





Discrete Musical Systems

Shlomo Dubnov



Methods of Proposed Music Systems Theory

Combination of DES and Machine Learning

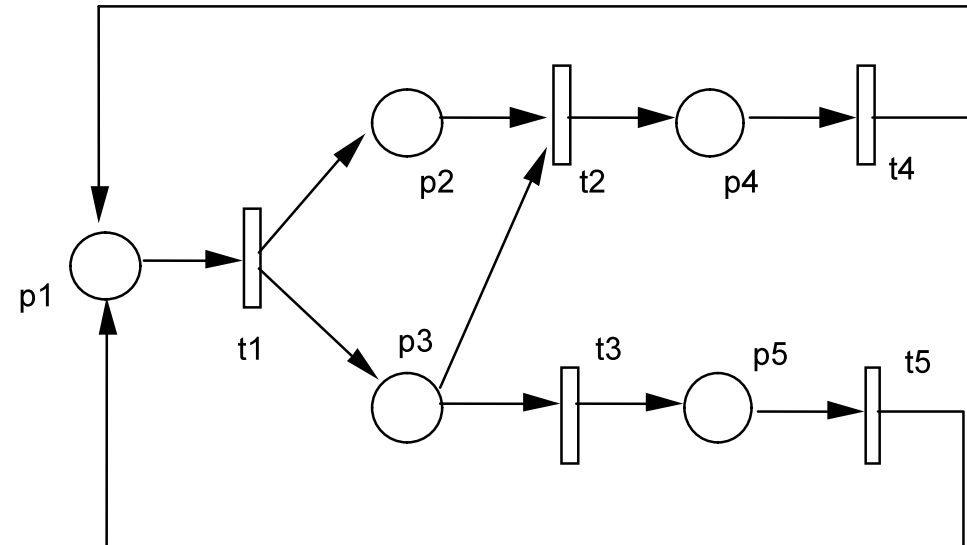
- DES:
 - Finite State Automata
 - Markov Processes
 - Petri Nets
- Machine Learning:
 - Learning FSA
 - Time Series Data mining
 - Reinforcement Learning *
 - Structure and Concurrency Modeling

Petri Nets and Music Structure

Show how to learn and control a large scale musical form using audio segmentation and Petri Nets

Places and transitions

- A **PETRI NET** is a bipartite graph which consists of two types of nodes: **places** and **transitions** connected by directed arcs.
- Place = circle, transition = bar or box.
- An arc connects a place to a transition or a transition to a place.
- No arcs between nodes of the same type.
- **Input and output places** of a transition
- **Input and output transitions** of a place



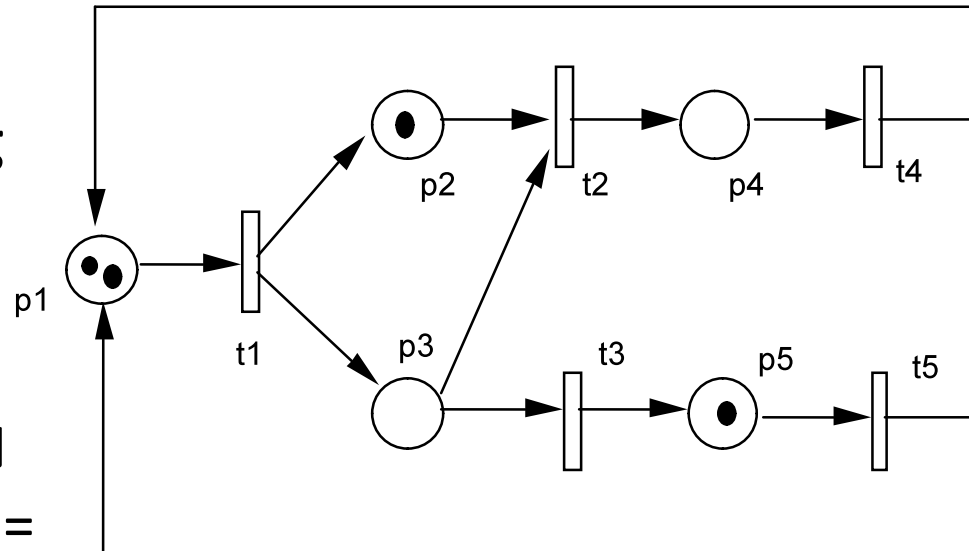
Token and marking system state

Each place **p_i** contains a number of **tokens**.

The distribution of tokens in the Petri net is called **marking** $M = (m_1, m_2, \dots, m_n)$ where $m_i = \#$ of tokens in place **p_i**

System State = marking of PN


The initial state of the system = initial marking M_0 .



System dynamics by transition firing

- A **transition** is said **enabled** (firable) if each of its input places contains at least one token. An enabled transition can fire.
- **Firing a transition** removes a token from each input place and add one token to each output place.
- Firing a transition leads to a new marking that enables other transitions.
- The dynamic behavior of the corresponding system = evolution of the marking and transition firings
- Convention: **simultaneous transition firings are forbidden.**

Firing Example



Formal definitions

Petri Nets

A Petri net is a five-tuple $PN = (P, T, A, W, M_0)$ where:

$P = \{ p_1, p_2, \dots, p_n \}$ is a finite set of places

$T = \{ t_1, t_2, \dots, t_m \}$ is a finite set of transitions

$A \subseteq (P \times T) \cup (T \times P)$ is a set of arcs

$W : A \rightarrow \{ 1, 2, \dots \}$ is a weight function

$M_0 : P \rightarrow \{ 0, 1, 2, \dots \}$ is the initial marking

$P \cap T = \Phi$ and $P \cap T = \Phi$

PN without the initial marking is denoted by N:

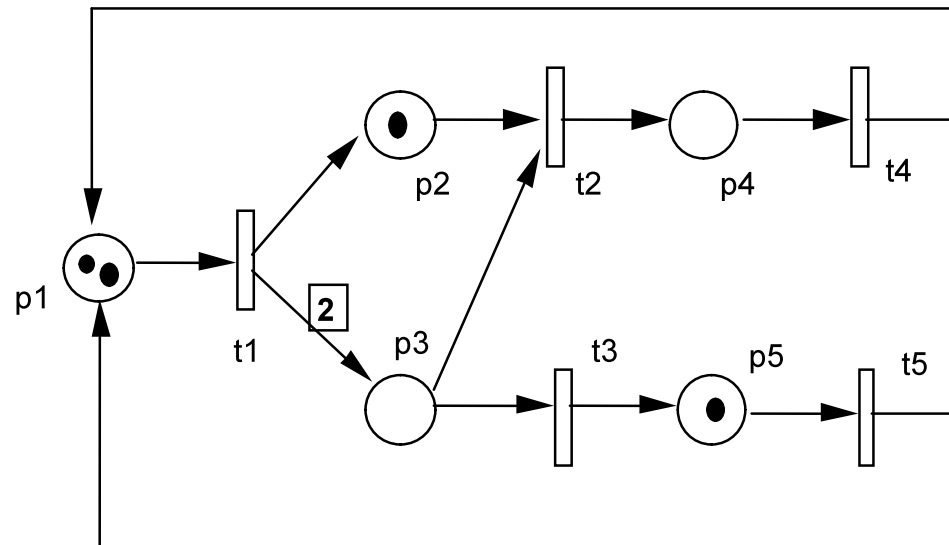
$N = (P, T, A, W)$

$PN = (N, M_0)$

A Petri net is said **ordinary** if $w(a) = 1, \forall a \in A$.

