

The aim of this thesis is to review the current state of machine learning in music composition and to train a computer on Beatles' songs using research project Magenta from the Google Brain Team to produce its own music. In order to explore the qualities of the generated music more thoroughly, we restrict ourselves to monophonic melodies only. We train three deep learning models with three different configurations (Basic, Lookback, and Attention) and compare generated results. Even though the generated music is not as interesting as the original Beatles, it is quite likable. According to our analysis based on musically informed metrics, generated melodies differ from the original ones especially in lengths of notes and in pitch differences between consecutive notes. Generated melodies tend to use shorter notes and higher pitch differences. In theoretical background, we cover the most commonly used machine learning algorithms, introduce neural networks and review related work of music generation.