KTH ROYAL INSTITUTE
OF TECHNOLOGY

Tunes from the Ai Frontiers

Bob L. T. Sturm

Speech, Music and Hearing Division

School of Electronic Engineering and Computer Science

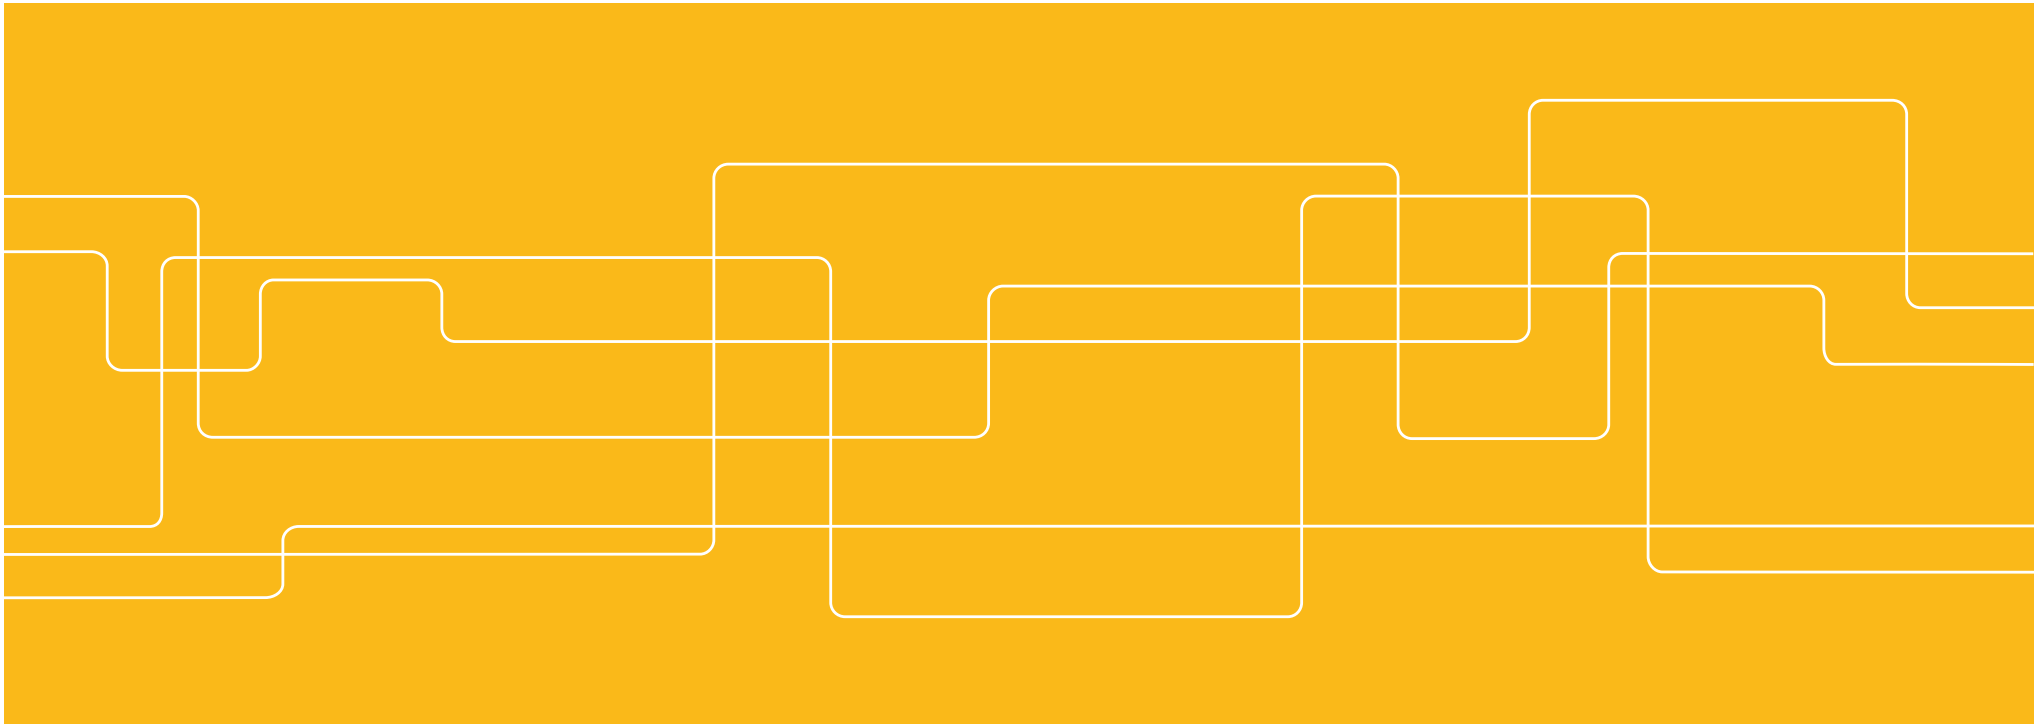

Machine learning is:

- Methods for making machines learn from data.
- One of the hotest areas of research in Ai today!
- Hugely successful due to a confluence of data, computational resources, advancements in efficient and effective algorithms, and massive amounts of capital.

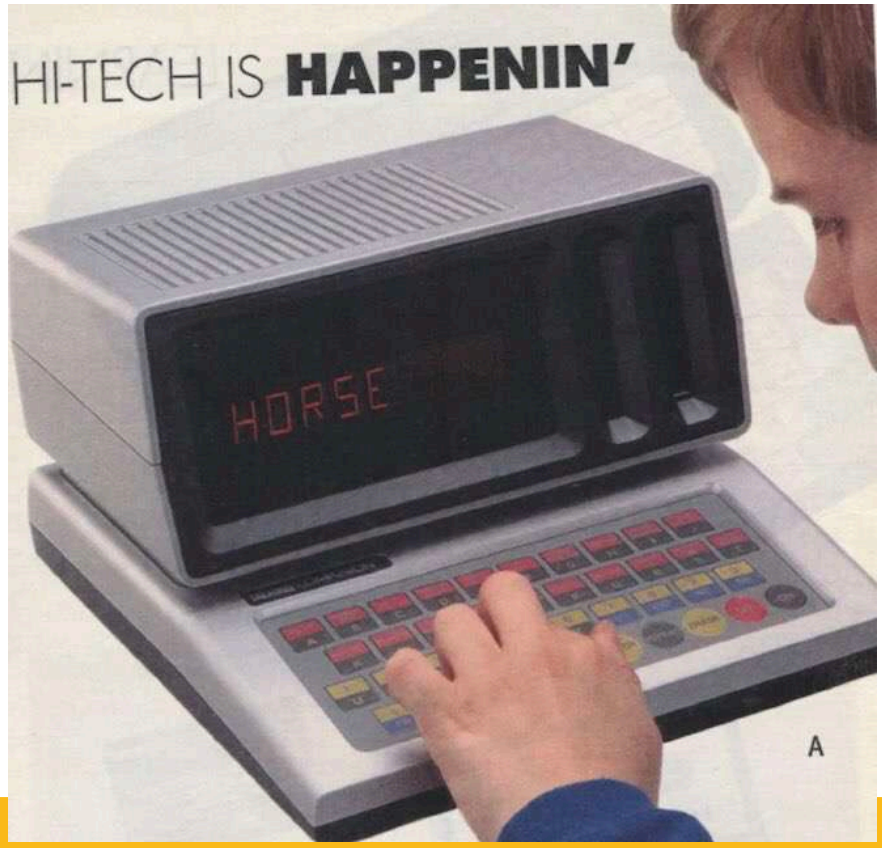

Folk music is:

- An historically motivated categorization of musical practice among “common” people sharing a regional identity
- Often accompanying other practices, such as dancing or story telling
- Passed aurally with provenance

Johnny O’Leary

Ai Folk?

Faux Folk?

trAiditional Music?

Machine Folk

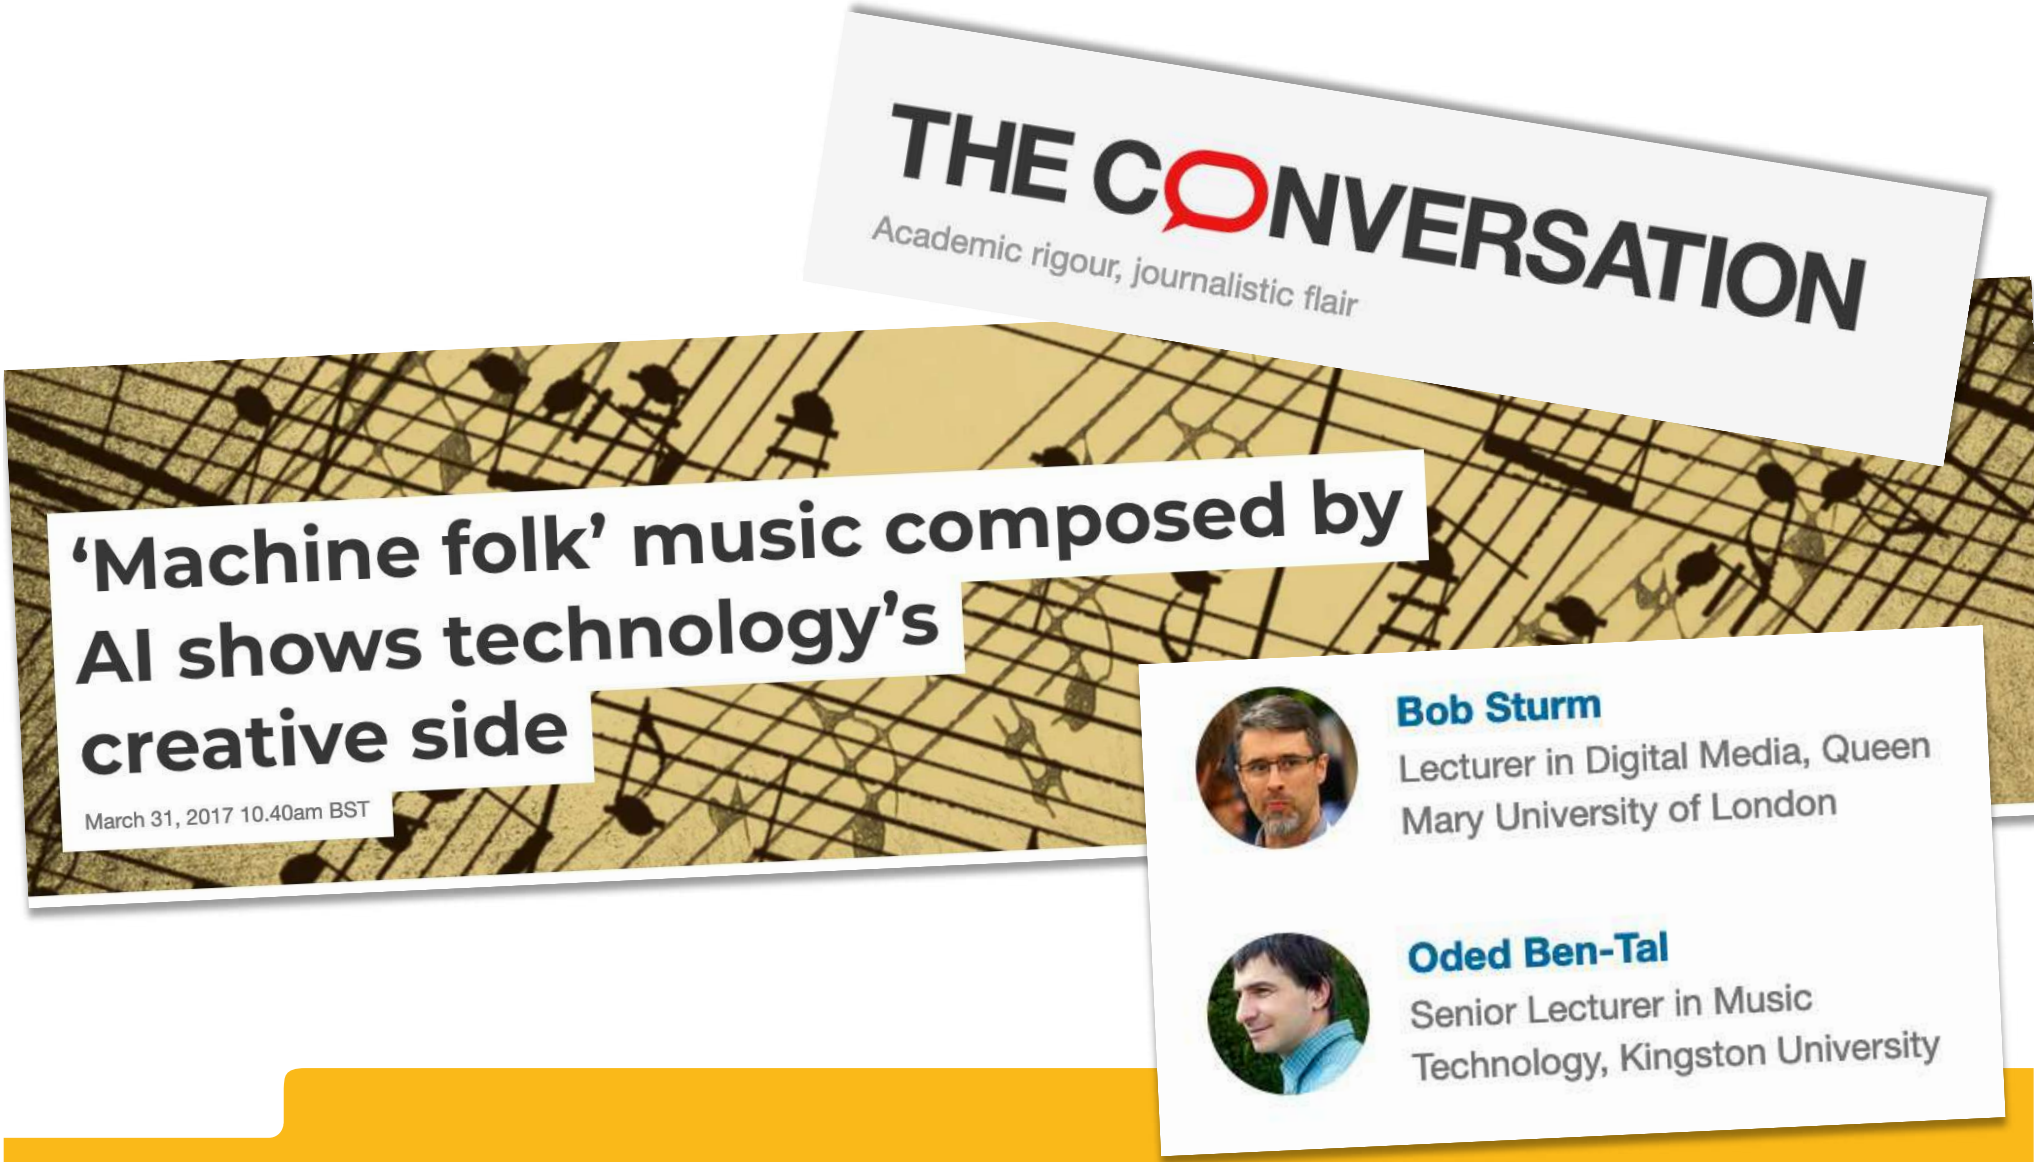

https://theconversation.com/machine-folk-music-composed-by-ai-shows-technologys-creative-side-74708

An Example: “The Boys of Ballinaburre”

```
<s> M:6/8 K:Edor |: B E E B E E | F E
D D E F | G E E B E E | B d e d B A |
B E E B E E | F E D D E F | G A B A F
D | E F E E 3 :| |: B e e e f e | d B
c d B A | B e e e d e | f d f d 2 e |
B e e e f e | d B c d B A | G A B d B
A | F E D E 3 :| </s>
```

An Example: “The Boys of Ballinaburre”

https://youtu.be/qcOZZpfSM_E

Some background

How did I get here?

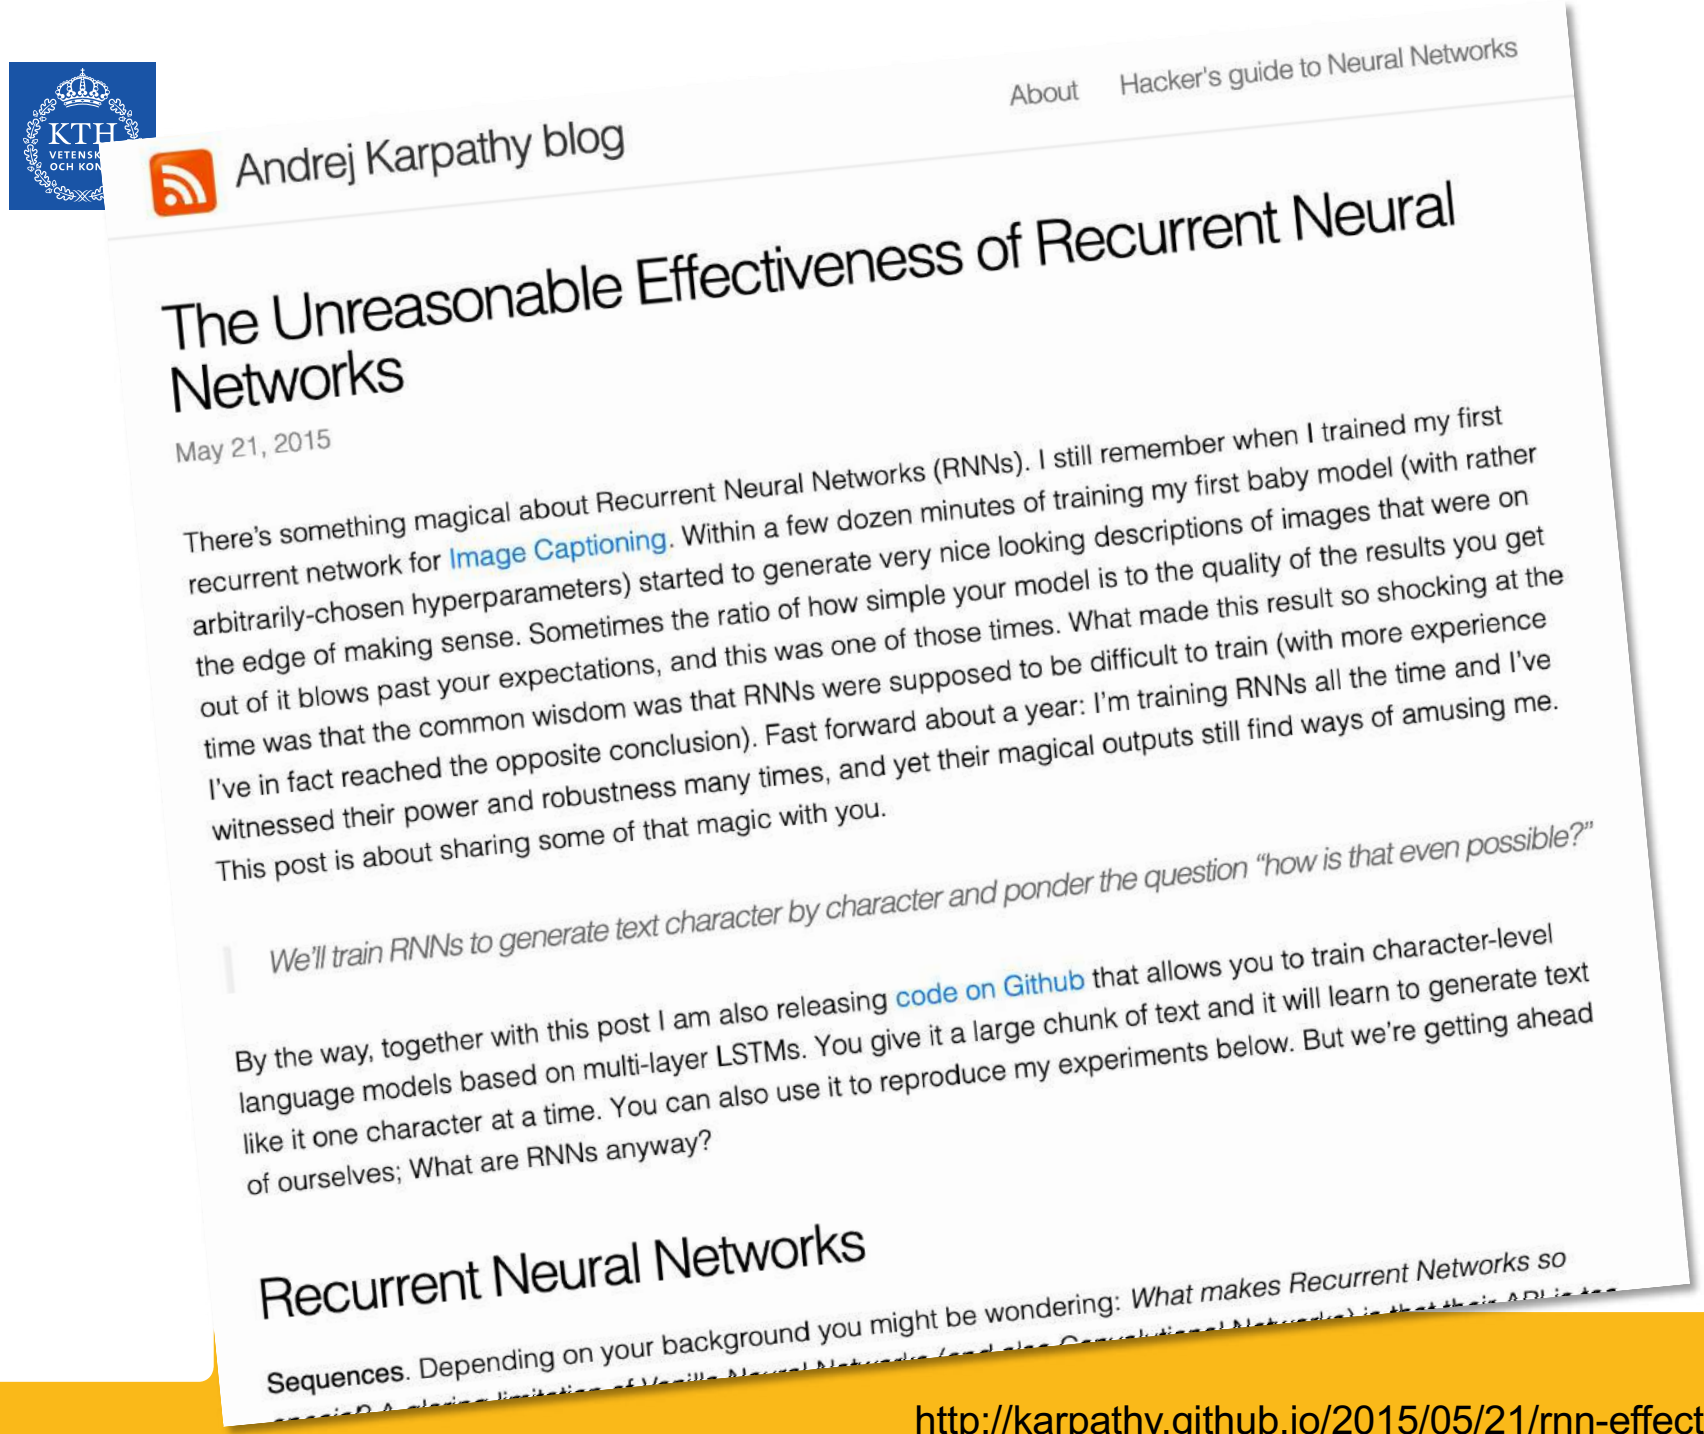

Andrej Karpathy blog

About Hacker's guide to Neural Networks

The Unreasonable Effectiveness of Recurrent Neural Networks

May 21, 2015

There's something magical about Recurrent Neural Networks (RNNs). I still remember when I trained my first recurrent network for Image Captioning. Within a few dozen minutes of training my first baby model (with rather arbitrarily-chosen hyperparameters) started to generate very nice looking descriptions of images that were on the edge of making sense. Sometimes the ratio of how simple your model is to the quality of the results you get out of it blows past your expectations, and this was one of those times. What made this result so shocking at the time was that the common wisdom was that RNNs were supposed to be difficult to train (with more experience I've in fact reached the opposite conclusion). Fast forward about a year: I'm training RNNs all the time and I've witnessed their power and robustness many times, and yet their magical outputs still find ways of amusing me. This post is about sharing some of that magic with you.

> *We'll train RNNs to generate text character by character and ponder the question "how is that even possible?"*

By the way, together with this post I am also releasing code on Github that allows you to train character-level language models based on multi-layer LSTMs. You give it a large chunk of text and it will learn to generate text like it one character at a time. You can also use it to reproduce my experiments below. But we're getting ahead of ourselves; What are RNNs anyway?

