goto; copenhagen

GOTO Copenhagen 2019

Conference Nov. 18 - 20

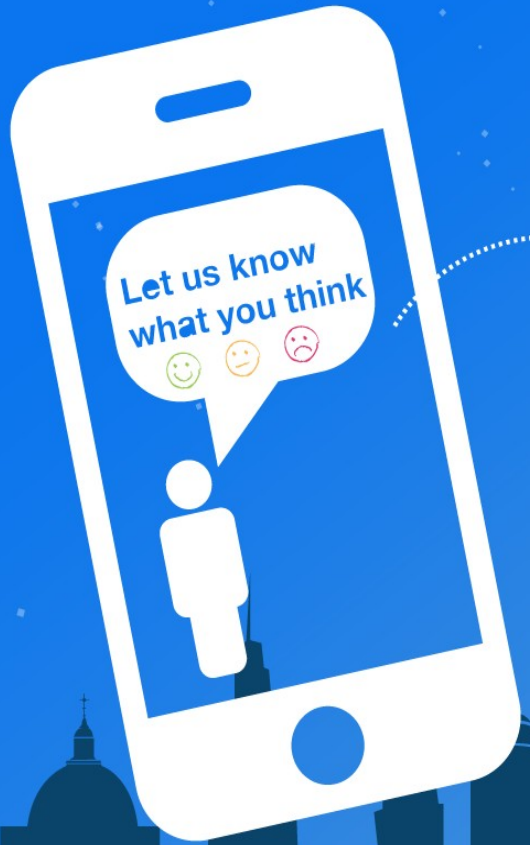

Click ‘Rate Session’ to rate session and ask questions.

Enjoying the conference background music?

BachBot

BachBot.com

0:04 0:14

1	Sample-1	▶ 2,244
2	Happy Birthday	▶ 2,091
3	Twinkle Twinkle Little Star	▶ 6,524
4	Sample-2	▶ 1,304
5	Sample-3	▶ 800

View 8 tracks

Like Share More

https://soundcloud.com/bachbot

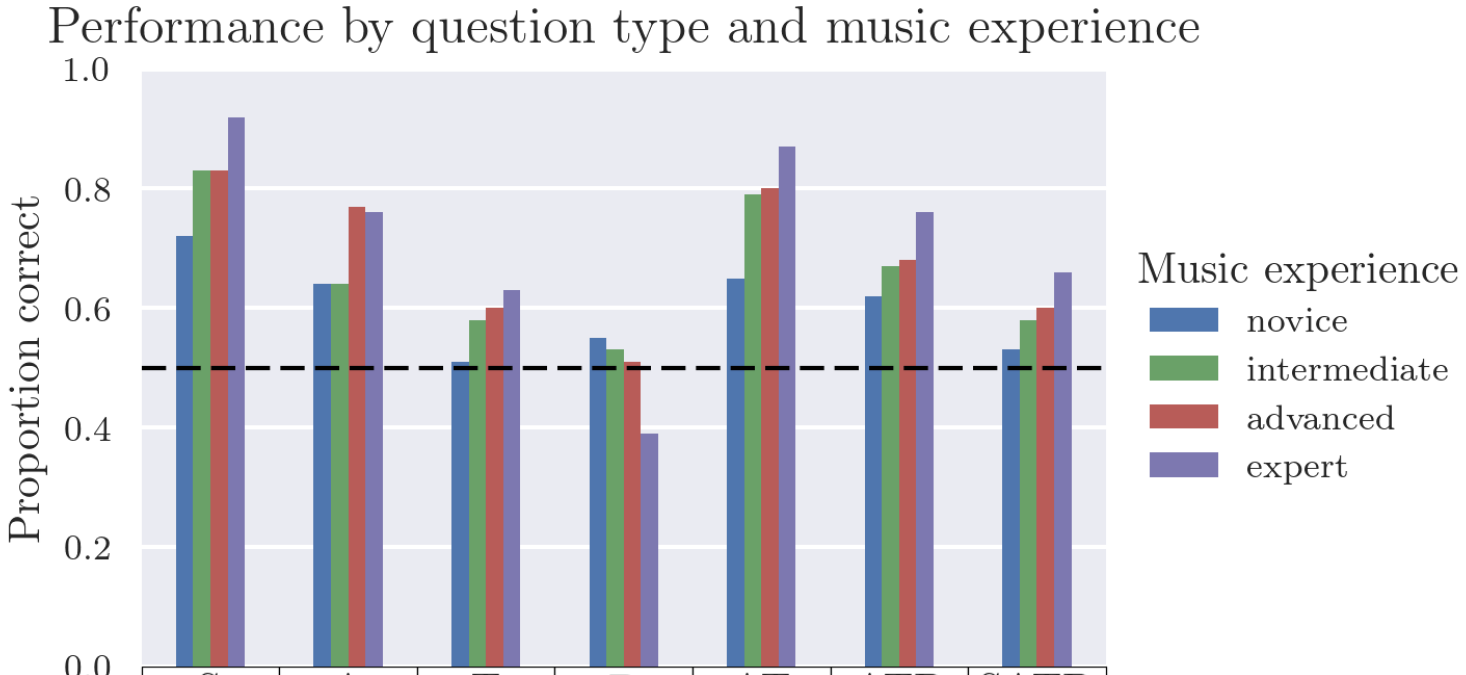

	S	A	T	B	AT	ATB	SATB
novice	0.72	0.64	0.51	0.55	0.65	0.62	0.53
intermediate	0.83	0.64	0.58	0.53	0.79	0.67	0.58
advanced	0.83	0.77	0.6	0.51	0.8	0.68	0.6
expert	0.92	0.76	0.63	0.39	0.87	0.76	0.66

Results from musical Turing test at bachbot.com

BachBot

Composing Bach Chorales Using Deep Learning

Feynman Liang

GOTO Copenhagen, 18 November 2019

The big questions

- Where is the frontier of computational creativity?
- How much has deep learning advanced automatic composition?
- How do we evaluate generative models?

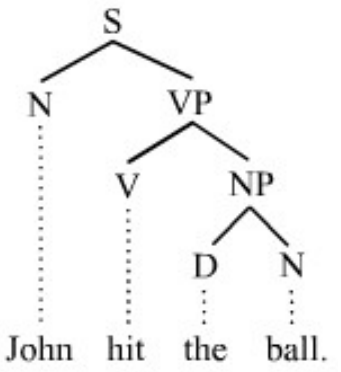

vs

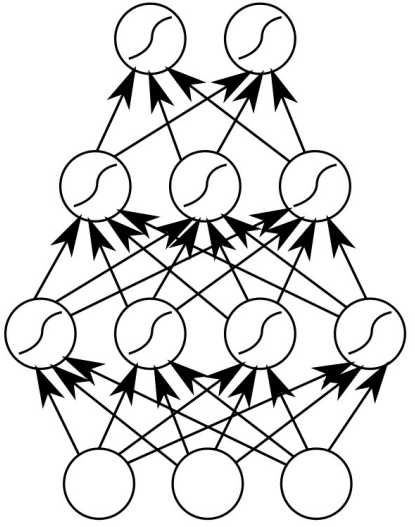

TL;DR

- Deep recurrent neural network model for music capable of polyphonic automatic stylistic composition and harmonization
- Learns music theory without prior knowledge
- In a musical Turing test with 1779 participants, performance is only 5% better than random chance

Open source; try it for yourself!

feynmanliang.github.io/bachbot-slides

BUILD YOUR OWN BACHBOT!

- Install `docker` from docker.com
- Pull the image:

```
docker pull fliang/bachbot:aibtb
```

- Prepare data, train model:

```
docker run --name bachbot -it fliang/bachbot:cor
bachbot datasets prepare
bachbot datasets concatenate_corpus scratch/BWV-
bachbot make_h5
bachbot train
```

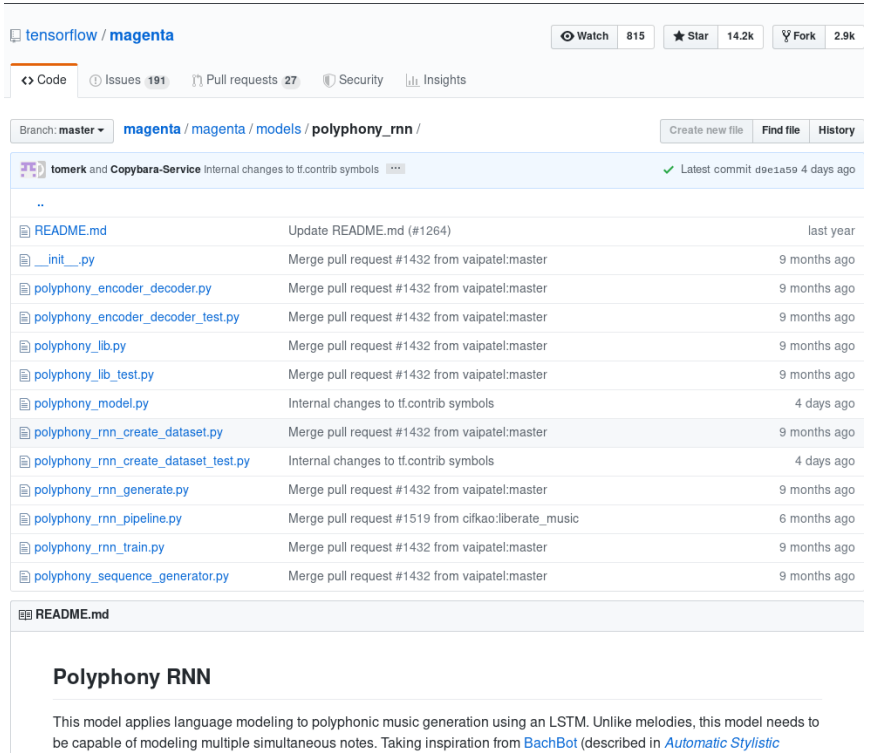

Overview

- Sequence modelling for music
- Deep recurrent neural networks (RNNs)
- Discoveries and demos

Motivating example for sequence modelling

Sequence modelling of text

The quick brown fox jumps _____

Question: What word comes next?

