

Deep Learning Techniques for Music Generation

1. Introduction to Computer Music

Jean-Pierre Briot

`Jean-Pierre.Briot@lip6.fr`

Laboratoire d'Informatique de Paris 6 (LIP6)
Sorbonne Université – CNRS



Programa de Pós-Graduação em Informática (PPGI)
UNIRIO



Computer Music

Why Using Computer for Music

- **Bad Reasons (Fears)**

- Lead human musicians to unemployment
- Lower the quality of music



- **Good reasons**

- Facilitate storing, indexing, delivering and sharing of music (MIDI, MP3, Spotify...)
- New instruments and interaction (Synthesizers, Interactive music performances...)
- New sounds (Synthesizers and Signal processing)
- Analysis tools and algorithms (Spectrum, Patterns...)
- Initiation and Education (Band in the box, Garage Band...)



- **Production**

- Partially automate tasks (mixing, etc.)

- **Composition, Analysis and Arrangement**

- Algorithmic composition
- Harmonization
- Analysis

Why Using Computer for Music


- Vast Associative Memory
 - More systematic than Human memory
- Representation of Musical pieces, Style, Patterns...
- Associations and Correlations
- Knowledge (Theory, Rules, Heuristics...)
- Can Help Human musicians
- Human musicians rarely compose from scratch – They steal from others
 - Consciously
 - » Plagiat, Citation...
 - Unconsciously
 - » Influence
 - Recombinations
 - Historical Evolution
 - » Modal monophonic -> Polyphonic (Counterpoint) -> Tonal Music (Harmony) -> Extended Harmony (Debussy, Jazz...)
 - Ruptures (Dodecaphonism, Free Jazz)
 - » Rare and often transient

Autonomous vs Assistance/Interactive Music Making

Autonomous vs Assistance Music Making

- Autonomous Generation/Interpretation


- Turing Test
- Ambient Music



CompressorHeads (no AI Inside 😊)

- Assistance to Human Composers and Musicians

- Propose
- Refine
- Analyze
- Harmonize
- Produce
- ...



Bach Chorales Turing Test


- Autonomous Artificial Musicians

- Music Composition Turing test

- Imitation Game Scenario [Turing, 1950]
- Designed by A. Turing to explore the question "Can Machines think?"



(A) J. S. Bach



(B) DeepBach [Hadjeres et al., 2017]



?



(C) Listener

- To evaluate artificial composers techniques
- To explore music cognition

A. M. Turing (1950) Computing Machinery and Intelligence. *Mind* 49: 433-460.

COMPUTING MACHINERY AND INTELLIGENCE

By A. M. Turing

1. The Imitation Game

I propose to consider the question, "Can machines think?" This should begin with definitions of the meaning of the terms "machine" and "think." The definitions might be framed so as to reflect so far as possible the normal use of the words, but this attitude is dangerous. If the meaning of the words "machine" and "think" are to be found by examining how they are commonly used it is difficult to escape the conclusion that the meaning and the answer to the question, "Can machines think?" is to be sought in a statistical survey such as a Gallup poll. But this is absurd. Instead of attempting such a definition I shall replace the question by another, which is closely related to it and is expressed in relatively unambiguous words.

The new form of the problem can be described in terms of a game which we call the 'imitation game.' It is played with three people, a man (A), a woman (B), and an interrogator (C) who may be of either sex. The interrogator stays in a room apart from the other two. The object of the game for the interrogator is to determine which of the other two is the man and which is the woman. He knows them by labels X and Y, and at the end of the game he says either "X is A and Y is B" or "X is B and Y is A." The interrogator is allowed to put questions to A and B thus:

C: Will X please tell me the length of his or her hair?

Now suppose X is actually A, then A must answer. It is A's object in the game to try and cause C to make the wrong identification. His answer might therefore be:


"My hair is shingled, and the longest strands are about nine inches long."

In order that tones of voice may not help the interrogator the answers should be written, or better still, typewritten. The ideal arrangement is to have a teleprinter communicating between the two rooms. Alternatively the question and answers can be repeated by an intermediary. The object of the game for the third player (B) is to help the interrogator. The best strategy for her is probably to give truthful answers. She can add such things as "I am the woman, don't listen to him!" to her answers, but it will avail nothing as the man can make similar remarks.

We now ask the question, "What will happen when a machine takes the part of A in this game?" Will the interrogator decide wrongly as often when the game is played like this as he does when the game is played between a man and a woman? These questions replace our original, "Can machines think?"


