



MIRtoolbox:

Sound and music analysis of audio recordings using Matlab


MUS483 I, Olivier Lartillot
Part II, 9.11.2017

Part 2

- Rhythm, metrical structure
- Tonal analysis
- Segmentation, structure
- Statistics
- Music & emotion
- Advanced use

mirevents

event detection function



- `mirevents('ragtime.wav')`

- `mirpeaks(mirsum(mirspectrum(..., 'Frame')))`

- `mirevents('ragtime.wav', Filter)`

- `mirpeaks(mirsum(mirenvelope(mirfilterbank(..., 'NbChannels', 40))))`

- `mirevents('ragtime.wav', SpectralFlux)`

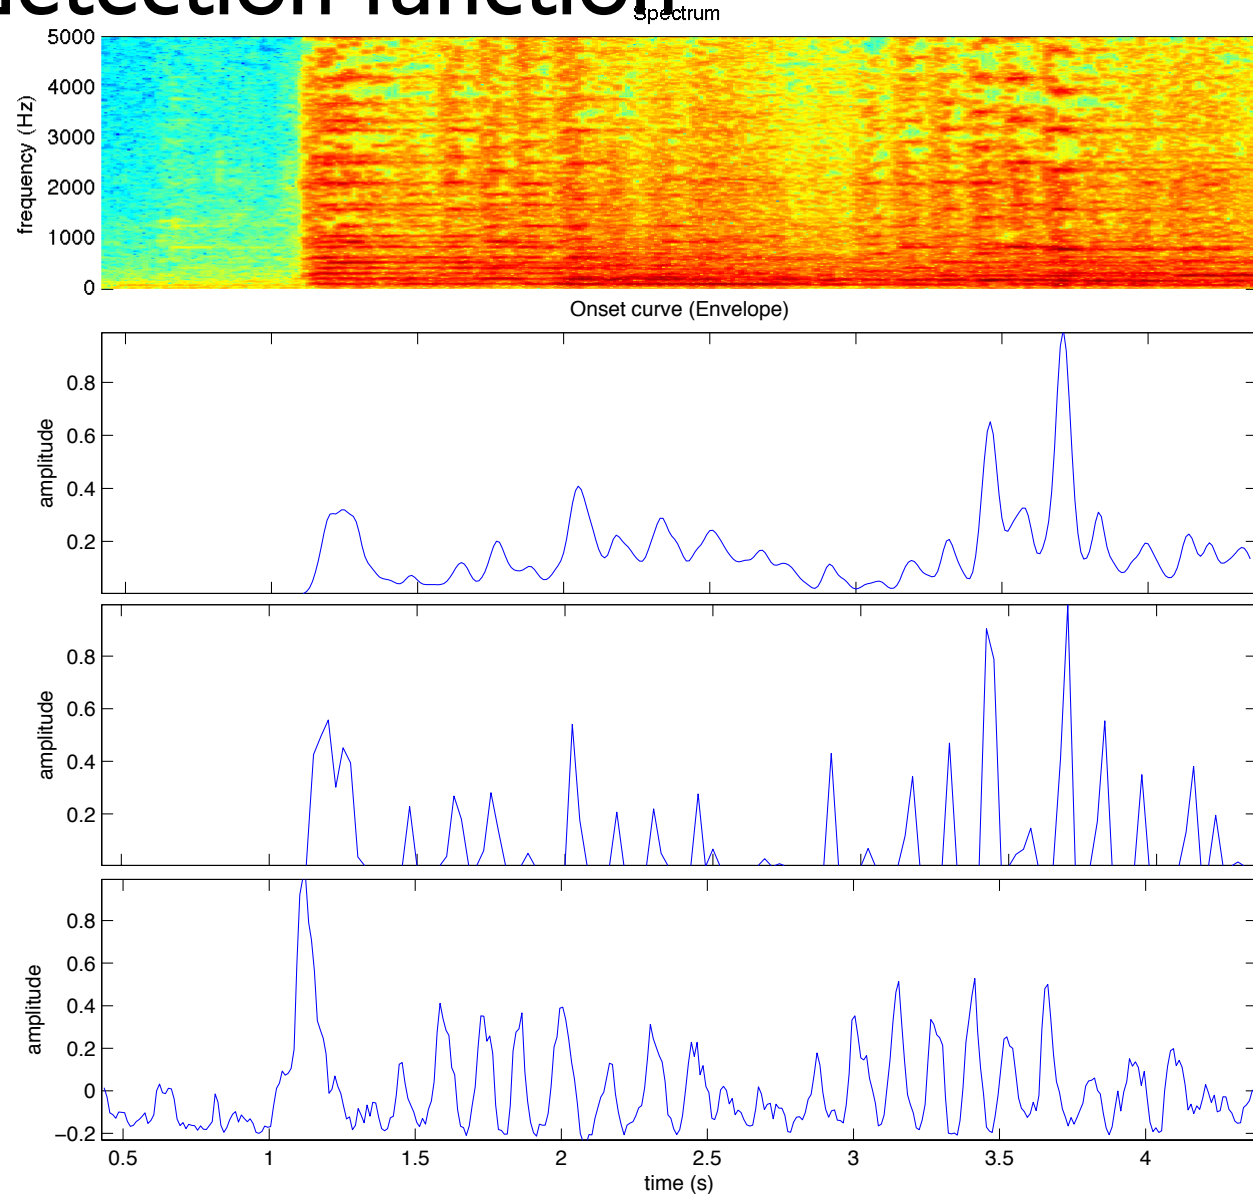
- `mirpeaks(mirflux(..., 'Inc', 'Halfwave'))`

`mirevents(...,
Contrast,
...)`

mirevents

event detection function


- ‘*Envelope*’, ‘*Filter*’: changes in dynamics
- ‘*SpectralFlux*’: global spectral changes
- ‘*Emerge*’: local changes in particular frequency regions