Sequence modelling of text

The quick brown fox jumps _____

The quick brown fox jumps over
The quick brown fox jumps around
The quick brown fox jumps lackadaisically

Sequence modelling of text

The quick brown fox jumps _____

P(over | the quick brown fox jumps) = 75%
P(around | the quick brown fox jumps) = 24%
P(lackadaisically | the quick brown fox jumps) = 1%

Sequence modelling of text

The quick brown fox jumps _____

P(over | the quick brown fox jumps) = 75%
P(around | the quick brown fox jumps) = 24%
P(lackadaisically | the quick brown fox jumps) = 1%

Question: Any potential problems?

The 2-gram sequence model

The quick brown fox jumps _____

P(over | jumps) = 90%
P(around | jumps) = 8%
P(behind | jumps) = 2%

Generating text using the 2-gram sequence model:

I
I am
am enjoying
enjoying GOTO
GOTO Copenhagen
I am enjoying GOTO Copenhagen

n-gram models trained on Hamlet

- 1-gram
  - To him swallowed confess hear both. Which. Of save on trail for are ay device androte life have
- 2-gram
  - Why dost stand forth thy canopy, forsooth; he is this palpable hit the King Henry. Liveking. Follow.
- 4-gram
  - King Henry. What! I will go seek the traitor Gloucester. Exeunt some of the watch. Agreat banquet serv’d in;

Examples taken from Dan Jurafsky, https://web.stanford.edu/~jurafsky/slp3/slides/LM_4.pdf

Music Primer

Modern music notation

Pitch: how “high” or “low” a note is

Duration: how “long” a note is

Polyphony: multiple simultaneous voices

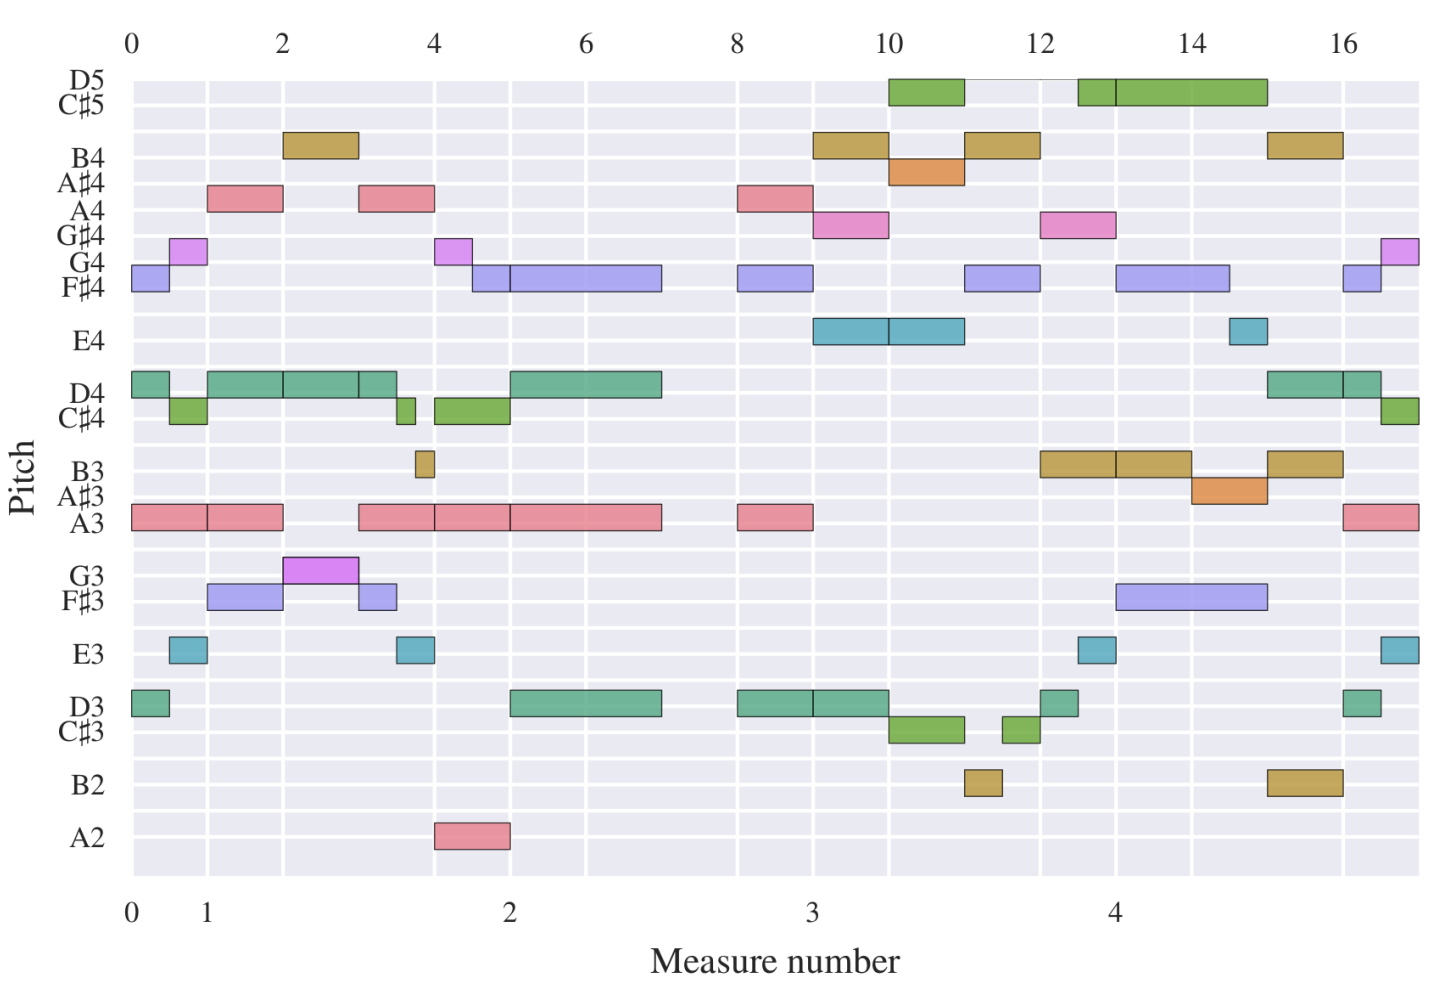

Piano roll: convenient computational representation

Fermatas and **phrasing**

From Bach Chorales to Sequences

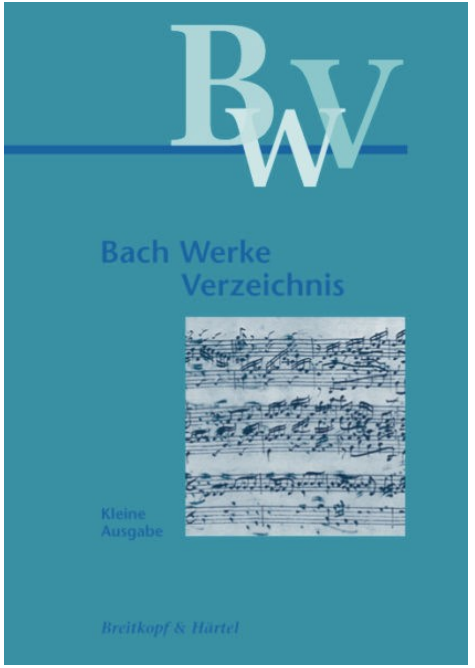

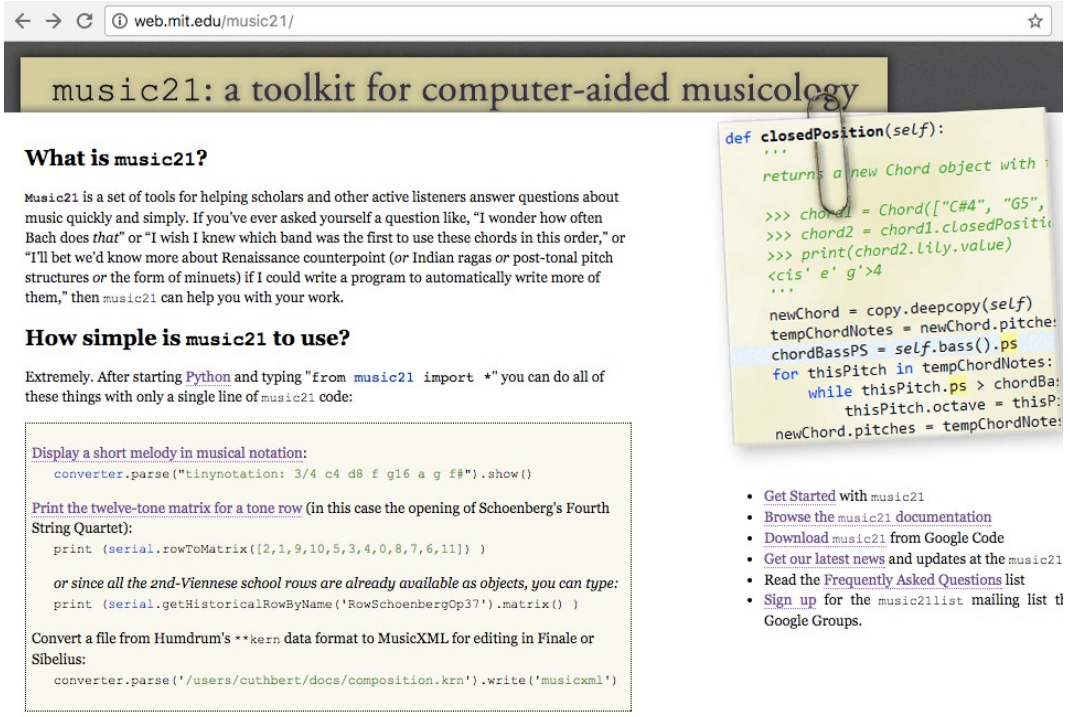

http://web.mit.edu/music21/

Transpose to Cmaj/Amin (convenience)
Quantize to 16th notes (computational)

