The strongest metric of co-composition available on https://folkrnn.org is whether *Initial ABC* contains a fragment of the previous generated transcription. This suggests the user has identified an interesting or useful portion, and wishes to prime the next generation with it. Testing for phrases comprising five characters or more (e.g. five notes, or fewer with duration values), we find this happened 283 times, i.e. 2% of the time.

We find 239 of the 'iterative' folkrnn.org transcriptions archived to https://themachinefolksession.org, such as ‘The Green Electrodes’: https://themachinefolksession.org/tune/294. This was generated by a user on folkrnn.org in the key of C Dorian, who also gave it the title when they archived it. The user submitted a ‘setting’ which transposed it the key E Dorian, but otherwise was unchanged. This shows one limitation of folkrnn.org, which is that all transcriptions are generated in a variant of C (a consequence of an optimisation made while training the model on the corpus of existing tunes). It also shows that the manual editing features of themachinefolksession.org have been used by people to work around such a limitation.

Direct evidence of user intent can be seen in 'Rounding Derry' (https://themachinefolksession.org/tune/587). This user generated 'FolkRNN Tune №24807' on a fresh load of folkrnn.org, i.e., using default parameters and a randomised seed. The user played this tune twice, and then selected the ABC phrase 'C2EG ACEG|CGEG FDB,G,' and entered this as initial ABC. The user generated the next iteration, played it back, named it and archived the result on themachinefolksession.org. There, the user writes, *“Generated from a pleasant 2 measure section of a random sequence, I liked this particularly because of the first 4 bars and then the jump to the 10th interval key center(?) in the second section.”*

Taking themachinefolksession.org as a whole during the first 235 days of activity, 551 tunes were archived, of which 80% were generated by folkrnn.org. Of these 551 tunes, 15% have had further iterations contributed, with some tunes having more than one. These two websites continue to document human-machine co-creations. As of August 2019, themachinefolksession.org currently hosts 65 recordings in total, and 731 tunes; and folkrnn.org has generated a total of 35,249 transcriptions.

3 “Let’s Have Another Gan Ainm”: An Experimental Traditional Album

We recorded a 45-minute album [18] at the Visconti studio, Kingston University, in January 2018 with a team of professional musicians. The challenge was to make an album that could be considered successful as an album of Irish traditional dance music. To do this, we hired Daren Banarsë [19], a musician we have worked with in several concerts to perform AI-generated material in real musical contexts. The symbolic representation used by our models does not capture the critical nuances of Irish traditional music, and so using experienced performing musicians is a good way to judge the suitability of output generated by our models. This album was an extension of our experience with the musicians in concert, and was aimed at reaching wider dissemination and within a context relevant to the specific domain from which our training data comes.

“Let’s Have Another Gan Ainm” contains 31 tunes, 20 of which come from material generated by our models [20]. Banarsë curated material from 100,000 transcriptions that we have assembled in 34 volumes (https://goo.gl/1rRmwL). In practice he only took material from six of those volumes, but made changes to all of them. Though they tend to be small edits [21], some changes are musically significant. Banarsë identified improving the musical flow as a major reason for his edits. Many changes are at link points: adding first and second endings to enable linking backward for repeats and forward to the second phrase; and changing the end of a tune for smoother transition to the next one. He also corrected some ‘mistakes’, e.g., a few bars with missing eighth notes (which also occur frequently in the training data).

Another aspect of Banarsë's editing is the balance between conformity to common patterns and the inclusion of unique or special features that will stand out in a tune. In some instances he reinforced repetition of patterns to improve the structure (e.g., in the B part of tune #2375, which led to the second *Gan Ainm* in the first track). In other cases he changed some notes to make the tune more special when he deemed it was too mundane. Figure 3 shows a transcription generated by a folkrnn model, and the changes Daren made to create the third *Gan Ainm 3* in track 3.

Figure 3 - The top staff shows the tune generated by a folkrnn model, and the bottom staff shows the changes made (in red) to create the thirds *Gan Ainm* in track 3 of "Let's Have Another Gan Ainm".

Daren explains the changes he made to the opening:

> Bars 1, 2 and 3 are each made up of a mini call and response -- 2 beats call, 2 beats response. I thought the 3rd response was too similar to bar 2, starting on a A, and not really seeing anything interesting. My rewritten response provides a mix between an inversion of the call, and a more interesting end to the 4 bar phrase.

Some additional changes happened in the recording session itself when a variant played by a musician was taken up by the others as it was judged to be better than the notated version (e.g. the end of the B part in the first *gan ainm* in track 3).

