Deep Learning Techniques for Music Generation
Additional Material (9)

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

Recent Creations

Electro Dance-Pop Music

- YΔCHT (Young Americans Challenging High Technology)
- Chain Tripping Album, 30 August 2019
- Composed with Magenta MusicVAE

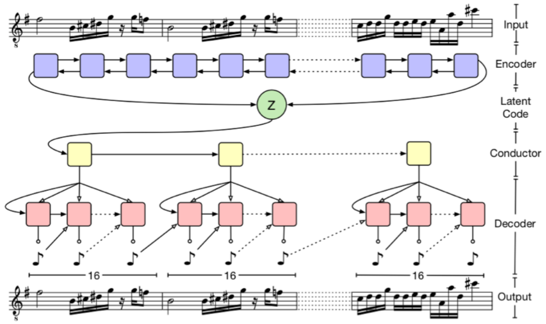

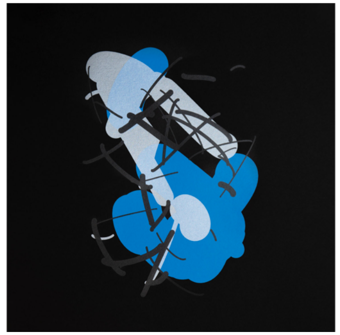

I'm so in love
I can feel it in my car
I can feel it in my heart,
I can feel it so hard
I want your phone to my brain
I want you to call my name
I want you to do it too
Oh, won't you come, won't you come
Won't you work on my head
Be my number nine

(Downtown) Dancing

Loud Light

YΔCHT + Magenta – Chain Tripping Album

- Melody/Chords/Rhythm Loops
  - MusicVAE (VRAE)
  - Training Corpus: Previous music by YΔCHT
- Lyrics
  - LSTM
  - Training Corpus: YΔCHT + Liked Lyrics
- Sounds
  - Nsynth (Signal VAE)
- Images and Videos
  - GAN

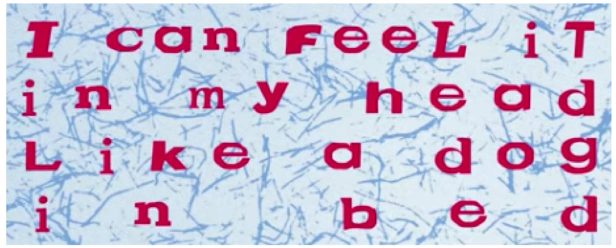

https://arstechnica.com/gaming/2019/08/yachts-chain-tripping-is-a-new-landmark-for-ai-music-an-album-that-doesnt-suck/

YΔCHT + Magenta – Chain Tripping Album

- Rules:
  - Every new song interpolated from existing YΔCHT melodies
  - 4 measures-long loops
  - Cannot add any note, harmony
  - Only substractive or transpositional changes
  - Structure and collage allowed
  - Assignment (to vocal, bass line…)
- Human Production and Arrangements

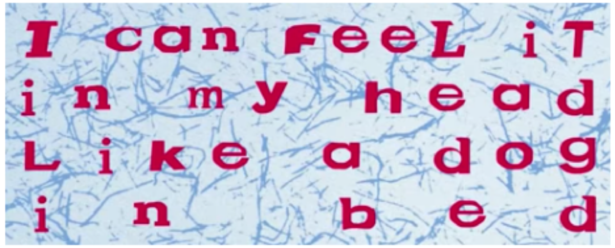

https://www.youtube.com/watch?time_continue=1378&v=pM9u9xcM_cs&feature=emb_logo

History Revisited

Deep Learning

- Boom Since 2012 (Imagenet Breakthrough)


- Image Recognition

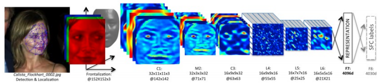

- Weather Prediction

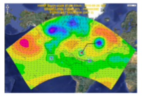

- Translation

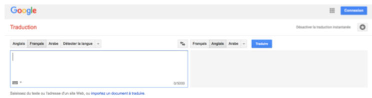

- Speech Recognition

- Speech Synthesis

- Source Separation

- Music Creation

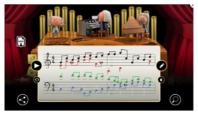

- Image Creation

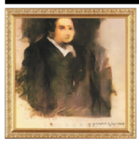

Deep Learning

- Overwhelming Success

- Simple Basic Receipt
  - Linear/Logistic Regression
  - Loss Function Minimization

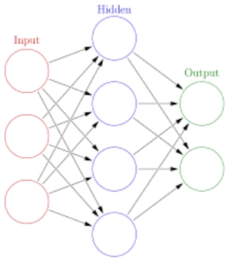

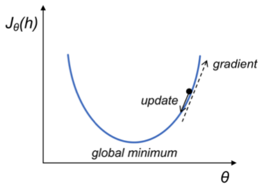

- Technical Improvements (since First Neural Networks)
  - Backpropagation, LSTM, Batch Normalization…
  - Loss Function Wide Application
    - » Meta-Level, ex: LSTM
    - » Constraints, ex: VAE

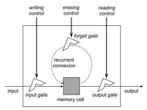

  - Optimized Implementations/Platforms
- Scale+
  - CPU
  - Data

Neural Networks in One Slide

Principle – Error Prediction/Classification Feedback

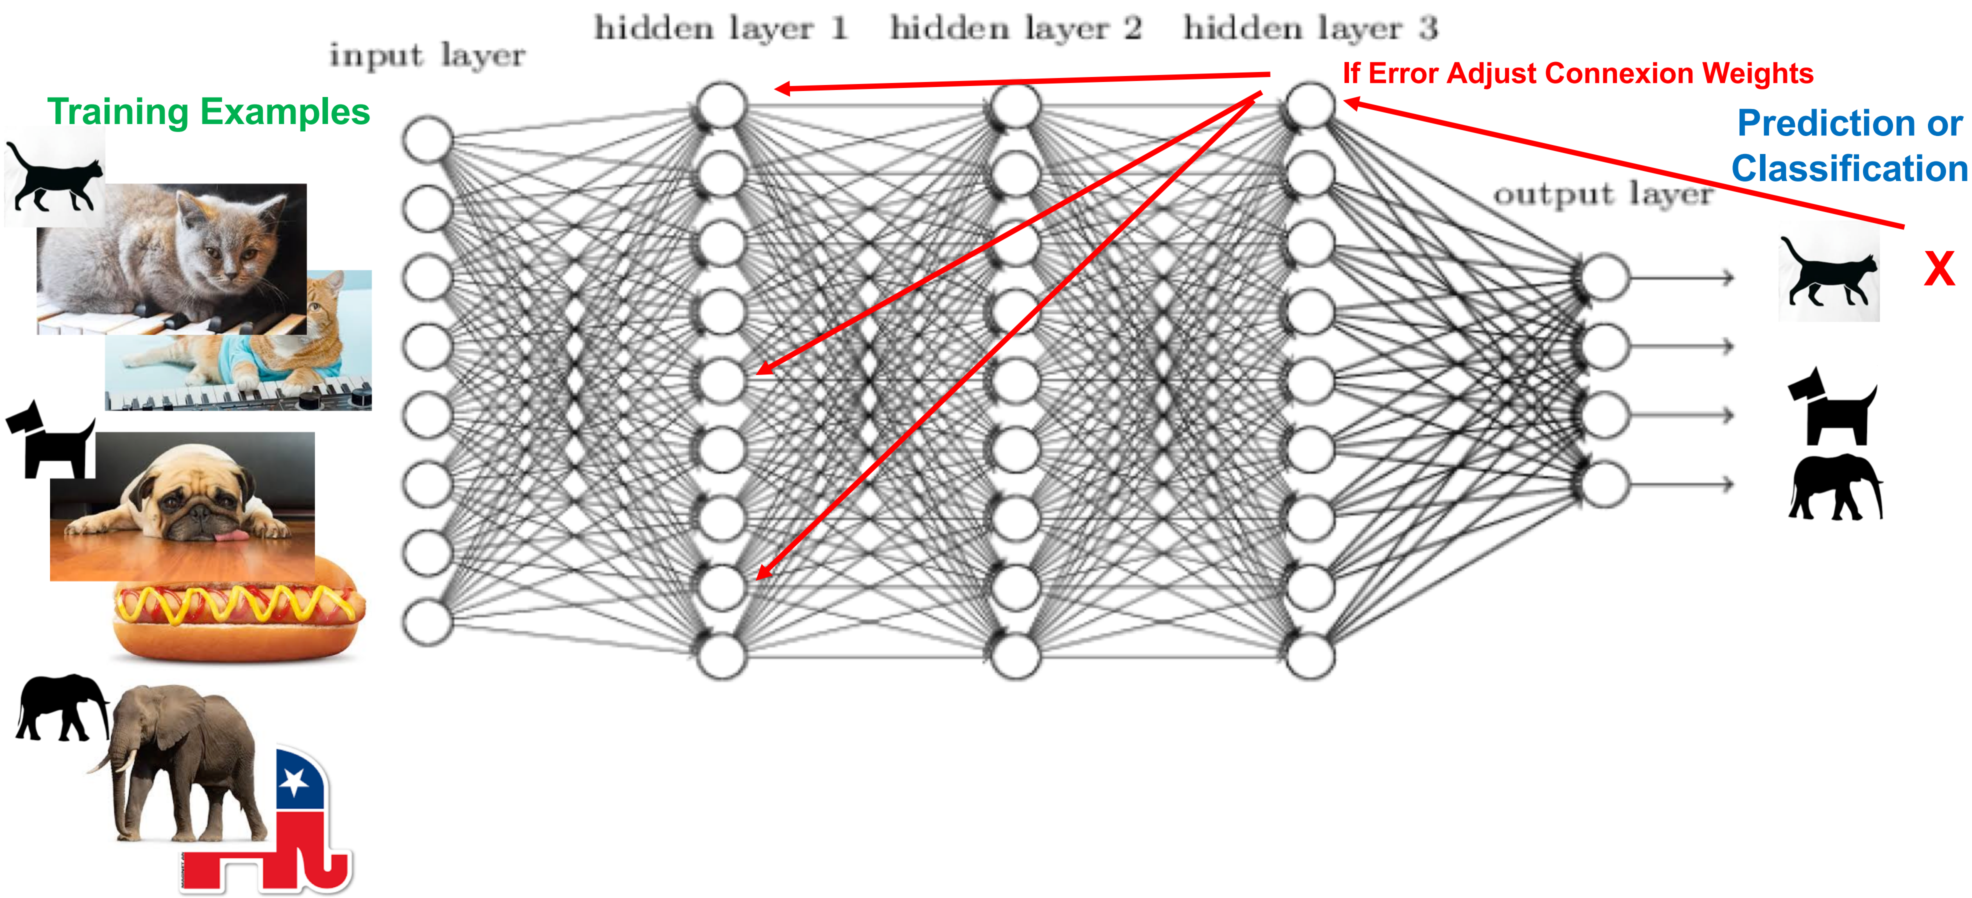

Neural Networks in ~~One~~ Two Slides

Principle – Error Prediction/Classification Feedback

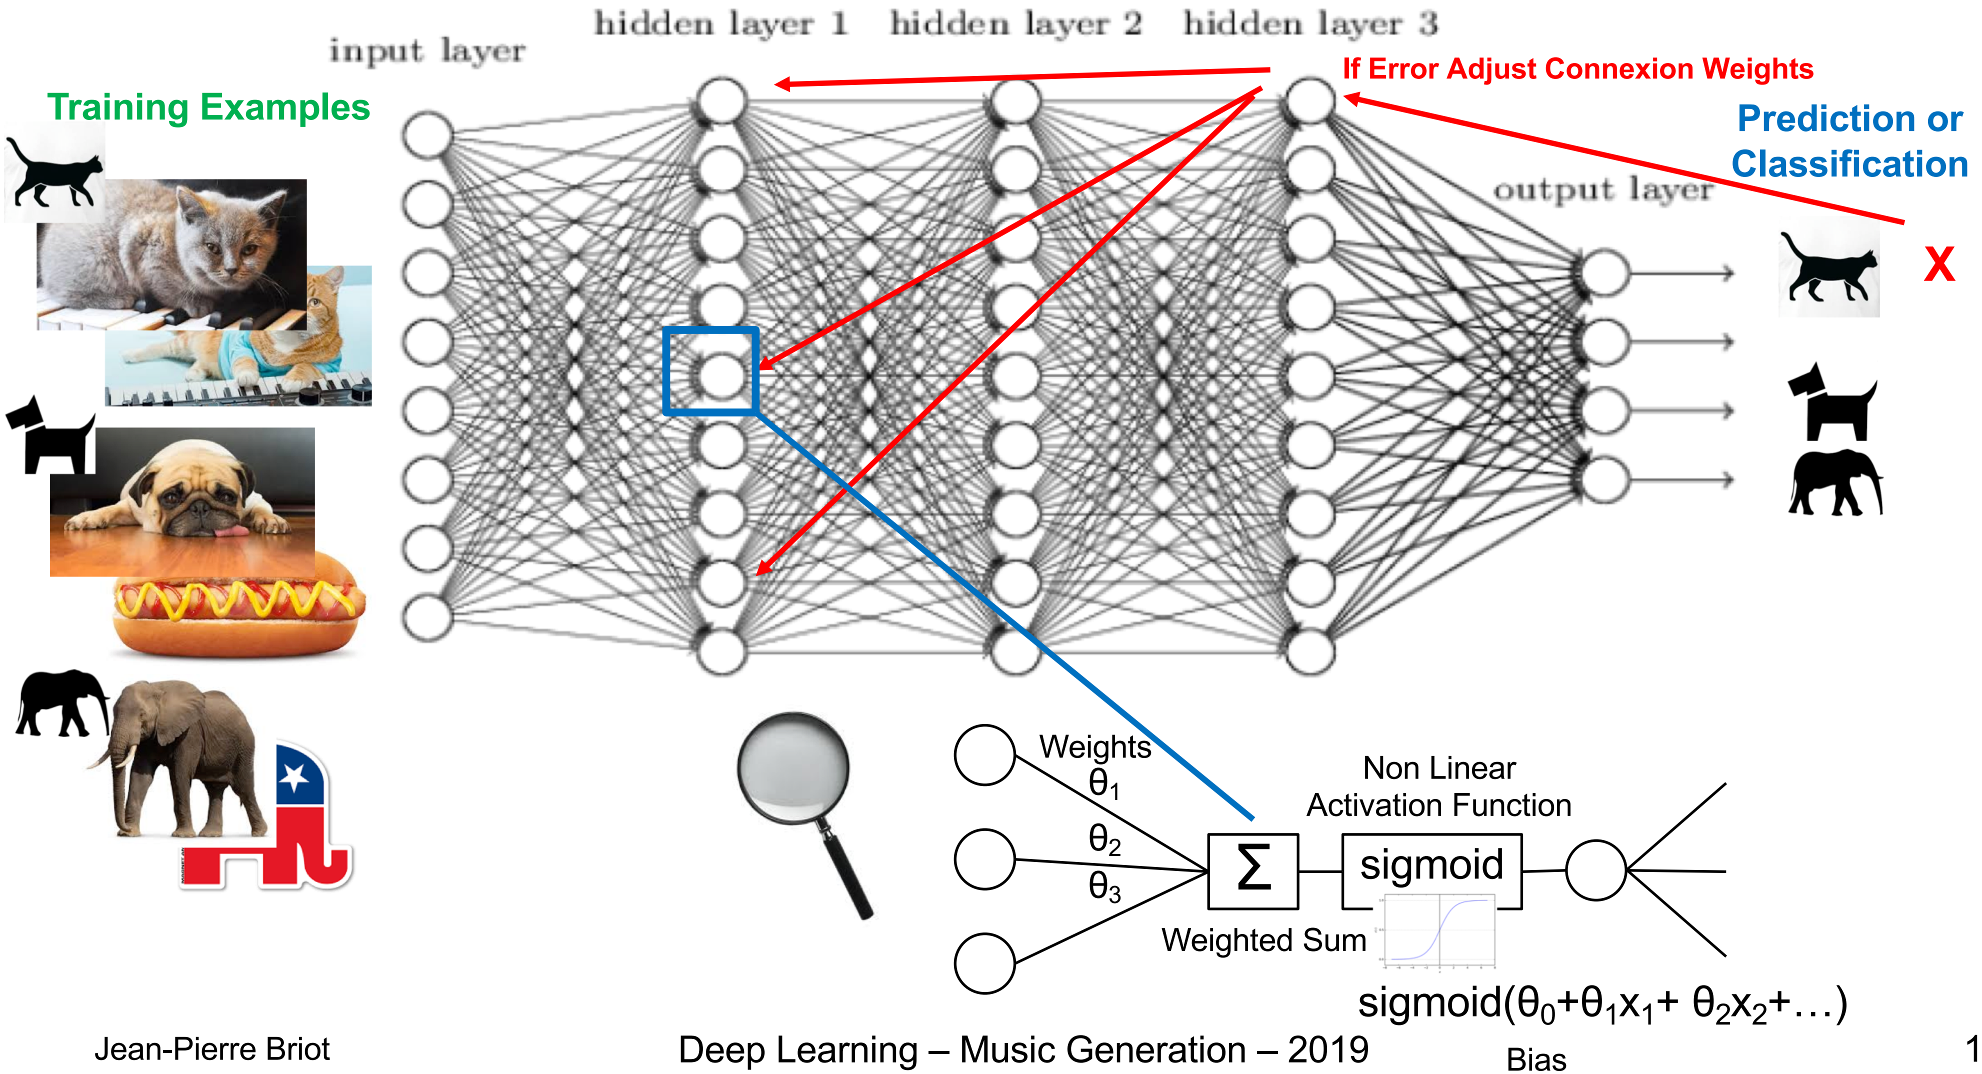

Number of Scientific Papers about Neural Networks and Music (Generation, Classification…) [Pons, 2018]

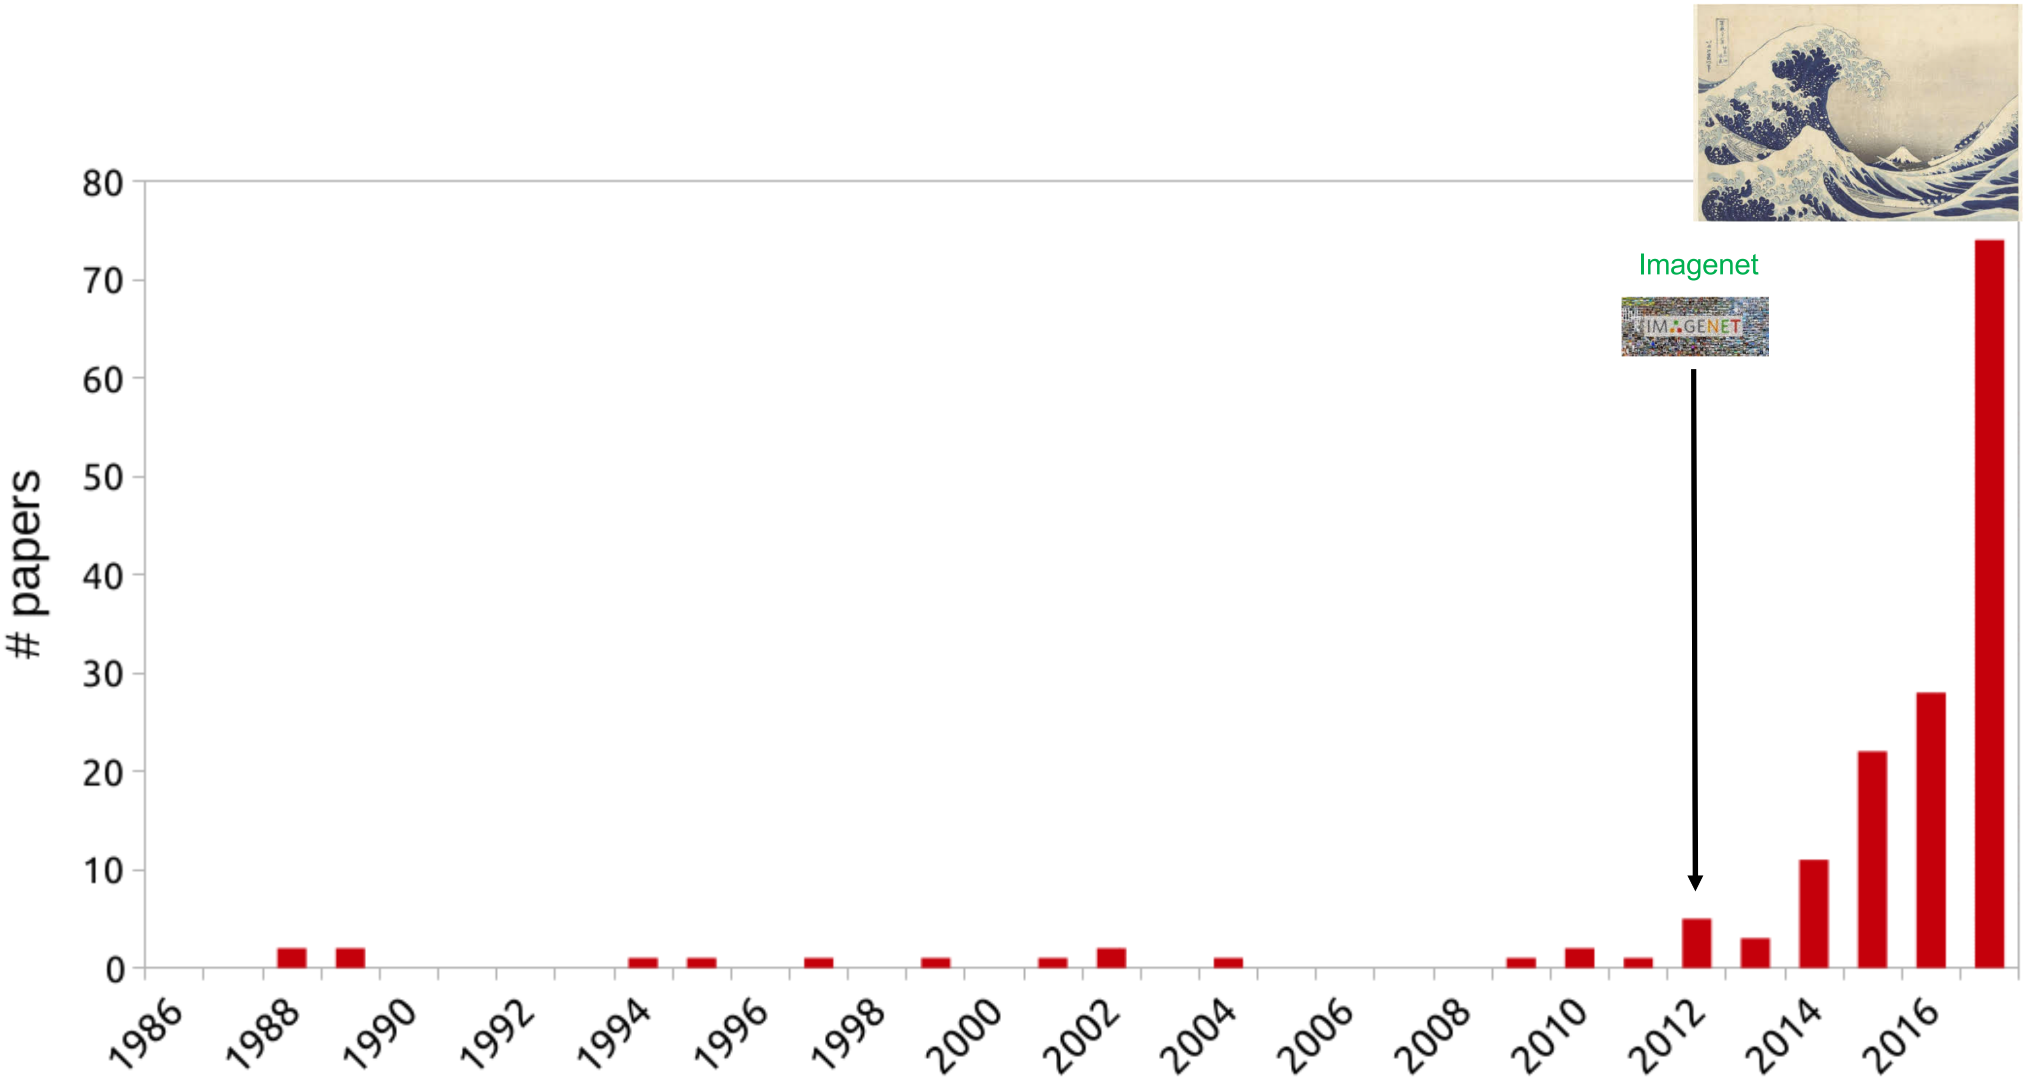

[Pons, 2018]

#Citations

Google Scholar

Samedi 19 octobre 2019

	#Citations	*Year*
Deep learning techniques for music generation-a survey JP Briot, G Hadjeres, F Pachet arXiv preprint arXiv:1709.01620	85	2017

» Computer Science » Artificial Intelligence

Computational Synthesis and Creative Systems

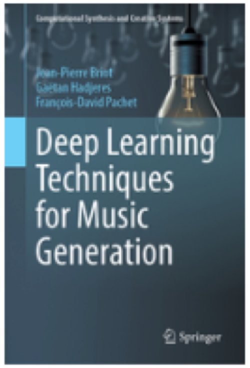

Deep Learning Techniques for Music Generation

Authors: **Briot**, Jean-Pierre, **Hadjeres**, Gaëtan, **Pachet**, François-David

Neural Networks Evolution

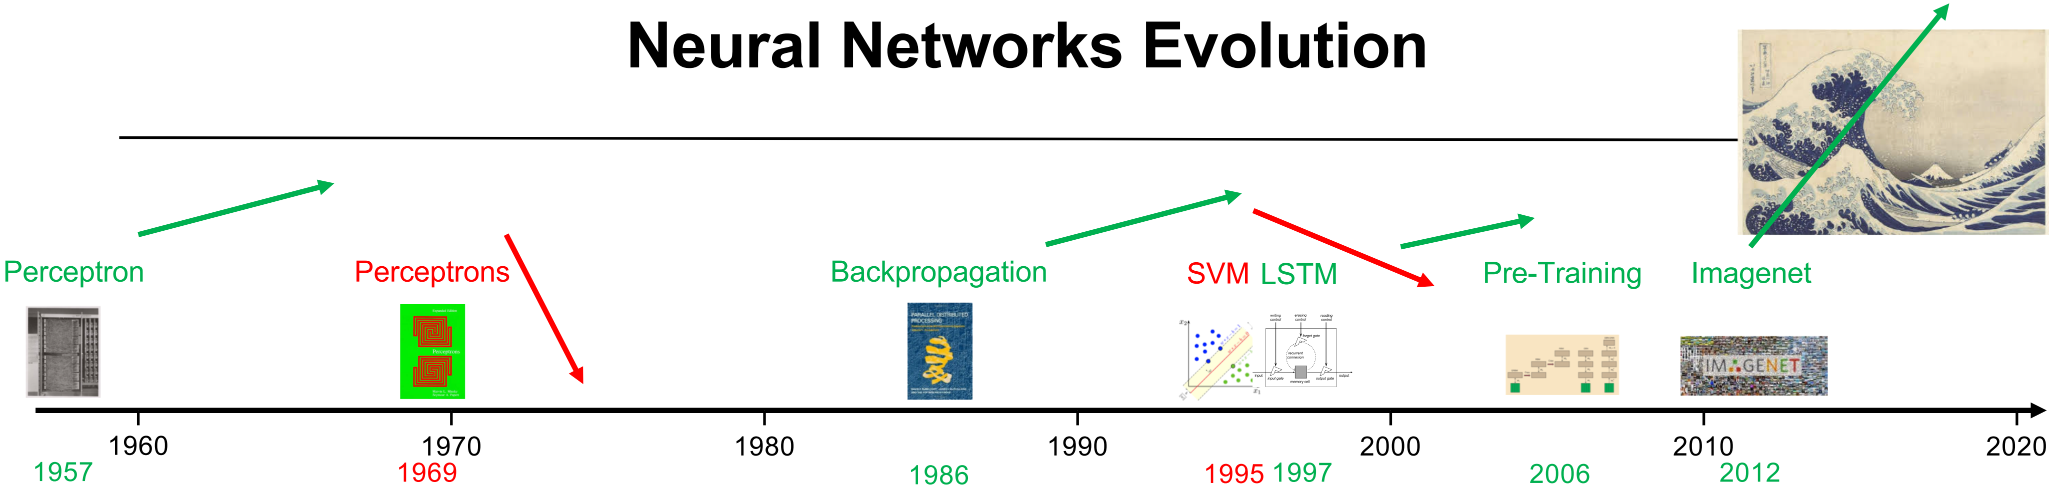

Symbolic vs Connexionist AI – History

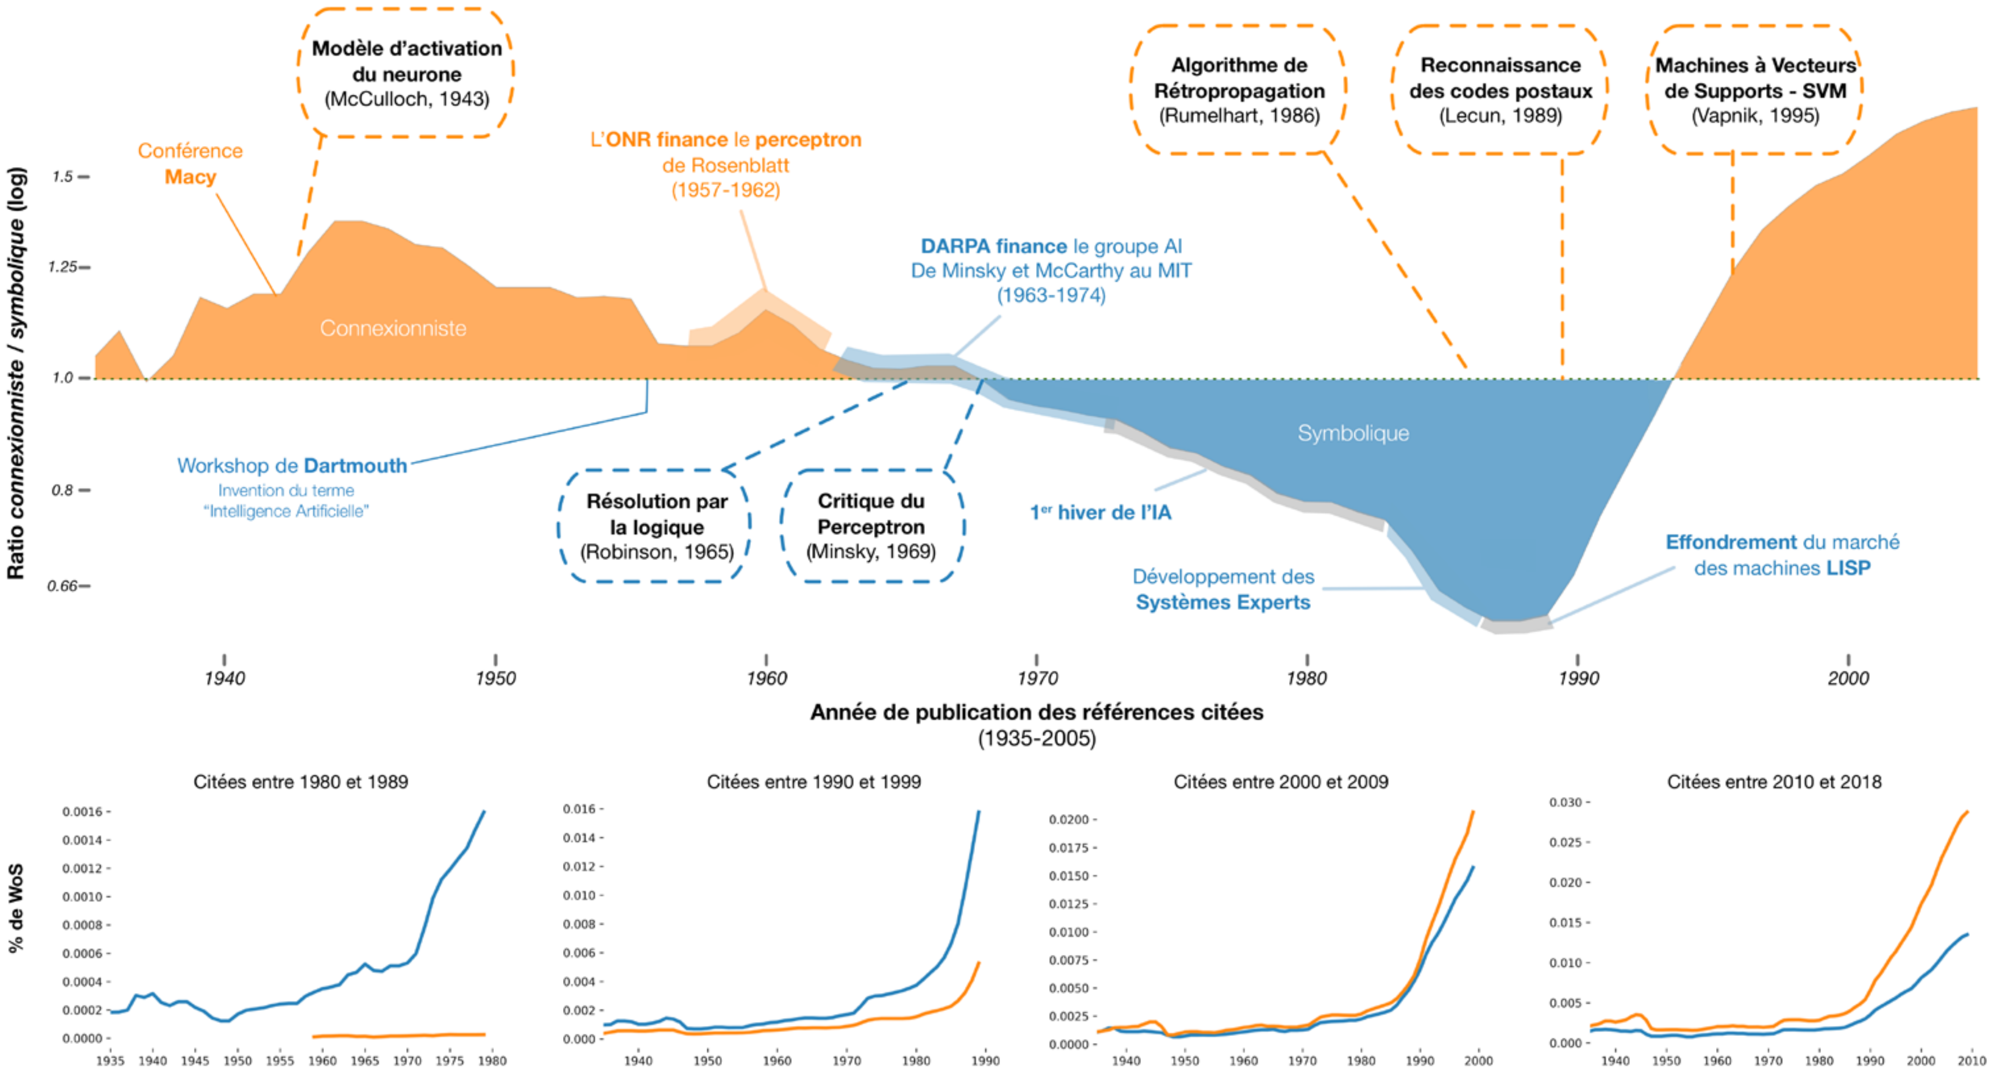

[Cardon et al., 2018]