Transposition preserves relative pitches

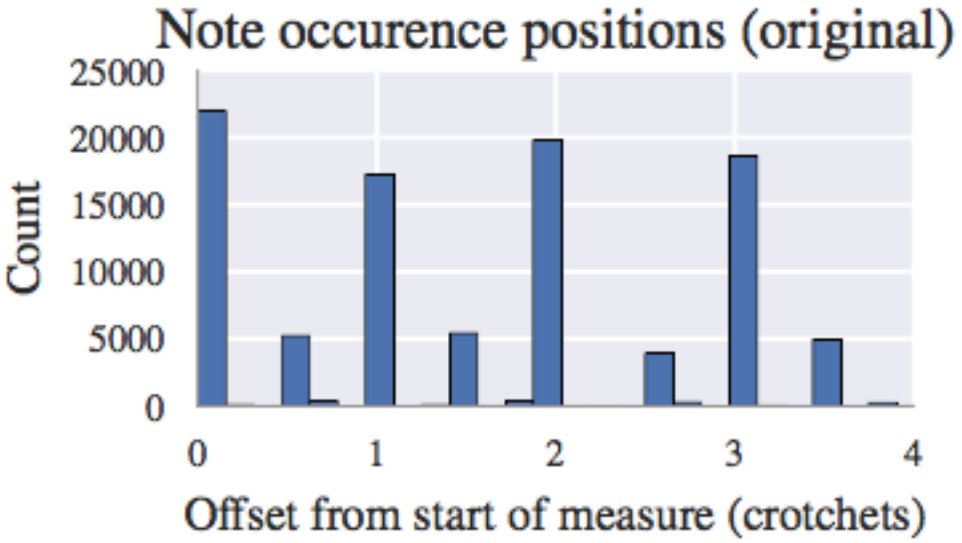

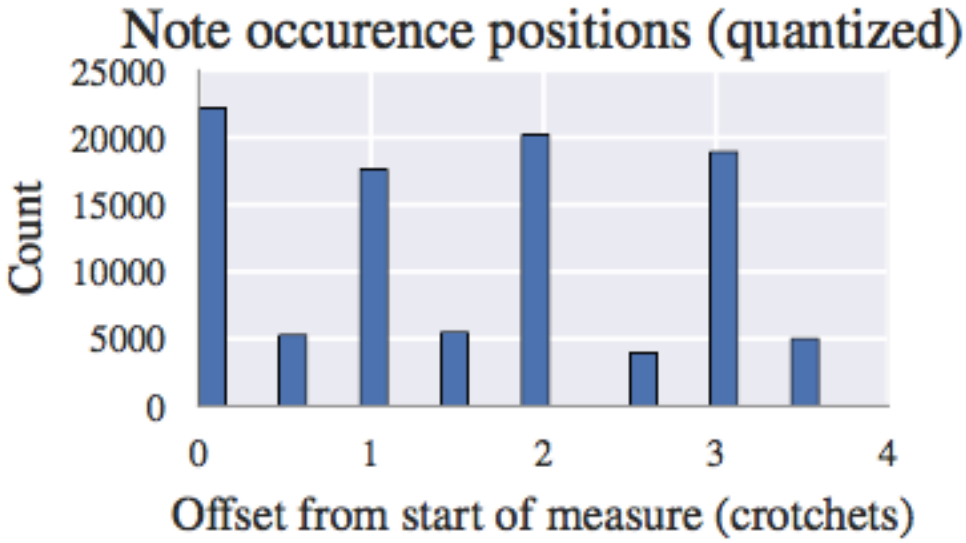

Quantization to 16th notes: affects less than 0.2% of dataset

Handling polyphony

Question: How many chords can be constructed from 4 voices, each with 128 pitches?

Handling polyphony

Question: How many chords can be constructed from 4 voices, each with 128 pitches?

Answer: $O(128^4)$! Data sparsity issue…

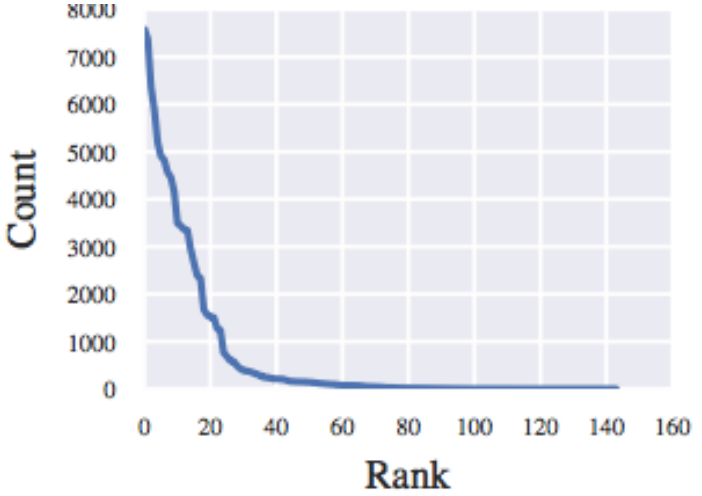

Serialize in SATB order

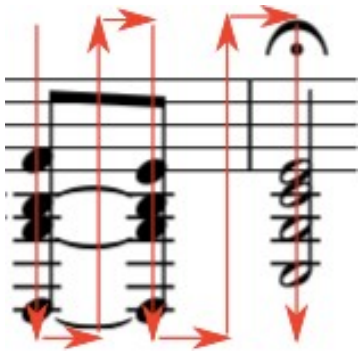

START
(62, True)
(56, True)
(52, True)
(47, True)
|||
(60, False)
(56, True)
(52, True)
(47, True)
|||

(.)
(60, False)
(57, False)
(52, False)
(50, False)
|||
(.)
(60, True)
(57, True)
(52, True)
(50, True)
|||

$O(128^4) \Rightarrow O(128)$ vocab. size!

Recurrent neural networks (RNNs)

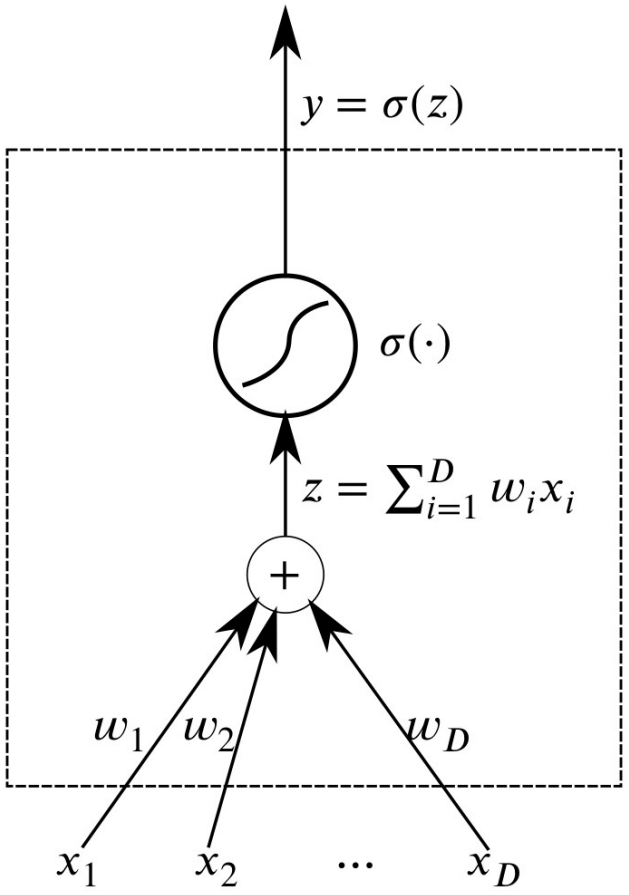

Neuron

Input **x**, output **y**, parameters **w**, activations **z**

$$\sigma\left(\boldsymbol{W}_{L,L-1}\sigma\left(\boldsymbol{W}_{L-1,L-2}\cdots\sigma\left(\boldsymbol{W}_{2,1}\boldsymbol{x}\right)\cdots\right)\right)$$

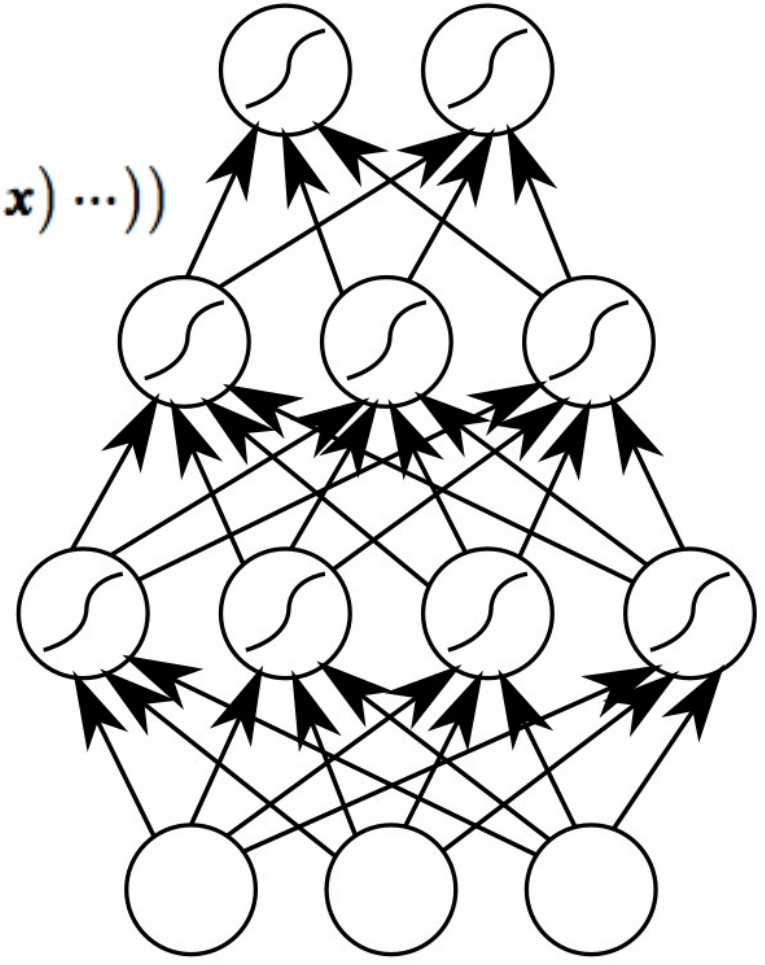

Output Layer

Hidden Layers

Input Layer

Feedforward neural network

$$\boldsymbol{h}_t = \boldsymbol{W}_{xh}\sigma_{xh}\left(\boldsymbol{x}_t\right) + \boldsymbol{W}_{hh}\sigma_{hh}\left(\boldsymbol{h}_{t-1}\right)$$

$$\boldsymbol{y}_t = \boldsymbol{W}_{hy}\sigma_{hy}\left(\boldsymbol{h}_t\right)$$