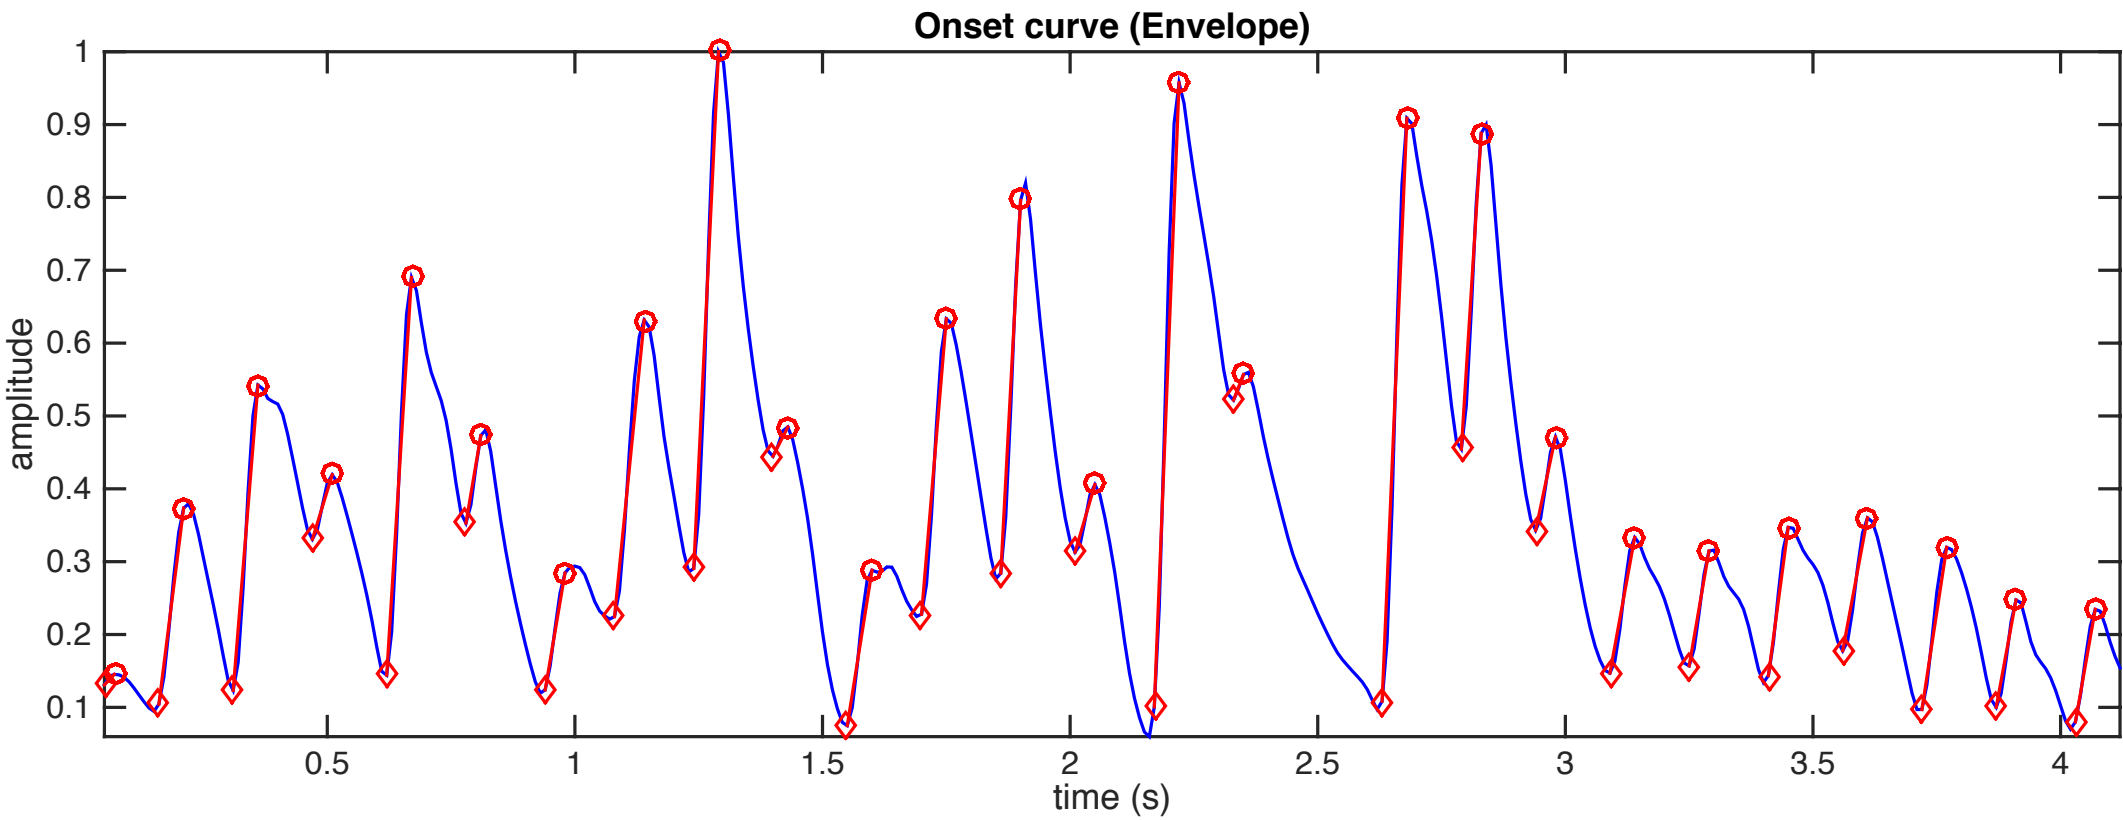
mirevents

event detection function

- ‘*Envelope*’, ‘*Filter*’: changes in dynamics
- ‘*SpectralFlux*’: global spectral changes
- ‘*Emerge*’: local changes in particular frequency regions




mirevents(..., 'Attack')



mirattackslope


average slope of note attacks



- $o = \text{mirevents}(\text{'ragtime.wav'}, \text{'Attacks'})$
- $\text{mirattackslope}(o)$

mirattackleap

amplitude of note attacks



- $o = \text{mirevents}(\text{'ragtime.wav'}, \text{'Attacks'})$
- $\text{mirattackleap}(o)$

Tempo estimation?

- 'george.wav'
- What tempo in BPM? You can tap on the beat while listening:
 - <http://www.all8.com/tools/bpm.htm>
- How to estimate tempo using the *MIRtoolbox* operators presented last week?

mirtempo

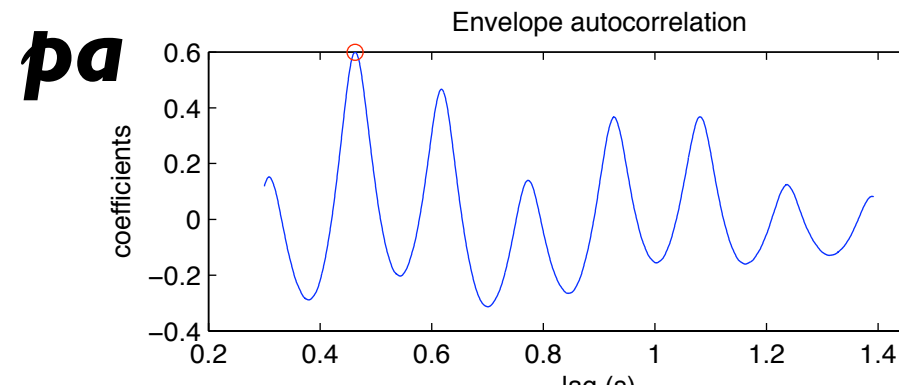

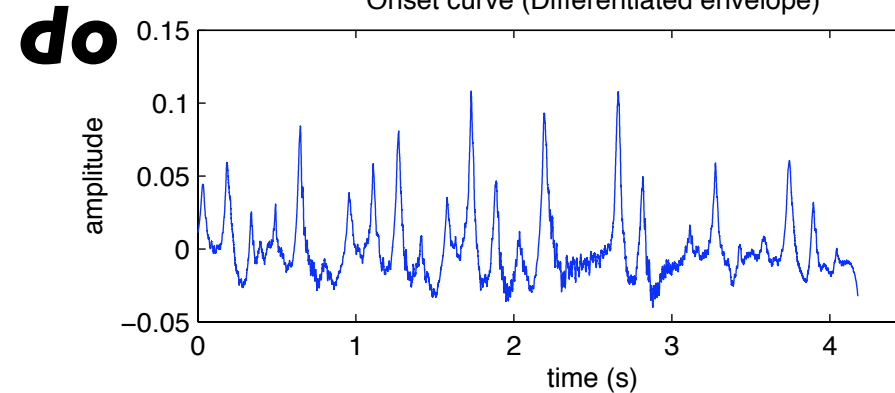
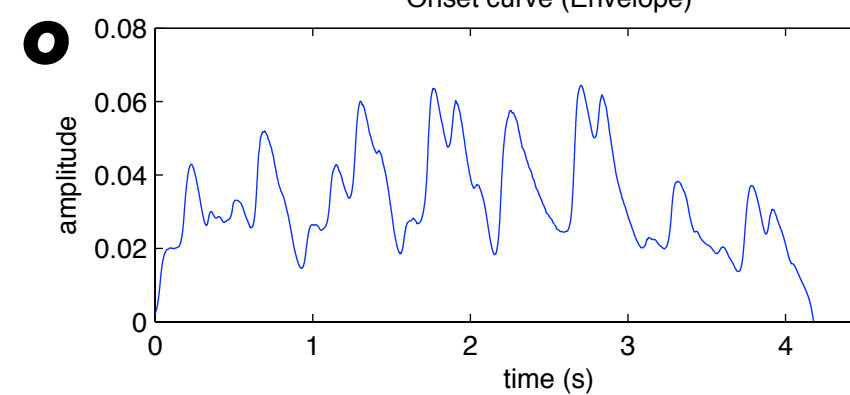
tempo (in beats per minute)