We released “Let’s Have Another Gan Ainm” in March 2018 with the following information:

> *During the Summer of 2017, three generations of the Ó Conaill family gathered at the family home in Roscommon to celebrate the life and legacy of Dónal Ó Conaill. The late father and grandfather to the Ó Conaill family, Dónal was quietly dedicated to the tradition, and known for collecting local tunes without names which he passed on to his family. His daughters, Caitín and Ùna, are joined by their children and family friends to make a recording of the best of these tunes, along with some of Dónal’s personal favourites.*

We disguised the role of the computer in order to garner reactions and opinions about the album and not the technology [22]. In some circumstances, reactions could be positively biased when the result sounds better than the listener thought possible for a machine. In other circumstances, people could be prejudiced about music created by machines. The latter is clearly evinced by comments made on a Daily Mail article about our work [23]. The journalist embedded a brief excerpt of computer-generated traditional music. Reader comments ranged from negative (“Until they find a way to inject heart and soul into a computer it won't happen.”, “Totally lifeless without warmth.”) to hostile (“Stupid idea, stupid outcome.“ “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.”). In fact, the journalist accidentally excerpted a real tune, but many commenters heard a “robotic Irish Jig”.

Reviewers of our album were positive and clearly heard the music sitting comfortably within the tradition from which the training data comes. Referring to the backstory we posted, One reviewer [24] wrote: “Caitlín and Ùna Ó Conaill and her families and friends have done lovers of Irish traditional music an immense favour by allowing us this snapshot of a family reuniting to make delightful music.”

When we contacted the reviewers to reveal the true nature of the album no one reacted negatively. We received interesting comments from one expert, Kevin McDermott, having listened to the album again after we revealed the true story. He still considered most of the tunes believable, some of them very successful while identifying two as odd or failures. In his email to us he could relate specific tunes to different sub-domains of the tradition. Referring to track 6 in the album “the ascent to the high note in the turn sounds like stuff young composers like the lads in Socks In The Frying Pan are writing”; and on track 10 “the second [gan ainm in the set] is spot-on: a fine traditional jig which bears all the hallmarks of one from the late 18 to the mid-19C”.

This process primarily engaged experts in the particular style of music on which the model was trained. They contributed from the expertise as performers/arrangers and as reviewers of the final outcome. We can see that while our models are generally rather successful for the style they can be improved by: (1) better handling of particular local context such as the musical meaning of 1st and 2nd ending; and (2) helping human users hone in on outputs suitable for their needs. We can also see that experts make finer distinctions about different parts of this musical corpus. A data collection of 23,000 examples is rather large compared to similar works (such as building a model on the 371 chorale harmonizations of Bach), which helps the relative success of the model. But at the same time perhaps finer nuances are also lost by aggregating all the transcriptions together.

4 Going beyond the traditional context of folkrnn

Curating and editing generated transcriptions – as Banarsë did in creating "Let's Have Another Gan Ainm" and other musicians we engaged in various concerts – is one mode of using our models, but working interactively with them can move several steps away from their roots in folk traditions. In his piece *Bastard Tunes*, Ben-Tal[15] used the different generation parameters to pull it away from the conventions of the tradition and used the result as pre-compositional material.

The parameters available for controlling the generation process do not have a direct and easily predictable effect on the output. Setting the mode and meter have the obvious stated outcome but also some less obvious ones since there are fewer examples for the model to learn from for 9/8, and even fewer for 9/8 in Mixolydian mode. The model is also highly non-linear which means that the interaction between the different initialisation parameters is opaque, but sometimes creatively fruitful. The temperature parameter has the most obvious effect: very low temperature (such as 0.1) will mostly yield very repetitive sequences. High temperature can have dramatic effects and can easily seem like a parody of "new music". Since increasing the temperature flattens the output distribution, at the theoretical limit all symbols become equally likely and equal. Changing the seed allows for re-generating from the same initial settings, which is also a useful compositional tool.

Steering the generative process using these parameters to produce outputs that the composer judges as useful is not straightforward. While his initial interaction with the system was mostly trial and error, after generating many hundreds of outputs (and discarding the vast majority of them) Ben-Tal felt he was able to steer the process in directions that he found compositionally useful. This turned out to be mostly through initialising the generation process with combinations that are uncommon in the original data. These include the less common meters and modes, non-modal opening sequences, or even just long notes or rests (which are rare in these dance-based tunes) His pre-composition process became an interactive search for regions of the model's creative space where the stylistic conventions modelled through the data are sufficiently weak but not entirely erased.

Like in any creative work, what is useful is personal rather than rule bound. Co-creating with folkrnn is an act of imagination as well as iterative generation of transcriptions. This push and pull between the composer and the system can lead to new discoveries for the composer. For instance in bars 143-145 of the first movement of *Bastard Tunes* (fig. 3) the higher temperature settings led the model to produce a “Jazzy” moment. The ensuing composition process involved identifying this material and choosing to bring it out in the piece. Another composer might have found it out of place and decided to delete it or obscure it instead. The idea of composing with external constraints is, of course, not new or groundbreaking. But, as these bars illustrate, the constraints imposed by the system are not arbitrary but grounded in music. While the AI system only captures a limited aspect of musical practice, it still learned from traces of human musical activity.