Bach Chorales Turing Test

- February 2017, Dutch TV Channel
- Bach vs DeepBach Turing Test



Autonomous Music Making

- Symbolic or/and Audio Music Generation



— https://www.youtube.com/watch?time_continue=11&v=aUFMFSI5qM8

Autonomous Music Making

- Symbolic or/and Audio Music Generation
 - For Commercials and Documentaries
 - Create Royalty-free or Copyright-buyable Music
 - Based on Deep learning + Samples + Sound processing techniques
- + Business model
- Musical model

Interactive/Assistance

- Assistance to Human Composers and Musicians
 - Propose
 - Complete
 - Refine
 - Analyze
 - Harmonize
 - Adapt
 - Produce
 - Accompany
 - ...



- Ex: FlowComposer...

I try to find a logic in the melody, so I make some edits

 - Wide Memory (FlowComposer: more than 12.000 songs/lead sheets)
 - Harmony knowledge (ex: Harmonic analysis, Harmonization, Chord substitution...)

FlowComposer – Demo

The screenshot shows the FlowComposer web application in a browser. The address bar displays the URL: `lsdb.flow-machines.com/leadsheet.php?id=570e6d8658e33885067b23c6&edition=1&player=1`. The browser's tab is labeled "LeadsheetJS - Flow Machin". The page features a navigation menu with options like "Database", "Songs", "New song", "Copy song", "Chords", "Styles", "Sources", "Songsets", "Music Files", "Rechord", "Radio", and "About". A secondary menu includes "File", "Notes", "Suggest Chords", "Structure", "Harmonic Analysis", "Harmonizer", "Similarity Analysis", and "Composer". The "Composer" tab is active, showing a control panel with fields for "No songset", "No composer", and "Song", along with sliders for "Harmony", "Note Duration", "Chord Changes", and "Inspiration from". A "Compose" button is present. The main workspace is titled "Untitled" and "Unknown", displaying a musical staff with a treble clef and a 4/4 time signature. The staff contains a single note on the first line, marked with a blue box and the letter 'A'. A "History" panel on the right shows "Open song - Untitled". The bottom of the interface includes a control bar with buttons for "Render", "Download Midi", "Stop", "Play", "Volume", "Tempo" (set to 120), "Loop", and "Metronome".



https://www.youtube.com/watch?time_continue=5&v=SDnkX8v8caY

Objective and Evaluation [Pachet, 2019]



	Current Systems	Future Systems
	Autonomous Generalization-based	Augmentation/Assistance Creative-incentived
Objective	Create music	Create music not possible otherwise
Evaluation	Please the listener	Please the composer
Risk	Conventional	Surprising But meaningful

Autonomous Artificial Intelligence (AI) vs Intelligence Augmentation (IA)






Symbolic vs Connexionist AI – History



Symbolic vs Cybernetics – History



AI vs IA – History




Symbolic vs Audio

- Symbolic
 - Conceptual Level
 - Composition Level
 - Symbolic Representations
 - Knowledge Representation (Symbolic AI)
 - Note, Chord, Rest...
 - Harmony, Modes...
 - Digital
- Audio
 - Signal Level
 - Signal Processing
 - Spectrum, Fourier Transform, MFCC...
 - Analogic
- Complementarity
- Integration challenge

Computer Music

- Started in the 50's
- ILLIAC Suite [Hillier & Isaacson, 1957]



<https://www.youtube.com/watch?v=v5z3t2qz4qo>

An example of Stochastic Algorithmic Music

- Generate and Test
- Random Generation
- Handcrafted Filtering Rules

- Generation of Pieces of Music Matching the Rules (Constraints)
- (Inefficient Constraint Solving Strategy)

Also (Ex):

- Zenakis Stochastic Compositions

- Musikalisches Würfelspiel (Musical Dice Game) by Mozart
 - Stochastic Combination of pre-Existing 2 measure-long Musical Segments
 - Fixed Style (Austrian Waltz) and Tonality/Key
 - Expertise of the Composer

<https://mozart.qvwx.de/>


Mozart Dice Music

- Actually one of the First Documented Stochastic Music (1792)
- By Mozart (?) – Muzikalisches Würfelspiel (Dice Music)
- Fixed Style (Vienna Waltz) and Tonality
- 11 Pre-defined 1-Measure Segments for each (16) Measure
- Stochastic Combination/Concatenation
- $11^{16} = 45,949,729,863,572,161$ Possible Pieces



Musikalisches Würfelspiel
K.516f: 6:8:6:4:5:6:8:2:8:5:6:7:3:4:9:8

W. A. Mozart



Models

- **Pre-Defined Elements**
 - Combination (ex: Mozart Dice Music)
- **Rules**
 - Application
 - Filtering (ILIAC Suite), Generation, Harmonic Analysis...
- **Generative Grammars**
 - Valid Sentences generated by the Grammar
 - Harmonic Cadences Construction, Substitutions...
- **Constraints**
 - Constraint Solving Problem
 - Generation (ILIAC Suite)
 - Accompaniment: Harmonization, Counterpoint...
- **Markov Chains**
 - Generation (Random Walk, Constrained)
 - Style Imitation
- **Neural Models (Deep Learning)**
 - Prediction/Classification
 - Style Imitation