Roughly:

- **o** = *mirevents*('file.wav', 'Filter')
- **do** = *mirevents*(**o**, 'Diff')
- **ac** = *mirautocor*(**do**, 'Resonance')
- **pa** = *mirpeaks*(**ac**, 'Total', 1)

In short:


- **[t, pa]** = *mirtempo*('file.wav')



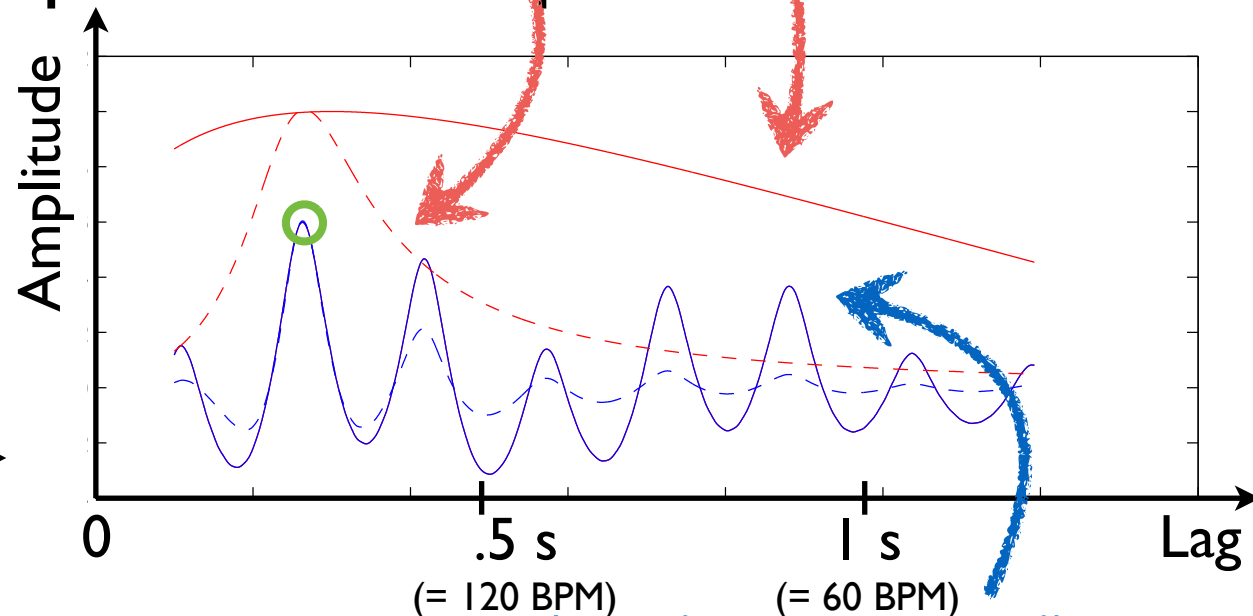
mirtempo

resonance curves

- *mirautocor*(..., **‘Resonance’**, **‘Toiviainen’**) (Toiviainen & Snyder, 2003)
- *mirautocor*(..., **‘Resonance’**, **‘vanNoorden’**) (van Noorden & Moelants, 2001)
- Emphasis on the best perceived tempi



mirautocor



mirautocor(..., **‘Resonance’**)

mirtempo

tempo estimation

- How to estimate tempo for such audio excerpt?
 - *'czardas.wav'*

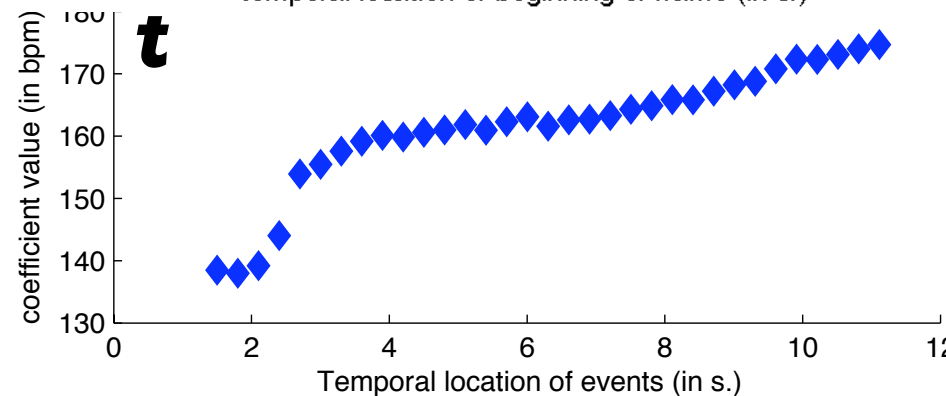
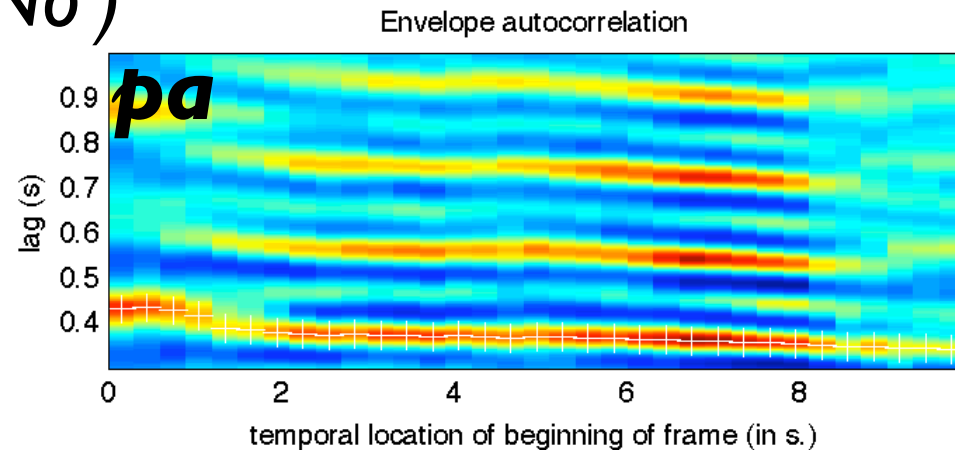
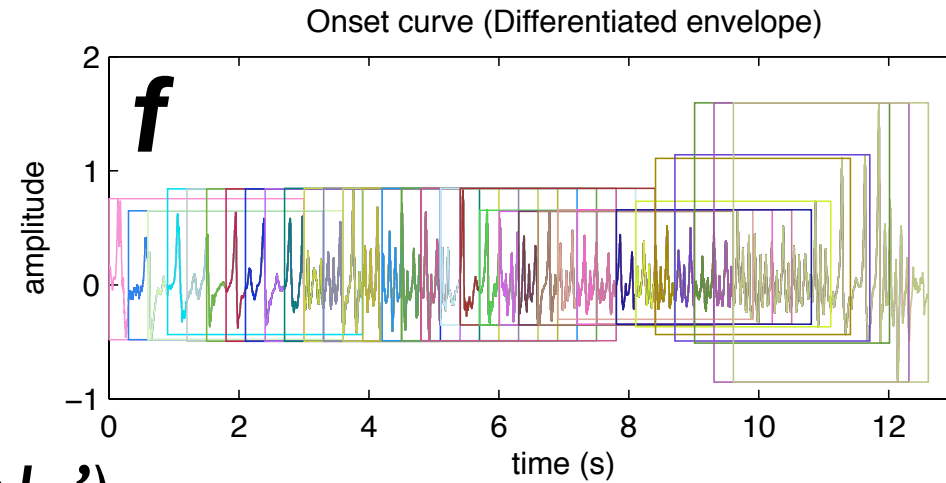
mirtempo

tempo (temporal evolution)