Graphic representation

Similar to that of ordinary PN but with default weight of 1 when not explicitly represented.



Transition firing

Rule 1: A **transition** t is **enabled** at a marking M if $M(p) \geq w(p, t)$ for any $p \in {}^{\circ}t$ where ${}^{\circ}t$ is the set of input places of t


Rule 2: An enabled transition may or may not fire.

Rule 3: **Firing transition** t results in:

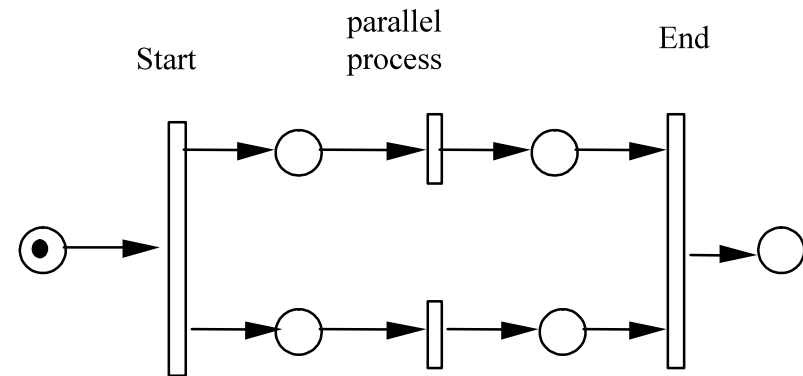
- removing $w(p, t)$ tokens from each $p \in {}^{\circ}t$
- adding $w(t, p)$ tokens to each $p \in t^{\circ}$ where t° is the set of output places of t

PN models of key characteristics


Precedence relation:



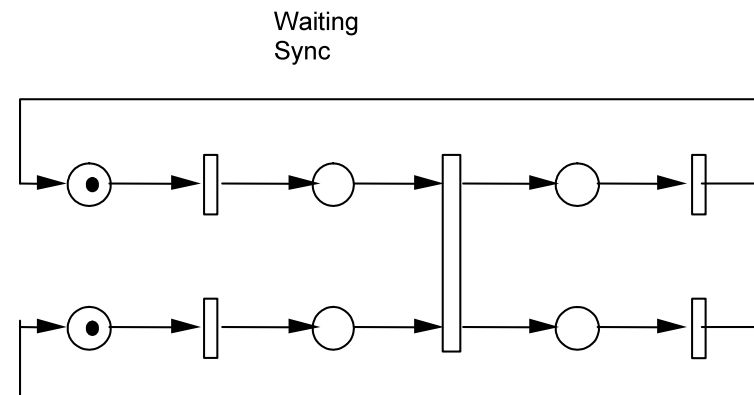
Parallel processes:



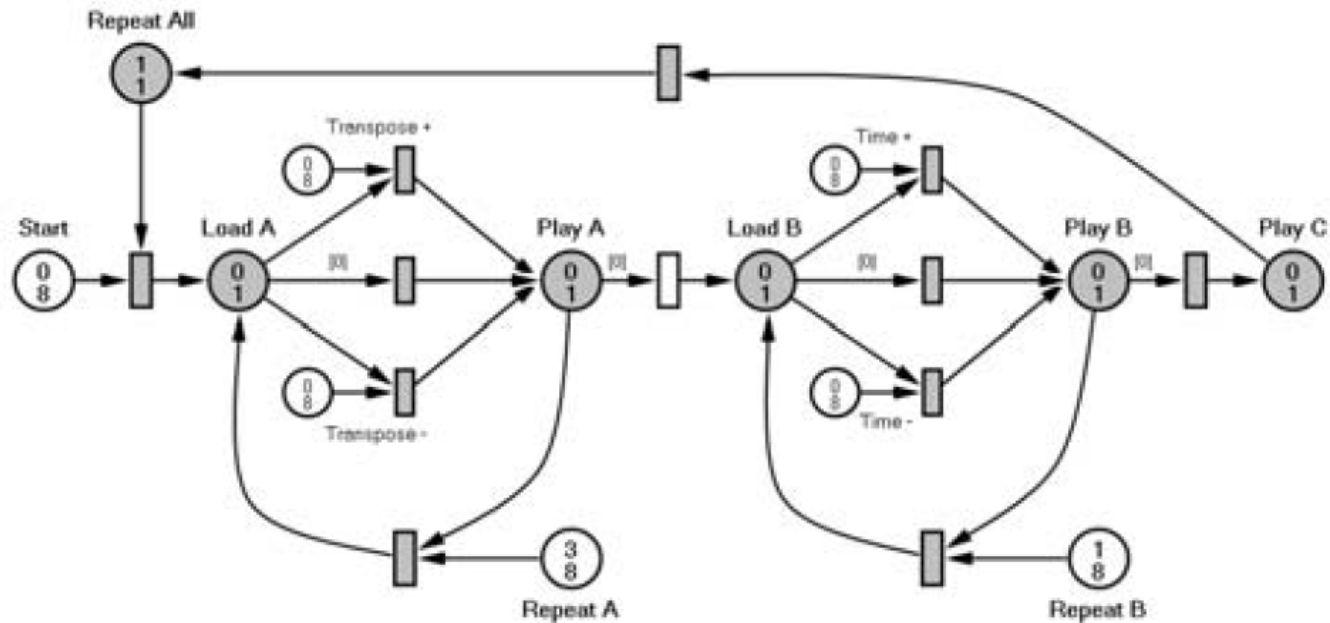
Alternative processes:



Synchronization:



