

Abstract

Optical Music Recognition (OMR) is the field of computationally reading music notation. This thesis presents, in the form of dissertation by publication, contributions to the theory, resources, and methods of OMR especially for handwritten notation. The main contributions are (1) the Music Notation Graph (MuNG) formalism for describing arbitrarily complex music notation using an oriented graph that can be unambiguously interpreted in terms of musical semantics, (2) the MUSCIMA++ dataset of musical manuscripts with MuNG as ground truth that can be used to train and evaluate OMR systems and subsystems from the image all the way to extracting the musical semantics encoded therein, and (3) a pipeline for performing OMR on musical manuscripts that relies on machine learning both for notation symbol detection and the notation assembly stage, and on properties of the inferred MuNG representation to deterministically extract the musical semantics.

While the the OMR pipeline does not perform flawlessly, this is the first OMR system to perform at basic useful tasks over musical semantics extracted from handwritten music notation of arbitrary complexity.