- **o** = *mirevents*('mysong', 'Detect', 'No')
- **do** = *mirevents*(o, 'Diff')
- **f** = *mirframe*(do)
- **ac** = *mirautocor*(f, 'Resonance')
- **pa** = *mirpeaks*(ac, 'Total', 1)


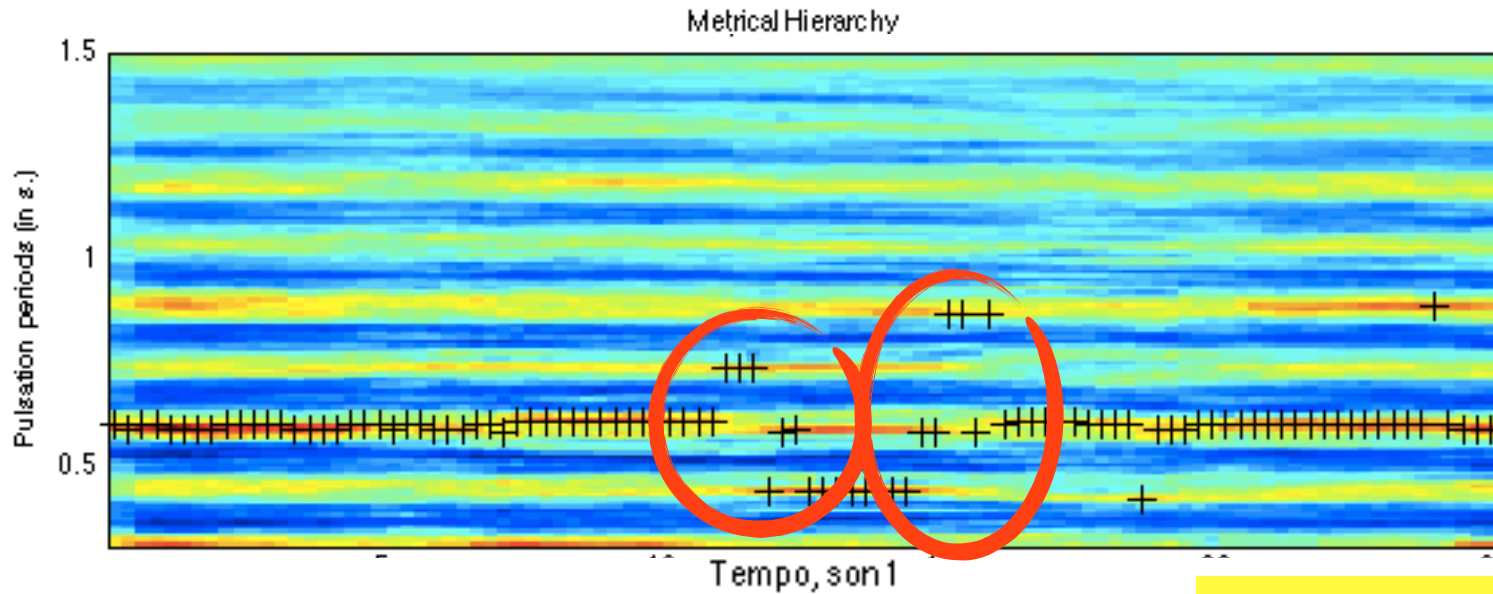
In short:

- **[t, pa]** = *mirtempo*('mysong', 'Frame')



Try for instance: `mirtempo('george.wav', 'Frame')`

mirtempo



Switch from one metrical level to another!

Try for instance: `mirmetre('george.wav')`


mirmetre

tracking all metrical levels

Metrical levels:

- Level N 's period is N times level 1's period.

(dominant)



mirmetre

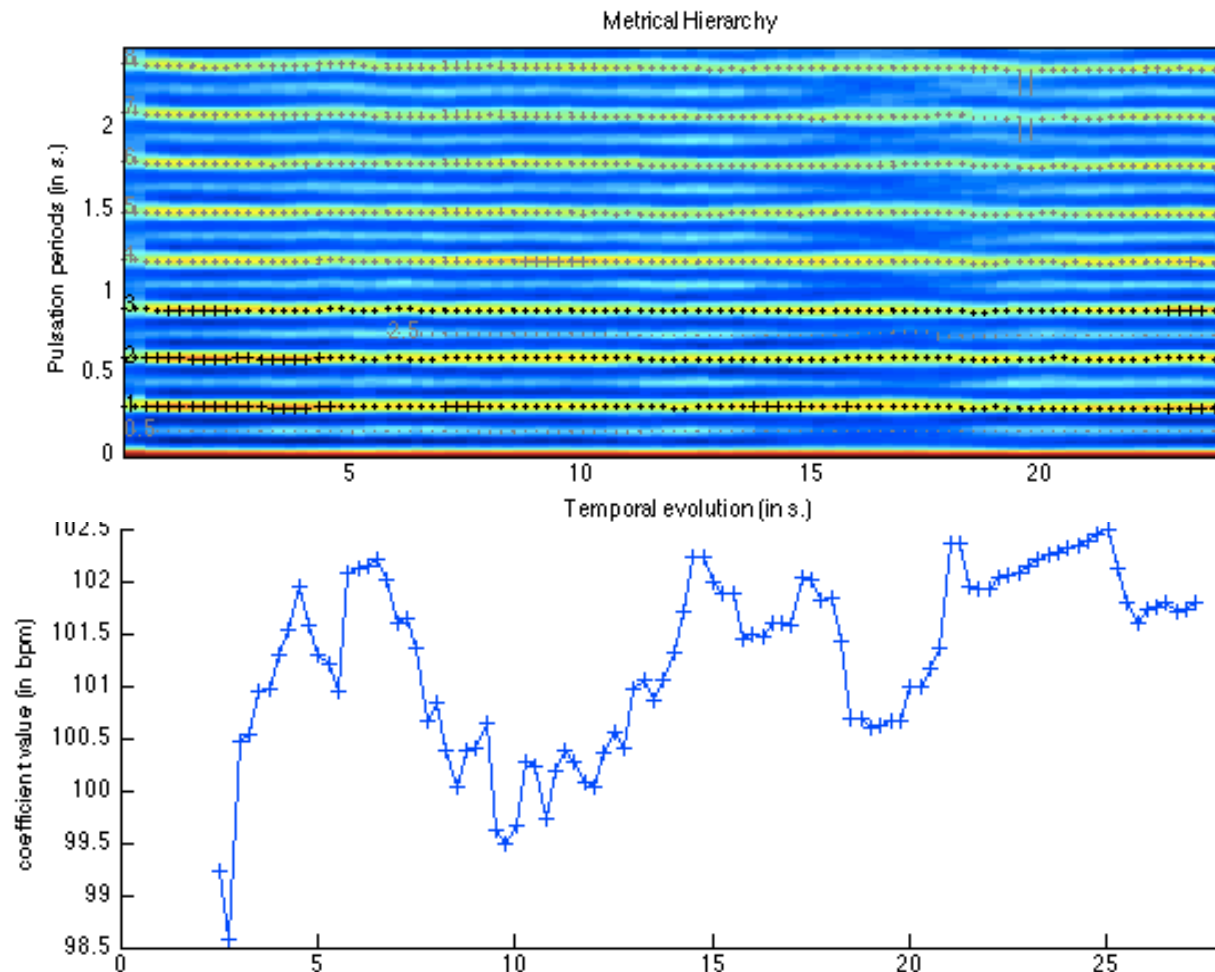
tracking all metrical levels

- $m = \text{mirmetre}(\dots)$
- $t = \text{mirtempo}(m)$

or:

