



Machine Learning and Data Mining for Computational Creativity

Hannu Toivonen

University of Helsinki

www.cs.helsinki.fi/hannu.toivonen



Self-determinism and creativity

- A purely preprogrammed generative system
 - only does what it was told to do
 - has little creativity
 - Adaptivity or self-determinism
 - Is necessary to attribute any creative autonomy or originality to a creative system
 - Transformative or meta-level creativity (cf. Boden, Wiggins) can be attributed with higher creativity
 - ...but how to build a system to deal with unanticipated cases?
- Opportunities for ML and DM



Data Mining and Machine Learning in Computational Creativity


– A review on the topic:

Toivonen and Gross: Data Mining and Machine Learning in Computational Creativity. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*. Accepted for publication, 2015.



ML and DM in CC


- Let's use a simple generate-and-test model to illustrate uses of machine learning (ML) and data mining (DM) in CC





Learning to evaluate

- Use ML to learn an evaluation function $\text{eval}(a)$ from training examples
 - E.g. a classifier that tells if the result is good
- Assuming a generator $\text{gen}()$ exists, its outputs are filtered by the trained classifier without explicit directions by the programmer





Learning to evaluate

An example system, DARCI (Ventura et al)

- Creates images that express an emotion
- Emotion detection is based on artificial neural networks trained by users of the system
- A genetic algorithm is used as generator `gen()`
 - Adapts to the evaluation/fitness function `eval()`
- <http://darci.cs.byu.edu/>
- "DARCI, draw me a happy picture!"



A happy image by DARCI, <http://darci.cs.byu.edu/>




Learning to evaluate

Bottlenecks in learning the eval() function

- Learning an evaluation (or fitness) function eval(a) can be very difficult
 - How does one evaluate the quality of a poem?
- Generating complex artefacts, i.e., writing (or learning) the function gen(), can be very hard
 - In practice, the generation step must be adaptive in order to be effective
- Pastiche generation, i.e., mere imitation of training examples rather than creativity



Learning to Generate



- Predictive models
- Generative models



Learning to Generate Using Predictive Models

1. Completion of partial artefacts

- Given some part of the artefact, predict the values of the remaining parts
- Based on training on complete artefacts

E.g. harmonization of music:

- Given a melody (possibly created by the system itself), choose suitable chords to accompany the melody



Learning to Generate Using Predictive Models

2. Reduce the task of generating complex structures to selection.

E.g. generation of accompaniment by running a classifier to pick a suitable chord, and then using (possibly automatically extracted) patterns to generate the exact accompaniment



Learning to Generate Using Predictive Models

3. Generate complex structures using instance-based techniques

- E.g. k-nearest neighbours and case-based reasoning
- avoids using models, decision structures, or patterns
 - can be difficult to specify or learn
 - could be restrictive.

Example: Corpus-based poetry by Toivanen et al.

- No explicit grammar, instances simply copied from a corpus



Learning to Generate Using Generative Models

Generative models (from ML and statistics) can be used more directly to generate artefacts

- E.g. Markov models for sequences such as text and music
- Artificial neural networks, with slight modification of weights (and keeping the input constant)



Mining patterns for creative tasks

1. Use data mining to discover patterns in, say, text
2. Utilize these patterns in a generation function `gen()`

Examples:

- Association-based creativity (Gross et al)
- Corpus-based poetry (Toivanen et al)



Mining patterns for creative tasks

Example: metaphor generation (Veale et al)

1. Extract similes (“strong as a bull”) from a corpus
 - Look for patterns of the form “T is as P as a V”
2. P (“strong”) is a typical property of V (“bull”) if the pattern “T is as strong as a bull” occurs often
3. To express “he is strong” in a metaphorical way, find a noun V for which “strong” is a typical property
 - Bull is found as a suitable V
4. Output “he is a V”, i.e., “he is a bull”

<http://ngrams.ucd.ie/metaphor-eye/>



Metaphor-Eye

Why are scientists like artists?