Neural Networks 4 Music Generation Evolution

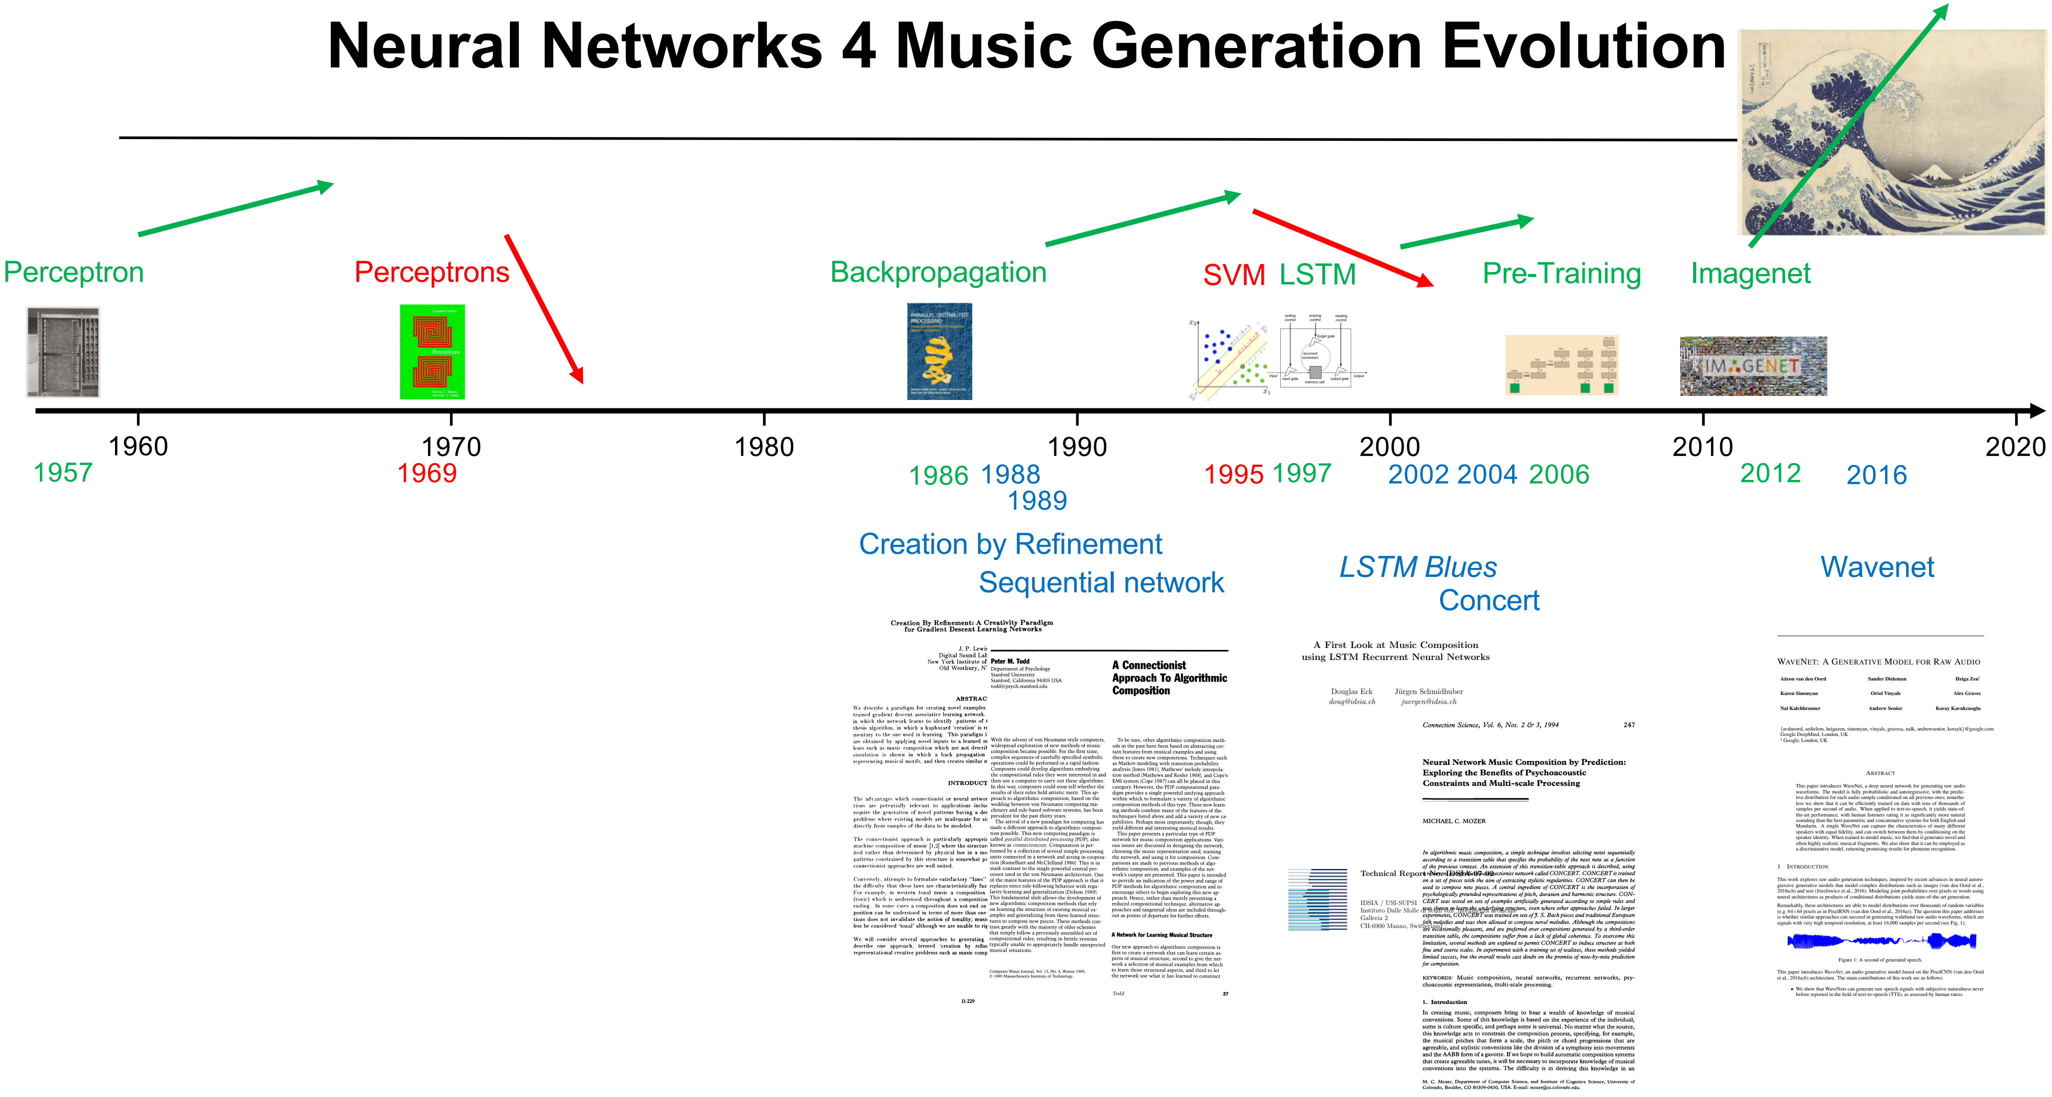

Neural Networks 4 Music Generation Evolution

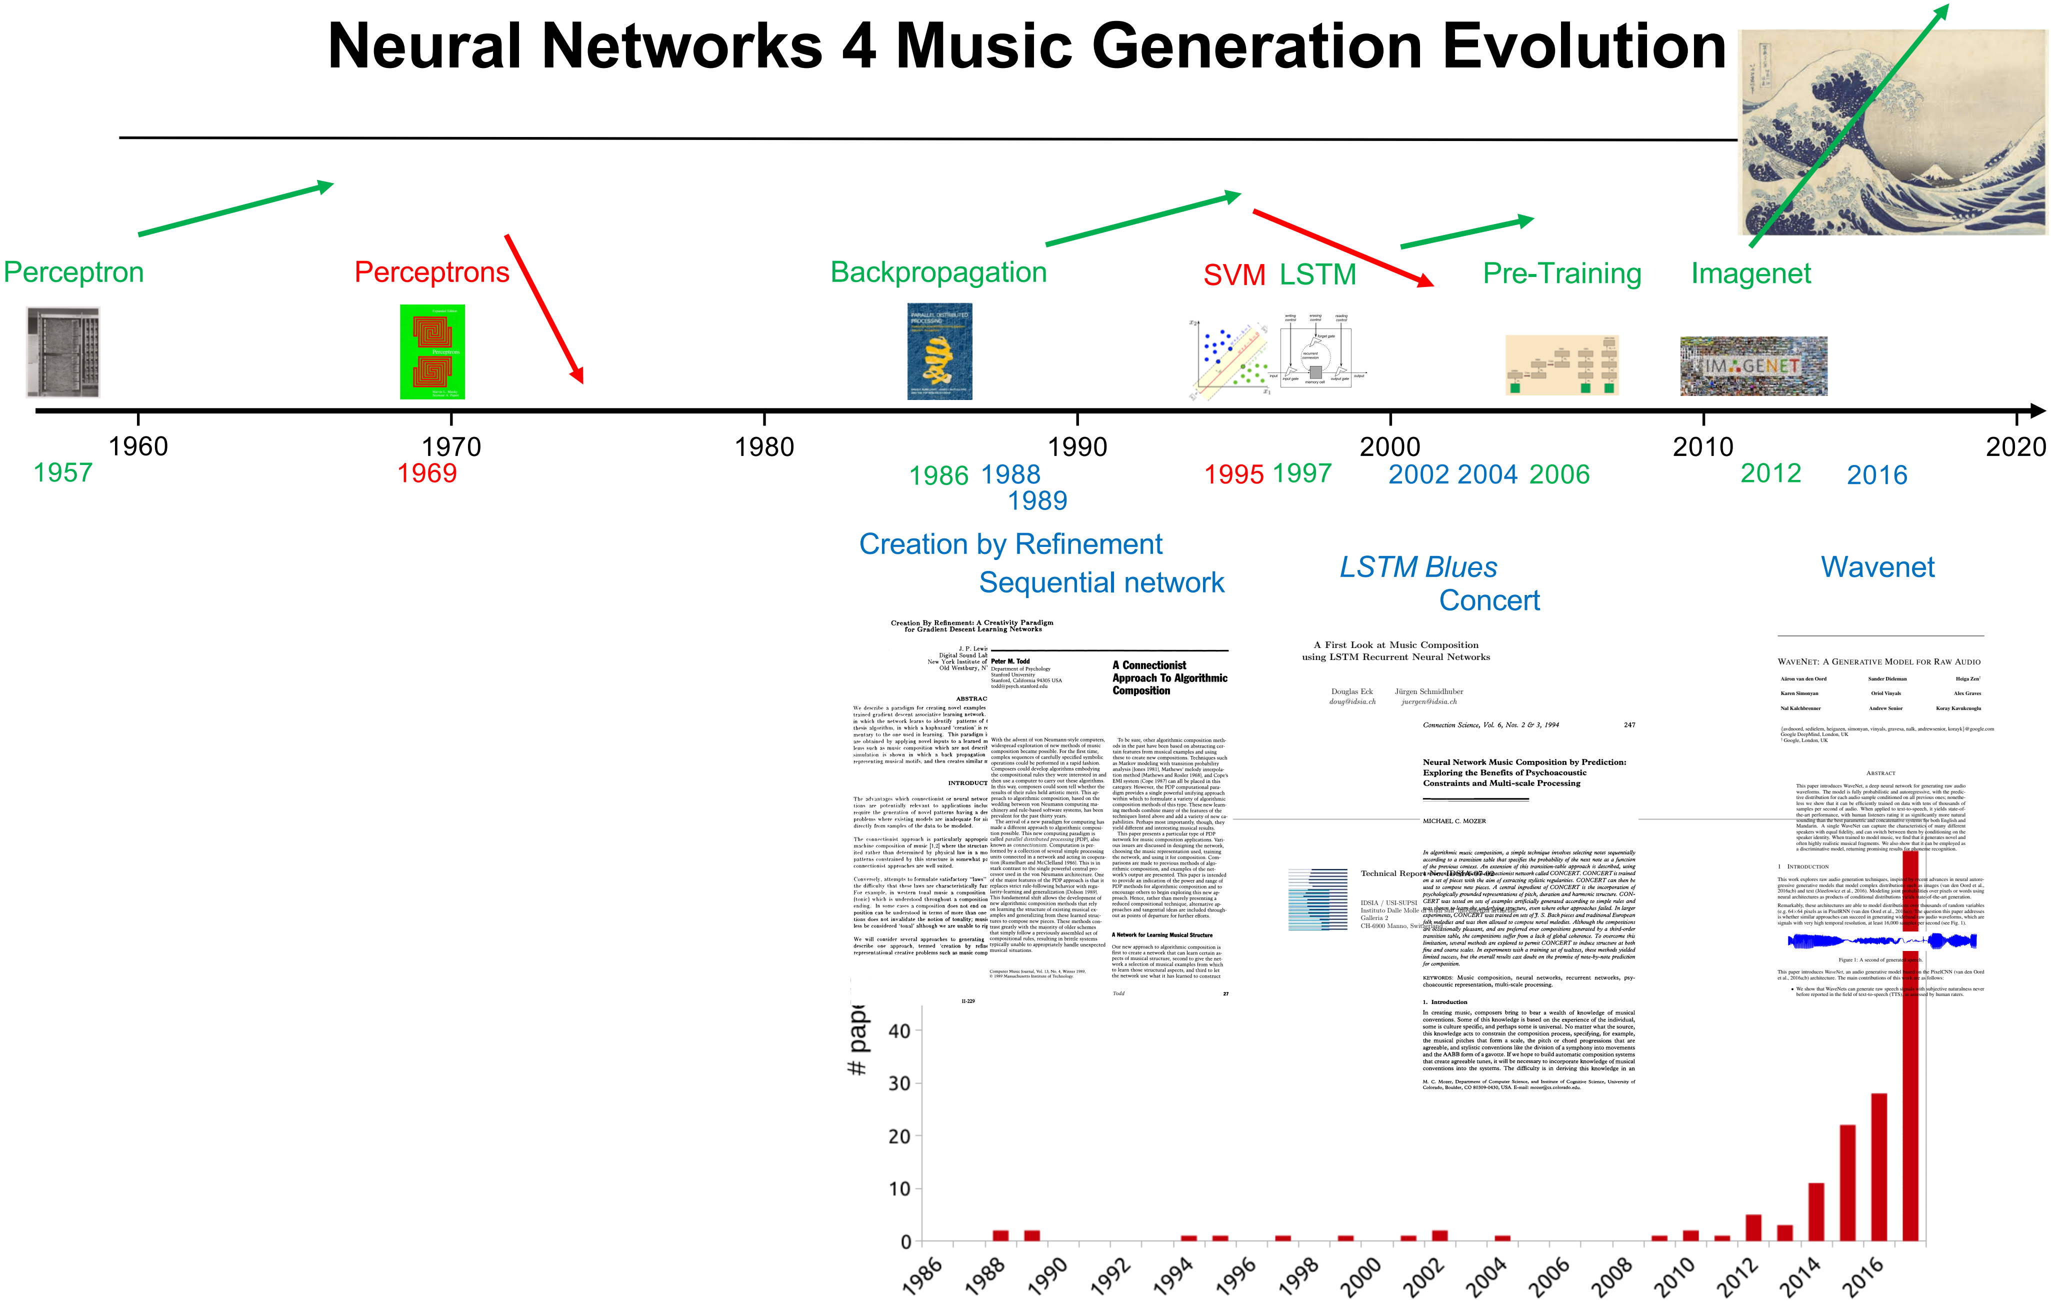

The Old Emperor Old Clothes

The Old Emperor Old Clothes (Neural Networks)

- Single Hidden Layer Neural Network
- Hand Made
- Technical Limitations
- Slow CPU
- Small memory
- Few Examples

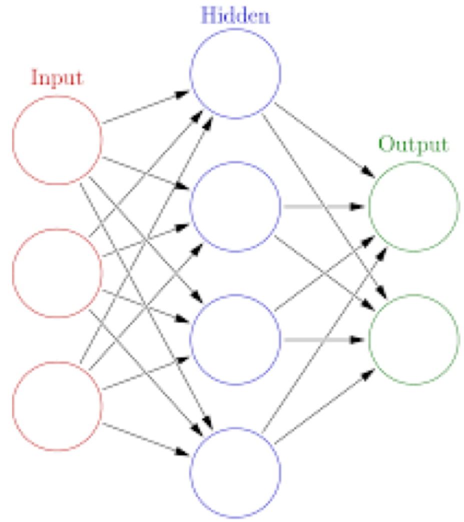

First Experiments in Using Artificial Neural Networks for Music Generation

1988–1989

- Lewis, J. P., Creation by Refinement: A Creativity Paradigm for Gradient Descent Learning Networks, International Conference on Neural Networks, San Diego, CA, USA, July 1988, pp. II-229–233.
- Todd, Peter M., A Sequential Network Design for Musical Applications, Proceedings of the 1988 Connectionist Models Summer School, CMU, June 1988, Touretsky, D., Hinton, G., Sejnowski, T. (eds), Morgan Kaufmann, pp. 76–84, 1989.
- Todd, Peter M., A Connectionist Approach to Algorithmic Composition, Computer Music Journal (CMJ), MIT Press, 13(4):27–43, 1989.

2004

- Mozer, M. C., Neural Network Music Composition by Prediction: Exploring the Benefits of Psychoacoustic Constraints and Multi-scale Processing, Connection Science, 6(2&3):247–280, 1994

Todd’s Architecture Variation [Todd, 1989]

Feedforward architecture
Iterative generation

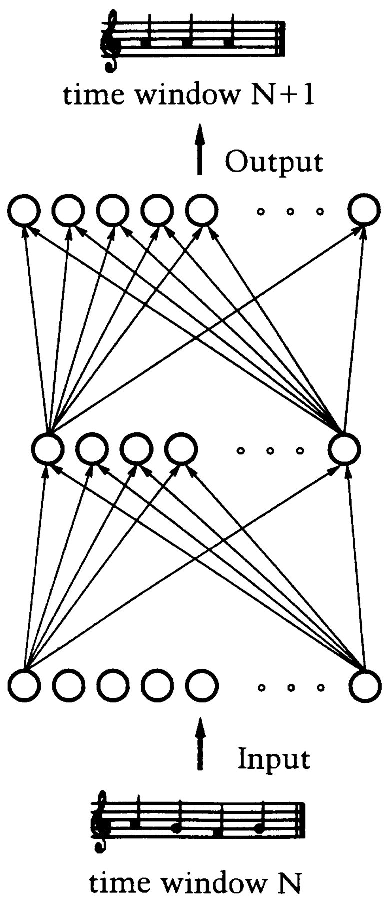

Recurrent architecture
Iterative generation

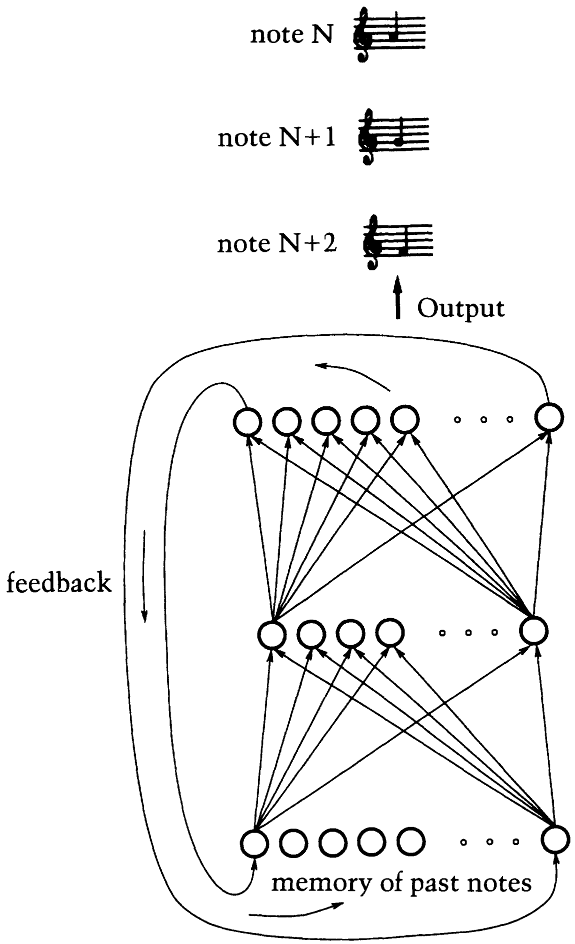

Recurrent + Conditioning architecture
Iterative generation

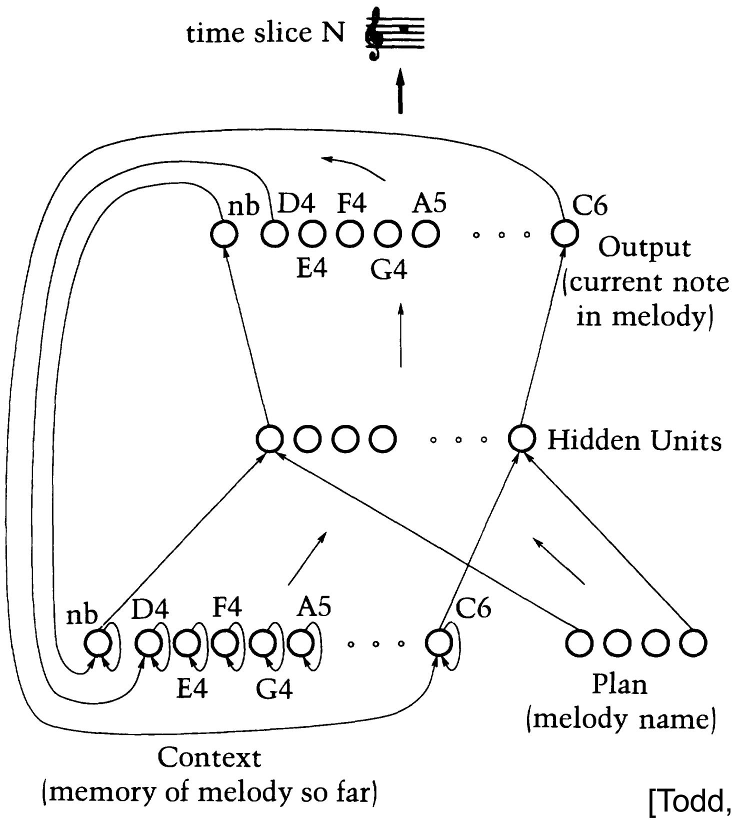

[Todd, 1988]

Todd's Conditioned Generation

Extrapolation

Original melody (plan 1.0)

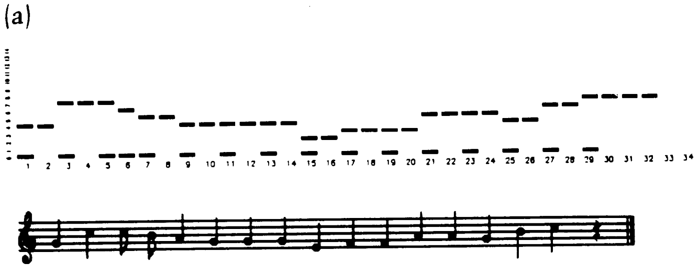

New melody (plan 0.0)

New melody (plan 2.0)

Interpolation

Original melody$_A$ (plan 0.0)

Original melody$_B$ (plan 1.0)

New melody$_{A\text{-}B}$ (plan 0.5)

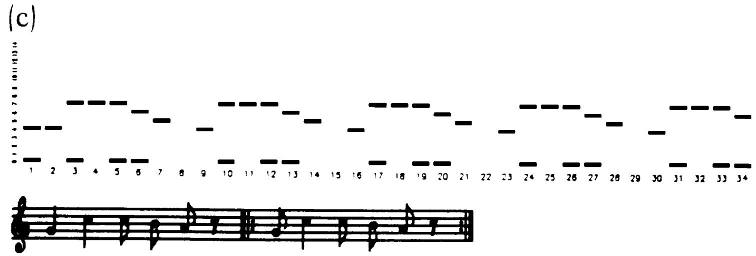

Todd's Architecture Prospects/Addendum (1/2) [Todd, 1989]

- Structure

One of the largest problems with this sequential network approach is the limited length of sequences that can be learned and the corresponding lack of global structure that new compositions exhibit. Hierarchically organized and connected sets of sequential networks hold promise for addressing these difficulties. Several ways of passing control back and forth between the interconnected networks will be described and the remaining issue of learning hierarchical structures will be addressed in this addendum.

- Hierarchy

One solution to these problems is first to take the sequence to be learned and divide it up into appropriate chunks (for instance, in the case of the sequence just presented, these could be A-B-C-D, E-E-E, A-B-C-D, and G-G). Next, train a sequential network to produce each of these subsequence chunks with a different plan. Finally, give this network the appropriate sequence of subsequence plans so that it will produce the chunks in the proper order to recreate the entire original pattern.