- $[t \ m] = \text{mirtempo}(\dots, \text{'Metre'})$

- The tempo curve is associated to one particular metrical level.




Influence of the
onset detection
method


mirmetre

tracking all metrical levels

‘Spectral
Flux’




‘Emerge’



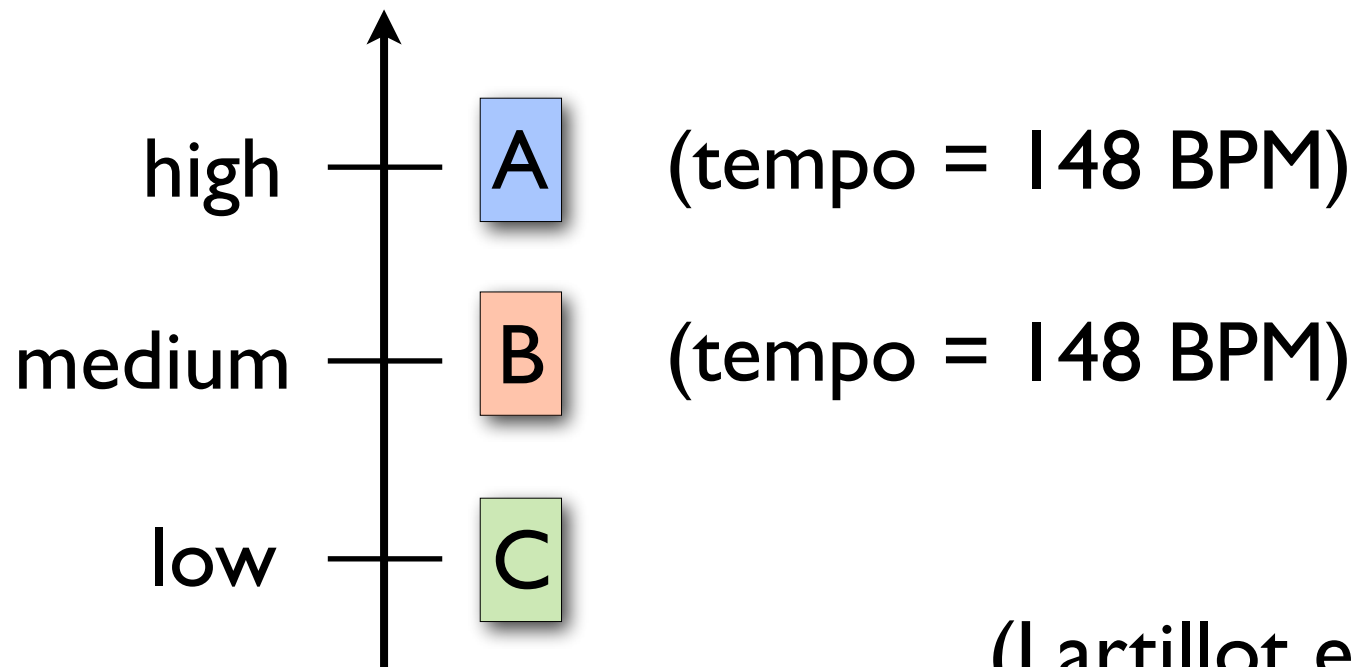
M. Bruch,
Violin Concerto
No. I in G minor,
op.26, Finale
(Allegro energico)

1 = 

.16 = 

Beat/rhythmic/metrical strength, clarity

- Subjective judgment:
 - How easily I can perceive the underlying pulsation in music.




(Lartillot et al., 2008)

mirp pulse clarity

beat strength

- Characterisation of the autocorrelation function



'MaxAutocor' **'MinAutocor'**

'KurtosisAutocor'

'TempoAutocor'

'EntropyAutocor'


(Lartillot et al., 2008)

- cf. also *beat strength* (Tzanetakis, Essl & Cook, 2002): variability of the autocorrelation curve throughout time

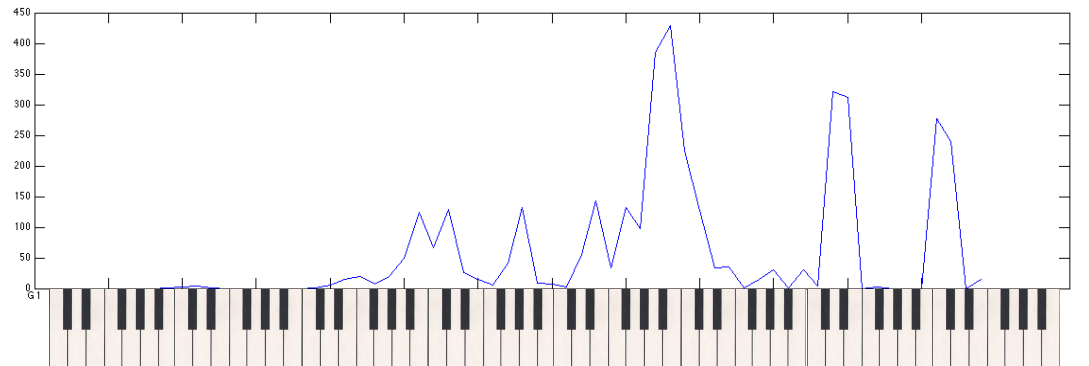
mirchromagram

energy distribution along pitches


- $\mathbf{s} = \text{mirspectrum}(\mathbf{a})$



- $\mathbf{c} = \text{mirchromagram}(\mathbf{s},$
'Wrap', 'no')




- $\mathbf{c} = \text{mirchromagram}(\mathbf{c},$
'Wrap', 'yes')




mirkeystrength

probability of key candidates

mirchromagram




C major




C minor

C# major




C# minor

D major



D minor


Cross-Correlations



mirkeystrength

probability of key candidates

- Chromagram compared to typical chromagrams representing each possible key (or mode).
- Detection of the most probably key (or mode).