- Scientists
 - ...develop ideas like artist
 - ...explore ideas like artist
 - ...acquire skills like artist
 - ...spread ideas like artist
 - ...nurture ideas like artist
 - ...develop techniques like artist



Transformational Creativity Using Data Mining and Machine Learning

Wiggins suggests uses of ML/DM:

- Automatic adaptation of R or T
 - To remedy aberration: use aberrant concepts as positive or negative examples, depending on their value
 - To remedy generative uninspiration: use positive (and negative) examples received from outside
- Automatic adaptation of E
 - Use feedback and evaluations received from outside (not covered by Wiggins)

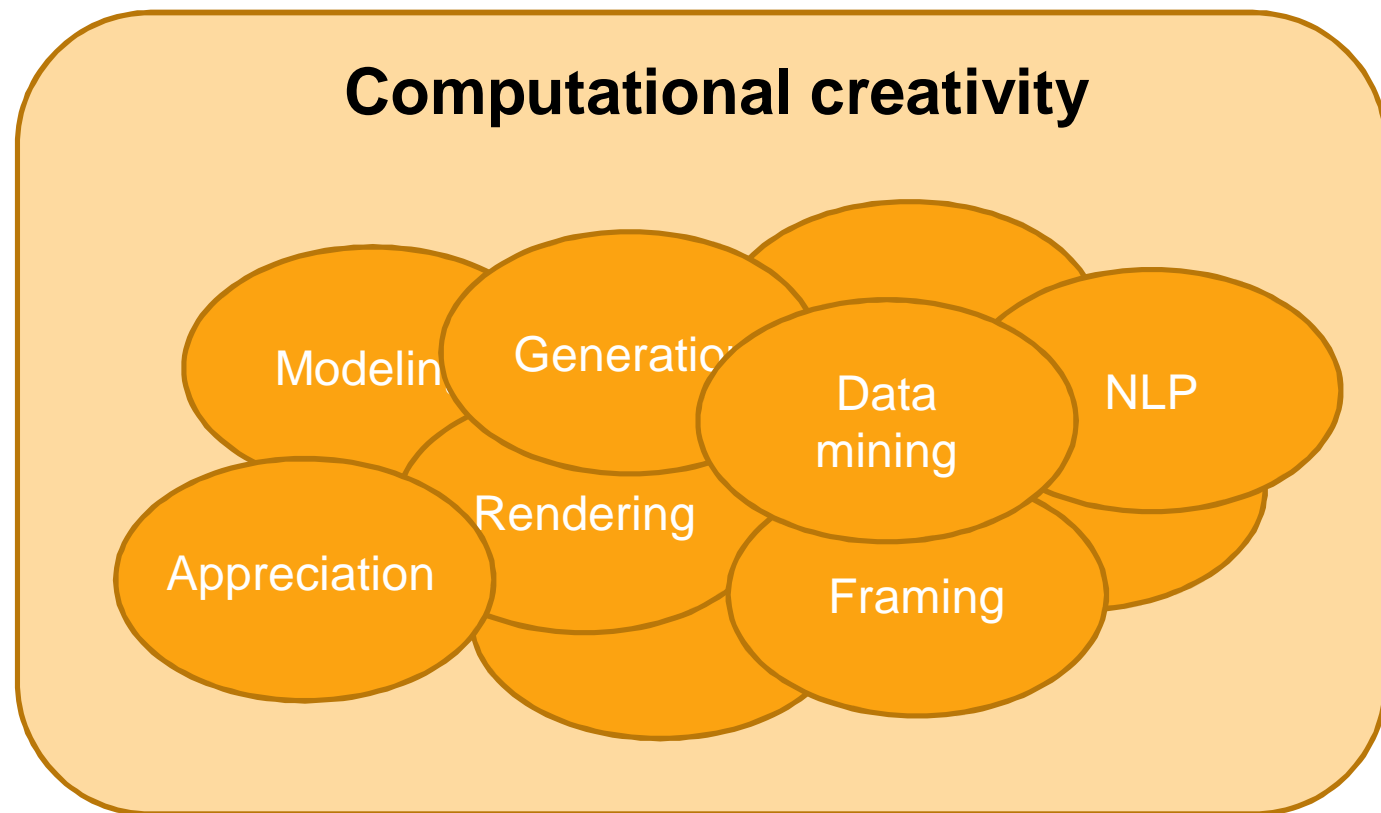


Computational Creativity for Data Analysis

Tulilaulu et al (2012 + manuscript)

