

Machine Learning Methods for Music Discovery and Recommendation

Douglas Eck

Google Research

douglas.eck@gmail.com

research.google.com/

Abstract. In this talk I will relate current work at Google in music recommendation to the challenge of automatic music annotation (“autotagging”). I will spend most of the talk looking at (a) signal processing and sparse coding strategies for pulling relevant structure from audio, and (b) training multi-class ranking models in order to build good music similarity spaces. Although I will describe some technical aspects of autotagging and ranking via embedding, the main goal of the talk is to foster a better understanding of the real-world challenges we face in helping users find music they’ll love. To this end I will play a number of audio demos illustrating what we can (and cannot) hope to achieve by working with audio.

Bio

Douglas Eck is a research scientist at Google in Mountain View, California. His current focus is on machine learning models and user interfaces for music discovery and recommendation. This involves not only algorithm development but also user studies and data analyses to better understand what listeners want from a music service. Before coming to Google, Douglas was an associate professor in Computer Science at University of Montreal where he worked in related areas such as meter and beat induction, automatic tagging of music tracks and expressive timing and dynamics in music performance.