(Simplistic) Example of Generative Grammar for Blues

8-Measures	->	First-4-Measures	Second-4-Measures
First-4-Measures	->	Opening-Cadence	Opening-Cadence
	->	Opening-Cadence'	Opening-Cadence
Second-4-Measures	->	Middle-Cadence	Opening-Cadence
Opening-Cadence	->	I I	
	->	I V	
Opening-Cadence'	->	I III	
	->	I IV	
Middle-Cadence	->	I IV	
	->	I V	
	->	IV I	

[John-Laird, 1991]

Ex:

| I | I | I | I | IV | I | I | V |





| I | IV | I | V | IV | I | I | I |




Markov Chains

- Hypothesis (Markov)
- Past is unnecessary to model the future
- $P(s_{t+1} \mid s_1, s_2 \dots, s_t) = P(s_{t+1} \mid s_t)$
- States
- Probabilities for transition to other (or same) State
- Ex: Weather Modeling





Markov Chains

- Other Example:
- Modeling of Market Evolution ↑ ↓



(Simplistic) Example of Markov Model



Models



- Cellular Automata
 - Evolution
- Genetic Algorithms
 - Selection
- Case-based Reasoning
 - Similarity and Adaptation
- Planning
 - Path (Melody, Chord Sequence...) Construction
- ...

Importance of Randomness (Stochasticity)

- Randomness/Stochasticity
 - Governed by Probabilities
- To avoid Determinism
 - Otherwise Infinitely Repeats the Same Generation
- To be able to Generate Various Musical Pieces from a single Generative Model
- Ex: Mozart Dice-Music, Xenakis...
- Incorporated in most of Algorithmic Compositions

Also Models of Synthesis/Sound


- Additive
 - Sinusoids/Fourier¹
- Subtractive
 - Filtering
- Additive+Operators
 - Analog Synthesizers
- Physical
 - ex: Singing Voice, String...
- Frequency Modulation
 - Ex: Yamaha DX7
- Sampling
 - Initial Analog Version : Musique Concrète
- Granular
 - Arbitrary Objective
- Concatenative (Musaicing)
 - Source Objective
- ...



Handcrafted vs Learnt Models

- **Handcrafted**
 - Tedious
 - Error-Prone
- **Automatically Learnt (Induction)**
 - Markov models
 - Neural models
- **Style Automatic Learned from a Corpus** (Composer, Form, Genre...)

- **Machine Learning Techniques**
 - Neural Networks, **Deep Learning**, **Reinforcement Learning**
 - (and other models/techniques)



Flow Machines [Pachet et al. 2012]

Reorchestration of Ode of Joy Using Various Techniques (Flow Machines)

Ode to Joy in several styles


<https://www.youtube.com/watch?v=buXqNqBFd6E>

Acquiring Models

- Handcrafted
 - Tedious
 - Error-Prone
- Automatically Learnt (Induction)
 - Ex:
 - Markov models
 - Neural models
- Style Automatic Learned from a Corpus (Composer, Form, Genre...)
- Machine Learning Techniques
 - Neural Networks, Deep Learning, Reinforcement Learning (and other models/techniques)

Modeling of Beethoven's Ode of Joy

- Beethoven "Ode to Joy" (9th Symphony)



What note follows a G ?

- G -> E 1 times : 0.14
- G -> F 3 times : 0.43
- G -> G 3 times : 0.43

$$P(n_{t+1}=X \mid n_t=G)$$

0.14

0.43

0.43

Higher-Order Markov

- Order 1: $P(s_{t+1} \mid s_1, s_2 \dots, s_t) = P(s_{t+1} \mid s_t)$
- Order 2: $P(s_{t+1} \mid s_1, s_2 \dots, s_t) = P(s_{t+1} \mid s_{t-1}, s_t)$
- Order 3: $P(s_{t+1} \mid s_1, s_2 \dots, s_t) = P(s_{t+1} \mid s_{t-2}, s_{t-1}, s_t)$
- ...
- Variable Order
- Limit: Plagiat

Higher-Order Markov



[Roy and Pachet, 2017]

References

- Gerhard Nierhaus, Algorithmic Composition: Paradigms of Automated Music Generation, Springer, 2009.
- David Cope, The Algorithmic Composer, A-R Editions, 2000.
- Roger T. Dean and Alex McLean, The Oxford Handbook of Algorithmic Music, Oxford Handbooks, Oxford University Press, 2018.
- Curtis Roads, The Computer Music Tutorial, MIT Press, 1996.