

Who Makes the Music?

Duquesne University | Mary Pappert School of Music | Room 322

Tuesday, 13 October 2026 | 7:30 p.m.



1. Welcome: Dr. Paul Miller (music)
2. Live performances of pieces: Dr. Kevin Sun (piano) and students + Dr. Miller (violin)
 - a. Jesu Leiden, Pein und Tod – J.S. Bach Harmonization – Riemenschneider 83
 - b. Jesu Leiden, Pein und Tod – J.S. Bach Harmonization – Riemenschneider 106
 - c. Jesu Leiden, Pein und Tod – Deepbach Harmonization
 - d. Jesu Leiden, Pein und Tod – ChoraleHarmonizer Harmonization
 - e. Jesu Leiden, Pein und Tod – Coconet Harmonization
 - f. Mozart: K. 331/i, first 8 bars – original and MIDI-LLM completion
 - g. Chopin: Mazurka op. 6/1, opening – original and AMT infill
 - h. Schumann: Waldszenen (“Eintritt”), opening – original and infill by MIDI-LLM
 - i. Berg: Sonata op. 1, opening – original and completion by MIDI-LLM
 - j. Rule-Guided Diffusion model – 8 bars of original music
3. Electronically rendered and AI-generated music
 - a. Text2Midi (MIDI in MuseScore)
 - b. Bark: precursor to Suno
 - c. Suno AI-generated song: “Weary Theory”
4. Short demonstration of Bach+, an LLM model based on Qwen, trained to answer basic questions about music theory concepts in the Bach chorale corpus: Dr. Miller
5. Statement by Dr. Raffaele Romagnoli (Computer Science)
6. Statement by Dr. Jeffrey McCurry (Philosophy)
7. Statement by Dr. Jacqueline Lipton (Law)
8. Open floor: questions from the audience and participants
9. Thank you and farewell

Nine Generative AI Models for Music

Model	Release	General generation strategy
DeepBach	2016–2017	Non-sequential probabilistic generation using pseudo-Gibbs sampling. A neural graphical model repeatedly resamples individual notes/voices conditioned on surrounding musical context. It is especially suited to Bach-style four-part chorales and allows some notes, rhythms, or cadences to be fixed as constraints. https://github.com/Ghadjeres/DeepBach
ChoraleHarmonizer	c. 2026	Hybrid Transformer + explicit musical constraints + search. A relatively small T5-style sequence-to-sequence transformer generates alto, tenor, and bass beneath a supplied soprano. Hard constraints, sampling penalties, phrase-wise best-of-N selection, a voice-leading rule engine, and a repair/resampling loop steer the result toward Bach-like voice leading. This model is <i>especially</i> good at avoiding common voice-leading errors, but it is not so great at writing convincing harmonizations as DeepBach. https://github.com/TimoKellerMath/ChoraleHarmonizer
Coconet	2017–2019	Convolutional masked reconstruction with blocked Gibbs sampling. Parts of a polyphonic score are masked, and the neural network repeatedly predicts/replaces missing musical material. Because generation is not strictly left-to-right, the model can revise notes anywhere in the musical texture. https://magenta.withgoogle.com/coconet
Anticipatory Music Transformer (AMT)	2023	Autoregressive Transformer with an “anticipation” representation. Future musical events or constraints are encoded so that a normally left-to-right transformer can generate material that must connect with music occurring later. This makes it especially useful for infill, bridge generation, and constrained continuation. This was especially good at generating bridges, sometimes. https://github.com/jthickstun/anticipation
MIDI-LLM	2025	General-purpose LLM adapted directly to MIDI. The vocabulary of a text LLM is expanded to include MIDI tokens, and the model is trained in stages so that the same LLM can interpret natural-language prompts and autoregressively generate multitrack MIDI. This model was also quite good, about 5% of the time. https://github.com/slSeanWU/MIDI-LLM
Rule-Guided Music Diffusion	2024	Latent diffusion plus rule-based guidance. A diffusion model gradually transforms noise into a symbolic piano-roll representation. Its stochastic control guidance mechanism can steer generation according to non-differentiable musical rules such as note density or chord progression without retraining the base model. This can only generate about 10.25 seconds of music at a time – roughly 8 bars – on the Apple MPS. With CUDA calls this could be potentially increased. https://github.com/yjhuangcd/rule-guided-music
Text2MIDI	2024–2025	Text-conditioned autoregressive transformer. A pretrained language-model encoder converts a prose caption into a representation that conditions a separate Transformer decoder, which then generates a sequence of MIDI tokens from left to right. This frequently proved to generate a lot of junk, but occasionally there are vague whiffs of decent music buried in the output. https://github.com/amaai-lab/text2midi
Bark	2023	Hierarchical text-to-audio generation using semantic and acoustic tokens. Bark uses a text-to-semantic transformer, an autoregressive semantic-to-coarse model, and a coarse-to-fine model to generate EnCodec audio-codebook tokens, which are decoded into a waveform sampled at 24kHz. It can generate speech, singing, nonverbal vocal sounds, background audio, and occasional accompaniment. In our tests it is highly stochastic: some outputs are strikingly convincing while others omit prompt elements, lose speaker identity, fall into short repetitive loops, or even collapse into noise. This was developed by Suno, and was an earlier open-source precursor. https://github.com/suno-ai/bark
Suno	2023-present	Text-conditioned generative audio model. A commercial, proprietary system that generates complete audio recordings—including instrumental accompaniment, vocals, and lyrics—from natural-language descriptions and other user-supplied musical material. Unlike MIDI- or symbolic-music models, Suno produces finished audio directly. Its detailed model architecture and training methodology have not been publicly disclosed. This is an astonishing model that produces pretty bland results at first, but can be steered towards much more sophisticated outputs. Probably trained on content illegally. https://suno.com/discover