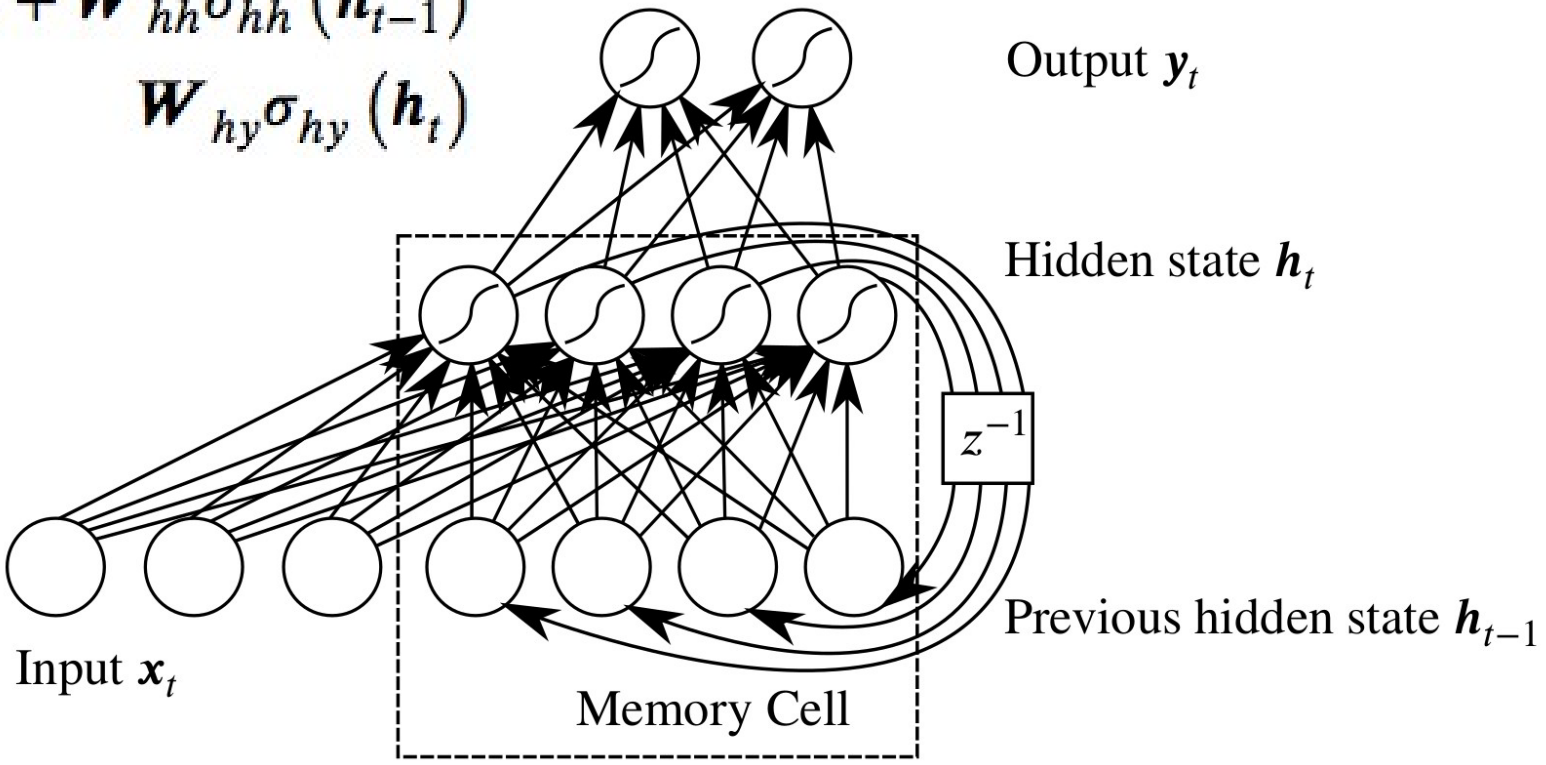

Memory cell

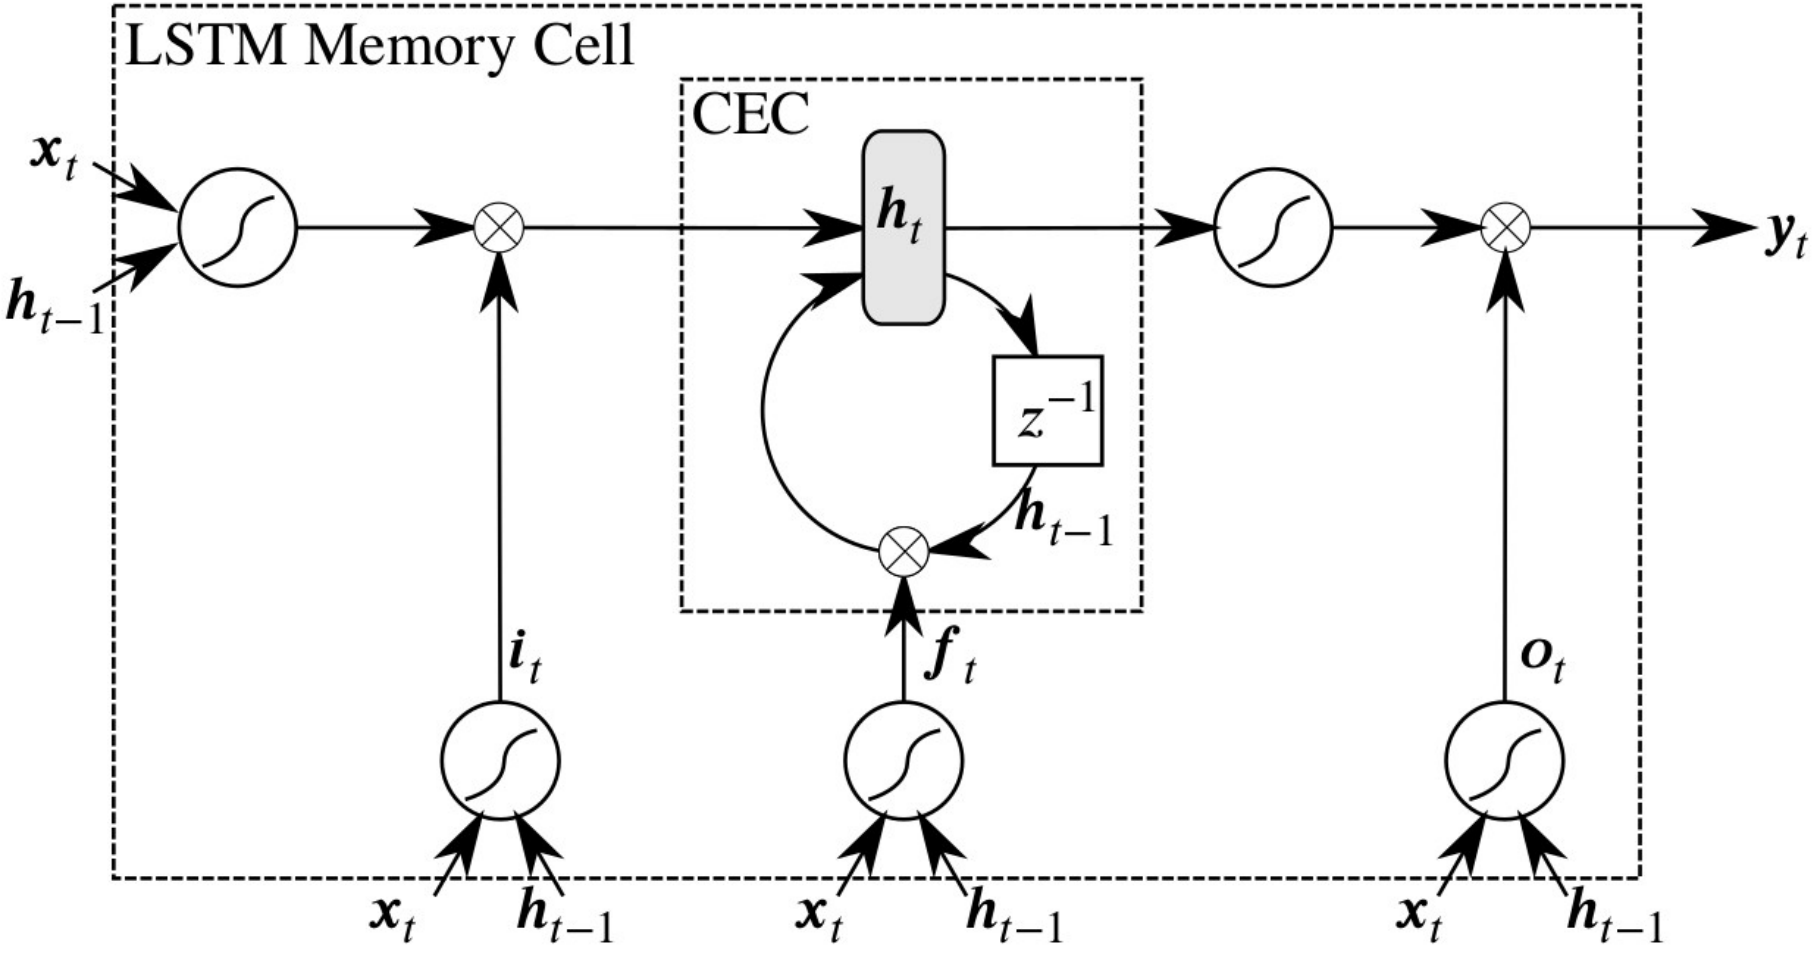

Long short-term memory (**LSTM**) cell

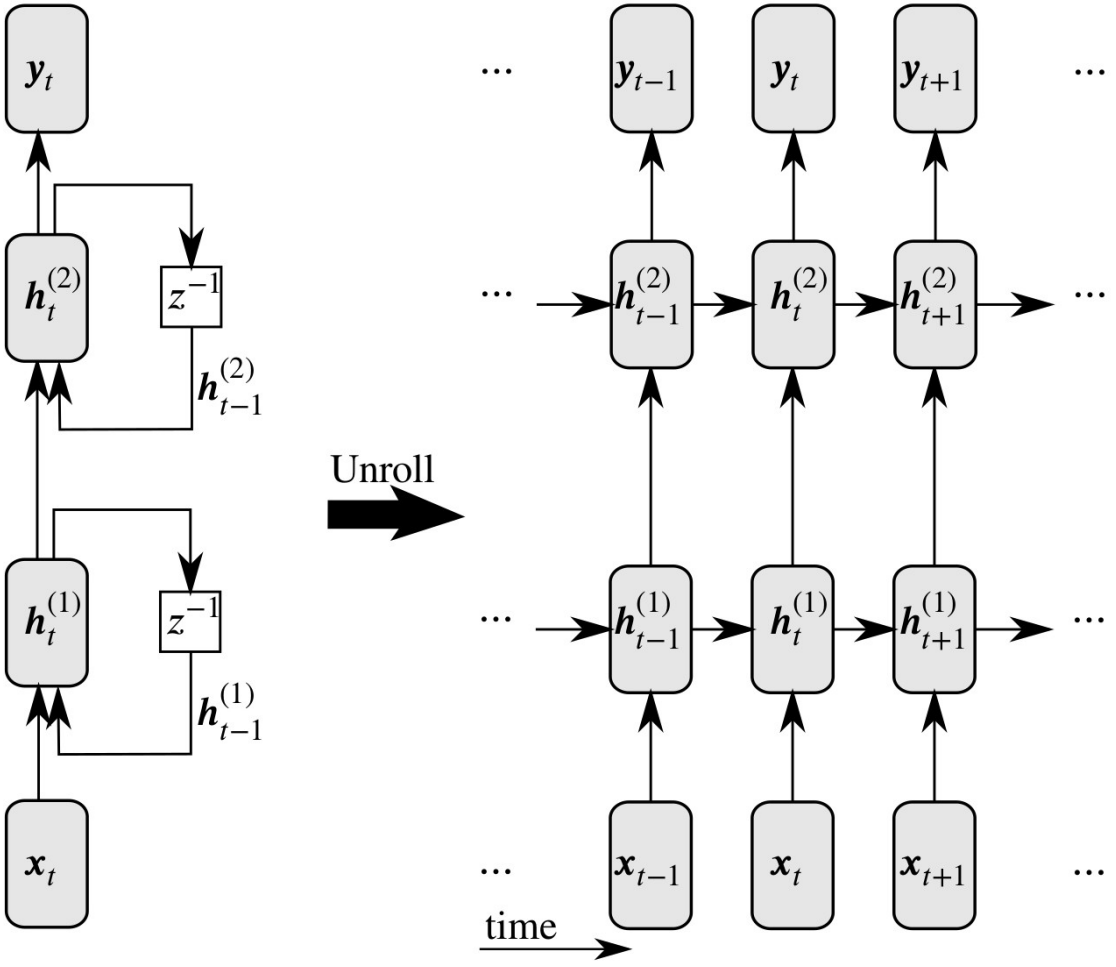

Stacking memory cells to form a deep RNN
Unrolling for training

Sequential prediction training criteria

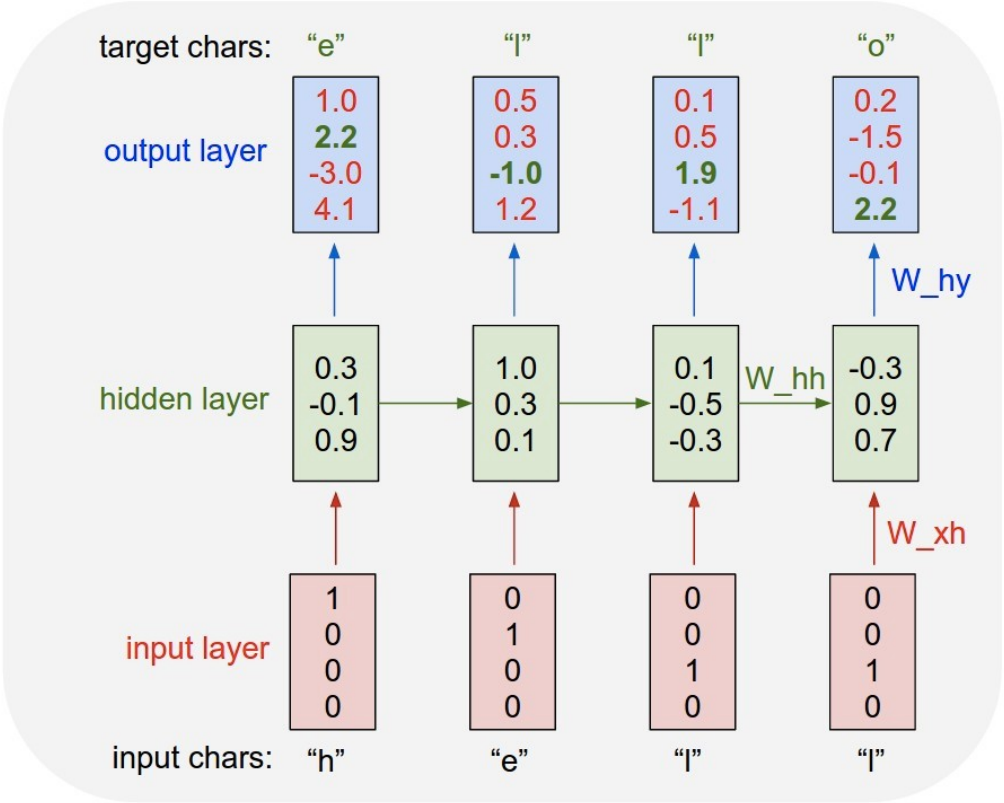

https://karpathy.github.io/2015/05/21/rnn-effectiveness/