Music Modeling with PN



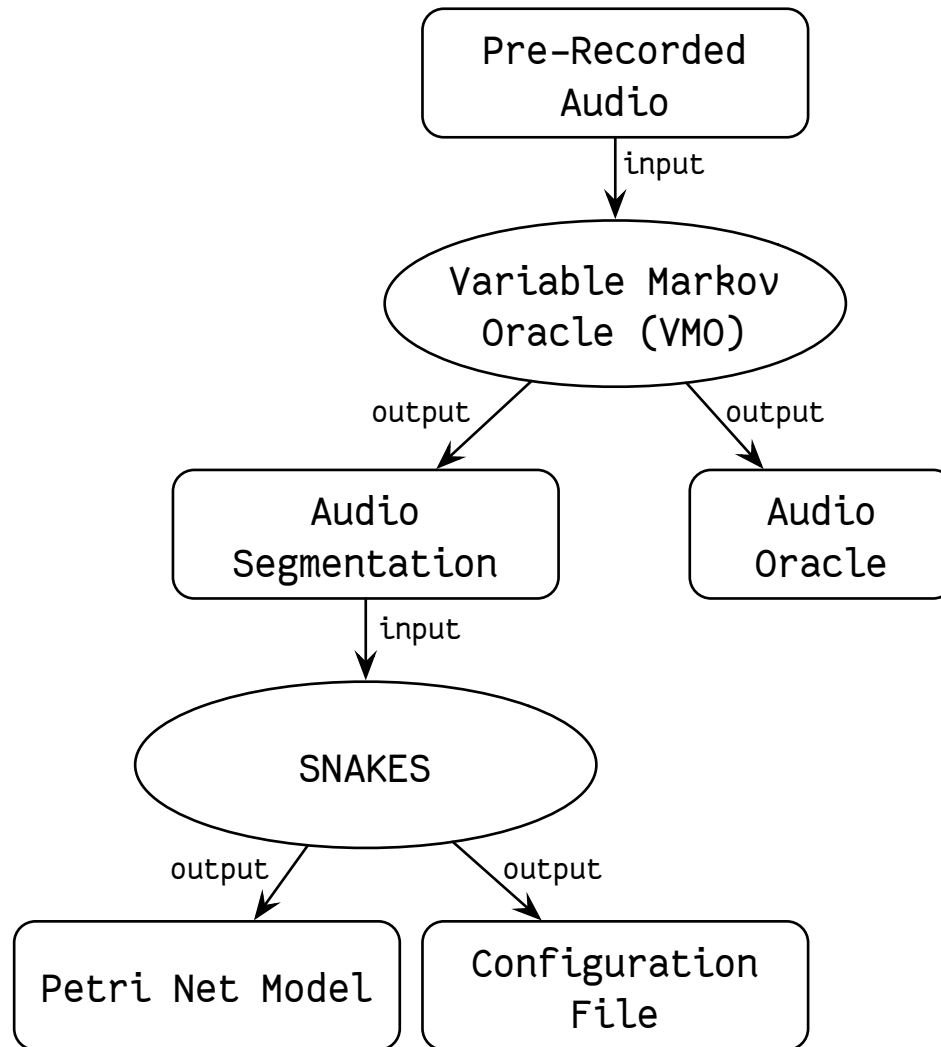
Frank Zappa "Peaches and Regalia" analysis by A. Baratè

Automatic Modeling of Musical Structure using PN

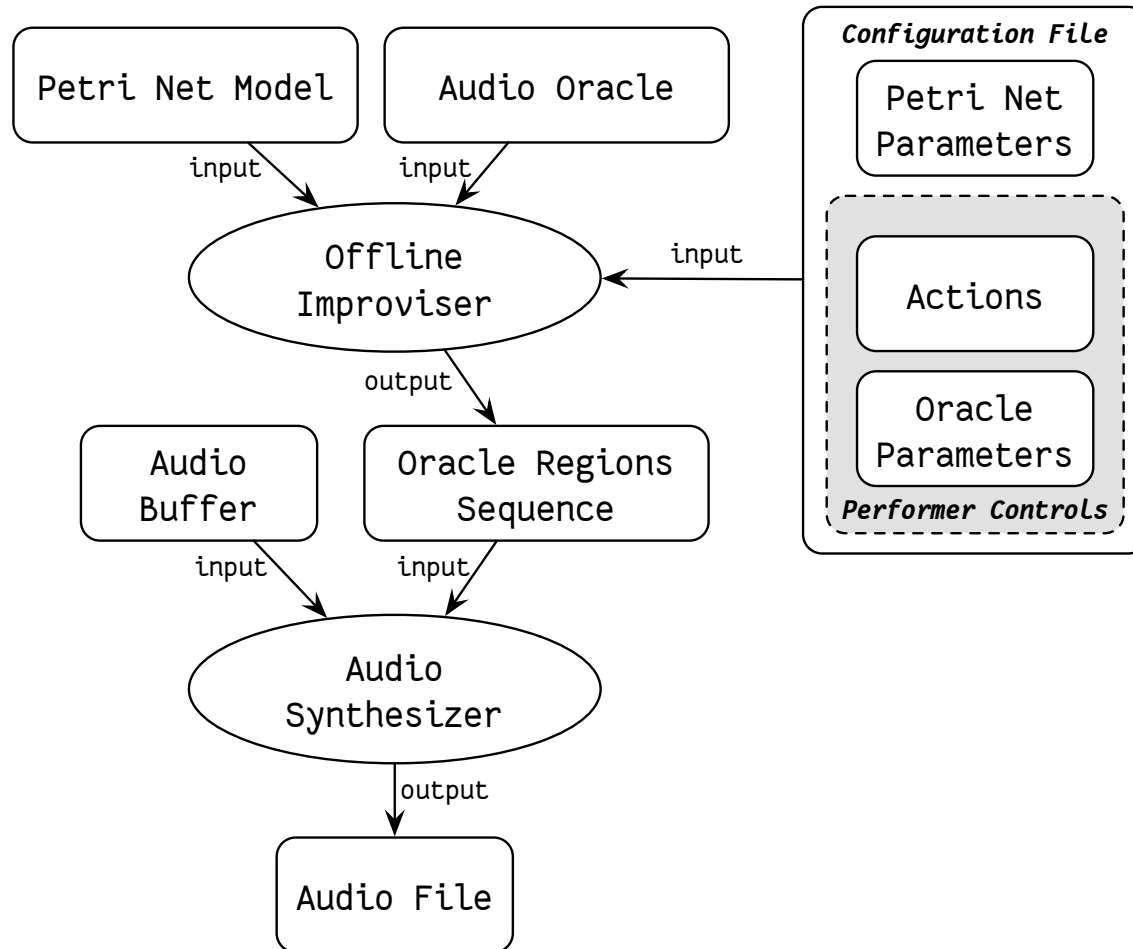
Overview of the system

- PyOracle: https://gitlab.com/himito/PyOracle_I-score
- i-score: <https://github.com/himito/i-score>
- VMO-Score: https://himito.github.io/vmo_i-score_generator