Recurrent Neural Networks

Sequences. Depending on your background you might be wondering: *What makes Recurrent Networks so*

http://karpathy.github.io/2015/05/21/rnn-effectiveness/

A basic recurrent network

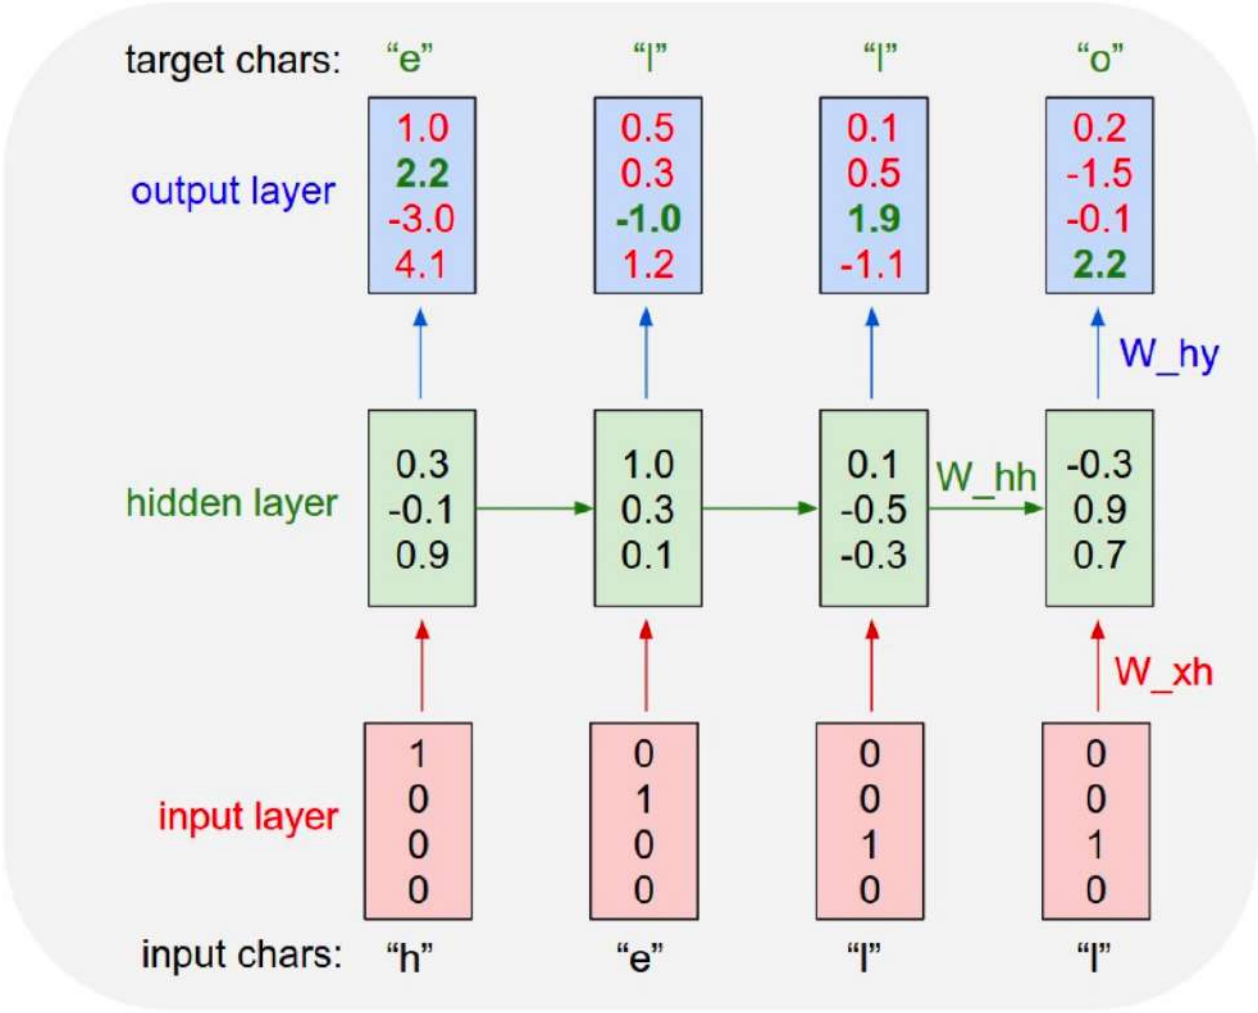

http://karpathy.github.io/2015/05/21/rnn-effectiveness/

http://karpathy.github.io/2015/05/21/rnn-effectiveness/

Source Code

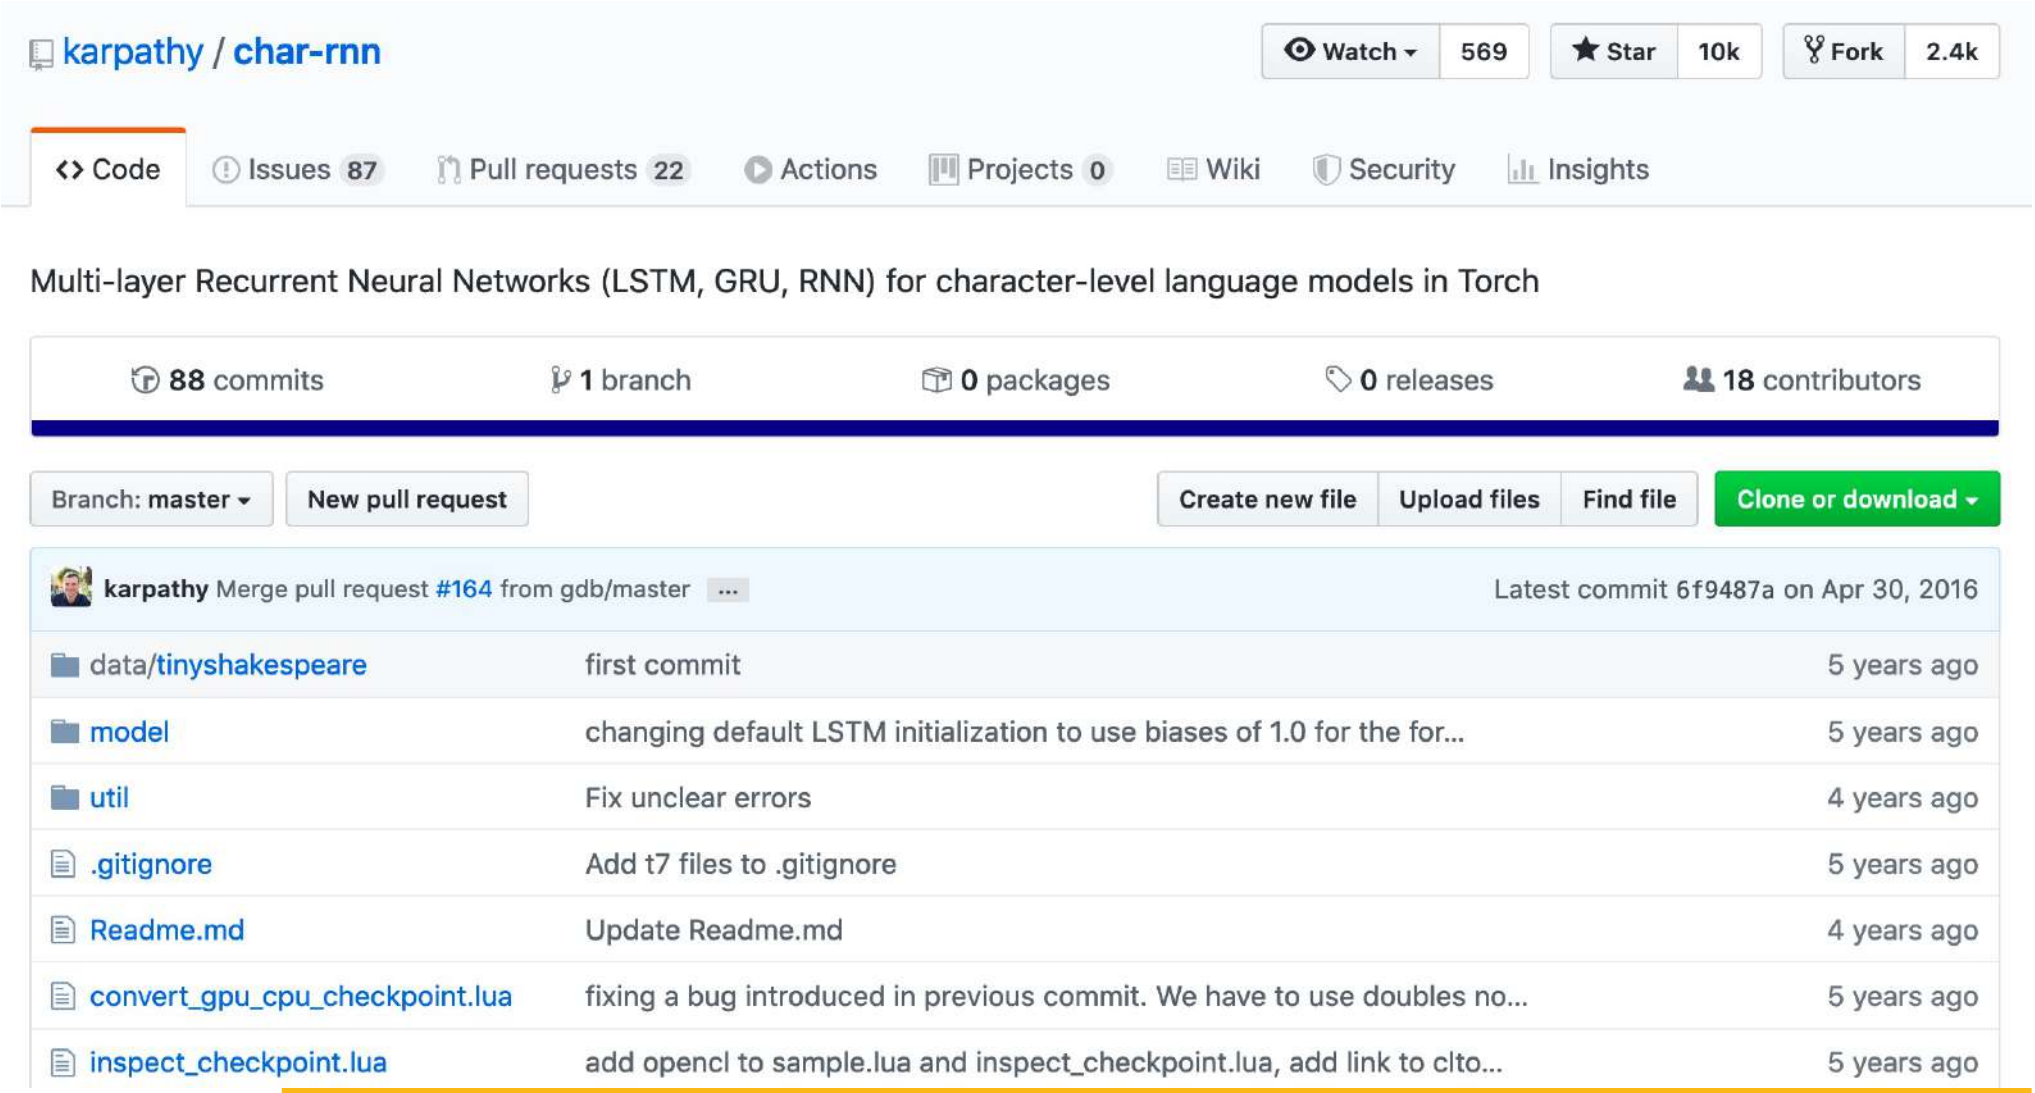

https://github.com/karpathy/char-rnn

First experiments (1180 tunes)

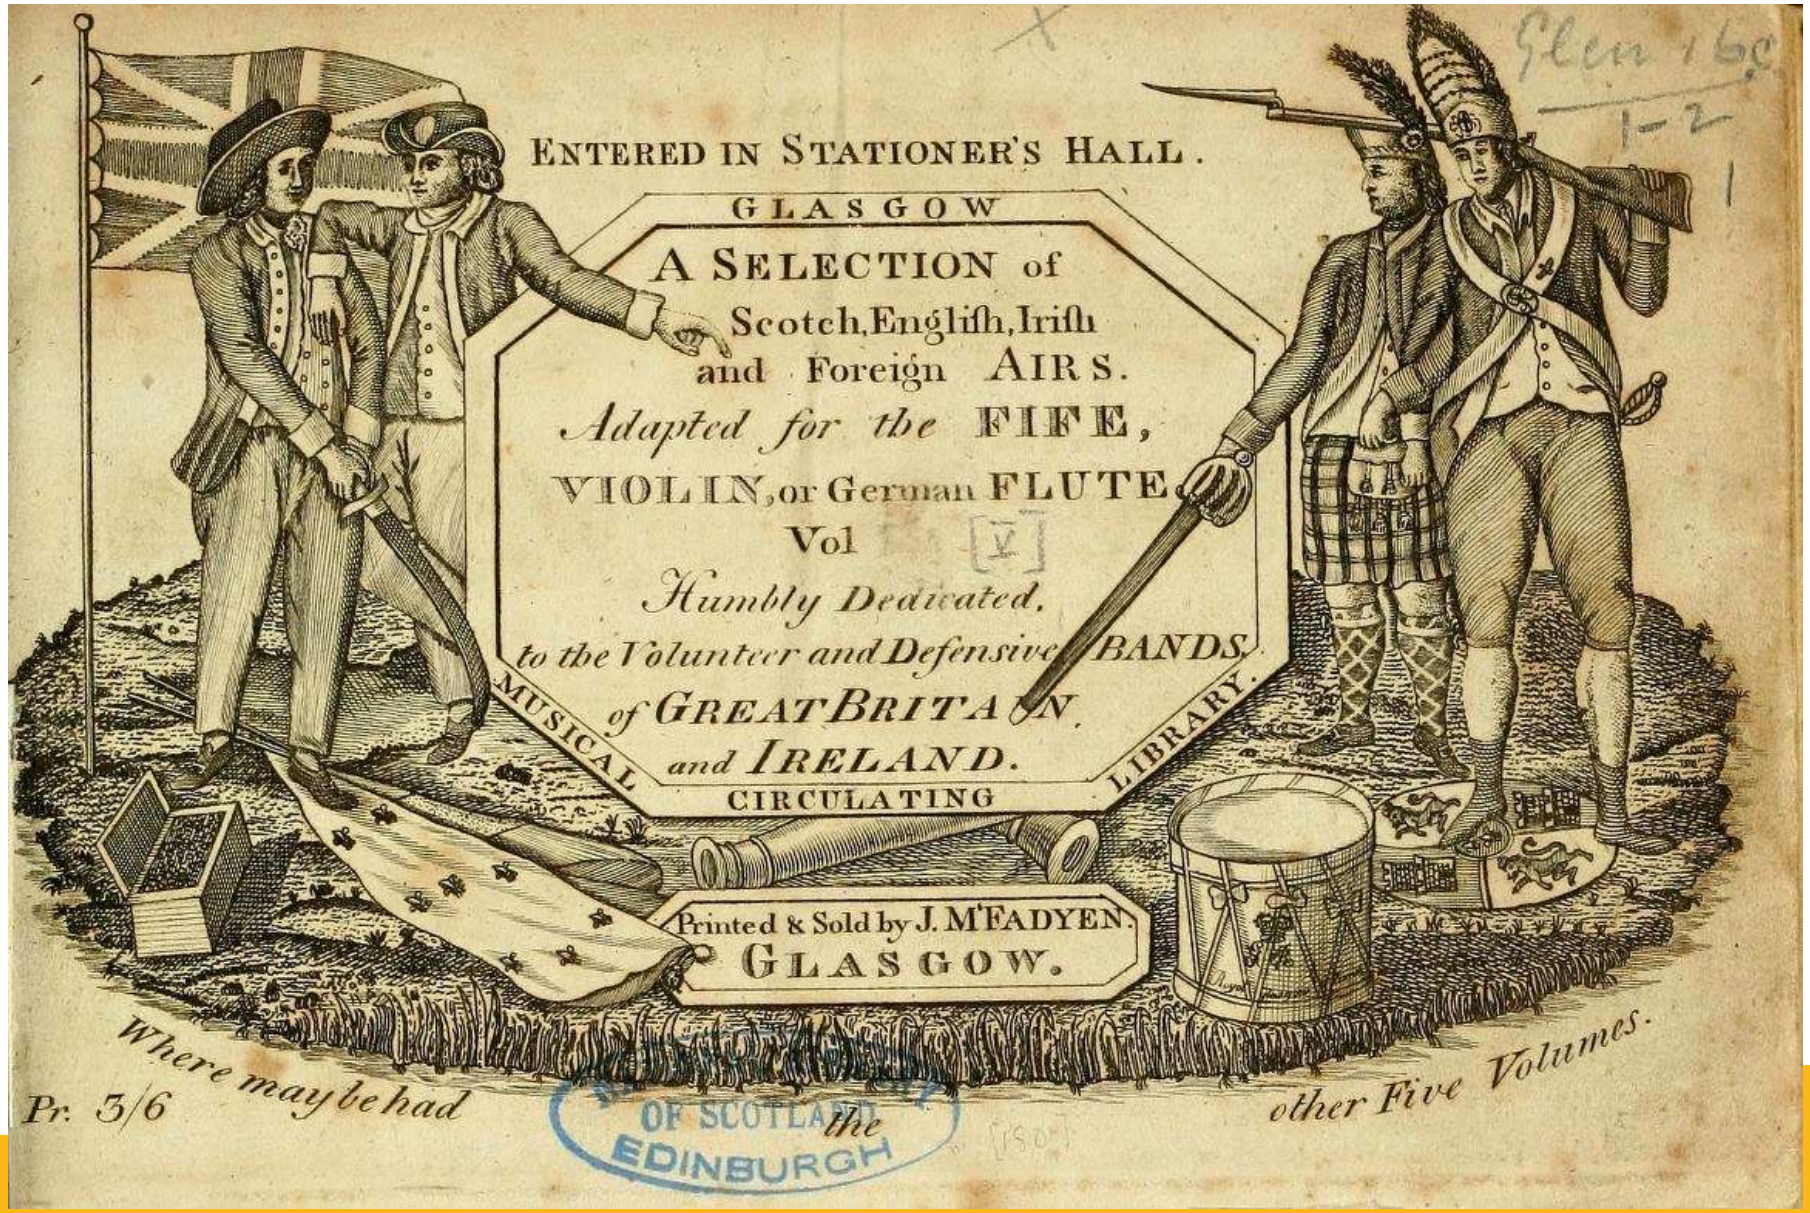

folk-rnn (v1): architecture

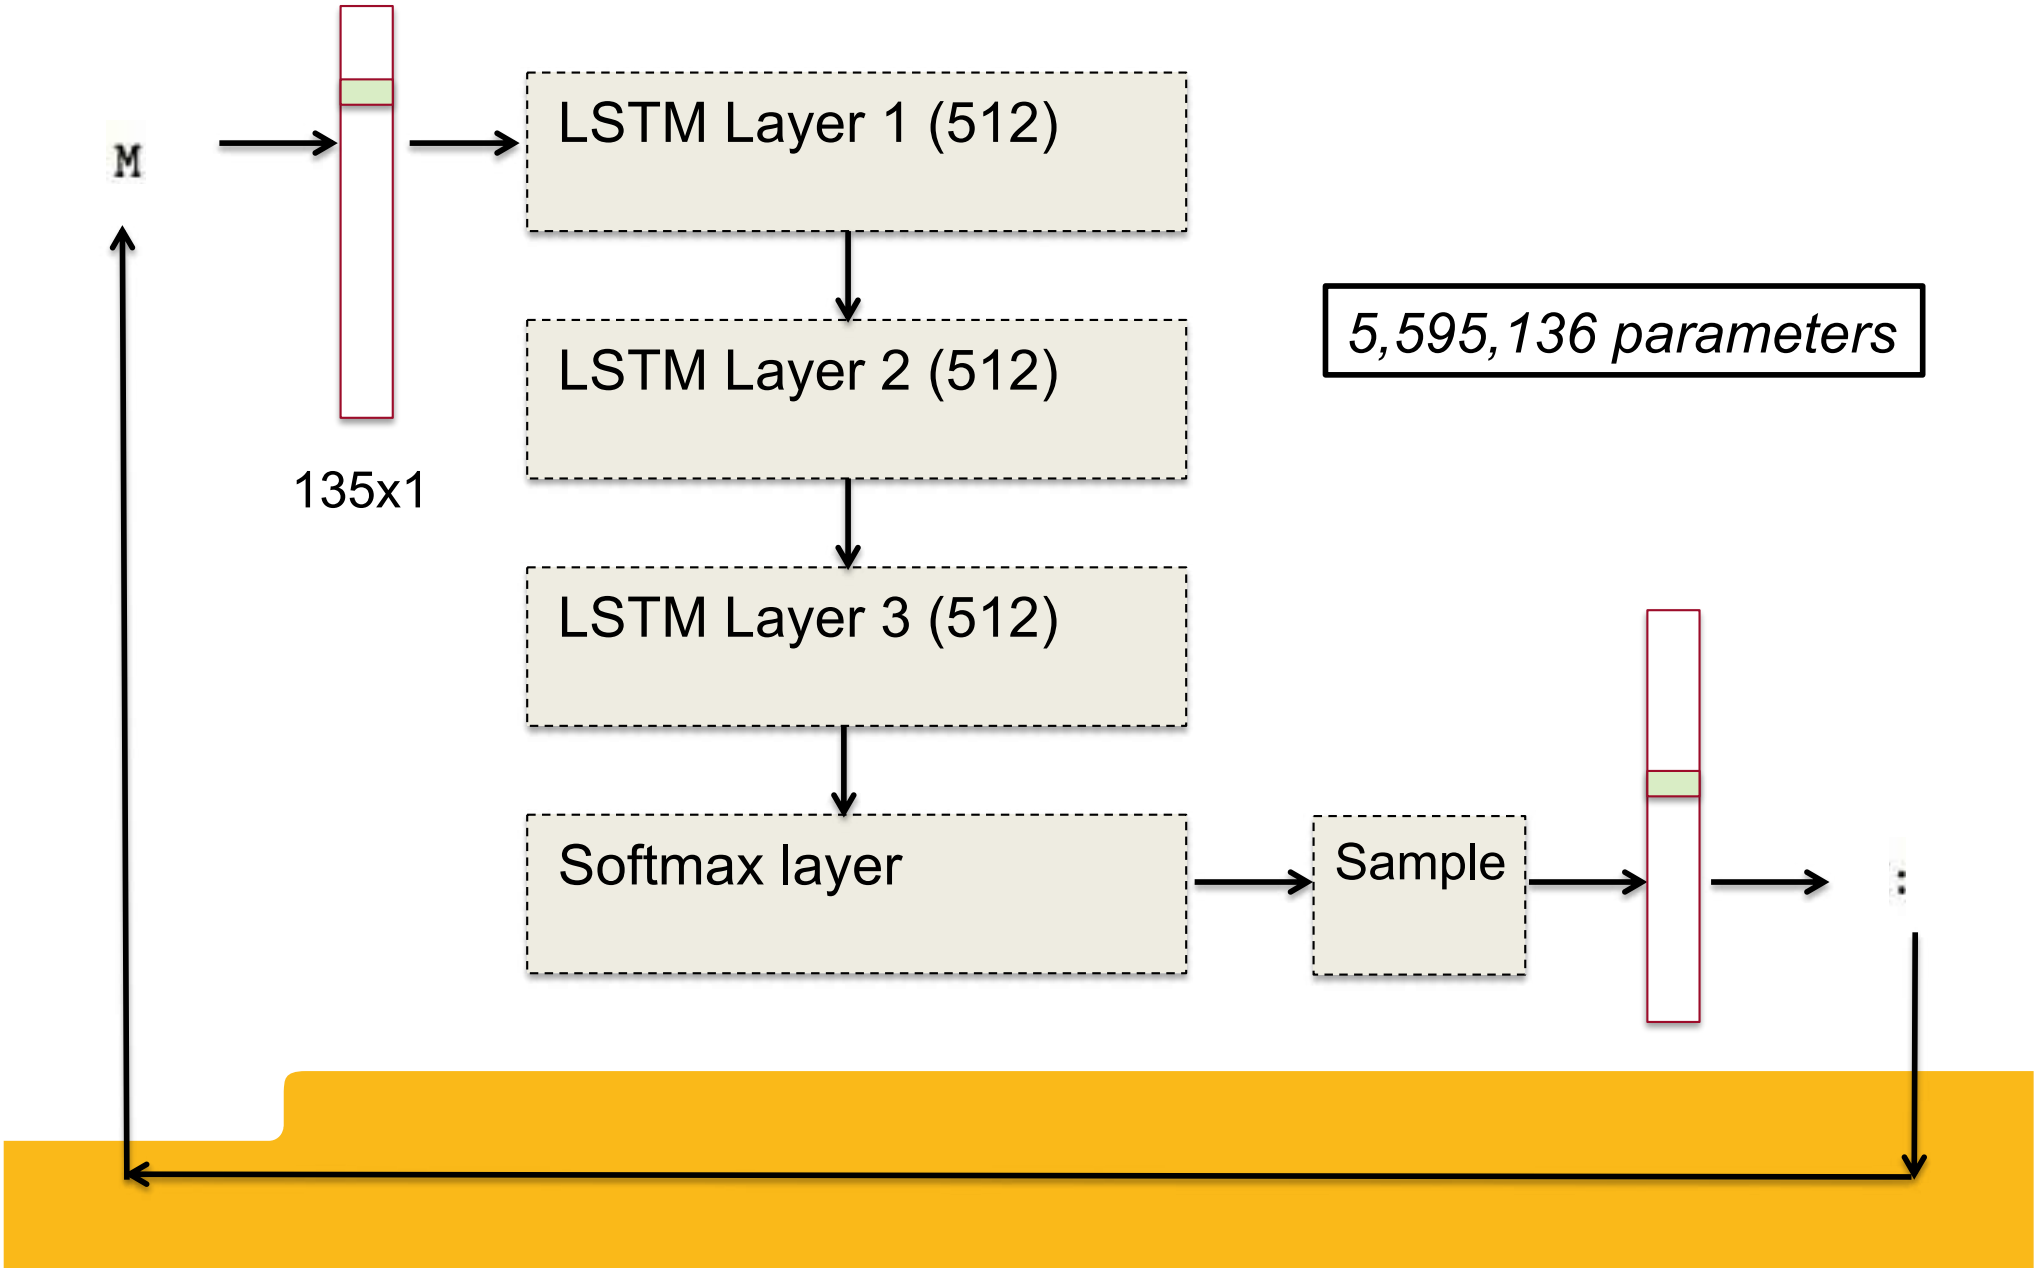

Meh

Let’s use more data!

