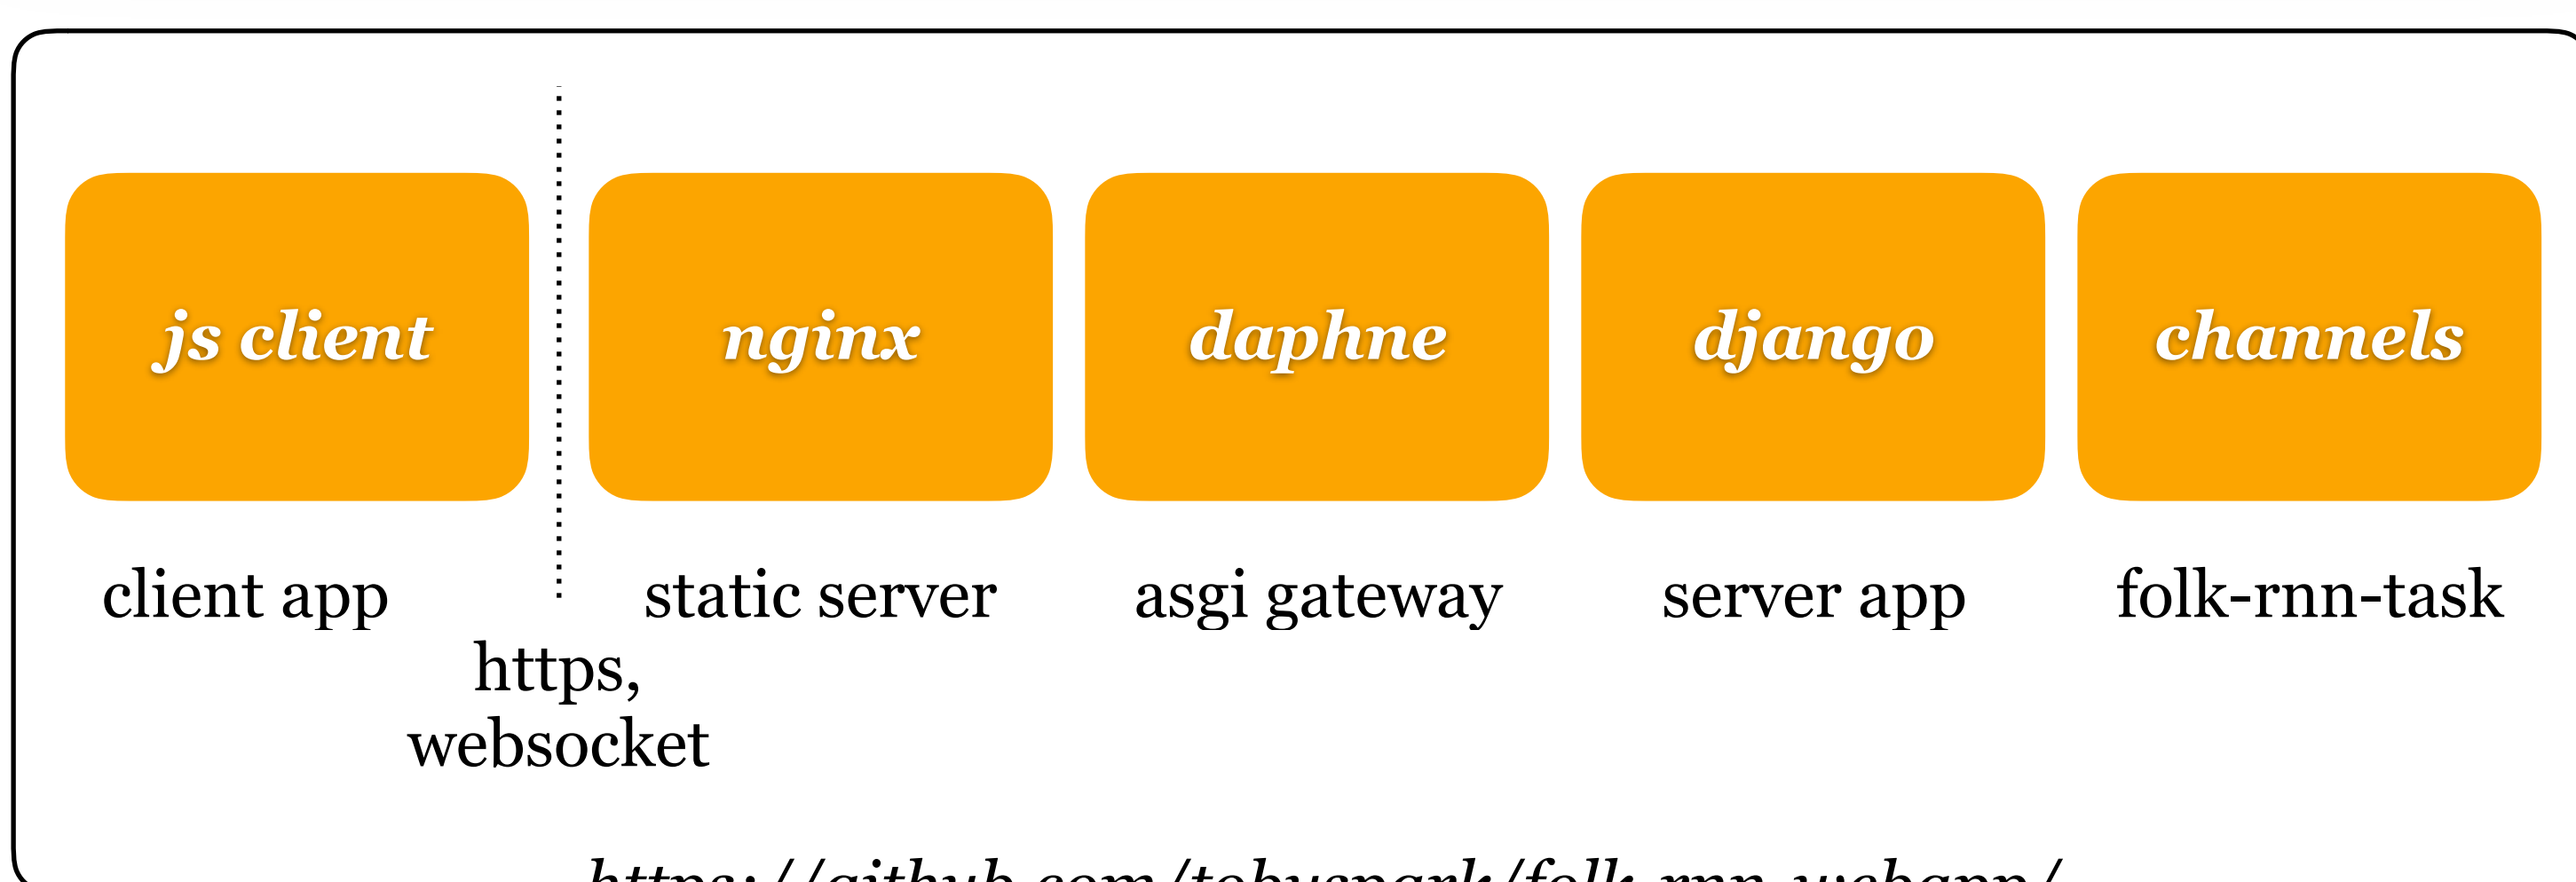
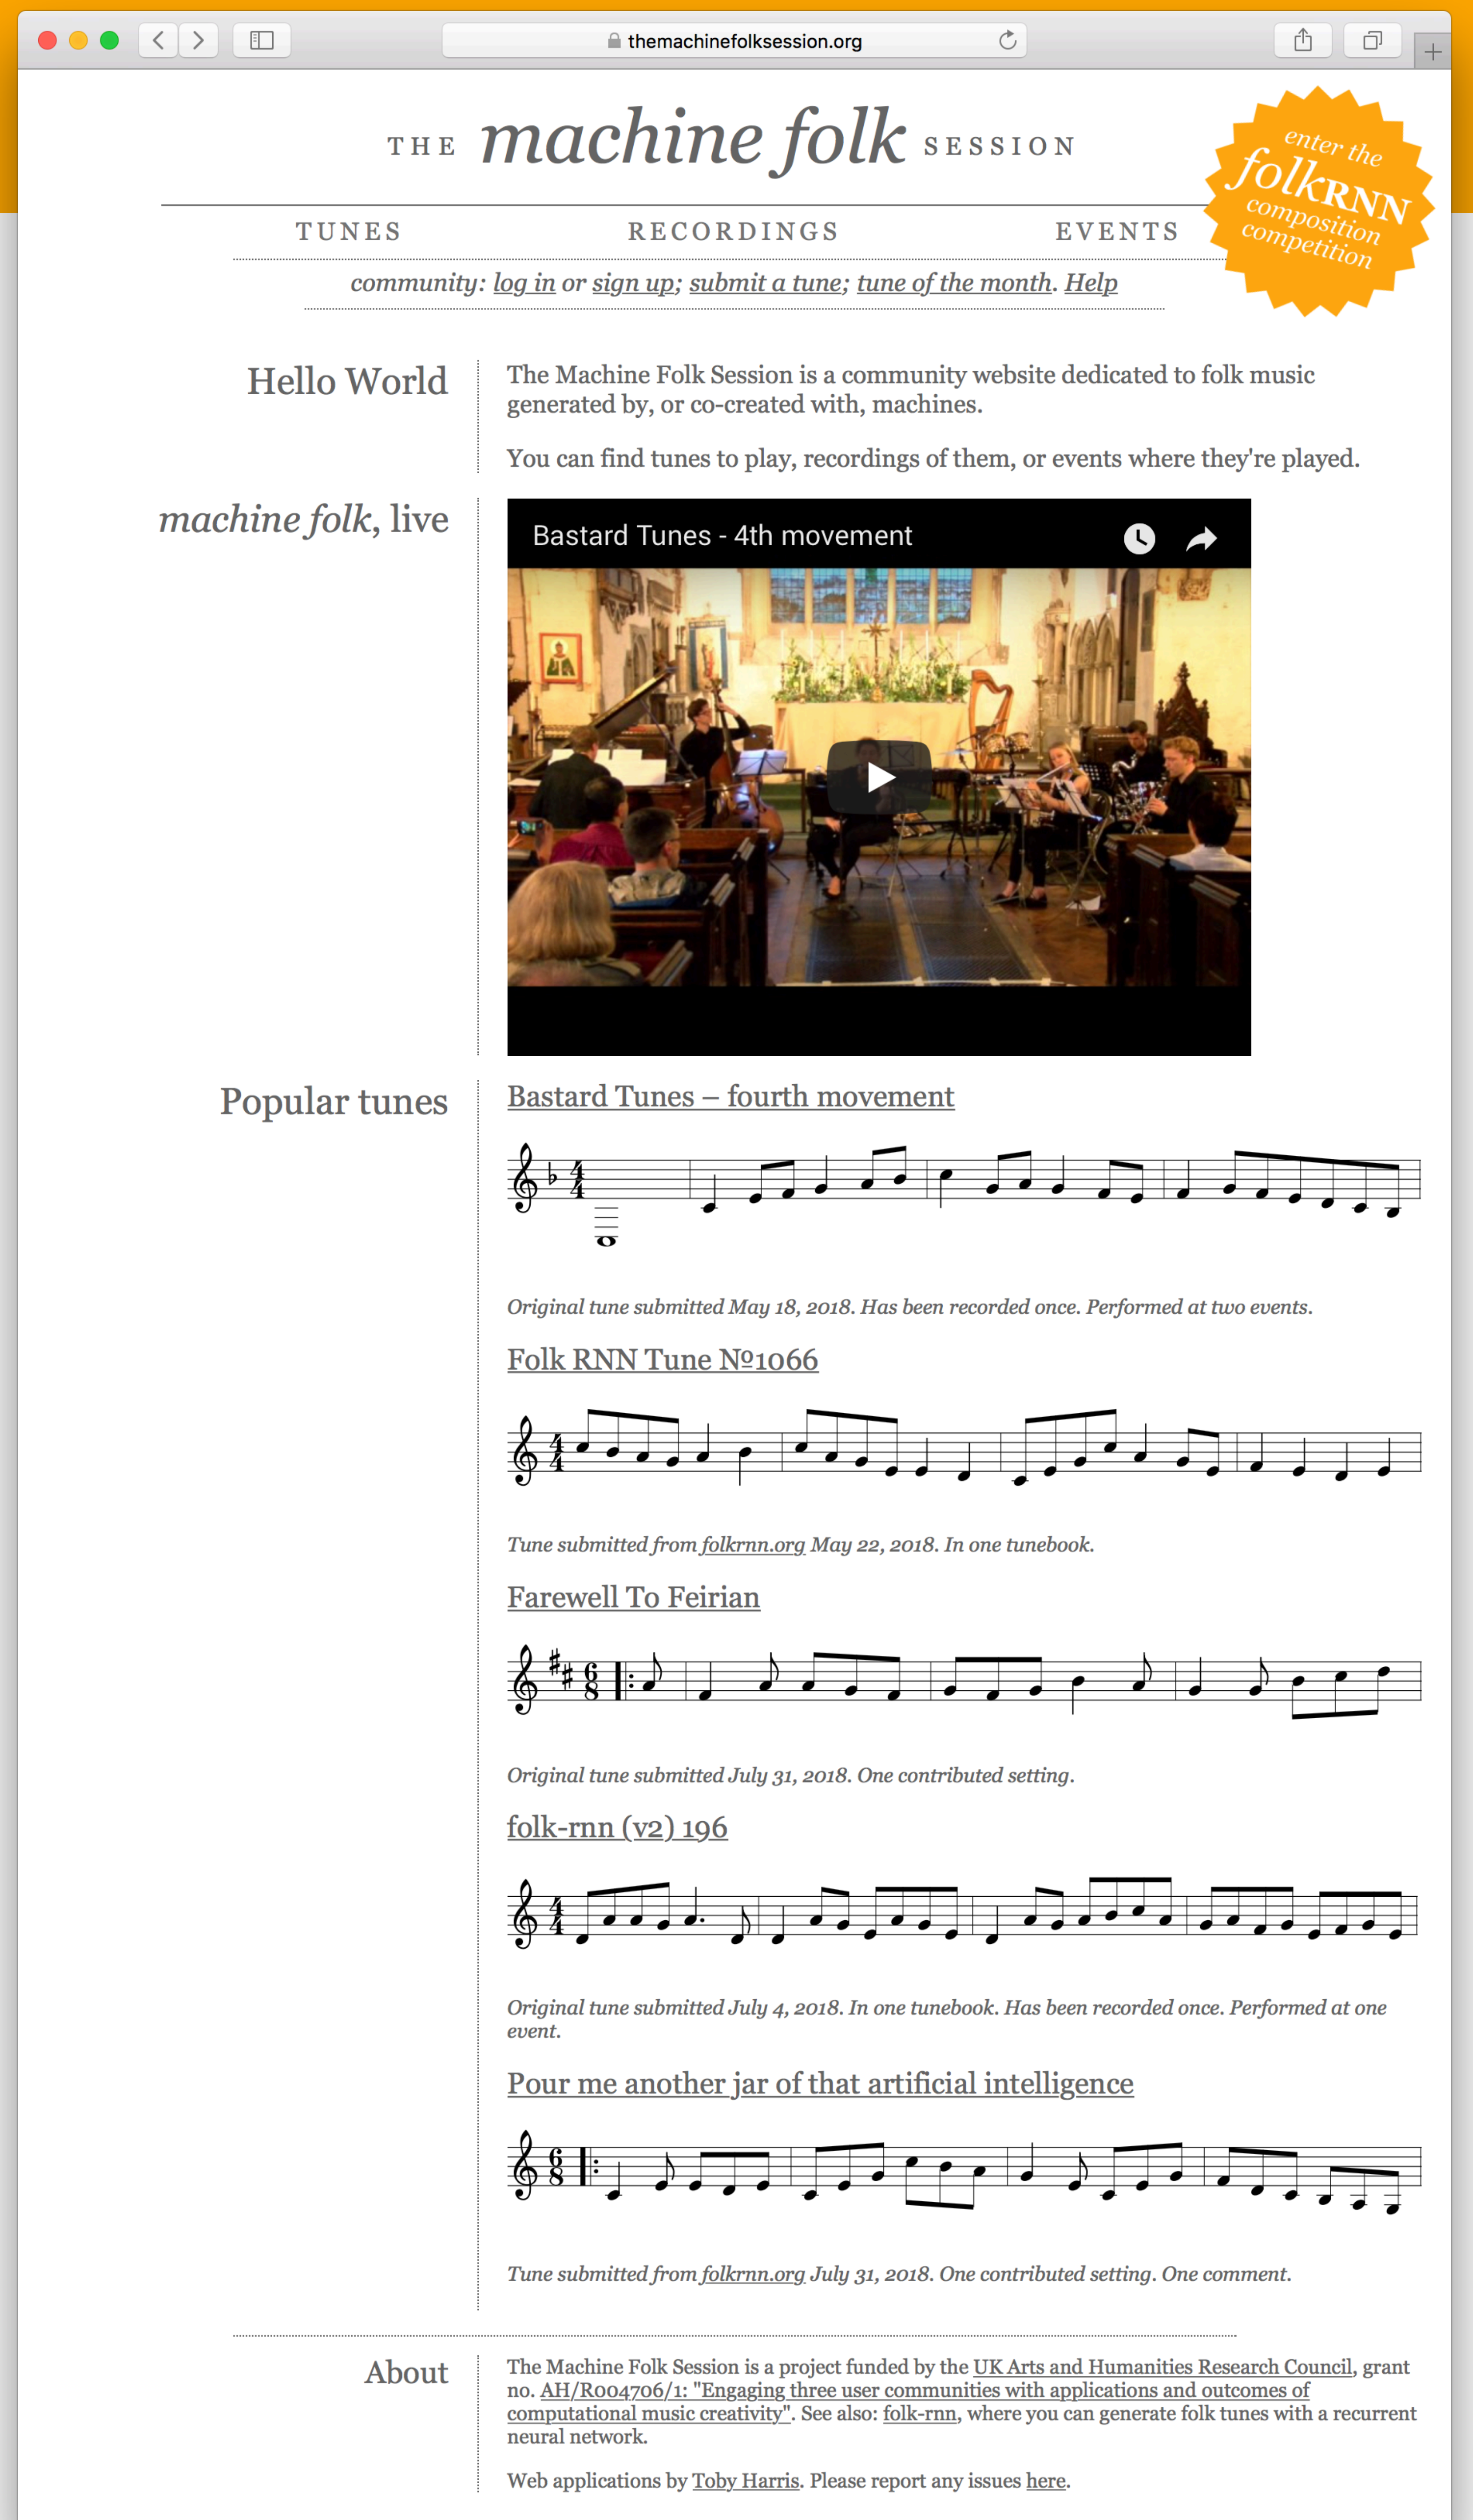
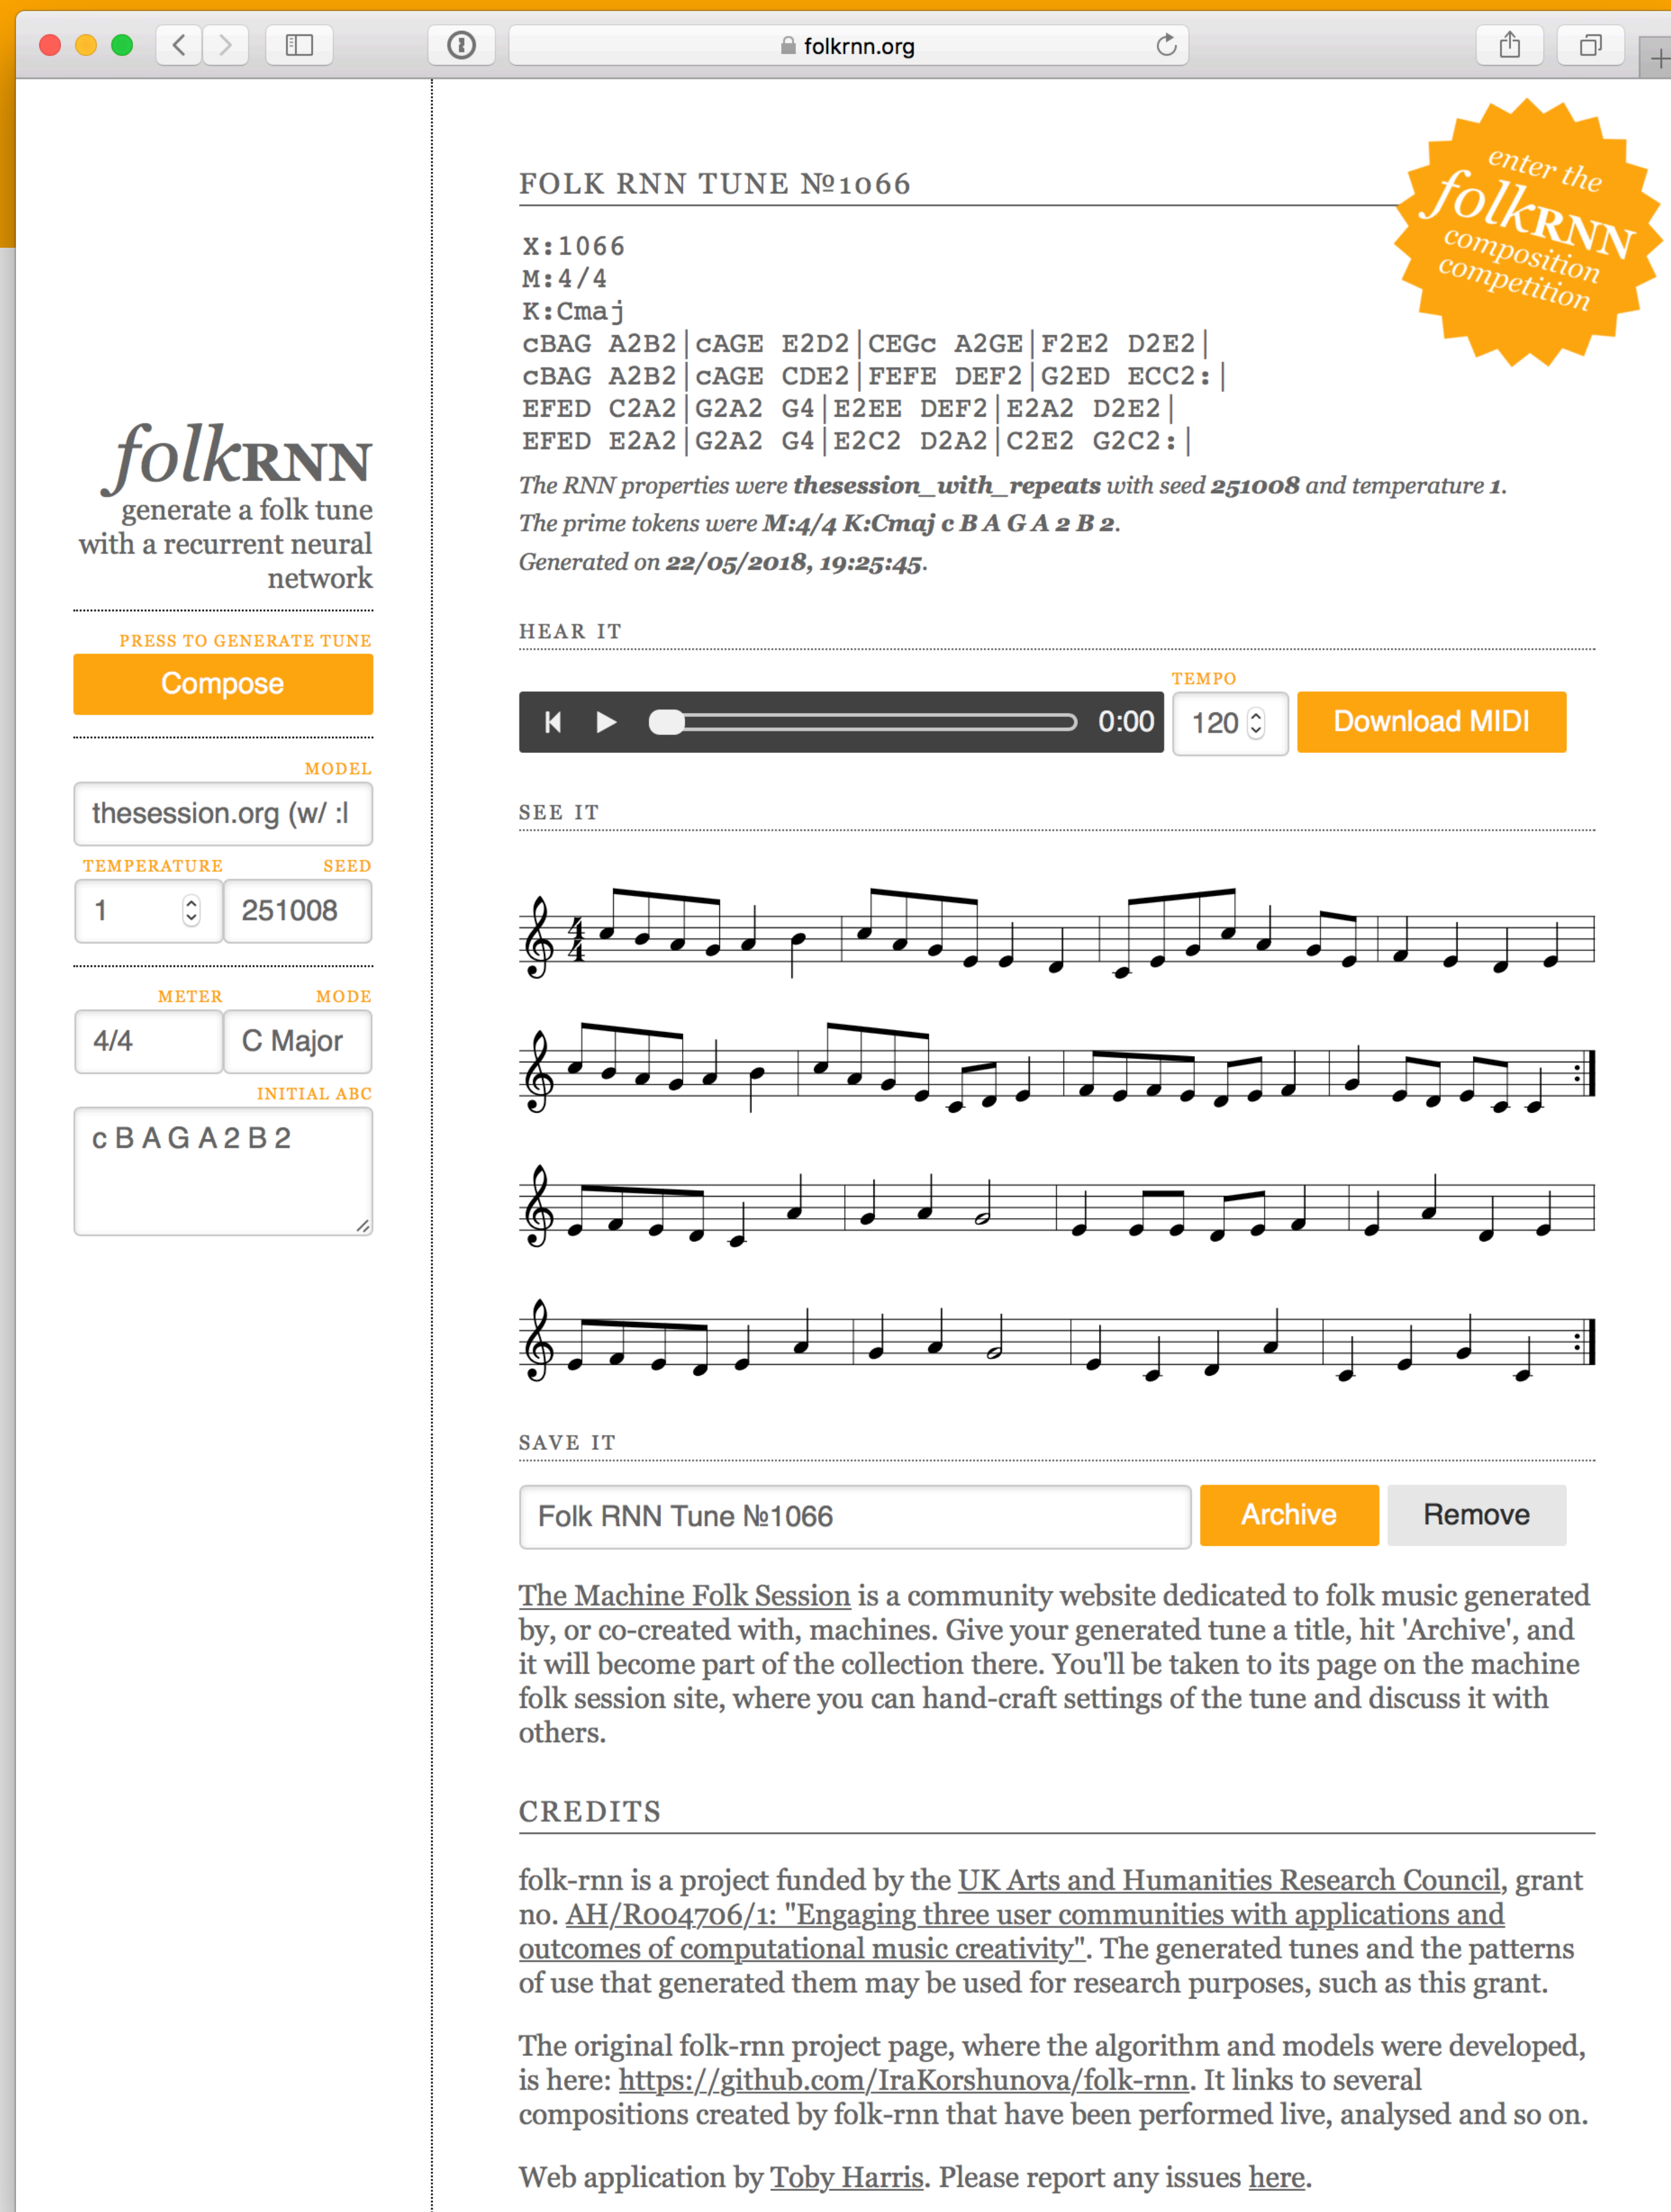


folkRNN

generate a folk tune with a recurrent neural network



- 1 <https://folkrrnn.org> – live since March 2018
A web-based implementation of a generative machine learning model trained on transcriptions of folk music from Ireland and the UK.

- 2 <https://themachinefolksession.org> – live since June 2018
An online repository of work created by, or co-created with, machines.

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, 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. The specific kind of model we use consists of three hidden layers of long short-term memory (LSTM) units. We trained this model on over 23,000 transcriptions crowd-sourced from an online community devoted to these kinds of folk music (<http://thesession.org>). Several compositions created with our application have been performed so far, and recorded and posted online. We have also organised 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.

This work is funded by the UK Arts and Humanities Research Council, grant no. AH/R004706/1: "Engaging three user communities with applications and outcomes of computational music creativity".

Matthew Tobias Harris

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

Oded Ben-Tal

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

Bob L. Sturm

Royal Institute of Technology KTH
Lindstedtsvägen 24, SE-100 44 Stockholm, Sweden
bobs@kth.se