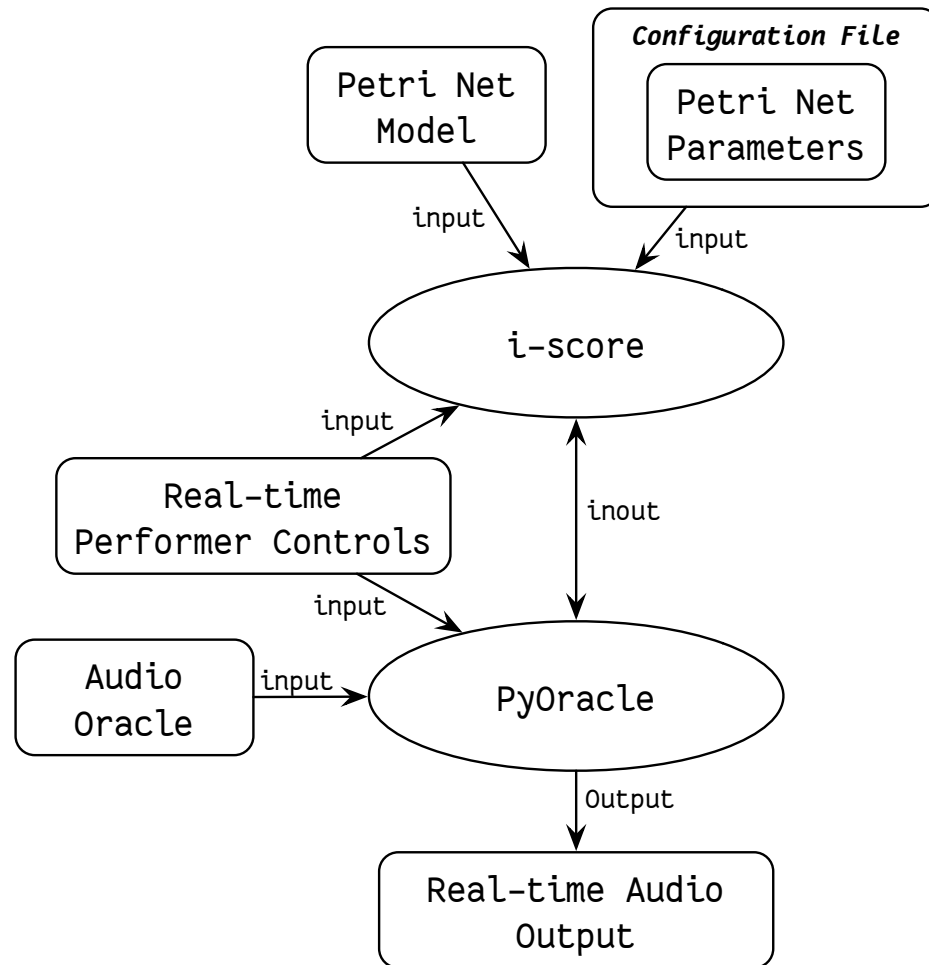
Overview of the system




Offline Improvisation



Real Time Improviser




Segmentation




Composition with PN

```
1  # file: configuration.yml
2
3  conditions:
4    - transition : 't0'
5      time-min   : 0.0
6      time-max   : 3.0
7      condition  : '/device/key == 10'
8
9    - transition : 't1'
10     ...
```

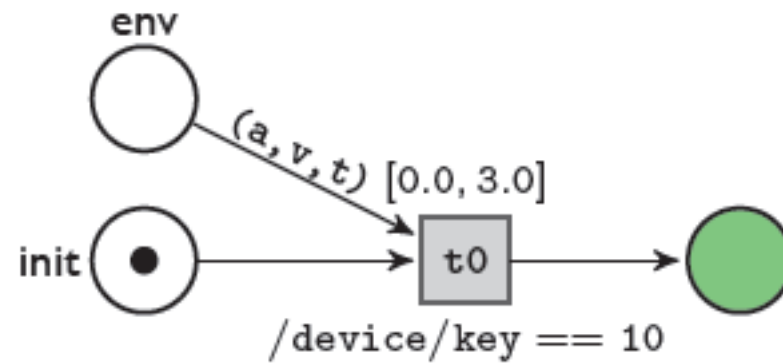
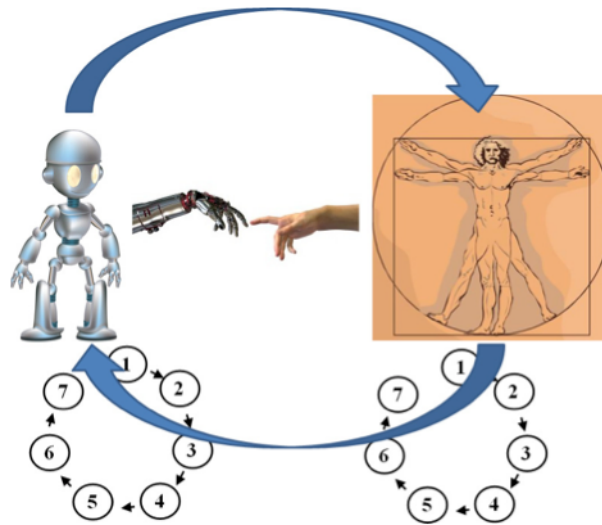


Improvisation with PN

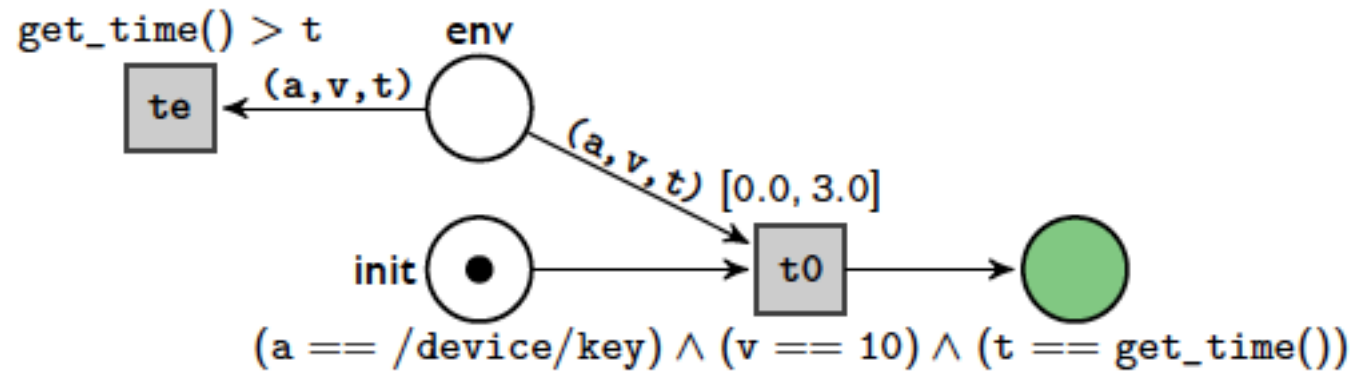
```
1  # file: configuration.yml
2
3  actions:
4    - address : '/volume/sensor/pos_x'
5      value   : 10
6      time    : 250
7
8    - address : ...
```