J. S. Bach -- Chorale

Er nahm alles wohl in acht (#28 from John Passion)

also *Jesu Leiden, Pein und Tod* (#61, #83, #106 in Riemenschneider)

Jesu Leiden, Pein und Tod *

83.



Jesu Leiden, Pein und Tod

106.



J. S. Bach – Chorale

Jesu Leiden, Pein und Tod (#61, #83, #106 in Riemenschneider)

Deepbach harmonization (temperature = 0.8, slightly conservative)

Tune

Keyboard

The first system of the musical score. The 'Tune' part is on a single staff in G major (three sharps) and 4/4 time, featuring a melody with a half note, a quarter note, and a half note. The 'Keyboard' part consists of two staves (treble and bass clef) in the same key and time, providing a harmonic accompaniment with chords and moving lines.

5

The second system of the musical score, starting at measure 5. It continues the melody and keyboard accompaniment from the first system, with the keyboard part showing more complex chordal textures and moving lines.

10

The third system of the musical score, starting at measure 10. The melody continues with a half note and a quarter note, while the keyboard accompaniment provides a steady harmonic foundation.

14

The fourth system of the musical score, starting at measure 14. This system concludes the piece with a final cadence, featuring a half note and a quarter note in the melody and a final chord in the keyboard.

J. S. Bach – Chorale

Jesu Leiden, Pein und Tod (#61, #83, #106 in Riemenschneider)

Choraleharmonizer ("ausgewogen") with fermatas marked.

Tune

Keyboard

The first system of the musical score. The 'Tune' part is on a single staff in treble clef, key of D major (two sharps), and 4/4 time. It contains 8 measures, with fermatas marked over the 3rd and 7th measures. The 'Keyboard' part consists of two staves (treble and bass clef) in the same key and time signature, providing a harmonic accompaniment for the tune.

6

The second system of the musical score, starting at measure 6. The 'Tune' part continues on a single staff, with fermatas marked over measures 7, 8, and 9. The 'Keyboard' part continues with its harmonic accompaniment across two staves.

12

The third system of the musical score, starting at measure 12. The 'Tune' part continues on a single staff, with fermatas marked over measures 13, 14, and 15. The 'Keyboard' part continues with its harmonic accompaniment across two staves, concluding the piece with a double bar line.

J. S. Bach – Chorale

Jesu Leiden, Pein und Tod (#61, #83, #106 in Riemenschneider)

Coconet – convolutional neural network with masked completion ("inpainting")

Tune

Piano

5

10

14

Repaired an errant D3 here

This musical score is for a chorale by J. S. Bach, titled 'Jesu Leiden, Pein und Tod' (numbers 61, 83, and 106 in Riemenschneider's collection). The score is presented in a system with three staves: a single staff for the 'Tune' and a grand staff (treble and bass clef) for the 'Piano' accompaniment. The key signature is D major (two sharps) and the time signature is 4/4. The score is divided into four systems, with measure numbers 5, 10, and 14 indicated at the beginning of their respective systems. The notation includes various musical symbols such as notes, rests, accidentals, and phrasing slurs. A specific annotation, 'Repaired an errant D3 here', points to a note in the piano's right hand at measure 11, indicating a correction made by the Coconet model.