Krumhansl, *Cognitive foundations of musical pitch*. Oxford UP, 1990. Gomez, "Tonal description of polyphonic audio for music content processing," *INFORMS Journal on Computing*, 18-3, pp. 294–304, 2006.

mirkey

tonality estimation

mirpeaks(mirkeystrength(...))

mirkeystrength


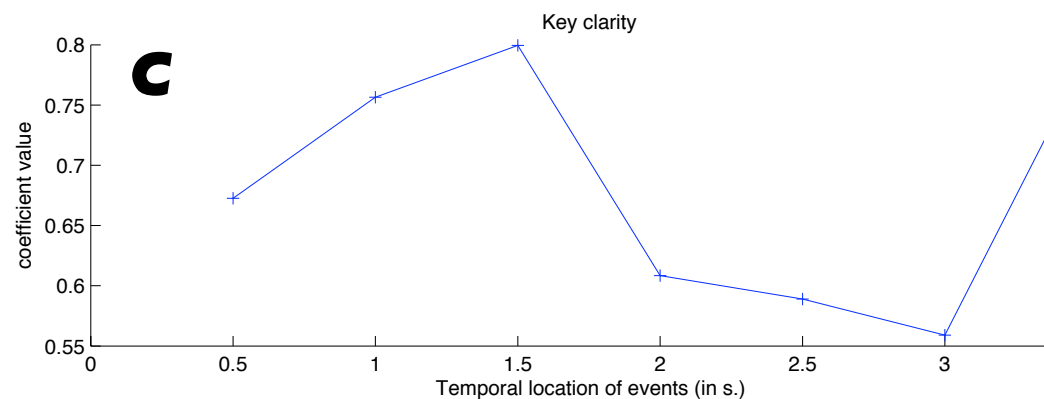



- **[k c s] = mirkey(...)**

mirkey

tonality estimation

- $[k \ c \ s] = \text{mirkey}(\dots, \text{'Frame'})$




mirkey


tonality estimation




Keys




monteverdi.wav



tiersen.wav



beethoven9.wav




schoenberg.wav

- $[k \ c \ s] = \text{mirkey}(\dots, \text{'Frame'})$

**clear
tonality**


unclear



Tchaikovski, *Swan Lake*, Act one

clear

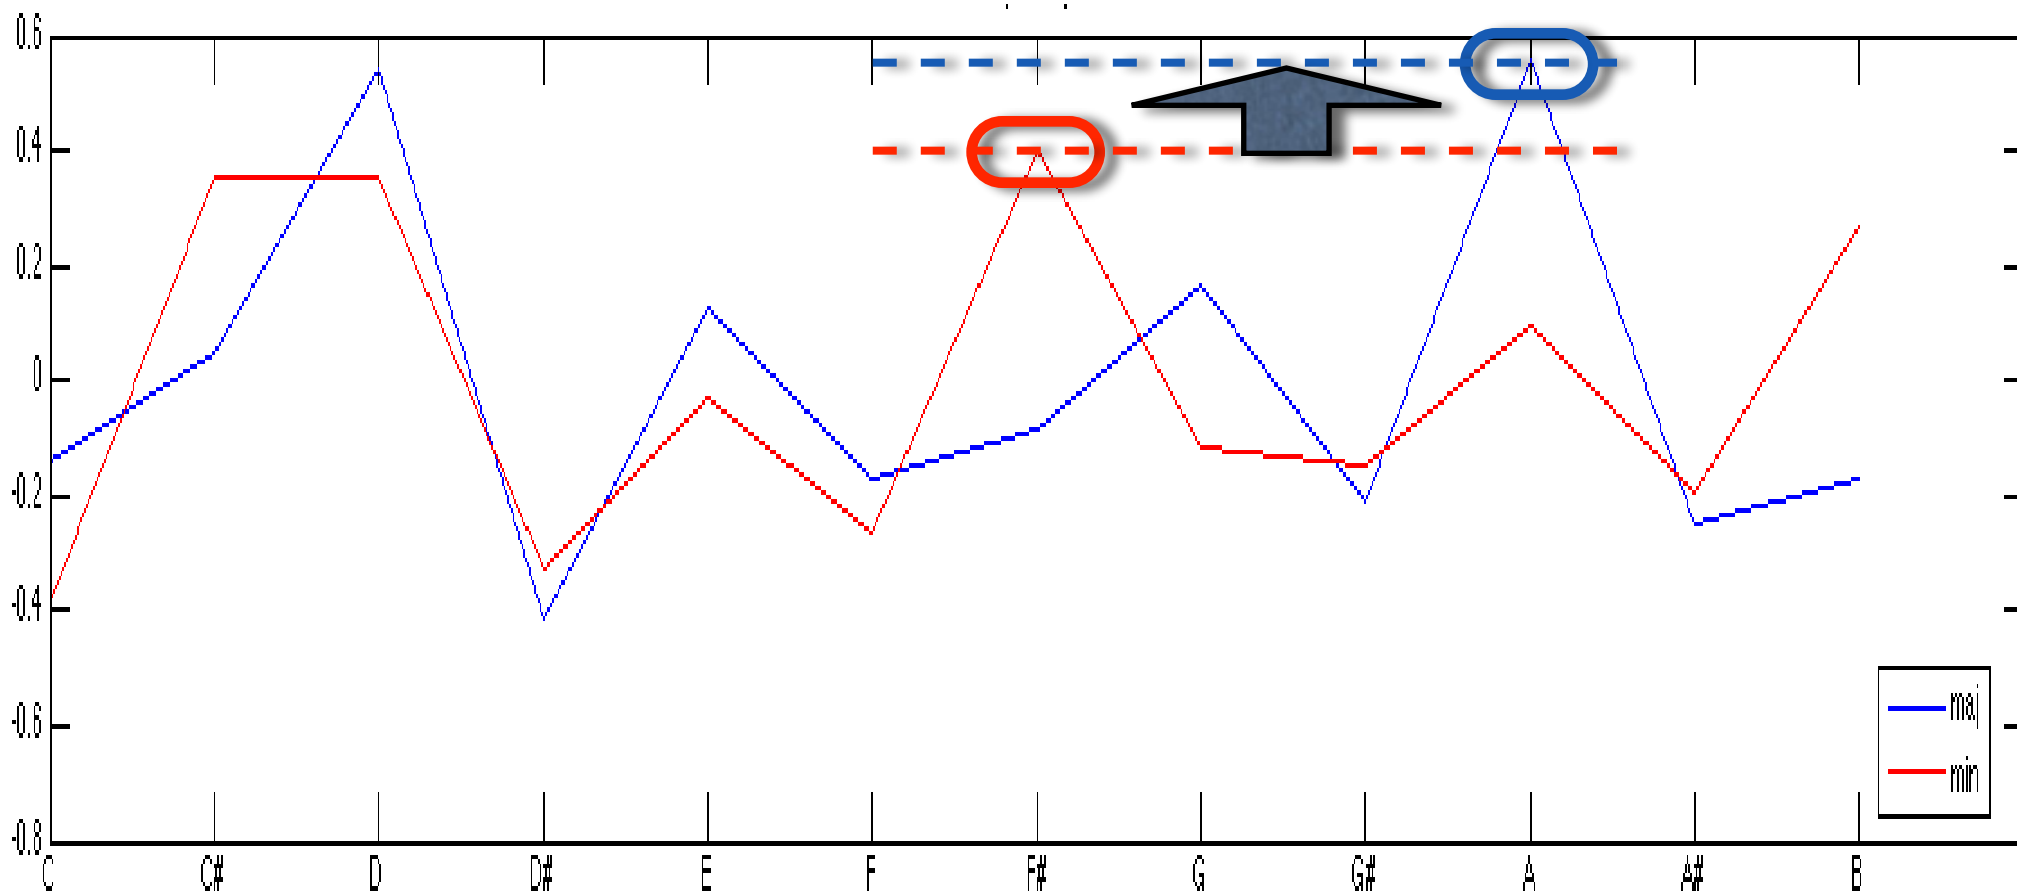
unclear



Prokofiev, Violon concerto No. in D major, *Scherzo: Vivacissimo*

mirmode


mode estimation



mirmode

mode estimation


major



minor

Schubert, Piano Trio No.2 in E-flat major, *Andante con moto*

major




minor

Mozart, Piano Concerto No.23 in A major, *Adagio*

mirkeysom

self-organizing map projection of chromagram



Toiviainen & Krumhansl, "Measuring and modeling real-time responses to music: The dynamics of tonality induction", *Perception* 32-6, pp. 741–766, 2003.

