

Learning to Groove with Inverse Sequence Transformations

Jon Gillick, Adam Roberts, Jesse Engel,
Douglas Eck, David Bamman

contact: jongillick@berkeley.edu



Google AI / Magenta

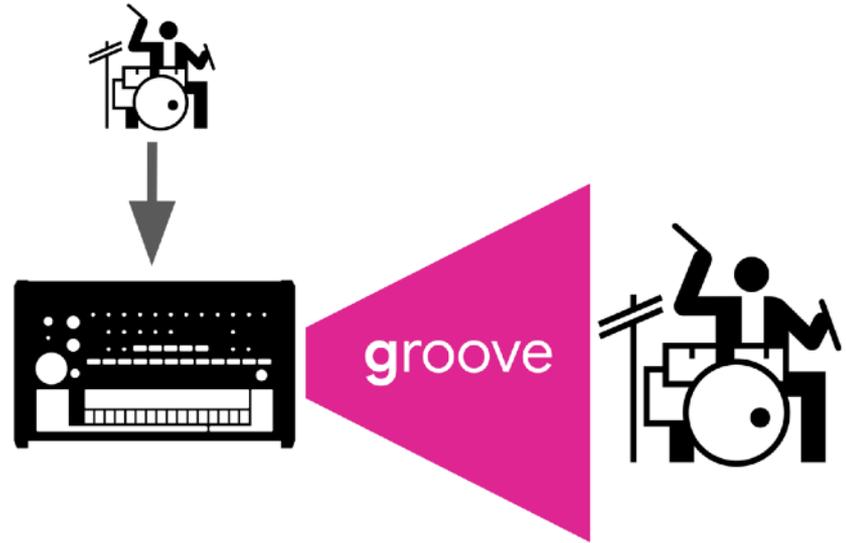


UC Berkeley
School of Information



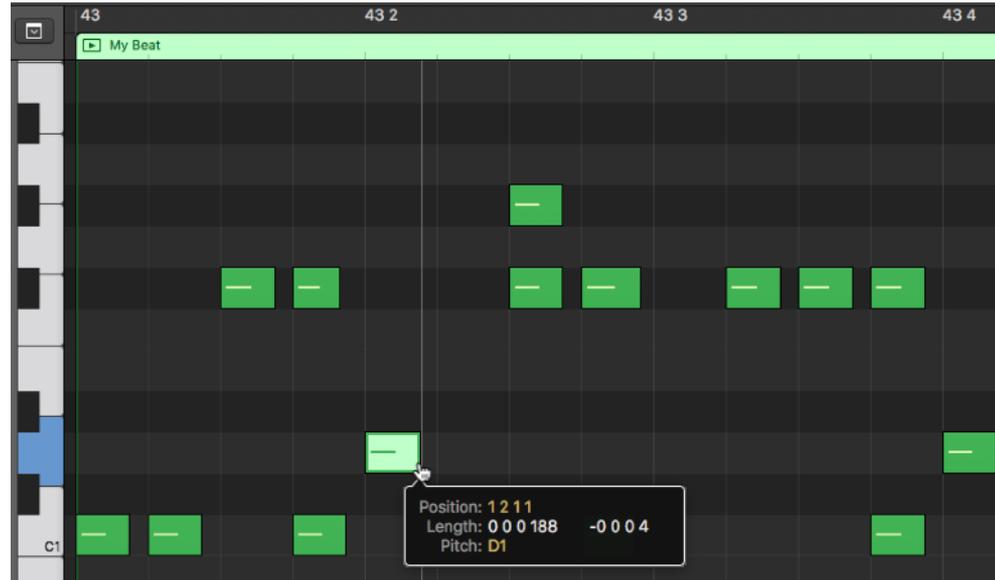
Questions

- How well can we model drum performances with machine learning?
- Can we use these models to make practical tools that give control to users?



Challenges in Editing Electronic Drums

- It is time consuming to edit the precise timing and volume of each note.
- Our ears connect with human performances.
- Not everyone can play drums, and recording drum kits is challenging and expensive.



