




Artificial Intelligence in music composition

Bruno Isaac Jiménez Castro


Federico Michelotto

Late 18th century: "musical dice game"
(*Musikalisches Würfelspiel*)



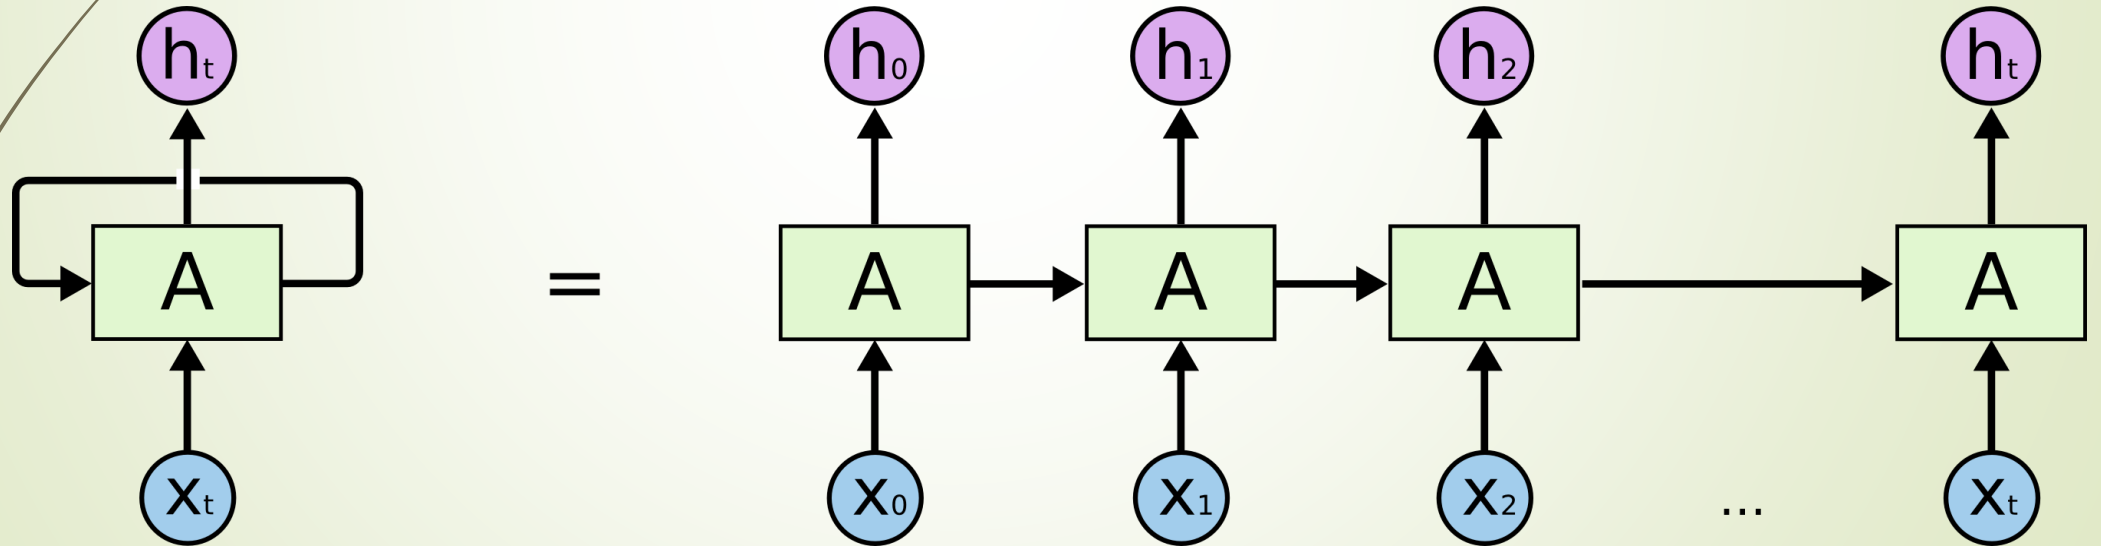
1950s: Markov chain

- State
- Probability



1980s: Recurrent Neural Networks (RNNs)

- Chain of Neural Networks
- Allow information to persist



An unrolled recurrent neural network




2002: “Long Short-term memory” (LSTM)


- Evolution of the standard RNN
- A more complex Neural Network to control how to storage past information
- A common LSTM unit is composed of a **memory cell**, an **input gate**, an **output gate** and a **forget gate**. The cell remembers values over fixed time intervals and the three *gates* regulate the flow of information into and out of the cell