Figure 4: Bars 143-147 from the first movement of Oded Ben-Tal’s *Bastard Tunes*. Note the surprisingly Jazzy part produced in the generation process and given to the piano.

To further stimulate interest in our models, over the summer of 2018 we organised a composition competition. Submissions included both a score for a set ensemble (flute, clarinet, violin, cello, piano) and an accompanying text describing how folkrnn.org contributed to the composition of the work. The judging panel – the first author was joined by Prof. Elaine Chew and Prof. Sageev Oore – considered the musical quality of the submission as well as the creative use of the model. The winning piece, *Gwyl Werin* by Derri Lewis, was performed by the New Music Players at a concert organised in partnership with the 2018 O’Reilly AI Conference in London. Lewis said he didn’t want to be ‘too picky’ about the tunes, but rather selected a tune to work from after only a few iterations with folkrnn.org. He did not use the tune as a melodic line directly in its generated form in the piece. Rather he describes treating the generated tune as a tone row and composing harmonic, melodic and motivic material out of it.

Both composers used folkrnn in a manner consistent with what Lubart described as “computer as pen pal” [25]. The process is still one-sided in this case: the computer generating ideas and the composer choosing, modifying, or asking for new ideas. One possible improvement of music AI tools would be to turn this into an interactive process where the computer can evaluate the individual choices and adapt what is offered in response. Though there is no simple or straightforward way of implementing this functionality.

5 Conclusions

Given its successes, AI will continue to be applied to and impact the domain of music. Our work demonstrates that there is an audience willing to engage with music AI. Both professional and amateur musicians found ways of including the models we developed in their musical activities. The web interface we deployed in folkrnn.org is a friendlier user interface that running the computer code directly. However, as we learn more about what different users find useful (and not) we aim to improve the usability of our system. We see evidence of interactively searching for creative ideas with the system – through the iterative process of generating and altering parameters. As Lomas [26] observed, the aims of the creative search in such circumstances include exploring the conceptual space, identifying fruitful locations, refining ideas, and seeking novelty. Translating his methods from the visual to the audio domain and from numeric to symbolic data is not straightforward. However, one of our future aims is to develop an ‘artificial critic’. We do not envision an aesthetic evaluator but rather an assistant that could facilitate the exploration of the creative space of the algorithm. Using similarity ratings (though this is not a trivial question for music [27]) could help a user map out the different regions of the space and hone in on more relevant ones (to them). Conversely dis-similarity can be leveraged when a user decides to look for novelty or contrast. The assistant could also filter out completely unacceptable outputs based on users’ input for the immediate task, for example, by building a ‘stylistic conformity’ sieve which allows a more nuanced control of the model’s adherence to the conventions in the training data. Of course, individual users may prefer outputs that conform or those that deviate from the style of the training set.

More broadly, it is imperative that creative AI researchers engage more thoroughly with a variety of practitioners. AI has the potential to augment human creativity and we believe such co-creative approach is more fruitful than a focus on replicating (and thereafter replacing) human creativity. Such an approach to AI development is more fruitful not only for the domain of artistic creation but also for the AI researchers. For instance, creative interrogation of our system (see section 3.3 ‘Nefarious Testing’ of [11]) revealed that the ‘intelligence’ of our AI system is rather shallow. Making a technology accessible to a wider audience can also reveal new avenues for development. At the same time demonstrating the co-creative potential of AI will also help allay some fears of this new technology. The human-vis-machine narrative makes good headlines but fuels the fear that machines will take over the world.

Acknowledgments: AHRC UK project no. AH/R004706/1 “Engaging three user communities with applications and outcomes of computational music creativity”