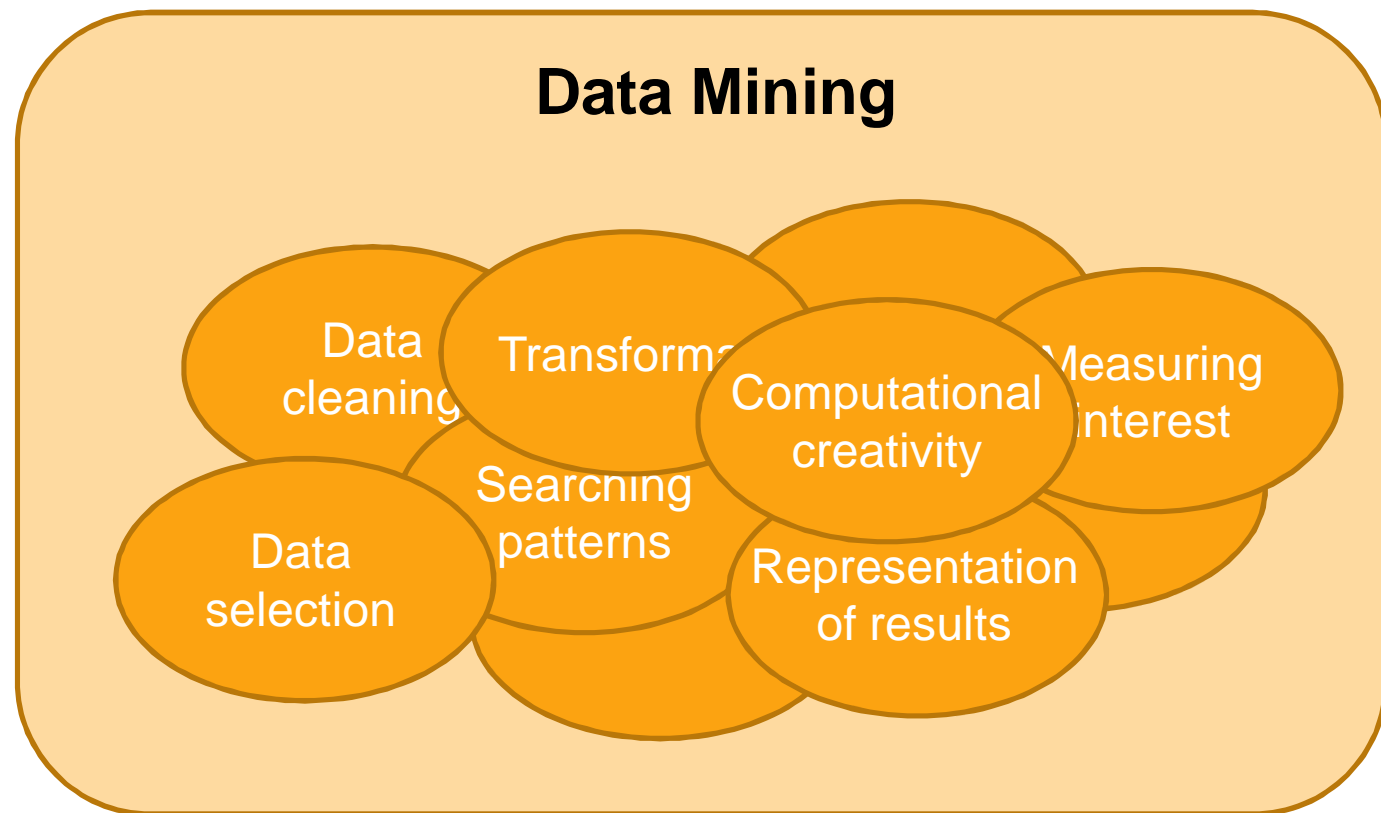


Data mining for creativity





Creativity for data mining





ICT as an amplifier of productivity

Faster, cheaper,
everywhere,
all the time

Burden on users,
stress, lack of free
time

ICT as an amplifier of mental satisfaction

Positive emotions,
mental wellbeing

Support creativity,
learning, meaningful
leisure time



*Can data be turned into a
subjective, esthetic experience?*

- Given **your** data, could **you** enjoy listening to it as music?
- Could you even feel **joy of creativity**?



Data Musicalization

- Using given data as input,
- automatically compose a novel piece of music



New goals for data analysis

Traditional data analysis and representation (statistics, visualization, sonification, etc.):

- Transfer of information
- Objectivity
- (Cognitive emphasis)

Data musicalization:

- Experiences involving feelings and emotions
- Subjective
- (Affective emphasis)



Motivation for data musicalization

- Esthetic pleasure
- Joy of creativity
- Transfer of (some) information in a novel way
 - Possibly in the background, unconsciously
 - Like sonification, but musical
- Building an emotional attachment to an application



Sleep musicalization

- An application of data musicalization
- A piece of music is composed from a night's sleep
- *Goal: help and motivate users track their sleep and eventually improve their sleep*
- Two phases
 1. Data analysis:
from sensor data to sleep measurements
 2. Composition algorithm:
from sleep measurements to music



Musicalization vs. sonification

Sonification:

- Maps data to sounds
- Not necessarily musical or creative
- Focus on conveying information objectively
- Novelty here: musicality, subjective experience




Phase 1, data analysis

From sensor data to sleep measurements



Sensor data




A thin piezo-electric force sensor placed under the mattress topper (from Beddit Ltd)

Force signal with sample rate 140 Hz (here a 12-second signal excerpt)