http://thesession.org

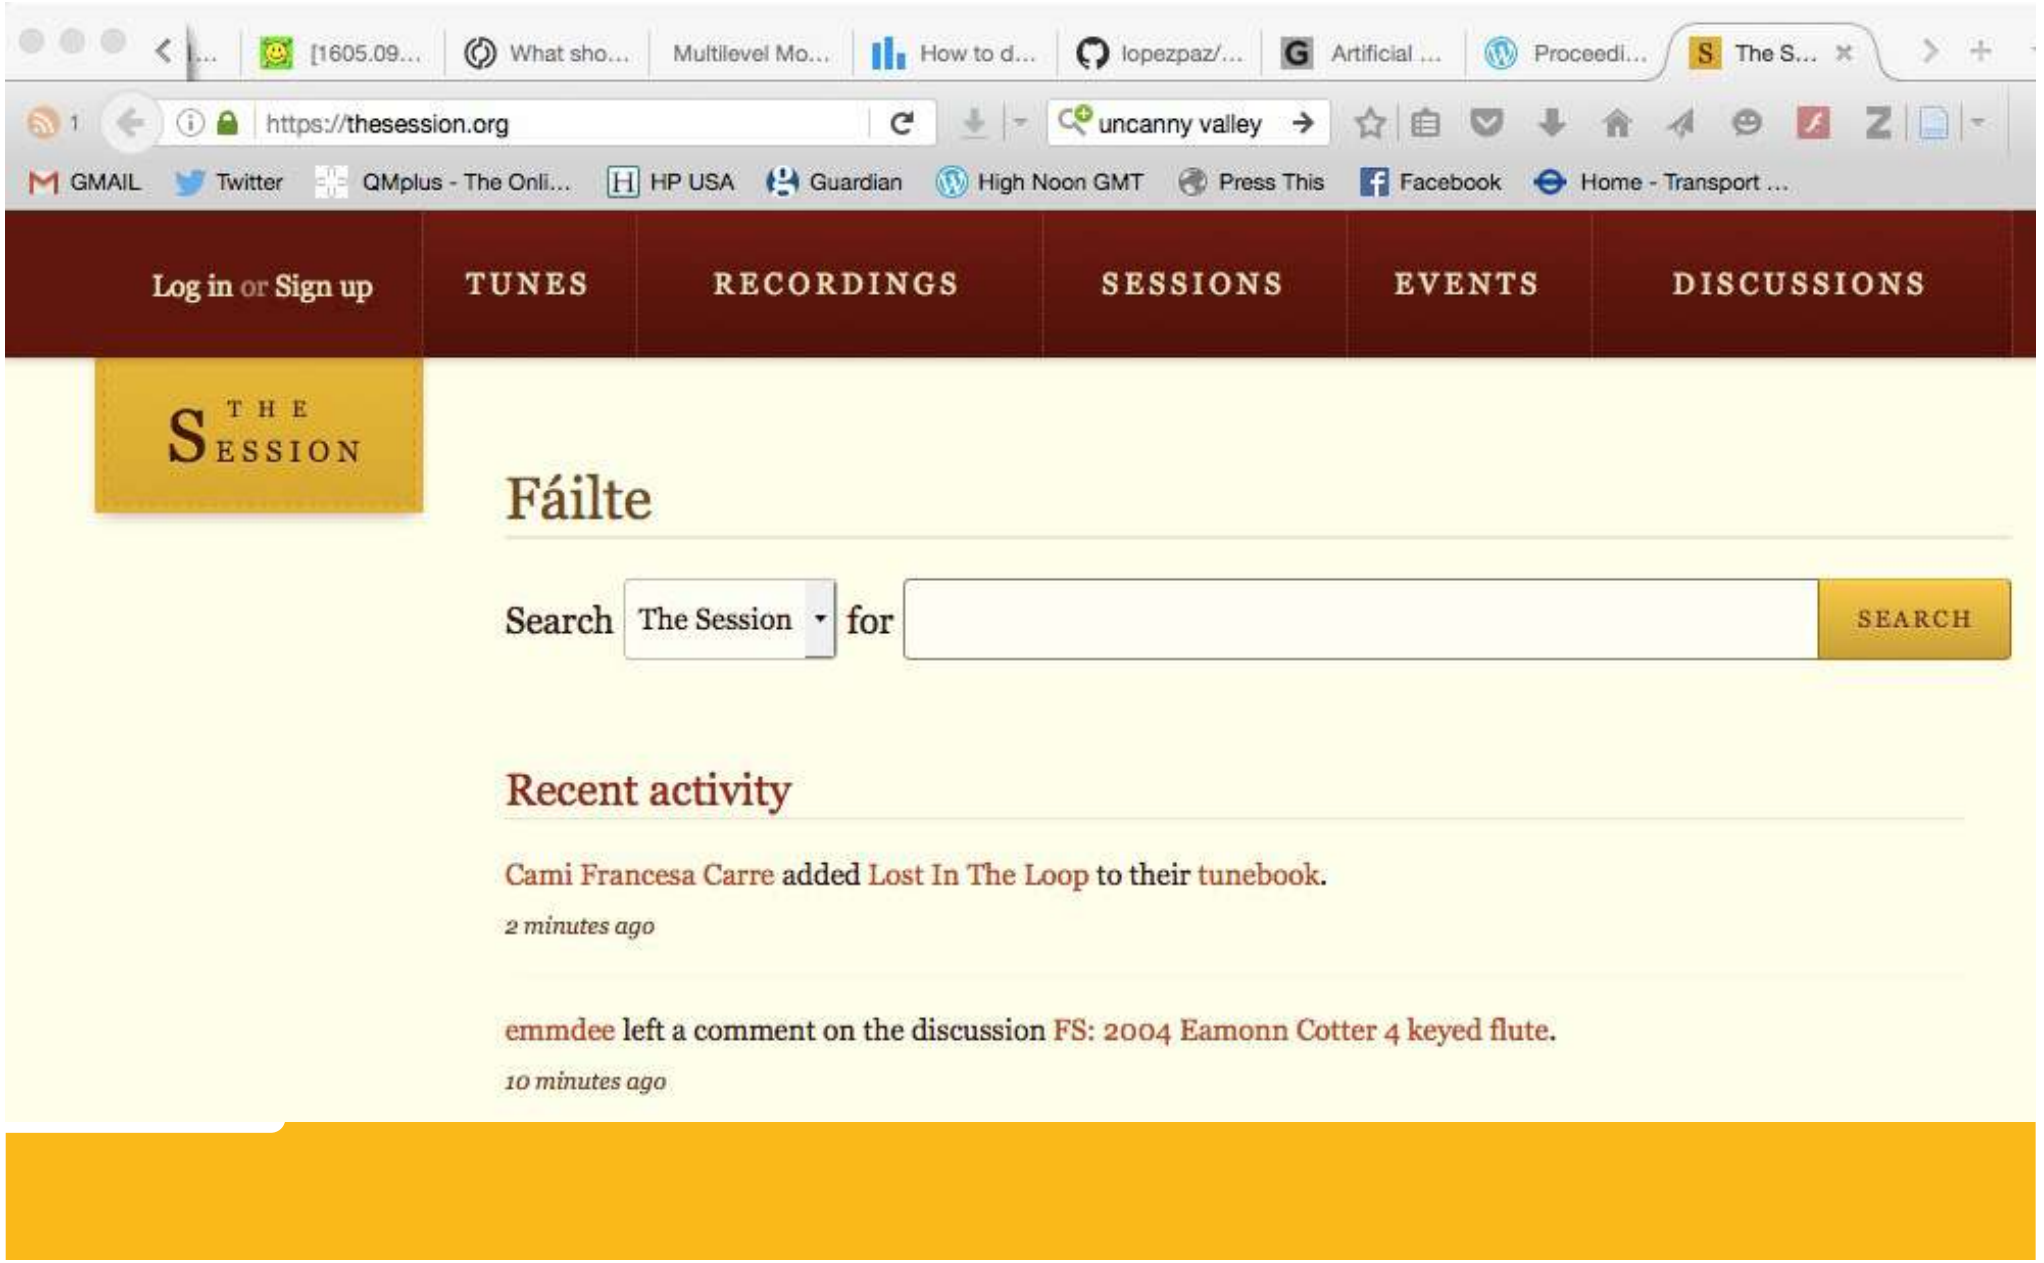

Example transcription from *thesession.org*

27,27,"Drowsy Maggie","reel","4/4","Edorian","|:E2BE dEBE|E2BE AFDF|E2BE dEBE|BABc dAFD:| d2fd c2ec|defg afge|d2fd c2ec|BABc dAFA| d2fd c2ec| defg afge|afge fdec|BABc dAFD|","2001-05-21 03:47:39","Jeremy"

Example transcription from *thesessio*

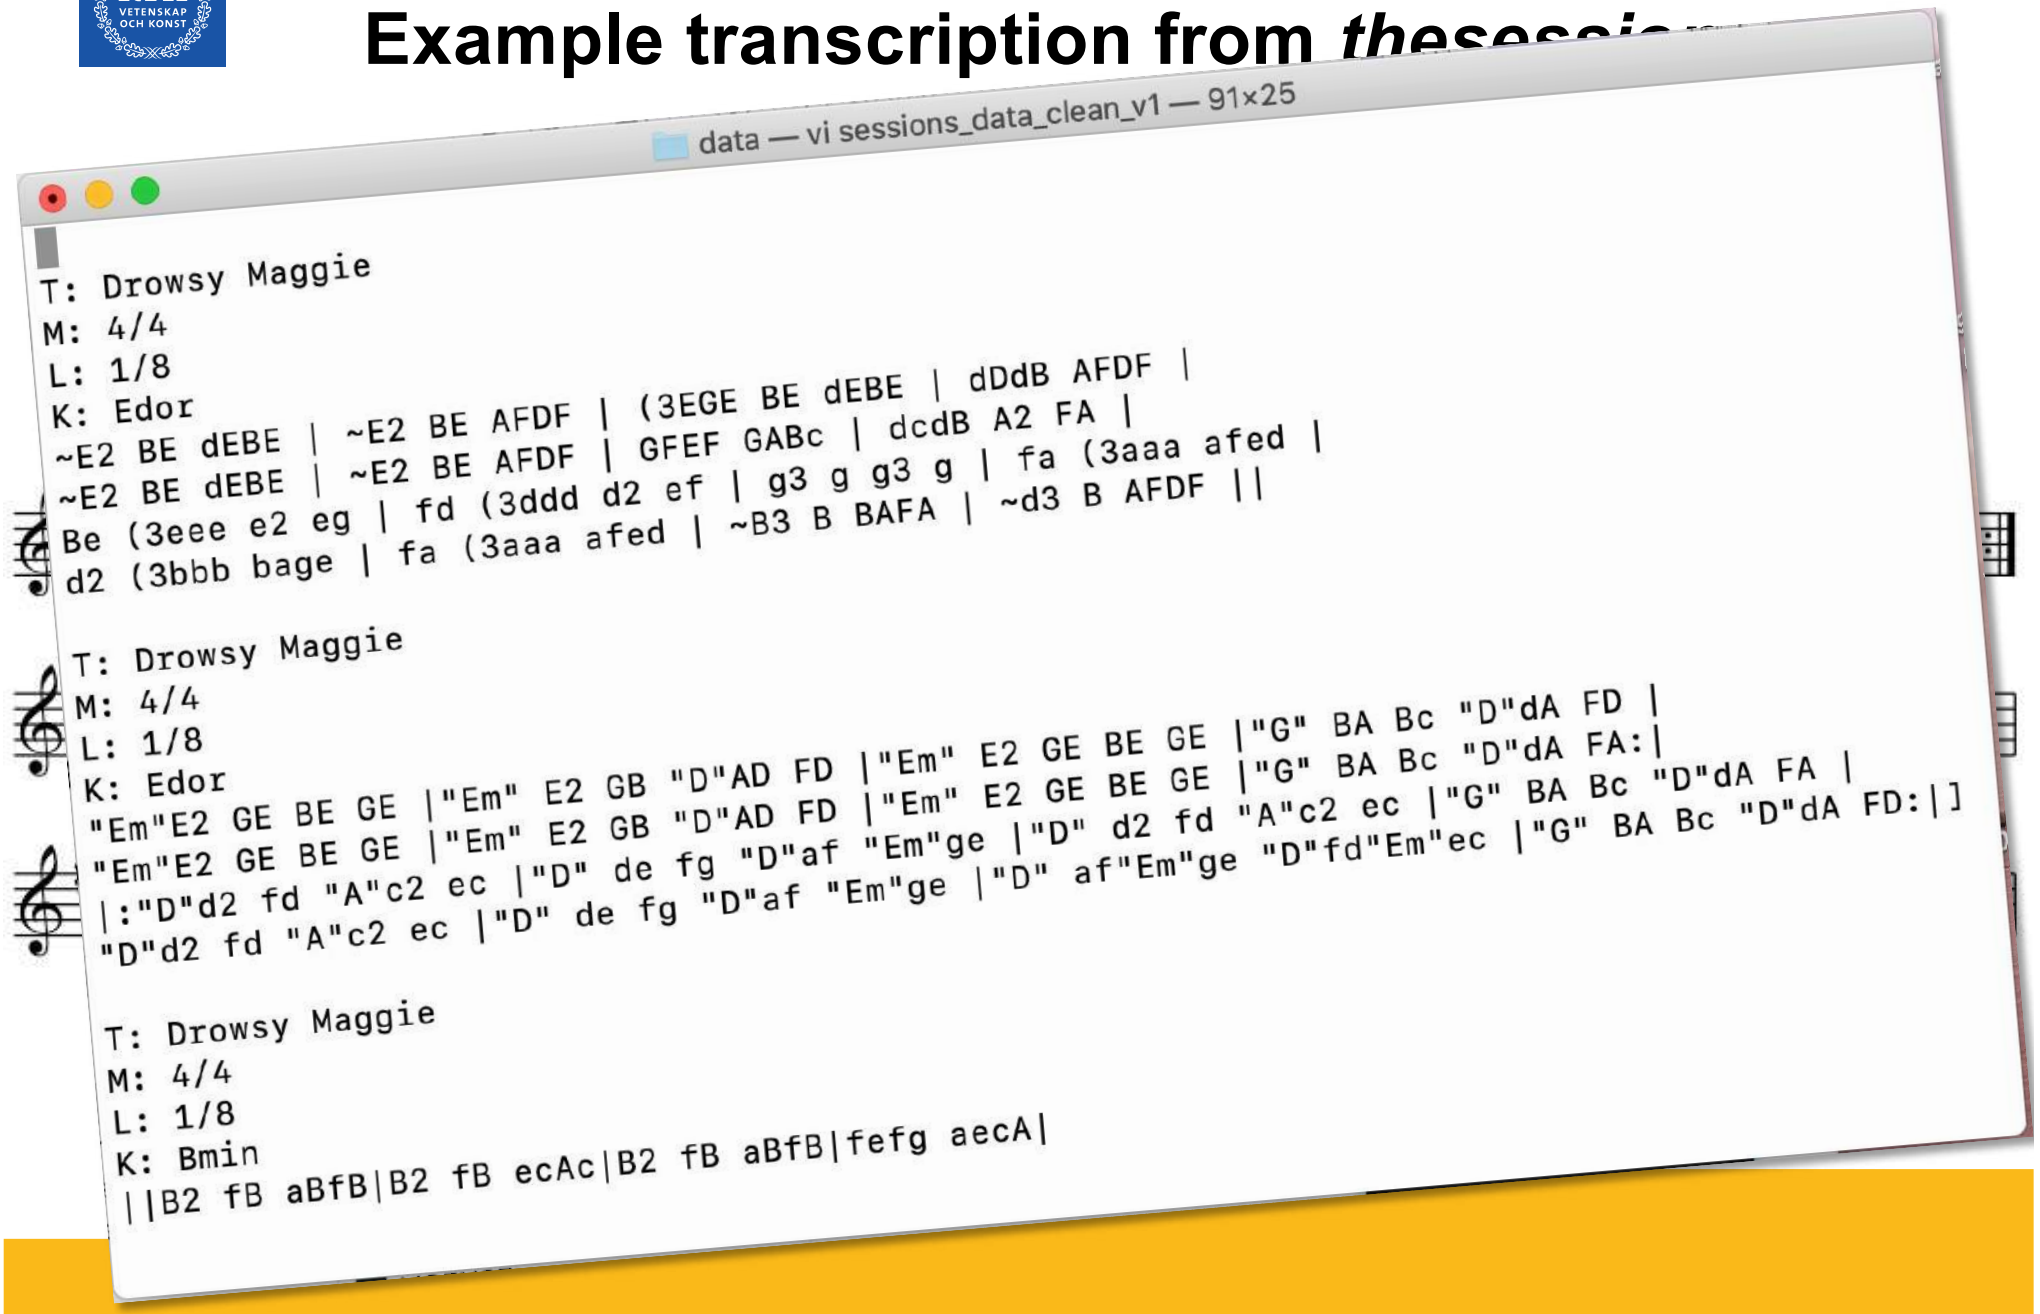

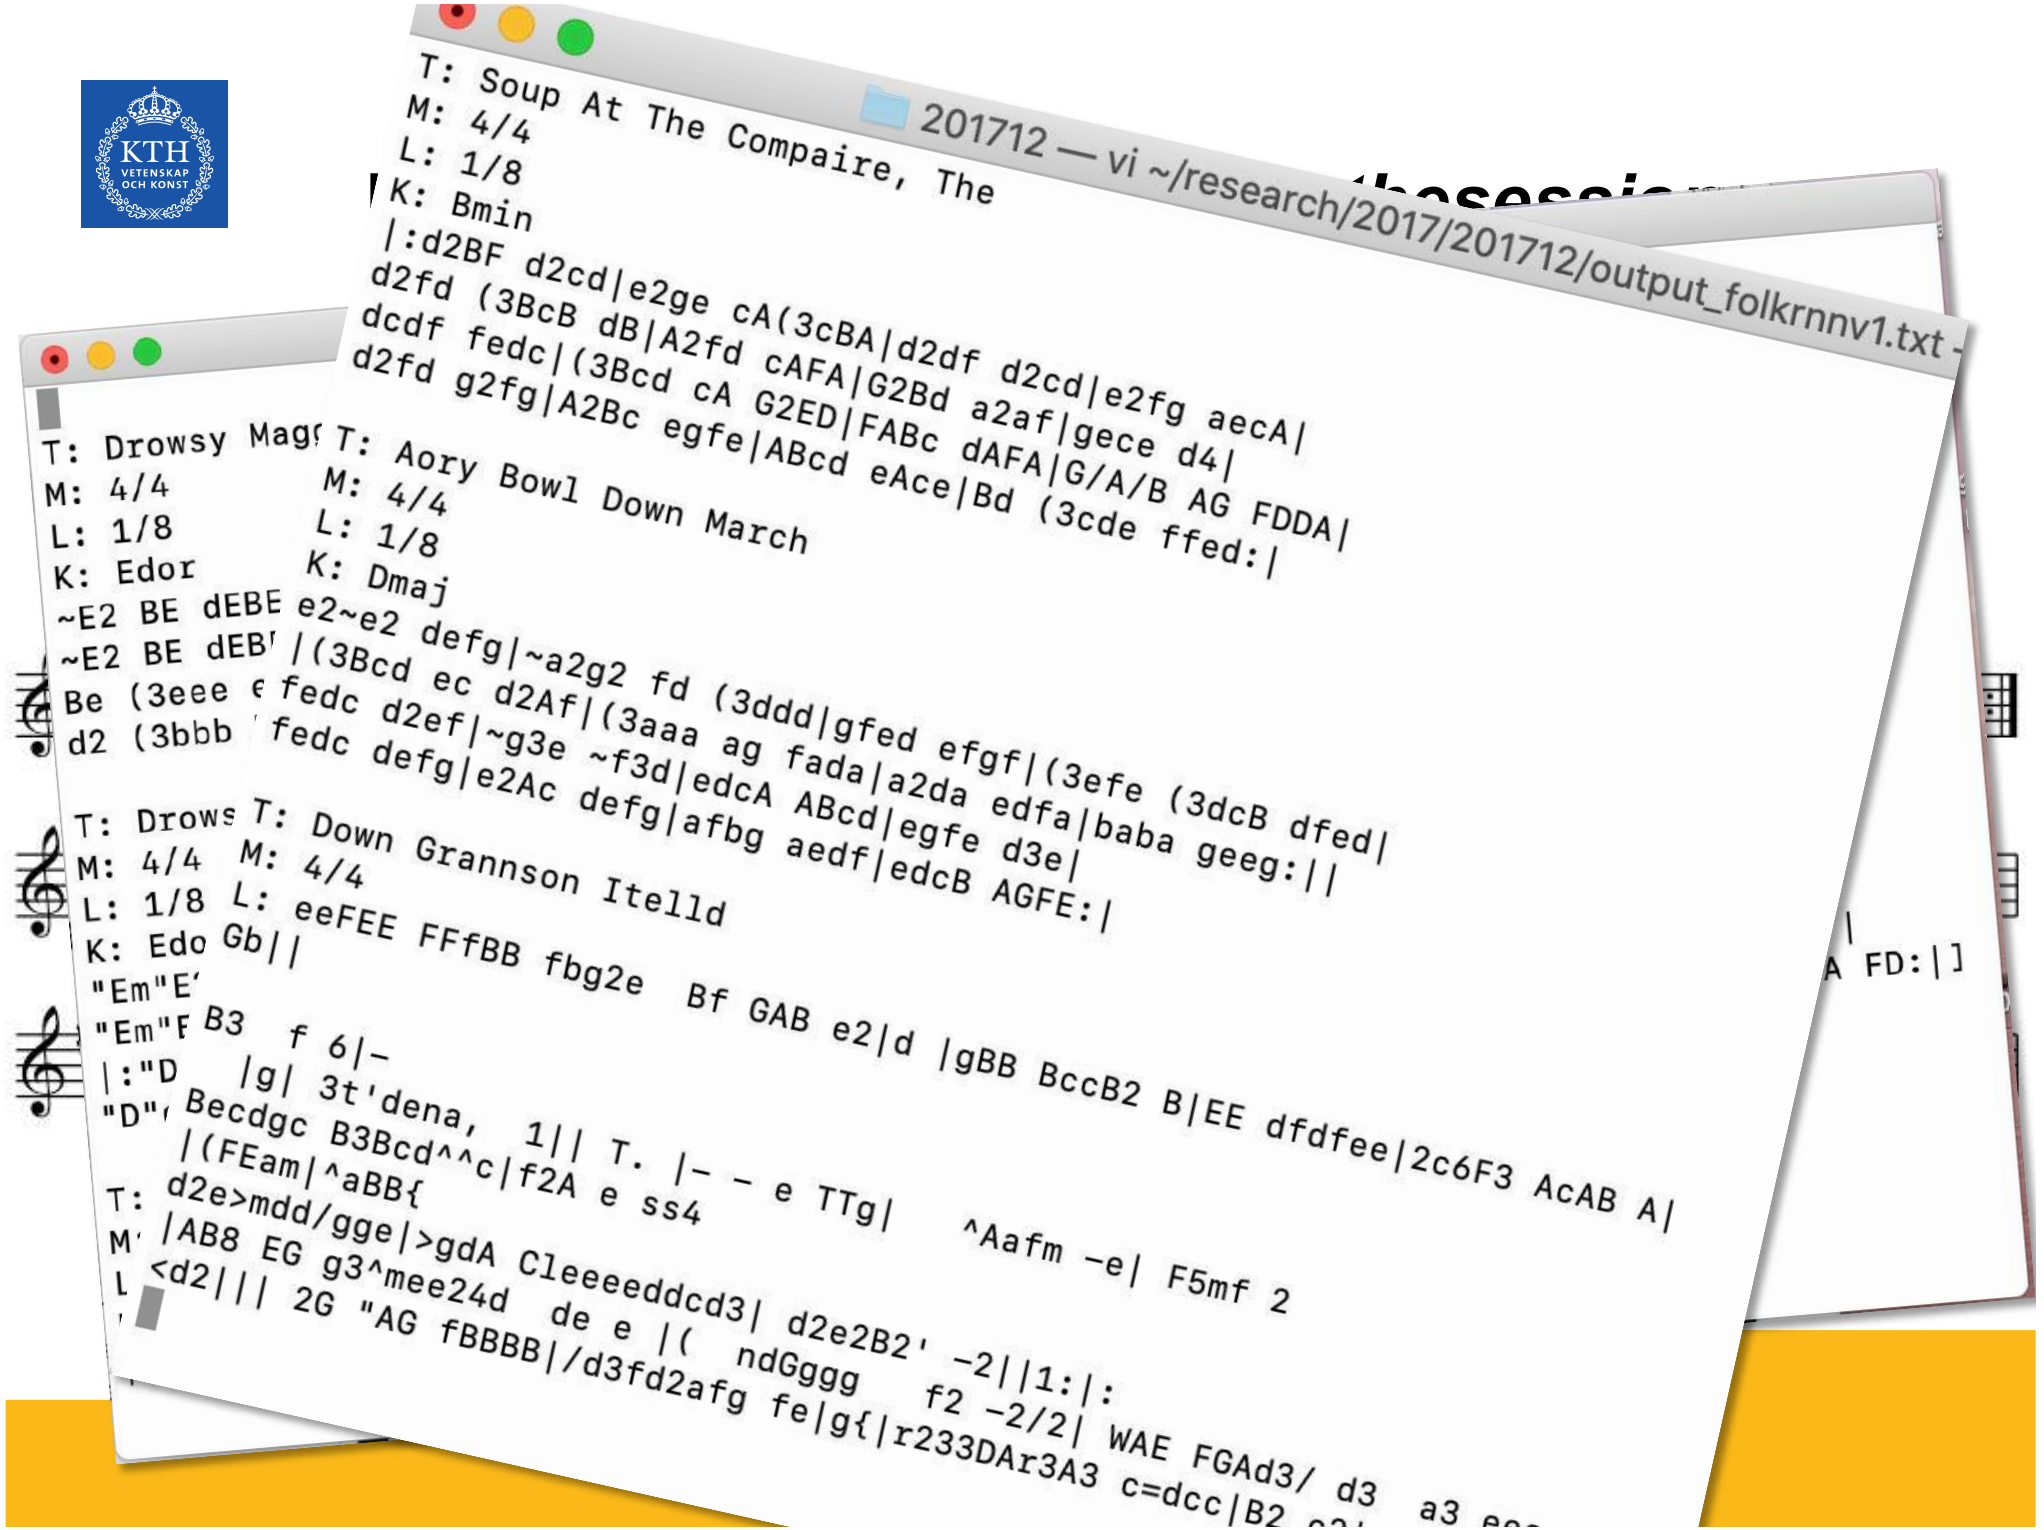
201712 — vi ~/research/2017/201712/output_folkrnnv1.txt
T: Soup At The Compaire, The
M: 4/4
L: 1/8
K: Bmin
|:d2BF d2cd|e2ge cA(3cBA|d2df d2cd|e2fg aecA|
d2fd (3BcB dB|A2fd cAFA|G2Bd a2af|gece d4|
dcdf fedc|(3Bcd cA G2ED|FABc dAFA|G/A/B AG FDDA|
d2fd g2fg|A2Bc egfe|ABcd eAce|Bd (3cde ffed:|
T: Aory Bowl Down March
M: 4/4
L: 1/8
K: Dmaj
T: Down Grannson Itelld
M: 4/4
L: eeFEE FFfBB fbg2e Bf GAB e2|d |gBB BccB2 B|EE dfdfee|2c6F3 AcAB A|
T: Drowsy Mag
M: 4/4
L: 1/8
K: Edor
~E2 BE dEBE e2~e2 defg|~a2g2 fd (3ddd|gfed efgf|(3efe (3dcB dfed|
~E2 BE dEB |(3Bcd ec d2Af|(3aaa ag fada|a2da edfa|baba geeg:||
Be (3eee e fedc d2ef|~g3e ~f3d|edcA ABcd|egfe d3e|
d2 (3bbb fedc defg|e2Ac defg|afbg aedf|edcB AGFE:|
T: Drows
M: 4/4
L: 1/8
K: Edo
"Em"E
"Em"F B3 f 6|-
|:"D |g| 3t'dena, 1|| T. |- - e TTg| ^Aafm -e| F5mf 2
"D" Becdgc B3Bcd^^c|f2A e ss4
|(FEam|^aBB{
T: d2e>mdd/gge|>gdA Cleeeeddcd3| d2e2B2' -2||1:|:
M: |AB8 EG g3^mee24d de e |(ndGggg f2 -2/2| WAE FGAd3/ d3 a3
L: <d2||| 2G "AG fBBBB|/d3fd2afg fe|g{|r233DAr3A3 c=dcc|B2
A FD:|]