Mozart: K. 331/i

Original

Sonate in A

KV 331 (300ⁱ)

Wien (Salzburg), 1783

Andante grazioso

First system of the original score, measures 1-4.

5

Second system of the original score, measures 5-8.

Completed by MIDI-LLM (first four bars supplied)

First system of the MIDI-LLM completion, measures 1-4.

Second system of the MIDI-LLM completion, measures 5-8.

Chopin: Mazurka op. 6/1

Original

FR. CHOPIN
Op. 6 Nr 1

♩ = 132

1 *p*

cresc.

rubato

decresc.

legato

rit.

p

pp

15

The musical score for Chopin's Mazurka op. 6/1 is presented in four systems. The first system (measures 1-4) begins with a tempo marking of ♩ = 132 and a piano (p) dynamic. It features a triplet of eighth notes in the right hand and a bass line with triplets and slurs. The second system (measures 5-8) includes a decrescendo (decresc.) and a rubato section. The third system (measures 9-12) shows a crescendo (cresc.) followed by a ritardando (rit.) and a piano (p) dynamic. The fourth system (measures 13-15) concludes with a pianissimo (pp) dynamic and a final triplet. The score is marked with various fingerings and articulations throughout.

Chopin: Mazurka op. 6/1

Infill by Stanford's Anticipatory Music Transformer (AMT).

The first four bars and last three full bars were supplied.

The image displays a musical score for Chopin's Mazurka op. 6/1, specifically measures 1 through 15. The score is written for piano in 3/4 time, with a key signature of two sharps (F# and C#). The notation is presented in three systems, each with a grand staff (treble and bass clefs).
- The first system (measures 1-6) shows the initial four bars supplied and the first three bars of the AMT infill. It features a prominent triplet in the right hand of measure 2 and another in measure 4.
- The second system (measures 7-12) continues the infill. Measure 7 is marked with a '7' above the staff. The right hand contains a triplet in measure 9.
- The third system (measures 13-15) shows the final three full bars supplied. Measure 13 is marked with a '13' above the staff. The piece concludes with a double bar line in measure 15.
Throughout the score, various musical notations are used, including eighth notes, quarter notes, and chords, with dynamic markings like 'z' (zaccato) appearing in measures 2, 4, 8, and 10.

Schumann: Waldszenen, "Eintritt"

Original

Nicht zu schnell. M.M. ♩ = 132.

1.

Infill by MIDI-LLM

first two bars and last last two bars supplied, but the last two are bars 3-4 of original

Alban Berg
Piano Sonata op. 1

SONATE.

Alban Berg, Op. 1.

Mäßig bewegt.

p *accel.* *rit.* *a tempo*

accel. e cresc. *stringendo* *f* *molto rit.*




Rascher als Tempo I.

rit. e dim. *ff* *r. H.* *l. H.* *sfz* *espressivo* *r. H.* *pp* *p*

betato

Alban Berg

Piano Sonata op. 1
Completion by MIDI-LLM



Rule-Guided Music Diffusion

<https://github.com/yjhuangdc/rule-guided-music>

Unlike all other generated content, this particular music had to be interpreted and slightly modified from the raw output. I interpreted some very fast notes as grace notes, and slightly altered a few awkward rhythms. The `target_rules` dictionary in the `RGM-settings.yml` file specifies four things: pitch histogram, vertical & horizontal density, and a rudimentary RN harmonic goal for each temporal window (here, measure). The output correlates strongly to these desired outcomes.