Some Components of a Performance

Beat



≈

Score



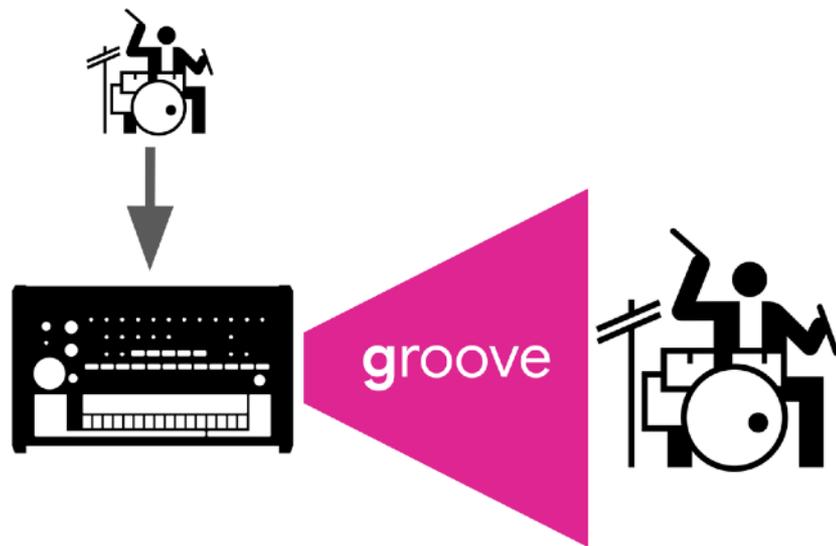
+

Groove



Contributions

- We build Machine Learning models that condition on **either** a score or a groove, **generating** the other.
- We collected and released the **Groove MIDI Dataset** of professional drum performances for modeling.



Models



Humanize



Model Architecture: Variational Autoencoder (**VAE**) or Variational Information Bottleneck (**VIB**) with recurrent encoders/decoders