Sorpike’s Cat

folk-rnn (v1) + Sturm

https://youtu.be/uSPDnew-7sY

100,000 tunes in 34 volumes

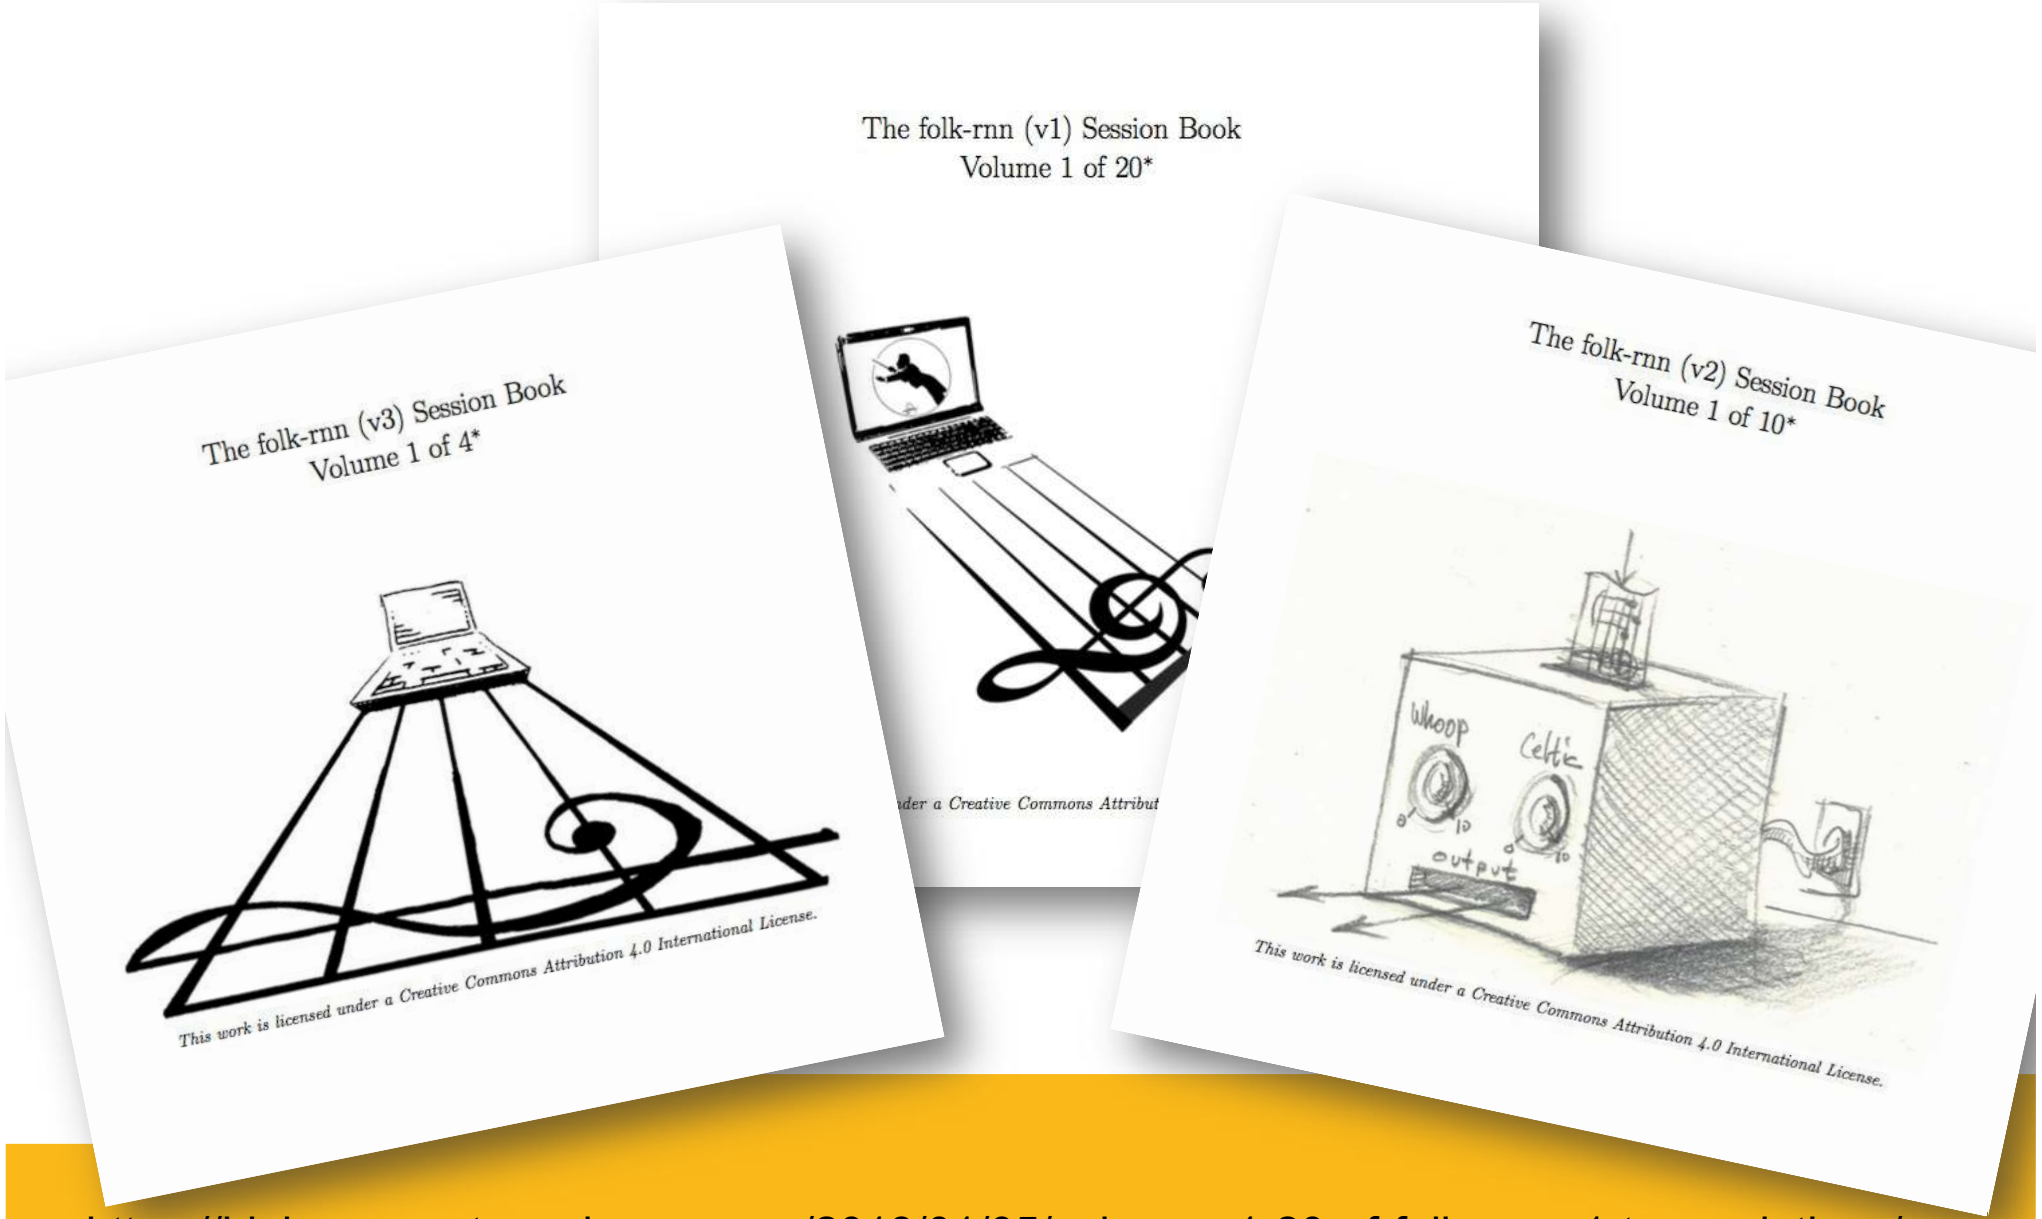

https://highnoongmt.wordpress.com/2018/01/05/volumes-1-20-of-folk-rnn-v1-transcriptions/

The Endless Traditional Music Session (2015)

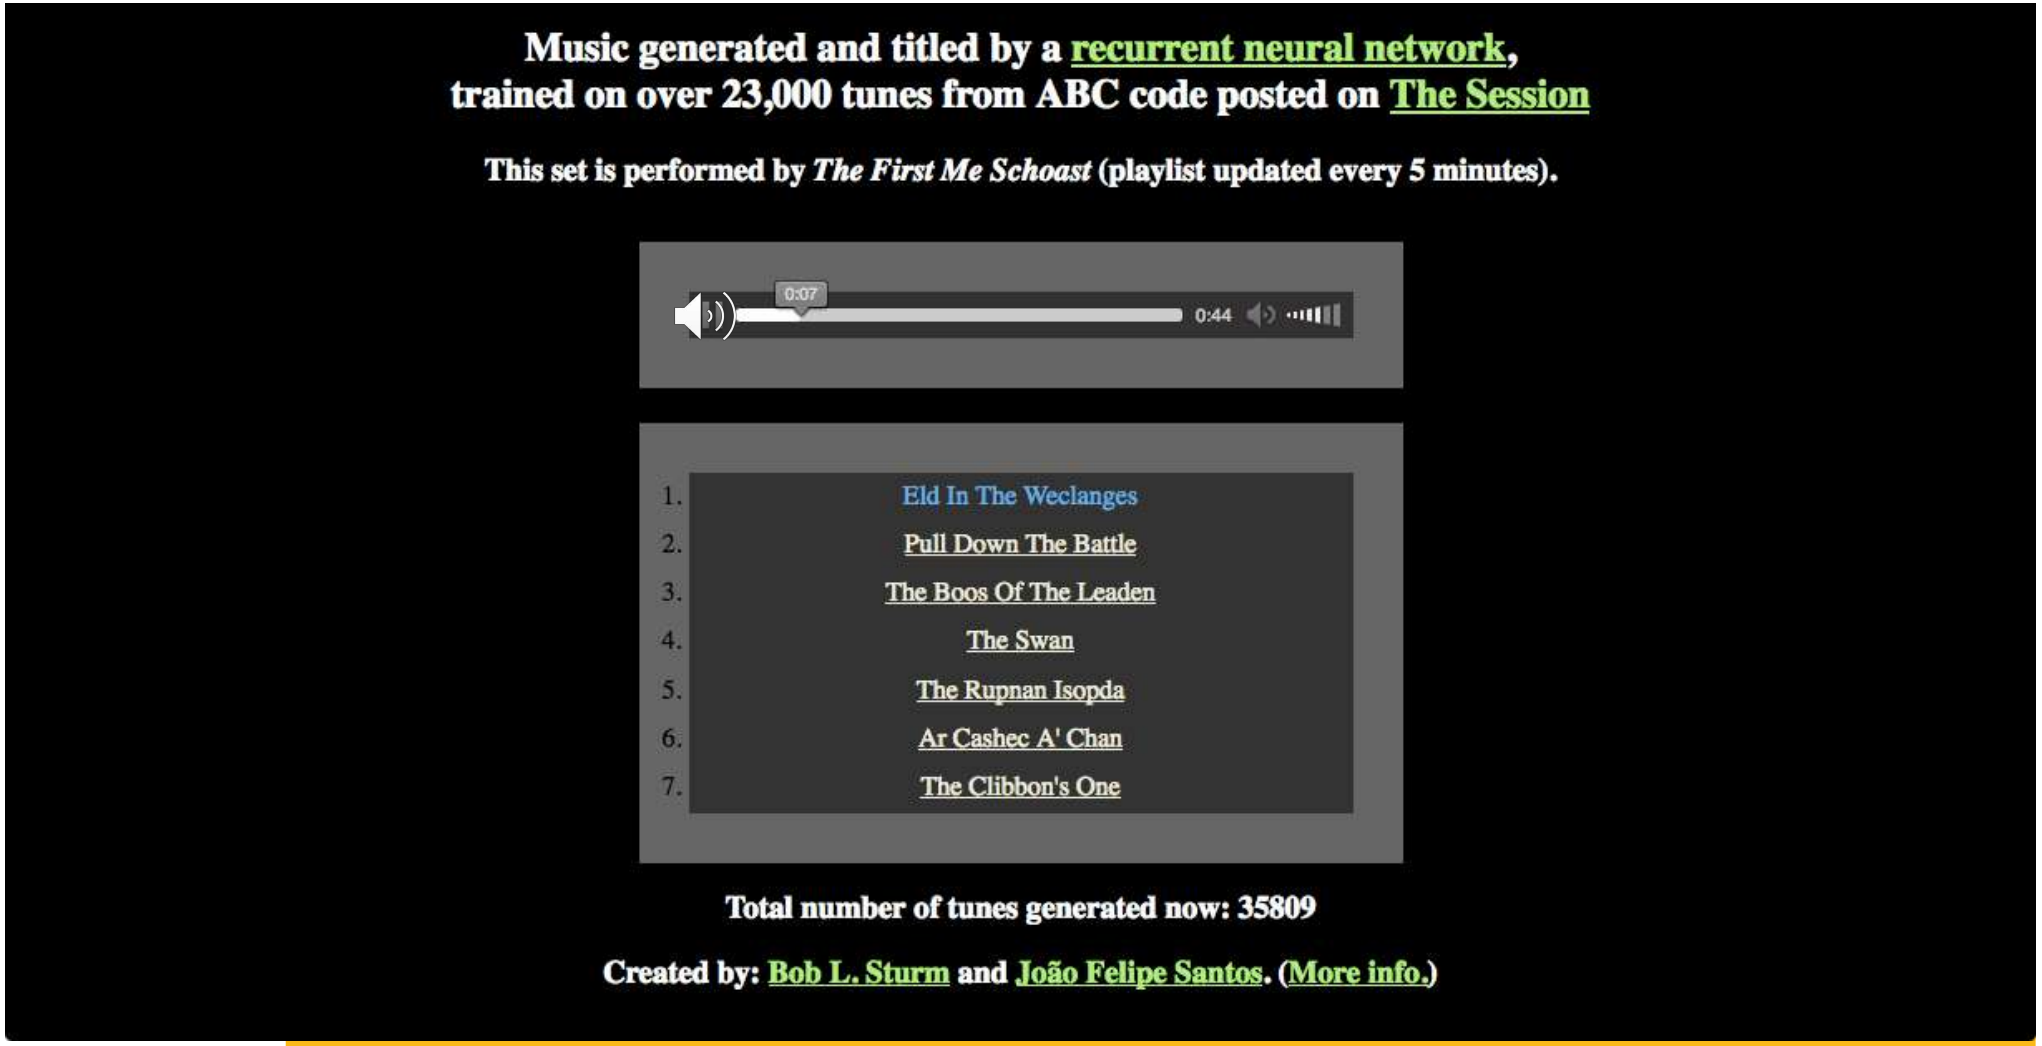

Page no longer works ☹.

https://soundcloud.com/sturmen-1/on-hold-millennial-whoop-reel

http://thesession.org

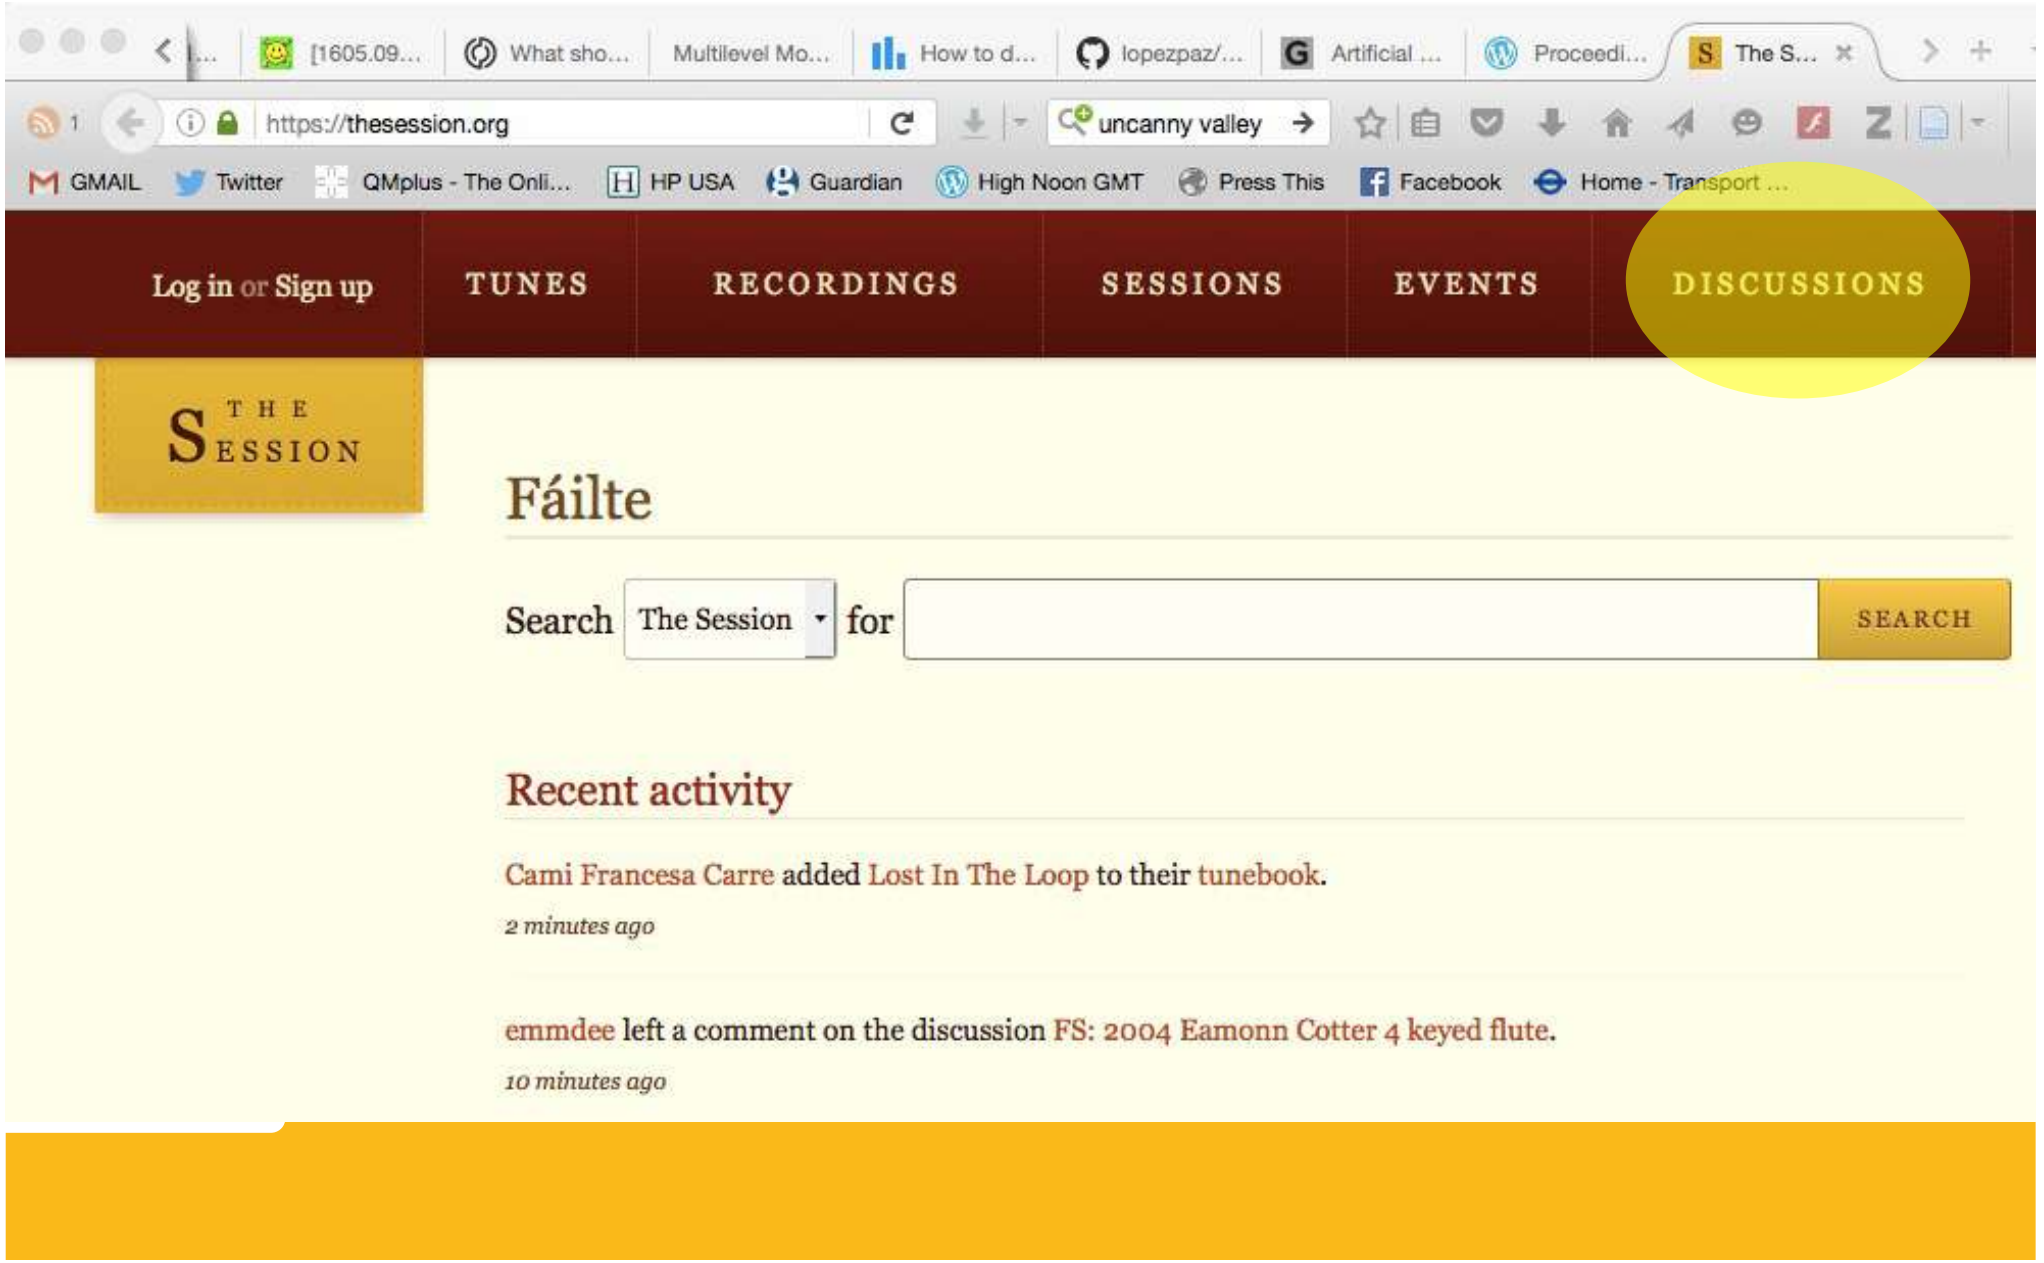

Comments @ thesession.org

Re: On computer generated tunes

Jeez. A computer that noodles. That's all we need.

 # Posted by Mark M 10 months ago.

Re: On computer generated tunes

This has to be a windup, they're terrible!!

hahaha

 # Posted by irishfiddleCT 10 months ago.

Re: On computer generated *sample* tunes

Frankly when I listened to the samples the percussion was the only bit which may have redeeming value. The tunes, as they were, weren't. Tunes that is.
Basically it's crude turntabling without the sense of a musician familiar with the significance of various motifs & phrases.

 # Posted by AB 10 months ago.

Re: On computer generated tunes

Teach it to dance first?! 😕

 # Posted by ceolachan 10 months ago.

"I think it's reckless to send 3,000 machine-created [tunes] into the world." -- *thesession.org* user Ergo

https://thesession.org/discussions/37800

A vanilla LSTM with a special vocabulary

https://github.com/IraKorshunova/folk-rnn

bobsturm committed on **GitHub** Update README.md — Latest commit 52a7d37 22 hours ago

configurations	bugfix	a year ago
data	change to readme	a year ago
metadata	Updated metadata with newer trained model	a year ago
samples	add metadata and samples	2 years ago
soundexamples	Update README.md	ago
.gitignore	clean before checkout	ago
LICENSE	license	ago
README.md	Update README.md	
data_iter.py	added 1hot option	
logger.py	bugfix	
sample_rnn.py	added terminal output, fixed initialisation	
train_rnn.py	bugfix	ago

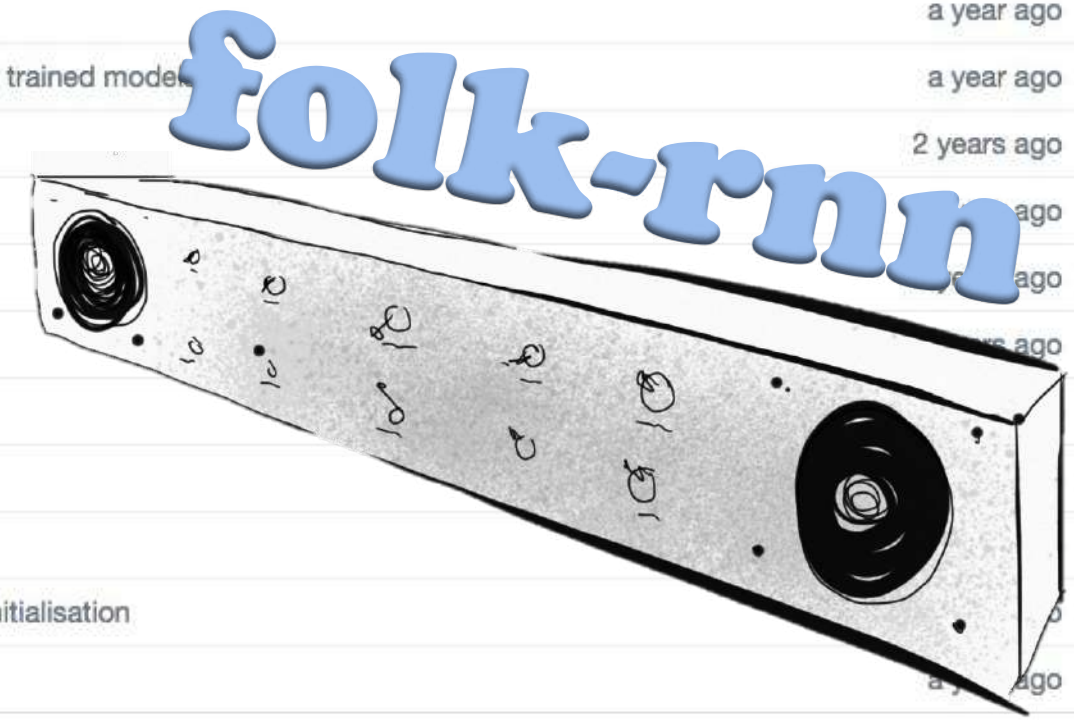

Sturm, et al., “Music transcription modelling and composition using deep learning,” in *Proc. Conf. Computer Simulation of Musical Creativity*, 2016.