Analysis of sensor data





Sleep analysis results





Phase 2, composition

From sleep measurements to music



Composition of music from sleep

Design principles/goals:

- Compose a novel piece of music
- Use sleep measurements to guide the composition
- Produce music, not sonification
- Make the music reflect properties of sleep
- Compress 8 hours of sleep to couple of minutes
- Use simple methods, build a proof of concept



Composition in 6 steps

1. Generation of harmonic progression (chord seq.)
2. Melody generation
3. Generation of rhythm
4. Generation of accompaniment
5. Adjusting volume
6. Regulation of tempo



Correspondence between sleep and music

In a nutshell:

- Deeper sleep corresponds to calmer music
- Lighter sleep corresponds to livelier music
- Each sleep stage has a different accompaniment (and theme)
- Tempo is regulated by the heart rate
- More activity corresponds to louder music



sleepmusicalization.net

- Sleep musicalization is available as a public web app at <http://sleepmusicalization.net>
- You can have your sleep measurements composed into music (if you have a Beddit sleep sensor), or
- Listen to songs published by other users



sleepmusicalization.net


Sleep musicalization

Composed songs

Sleep musicalization

Perceive your sleep as a unique musical experience!

Musicalization turns data into a genuine piece of composed music.

 Try it on your Beddit data!

Introduction

Sleep musicalization is a novel way of perceiving and experiencing sleep measurement data. The goal is to help users understand and analyze their sleeping patterns and eventually improve their sleep.

The musicalization process follows musicological principles when composing a melody, designing the rhythm and changes in tempo, arranging the accompaniment, and playing out the music at different levels of volume. These aspects are inspired but not dictated by the data. The result of musicalization of eight hours of sleep is an original piece of couple of minutes of music.