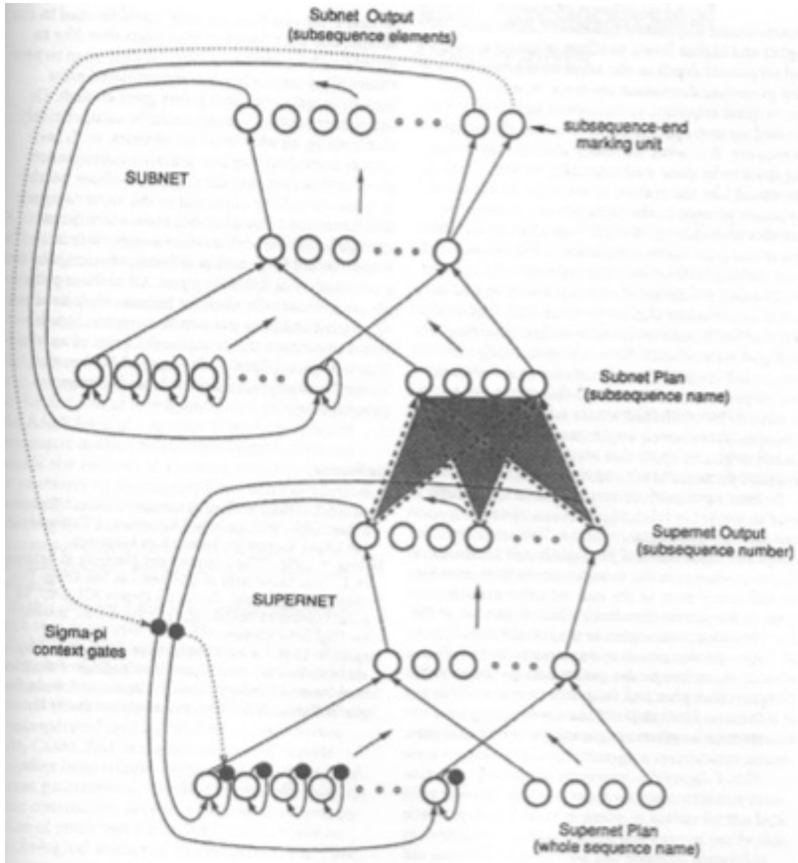

- Multiple Time/Clocks

Of course, one way to present this subsequence-generating network with the appropriate sequence of plans is to generate *those* by another sequential network, operating at a slower time scale. Then,

Todd's Architecture Prospects/Addendum (2/2) [Todd, 1989]

- Precursor of
- Hierarchy
  - Ex: MusicVAE [Roberts et al., 2018]

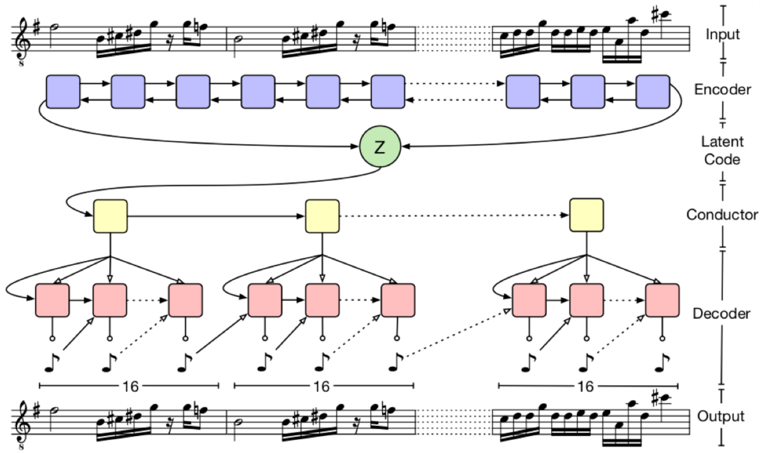

- Multiple Time/Clocks
  - Ex: Clockwork RNN [Koutnik et al., 2014]

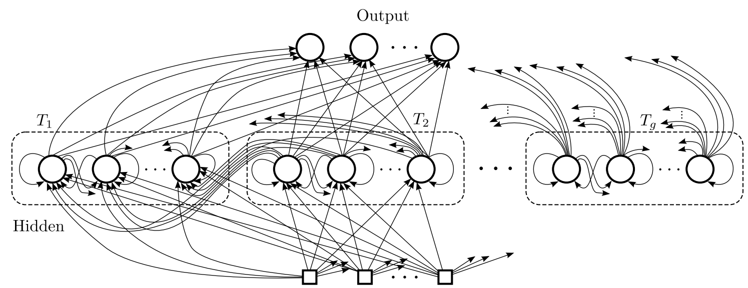

  - SampleRNN [Mehri et al., 2017]

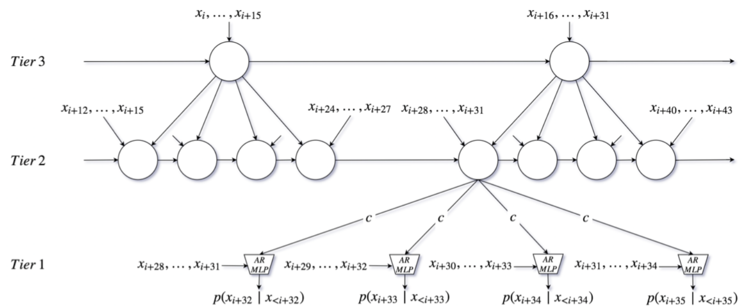

Lewis' Creation by Refinement (1/4) [Lewis, 1988]

- Training on 30 Manually Generated 5-Note Melodies
- 7 Possible Notes (from C to B, without alteration)
- Well Formed
  - Possible Intervals:
    - » Unison, 3rd, 5th,
    - » Scale Degree Stepwise Motion
- Poorly Formed
  - Excessive Motion or Excessive Repetition
- Binary Classification Training
  - Well or Poorly Formed

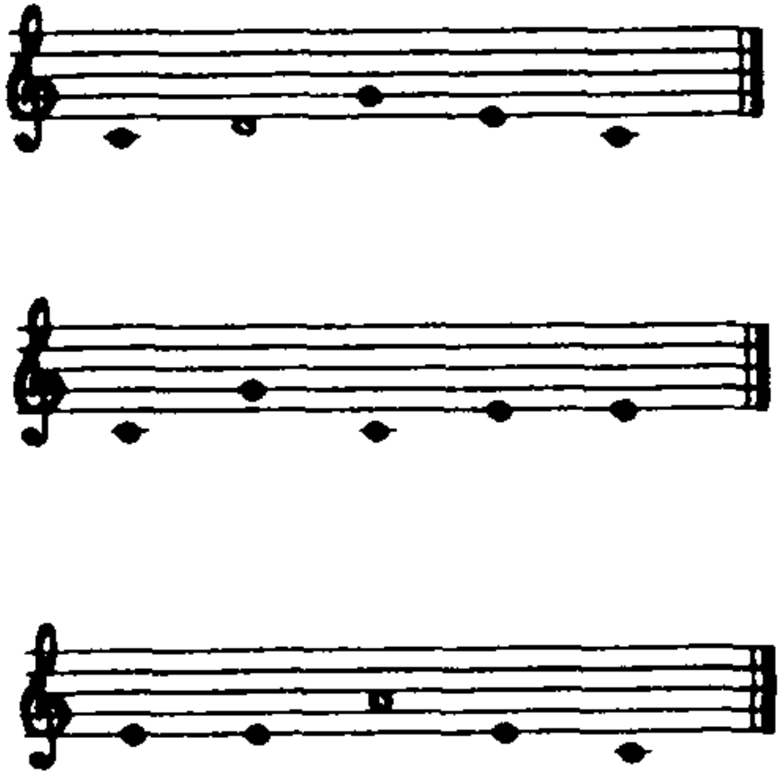

Ex. of Training Examples

Lewis' Network Architecture

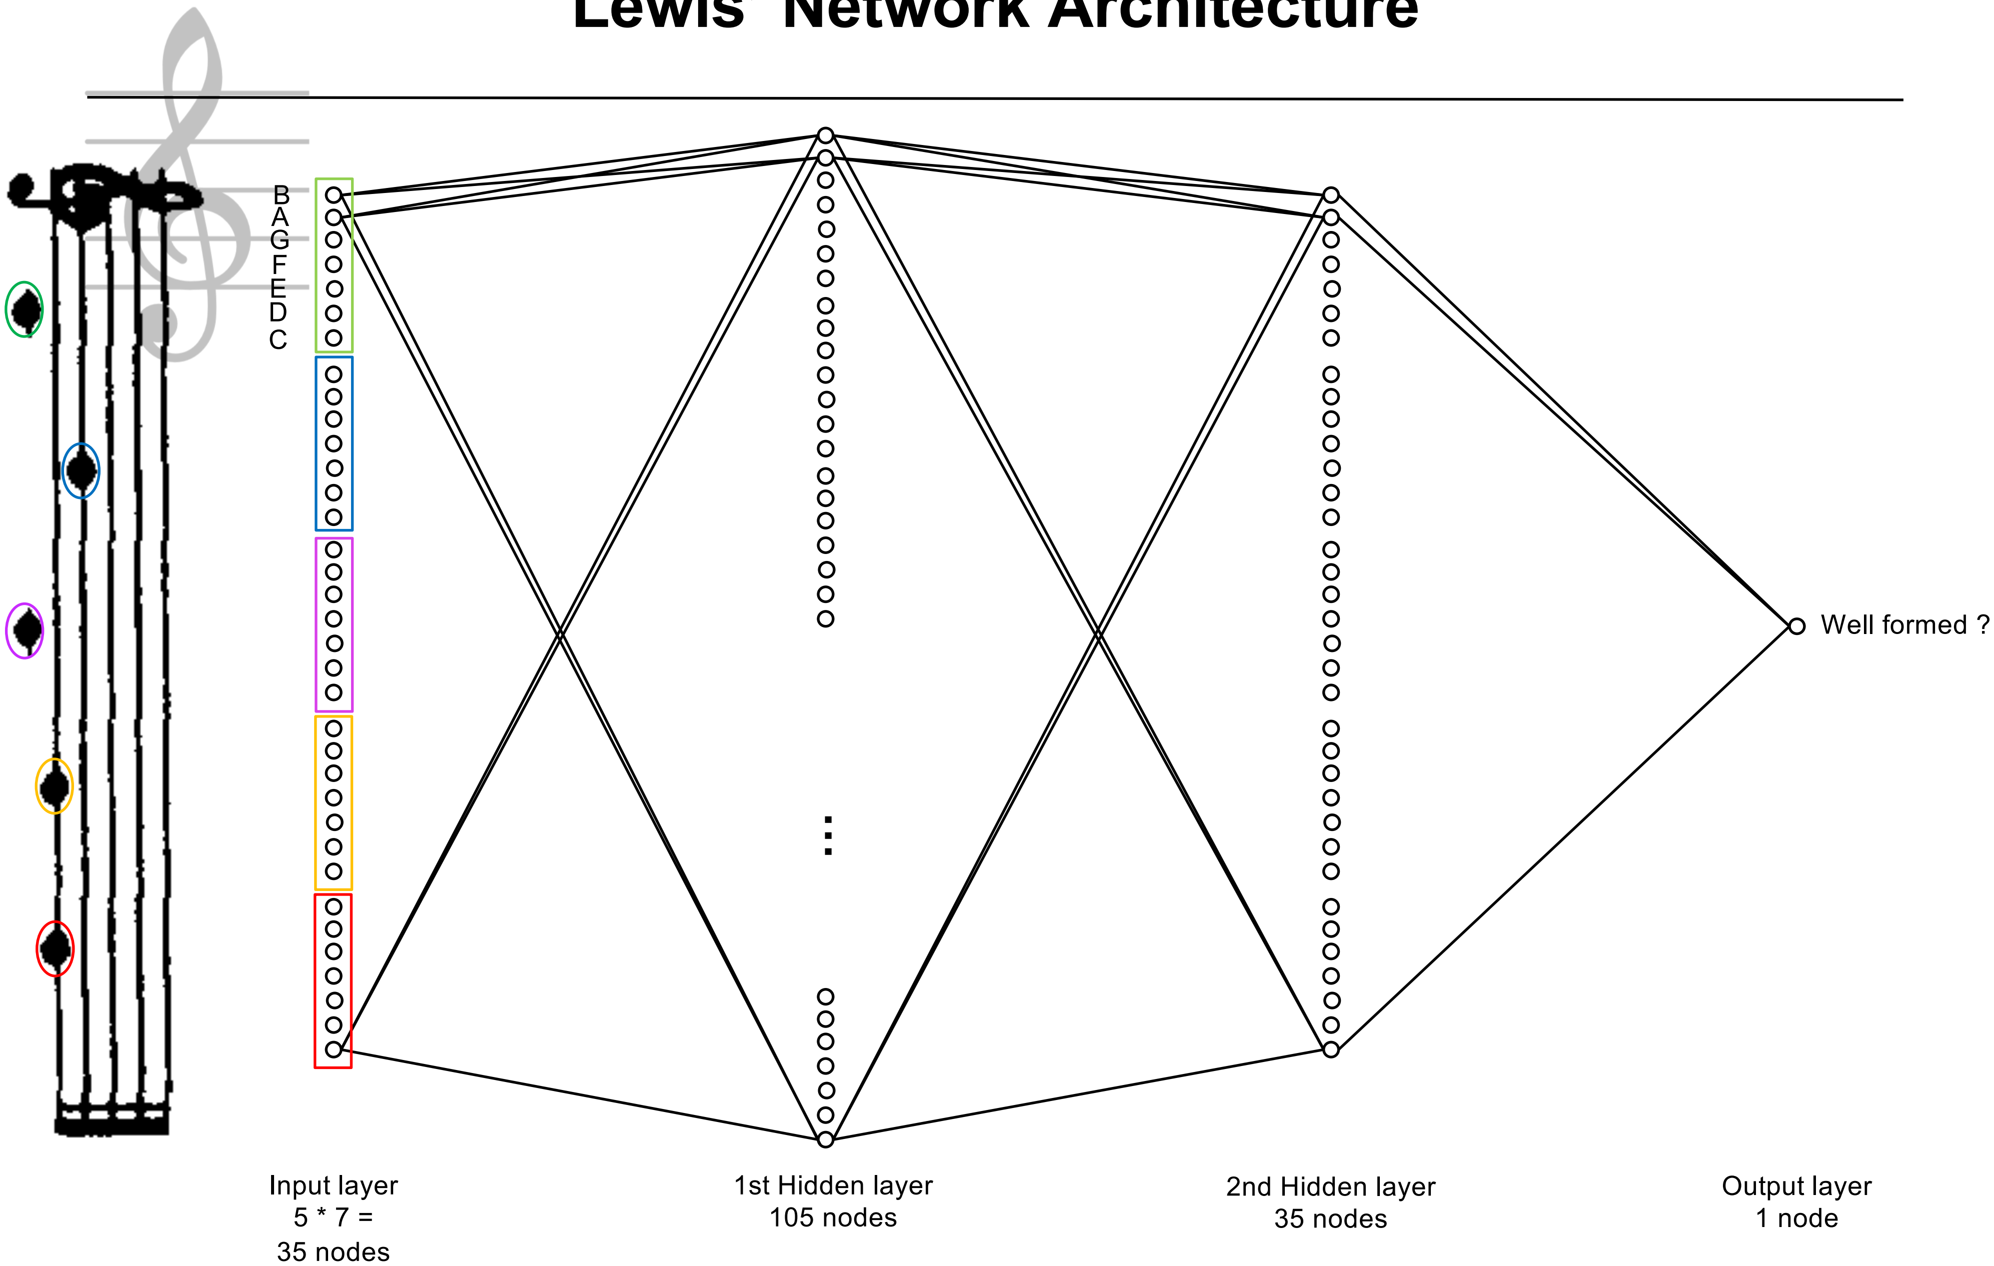

Lewis' Creation by Refinement (1/6)

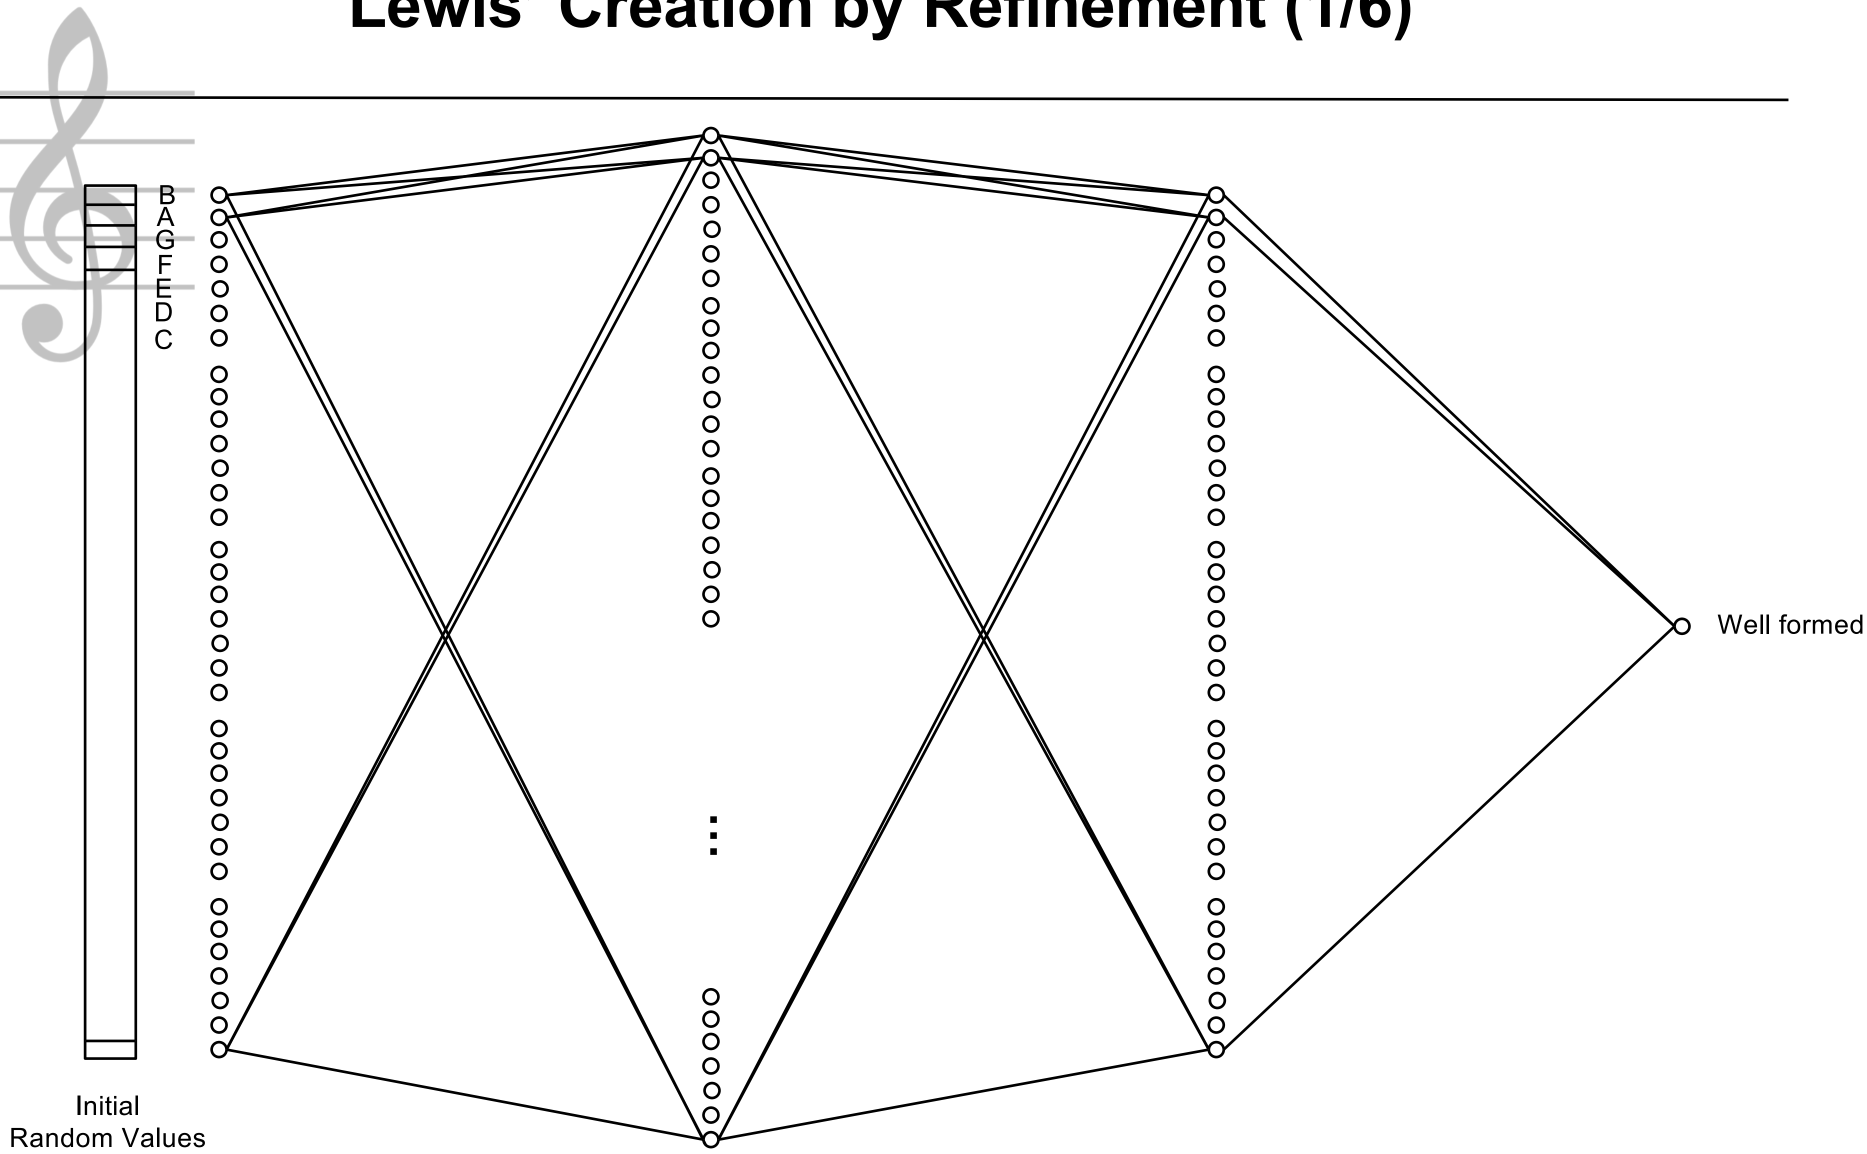

Lewis' Creation by Refinement (2/6)

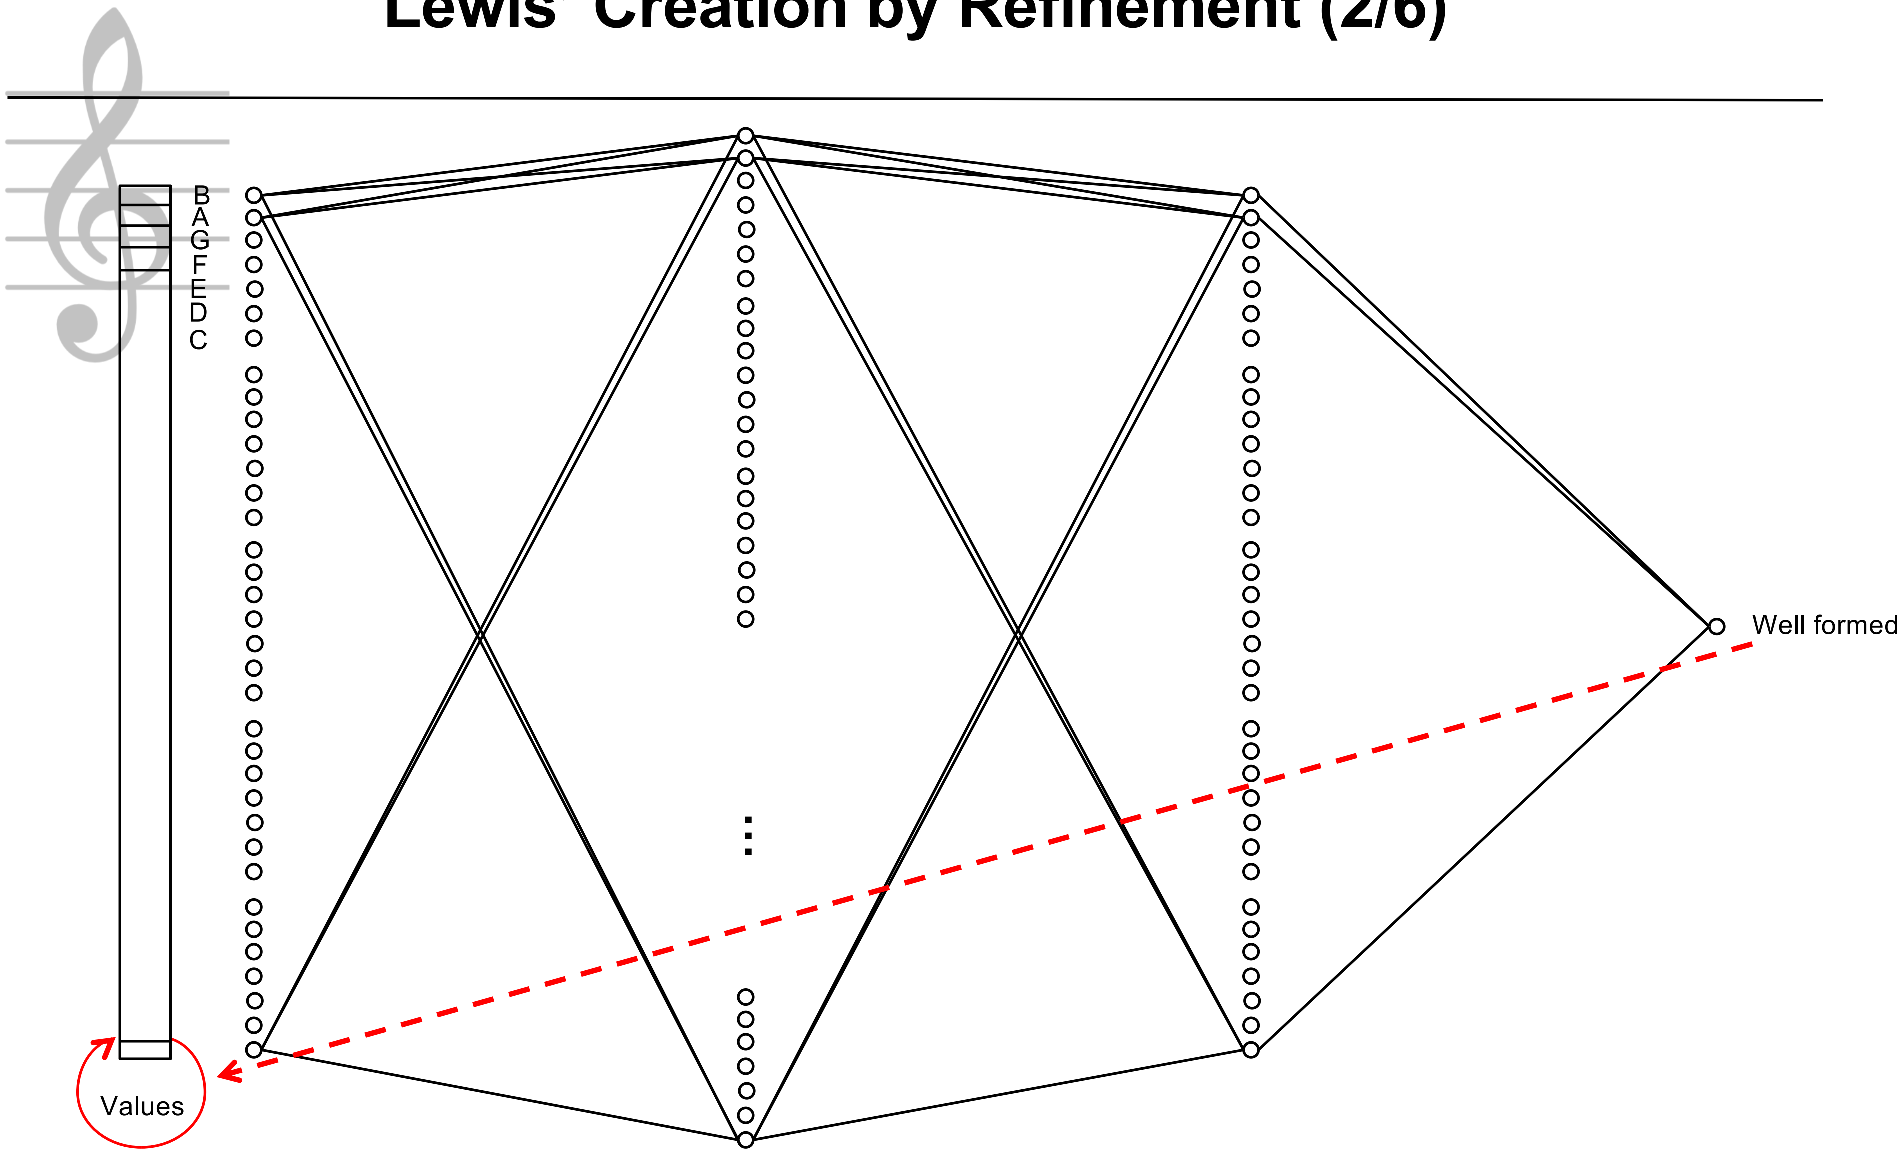

Input Values are Incrementally Manipulated

Under the Control of a Gradient Descent on Error in Predicted Well Formed

Lewis' Creation by Refinement (3/6)

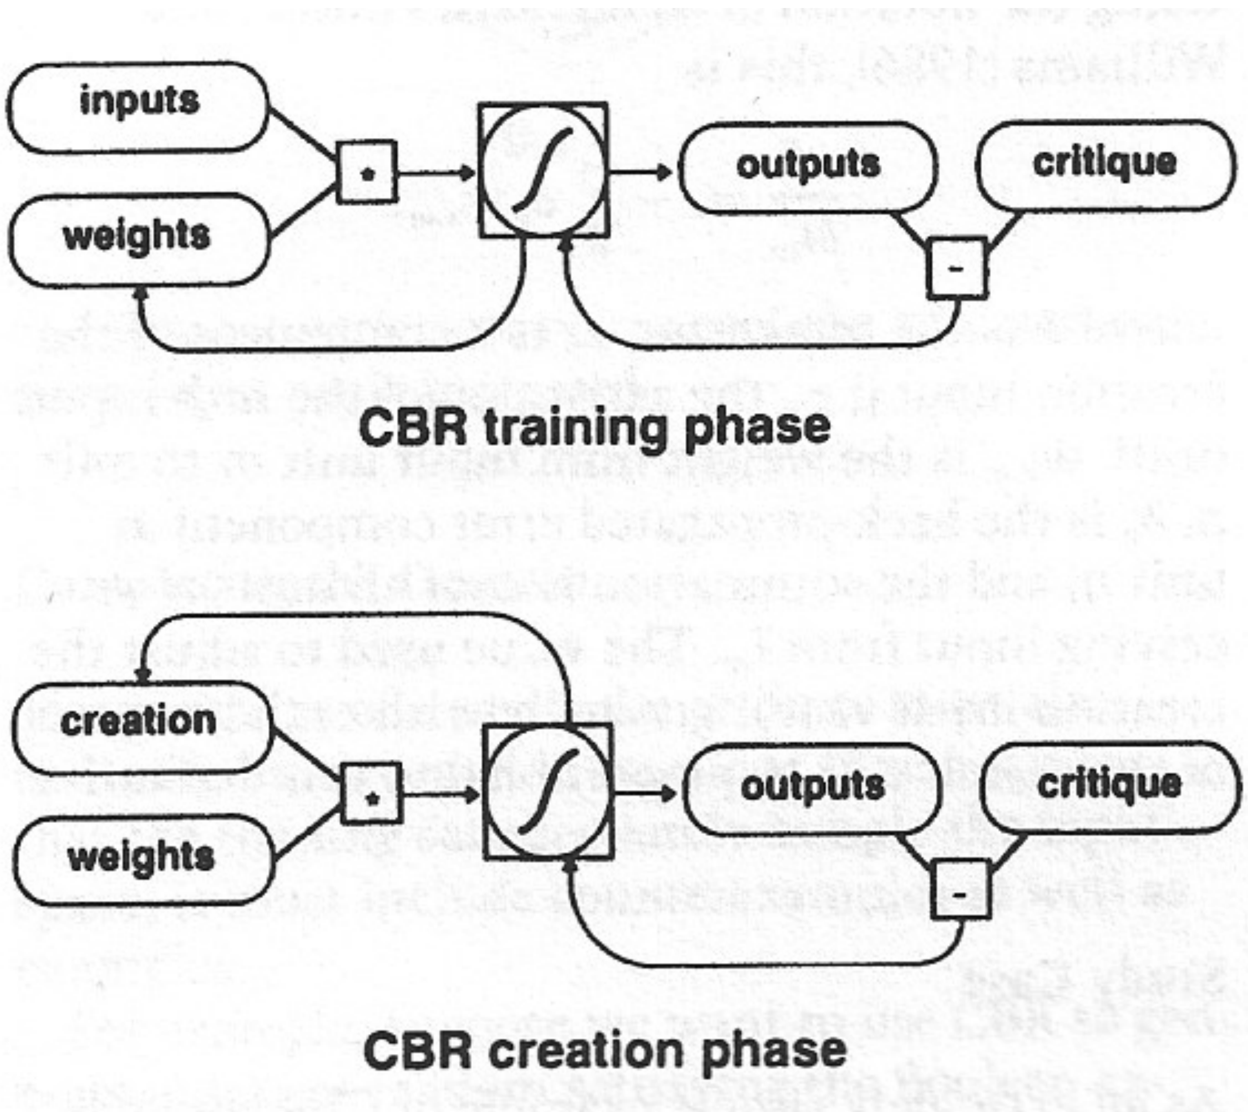

Lewis' Creation by Refinement (4/6)

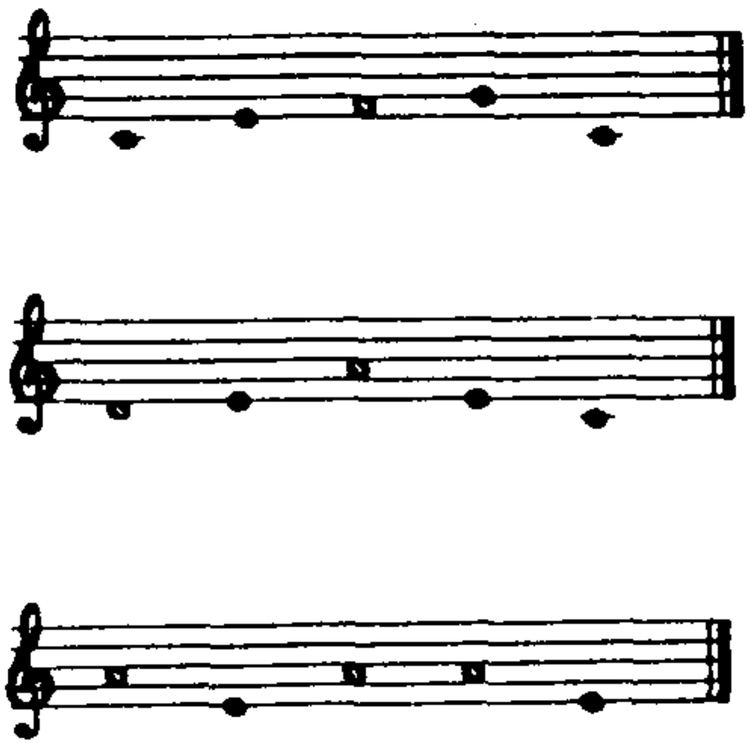

Ex. of Melodies Created by Refinement

- The Network Learned Preference for Stepwise and Triadic Motion

Lewis' Creation by Refinement (5/6)

- Attention

Attentional CBR

In order to partition a large problem into manageable subproblems, we need to provide both an attention mechanism to select subproblems to present to the network and a context mechanism to tie the resulting subpatterns together into a coherent whole. A context mechanism can be provided by context inputs, which during the creation phase are clamped to the values of the surrounding and previously constructed pattern. As an example, to produce elaborations on a short phrase, the training set inputs would consist of sample phrases paired with corresponding embellished phrases (possibly using a suitable null-note representation to allow different phrase lengths), and the critique would (as usual) consist of some critique of the character of the embellishment. In the creation phase, the embellished inputs would be set to random values, but the context inputs would be clamped to the phrase itself.

- Hierarchy

The author's experiments have employed *hierarchical* CBR. In this approach, a developing pattern is recursively filled in using a scheme somewhat analogous to a formal grammar rule such as $ABC \rightarrow AxByC$, which expands the string without modifying existing tokens. That is, three tokens (for example, musical notes) labeled A,B,C will be expanded with two additional tokens x,y inserted in the indicated positions. The expanded string $AxByC$ may be rewritten further using a suitable scheme.