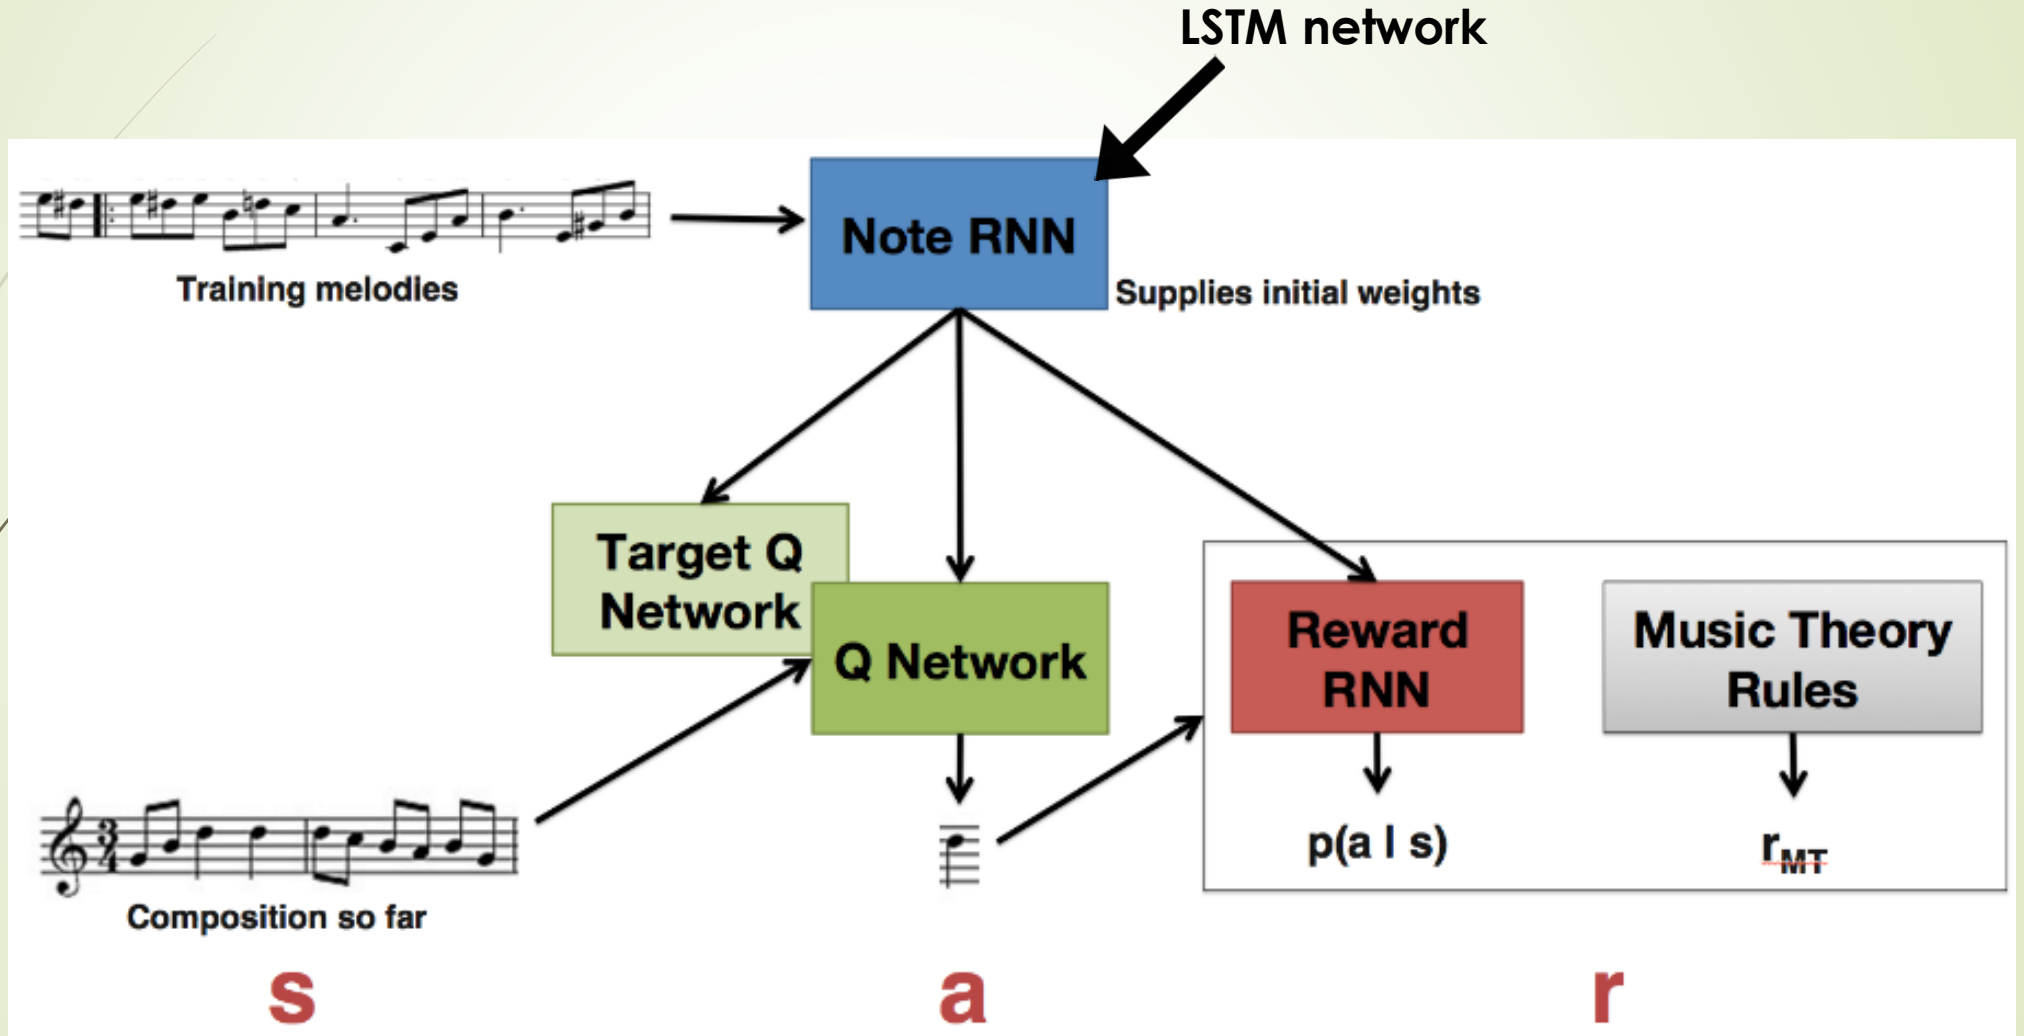
RNN



LSTM

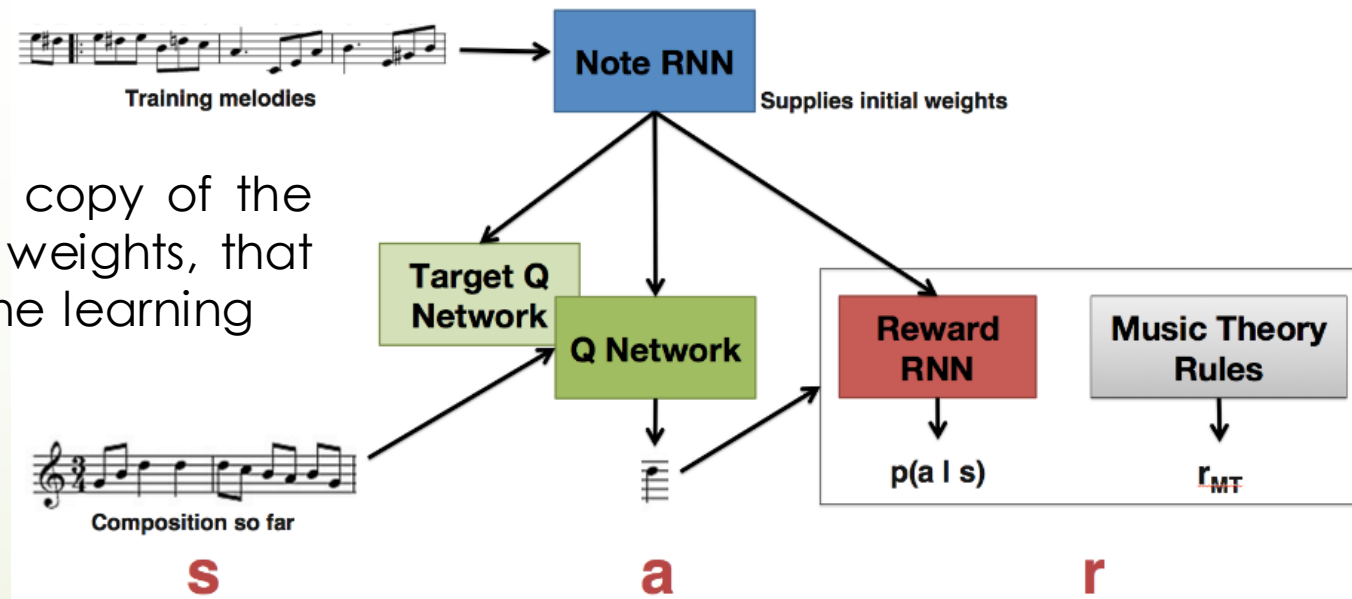


Composing music



1. Train the Note RNN with a lot of songs, and return initial weights to the RL models: Q-Network, Target-Q-Network and Reward RNN
2. The Q-Network generates one note at the time
3. Feed the Reward RNN with the note generated applying the music theory rules defined at the beginning
4. Calculate the reward and adjust the weights of the model and repeat from step 2