Musicalization of data provides a whole new way to experience data as a music. Music has a unique capability to invoke emotions, giving users a novel opportunity to perceive their data

Listen to latest samples

Sleeper Agent

Contributed by UFOPOLI

 Listen

Wild Trances

Contributed by UFOPOLI

 Listen

Deep dreams

Contributed by discovery

 Listen

Hannun viime yö

Contributed by discovery

 Listen

eva 09/11/12


 Listen



Listening to sleep music

Sleep musicalization

Composed songs




Deep dreams

Share this song

Copy-paste the following link to email, discussion, etc:

<http://sleepmusicalization.net/song/wOqbl1icfDNE>



Sleep stages visualized

The hypnogram on the left shows visualized sleep stages.

Movements during the sleep

The actigram on the left shows the amount of movements during the night.


My songs

There are no songs yet

Latest songs

Sleeper Agent

Contributed by UFOPOLI



Wild Trances

Contributed by UFOPOLI





A different example

Sleep musicalization

Composed songs

Disc's sleep

Sign out




Sleep with sick children

...is like running a marathon while you should sleep.

Share this song

Copy-paste the following link to email, discussion, etc:

<http://sleepmusicalization.net/song/8hQRR20FZ1xX>



Sleep stages visualized

The hypnogram on the left shows visualized sleep stages.

Movements during the sleep

The actigram on the left shows the amount of movements during the night.

My songs

What is this?

Contributed by discovery

Up to 05/27/12

Latest songs

Sleeper Agent

Contributed by UFOPOLI

Wild Trance

Listen

Listen



Independent reviews



“[The songs] do have a uniquely personal and (dare I say it?) dreamlike feel”

- Leslie Katz, CNET




“[...] these songs left me with a sense of incompleteness, as if the piano-dominated music would be an inadequate representation of my dreams.”