Ex. of Melodies Created by Hierarchical Refinement (ABCD -> ABxCD scheme)

Lewis' Creation by Refinement (6/6)

- Reinforcement

Reinforcement CBR

Developing the training set is probably the most difficult aspect of employing CBR (and other supervised learning algorithms). In *reinforcement* CBR some or all of the training set is produced automatically, by completing the domain, rather than being compiled by the experimenter as in the standard supervised learning paradigm. In this scheme, the training phase is interrupted at intervals, and the creation phase is invoked. The resulting creations are evaluated by the experimenter and are added to the training set with a corresponding critique if they are judged to extend the existing training set. After the training set is extended, the net is retrained, followed by the accumulation of new examples, etc., until all sample creations are judged satisfactory by both the experimenter and the network.

Not Reinforcement learning

Created Melodies which are Liked are Added to the Training Set

Lewis' Creation by Refinement Pioneering (1/3)

- Precursor of
- Gradient Descent Input Manipulation [Briot et al., 2017]
- Ex: DeepHear [Sun, 2016]
  - Melody Consonant Accompaniment Creation

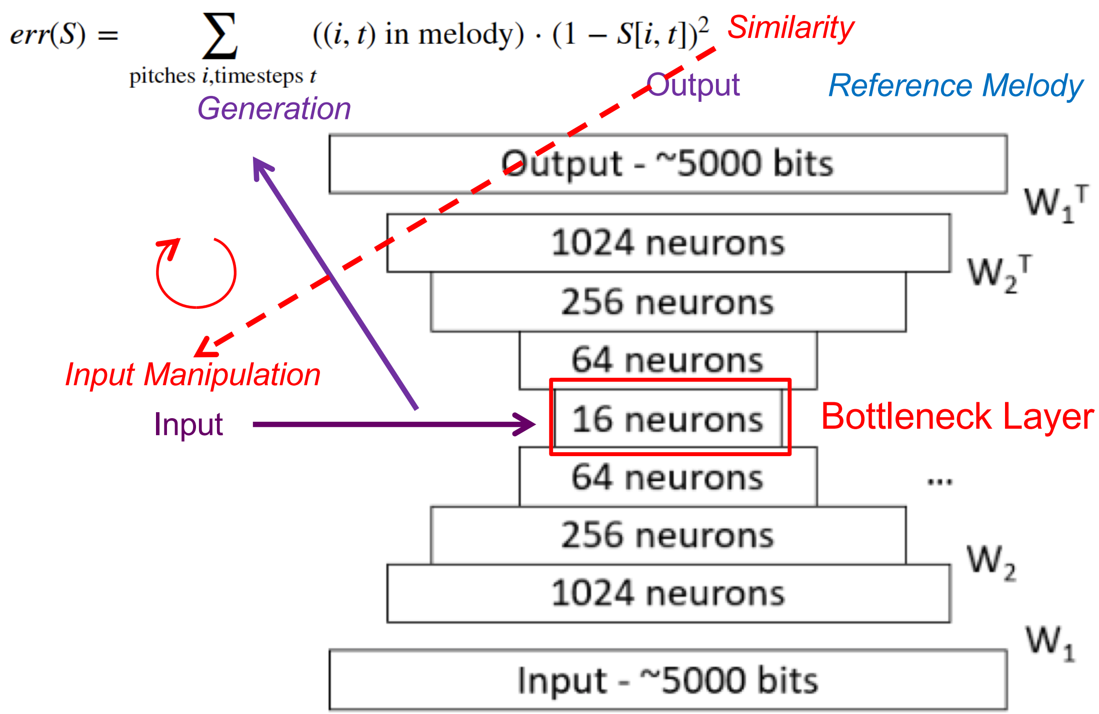

https://fephsun.github.io/2015/09/01/neural-music.html#

Lewis' Creation by Refinement Pioneering (2/3)

- Precursor of
- Gradient Ascent Input Manipulation [Briot et al., 2017]
- Ex: DeepDream [Mordvintsev et al. 2015]
  - Motif Detector Neuron Activation Maximization

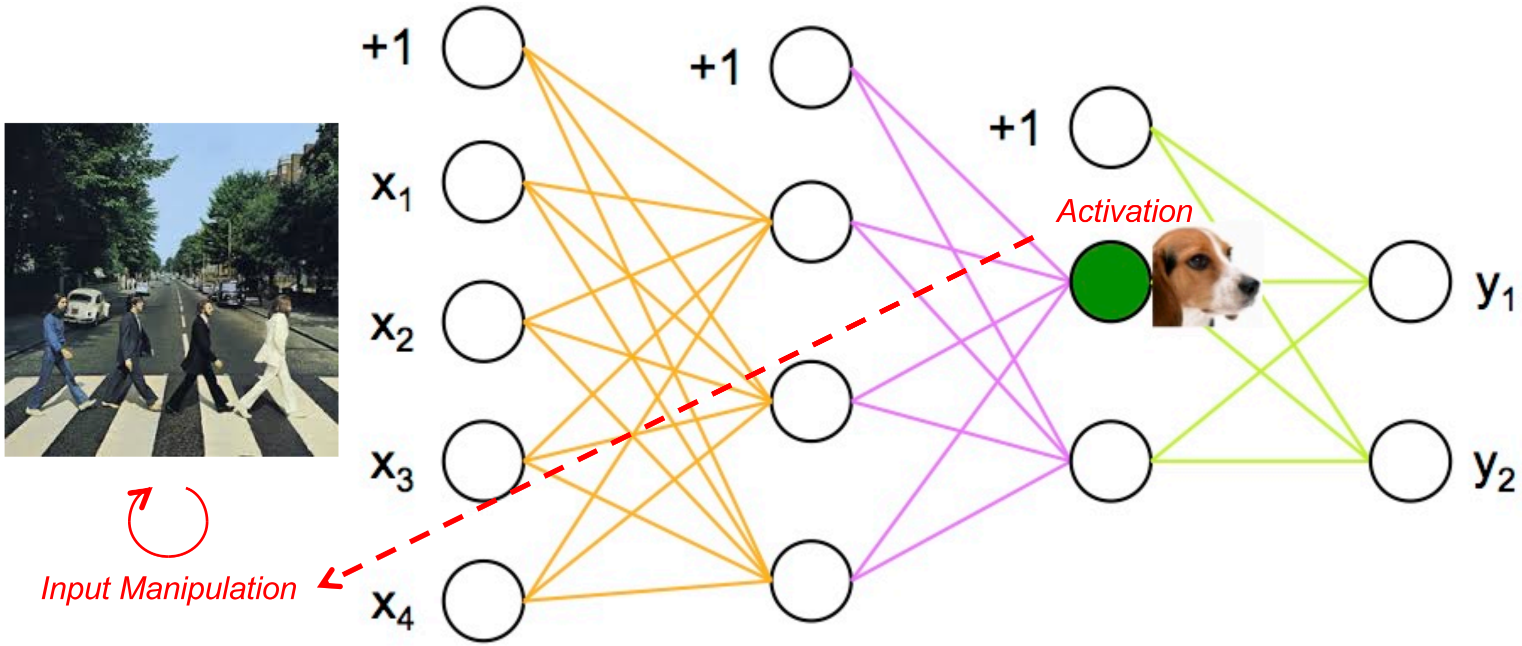

Lewis' Creation by Refinement Pioneering (3/3)

Initial Image

Deep Dream Image

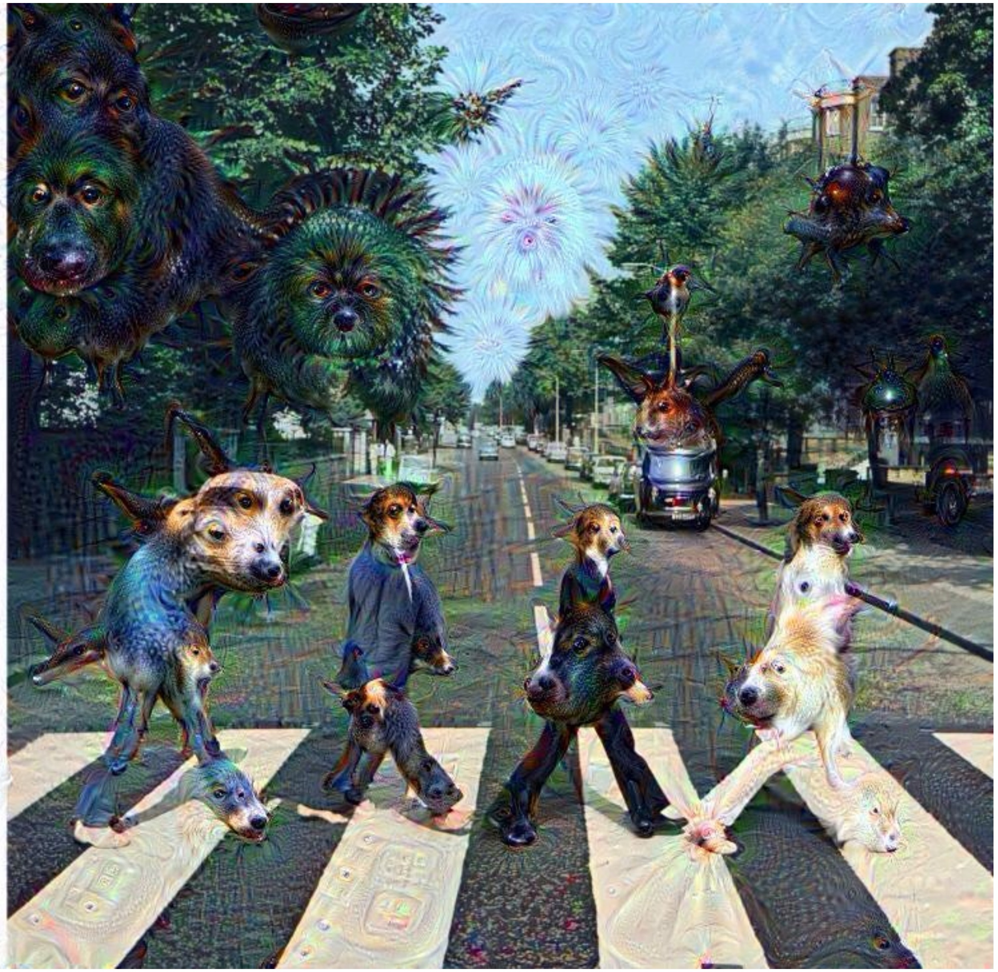

Structure Imposition (1/2) [Lattner et al., 2016]

- Constrained sampling, C-RBM [Lattner et al., 2016]
- Convolutional Restricted Boltzmann Machine (RBM)

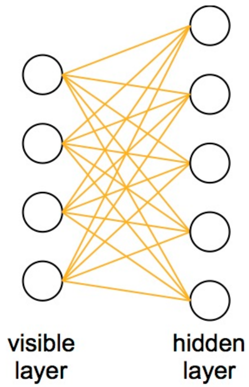

- Combination of:
  – Input Manipulation guided by Gradient Descent of current sample
    » to impose Higher-Level Structure/Constraints:
      - Structure (Structure Repetition, Ex: AABA), via Self-Similarity Matrix
      - Tonality, via Similarity of Distribution of Pitch-Classes
      - Meter (Rhythm Pattern/Signature and Beat Accent)

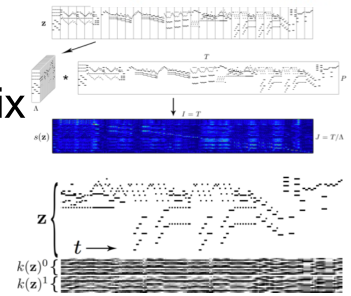

  – Sampling Control, by Selective Gibbs sampling (SGS)
    » at a Selected Low-Level (subset of variables)
    » to realign selectively the sample to the learnt distribution

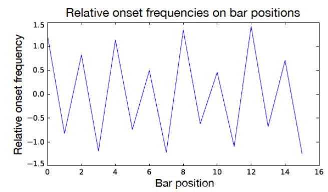

  – Alternate IP/GD and SGS, controlled by Simulated Annealing
  – But not exact as, e.g., Markov Constraints [Pachet & Roy, 2011]

Structure Imposition

- Structure (Repetition Structure, Ex: AABA)
  - Self-Similarity Matrix
  - For each Music Slice

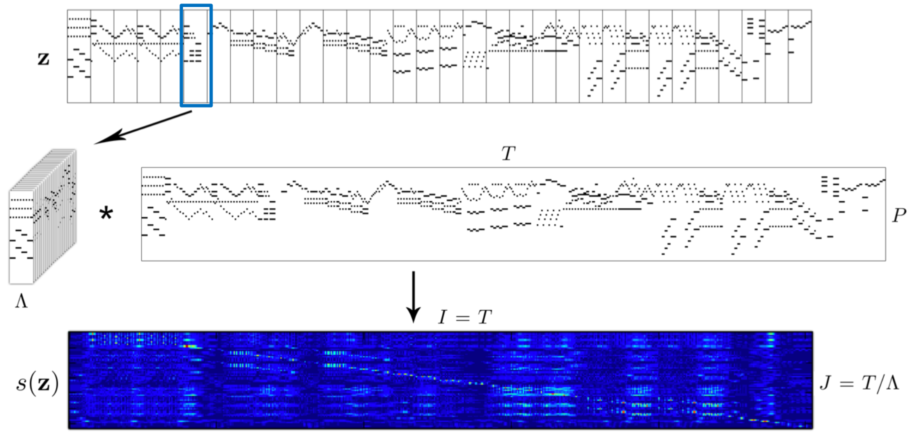

- Tonality, via Similarity of Distribution of Pitch-Classes
  - Key Estimation Vectors over Time

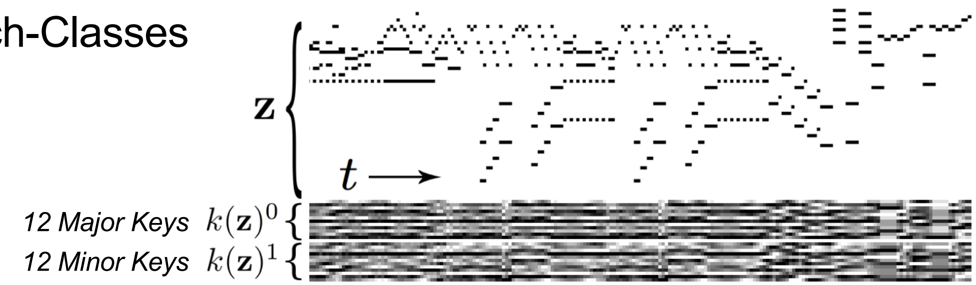

- Meter
  - Duration and Accent Patterns (ex: on 1st and 3rd Beats)
  - Via Relative Occurrence of Note Onsets

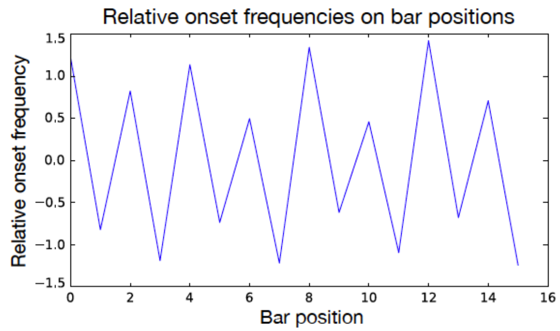

C-RBM [Lattner et al., 2016]

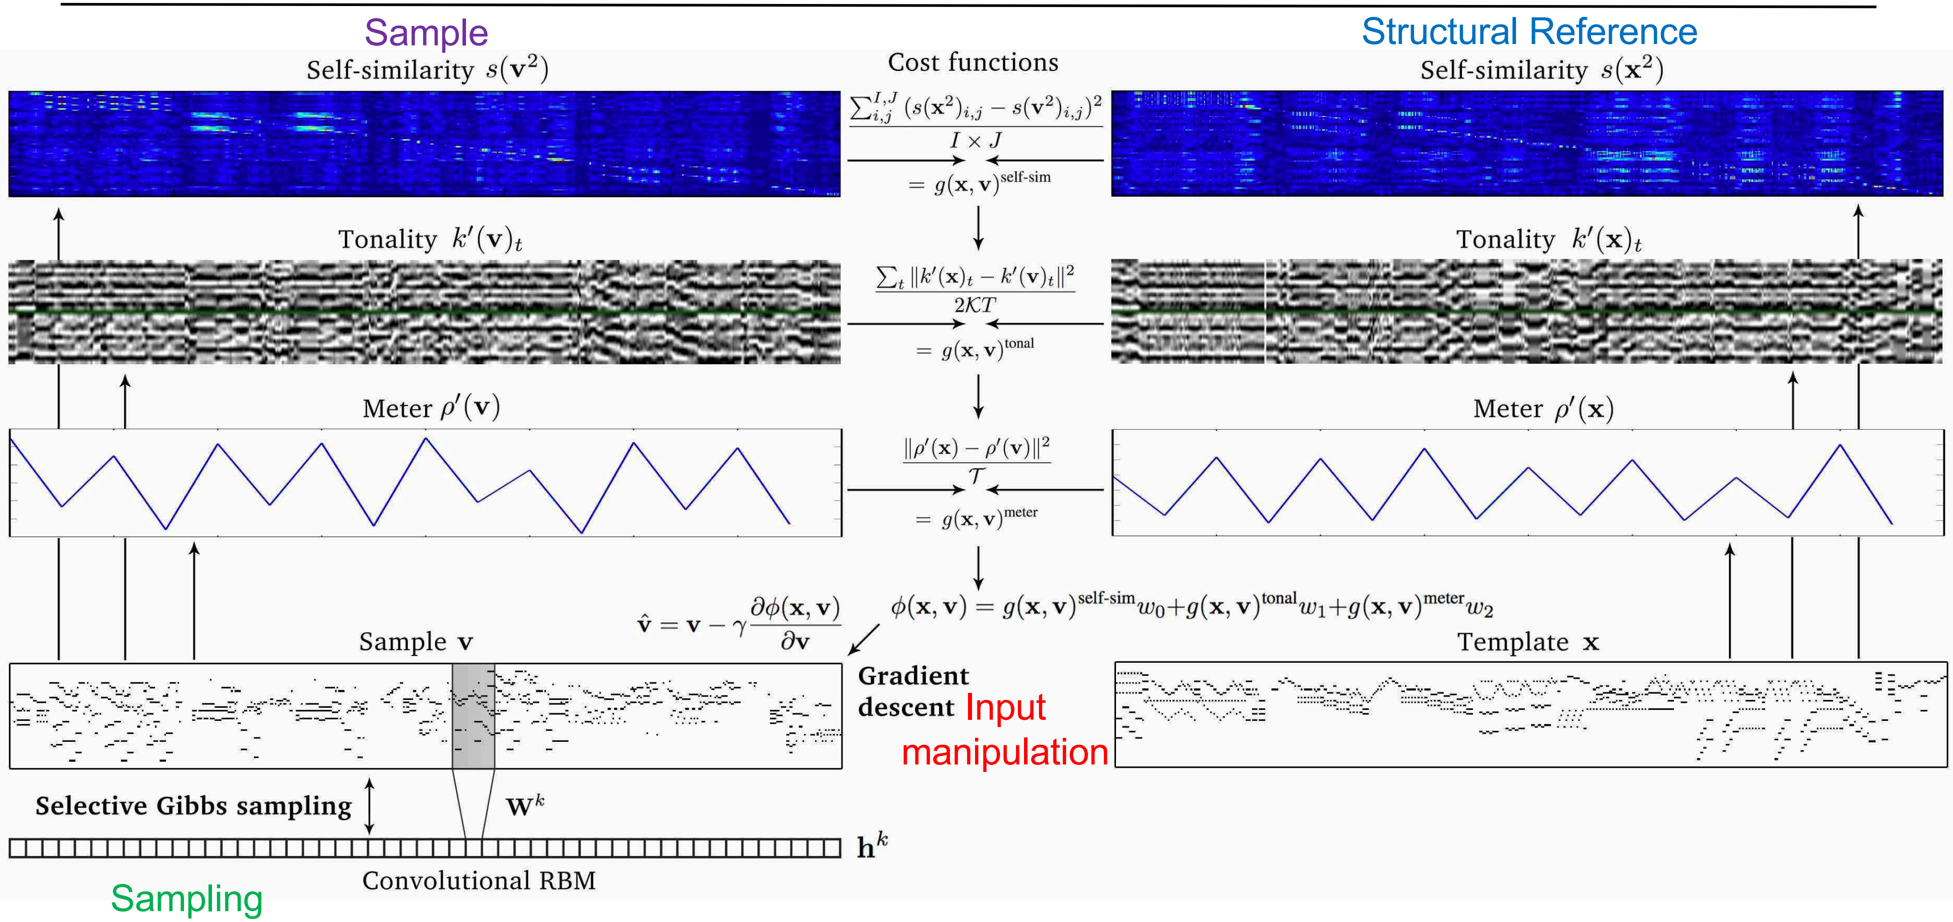

Both Manipulation and Sampling of Input because RBM's "Output" is its Input

https://soundcloud.com/pmgrbm

C-RBM Examples

- RNN-RBM Sample
- Unconstrained Sample
- Template Piece
- Constrained Sample

https://soundcloud.com/pmgrbm

Mozer's Rich Representation Model [Mozer, 1994]

Note/Harmony

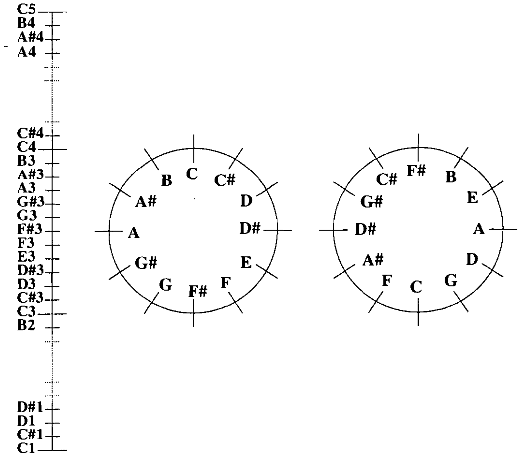

Pitch Height Chroma Circle Circle of Fifths

Pitch	PH	CC						CF					
C1	−9.978	+1	+1	+1	−1	−1	−1	−1	−1	−1	+1	+1	+1
F#1	−7.349	−1	−1	−1	+1	+1	+1	+1	+1	+1	−1	−1	−1
G2	−2.041	−1	−1	−1	−1	+1	+1	−1	−1	−1	−1	+1	+1
C3	0	+1	+1	+1	−1	−1	−1	−1	−1	−1	+1	+1	+1
D#3	1.225	+1	+1	+1	+1	+1	+1	+1	+1	+1	+1	+1	+1
E3	1.633	−1	+1	+1	+1	+1	+1	+1	−1	−1	−1	−1	−1
A4	8.573	−1	−1	−1	−1	−1	−1	−1	−1	−1	−1	−1	−1
C5	9.798	+1	+1	+1	−1	−1	−1	−1	−1	−1	+1	+1	+1
Rest	0	+1	−1	+1	−1	+1	−1	+1	−1	+1	−1	+1	−1

[Mozer, 2004]

Duration/Rhythm

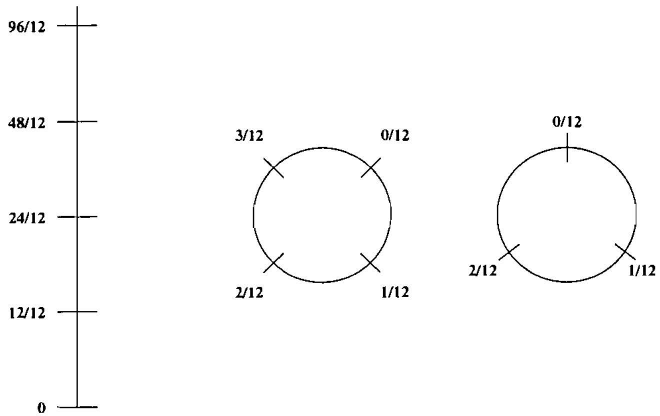

Duration Height
log(duration)

1/3 Beat Circle
mod(duration, 1/3)

1/4 Beat Circle
mod(duration, 1/4)

The Old Emperor New Clothes (Deep Networks/Learning)

The Old Emperor New Clothes

- Multiple Hidden Layers Neural Network

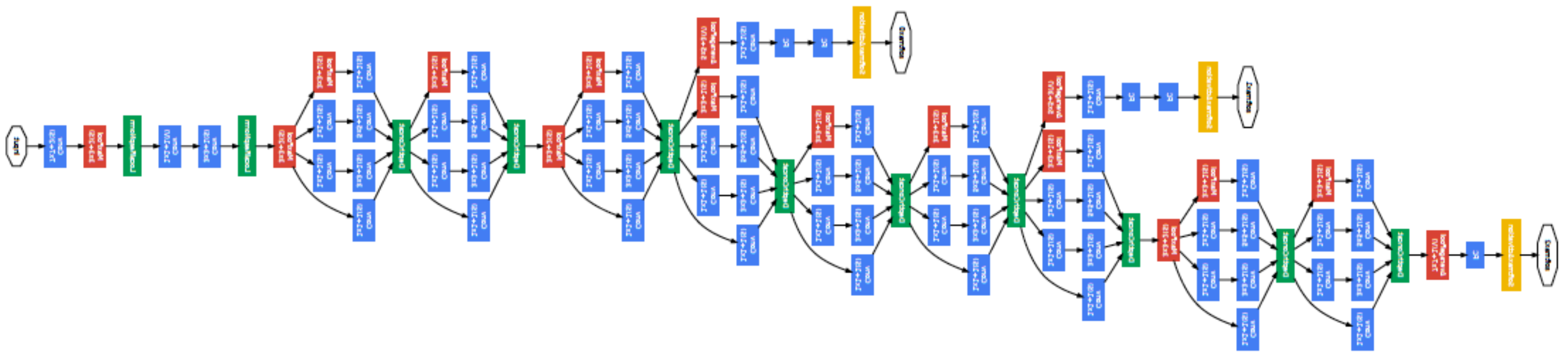

- Platforms
- Tecnical Advances
  - Pre-Training, Batch Normalization, Residual Learning…
- Fast CPUs
  - GPUs

- Large Memory
- Available Data

Power Increase

- Brute Force

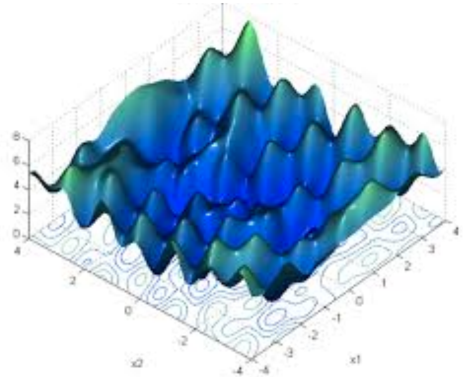

Loss Minimization

- Hypervitamined Brute Force

GPUs

TensorFlow PyTorch

But Not Only…

- Deep Architecture

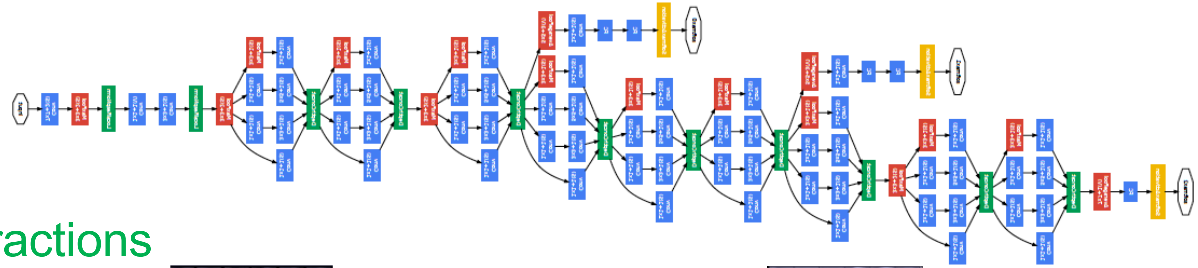

  - Multiple Levels of Abstractions

  - End-to-End Architecture

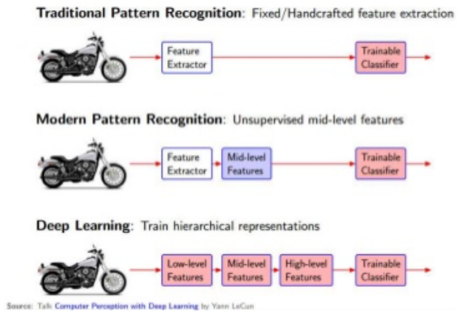

- New Architectures

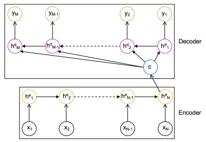

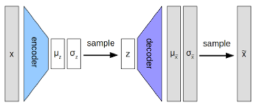

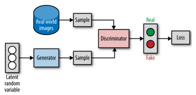

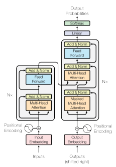

Why Deep ?