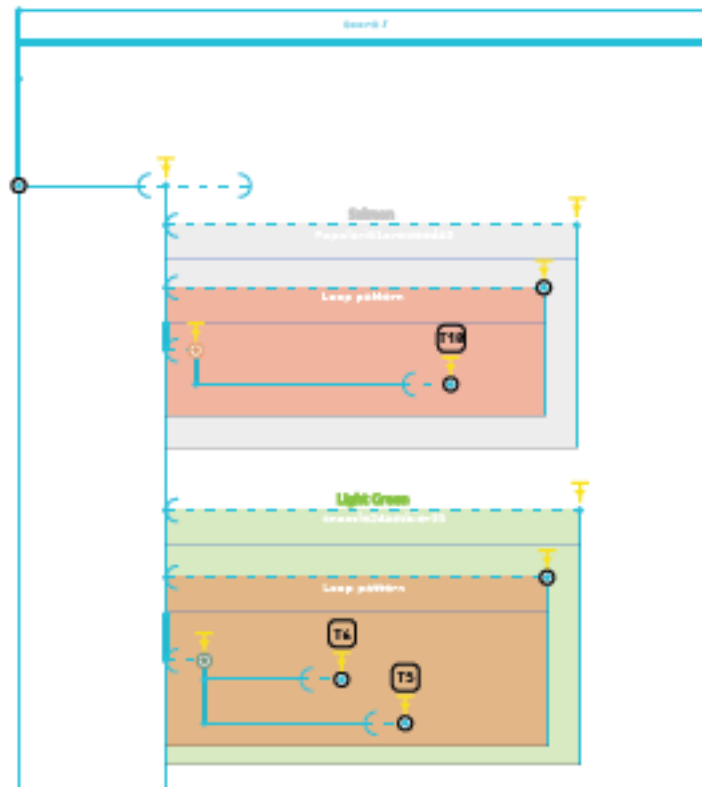
Modeling Interaction




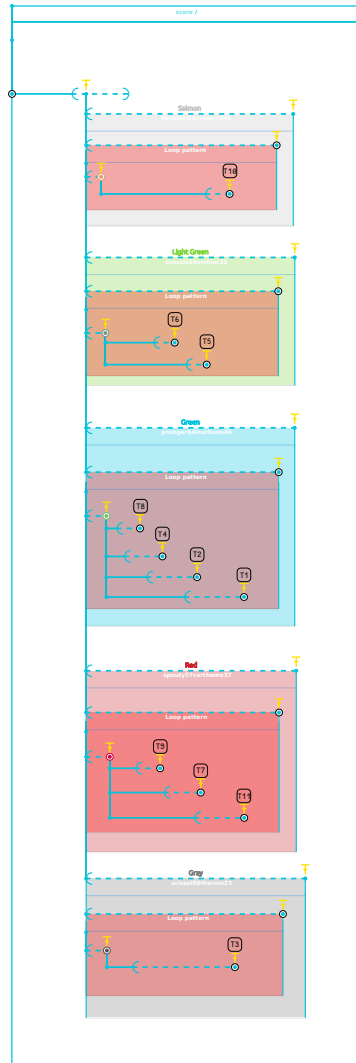
Adding Environment



i-score representation





PN to i-score mapping




Segmentation

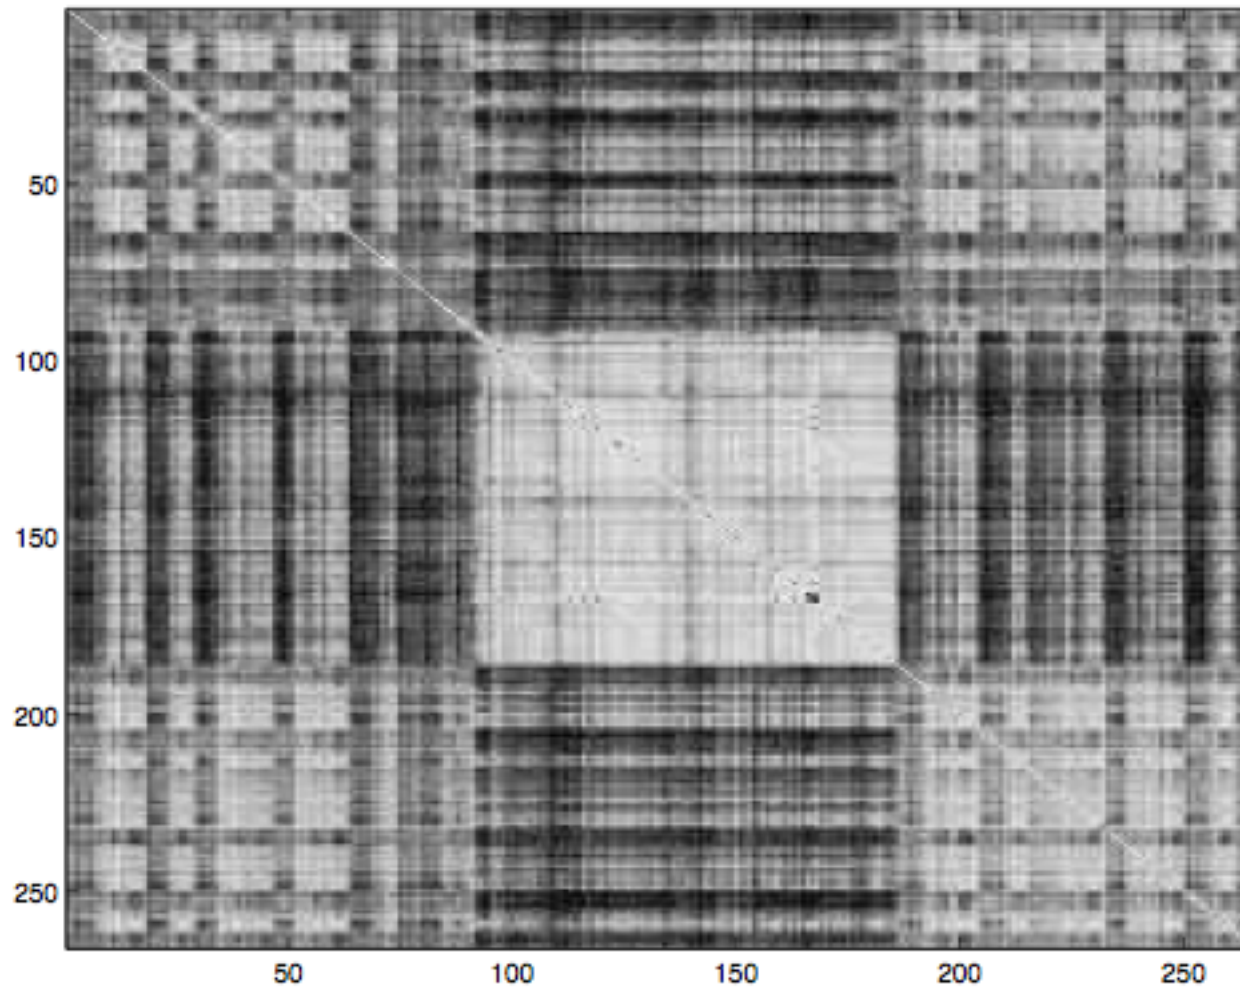
Using self-similarity (recurrence) to partition an audio file into inter-connected regions by spectral clustering