Training and optimizing BachBot

Training BachBot

RNN dynamics

$$\tilde{P}(x_{t+1}|x_t, h_{t-1})$$

Initial state (all 0s)

$$+ \quad h_0$$

Prob. distr. over sequences

$$y_t = \tilde{P}(x_t|x_{1:t-1})$$

Optimize the RNN parameters...

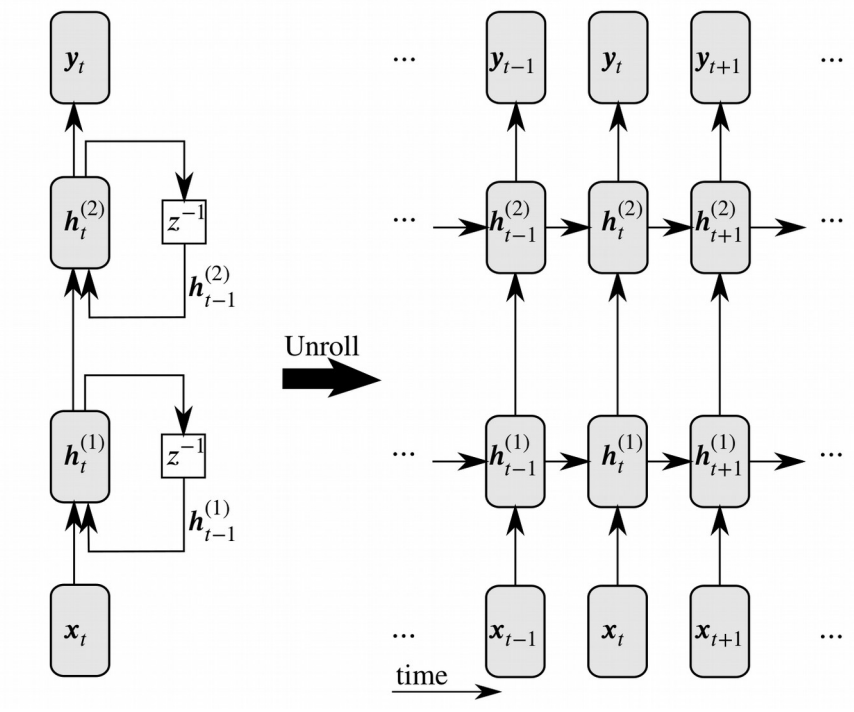

...in order to maximize the probability of the real Bach chorales.