- More Complex Models
- Learns better Complex Functions
- Hierarchical Features/Abstractions
- No Need for Handcrafted Features
  - (Automatically Extracted)

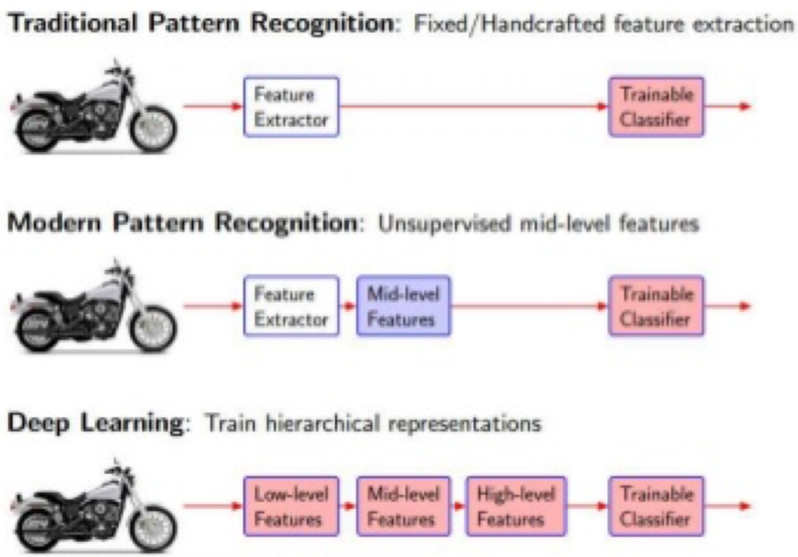

Source: Talk Computer Perception with Deep Learning by Yann LeCun

End-to-End Architecture

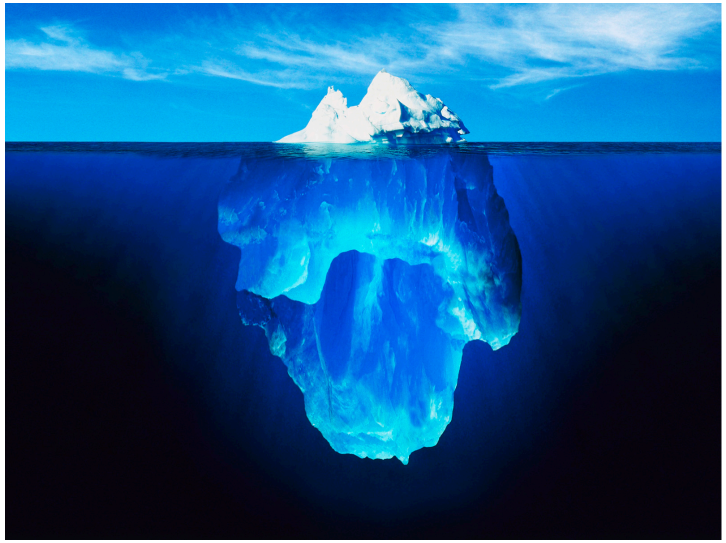

Distributed Representations

The Groundbreaking Start of Deep Learning

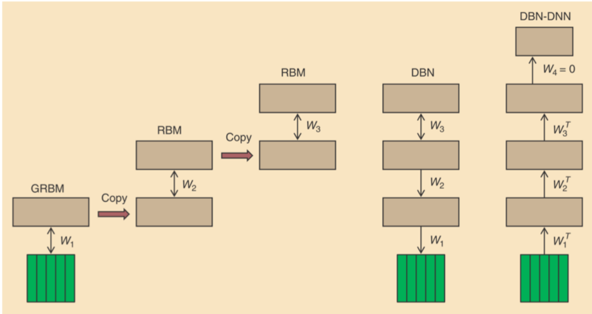

Pre-Training [Hinton et al. 2006]
Layer-Wise Self-Supervised Training/Initialization

Rank	Name	Error rate	Description
1	**U. Toronto**	0.15315	Deep learning
2	U. Tokyo	0.26172	Hand-crafted features and learning models. Bottleneck.
3	U. Oxford	0.26979	
4	Xerox/INRIA	0.27058	

ImageNet 2012 Image Recognition Challenge Breakthrough

WaveNet Audio End-to-End Generation [van den Oord et al., 2017]

- Van den Oord, A., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., Kalchbrenner, N., Senior, A., Kavukcuoglu, K., WaveNet: A Generative Model for Raw Audio, arXiv:1609.03499, December 2016.
- Waveform

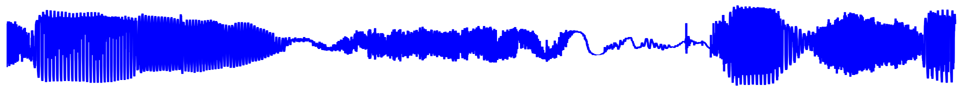

- End to end architecture

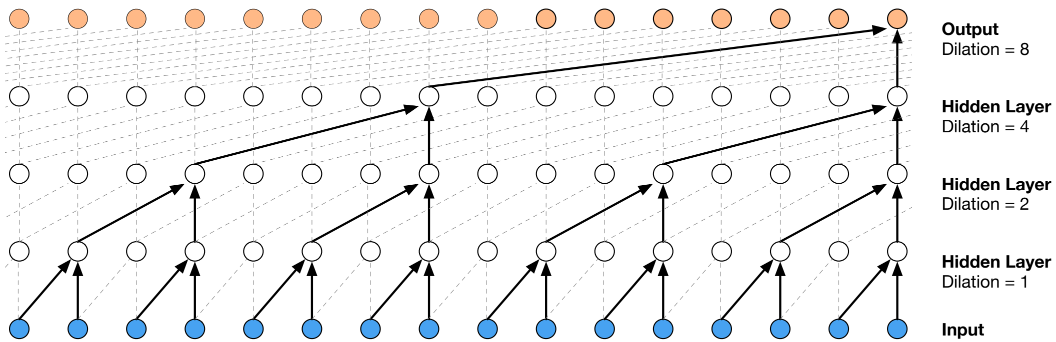

[van den Oord, 2016]

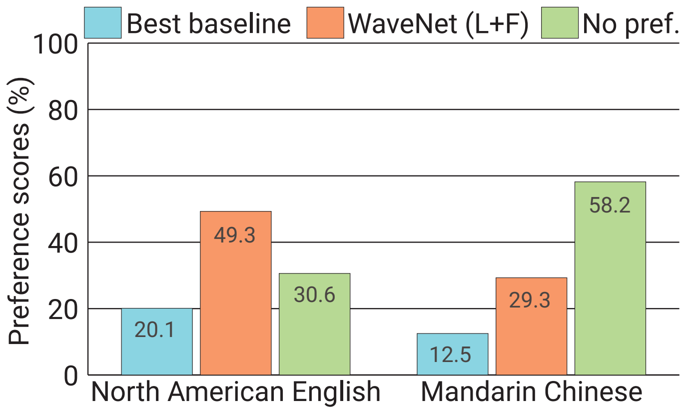

New Architectures

- New Architectures and Mechanisms
- RNN Encoder Decoder
- Variational Autoencoders
- Generative Adversarial Networks
- Transformer
- Attention Mechanism
- …

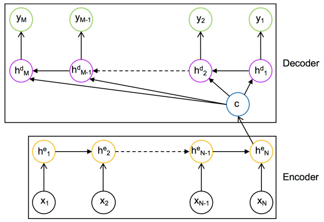

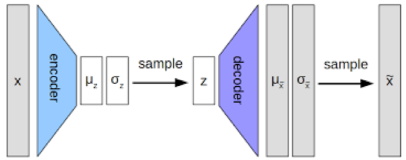

[Bechberger, 2018]

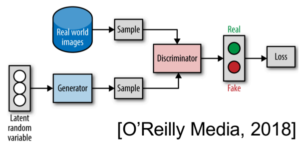

[O'Reilly Media, 2018]

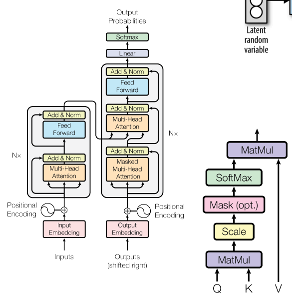

[Vaswani et al., 2017]

Phylogenetics

Deep Learning Phylogenetics

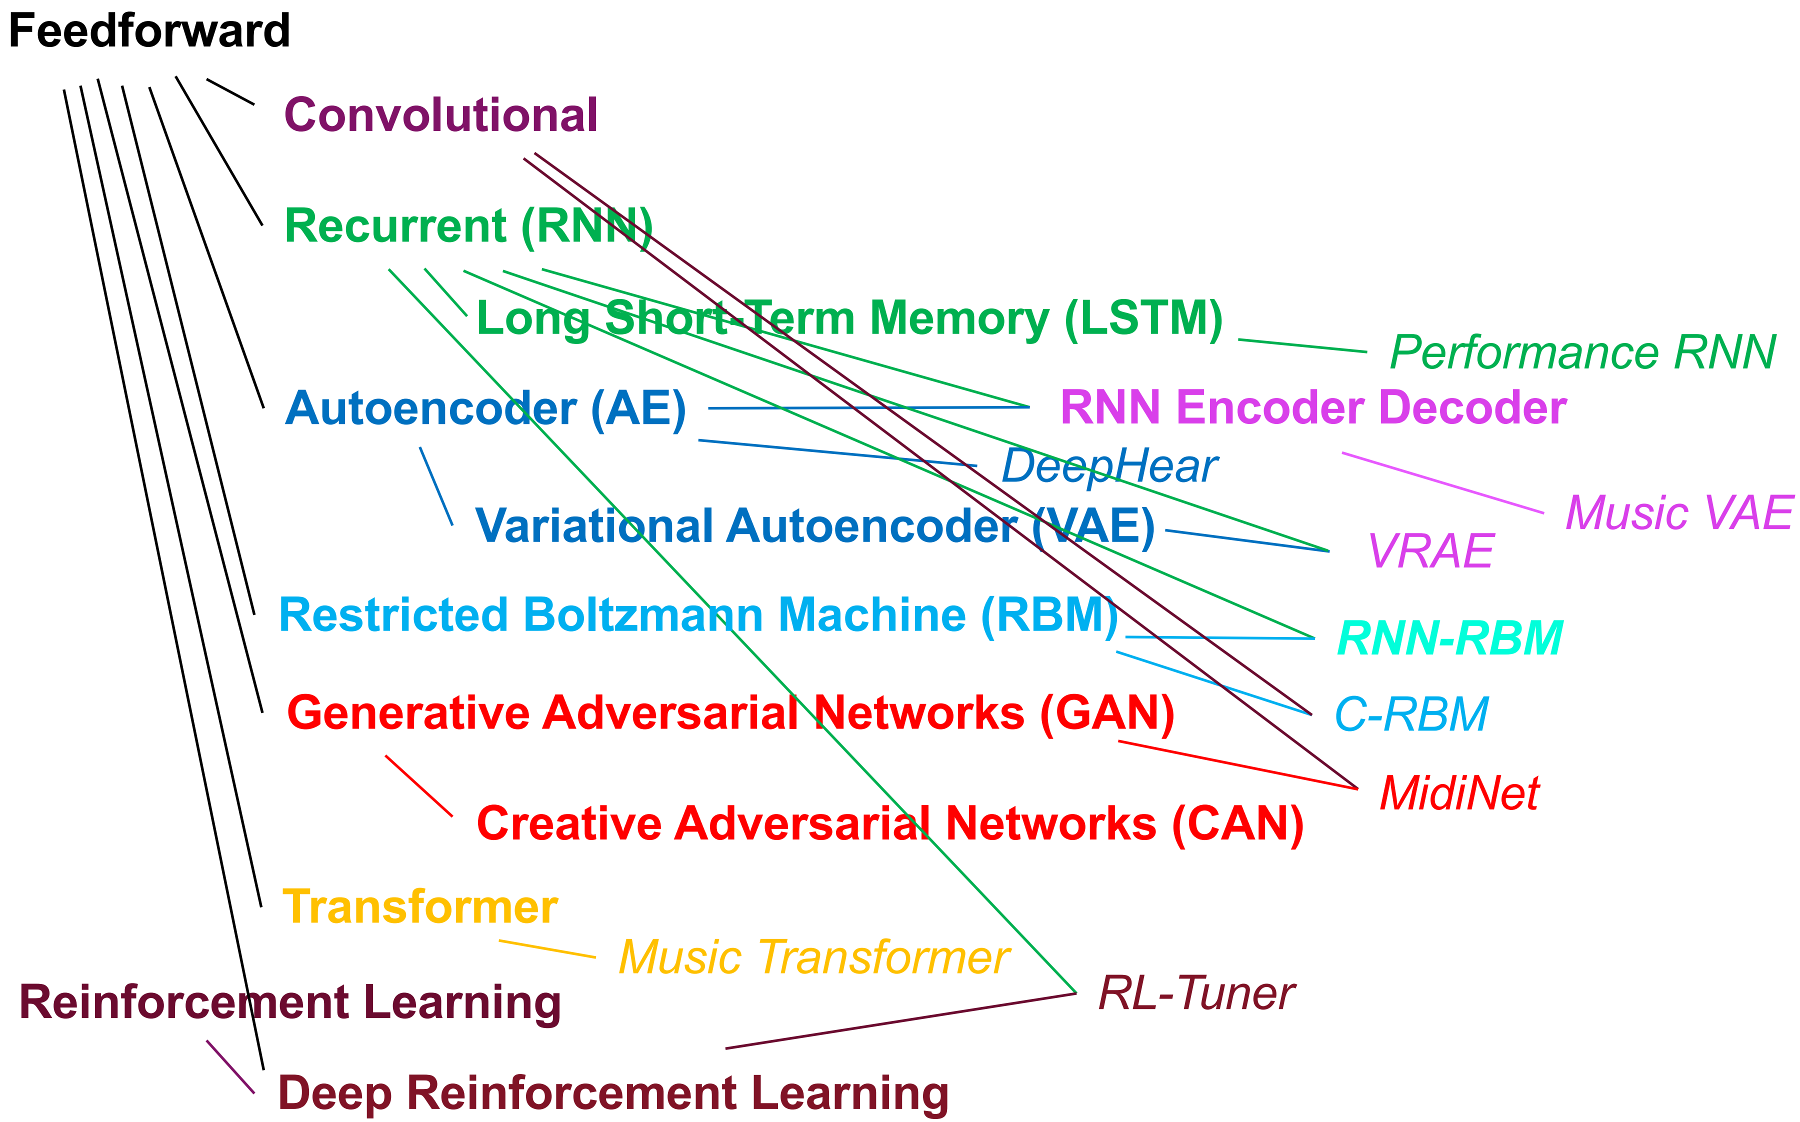

Deep Learning Phylogenetics

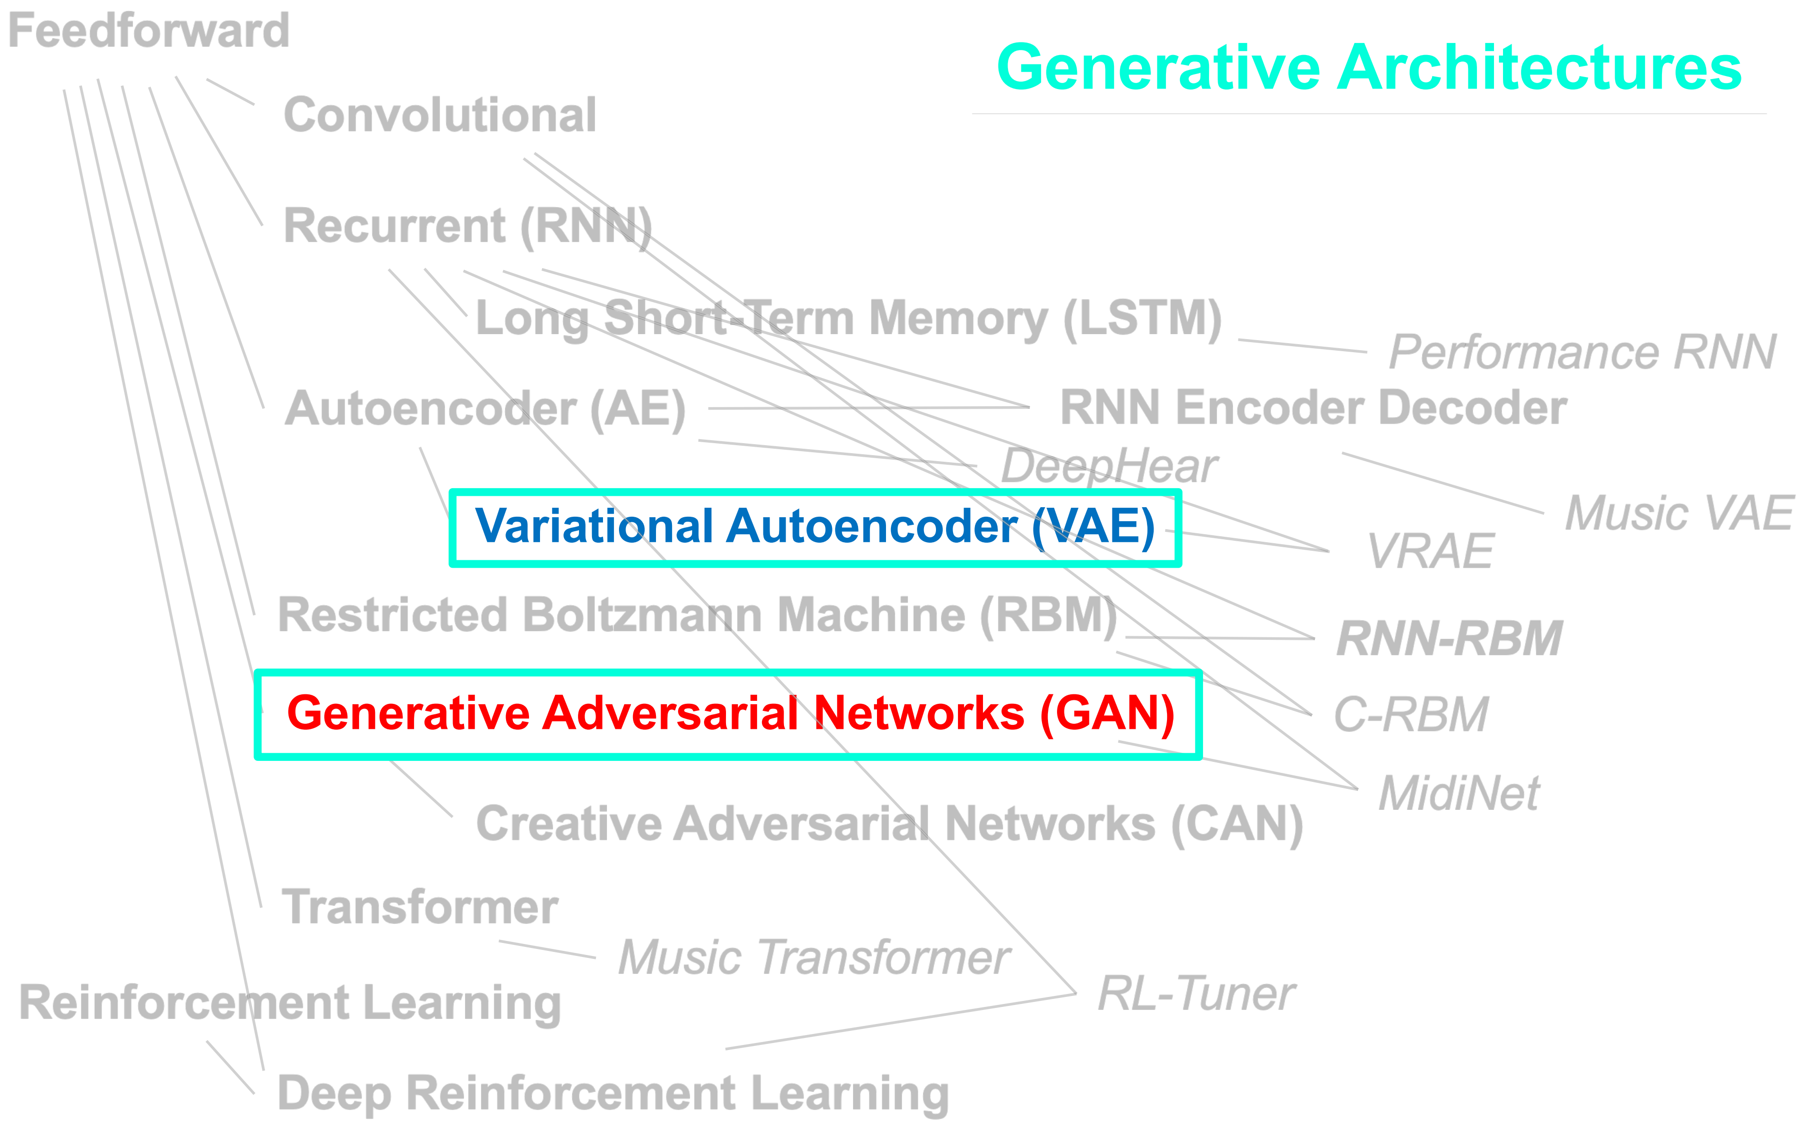

Artificial Intelligence and Machine Learning

Machine Learning and Artificial Intelligence

- Backfire (Irony) of History
- In 1960, Minsky and Papert founded AI (Artificial Intelligence) based on Concepts, Symbols, Logic, Reasoning…, Against Cybernetics (Feedback) and Connexionism (Neural Networks)
- In 1969, they "Killed" Connexionism/Neural Networks (Sound Critic of Perceptron)
- In 2006, Start of Deep Learning
- Now, AI is synonym of Deep Learning
- When Actually, Neural Networks are somehow based on Statistical (Correlation) Brute Force