Performance Notes

The rhythm of the inner voice can be somewhat free. The “suttering” effect of the grace note at the beginning is probably desirable. The long sustained notes don’t necessarily have to be explicitly held on the piano; they can be dropped out early if you wish. The last C chord could be forte. It seems that this generated fragment invites continuation...but we only get this much out of the model running on a Mac M5. The tempo ♩ = 126 is just a suggestion.



```
target_rules:
  pitch_hist: [0.28, 0., 0.12, 0., 0.18, 0., 0., 0.26, 0., 0., 0., 0.16]
  vertical_nd: [2., 2., 2., 3., 3., 2., 2., 2.]
  horizontal_nd: [5., 6., 7., 8., 8., 7., 6., 5.]
  chord_progression: [1, 1, 5, 5, 5, 5, 1, 1]
```

} This is the interesting part for us

```
guidance:
  vae: True # whether need to decode during sampling process
  nn: True # whether to use classifier
  scg: True
  method: classifier_guidance
  cond_fn:
    rule_names: ['pitch_hist', 'note_density', 'chord_progression']
    fns: ['grad_nn_zt_mse', 'grad_nn_zt_mse', 'grad_nn_zt_chord']
    classifier_scales: [400, 10., 20.]
    classifiers:
      num_classes: [12, 16, 8]
      names: ['DiTRotary-S/8-cls', 'DiTRotary-S/8-cls', 'DiTRotary-S/8-chord-cls']
      paths: ['loggings/classifier/pitch/model009999.pt', 'loggings/classifier/nd/model009999.pt', 'loggings/classifier/chord/model004999.pt']
  schedule: True
  t_start: 750 # when to start use guidance
  t_end: 0 # when to end use guidance
  interval: 1 # interval between successive guidance step
```

```
scg:
  num_samples: 16
  # weight for different losses
  pitch_hist: 40.
  note_density: 1.
  chord_progression: 2.
```

```
sampling:
  use_ddim: False
  diff_collage: False
  t_end: 0 # early stopping for ddpm
```

```
dc:
  type: circle
  overlap_size: 64 # how much overlap
  num_img: 1 # how many square images
```

Text2MIDI (2024)

Prompt:

“A piece in the baroque style in a minor key such as g minor.
This should be a minuet with an ABA form. It must be for keyboard only.”

$\text{♩} = 121$

Flute

Oboe

Clarinet

Bright Acoustic Piano

Electric Piano

Electric Bass

Slap Bass

Voice

Percussion

Piano

5

9

12

3

15

19

24

28

SUNO

Home

Explore

Create

Studio

Library

lauprellim

Earn Credits

Labs

Notifications

More

Upgrade to Pro

Songs

Speech

SimpleAdvancedv6-mini

+ Audio

+ Voice

+ Image

Epic dancehall song about strangers on the same train

+ Create

Workspaces > My Workspace

SearchFilters (3)NewList1

Thursday, Aug 13

3:25

Weary TheoryV4, 5-ALL

Pop-rock with driving mid-tempo four-on-the-floor, tight palm-muted guitar v...

...

3:33

Weary TheoryV4, 5-ALL

Pop-rock with driving mid-tempo four-on-the-floor, tight palm-muted guitar v...

...

1:00

Weary TheoryV5, 5 PREVIEW

Pop-punk with driving four-on-the-floor drums, palm-m...

Upgrade for full song

1:00

Weary TheoryV5, 5 PREVIEW

Pop-punk with driving four-on-the-floor drums, palm-m...

Upgrade for full song

4 songs

Weary Theory

lauprellim

0:073:24

🔊

🔍

👍

👎

💬

🔗

⋮

🔊

📱

1

0

0

Weary Theory

Add a Caption

Remix/Edit

Styles

Pop-rock with driving mid-tempo four-on-the-floor, tight palm-muted guitar verses, pre-chorus lifting on handclaps and stacked gang vocals, chorus opening wide with bright power chords and octave bass, bridge dropping to voice and toms before a final key change; lead vocal stays close-mic in verses with doubled chorus, short delay throws on the hook, crisp punchy mix with risers, snapped fills, and shimmering bell accents

[Verse 1]

Open the book, baby
Circle the notes in red
Half steps in my pocket
Whole steps in my head
Count it out, tap the desk
Lean in, don't fall asleep
Every scale's a staircase
We climb it in our sleep

[Pre-Chorus]

Build it from the ground up
Find the pattern, line by line
When the old rules click
You can hear them start to shine

[Chorus]

We're gonna learn music theory
Even if it makes us weary
We're gonna learn music theory
(keep it moving, keep it near me)
Major, minor, let it steer me
We're gonna learn music theory

[Verse 2]

G and C in the window
Follows where the song goes

This was my prompt: [