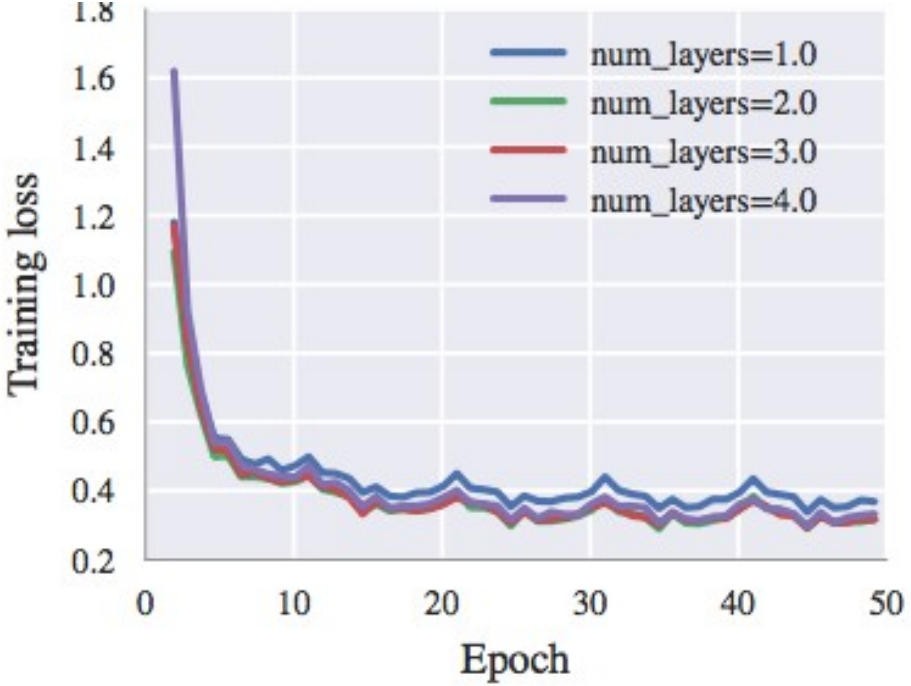

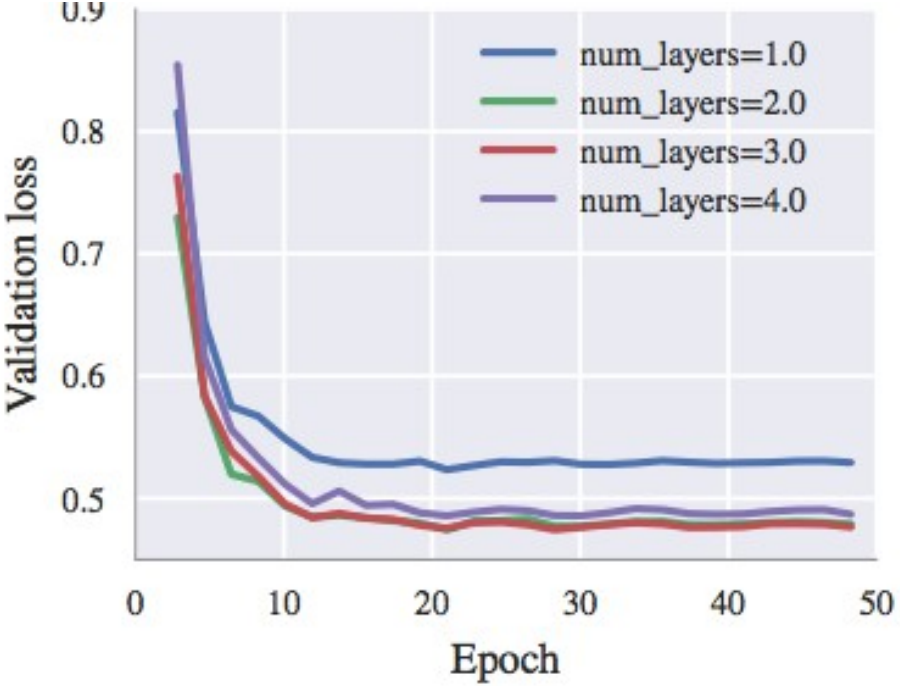

Deeper RNNs do better (to a certain point)

Samples from various points during training

100 minibatches:

500 minibatches:

2000 minibatches:

5000 minibatches:

Discoveries and Demos

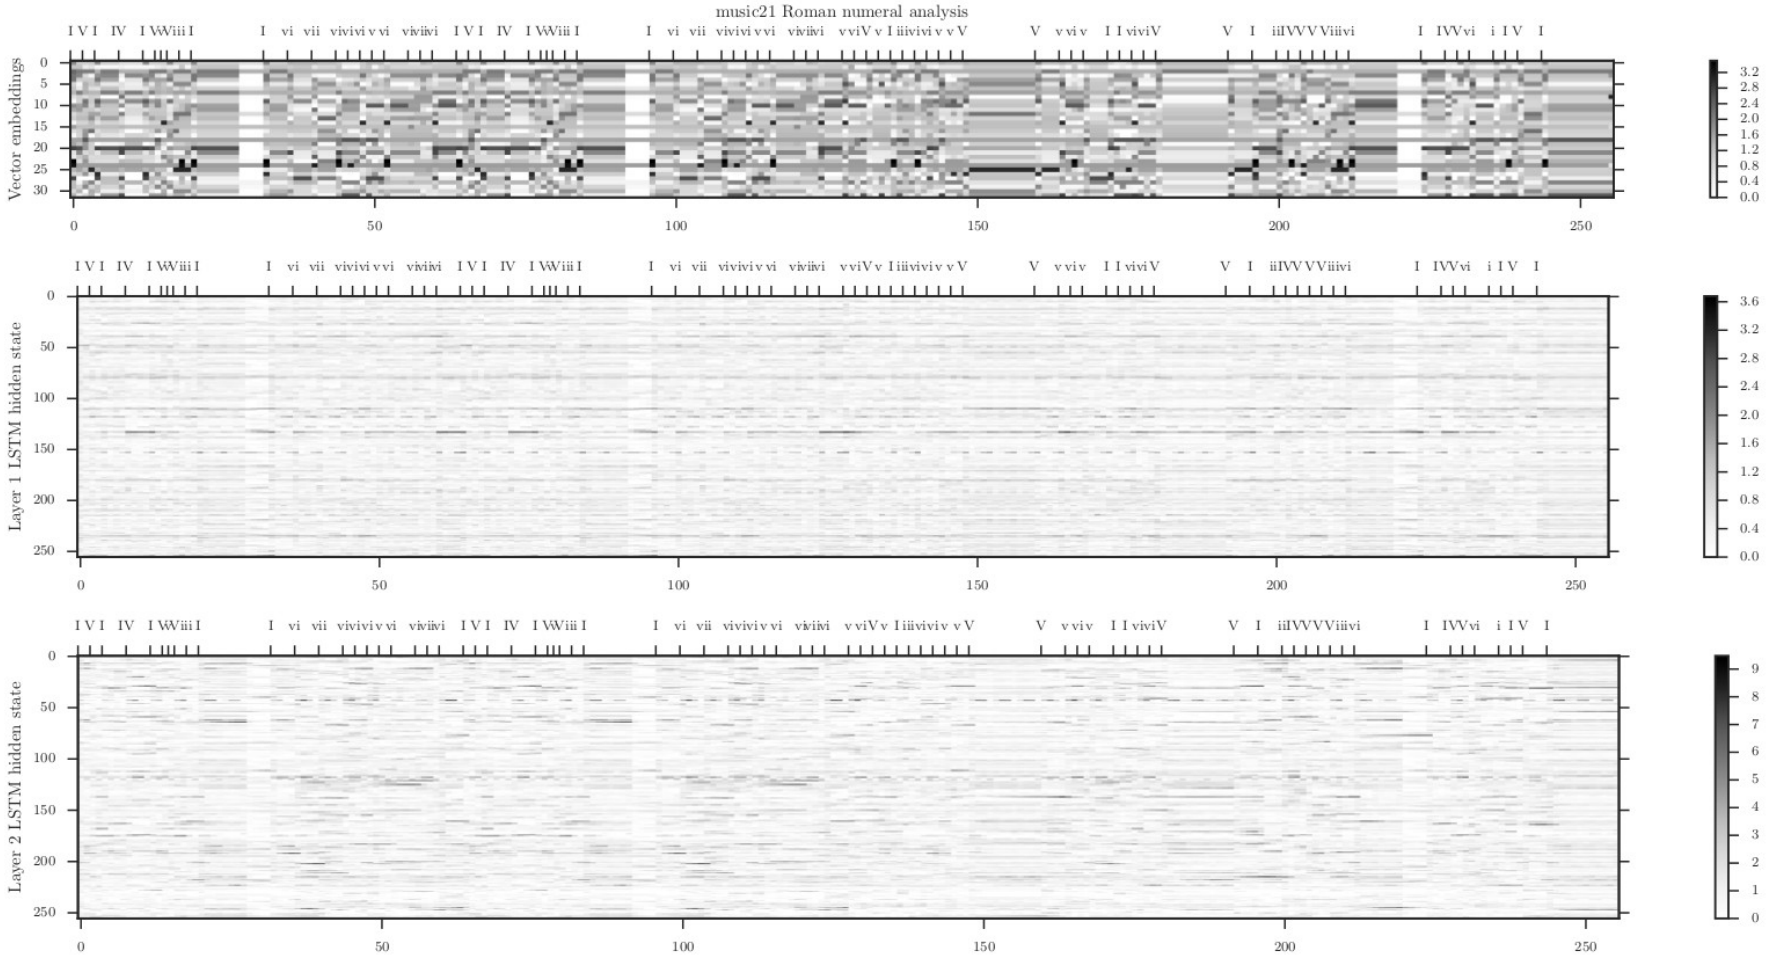

Hidden state is difficult to interpret!

Input and memory cell (layer 1 and 2)