Terminology

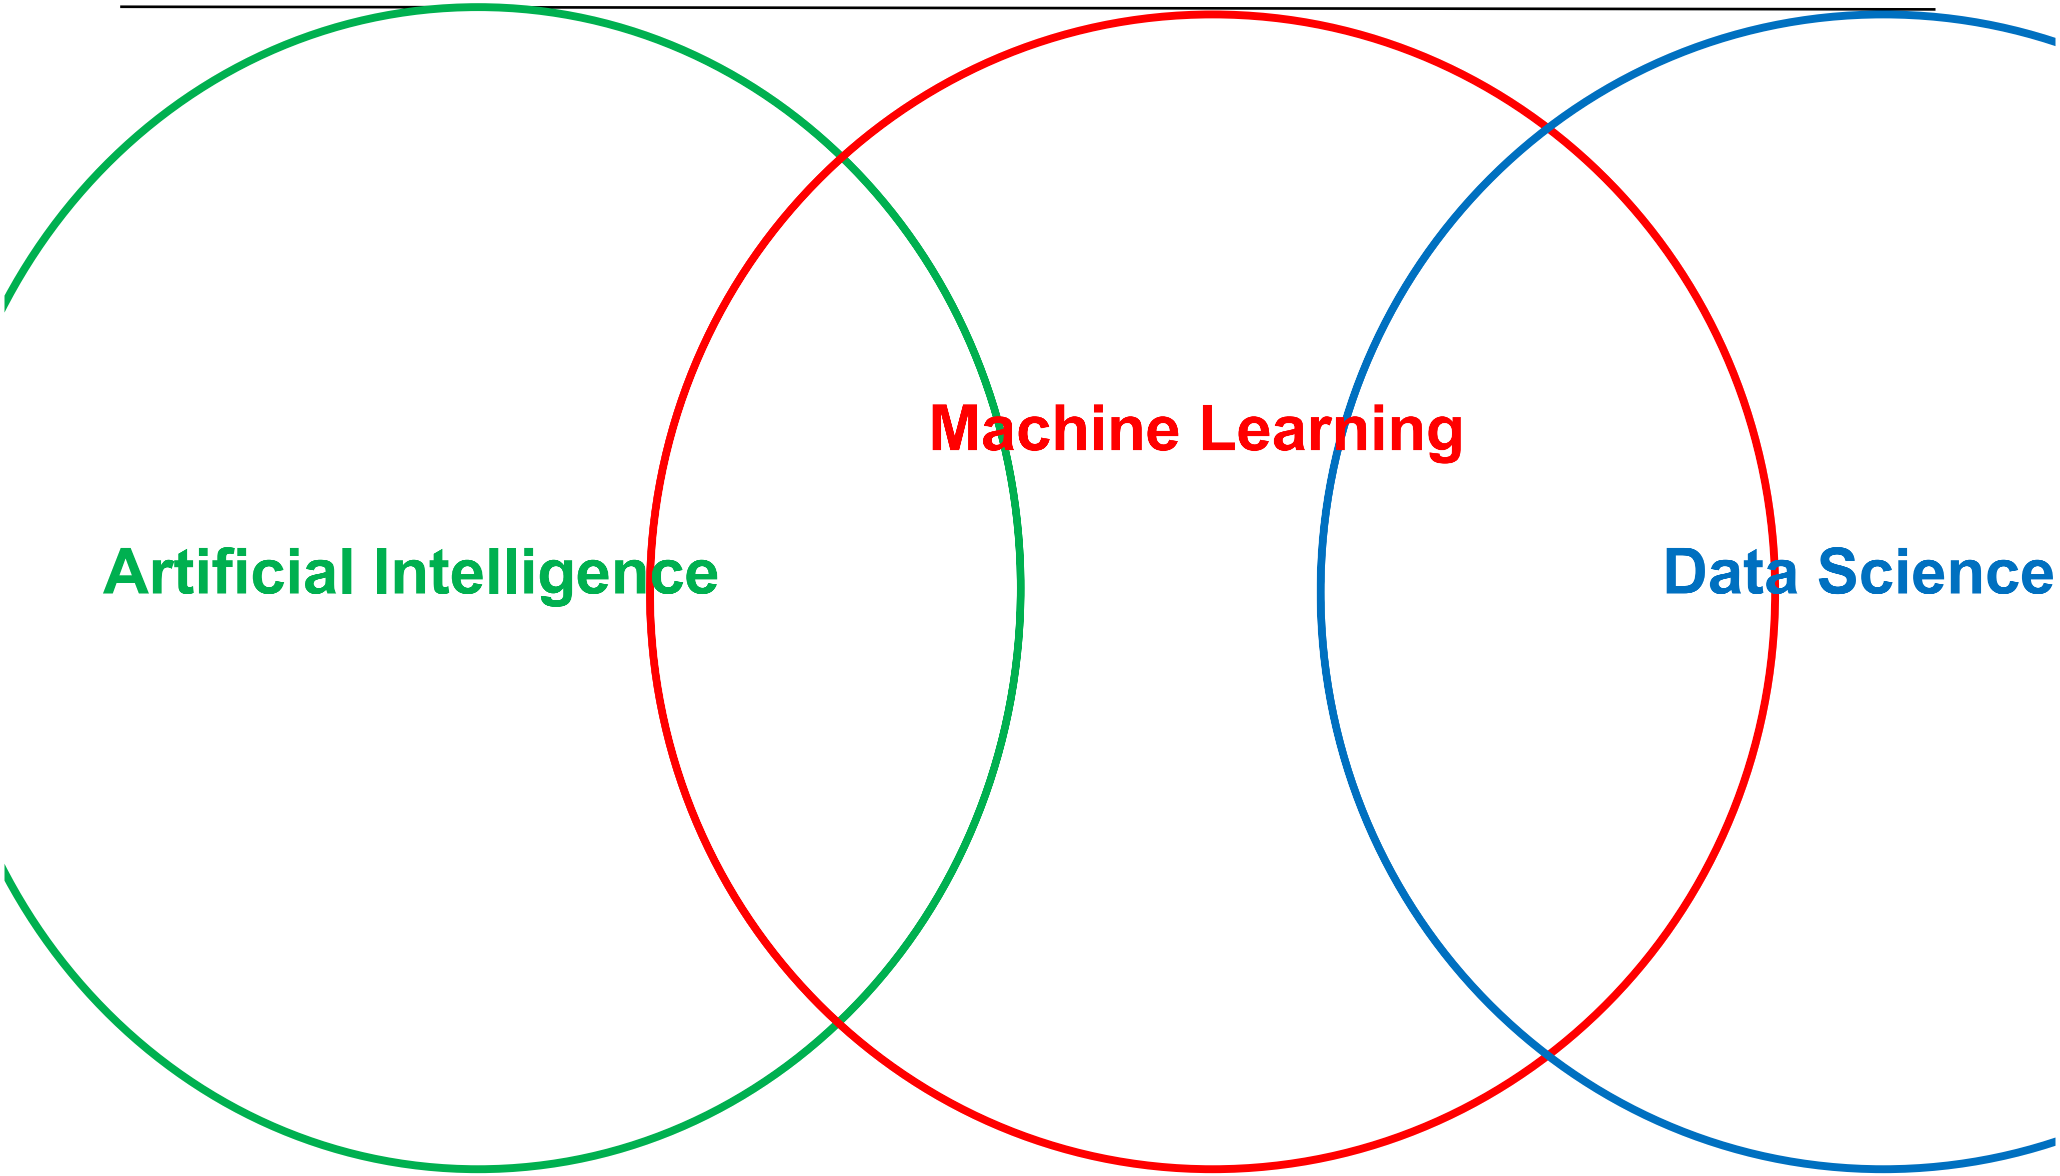

Terminology

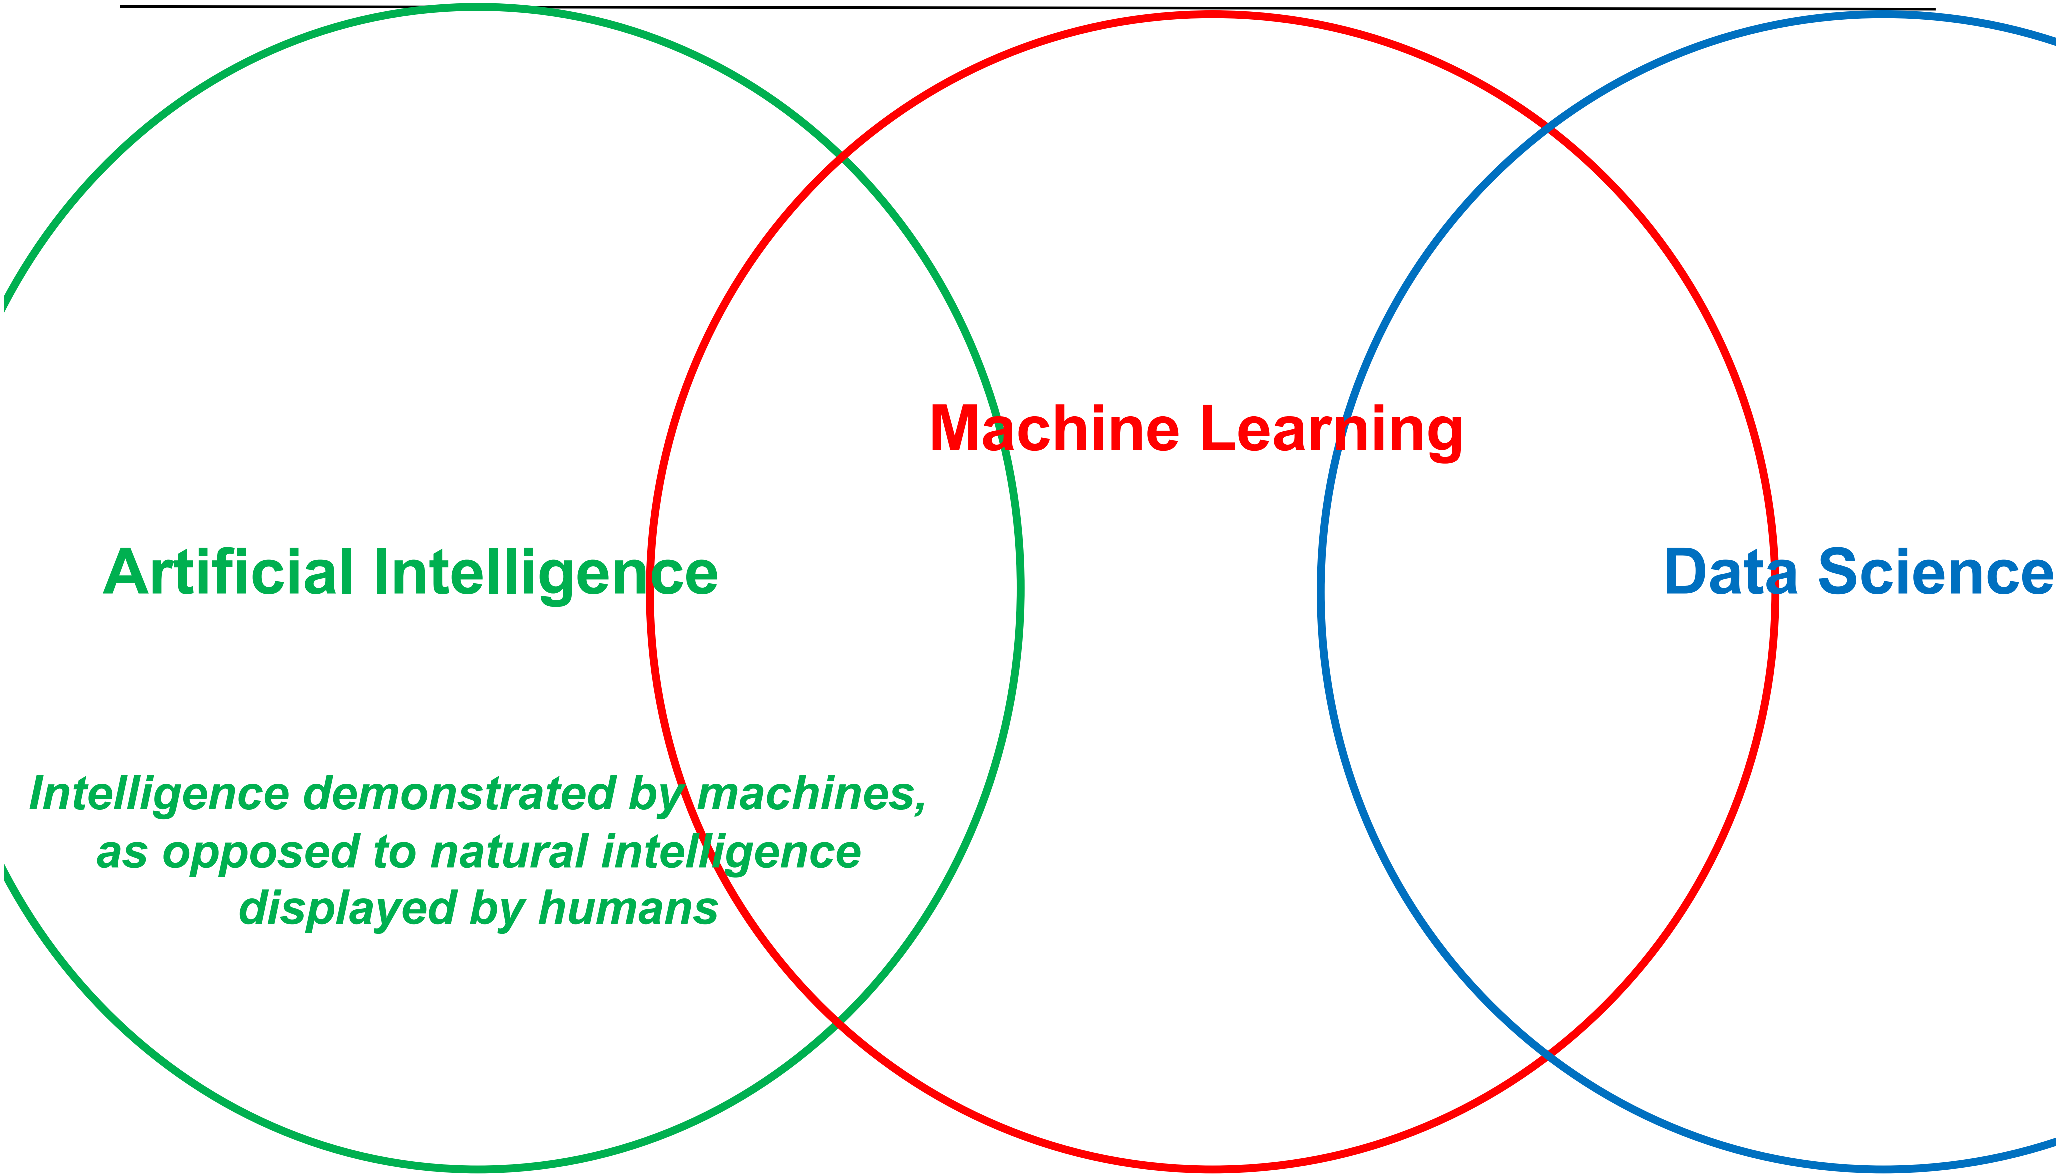

Terminology

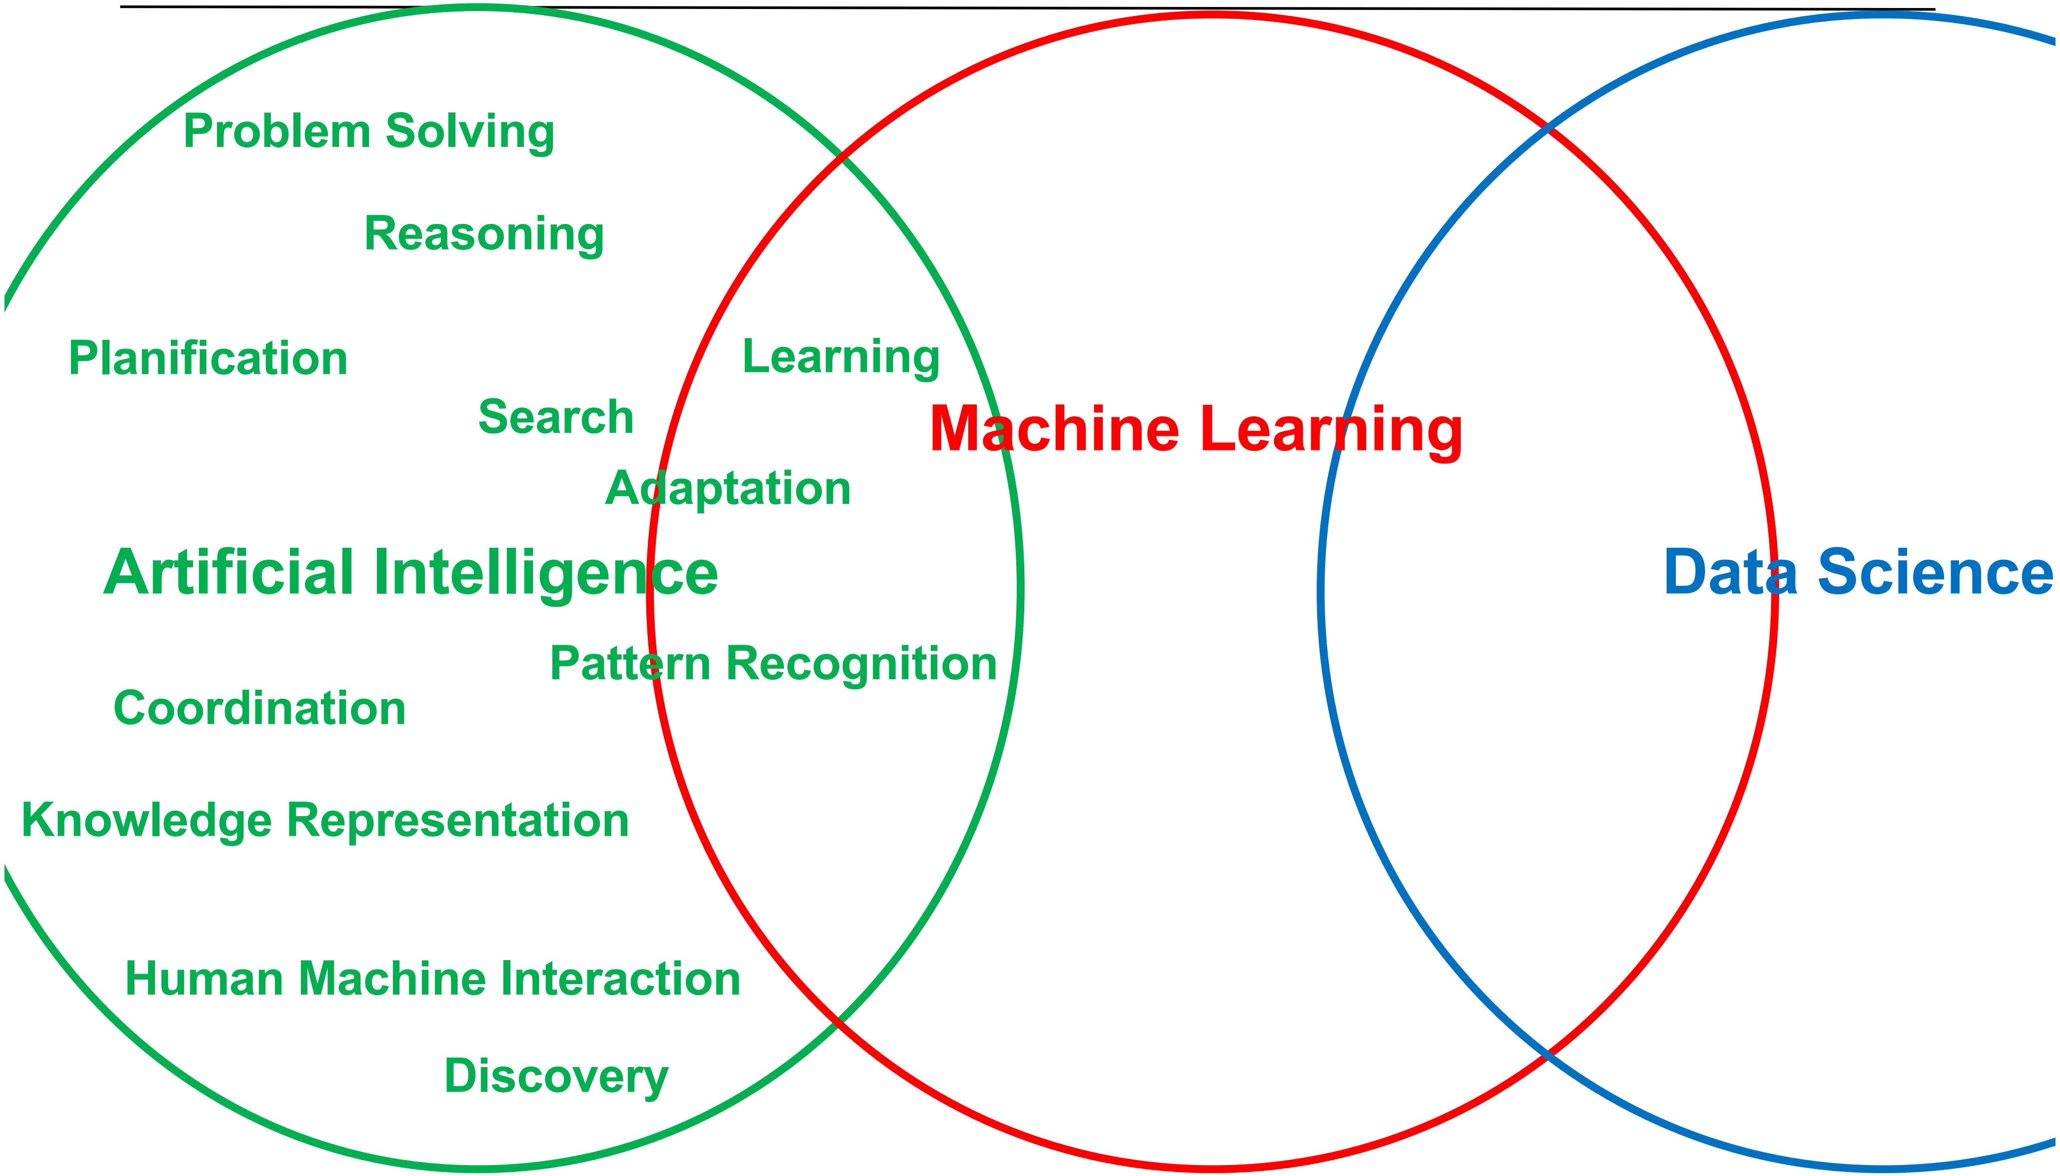

Terminology

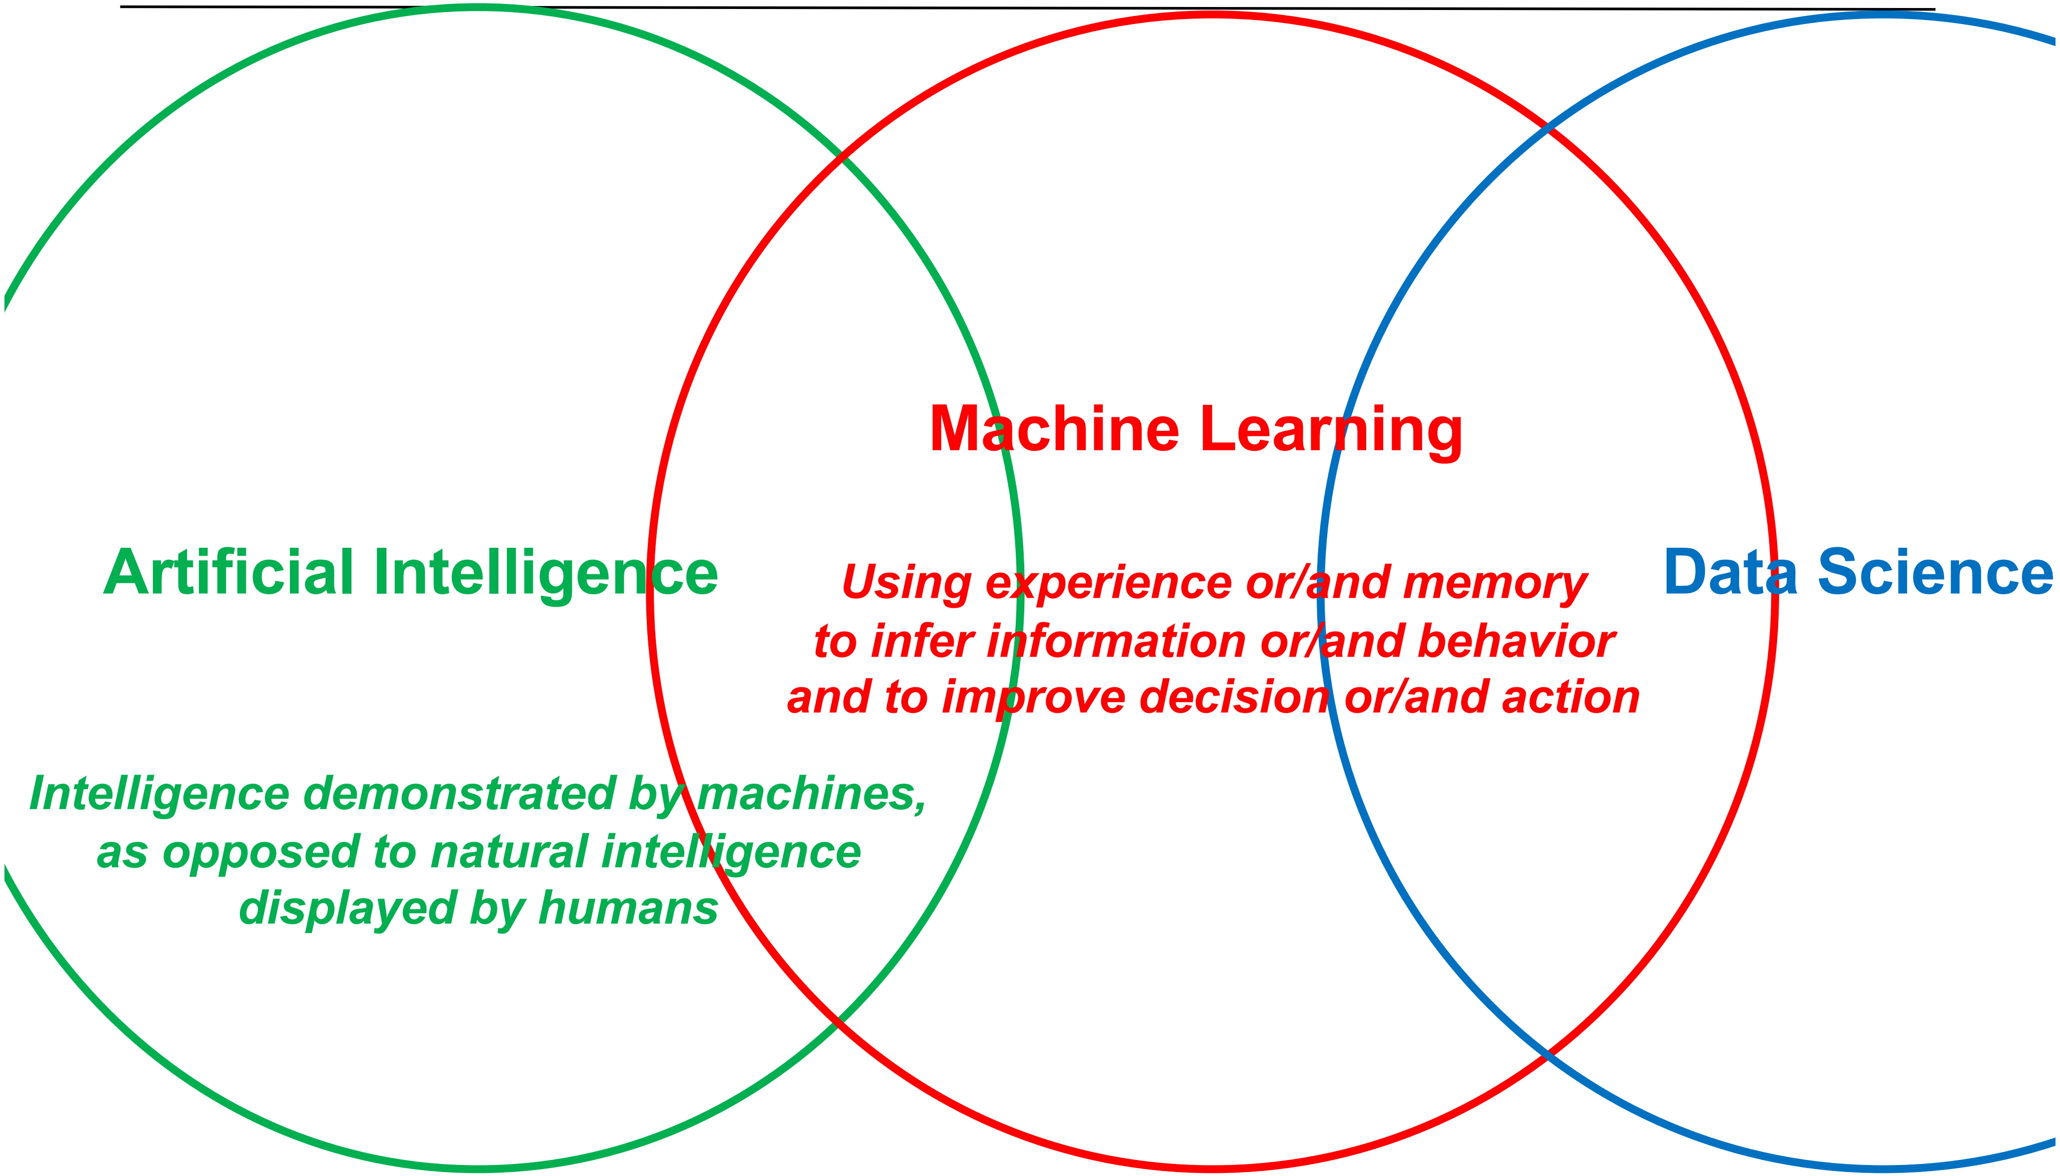

Terminology

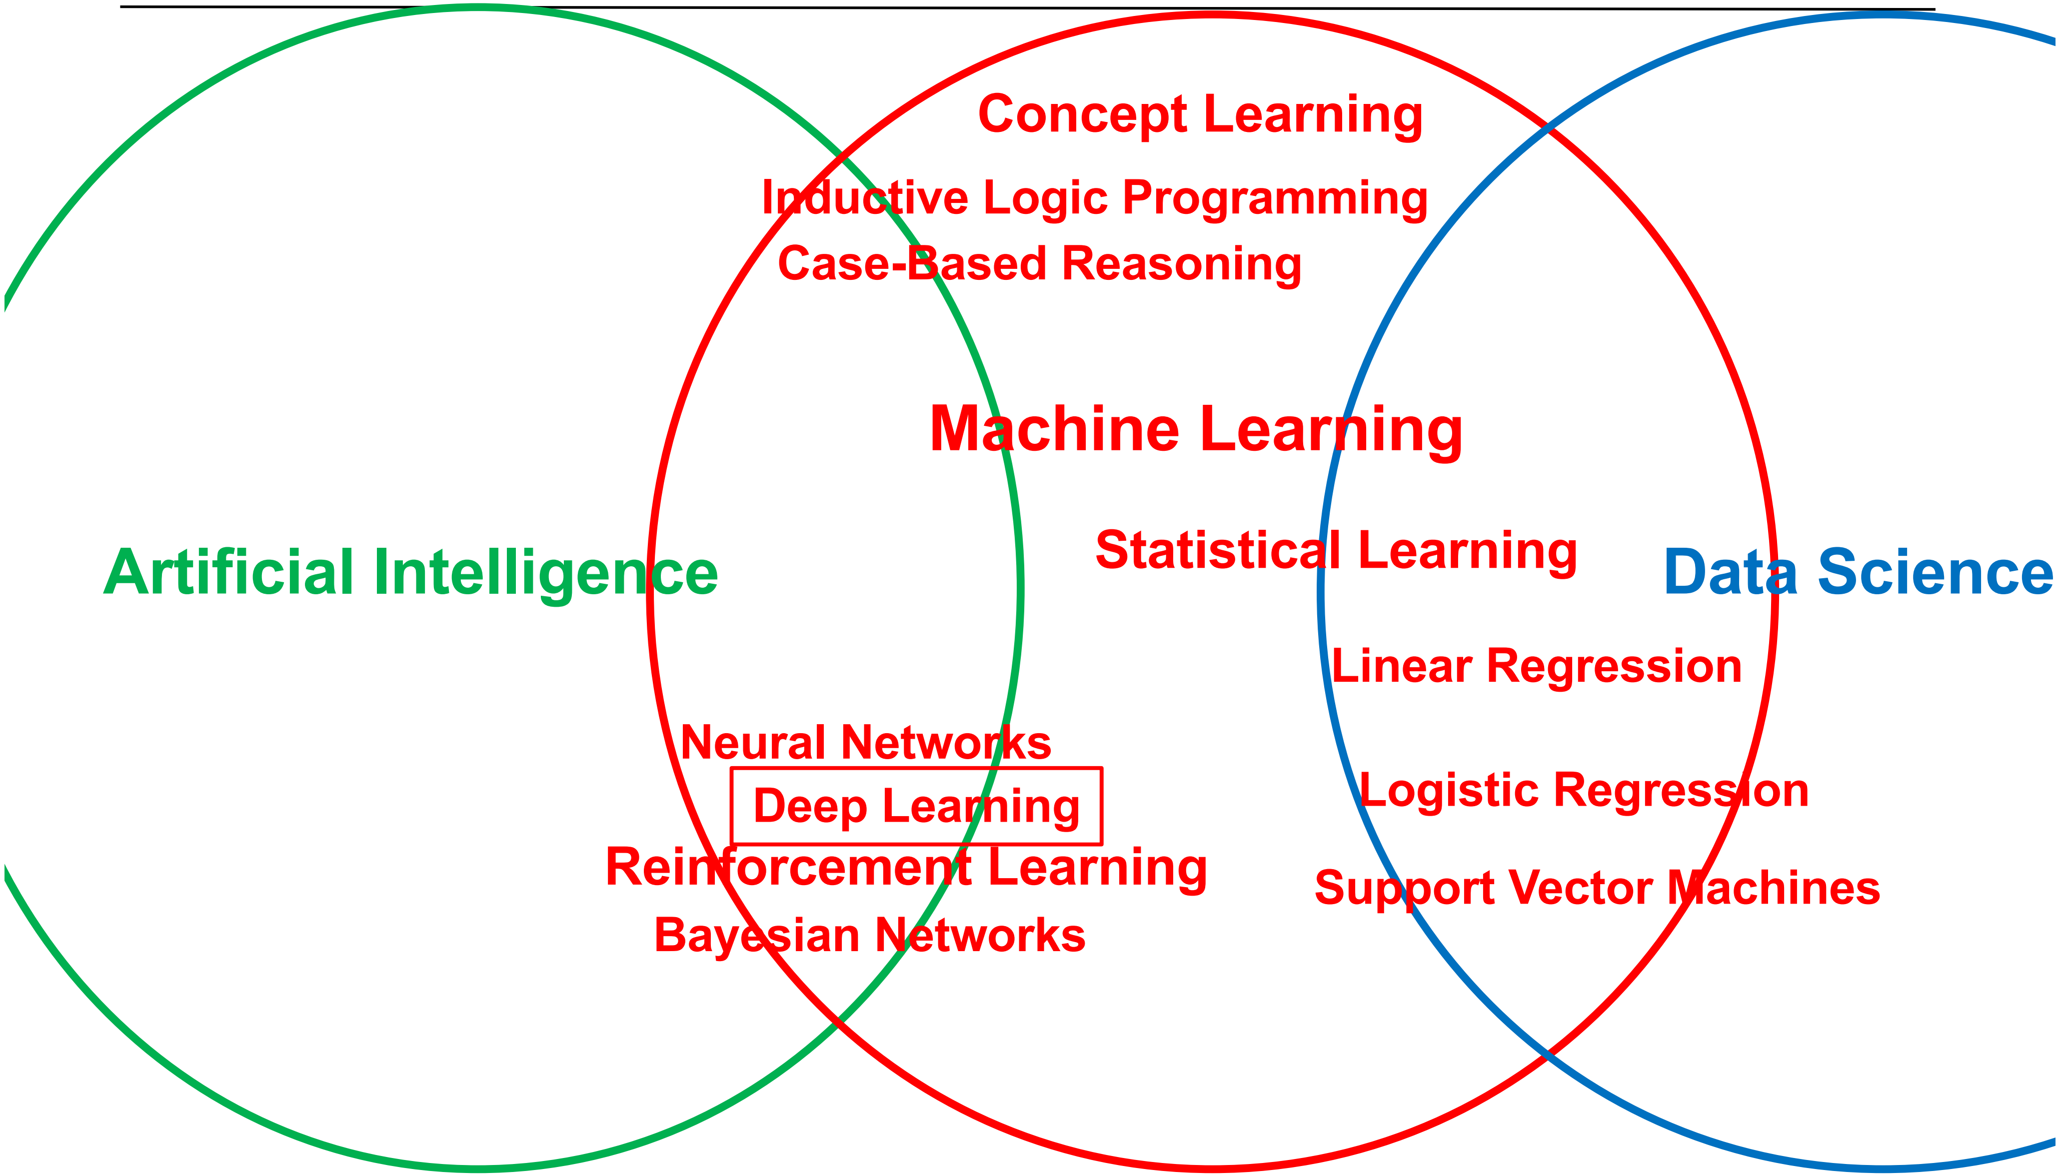

Terminology

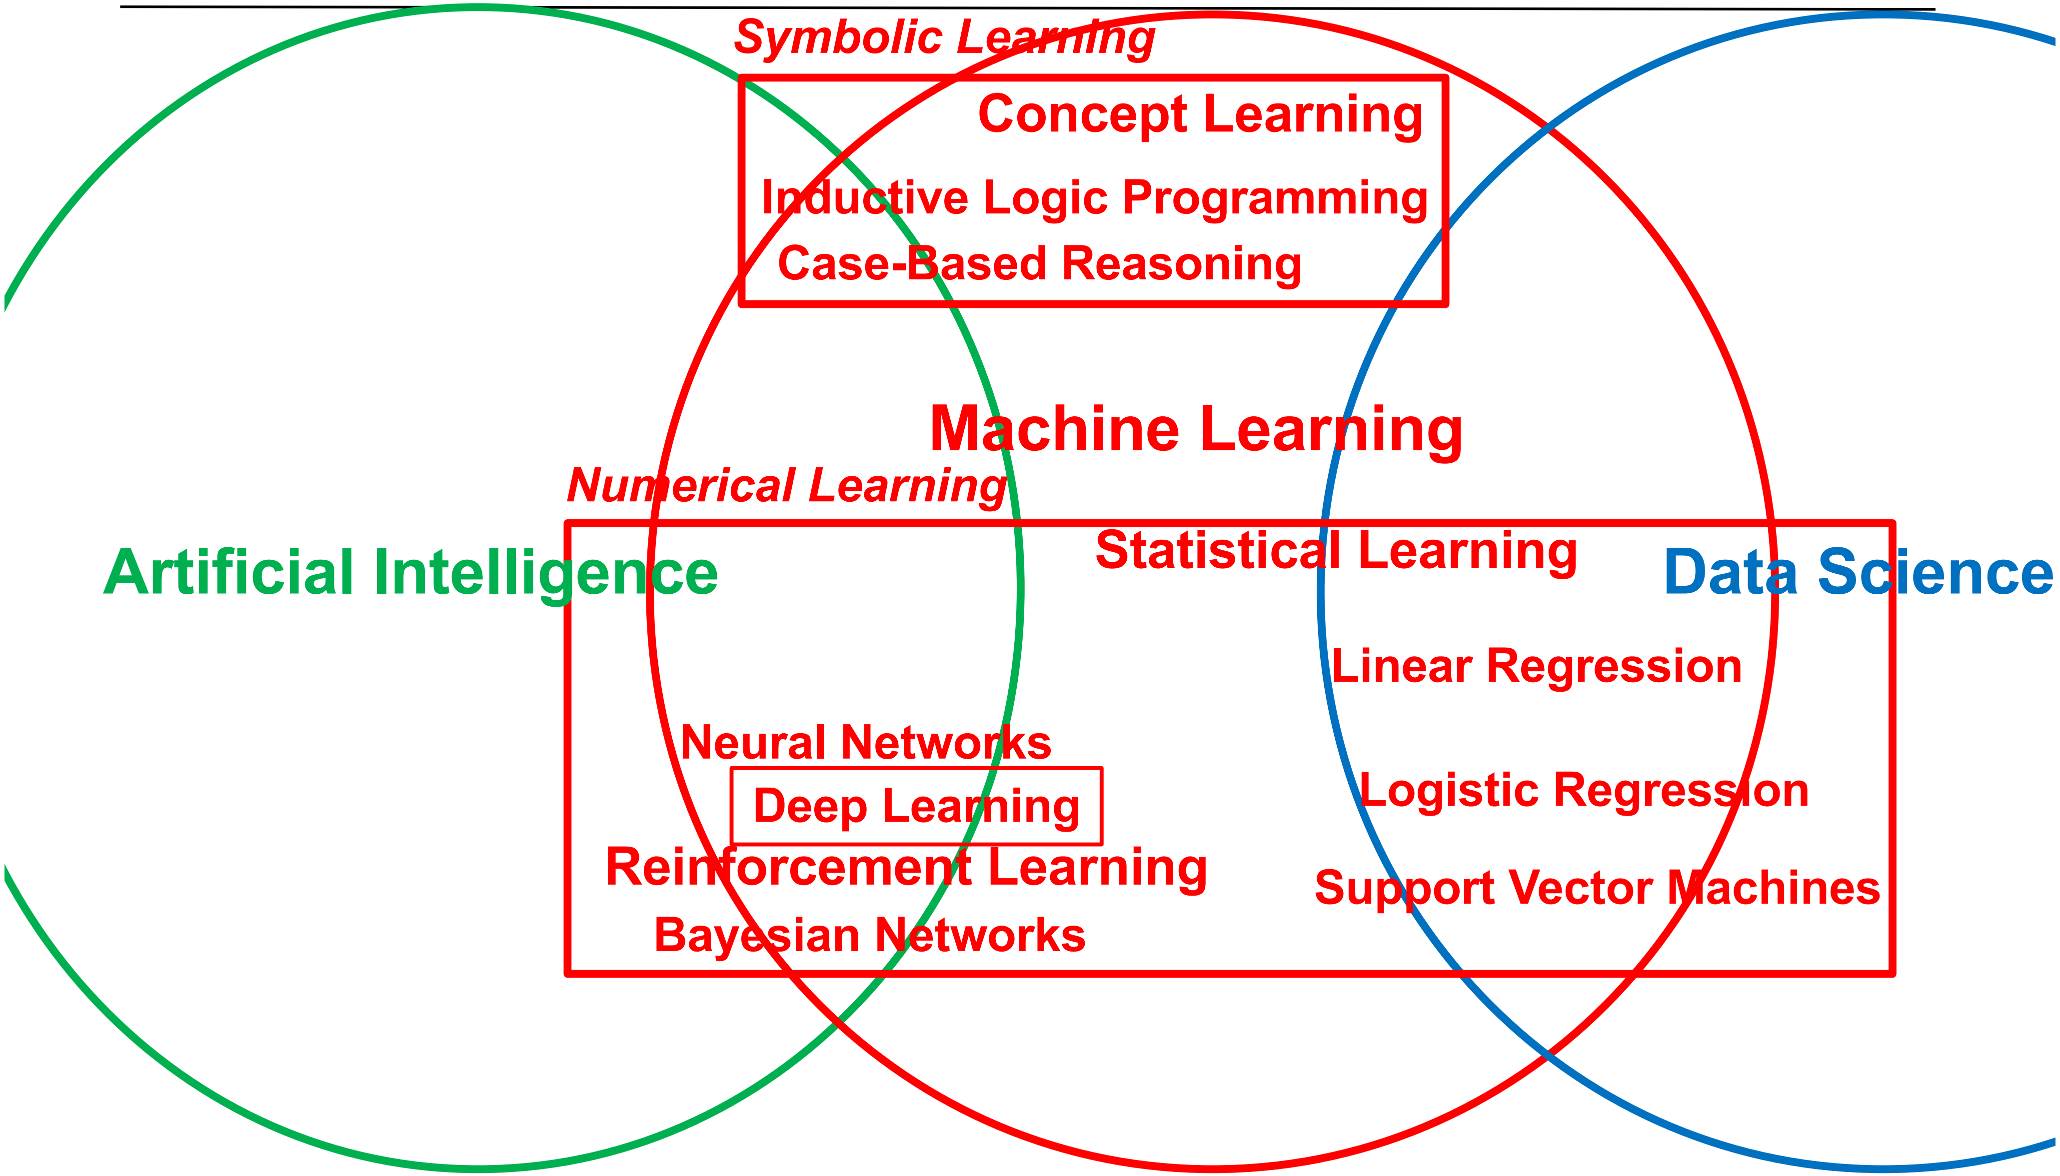

Terminology