Token types in the *folk-rnn* (v2) vocabulary

***Aim**: design a vocabulary that has little ambiguity, and efficiently represents transcriptions.*

1. Transcription (2) <s>, </s>
2. Meter (7) M:4/4 
3. Mode (4) K:Cdor 
4. Measure (5) | :| 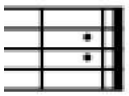
5. Pitch (85) A c ^f 
6. Duration (25) 2 
7. Grouping (9) (3 [ ] 

- *All titles, ornamentations, etc. removed*
- *All tunes transposed to root C*
- *137 tokens in total*

Example transcription from *thesession.org*

27,27,"Drowsy Maggie","reel","4/4","Edorian","|:E2BE dEBE|E2BE AFDF|E2BE dEBE|BABc dAFD:| d2fd c2ec|defg afge|d2fd c2ec|BABc dAFA| d2fd c2ec| defg afge|afge fdec|BABc dAFD|","2001-05-21 03:47:39","Jeremy"

extract, transpose, tokenise

M:4/4 K:Cdor |: C 2 G C B C G C | C 2 G C F D B, D | C 2 G C B C G C | G F G A B F D B, :| B 2 d B A 2 c A | B c d e f d e c | B 2 d B A 2 c A | G F G A B F D F | B 2 d B A 2 c A | B c d e f d e c | f d e c d B c A | G F G A B F D B, |

folk-rnn (v2): architecture

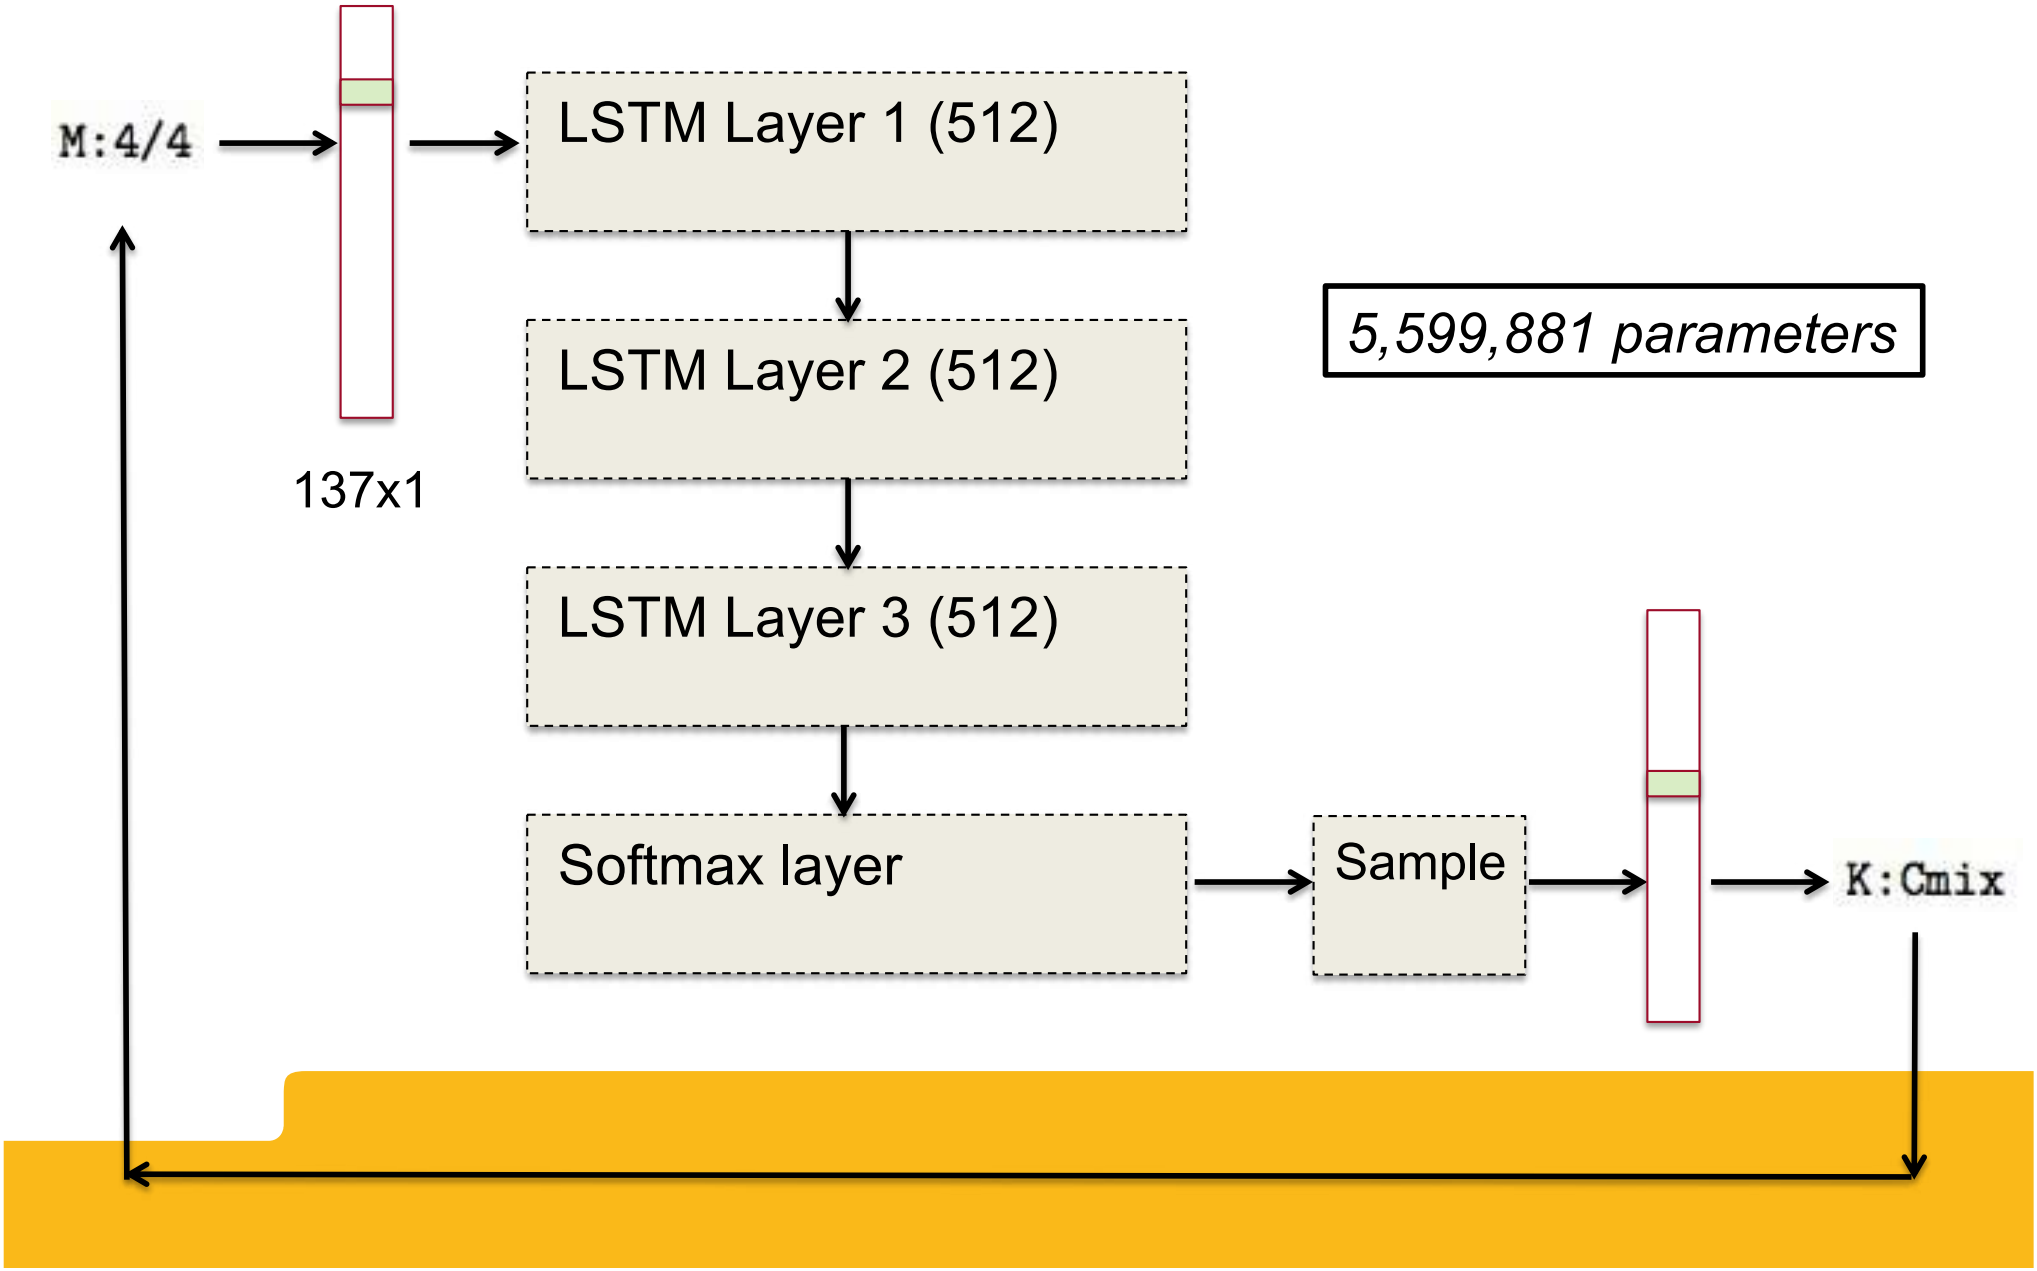

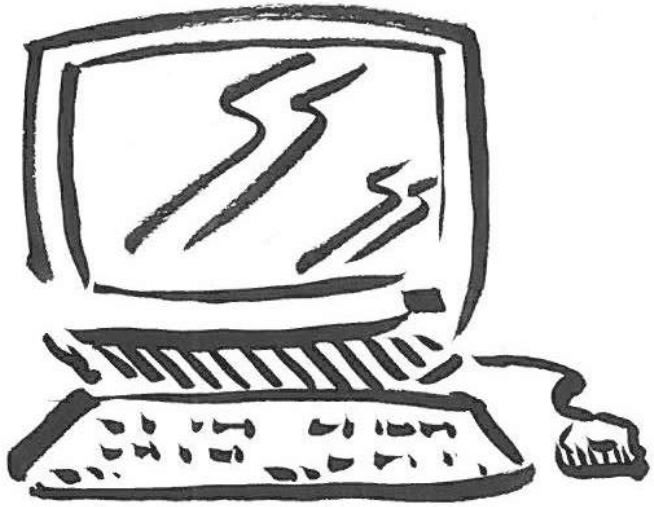

Human-in-the-loop composition

Deep learning for assisting the process of music composition (part 1)

Edit

Posted on August 11, 2015 by Bob L. Sturm

This is part 1 of my explorations of using deep learning for assisting the process of music composition. In this part, I look at some almost-winning output of a model trained by deep learning methods on over 23,000 folk tunes, and make improvements to produce a session-ready piece.

Deep learning for assisting the process of music composition (part 4)

Edit

Posted on August 15, 2015 by Bob L. Sturm

This is part 4 of my explorations of using deep learning for assisting the process of music composition. In this part, I look at the process of composition using a model built using deep learning methods on over 23,000 folk tunes. Part 1 is here. Part 2 is here. Part 3 is here.

https://highnoongmt.wordpress.com/2015/08/15/deep-learning-for-assisting-the-process-of-music-composition-part-4/

https://folkrnn.org/

folkRNN

generate a folk tune with a recurrent neural network

PRESS TO GENERATE TUNE

Compose

MODEL

thesession.org (w/ :|

TEMPERATURE 1

SEED 359615

METER 4/4

MODE C Major

INITIAL ABC

Enter start of tune in ABC notation

FOLK RNN TUNE №1580

enter the folkRNN composition competition

```
X:1580
M:4/4
K:Cmaj
cGGE GFEG|(3CCCCD EGEG|cAdc ecdB|ACDE G3A|
(3cccCG EGGC|EDCD (3EEEG2|cGGG e2dc|(3ABcdB c3d:|
|:e3d cege|f3d Bdga|gfeg ecde|fegc A2GF|
EGG2 cdeg|(3fffaf dafd|e2eg fdBc|dedB c3d:|
```

*The RNN properties were **thesession_with_repeats** with seed **441885** and temperature **1**.*

*The prime tokens were **M:4/4 K:Cmaj**.*

*Generated on **14/06/2018, 14:46:35**.*

HEAR IT

0:00 Download MIDI

SEE IT

Theme from “Close Encounters of the 3rd Kind”

<s> M:4/4 K:Cmaj d e c C G 2

Theme from “Close Encounters of the 3rd Kind”

folk-rnn (v2) + Sturm

https://youtu.be/ZN0Fi_oyu44

Two Short Pieces with an Interlude

1. March to the Mainframe
2. Interlude
3. The Humours of Time Pigeon

Chicken Bits and Bits and Bobs

Sturmen

We three layers o' hidd'n units are

3 years ago

neural network

Sturmen

It Came Out From a Pretrained Net

2 years ago

3:23

https://soundcloud.com/sturmen-1

“Bastard Tunes” by Oded Ben-Tal (2017)

https://youtu.be/YZ2jb0ksOm4

Failure can be very interesting

"The Humours of Time Pigeon"

Arranged for ensemble:

https://bit.ly/386wrnG

Daren Banarsë

Jennifer Walshe

John Hughes

Laura Agnusdei

Paudie O'Connor
Aoife Ní Chaoimh

Úna Monaghan

Torbjörn Hultmark

Luca Turchet

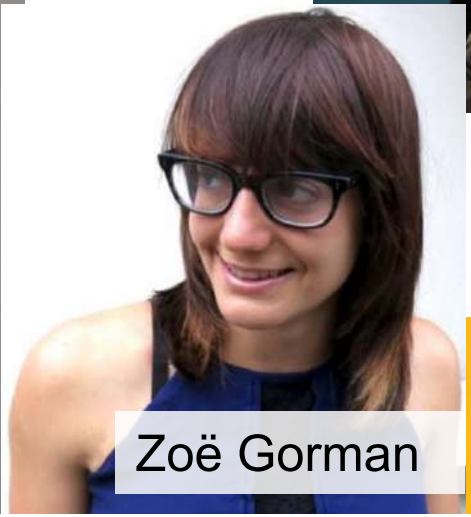

Zoë Gorman

Jennikel Andersson

Pieces for

Richard Salmon
organ

Music in the age of Artificial Creation

November 20 2017 7 PM
St Dunstan and All Saints, Stepney High Street E1 0NR

PARTNERSHIPS

Music composed by and co-composed with computers having "musical intelligence"

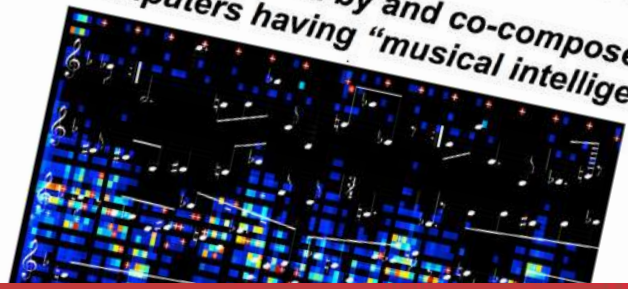

PRESENTED BY O'REILLY AND INTEL AI

BEIJING • NEW YORK • SAN FRANCISCO • LONDON

Artificial Intelligence CONFERENCE

PUT AI TO WORK
8-9 OCT 2018: TRAINING
9-11 OCT 2018: TUTORIALS & CONFERENCE
LONDON, UK

MUME 2017 CONCERT

The Musical Metacreation Concert at the Eighth International Conference on Comp tional Creativity, Atlanta, Georgia

Monday June 19th. Doors 7pm. Music 8pm.
Mammal Gallery, Atlanta.
Tickets $20.

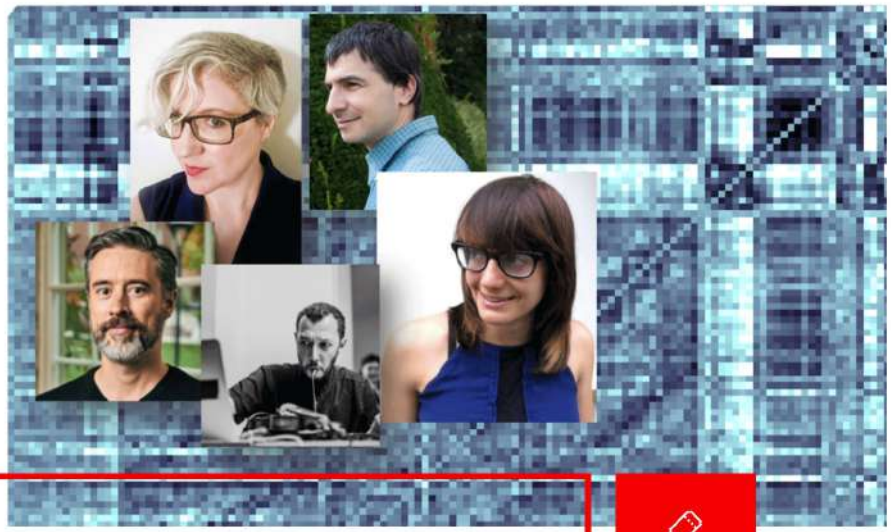

BOKA BILJETT

Musik/Data – Data/Musik

Klassiskt, Elektroakustisk musik

tisdag 15 oktober 20.00 - 21.30

Pieces for organ: The Glas Herry Comment & X:7153 by *folk-rnn* + *Deep Bach* (2017)	**Richard Salmon** *organ*
Traditional Irish Sets (with *folk-rnn* tunes in *italics*) • **Jigs** (The Cuil Aodha, The Dusty Windowsill, *The Glas Herry Comment*) • **Slow Reels** (Maghera Mountain, *X:2897*) • **Fast Reels** (The Rookery, *X:1068*, Toss The Feathers II)	**Daren Banarsë and Musicians**
March to the Mainframe, Interlude, The Humours of Time Pigeon by Bob L. Sturm + *folk-rnn* (2017)	**Ensemble x.y**
Ed SheerAI vs XenAkIs vs AIdele by Nick Collins (2017)	**Ensemble x.y**
3 morphed pieces from "A Little Notebook for Anna Magdalena" by J. S. Bach (1722) + MorpheuS (2017) **3 morphed pieces from "30 and 24 Pieces for Children"** by Kabalevsky (1937) + MorpheuS (2017)	**Elaine Chew** *piano*
Safe Houses by Úna Monaghan + *folk-rnn* (2017)	**Úna Monaghan**
The Choice by Úna Monaghan (2015) ***The Chinwag*** by Úna Monaghan (2015)	*Irish harp, concertina, electronics*
Pieces for organ: X:633 & The Drunken Pint by *folk-rnn* + *Deep Bach* (2017)	**Richard Salmon** *organ*
Chicken Bits and Bits and Bobs by Bob L. Sturm + *folk-rnn* (2017)	**Ensemble x.y**
Bastard Tunes by Oded Ben-Tal + *folk-rnn* (2017)	**Ensemble x.y**

YouTube GB

Search

32 subscribers | 6,520 views | Video Manager | https://www.youtube.com/channel/UC7wzmG64y2IbTUeWji_qKhA

The Bottomless Tune Box

Subscribe 32

The music on this channel arises from creative partnerships between musicians and a deep learning system that models music transcription ... Show more

Uploads Public

Set #3 (fast reels)
31 views • 1 week ago

Set #1 (jigs)
69 views • 3 weeks ago

"Chicken Bits and Bits and Bobs" by Bob L. Sturm + folk-rnn
68 views • 3 weeks ago

"Interlude" by Bob L. Sturm + folk-rnn
26 views • 3 weeks ago

"The Humours of Time Pigeon" by Bob L. Sturm + folk-rnn
41 views • 3 weeks ago

Created playlists Public

C4DM concert (QMUL Nov. 18 2016)

Partnerships concert (May 23 2017)

Two short pieces and an Interlude in Concert

Bastard Tunes in Concert

C4DM concert (QMUL Nov 23 2015)

An Unintentional Experiment!

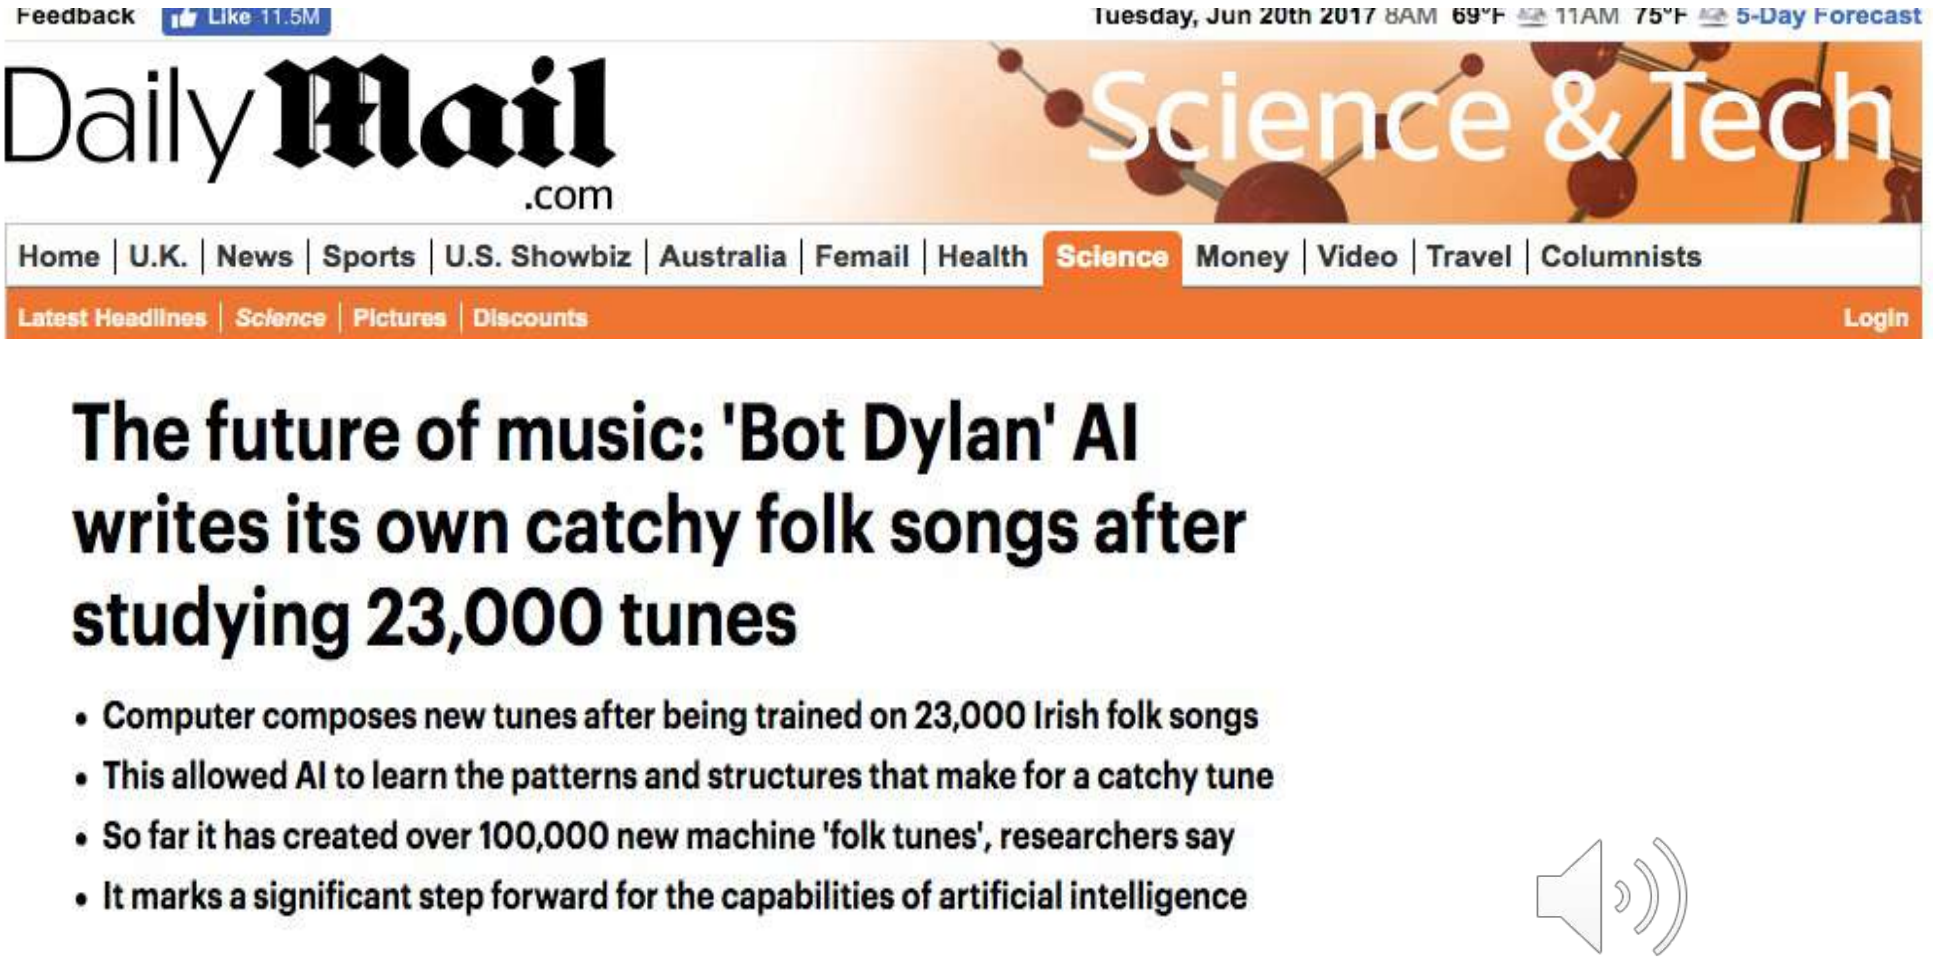

https://www.dailymail.co.uk/sciencetech/article-4544400/Researchers-create-computer-writes-folk-music.html

Duane_1981, Preston, United Kingdom, 11 months ago

It's sounds very neat. It's missing the "human" element.