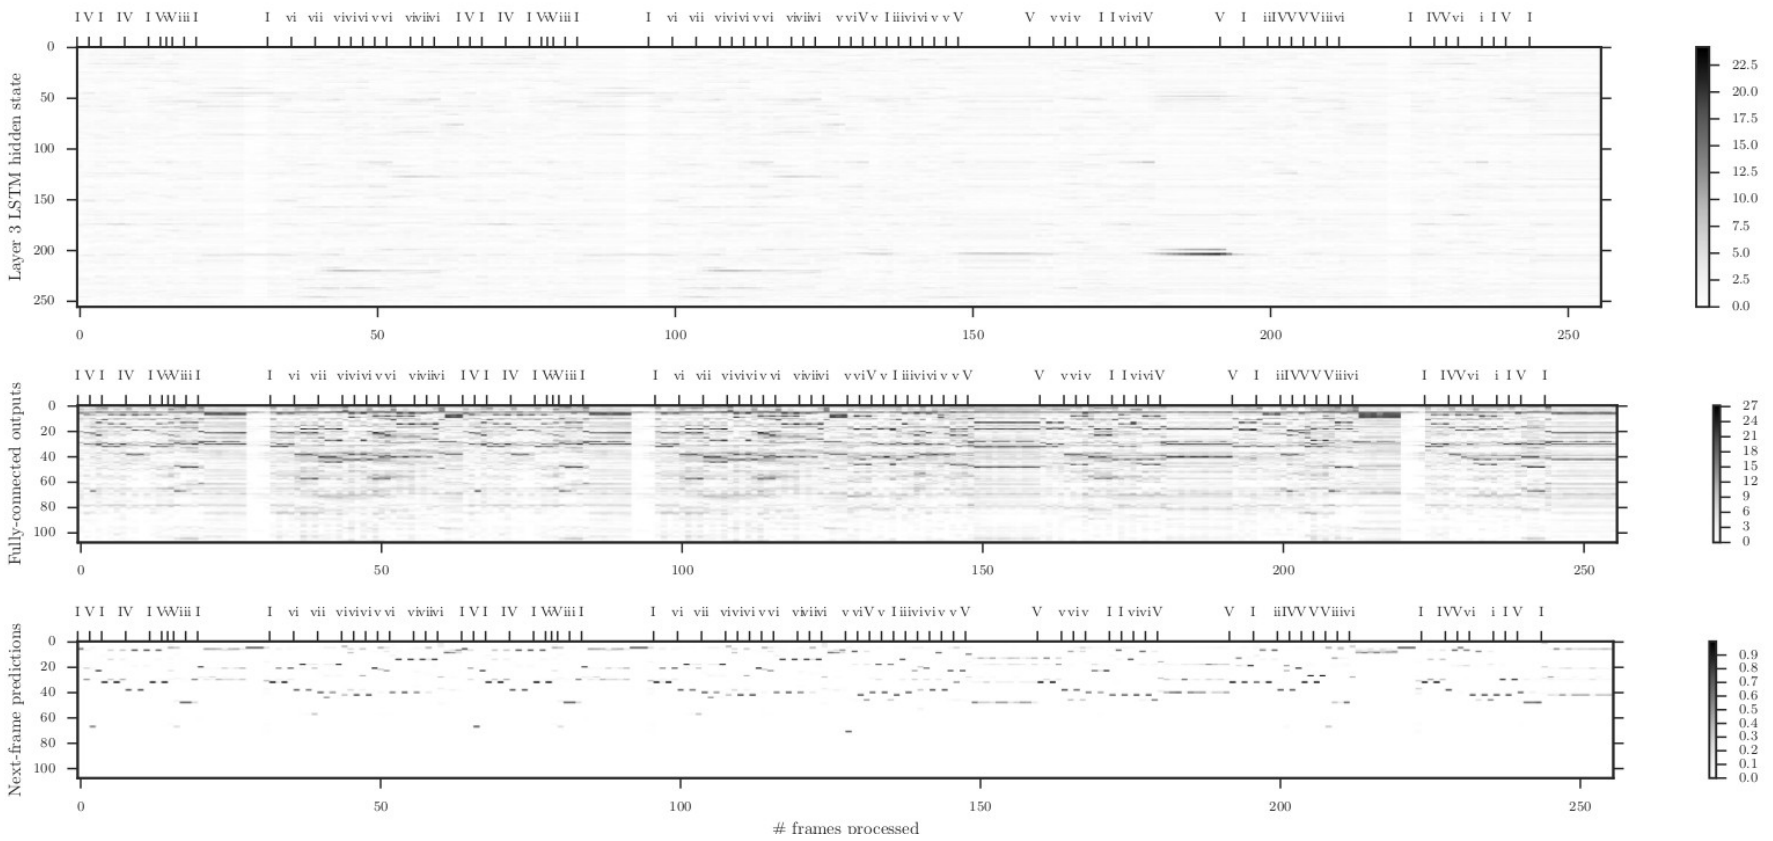

Layers closer to the output resemble piano roll (consequence of sequential training criteria)

Memory cell (layer 3), outputs, and predictions

Model learns music theory!

- L1N64 and L1N138: Perfect cadences with root position chords in tonic key
- L1N151: A minor cadences ending phrases 2 and 4
- L1N87 and L2N37: I^6 chords

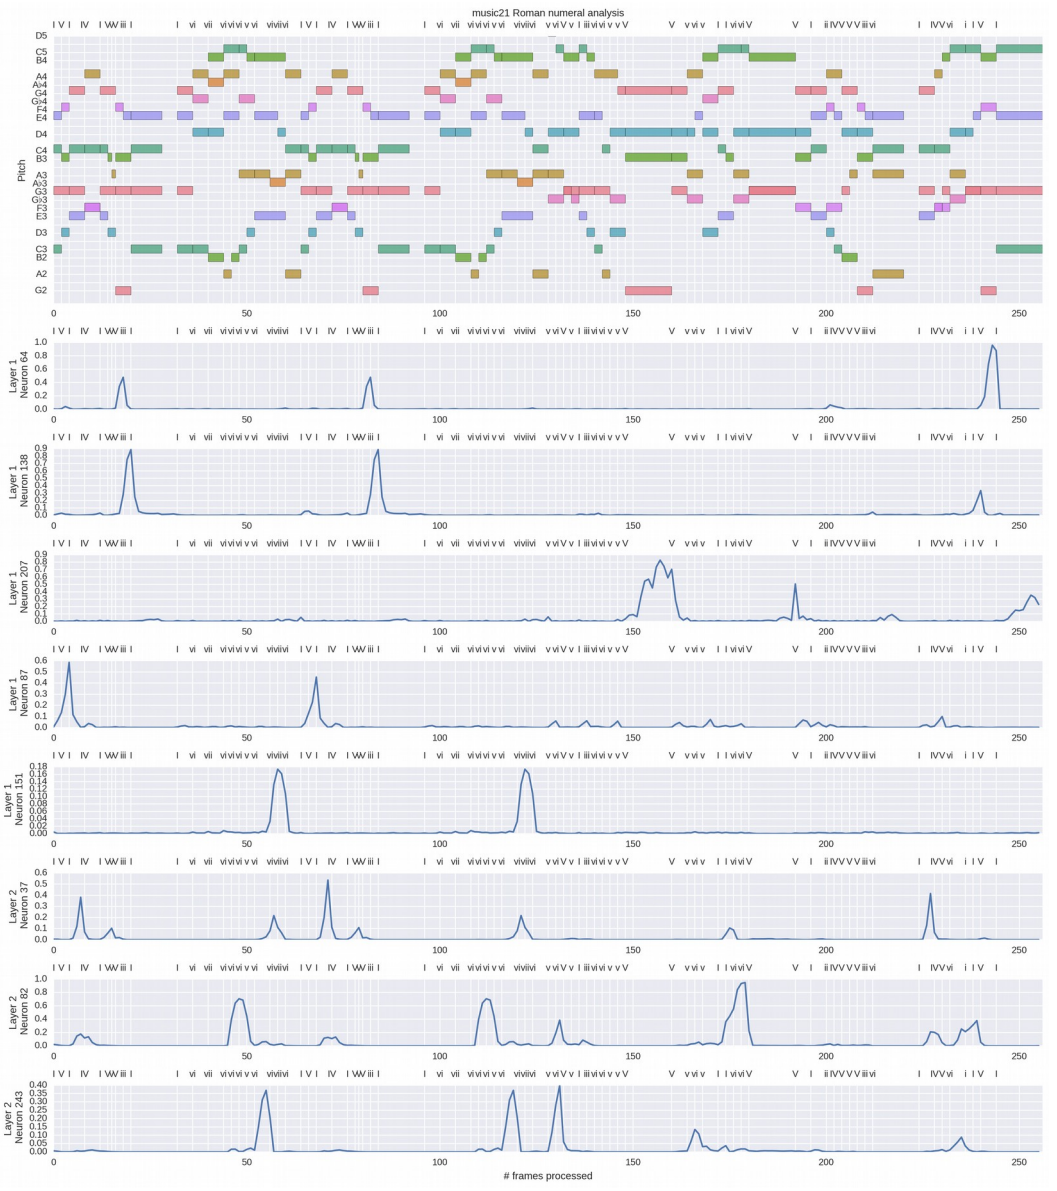

Harmonization: given a melody (here C major scale)