The Target-Q-Network is a copy of the Q-Network, but with fixed weights, that has the scope to stabilize the learning



Bellman optimality equation

$$Q(s_t, a_t; \pi^*) = r(s_t, a_t) + \gamma \mathbb{E}_{p(s_{t+1}|s_t, a_t)} \left[\max_{a_{t+1}} Q(s_{t+1}, a_{t+1}; \pi^*) \right]$$

π^* is the optimal policy, that say which action a take at state s

Loss function

$$L(\theta) = \mathbb{E}_{\beta} \left[\left((r(s, a) + \gamma \max_{a'} Q(s', a'; \theta^-) - Q(s, a; \theta)) \right)^2 \right]$$


Artificial Intelligence in music composition

- Mathematical approach of a human behaviour within an environment.
- Music depends mostly on human creativity.
- Can creativity be mathematically approached?



Rules

- All notes in a single key
- Tempo
- Do not repeat the same note >4
- Harmonic intervals
- Reward for motifs (musical ideas)





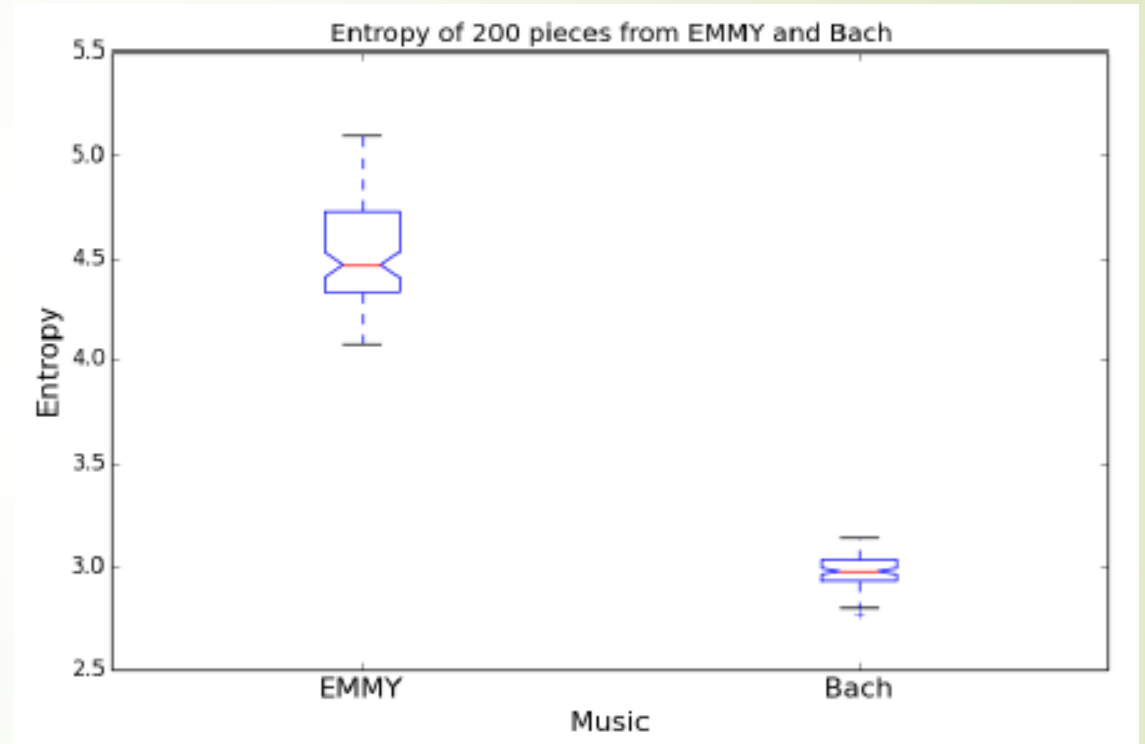
Emily Howell (90's)


- An AI able to mimic a style
- Expected to develop its own style



What is the difference then?

- Shannon entropy: unpredictability
- High entropy = surprisal element
- Low entropy = pattern





Thanks