References and Notes

1. Ariza, C. (2009) The interrogator as critic: The Turing test and the evaluation of generative music systems. *Computer Music J.* 33(2): 48–70.
2. Yang, L.-C. and Lerch, A. (2018) On the evaluation of generative models in music. *Neural Computing & Applications*.
3. Agres, K., Forth, J., and Wiggins, G. A. (2016). Evaluation of musical creativity and musical metacreation systems. *Computers in Entertainment*, *14*(3).
4. Wagstaff, K. L. (2012). Machine learning that matters. In *Proc. Int. Conf. Machine Learning*, pp 529–536, Edinburgh, Scotland.
5. Dannenberg, R. B., Thom, B., and Watson, D. (1997). A machine learning approach to musical style recognition. In *Proc. Int. Computer Music Conf.*, pp 344–347.
6. Pearce, M., Meredith, D., and Wiggins, G. (2002). Motivations and methodologies for automation of the compositional process. *Musicae Scientiae*, 6(2): 119–147.
7. Dubnov, S., Assayag, G., Lartillot, O., and Bejerano, G. (2003). Using machine-learning methods for musical style modeling. *Computer,* 36(10): 73–80.
8. Nierhaus, G. (2008). Algorithmic Composition: Paradigms of Automated Music Generation. Springer.
9. Fernández, J. D. and Vico, F. (2013). AI methods in algorithmic composition: A comprehensive survey. *J. Artificial Intell. Res.,* 48(1): 513–582.
10. Herremans, D., Chuan, C.-H., and Chew, E. (2017). A functional taxonomy of music generation systems. *ACM Computing Surveys*, 50(5): 1–30.
11. Sturm, B. L. and Ben-Tal, O. (2017) Taking the models back to music practice: Evaluating generative transcription models built using deep learning. *J. Creative Music Systems*, vol. 2.
12. Sturm, B. L. Santos, J. F. Ben-Tal, O. Korshunova, I. (2016) Music transcription modelling and composition using deep learning. in *Proc. Conf. Computer Simulation of Musical Creativity*, Huddersfield, UK.
13. http://abcnotation.com/
14. Mossmyr, S. Hallstrom, B. Sturm, B. L. Vegeborn, V. H. and Wedin, J. (2019) From Jigs and Reels to Schottisar och Polskor: Generating Scandinavian-like Folk Music with Deep Recurrent Networks. In *Proc. Sound and Music Computing,* Malaga, Spain.
15. Sturm, B. L. Ben-Tal, O. Monaghan, U. Collins, N. Herremans, D. Chew, E. Hadjeres, G. Deruty, E. and Pachet, F. (2018) Machine learning research that matters for music creation: A case study. *J. New Music Res*. 48(1): 36-55.
16. https://www.heise.de/newsticker/meldung/Missing-Link-Musik-ohne-Musiker-KI-schwingt-den-Taktstock-4224798.html

17. An approximate figure derived engagement metrics of the project's GitHub page https://github.com/IraKorshunova/folk-rnn
18. https://soundcloud.com/oconaillfamilyandfriends
19. http://www.darenbanarse.com
20. "Gan Ainm" translates to "no name" from Gaelic.
21. Sturm B. L. and Ben-Tal, O. (2018) Let's Have Another Gan Ainm: An experimental album of Irish traditional music and computer-generated tunes. *Tech. Report KTH.* http://kth.diva-portal.org/smash/record.jsf?pid=diva2%3A1248565&dswid=7147
22. Ethics approval for this deception was granted by Kingston University research ethics panel. See footnote 21.
23. https://www.mirror.co.uk/tech/bot-dylan-computer-using-artificial-10504774
24. Harley, D. Ó CONAILL FAMILY AND FRIENDS – Let's Have Another Gan Ainm (Digital release) Retrieved Aug. 19 2019 from https://folking.com/o-conaill-family-and-friends-lets-have-another-gan-ainm-digital-release/
25. Lubart, T. (2005). How can computers be partners in the creative process: classification and commentary on the special issue. *Int. J. Human-Computer Studies*, *63*(4-5): 365-369.
26. Lomas, A. (2018). On hybrid creativity. *Arts*, 7(3): 25
27. Flexer, A. and Thomas G. (2016). The problem of limited inter-rater agreement in modelling music similarity. *J. New Music Res.* 45(3): 239-251.

Biographies
Oded Ben-Tal is a composer and senior lecturer in music at Kingston University (UK). Ben-Tal studied composition at the Rubin Academy of Music in Jerusalem, followed by doctoral studies at Stanford University with Jonathan Harvey and Brian Ferneyhough. In addition to working with AI, he regularly uses his own intelligence to compose music. His music was featured in international festivals - Diffrazione in Florence, The New York City Electroacoustic Music Festival, and ME_MMIX in Palma, Majorca, and performed by musicians such as Matthew Barley, the New Music Players, and Plus-Minus ensemble.

Matthew Tobias Harris researches audiences and interaction, prompted by his art practice. For this project, he was a Postdoctoral Researcher at Queen Mary University of London, where he also helped teach a robot to perform stand-up comedy, and taught design to computer science and psychology students.

Bob L. Sturm is the principal scientist on the folkrnn project. He has been an enthusiast of Irish traditional music since living in Limerick, Ireland during the summer of 2000. Bob is an Associate Professor of Computer Science in the Speech, Music and Hearing research division of the KTH Royal Institute of Technology in Stockholm. His research is focused on making computers work "intelligently" with sound and music data. He also plays in sessions in Stockholm, where he runs a Learners' Session (that sometimes includes "machine folk").