Time (s)



Recurrence Matrix



$$d(i, j) = \frac{\langle X_i, X_j \rangle}{\|X_i\| \|X_j\|}.$$

SVD

$$\mathbf{X}_{n \times p} = \mathbf{U}_{n \times n} \mathbf{\Lambda}_{n \times p} \mathbf{V}_{p \times p}^T$$

where

$$\mathbf{U}^T \mathbf{U} = \mathbf{I}_{n \times n}, \quad \text{and} \quad \mathbf{V}^T \mathbf{V} = \mathbf{I}_{p \times p}.$$

SVD – Geometric Interpretation

‘spectral decomposition’ of the matrix

$$\begin{bmatrix} \mathbf{X} \end{bmatrix} = \begin{bmatrix} \mathbf{u}_1 & \mathbf{u}_2 \end{bmatrix} \cdot \begin{bmatrix} \lambda_1 & \emptyset \\ \emptyset & \lambda_2 \end{bmatrix} \cdot \begin{bmatrix} \text{---} & \mathbf{v}_1 & \text{---} \\ \text{---} & \mathbf{v}_2 & \text{---} \end{bmatrix}$$

SVD – Geometric Interpretation

‘spectral decomposition’ of the matrix:

$$\begin{array}{c} \leftarrow m \rightarrow \\ \left[\begin{array}{c} \mathbf{X} \end{array} \right] \\ \left. \begin{array}{l} \uparrow \\ n \\ \downarrow \end{array} \right\} \end{array} = \lambda_1 \mathbf{u}_1 \mathbf{v}_1^T + \lambda_2 \mathbf{u}_2 \mathbf{v}_2^T + \dots$$

SVD – Geometric Interpretation

‘spectral decomposition’ of the matrix:

The diagram illustrates the SVD decomposition of a matrix X . On the left, a large square bracket contains the matrix X . A vertical double-headed arrow to the left of the bracket is labeled n , indicating the number of rows. A horizontal double-headed arrow above the bracket is labeled m , indicating the number of columns. To the right of the matrix is an equals sign. To the right of the equals sign is a horizontal double-headed arrow labeled "r terms". Below this arrow is the expression $\lambda_1 \mathbf{u}_1 \mathbf{v}_1^T + \lambda_2 \mathbf{u}_2 \mathbf{v}_2^T + \dots$. Below the first term, an arrow points from the label $n \times 1$ to the vector \mathbf{u}_1 , and another arrow points from the label $1 \times m$ to the vector \mathbf{v}_1^T .

$$\begin{array}{c} \left[\begin{array}{c} \text{---} m \text{---} \\ \mathbf{X} \\ \text{---} m \text{---} \end{array} \right] \begin{array}{c} \longleftarrow \\ \longrightarrow \end{array} \\ \begin{array}{c} \updownarrow \\ n \end{array} \end{array} = \begin{array}{c} \longleftarrow \quad r \text{ terms} \quad \longrightarrow \\ \lambda_1 \mathbf{u}_1 \mathbf{v}_1^T + \lambda_2 \mathbf{u}_2 \mathbf{v}_2^T + \dots \\ \begin{array}{cc} \nearrow & \nwarrow \\ n \times 1 & 1 \times m \end{array} \end{array}$$

SVD – Geometric Interpretation

Approximation / dim. Reduction - by keeping the first few terms (how many?)

$$\begin{array}{c} \leftarrow m \rightarrow \\ \left[\begin{array}{c} \mathbf{X} \end{array} \right] \\ \left. \begin{array}{l} \uparrow \\ n \\ \downarrow \end{array} \right\} \end{array} = \lambda_1 \mathbf{u}_1 \mathbf{v}_1^T + \lambda_2 \mathbf{u}_2 \mathbf{v}_2^T + \dots$$

assume: $\lambda_1 \geq \lambda_2 \geq \dots$

$$\mathbf{X}_{n \times p} = \mathbf{U}_{n \times n} \mathbf{\Lambda}_{n \times p} \mathbf{V}_{p \times p}^T$$

where

$$\mathbf{U}^T \mathbf{U} = \mathbf{I}_{n \times n}, \quad \text{and} \quad \mathbf{V}^T \mathbf{V} = \mathbf{I}_{p \times p}.$$

$$\mathbf{X} = \mathbf{U} \mathbf{S}_x$$

\mathbf{U} – orth. basis vectors

$$\mathbf{S}_x = \mathbf{\Lambda}_x \mathbf{V}^T$$

$$\mathbf{D} = \mathbf{S}_x^T \mathbf{S}_x.$$


Property I: Eigenvectors S_d of \mathbf{D} are transpose of the expansion coefficients of \mathbf{X} , $S_d = S_x^T$.