- Nic Halverson, Discovery News

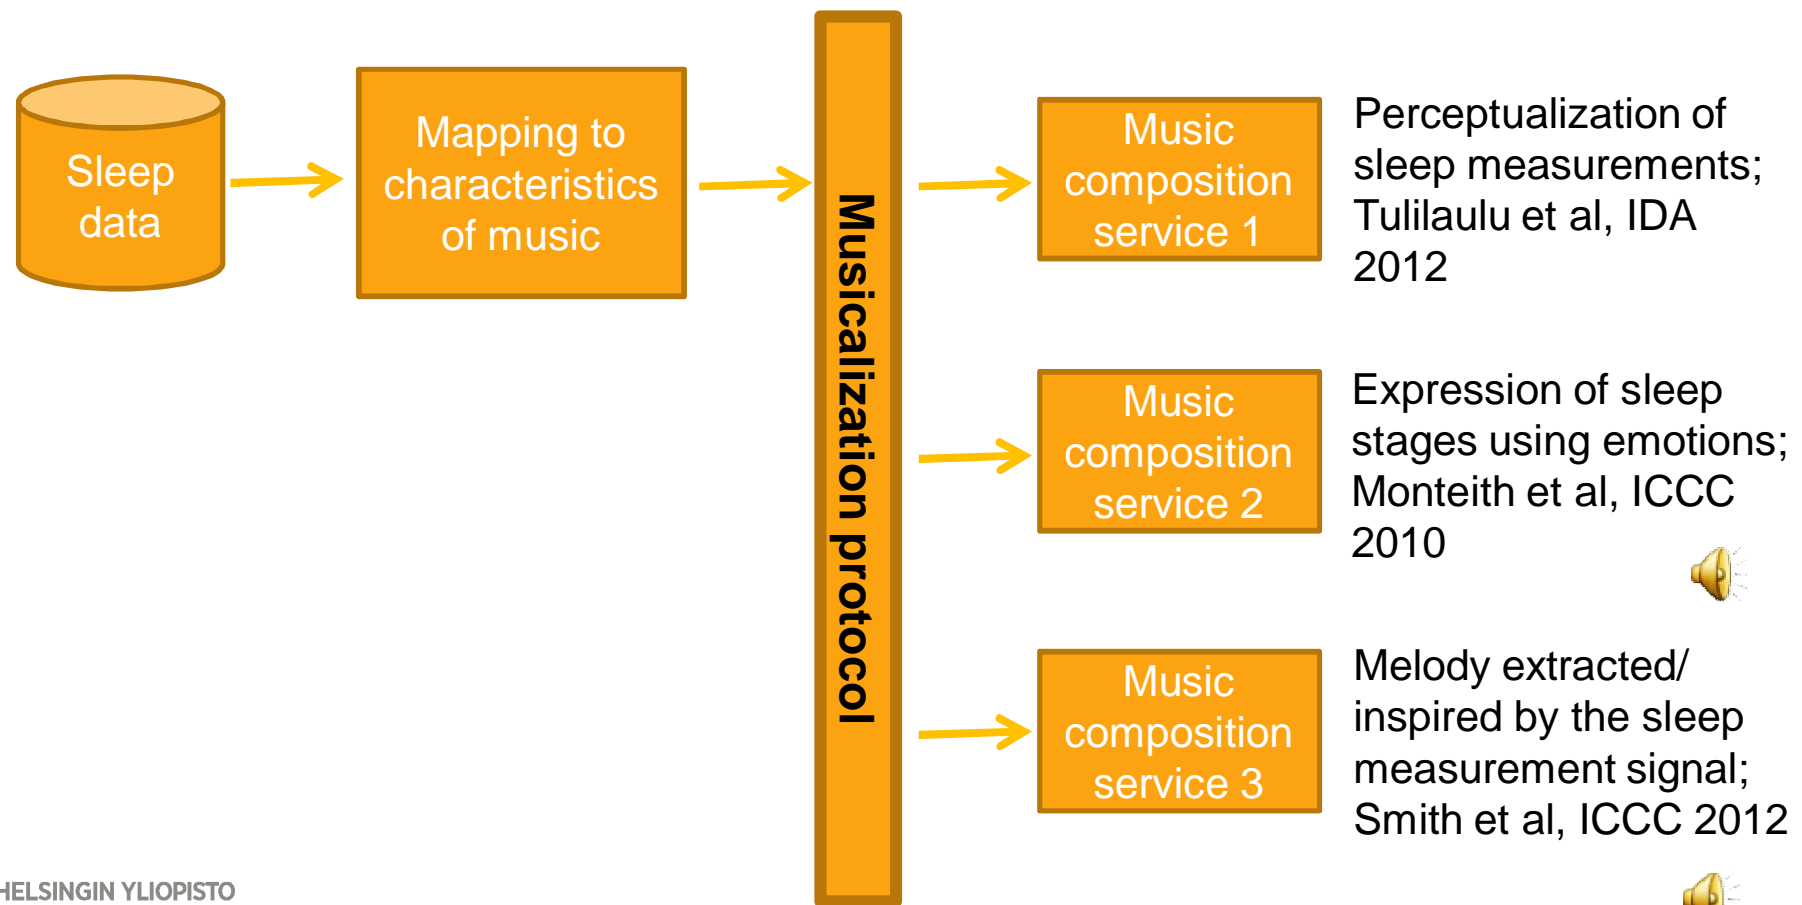


Modular musicalization?