PaxRomana, Novi, 11 months ago

That's it?!!! I'm not impressed.

paevo, USA, United States, 11 months ago

Sounds like a robotic Irish jig....

Mikeyt1941, London, Canada, 11 months ago

Totally lifeless without warmth. Mind you much human tuneless junk that passes for music today isn't much better.

Click to rate 11 1

Fabrice, Manchester, United Kingdom, 11 months ago

No no no.

rocksnoop1, dover, United Kingdom, 11 months ago

Isn't music robotic enough these days?

pen, somewhere, United Kingdom, 11 months ago

Let's make all humans redundant, brilliant! Has everybody really lost their soul?!

Radar Also, Hemet, 11 months ago

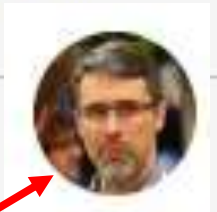

This computerized "AI" is just so non musically untalented lazy nerds can infiltrate the world of true musicians who love, created, and write the music from the joy, hurt, and life emanating from their hearts.

Click to rate 7 1

Evaluation: Unintentional experiments

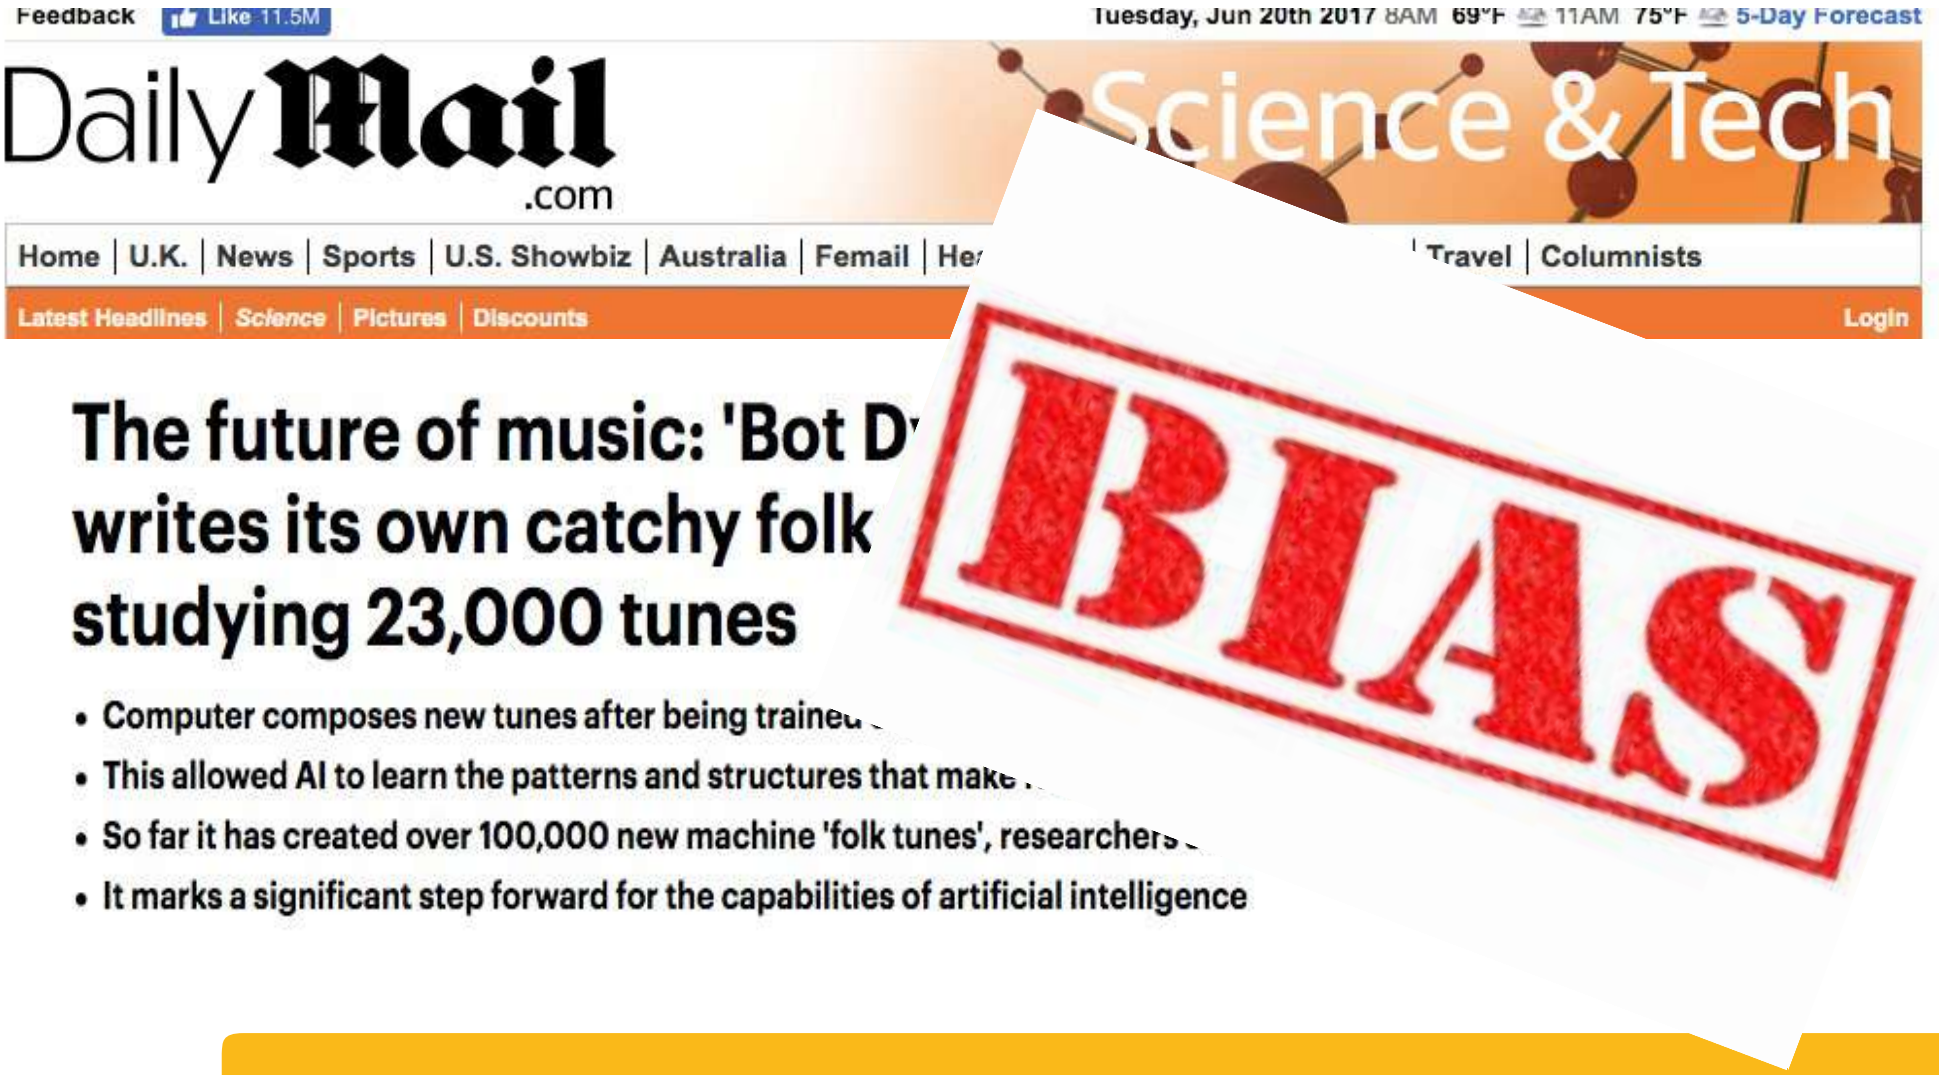

https://www.dailymail.co.uk/sciencetech/article-4544400/Researchers-create-computer-writes-folk-music.html

An Intentional Experiment

How difficult will it be for a professional musician to produce an album using material generated by our system that will be judged successful within the idiom of Irish traditional music?

100,000 tunes in 34 volumes

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


The reveal is
coming soon ...

“Let’s Have Another Gan Ainm”

Track listing:

1. Gan Ainm, Gan Ainm, Gan Ainm
2. The Drunken Landlady, Gan Ainm, Gan Ainm
3. Gan Ainm, Gan Ainm, Gan Ainm
4. Battle Of Aughrim, Gan Ainm, Lord Mayo
5. Gan Ainm, Gan Ainm, Tom Billy’s
6. Girls Of Banbridge, Gallowglass, Gan Ainm
7. The Blackbird, Gan Ainm, Mrs Galvin’s
8. Gan Ainm
9. Gan Ainm, Bunch of Green Rushes, Gan Ainm
10. Gan Ainm, Gan Ainm, Anthony Frowley’s
11. Gan Ainm, Toss the Feathers (II), Gan Ainm

https://soundcloud.com/oconaillfamilyandfriends

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*. Technical report, KTH.

“Let’s Have Another Gan Ainm”

- 31 tunes in total arranged in 11 sets
- material of 20 tunes generated by *folk-rnn* models
- 11 tunes come from traditional repertoire
- Banarsë had access to all 100,000 tunes, but he took tunes from only six of the 34 volumes
- Banarsë and the musicians were free to do as they wished
- **More info**:
  - Ben-Tal, Harris, and Sturm, “How music Ai is useful: Engagements with composers, performers, and audiences”, *Leonardo,* 2020.

Machines generating music is nothing new

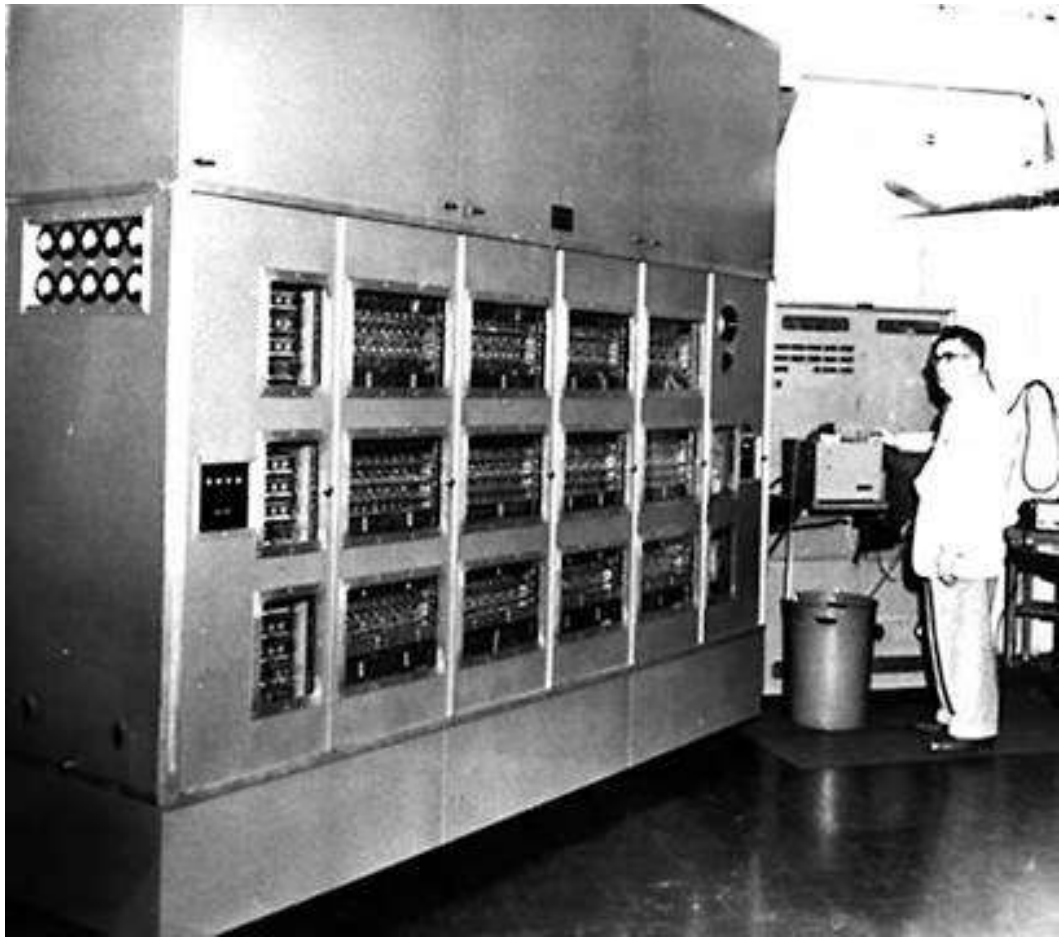

The Illiac Suite (1957)

Machines generating music is nothing new

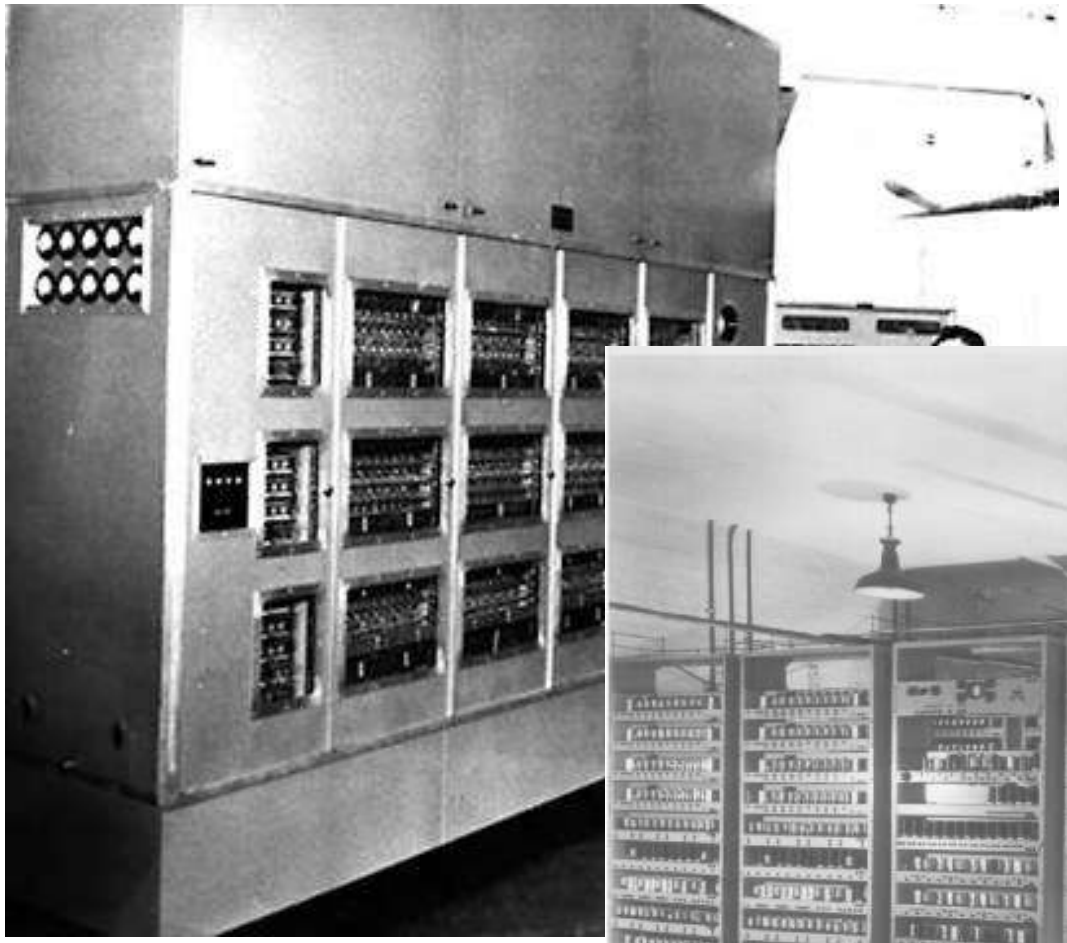

The Illiac Suite (1957)

Music from EDSAC (1960)

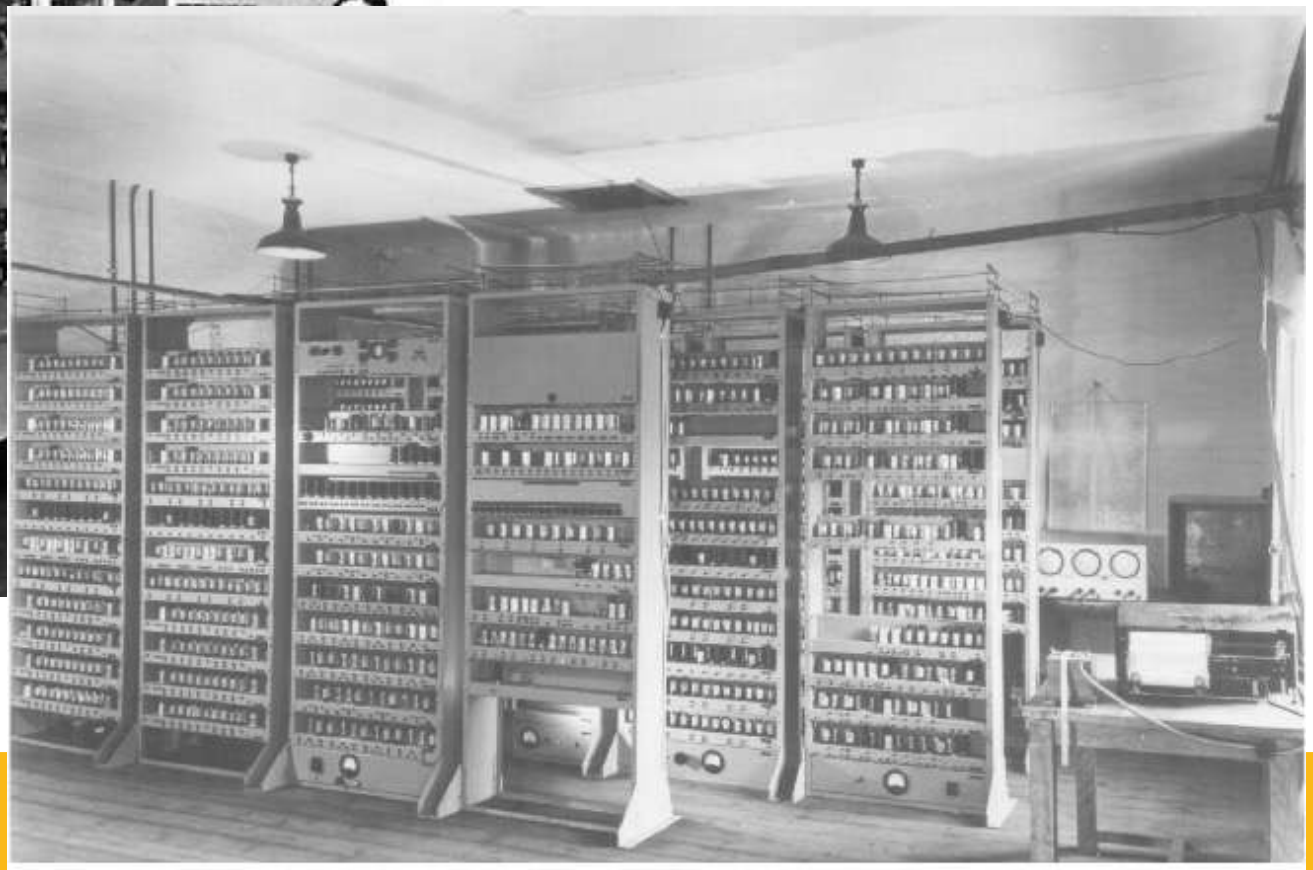

Machines generating music is nothing new

The Illiac Suite (1957)

Music from EDSAC (1960)