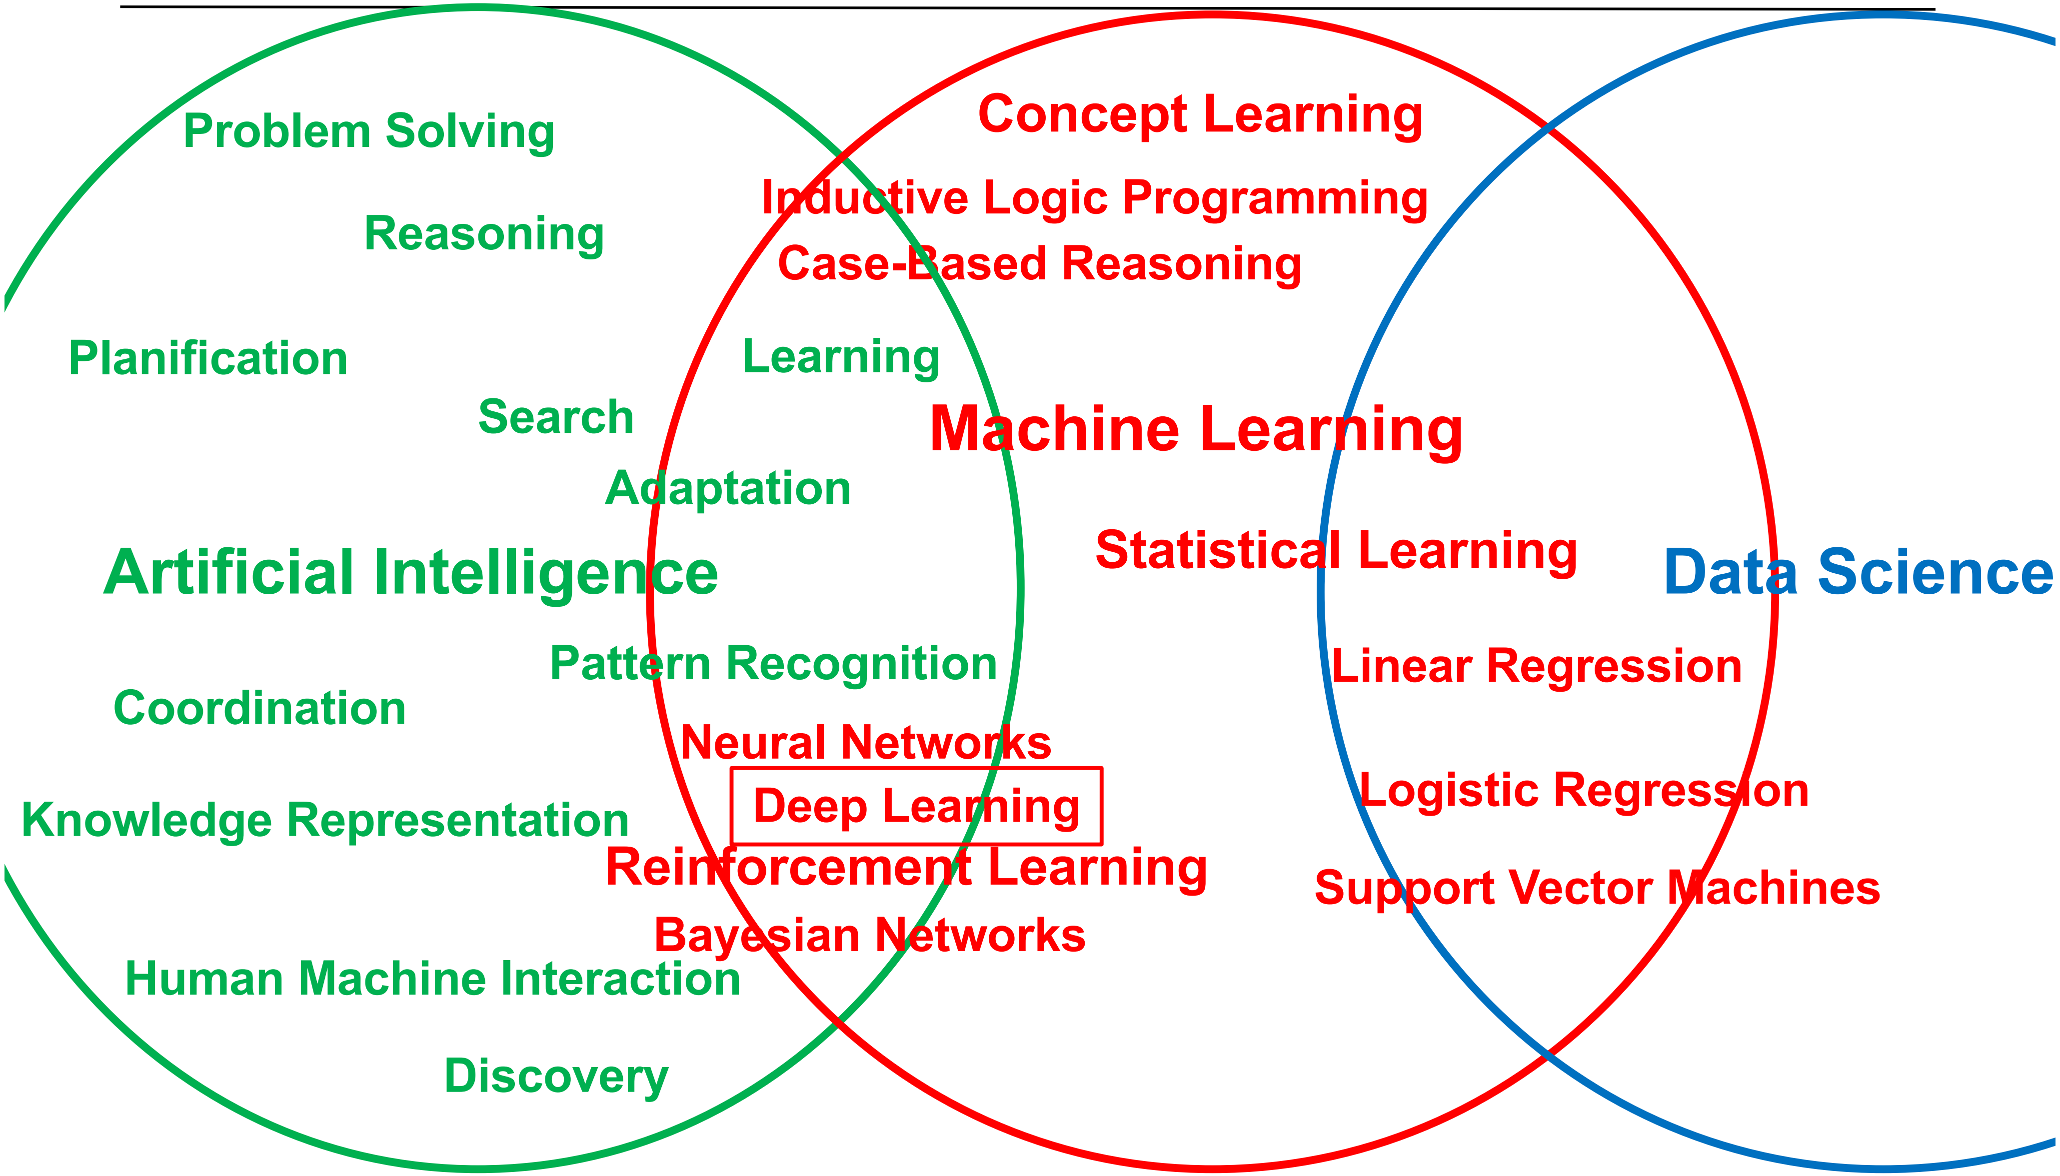

Terminology

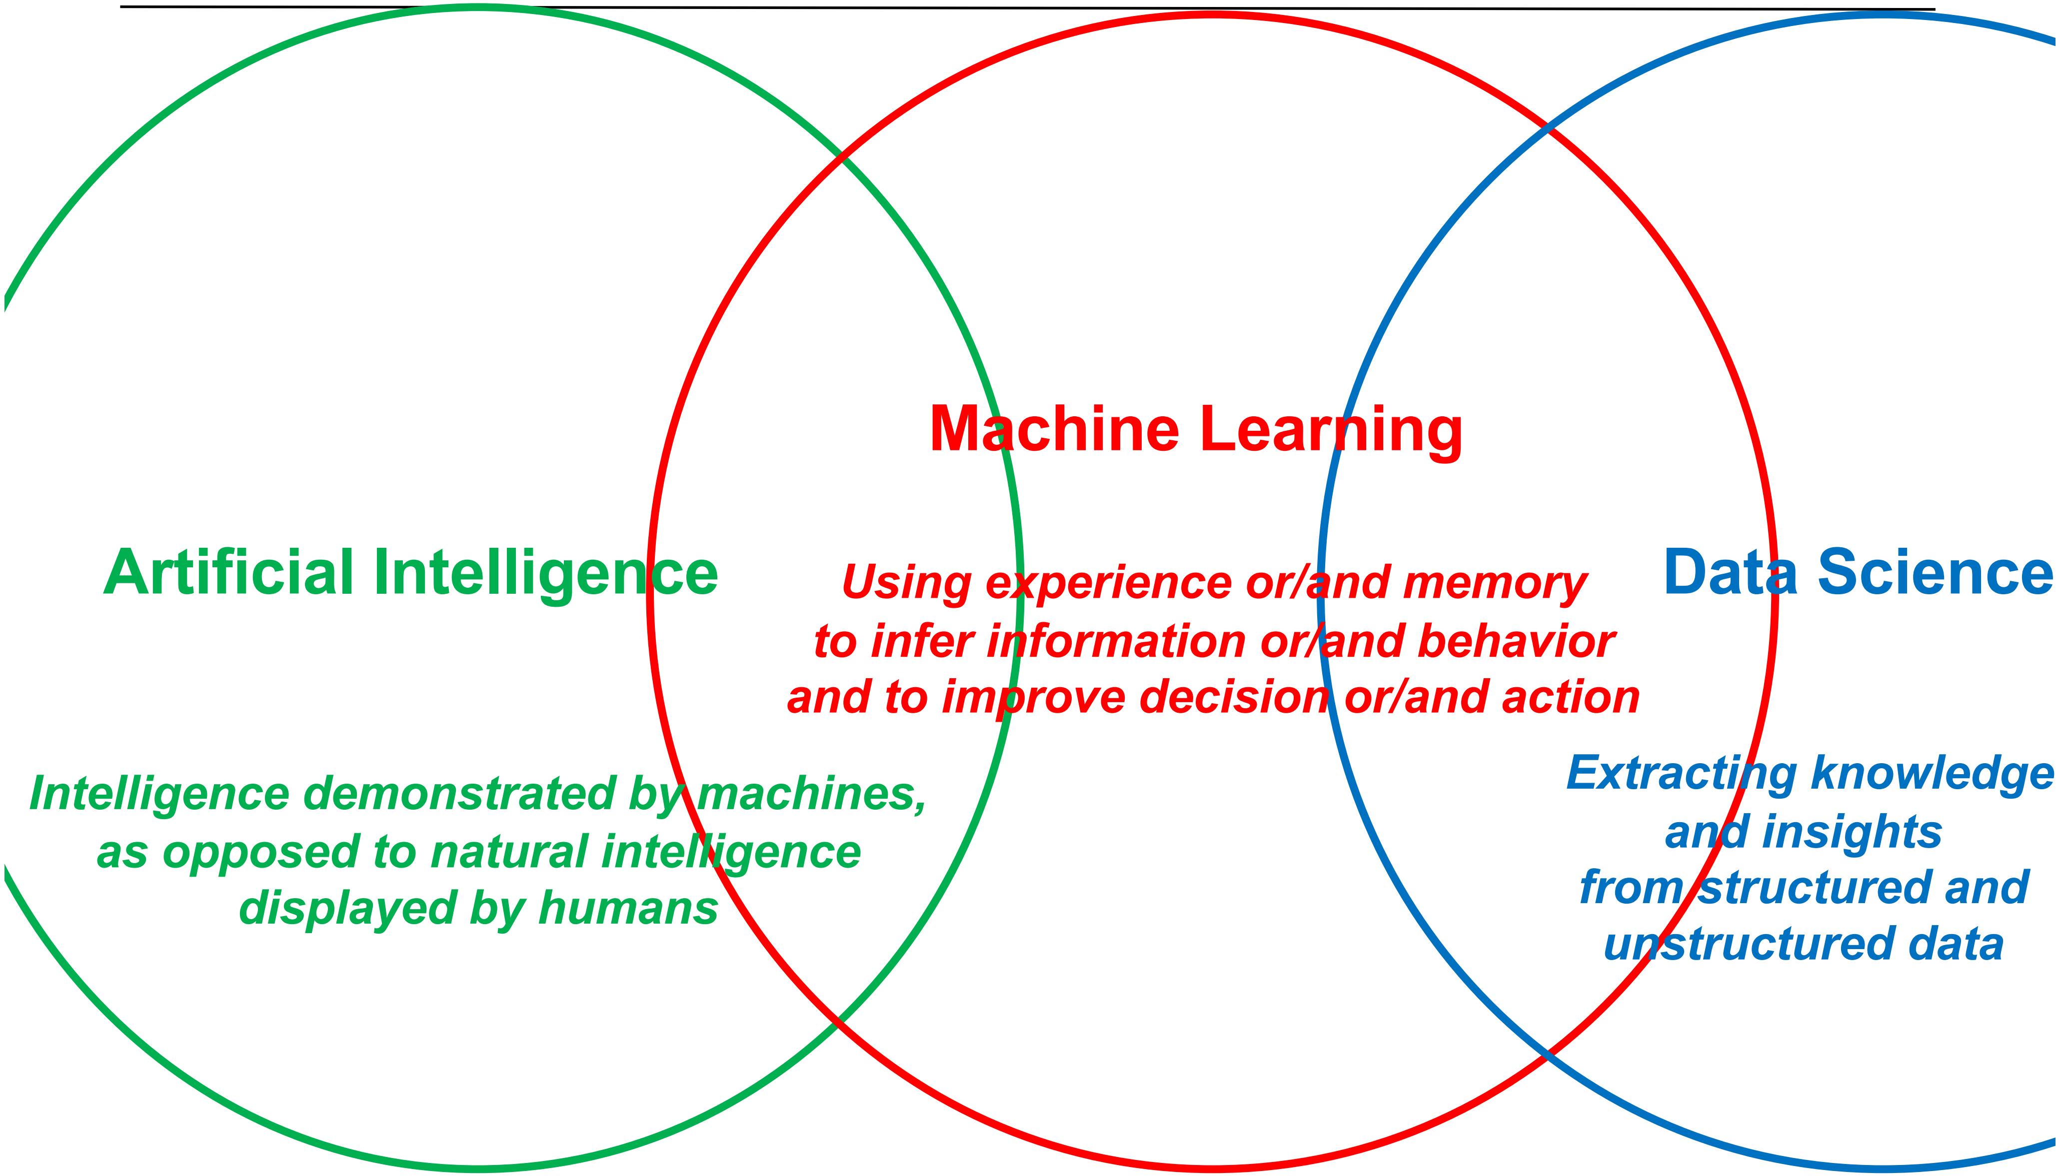

Terminology

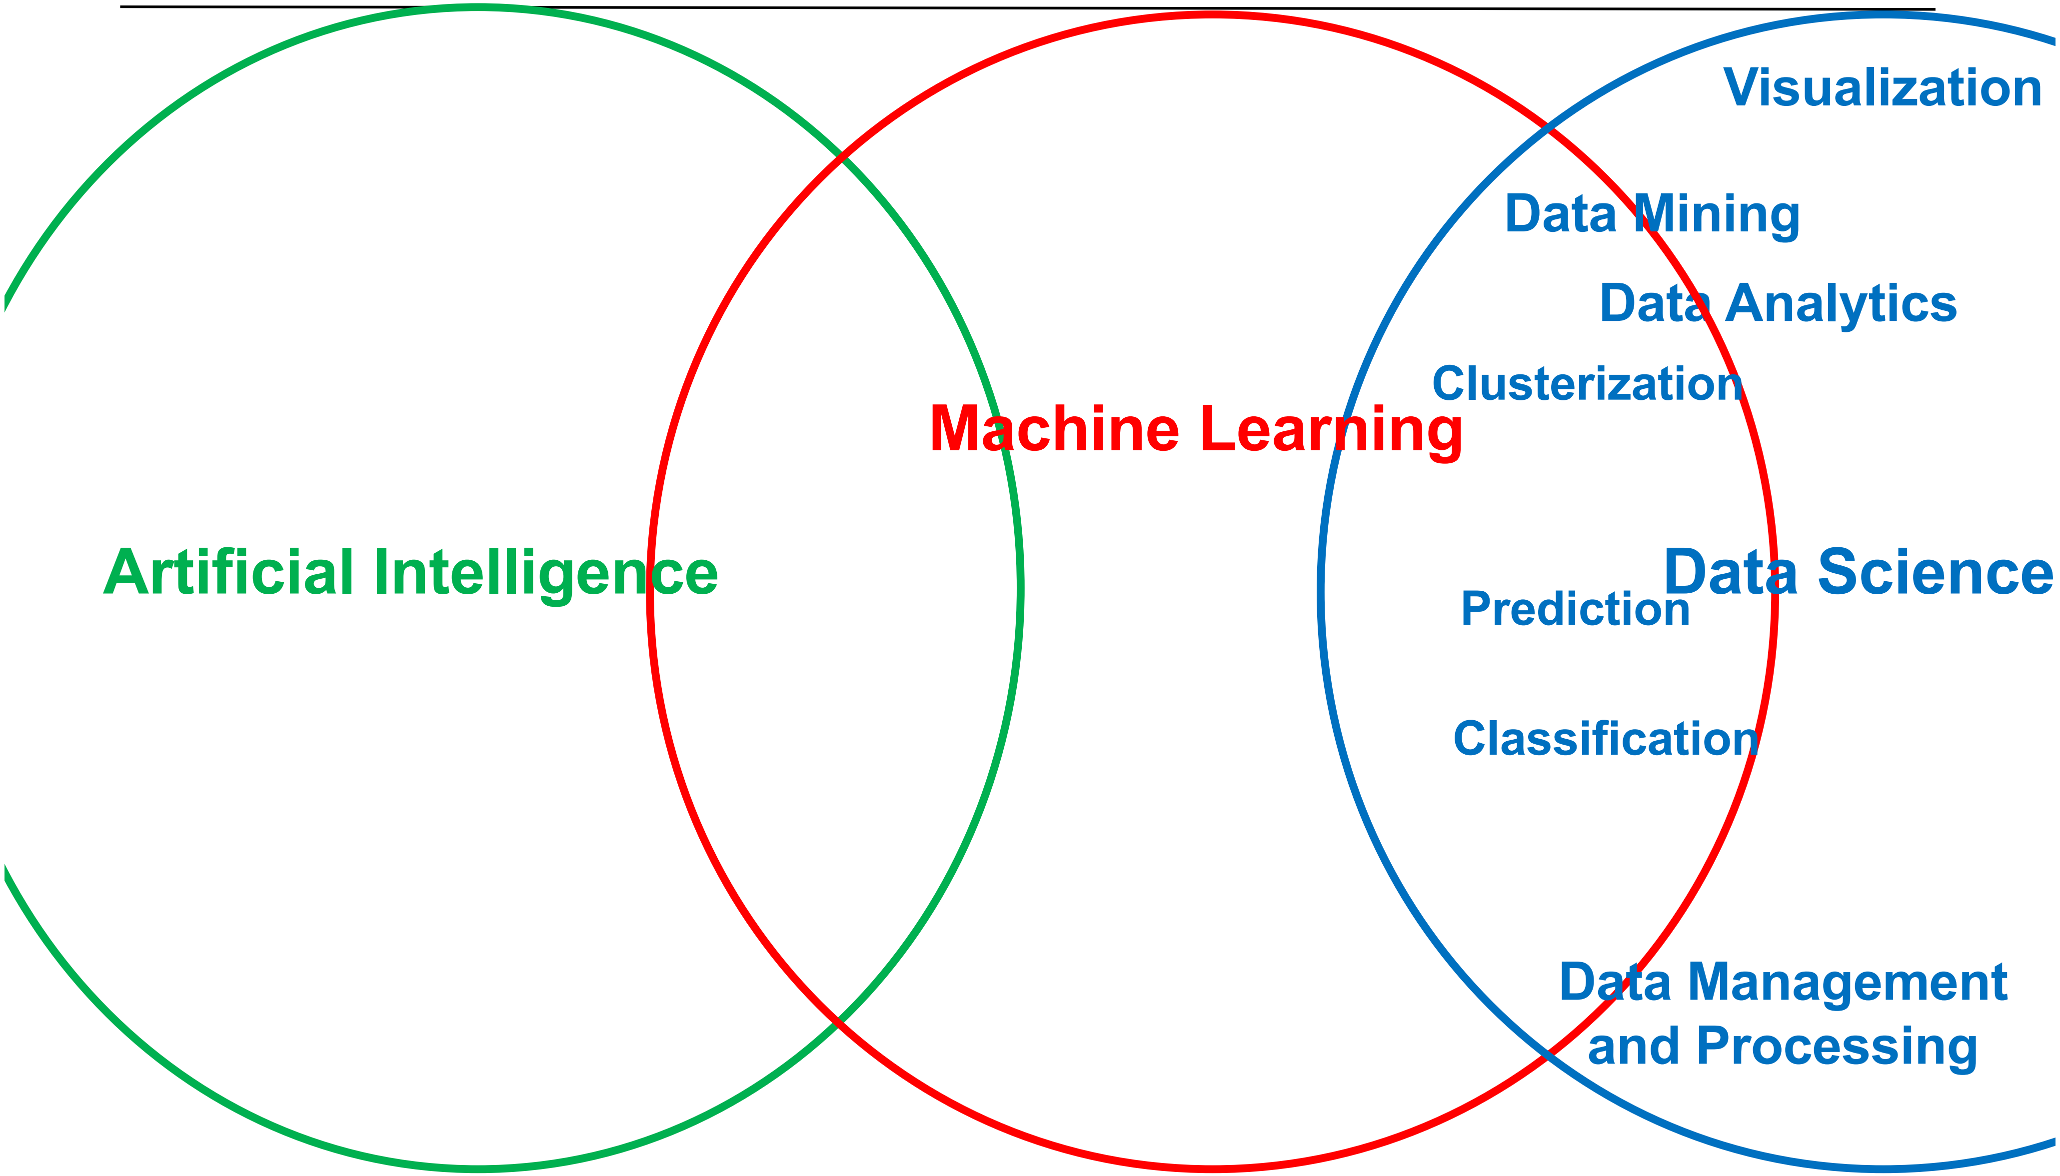

Terminology

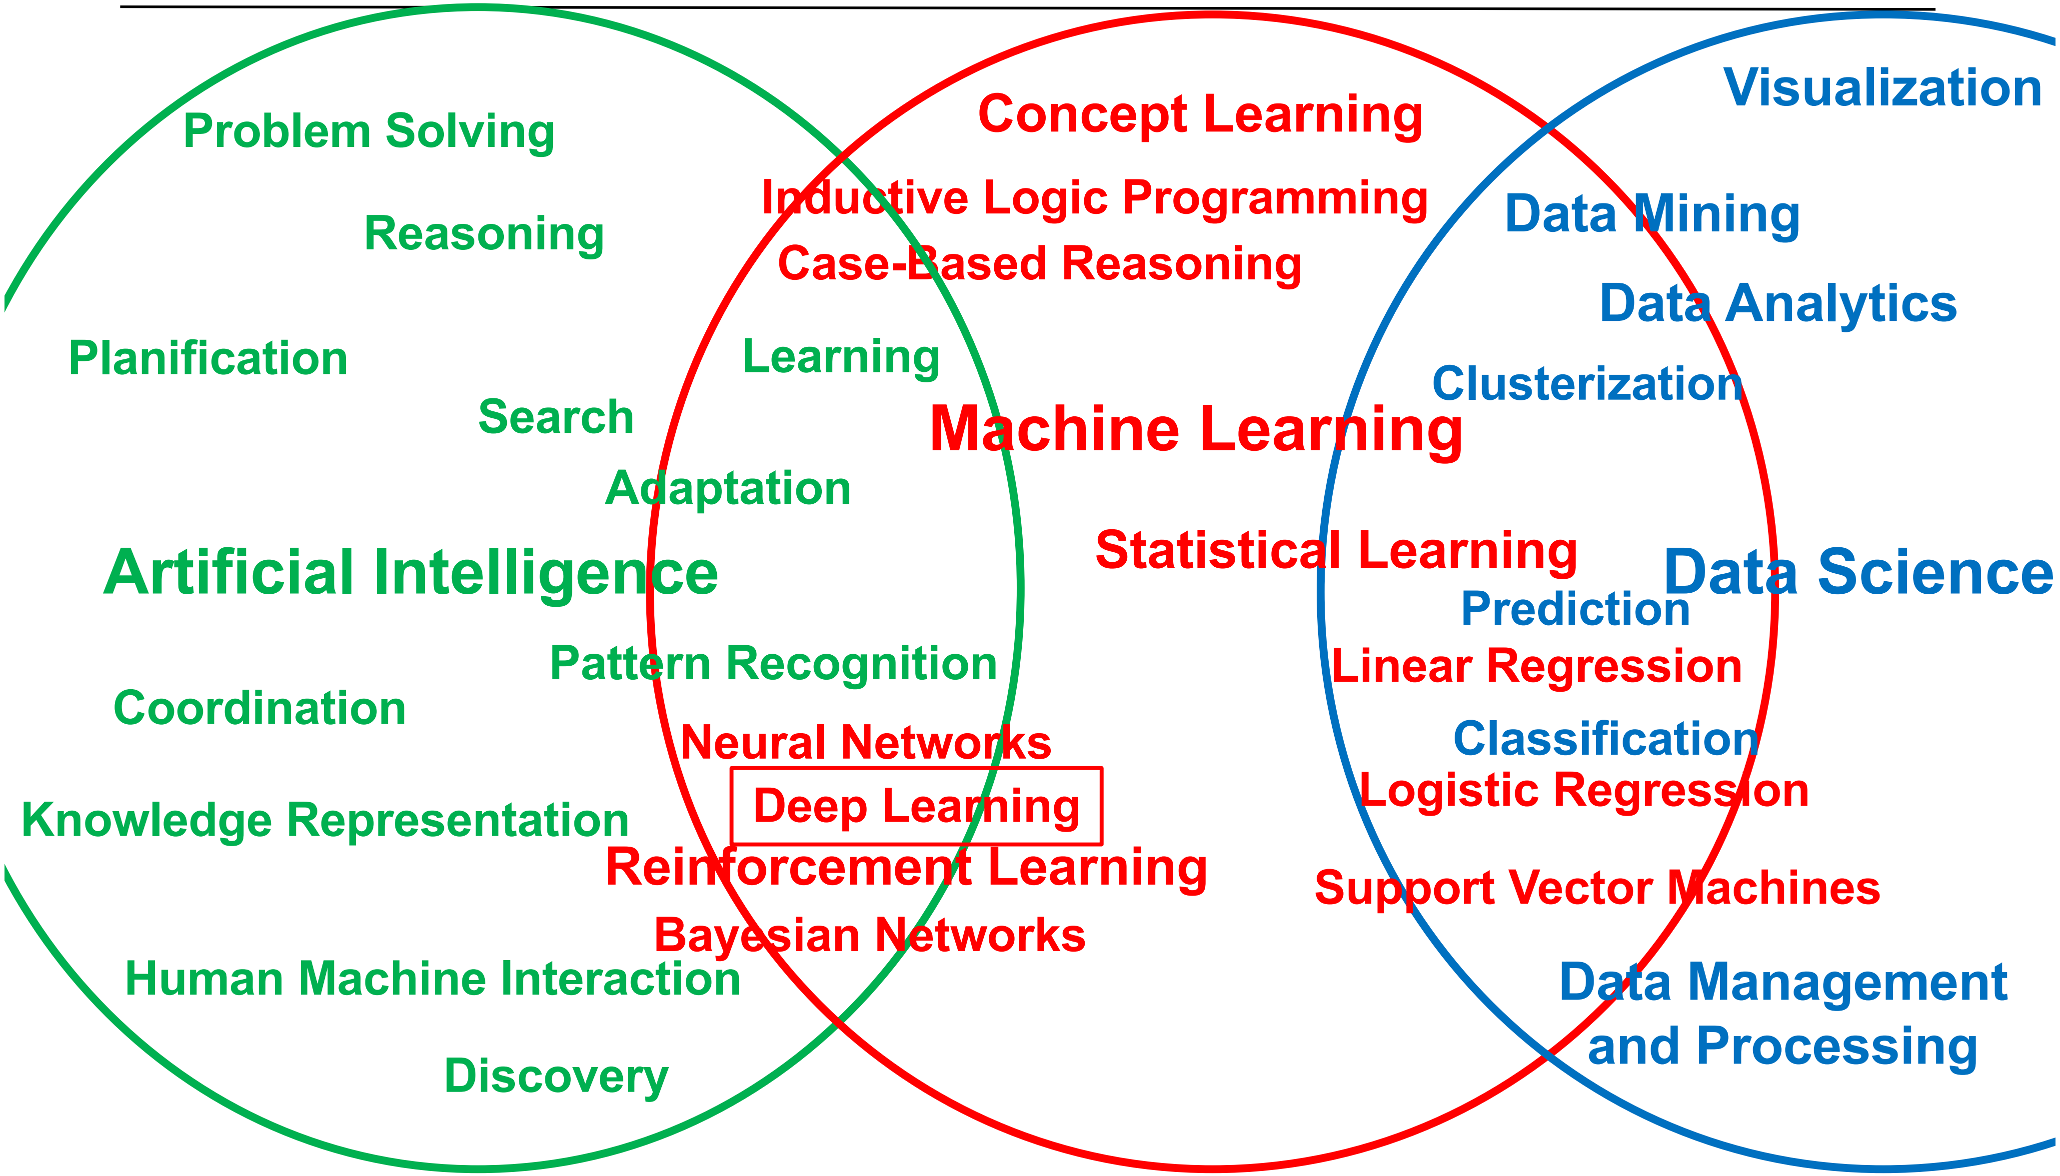

Terminology

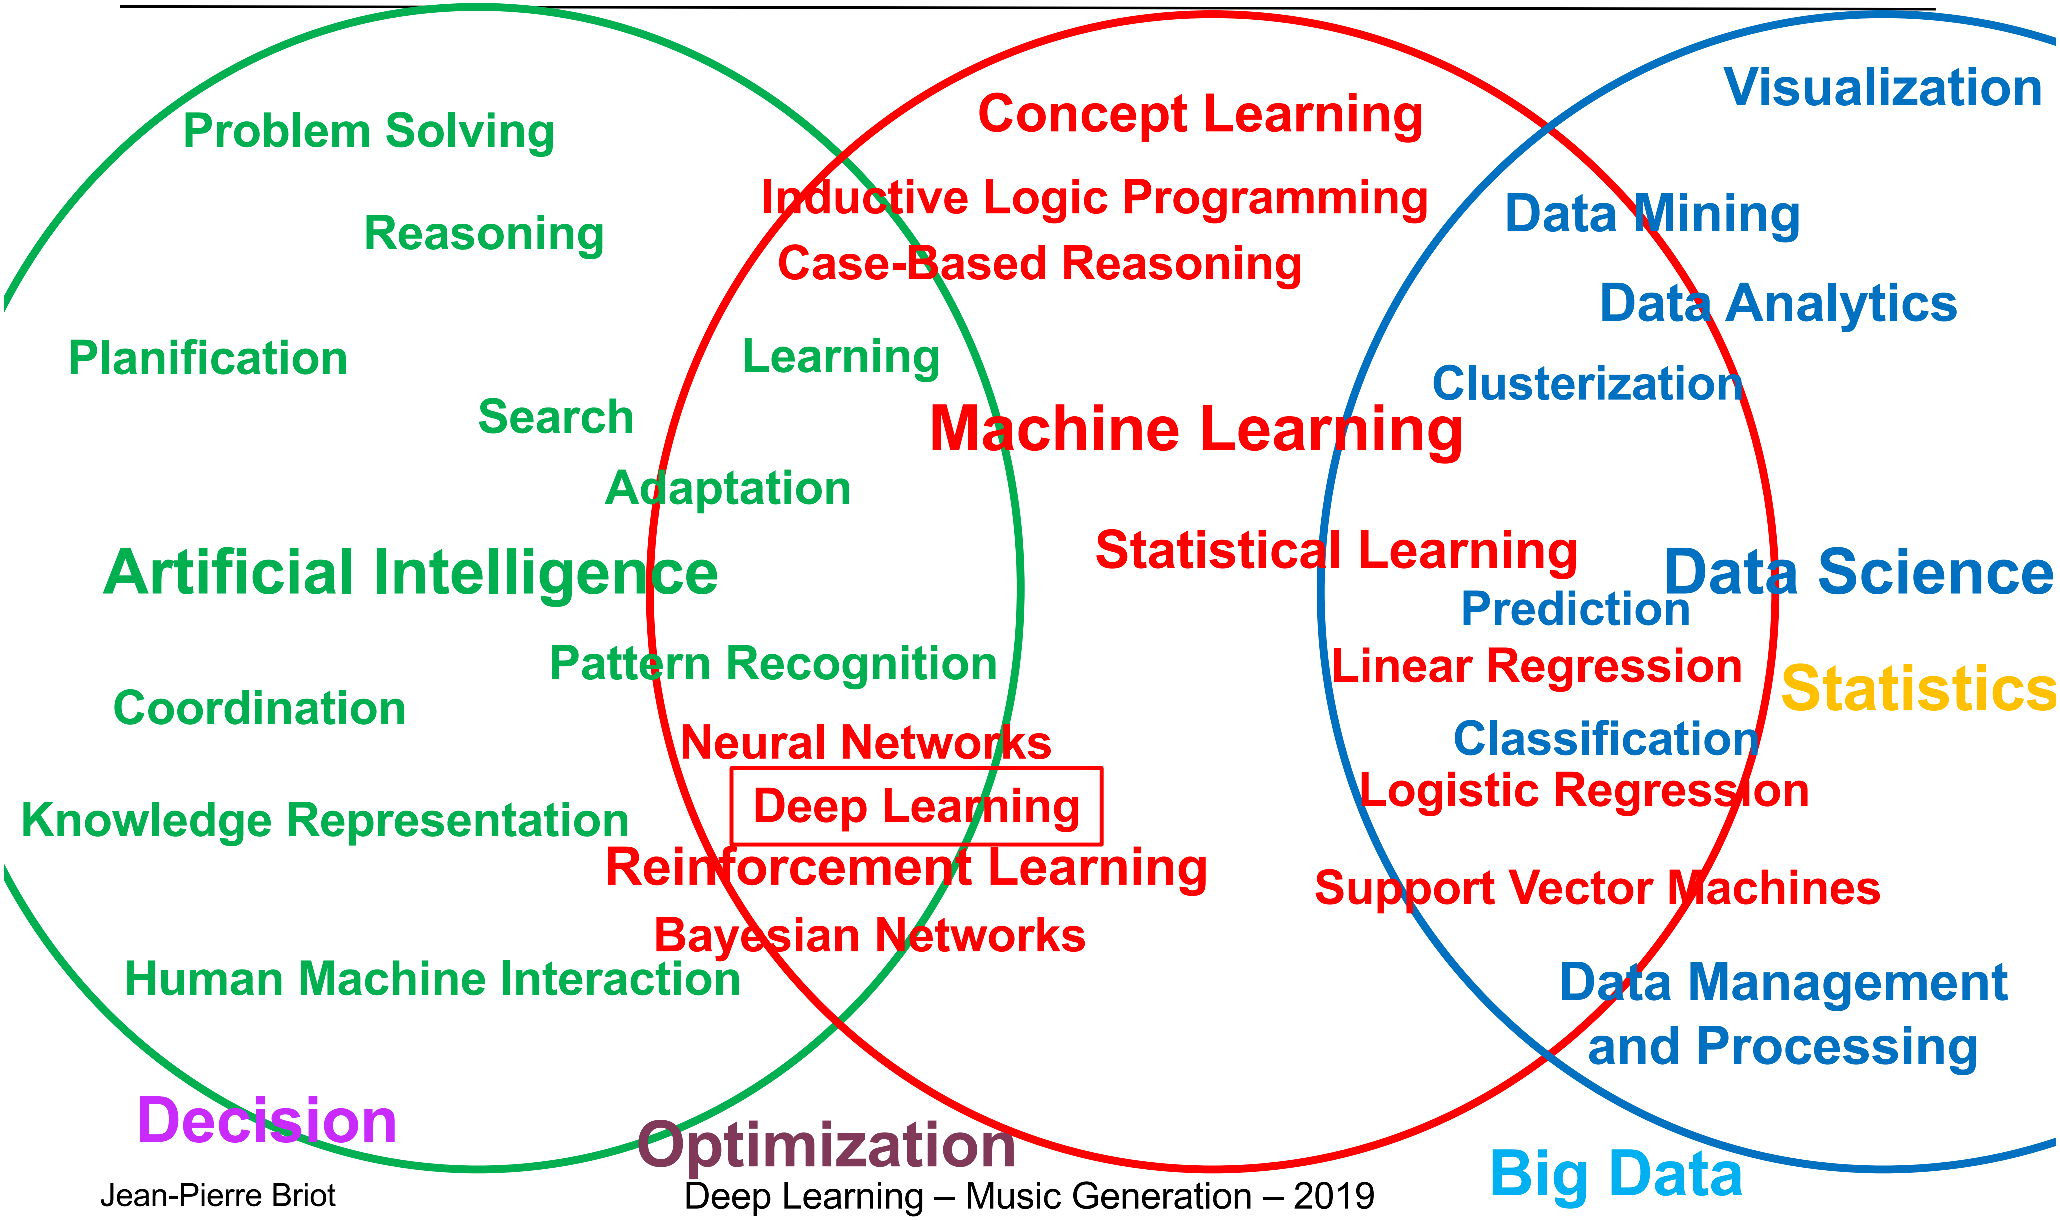

Correlation vs Causation

- Deep Learning Learns Correlations
- Does Function Mapping
- And Does it Very Well!