Relation between SVD and Self-Similarity based clustered

- Matrix D is a correlation matrix
- When self-similarity is computed using dot-product, D becomes a similarity matrix
- Spectral Clustering uses eigenvectors of a normalized similarity matrix to find objects in the data (to be explained next)
 - Side Note: the name “Spectral Clustering” comes from spectral decomposition of a matrix (eigenvectors) and has nothing to do with spectrum of the audio signal

Dimension Reduction / Audio Basis

Eigenvector Clustering



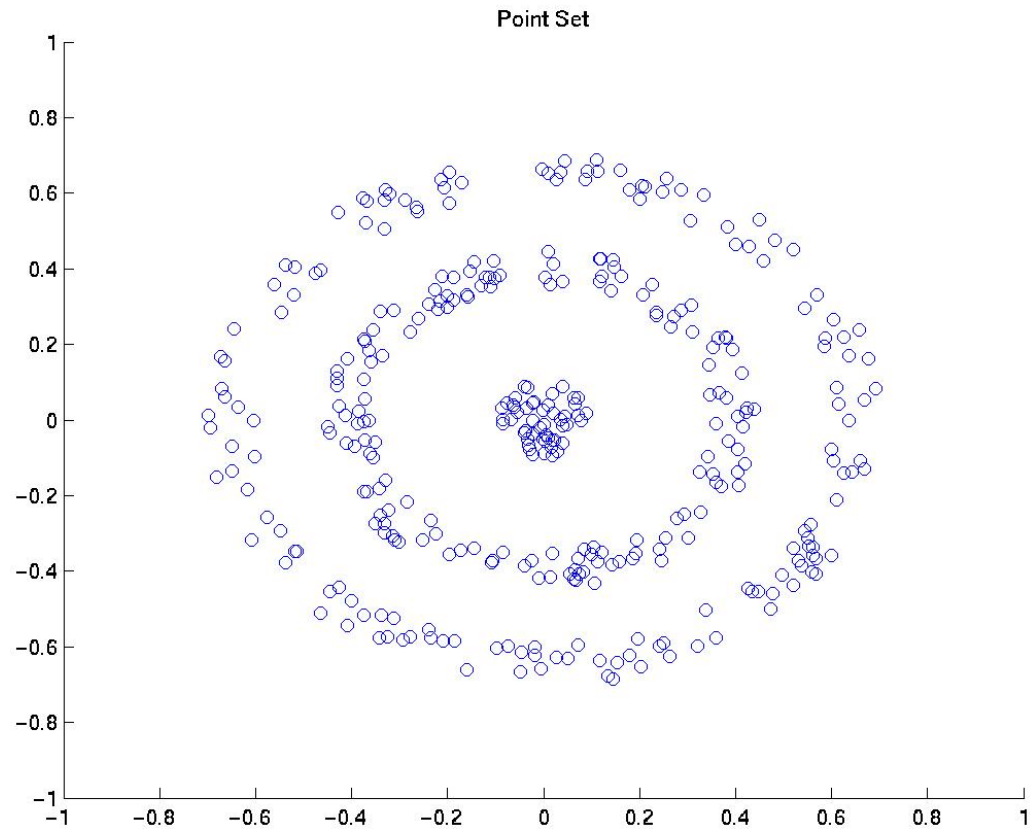
Spectral Clustering

- Transform the similarity matrix D to a stochastic matrix:

$$P = Z^{-1}D$$

- P_{ij} is the probability of moving from sound grain i to sound grain j in one step of a random walk
- Same eigenvectors: $y = yP$; eigenvalues: $\lambda = 1 - \lambda P$


Synthetic Example I




Note: Usual distance based clustering does not work for this type of data.

Distance Matrix


Euclidean Distance Matrix




The second generalized eigenvector




The first partition



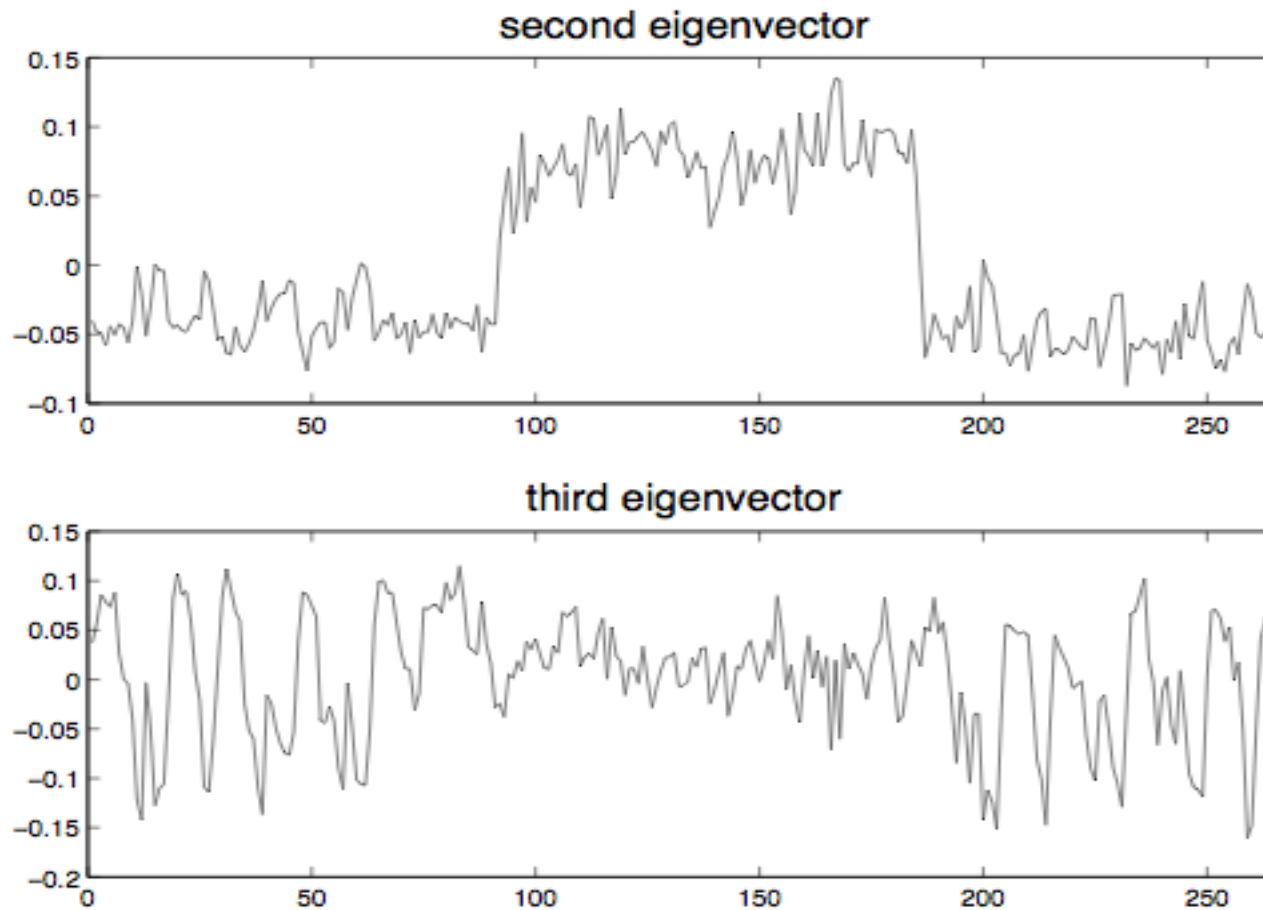
The second generalized eigenvector




The second partition




Example II: Audio Segmentation using eigenvectors



second and third eigenvector clusters



Example

Original 

Segmented




Example III:


The Angel of Death by Roger Reynolds

for piano, chamber orchestra and computer-processed sound


S-D VERSION




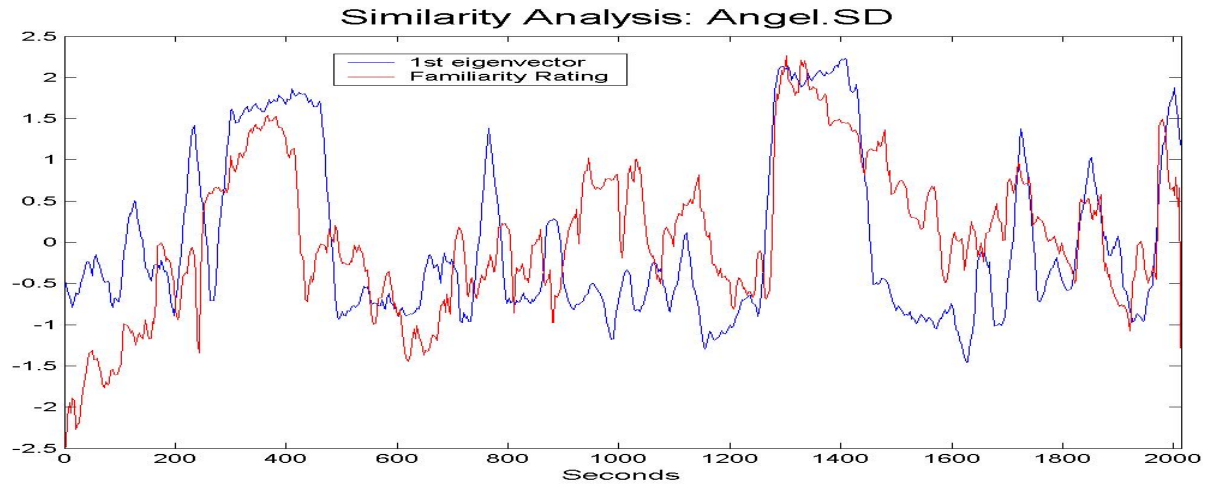
D-S VERSION



Familiarity vs. Recurrence

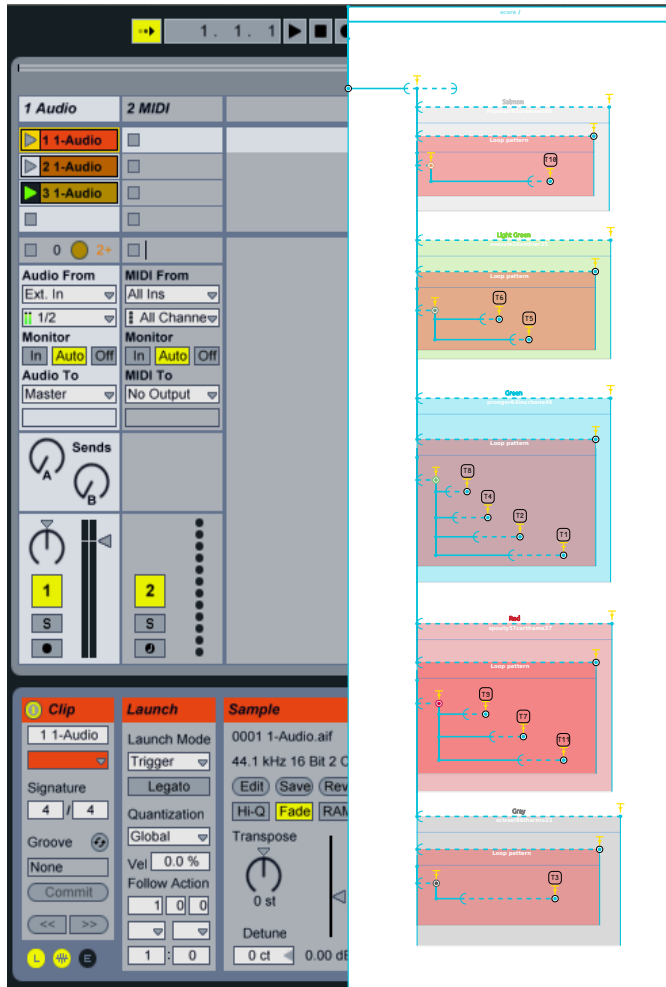


Matrix Eigenvector Profile



- **Segmentation** based on **Recurrence** (Similarity Matrix) can be done directly from the **first eigenvectors**
 - after proper normalization
 - automatic ways to find thresholds could be devised (such as k-means method in ICMC 2004 paper)
- If your **distance** function can be expressed as **dot-product**, then you can **use SVD** directly on the **data**
- segmentation is non-linear with respect to the original data space
- Spectral Clustering can be applied to graph derived from VMO analysis (not covered here – see references)

Summary: New type of DAW



- Automatic Clips Cut
- Improvise
- Follow conditions and actions
- Side-chaining query

References

- Foote, J. and M. Cooper (2001). Visualizing musical structure and rhythm via self-similarity. In Proceedings of the ICMC, pp. 419–422. ICMA.
- Lu, L., S. Li, L. Wenyin, and H. Zhang (2002). Audio textures. International Conference on Acoustics, Speech, and Signal Processing, 1761–1764.
- Shi, J. and J. Malik (2000). Normalized cuts and image segmentation. IEEE Trans. Pattern Anal. Mach. Intell. 22(8), 888–905.
- Dubnov, S. and Apel, T. (2004). Audio Segmentation by Singular Value Clustering, ICMC
- Wang, C. and Mysore, G.J (2016), Structural segmentation with Variable Markov Oracle and boundary adjustment, ICASSP
- Arias, J., Desainte-Catherine M. and Dubnov, S. (2016) Automatic Construction of Interactive Machine Improvisation Scenarios from Audio Recordings, International Workshop on Musical Metacreation, AAAI Conference on Computational Creativity

The End