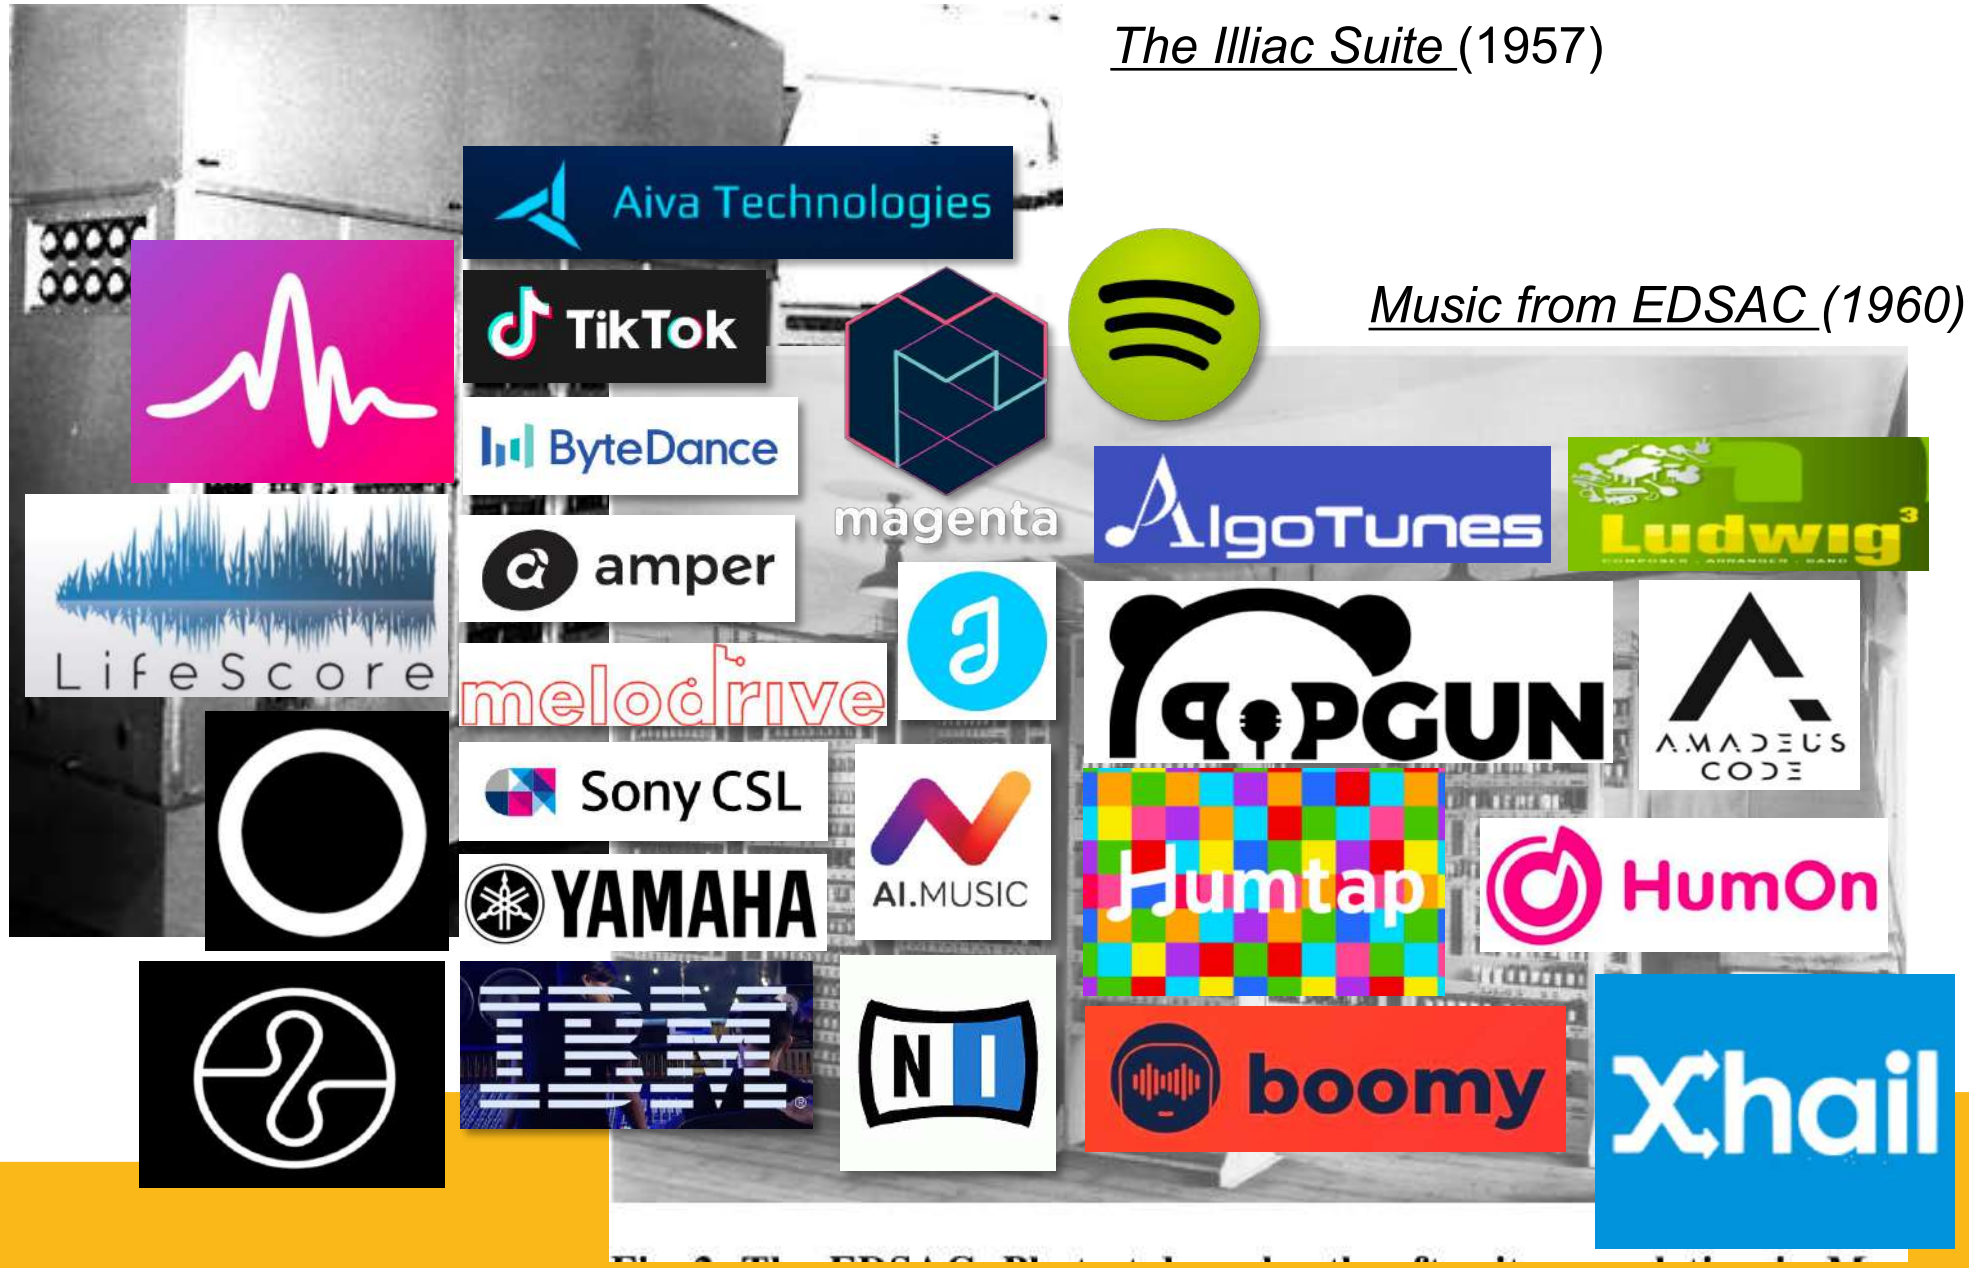

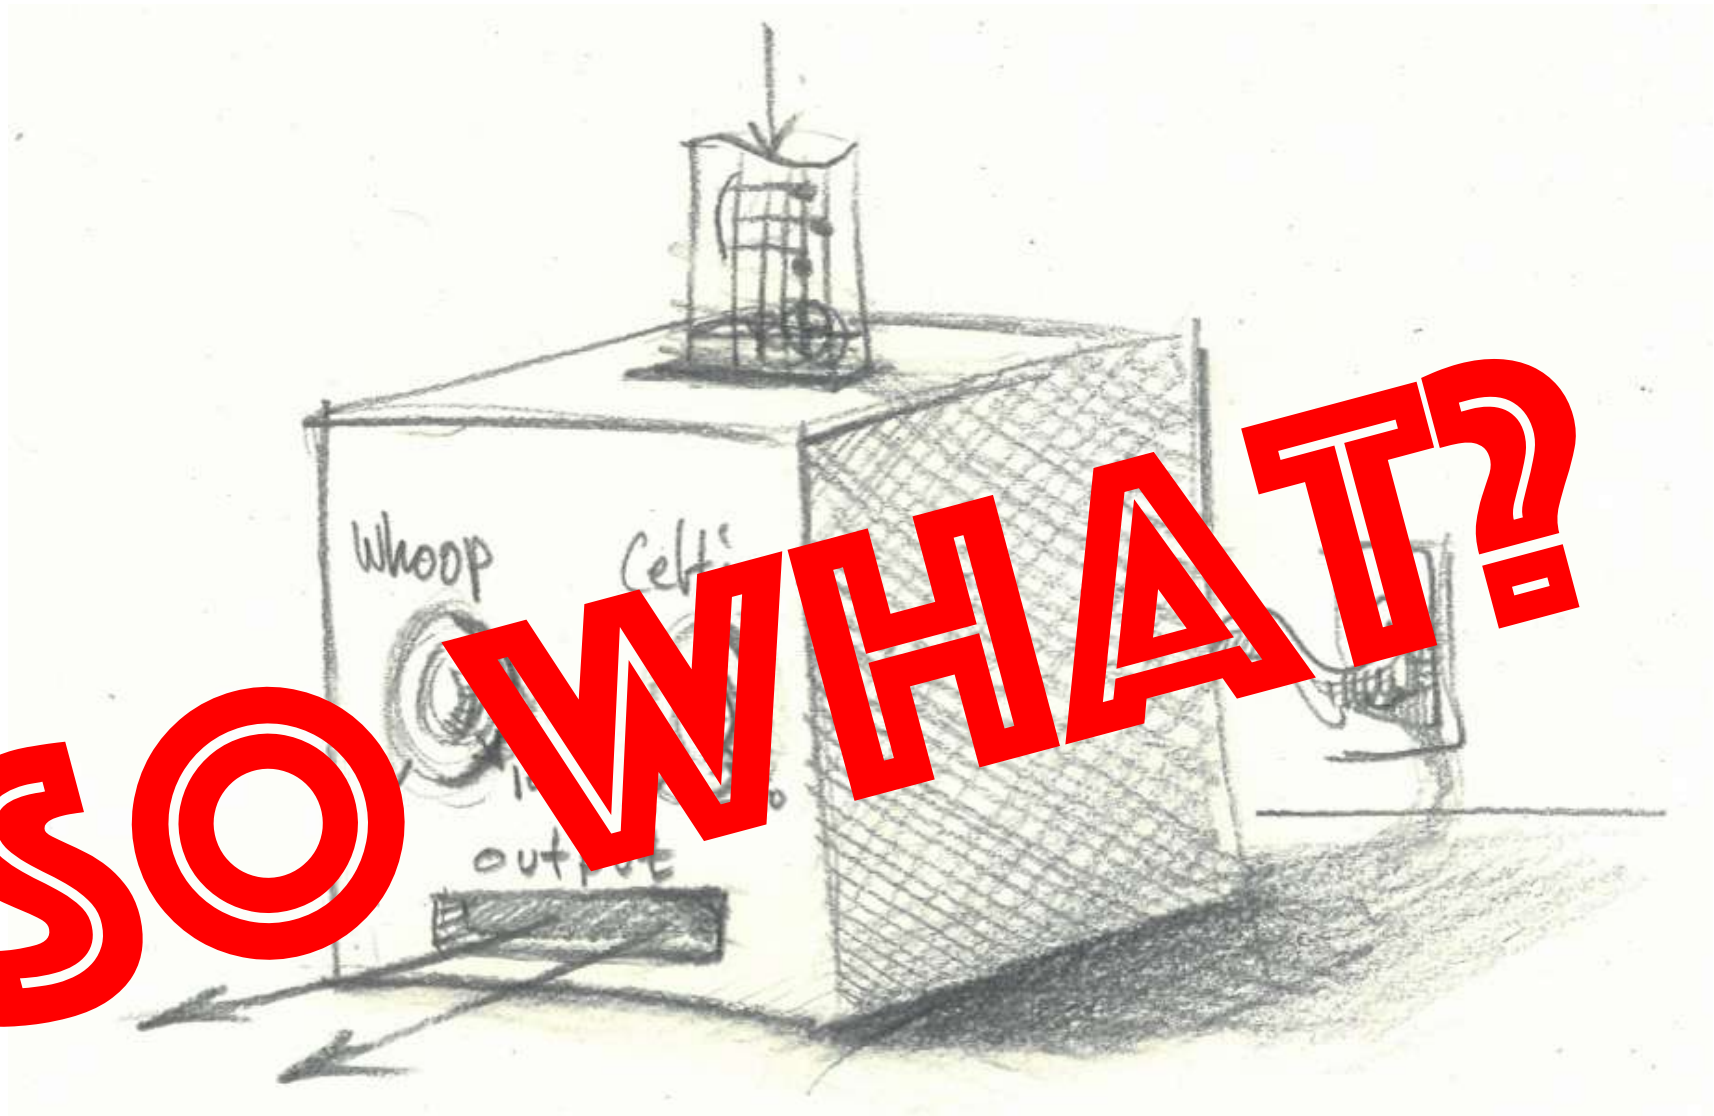
Whoop
output
SO WHAT?

Let’s get serious!

My Saturday morning fun developed into several interesting research questions:

- How far can we get with off-the-shelf (vanilla) ML?

Let's get serious!

My Saturday morning fun developed into several interesting research questions:

- How does the system work?
- Where is its knowledge in its 5.6 million parameters?

Sturm et al., "Machine learning research that matters for music creation: A case study," J. New Music Research 48(1): 36–55, 2018.

folk-rnn (v2): architecture

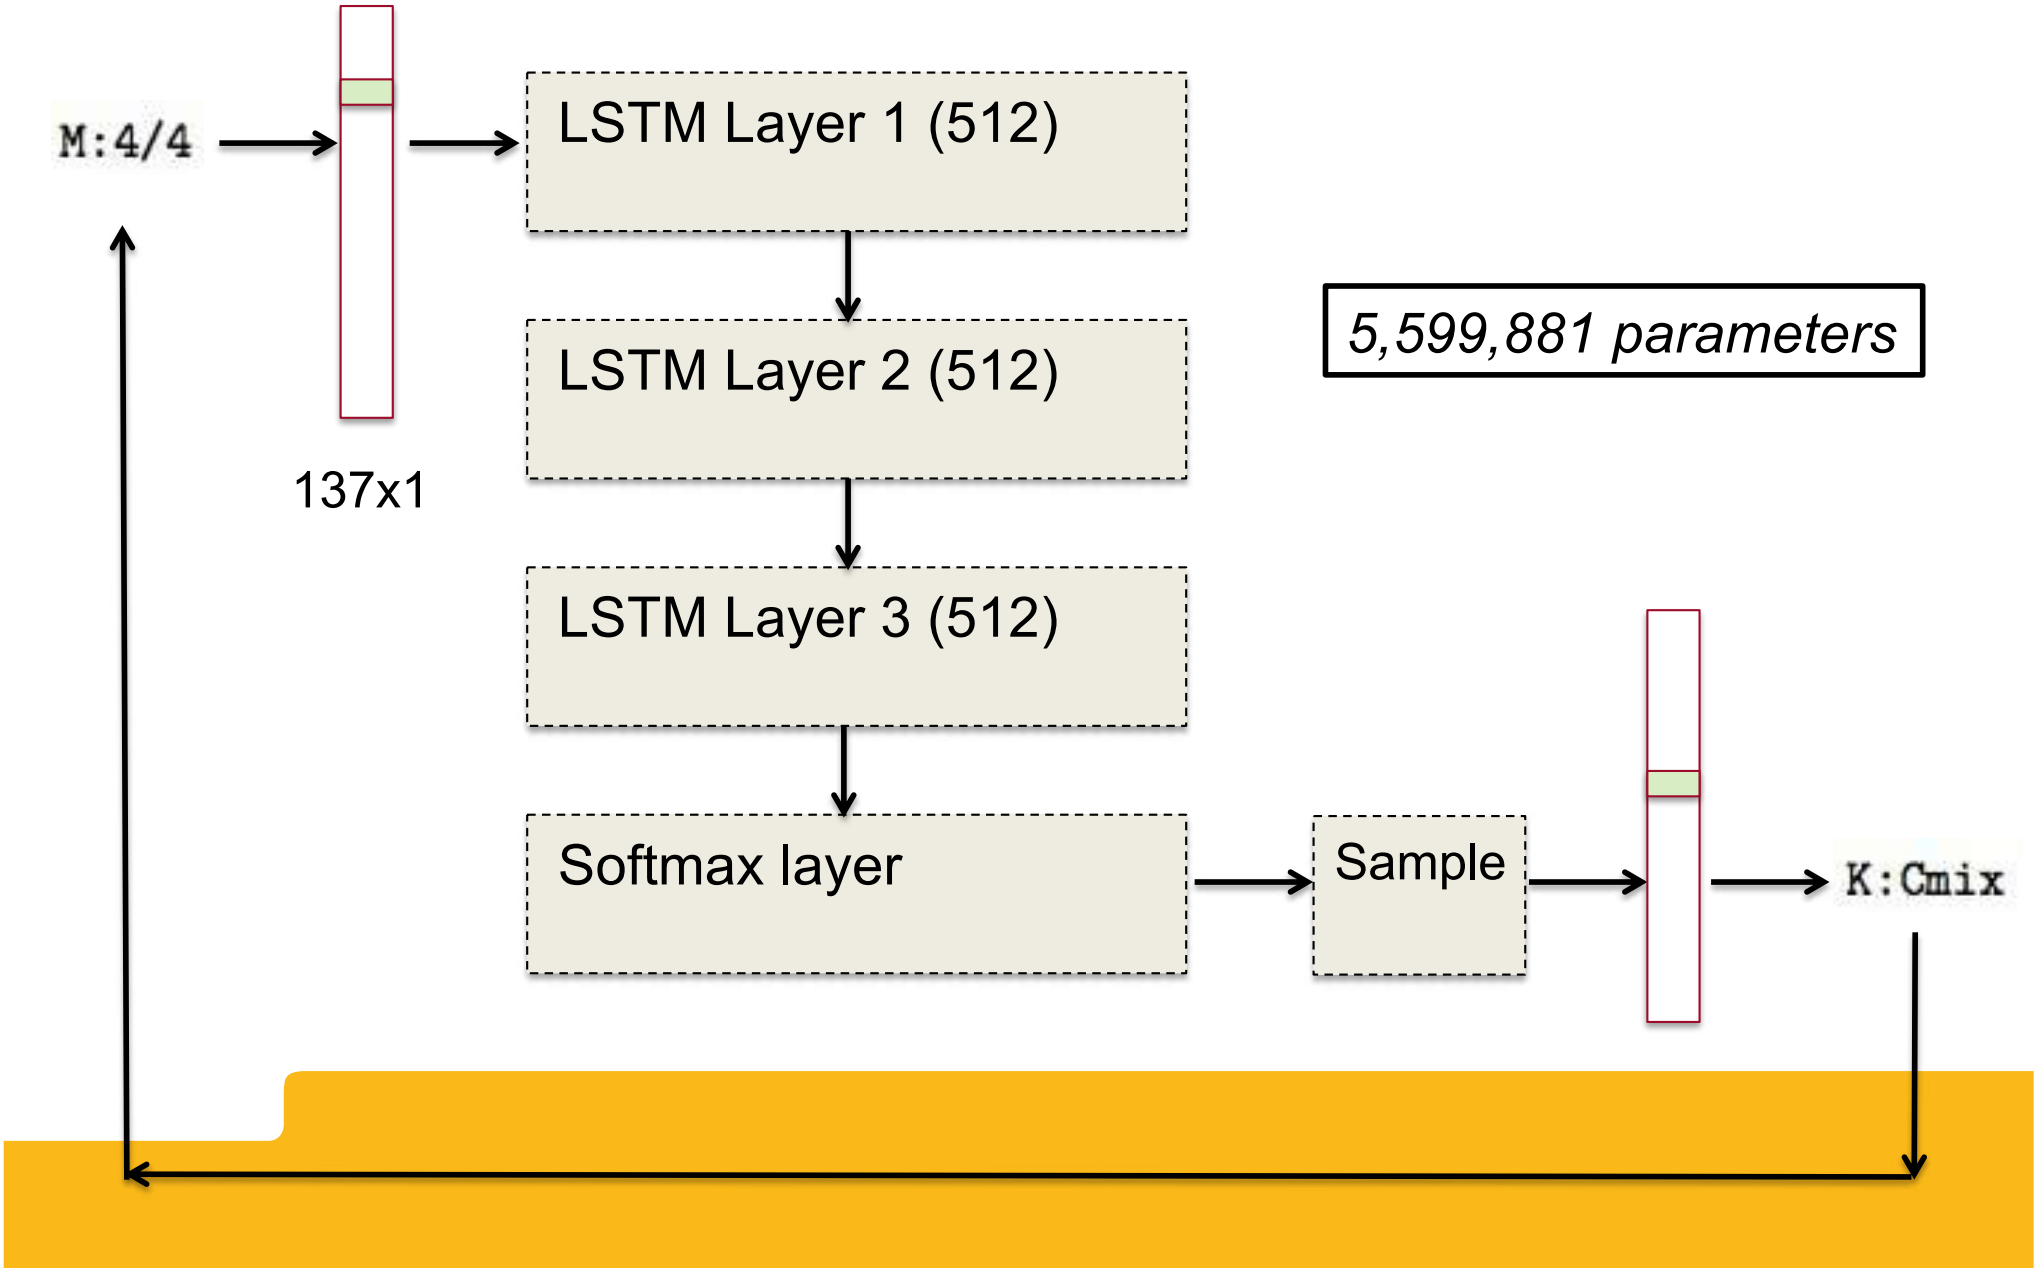

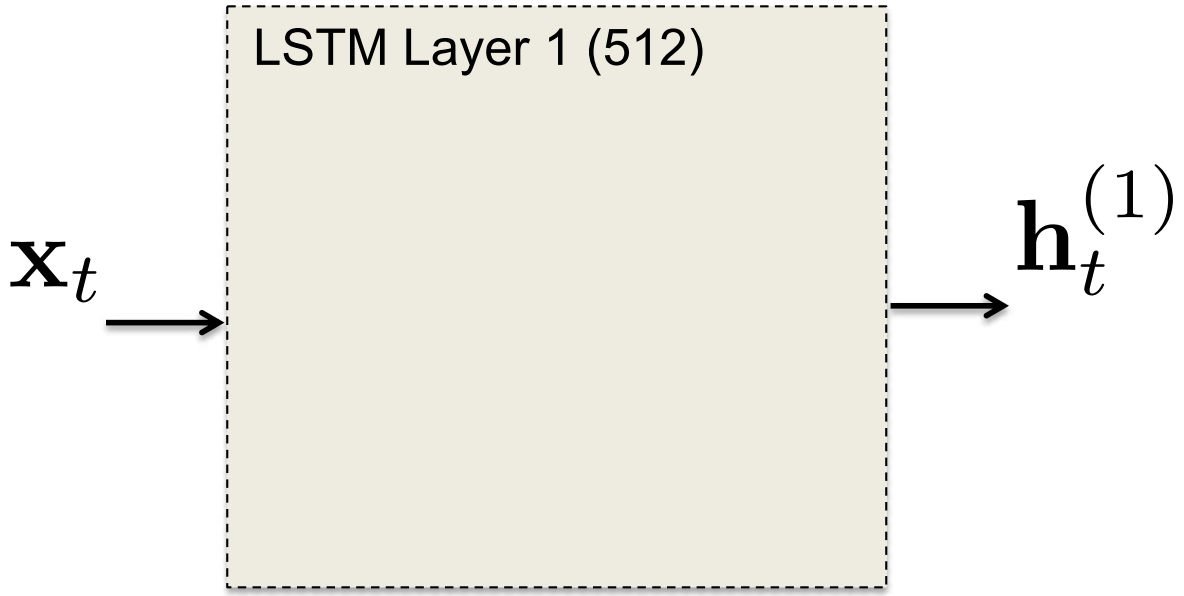

4 gates

$$\mathbf{i}_t^{(1)} \leftarrow \sigma(\mathbf{W}_{xi}^{(1)}\mathbf{x}_t + \mathbf{W}_{hi}^{(1)}\mathbf{h}_{t-1}^{(1)} + \mathbf{b}_i^{(1)})$$

$$\mathbf{f}_t^{(1)} \leftarrow \sigma(\mathbf{W}_{xf}^{(1)}\mathbf{x}_t + \mathbf{W}_{hf}^{(1)}\mathbf{h}_{t-1}^{(1)} + \mathbf{b}_f^{(1)})$$

$$\mathbf{o}_t^{(1)} \leftarrow \sigma(\mathbf{W}_{xo}^{(1)}\mathbf{x}_t + \mathbf{W}_{ho}^{(1)}\mathbf{h}_{t-1}^{(1)} + \mathbf{b}_o^{(1)})$$

$$\mathbf{c}_t^{(1)} \leftarrow \tanh(\mathbf{W}_{xc}^{(1)}\mathbf{x}_t + \mathbf{W}_{hc}^{(1)}\mathbf{h}_{t-1}^{(1)} + \mathbf{b}_c^{(1)}) \odot \mathbf{i}_t^{(1)} + \mathbf{f}_t^{(1)} \odot \mathbf{c}_{t-1}^{(1)}$$

hidden state

$$\mathbf{h}_t^{(1)} \leftarrow \tanh(\mathbf{c}_t^{(1)}) \odot \mathbf{o}_t^{(1)}$$

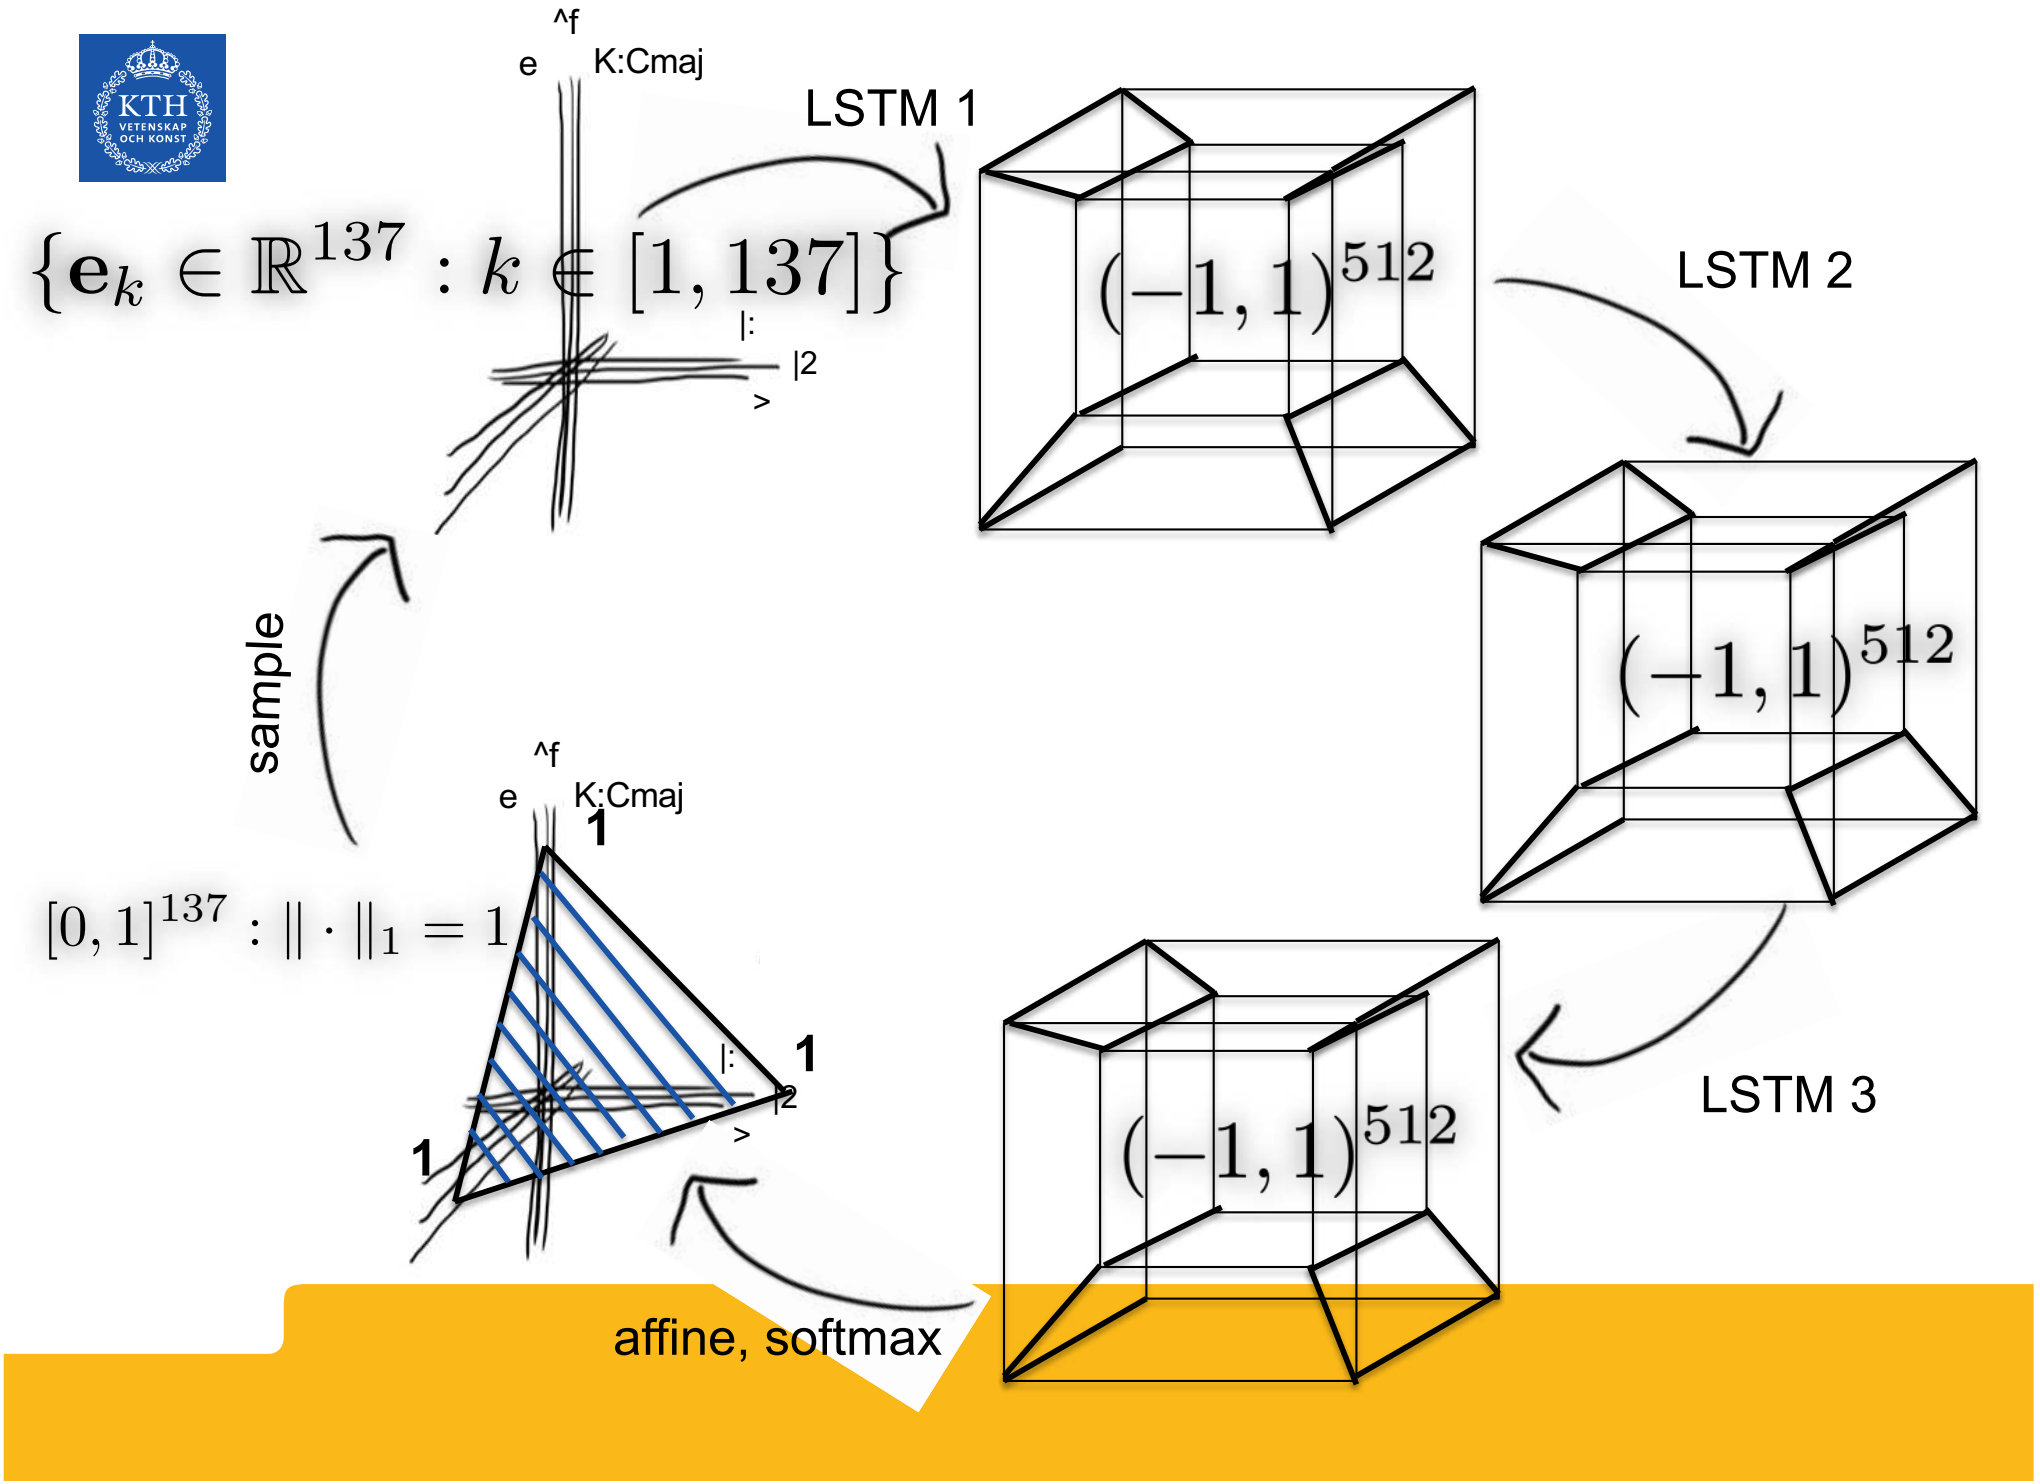
^f
e
K:Cmaj
|:
|2
>
LSTM 1
LSTM 2
LSTM 3
sample
affine, softmax
1
1
1