mirsimatrix

similarity matrix

- Any musical feature as input x :

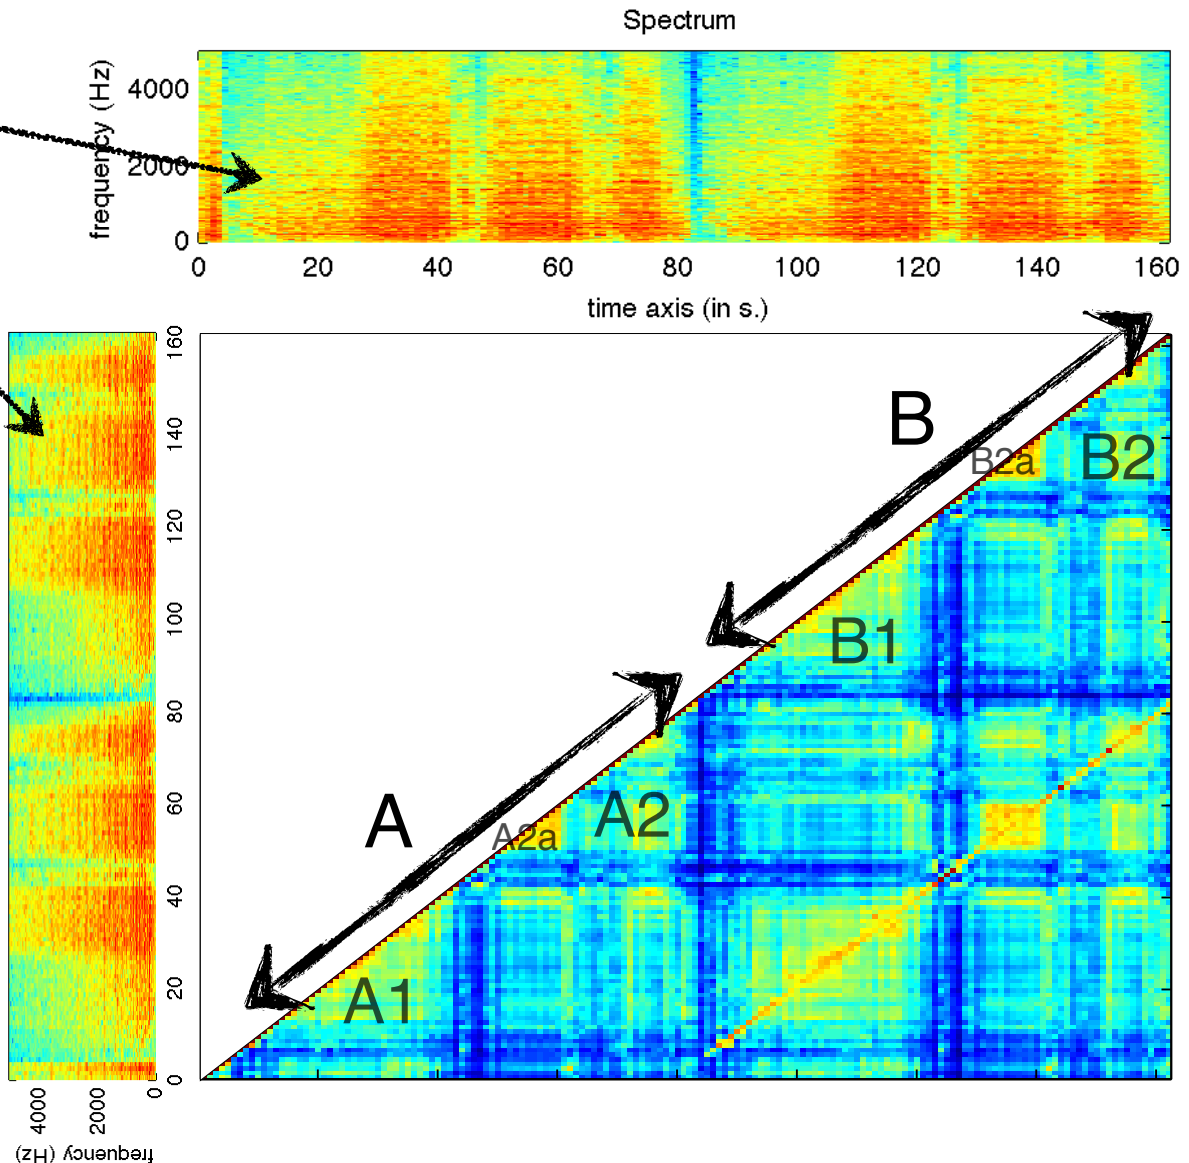
- Spectrum

- Timbre(MFCC, ...)

- Tonality (Chroma, key strength, ...)

- Rhythm, etc.

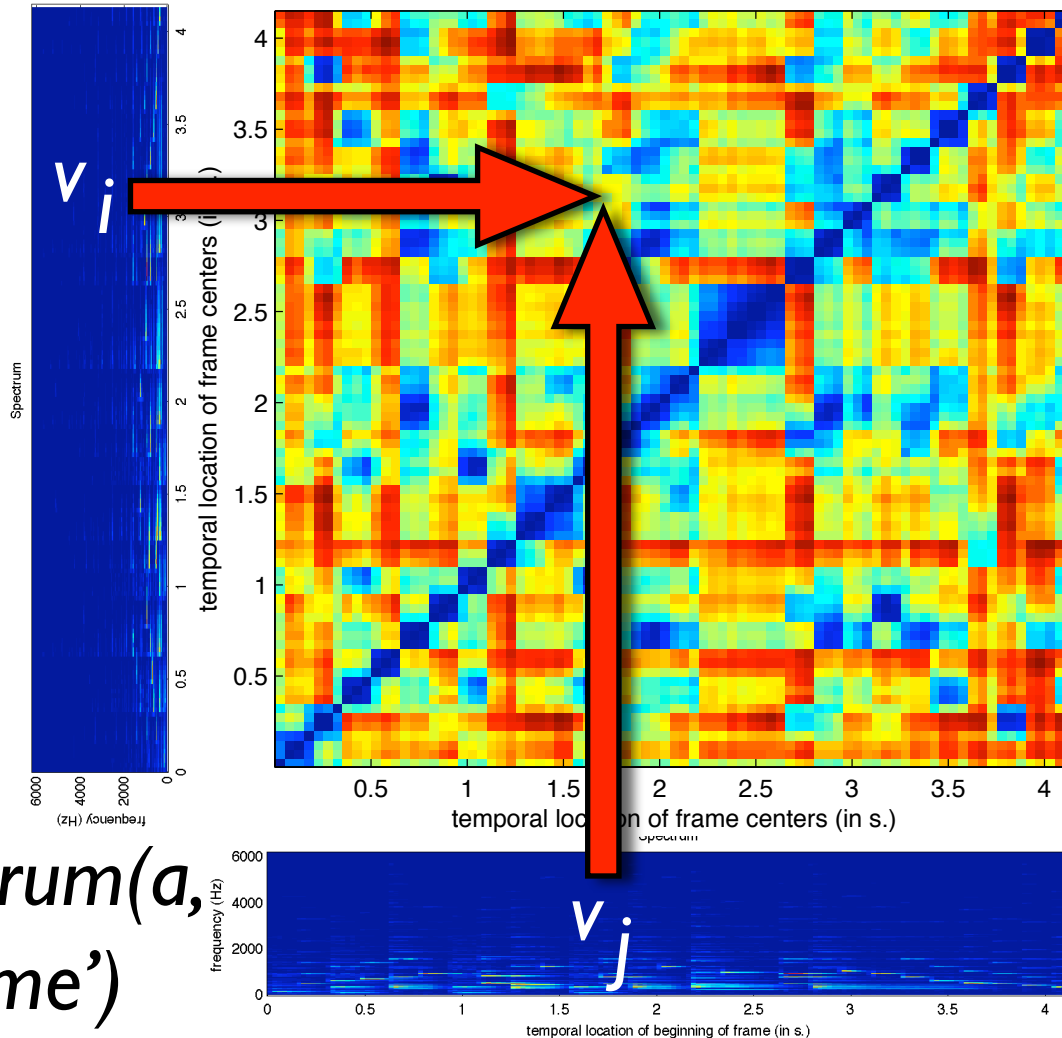
- Frames are compared using a distance measure (such as cosine distance)



