

Music-to-Knowledge (M2K): A Prototyping and Evaluation Environment for Music Information Retrieval Research

J. Stephen Downie
Graduate School of Library and
Information Science
University of Illinois at Urbana-
Champaign
jdownie@uiuc.edu

Andreas F. Ehmann
Department of Electrical and Computer
Engineering
University of Illinois at Urbana-
Champaign
aehmann@uiuc.edu

David Tcheng
National Center for Super Computing
Applications
University of Illinois at Urbana-
Champaign
dtcheng@ncsa.uiuc.edu

Categories and Subject Descriptors

H.3.4 [Information Storage and Retrieval]: *Systems and Software*

General Terms: Design, Evaluation

1. INTRODUCTION

This demonstration introduces the M2K (Music-to-Knowledge) music digital library (MDL) and music information retrieval (MIR) prototyping and evaluation environment to the information retrieval research community. M2K is being developed as part of the International Music Information Retrieval Systems Evaluation Laboratory project (IMIRSEL). The goal of the IMIRSEL project is the creation of standardized tools, datasets and TREC-like evaluations of MDL and MIR systems. M2K plays a central role in the upcoming Music Information Retrieval Evaluation eXchange (MIREX) contest being held over the spring and summer of 2005. MIREX participants will be using the M2K to prototype and then submit for formal evaluation their novel approaches to such MIR tasks as audio-based and symbol-based melody retrieval, artist identification, genre classification, etc. The MIREX participants will be meeting in September, 2005 at the 6th International Conference on Music Information Retrieval (ISMIR 2005) in London, UK to compare and contrast the success and failures of their techniques. MIREX information is available at: <http://www.music-ir.org/mirexwiki>.

2. M2K FEATURES

The computational backbone of M2K is the Data-to-Knowledge (D2K) machine learning and data mining environment developed by the Automated Learning Group (ALG) of the National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign (UIUC). The D2K framework is one where developers “wire” together computational modules into programs called *itineraries* which represent dataflow between modules. These itineraries can then be run, or nested within other itineraries and used as modules, allowing for the development of itineraries with arbitrary complexity.

The IMIRSEL team decided to base M2K on D2K because the D2K framework:

1. provides a visual, modular programming environment shown to cut development time (see Fig. 1);
2. supports distributed computing across multiple processors;
3. simplifies code reuse and sharing;
4. is written in Java for maximum portability; and,
5. includes a comprehensive set of data mining and retrieval modules including many used by IR/MIR researchers such as neural nets, Bayesian classification, decision trees, etc.

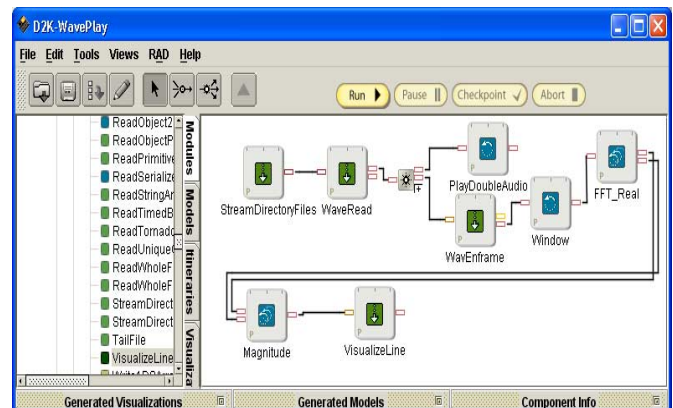


Figure 1. An M2K real-time spectral analyzer itinerary built for the exploration of audio music files prior to indexing (put together in less than five minutes!).

M2K is being released as an open source distribution to encourage the development of MIR-specific, and more general IR modules. Information about M2K is available at <http://music-ir.org/evaluation/m2k>. Information about D2K (needed to run M2K modules and itineraries) is available at <http://alg.ncsa.uiuc.edu/do/downloads>.

3. ACKNOWLEDGEMENTS

We thank the Andrew W. Mellon Foundation, and the National Science Foundation (grant numbers IIS-0340597 and IIS-0327371). Special thanks to Joe Futrelle, Kris West and Paul Lamere for their assistance in making M2K a reality. We thank the 2005 committees of the Joint Conference on Digital Libraries and SIGIR for permission to cross-present this demonstration.