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

Sturm

ity of London UK

Ab

A *folk-rnn* model is a lor
(LSTM) that generates
evaluated these models
statistical analyses of
use in music practice
their behaviours preci
knowledge is essentia
brating them, and br
paper, we analyse th
layer of a specific
key aspects of the
for instance, that
highly reliant on a
adjust the output
attenuate its prob

The *folk-rnn* soft

How Stuff Works: LSTM Model of Folk Music Transcriptions

Bob L. Sturm [1]

We analyse the parameters of the first layer (L1) of an LSTM created by the *folk-rnn* software from crowd sourced music transcriptions (tunes) (Sturm et al., 2016).[1] The model has three hidden layers each of size 512 LSTM cells, and an input layer and softmax output layer of dimension 137. Each unit of the visible layers corresponds to one element of the vocabulary (token), e.g., `M:6/8` is 6/8 meter; `K:Cmaj` is C major mode; `C` is middle C; and `|` is a measure line. At step t, L1 transforms the input $\mathbf{x}_t \in \{\{0,1\}^{137} : \|\mathbf{x}_t\|_1 = 1\}$ into a hidden state $\mathbf{h}_t^{(1)} \in (-1,1)^{512}$ by (Graves, 2013):

$$\mathbf{i}_t^{(1)} \leftarrow \sigma(\mathbf{W}_{xi}^{(1)}\mathbf{x}_t + \mathbf{W}_{hi}^{(1)}\mathbf{h}_{t-1}^{(1)} + \mathbf{b}_i^{(1)}) \quad (1)$$

$$\mathbf{f}_t^{(1)} \leftarrow \sigma(\mathbf{W}_{xf}^{(1)}\mathbf{x}_t + \mathbf{W}_{hf}^{(1)}\mathbf{h}_{t-1}^{(1)} + \mathbf{b}_f^{(1)}) \quad (2)$$

$$\mathbf{o}_t^{(1)} \leftarrow \sigma(\mathbf{W}_{xo}^{(1)}\mathbf{x}_t + \mathbf{W}_{ho}^{(1)}\mathbf{h}_{t-1}^{(1)} + \mathbf{b}_o^{(1)}) \quad (3)$$

$$\mathbf{c}_t^{(1)} \leftarrow \tanh(\mathbf{W}_{xc}^{(1)}\mathbf{x}_t + \mathbf{W}_{hc}^{(1)}\mathbf{h}_{t-1}^{(1)} + \mathbf{b}_c^{(1)}) \odot \mathbf{i}_t^{(1)} + \mathbf{f}_t^{(1)} \odot \mathbf{c}_{t-1}^{(1)} \quad (4)$$

$$\mathbf{h}_t^{(1)} \leftarrow \tanh(\mathbf{c}_t^{(1)}) \odot \mathbf{o}_t^{(1)}.$$

direction of the fifth rsv it amplifies the probability of the five measure tokens. We find that if we approximate $\mathbf{W}_s$ by its first 30 singular vectors (reducing the subspace in which L3 operates from 137 to 30 dimensions) the model can still generate tunes with plausible local and global characteristics. If we only attenuate rsv 5 at L3 then the model is not able to correctly generate measure tokens. The model seems to rely on measure tokens to generate plausible tunes.

Since $\mathbf{x}_t$ is one-hot corresponding to a token, we can relate each column of $\mathbf{W}_{x*}^{(1)} \in \mathbb{R}^{512\times137}$ to one token, and groups of columns to token types. We want to know how these columns and the subspaces they span in $\mathbb{R}^{512}$ relate to one another. Figure 1 shows the angles between all pairs of columns of each gate matrix. We see clear structures in those of the out and cell gates. Diagonals occur only for the pitch tokens, and relate tokens at octaves and enharmonics. The strongest off-diagonal line in (c) relates pitch tokens over three octaves without

Let's get serious!

My Saturday morning fun developed into several interesting research questions:

- How can we evaluate the resulting models?
- How can such models assist/hinder music creation?

Involving practitioners, organizing workshops and concerts, composing ourselves and inviting others, cherry picking, …

The models definitely assist me! But mileage may vary.

Sturm et al., "Machine learning research that matters for music creation: A case study," J. New Music Research 48(1): 36–55, 2018.

Let’s get serious!

My Saturday morning fun developed into several interesting research questions:

- How can folk music improve ML research?
- What does this mean for traditional music?

MUSAiC: 2020-2025 @ KTH

Music at the Frontiers of Artificial Creativity and Criticism

Confronting questions and challenges at the frontier of the AI disruption/transformation of music

https://musaiclab.wordpress.com

European Research Council
Established by the European Commission

Ai Music Generation Challenge 2020

Build an Ai that generates the most plausible *double jigs*, as judged against the 365 published in O'Neill's *The Dance Music of Ireland* (1907).

- Up to two prizes awarded
- The panel of judges consisted of four (human) experts in Irish traditional music and performance
- A performance of the best AI jigs occurred at the 2020 Joint Conference on AI Music Creativity
- More information

Ai Music Generation Challenge 2020

Build an Ai ti… …gs,

- …
- …in
- …
- …

https://youtu.be/KSoSyoEx6hc

Ai Music Generation Challenge 2020

Build an Ai that generates the most plausible *double jigs*, as judged against the 365 published in O'Neill's *The Dance Music of Ireland* (1907).

Three aims:

- to promote meaningful approaches to evaluating music Ai;
- to see how music Ai research can benefit from considering traditional music, and how traditional music can benefit from music Ai research;
- to facilitate discussions about the ethics of music Ai research applied to traditional music practices.

Ai Music Generation Challenge 2021

Build an Ai that generates the most plausible …

https://github.com/boblsturm/aimusicgenerationchallenge2021

Ai Music Generation Challenge 2021

Build an Ai that generates the most plausible …

Slängpolskas!

https://github.com/boblsturm/aimusicgenerationchallenge2021

En själ-lös “svensk” låt!

folkRNN

generate a folk tune with a recurrent neural network

PRESS TO GENERATE TUNE

Compose

MODEL

folkwiki.se

TEMPERATURE 1

SEED 352794

METER 3/4

MODE D Minor

INITIAL ABC

Enter start of tune in ABC notation

https://themachinefolksession.org/tune/551

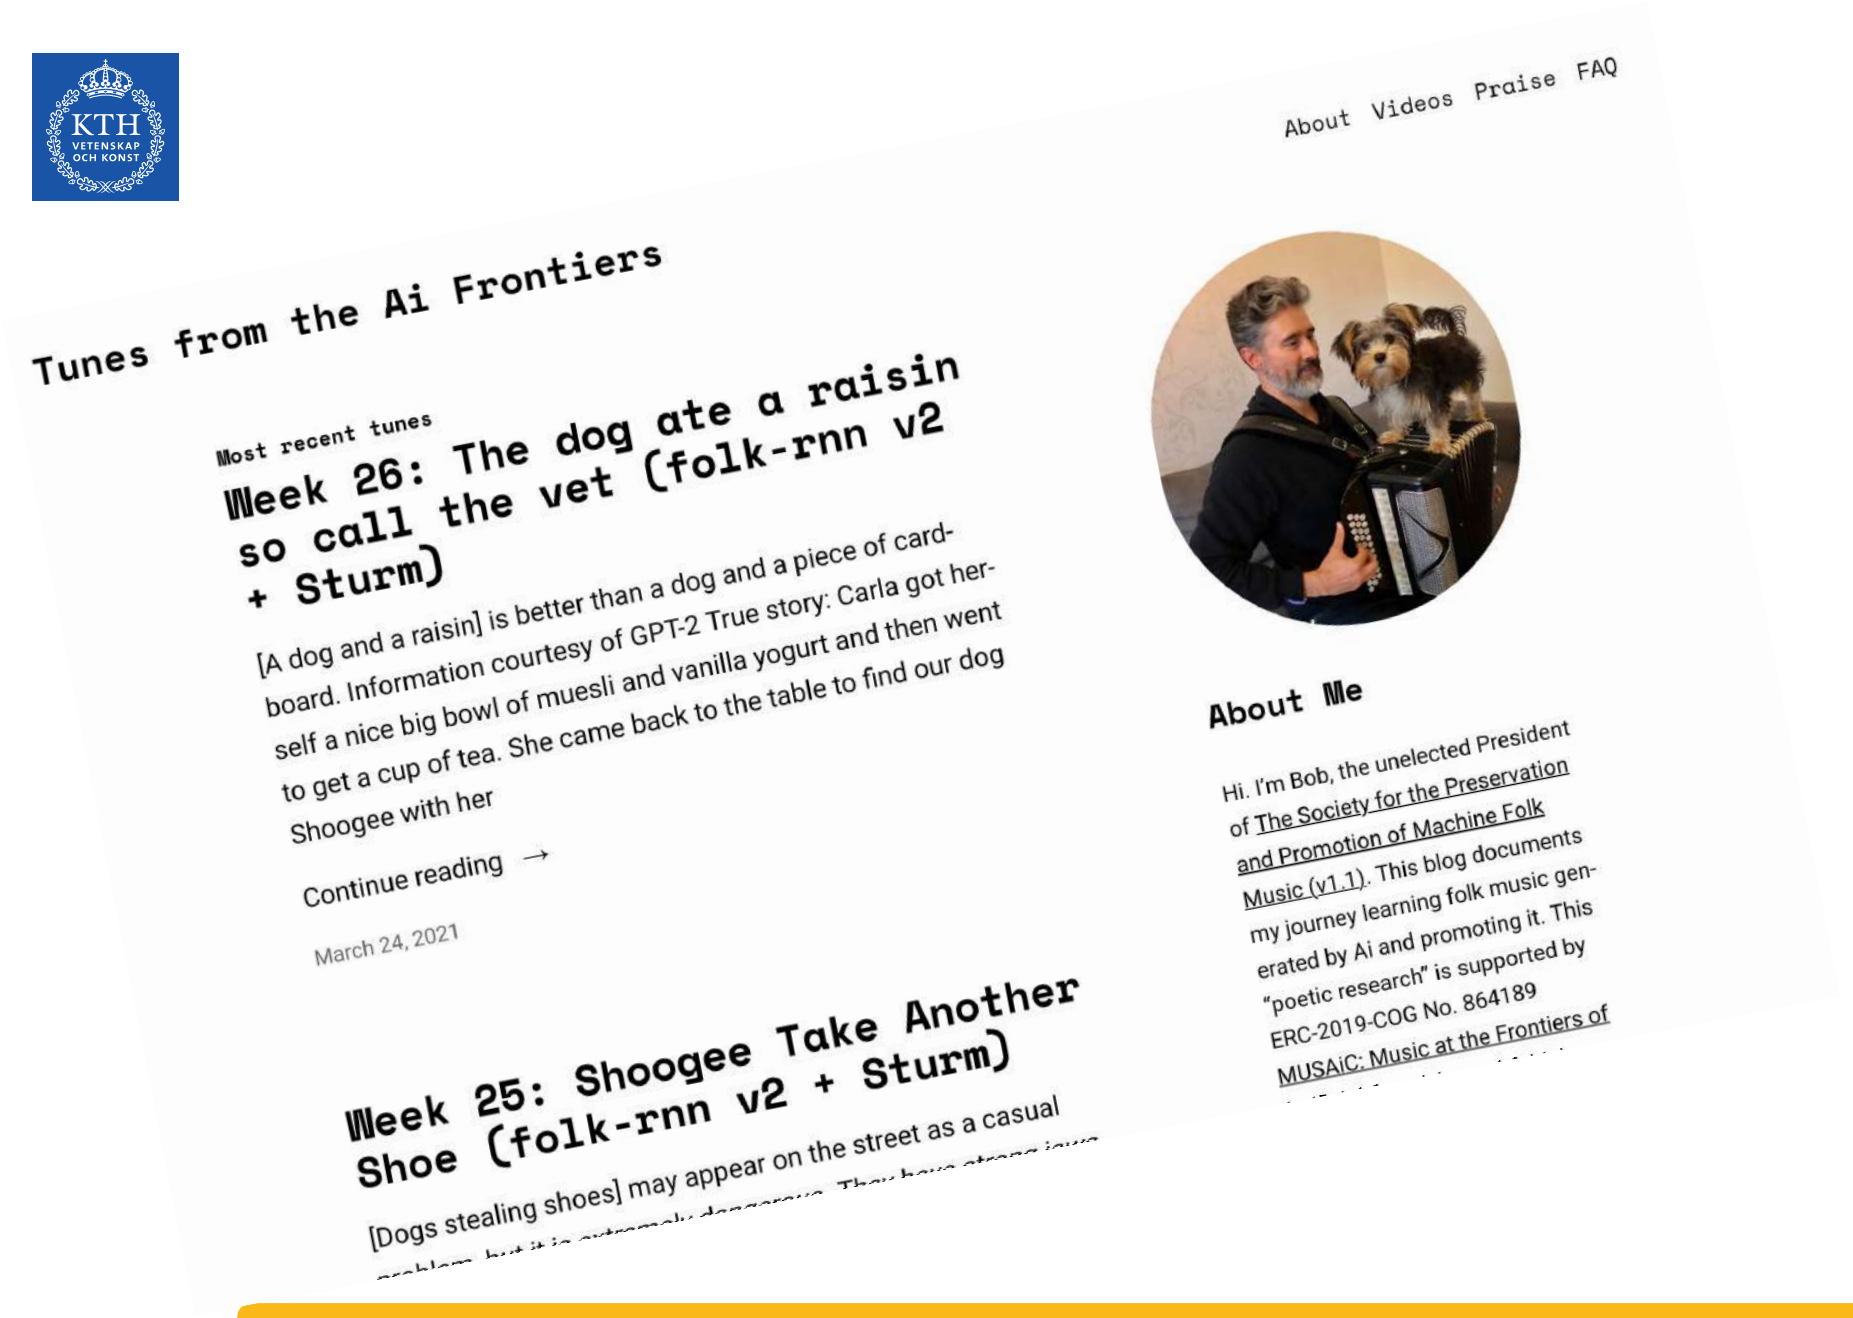

https://tunesfromtheaifrontiers.wordpress.com/

Want to do more?

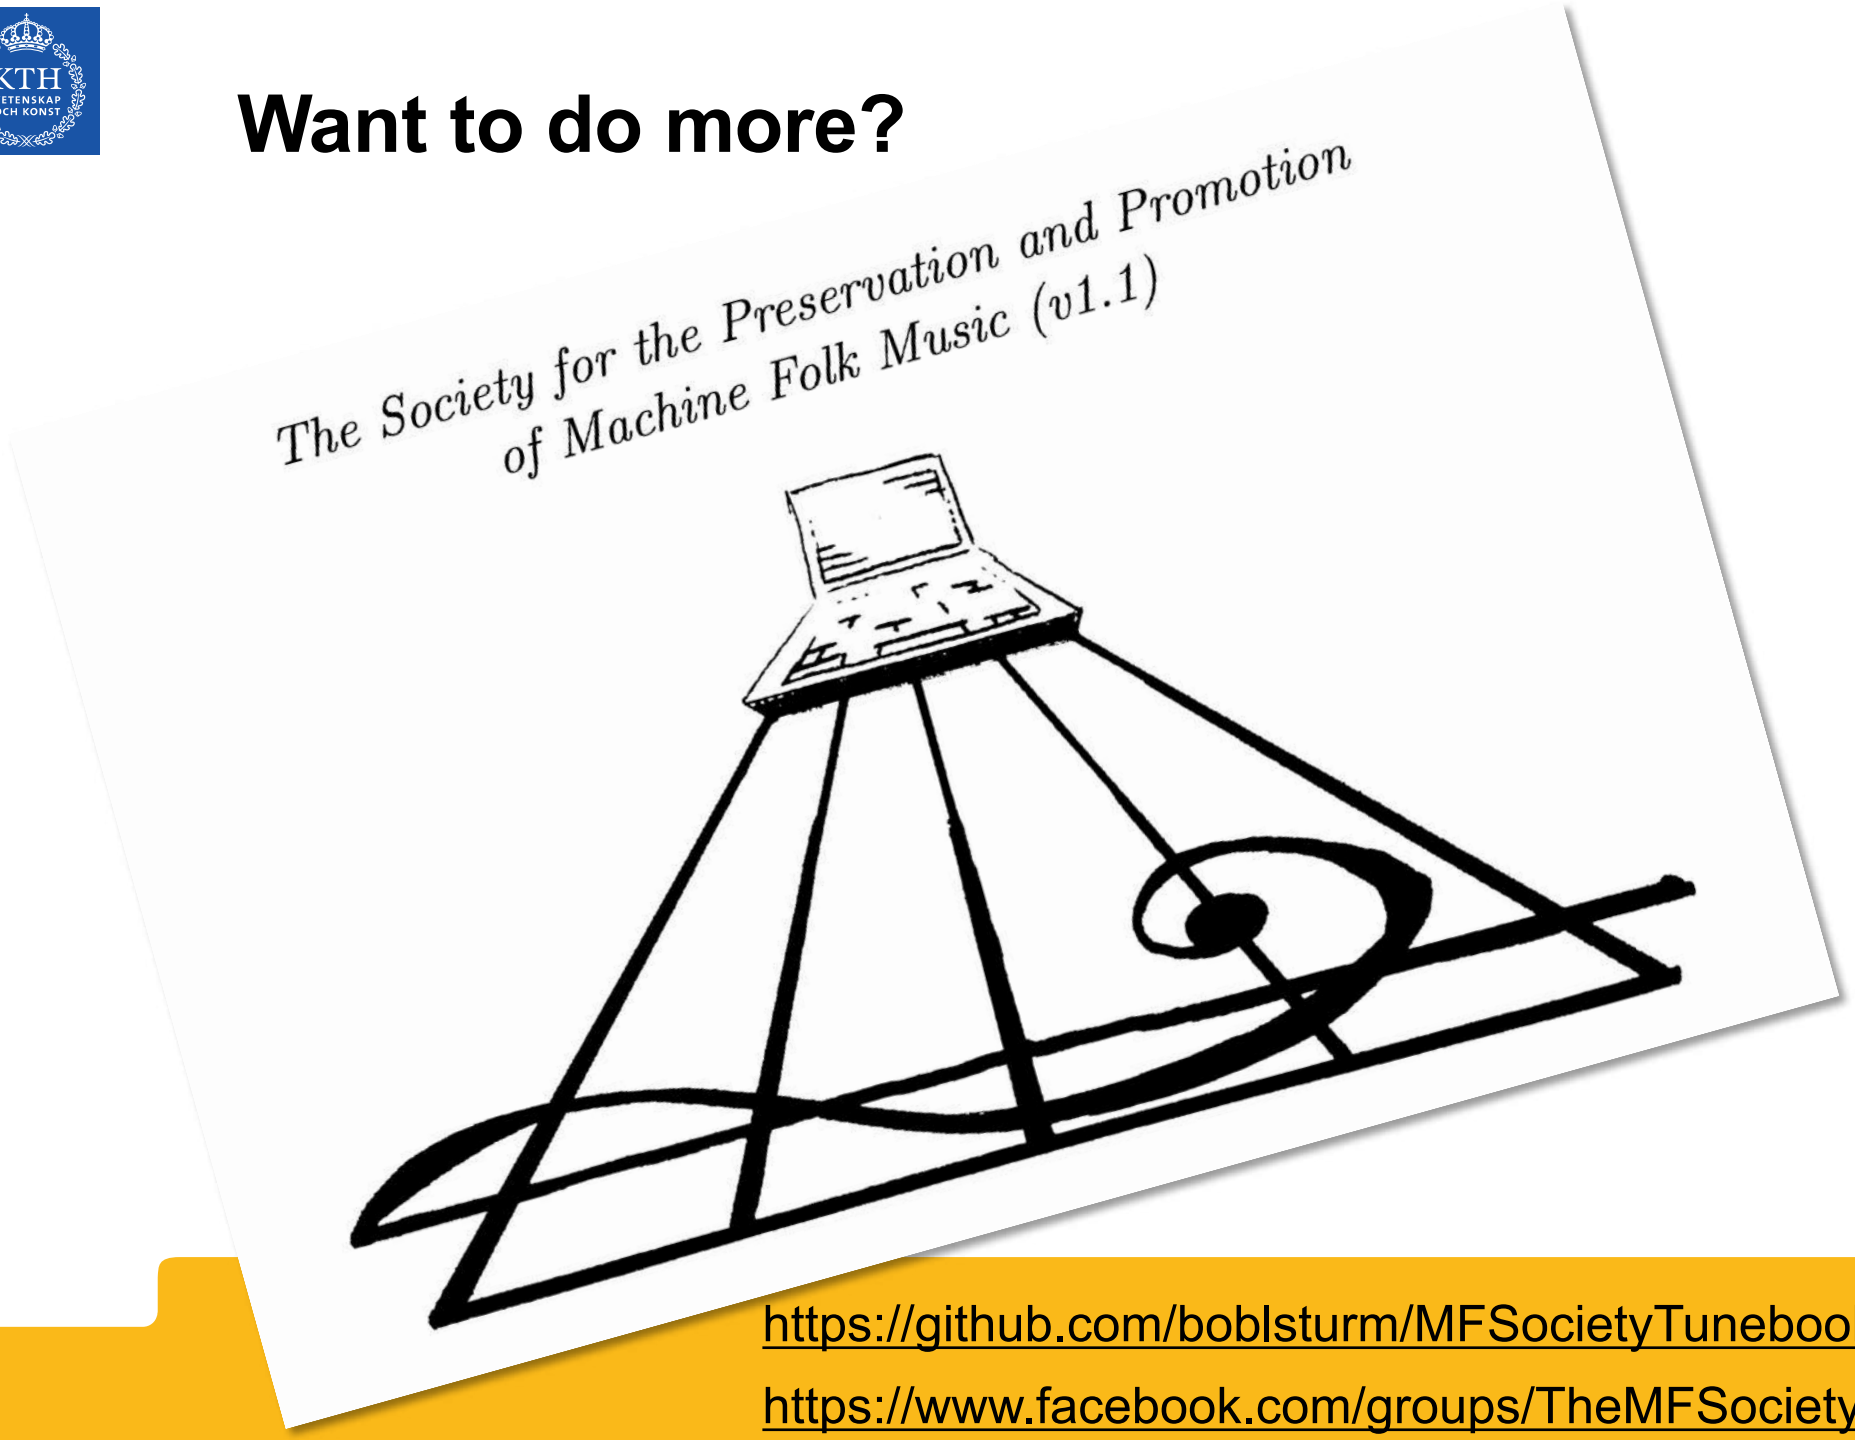

https://github.com/bobIsturm/MFSocietyTunebook

https://www.facebook.com/groups/TheMFSociety