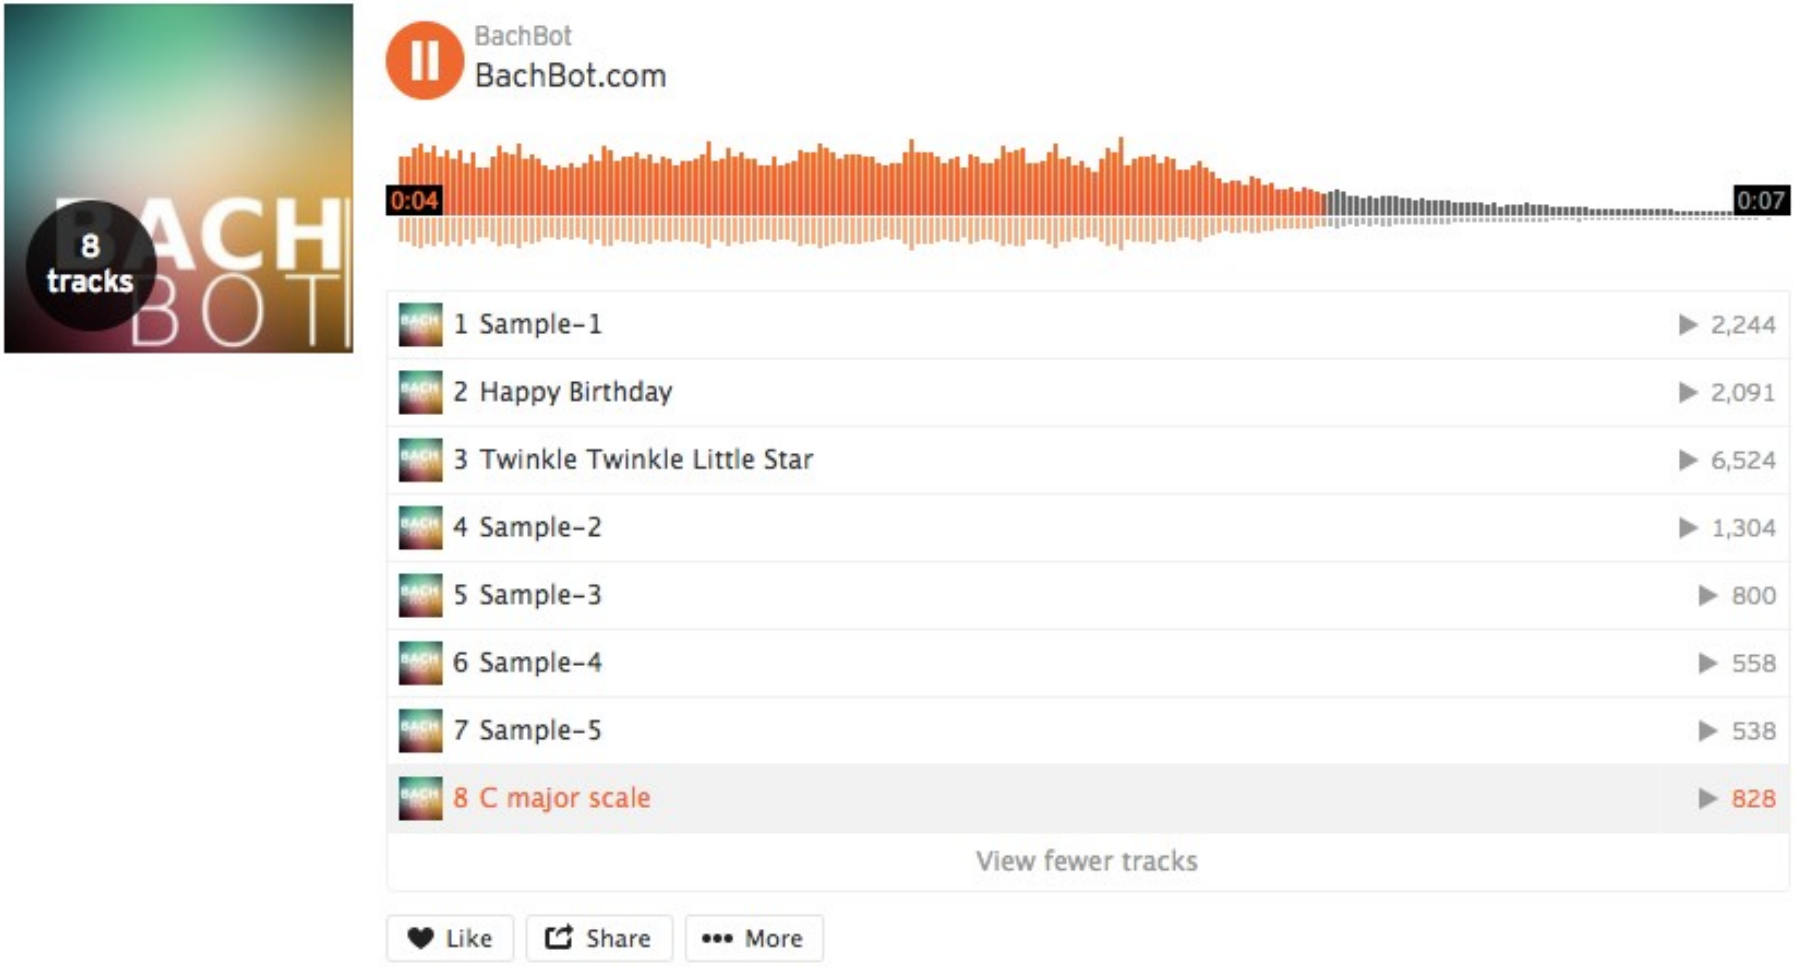

Harmonization: produce the accompanying parts

https://soundcloud.com/bachbot

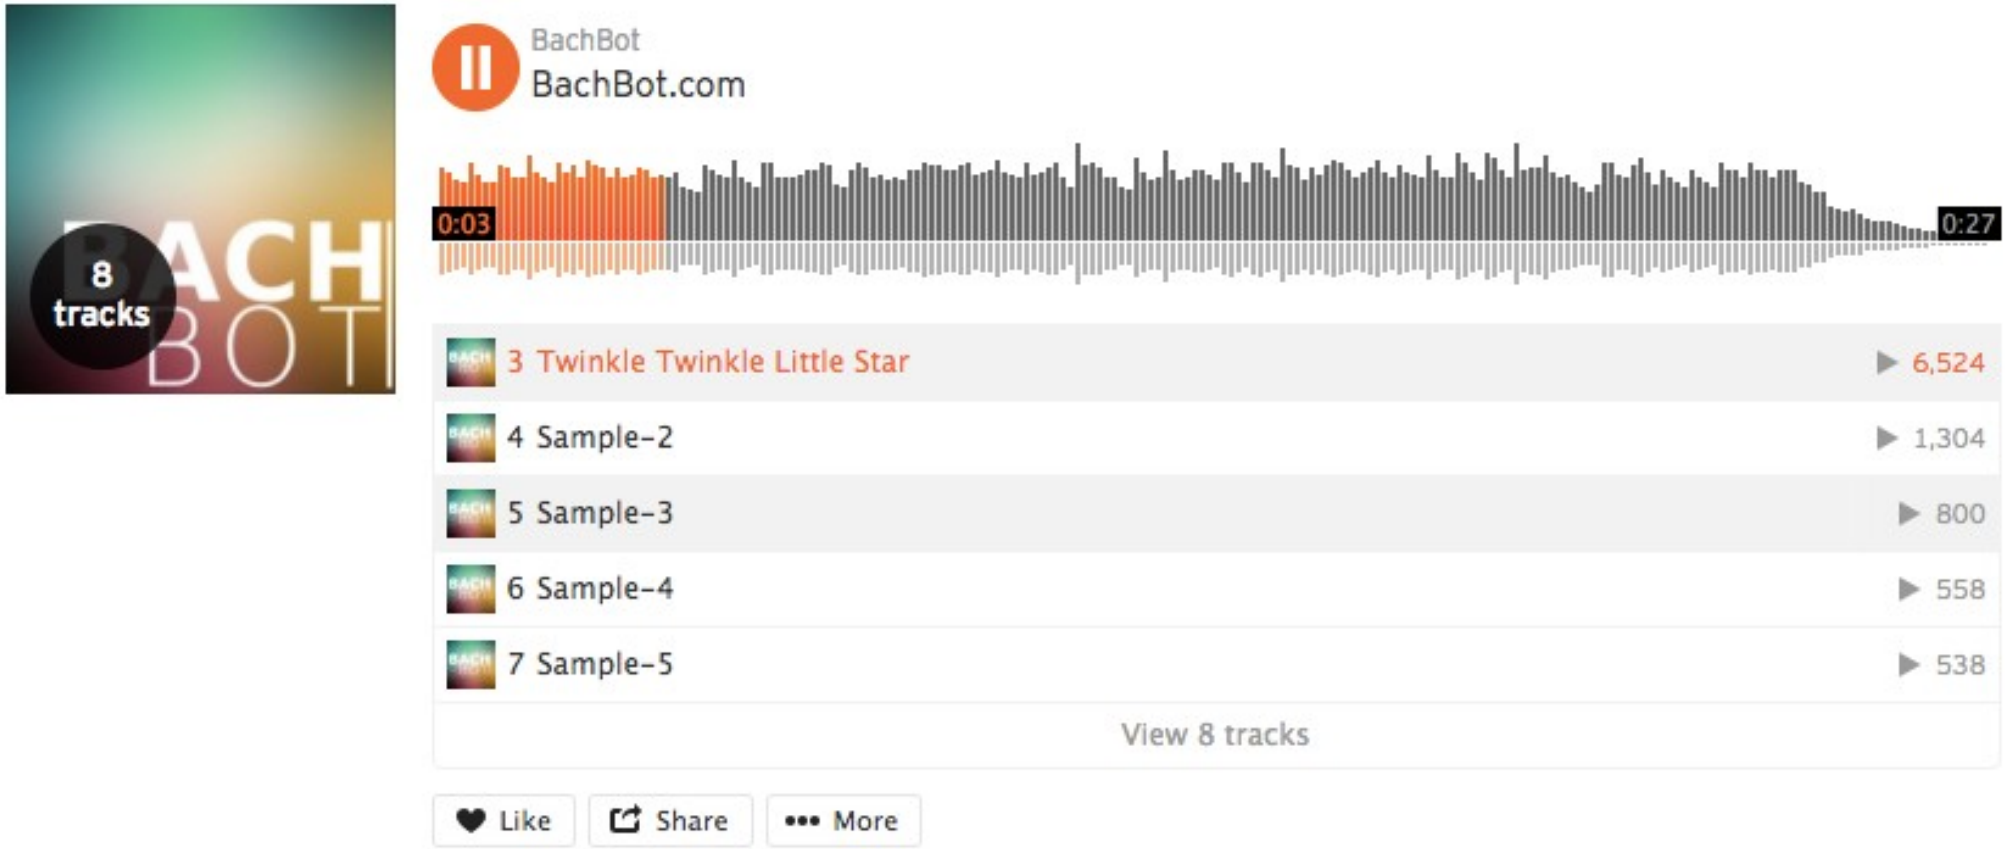

What if Bach remixed twinkle twinkle little star?

https://soundcloud.com/bachbot

Stand up if it’s your birthday

Thank You!

goto; copenhagen

GOTO Copenhagen 2019

Conference Nov. 18 - 20

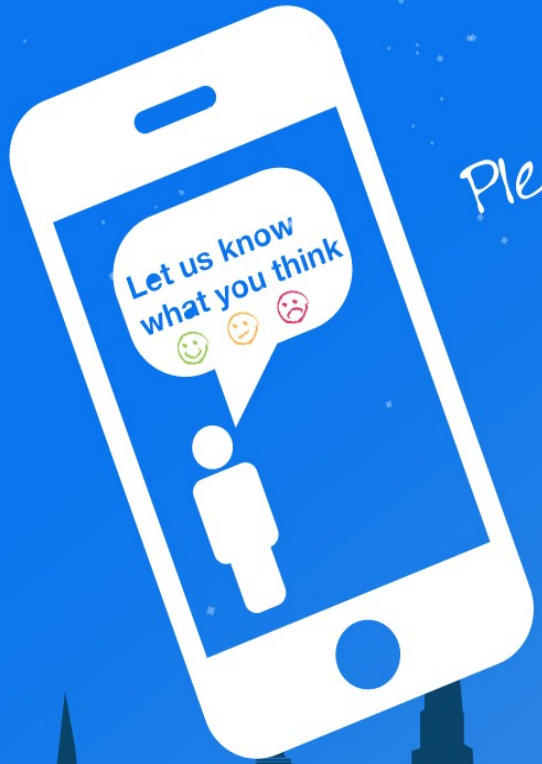

Please

Remember to rate this session

Thank you!