Models

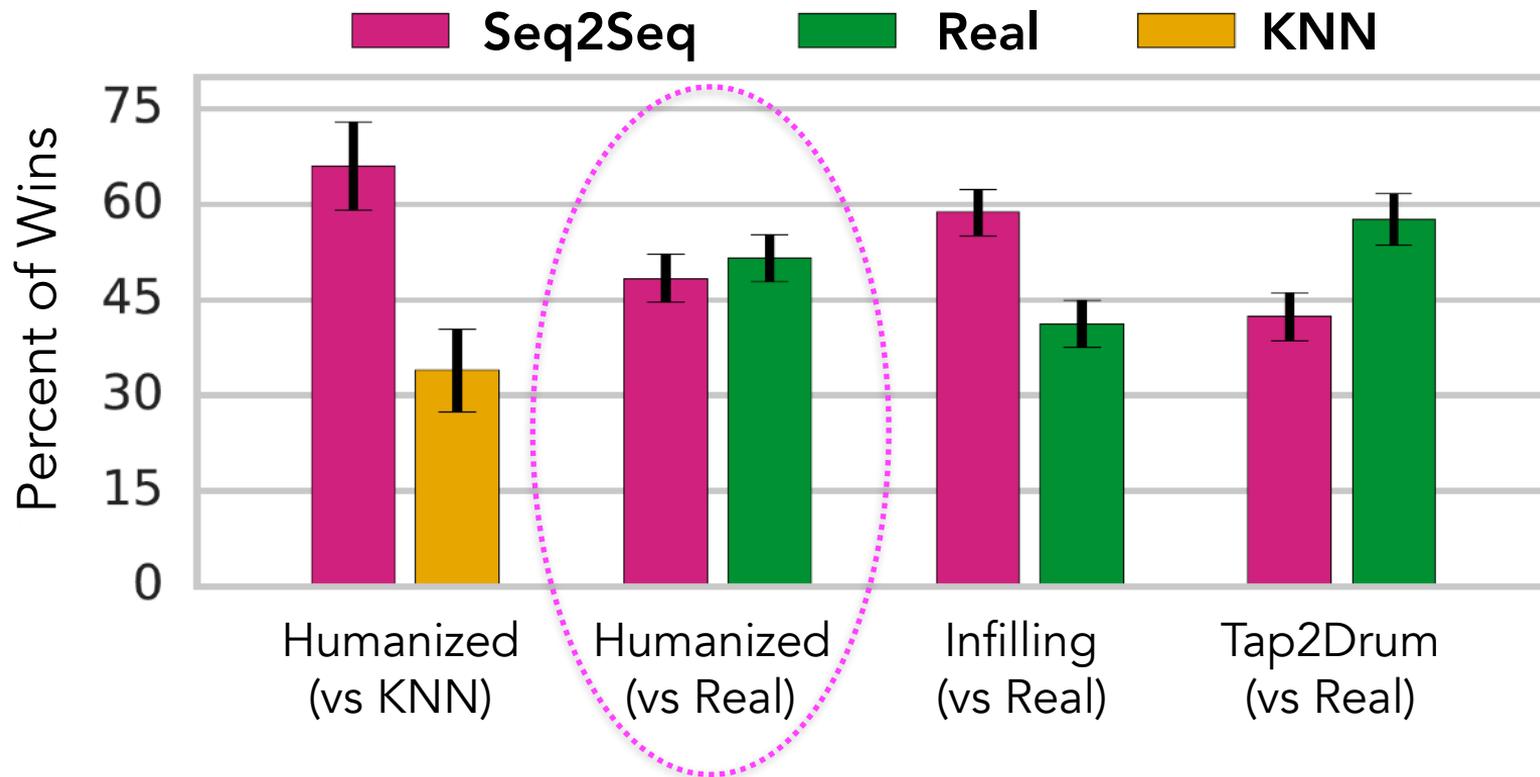


Tap2Drum



Model Architecture: Variational Autoencoder (**VAE**) or Variational Information Bottleneck (**VIB**) with recurrent encoders/decoders

Results: Listening Tests



Groove Model Demonstration

The image displays a software interface for a Groove Model demonstration, split into two main sections: a control panel on the left and a DAW interface on the right.

Control Panel (Left):

- GROOVE** (Title)
- Drums** (Red button)
- Input Clip** (Section header)
- MIDI Drums** (Red button with dropdown arrow)
- Beat** (Red button with dropdown arrow)
- Temperature** (Slider set to 1.0)
- Output clip to Clip Slot 2** (Text)
- Generate** (Large red button)
- Background image of a mountain landscape.
- Play button icon at the bottom left.

DAW Interface (Right):

- Window title: **Groove [GroVAE]**
- Transport: **Link Tap 120.00**, **4 / 4**, **1 Bar**, **13 . 4 . 1**
- Track list:
 - Magenta Studio** (Folder)
 - MIDI Drums** (Track)
 - Beat** (Clip)
 - Master** (Track)
 - 1
 - 2
 - 3
 - 4
 - 5
- Area: **Drop Files and Devices Here**
- Tempo: **120**
- Clip Properties (Bottom Left):
 - Clip**: **Beat**
 - Notes**: **Start** (1), **End** (1)
 - Signature**: **4 / 4**
 - Groove**: **Loop**
 - Name**: **None**
 - Comments**: **None**
 - Start**: **1**
 - End**: **1**
 - Length**: **0**
 - Position**: **1**
 - Bank**: **1**
 - Sub**: **0**
 - Pgm**: **2**
- Timeline: **1**, **1.2**, **1.3**, **1.4**, **2**, **2.2**, **2.3**, **2.4**
- Bottom status: **Insert Mark 1.1.1**, **MIDI Drums**, **1/16**

Drumify Model Demonstration

The image displays a digital audio workstation (DAW) interface, likely Ableton Live, showing a drum track and a sample editor. The interface is divided into several sections:

- Left Panel (Drumify Model):** Contains a "Drums" button, an "Input Clip" section with "2 Electric Piano" and "Electric Piano" buttons, a "Temperature" slider set to 1.0, an "Output clip to Clip Slot 2" label, and a "Generate" button. A background image of a mountain landscape is visible at the bottom.
- DAW Main Area:** Shows a multi-track arrangement with tracks for "MAGENTA", "Electric Piano", "Bass", "Drums", and "Bass 2". The "Drums" track is selected and contains a clip named "1/1 [Electric Pi]". The interface includes MIDI and audio routing options, volume faders, and solo/mute buttons for each track.
- Bottom Panel (Sample Editor):** Displays a clip named "Clip" with a sample of "4 Audio 0002 [2019]". The sample is a 44.1 kHz, 24-bit, 1-channel audio file. The editor shows the sample's waveform and various parameters such as "Signature" (4/4), "Groove" (None), "Transpose" (0 st), "Detune" (0.00 dB), "Warp" (Set), "Start" (1), "End" (3), "Beats" (1), "Position" (1), "Length" (2), and "Loop" (Off).

The DAW interface also shows a "Master" track with a tempo of 126 bpm and a volume fader. The bottom right corner of the DAW interface displays the page number "1/16".

Questions for the Future

- How do professionals and/or amateur musicians experience working with these tools?
- How can/should we facilitate collaborations with the expert creators (such as drummers) that enable this kind of research?
- What specifically do these models learn? What biases do they capture, and how does this inform future data collection?

Thank you!

g.co/magenta/groovae

Stop by our poster: 6:30pm,
Pacific Ballroom #242 for audio examples,
interactive demos, and more!



Google AI / Magenta



UC Berkeley
School of Information

Images

Drummer by Luis Prado from the Noun Project

Drum Machine by Clayton Meador from the Noun Project