Alternative sleep musicalizations

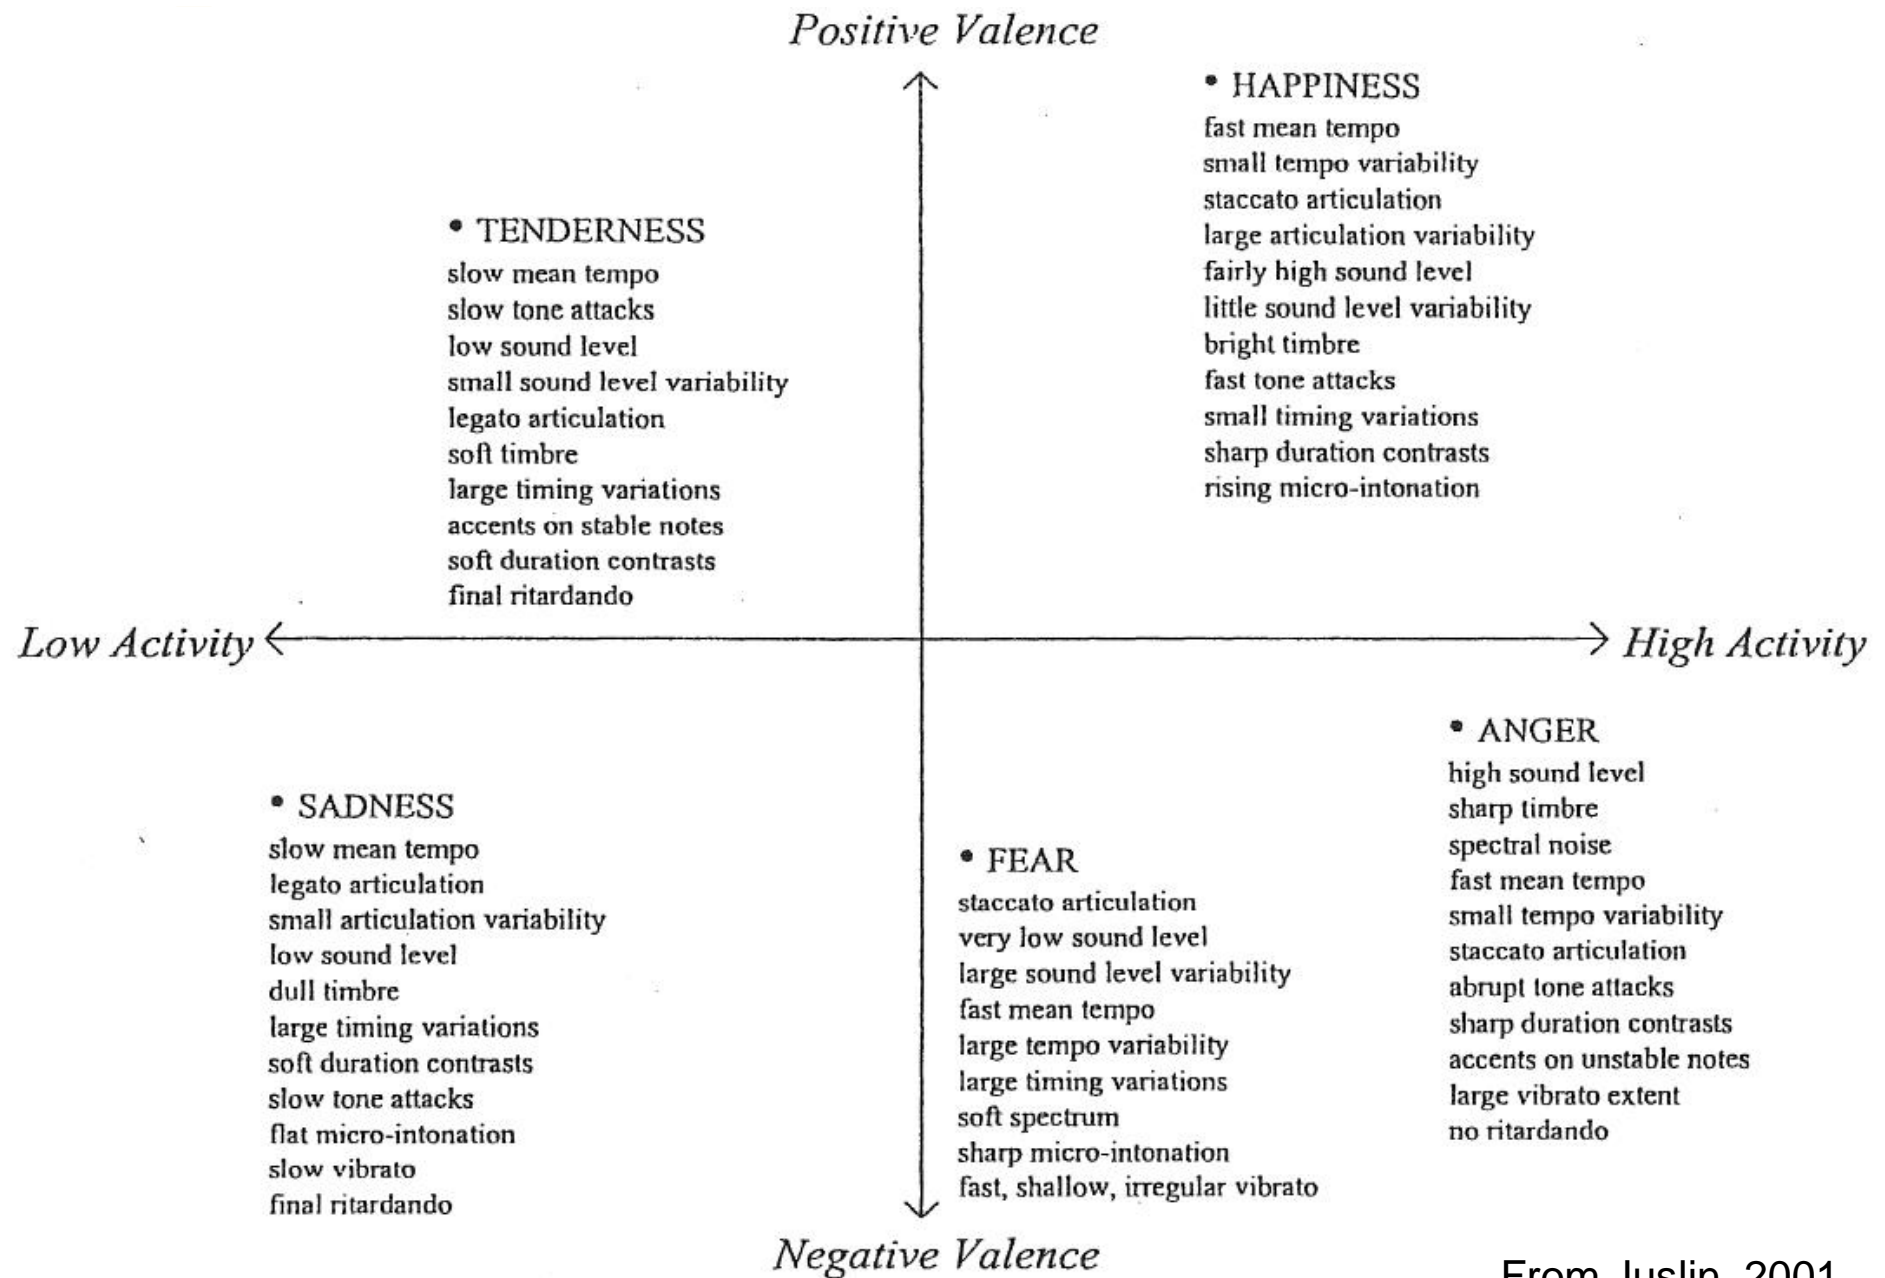




Future work in data musicalization



- Understand how subjective experiences can be affected by automatically composed music
- Develop principles for mapping properties of data to intended experiences
- Musicalization of patterns and other DM results



From Juslin, 2001



Applicability of data musicalization?

Where could data musicalization potentially be used?

The following characteristics may help:

- Personal data
 - Affectivity + creativity factors
- Data with a temporal dimension
 - Natural mapping to the temporal dimension of music