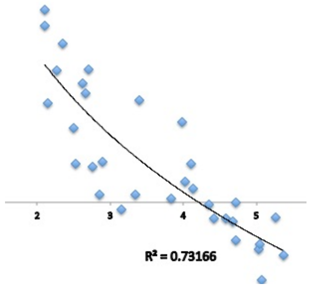

- It Creates a Predictive Model
- But not an Explicative Model

- Correlation $\not\Rightarrow$ Causation
- Still Missing Step

[Pearl and Mackenzie, 2018]

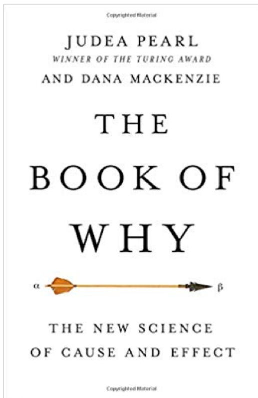

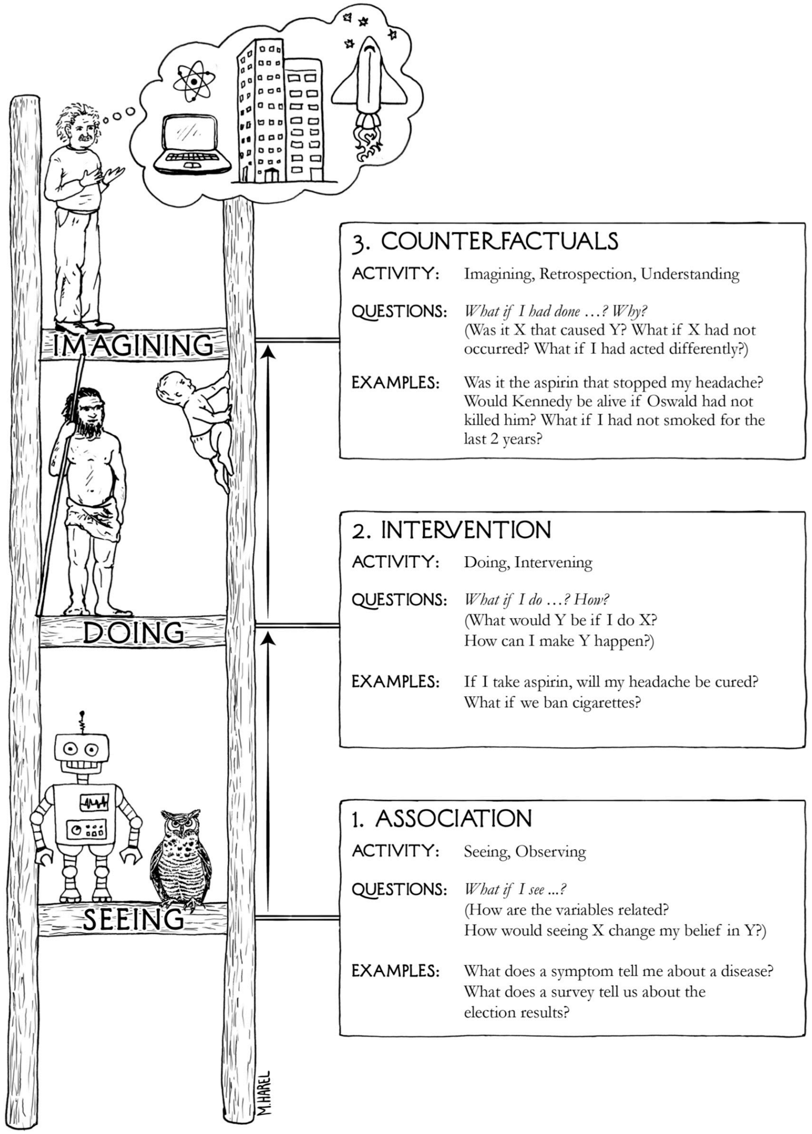

Ex. of Spurious Correlation (Confounding)

- Positive Correlation (for a country) between
  - Chocolate Consumption
  - Number of Nobel Prizes

- False Deduction/Causation:
  - More Chocolate -> More Nobel Prizes

- Common Cause: Country Wealthiness

- Chocolate <- Wealthiness -> Nobel Prizes

From Correlation to Causation

Causation Inference Engine [Pearl and Mackenzie, 2018]

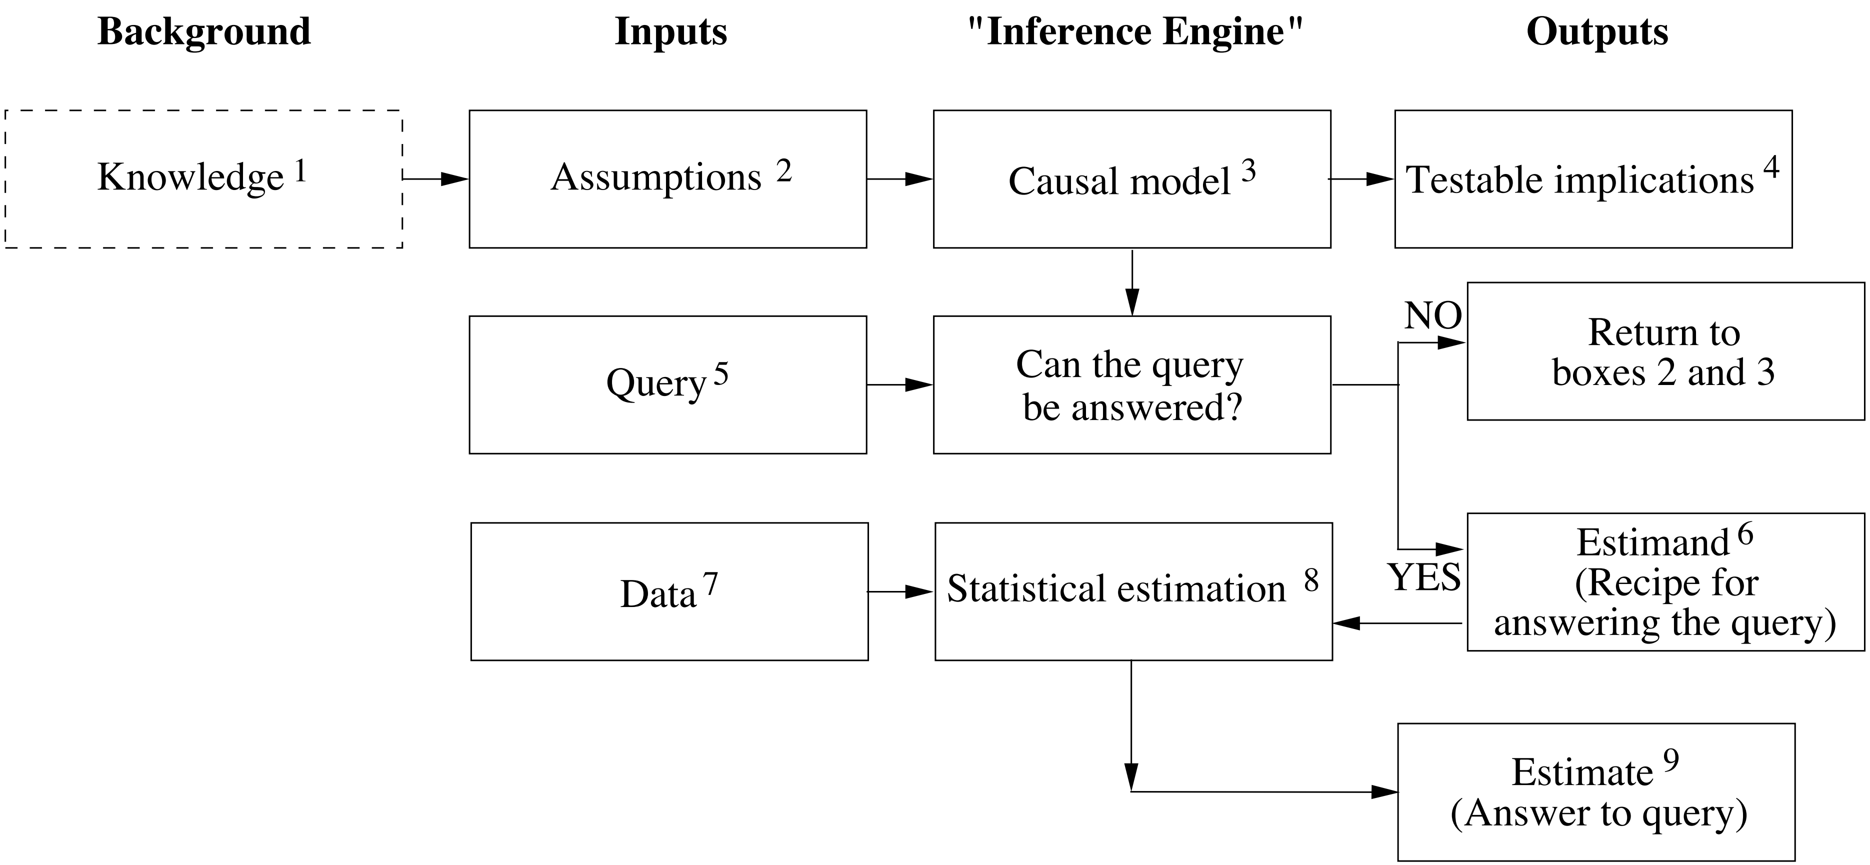

[Pearl and Mackenzie, 2018]

Modes of Creation

Handcrafted vs Learnt Models

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

Flow Machines [Pachet et al. 2012]

Artistic Content Generation Basic Cycle

- Curation
  - Collecting Examples (Training Set)
  - *Extensional Definition* of the Style
- Configuration
  - of the (Selected) Learning Model/Architecture
- Selection
  - Among Results Generated

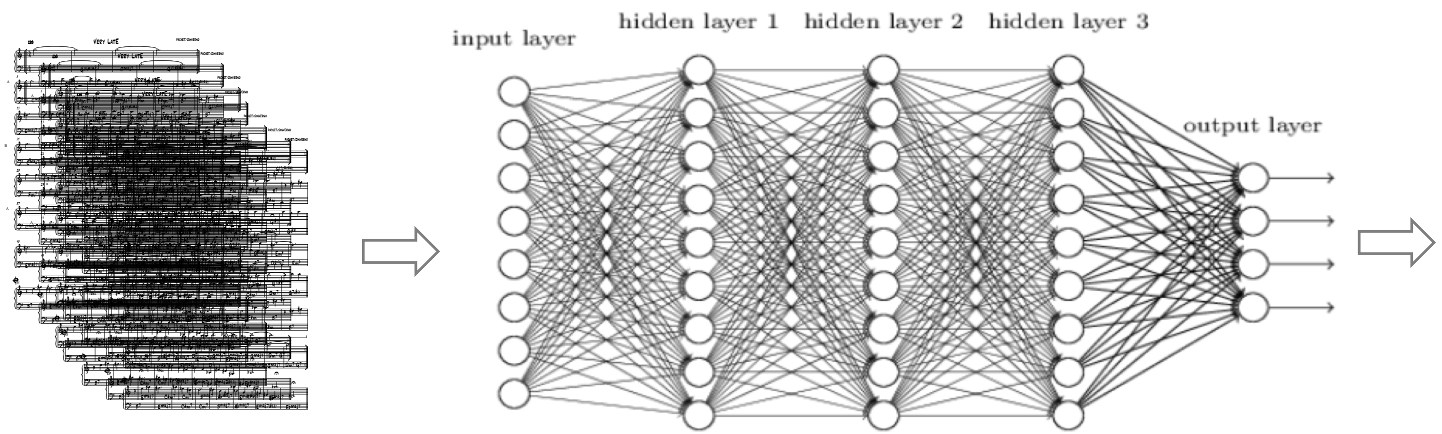

Curation Configuration Selection

Reorchestration of Ode of Joy by DeepBach (and other techniques [Flow Machines])

Ode to Joy in several styles

Autonomous vs Assisted Music Creation

- Autonomous Generation/Interpretation
  - Turing Test
  - Symbolic or/and Audio Music Generation
  - Parametrization/User Preferences (Style, Mood, etc.)
  - For Commercials and Documentaries
  - Create Royalty-free or Copyright-buyable Music
  - Ex:
    Jukedeck
    Aiva Technologies
    AMPER MUSIC

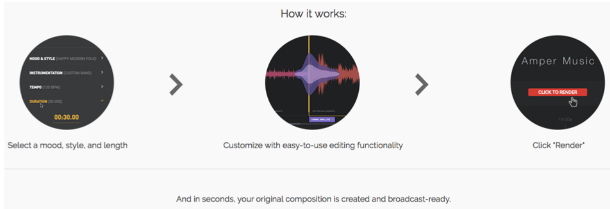

AMPER MUSIC

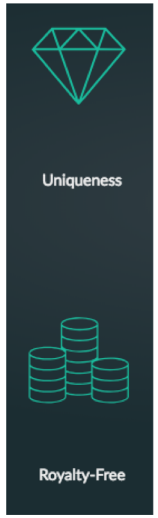

- Assistance to Human Composers and Musicians
  - Propose
  - Refine
  - Analyze
  - Harmonize
  - Produce
  - Ex: FlowComposer [Pachet et al., 2014]

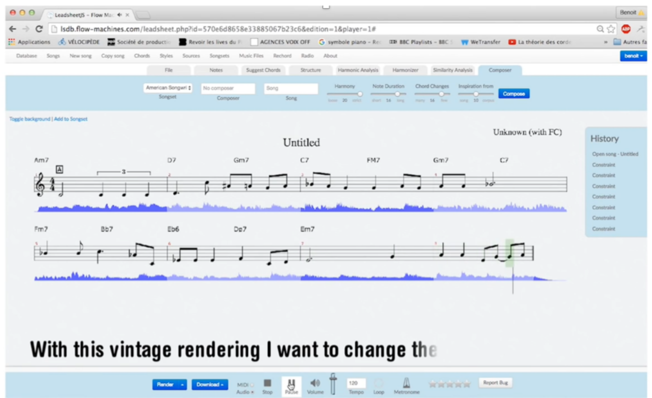

Objective and Evaluation [Pachet, 2019]

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

Some Preconcepts Against Deep Learning / AI

- No Emotion
  - Create Emotion to the Human Target ?
  - Or/And Internal Model of Emotion ?

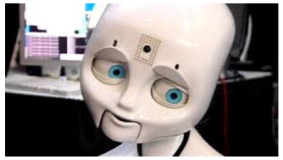

[Image: BBC]

[Karras et al., 2018]

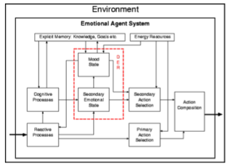

[Bryson et al., 2004]

- No Creativity
  - Exploratory
    - » AlphaZero used successful strategies yet unconsidered

  - Recombination
    - » Concept and Conjecture Discovery (ex: Numbers, Prime Numbers, Prime Numbers Decomposition) AM and Eurisko [Lenat, 1976; 1983]
    - » Style Transfer [Gatys et al., 2015]

  - Paradigm Reformulation
    - » Ex: Quantum Physics, Algebraic Geometry, Dodecaphonism…
    - » More difficult

Co-Creativity

- Co-Creation by Human(s)+Machine(s)
  - Ex: FlowComposer [Pachet et al., 2014]

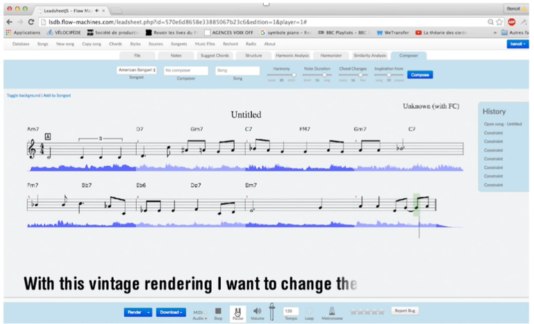

  - Continuator [Pachet, 2002]

  - Omax/DYCI2 [Assayag et al., 2003]

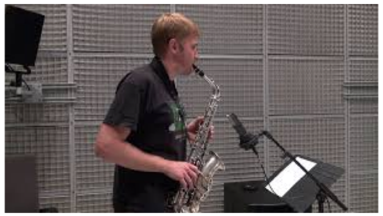

Autonomous vs Assisted Music Creation

"On the one hand, we have François Pachet's Flow Machines, loaded with software to produce sumptuous original melodies, including a well-reviewed album. On the other, researchers at Google use artificial neural networks to produce music unaided. But at the moment their music tends to lose momentum after only a minute or so."

[Creativity and AI: The Next Step – Combining two types of machine intelligence could open new frontiers of art, Arthur I. Miller, Scientific American, October 1, 2019]

Open Issues

- Structure
  - Ex: LSTM [Hochreiter & Schmidhuber, 1997]
  - Clockwork RNN [Koutnik et al., 2014]
  - SampleRNN [Mehri et al., 2017]
  - MusicVAE [Roberts et al., 2018]

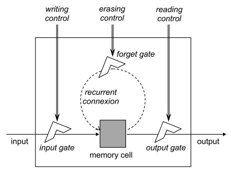

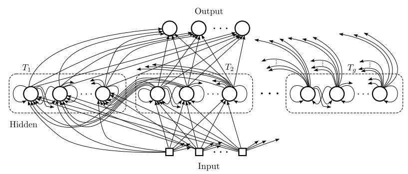

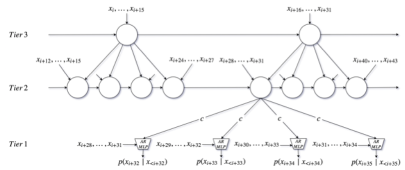

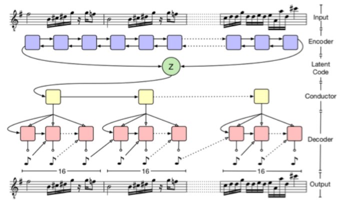

- Control
  - Tonality Conformance
  - Rhythm
  - Ex: C-RBM [Lattner et al., 2016]
  - Conditioning
  - Arbitrary Constraints

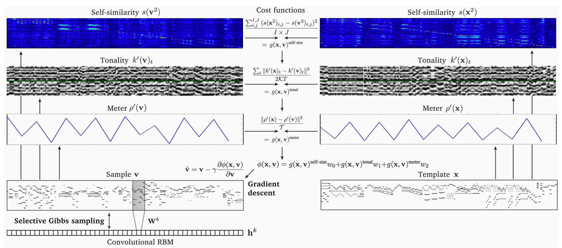

- Creativity Incentive
  - Vs Style Conformance
  - Ex: CAN [Elgammal et al., 2017]

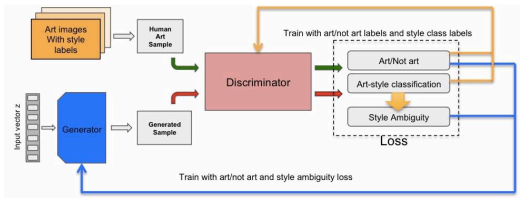

- Interactivity/Incrementality
  - Ex: DeepBach [Hadjeres et al., 2017]
  - Incremental Sampling

Style vs/and Control

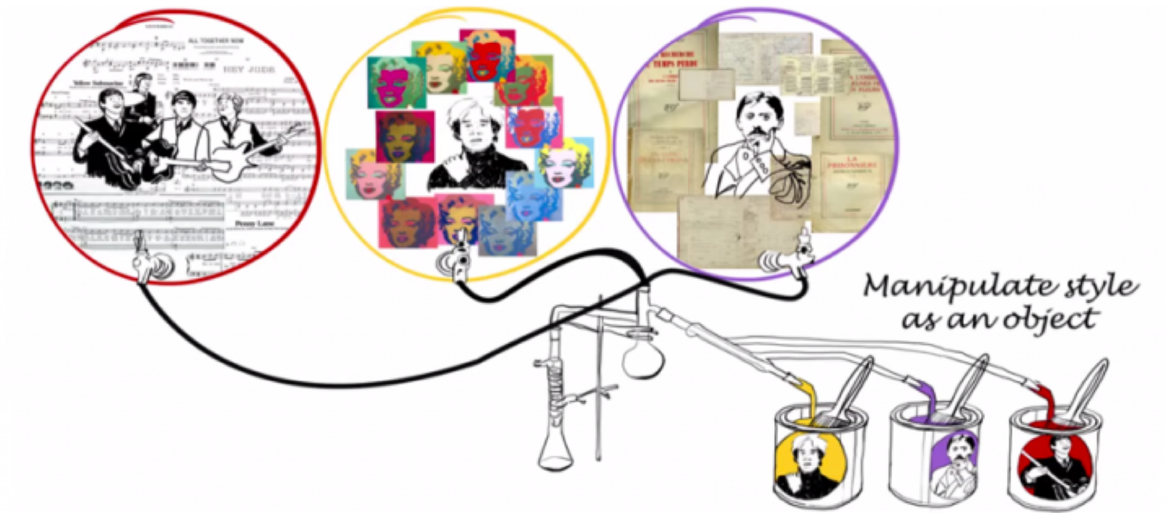

[Flow Machines]

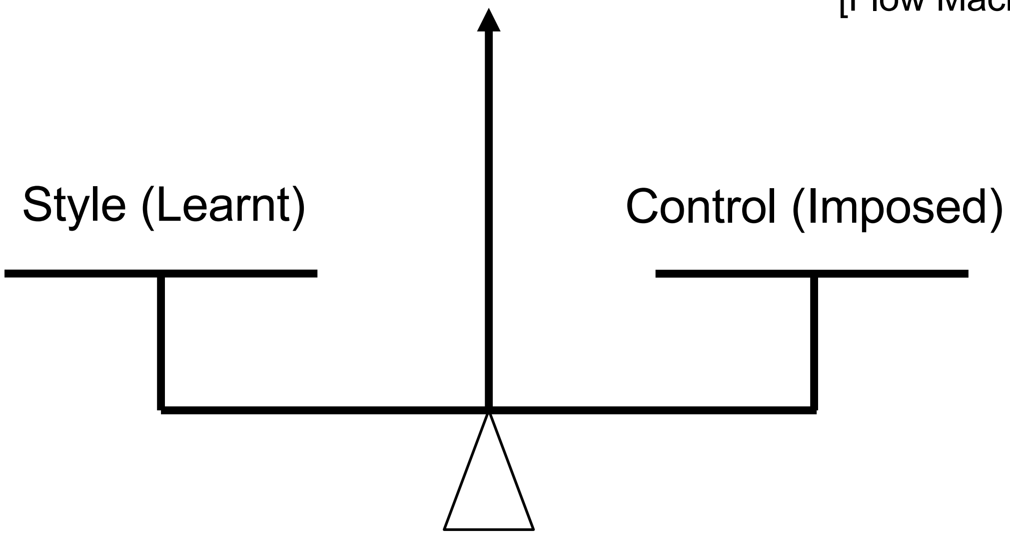

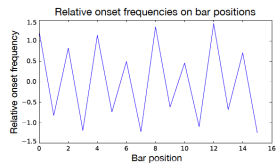

Style vs/and Originality

[Mimi & Eunice]

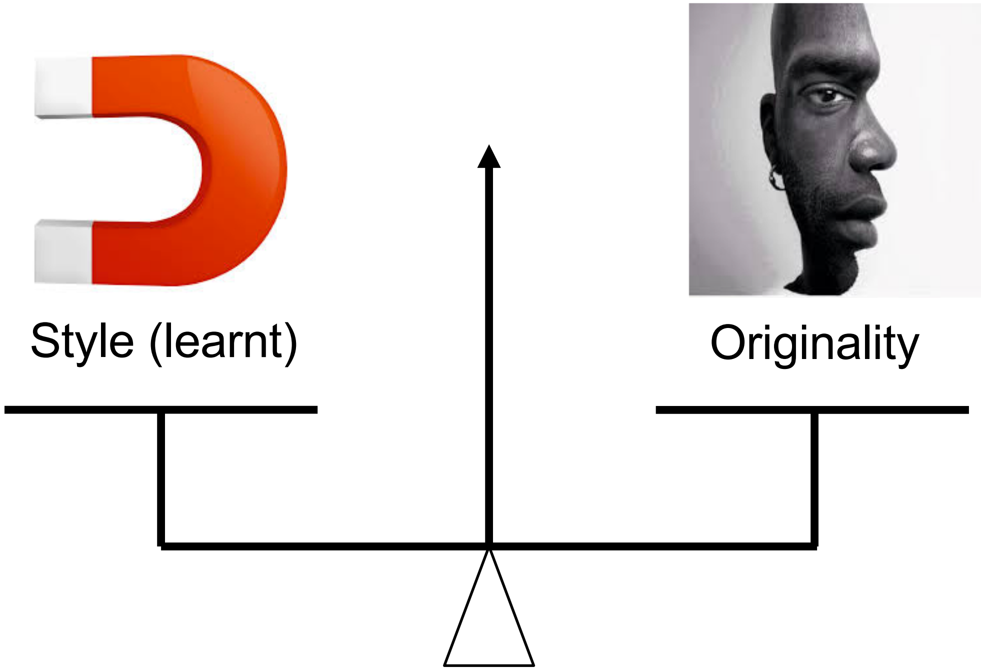

Conclusion

Conclusion/Prospects

- Deep Learning-based Music Generation
- Successes and Limits/Prospects

- Objective Loss Function Hypothesis
- Conformance Pros and Cons
- Control
- Context
- Explication

- Markov Models (and other Models) still Interesting
- Symbolic AI (GOFAI) still Necessary
- Automated Generation vs Human-Machine Co-Creation
- New Usages

Self-References for More Information

J.-P. Briot, G. Hadjeres, F.-D. Pachet, Deep Learning Techniques for Music Generation, Computational Synthesis and Creative Systems Series, Springer, 2019.

https://www.springer.com/br/book/9783319701622

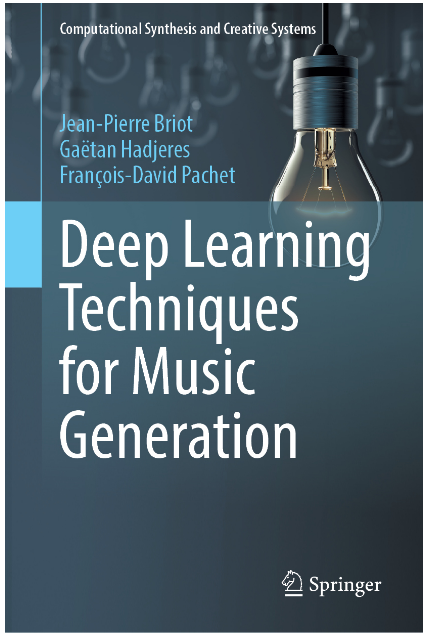

ArXiv version:

https://arxiv.org/abs/1709.01620

UNIRIO Course:

http://www-desir.lip6.fr/~briot/cours/unirio3/

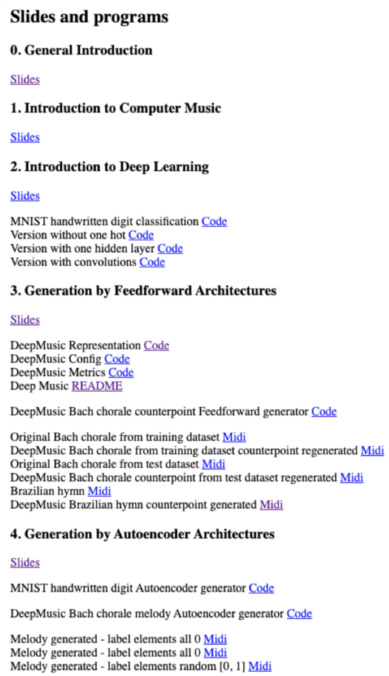

Slides and programs

0. General Introduction

Slides

1. Introduction to Computer Music

Slides

2. Introduction to Deep Learning

Slides

MNIST handwritten digit classification Code
Version without one hot Code
Version with one hidden layer Code
Version with convolutions Code

3. Generation by Feedforward Architectures

Slides

DeepMusic Representation Code
DeepMusic Config Code
DeepMusic Metrics Code
Deep Music README

DeepMusic Bach chorale counterpoint Feedforward generator Code

Original Bach chorale from training dataset Midi
DeepMusic Bach chorale from training dataset counterpoint regenerated Midi
Original Bach chorale from test dataset Midi
DeepMusic Bach chorale counterpoint from test dataset regenerated Midi
Brazilian hymn Midi
DeepMusic Brazilian hymn counterpoint generated Midi

4. Generation by Autoencoder Architectures

Slides

MNIST handwritten digit Autoencoder generator Code

DeepMusic Bach chorale melody Autoencoder generator Code

Melody generated - label elements all 0 Midi
Melody generated - label elements all 0 Midi
Melody generated - label elements random [0, 1] Midi

(Some) Other References

- Jordi Pons, Neural Networks For Music: A Journey Through Its History, October 2018, https://towardsdatascience.com/neural-networks-for-music-a-journey-through-its-history-91f93c3459fb
- Ian Goodfellow, Yoshua Bengio and Aaron Courville, Deep Learning, MIT Press, 2018
- Andrew Ng, Machine Learning Yearning, Deeplearning.ai
- Tom Mitchell, Machine Learning, McGraw Hill, 2017
- Pedro Domingos, The Master algorithm, Basic Books, 2015
- Judea Pearl and Dana Mackenzie, The Book of Why, Penguin Books, 2018
- 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

Thank You – Questions