mirsimatrix

dissimilarity matrix

Dissimilarity matrix



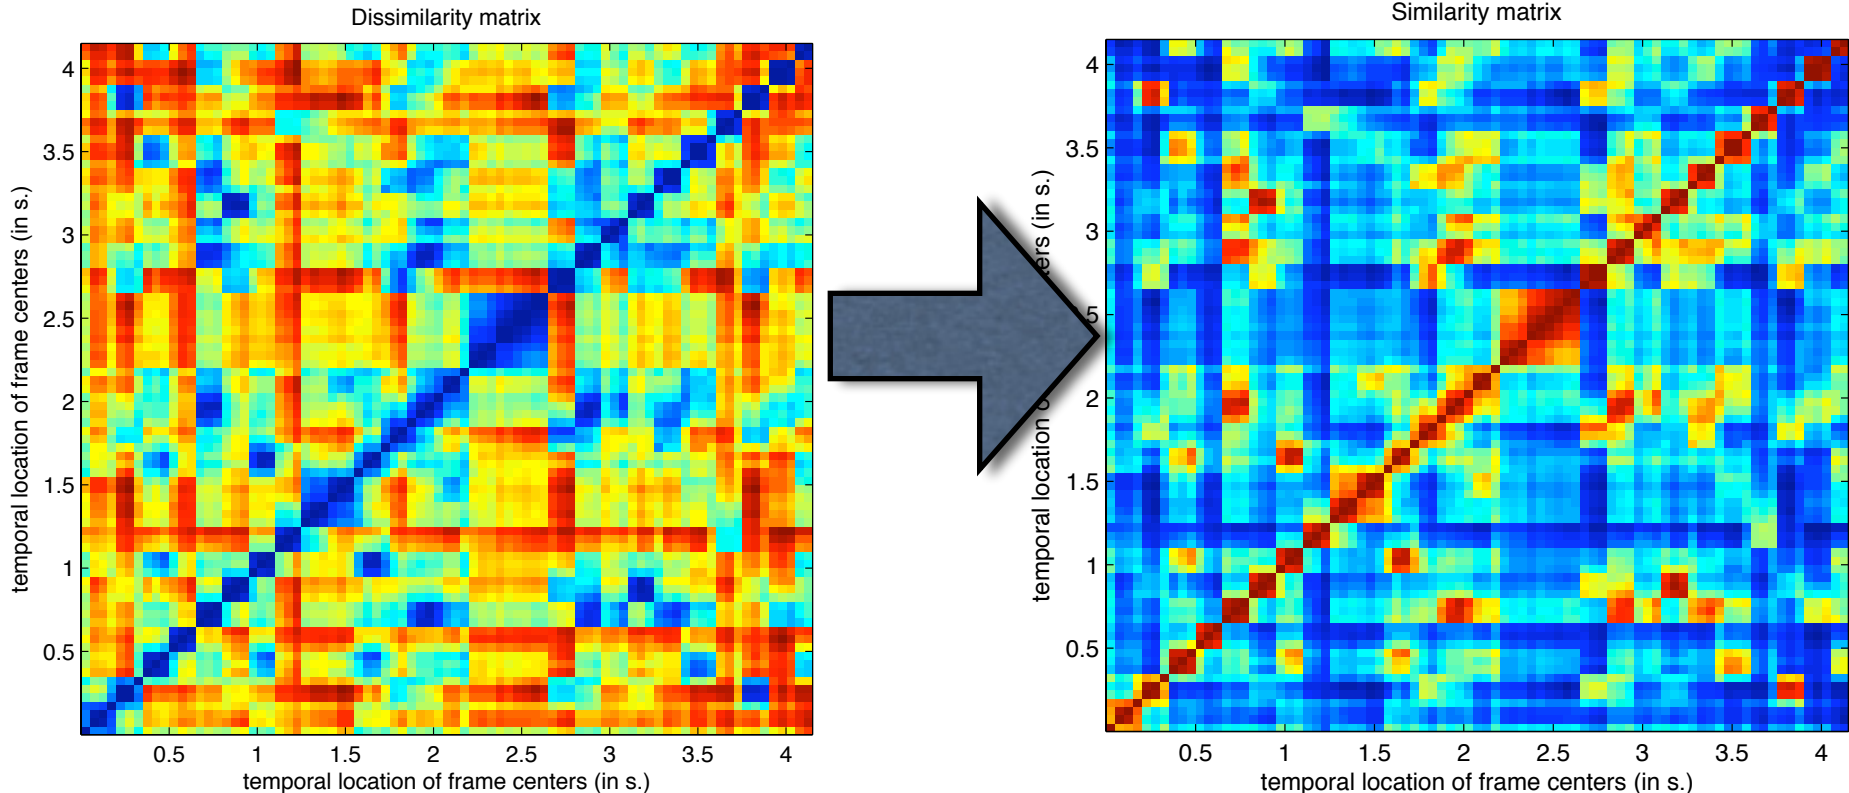
mirsimatrix(*a*,
‘Dissimilarity’)

mirsimatrix(...,
‘Distance’,
‘cosine’)

$$d_{cos}(v_i, v_j) = \frac{\langle v_i, v_j \rangle}{|v_i||v_j|}$$

mirspectrum(*a*,
‘Frame’)

mirsimatrix similarity matrix



mirsimatrix(*a*, **'Similarity'**, **'exponential'**)

$$d_{exp}(v_i, v_j) = \exp(-d_{cos}(v_i, v_j))$$

Foote, Cooper. "Media Segmentation using Self-Similarity Decomposition",
SPIE Storage and Retrieval for Multimedia Databases, 5021, 167-75.

mirsimatrix

similarity matrix



- For instance:
 - $s = \text{mirspectrum}(\dots, \text{'Frame'}, \dots)$
 - $\text{mirsimatrix}(s)$

- Observe the structure of this excerpt along different musical dimensions:
 - *'george.wav'*

mirnovelty

novelty, kernel method

- Convolution with checkerboard kernel along the diagonal of *mirsimatrix*
- **Peaks** indicate transitions between phases




mirnovelty

novelty, kernel method

- *mirnovelty*(..., '**Kernel**', N)
 - where N is the kernel size
- Observe the structure of this excerpt using different kernel sizes:
 - 'george.wav'

mirnovelty


novelty, kernel method



Kernel size:

- 64 frames

- 16 frames




mirsegment


segmentation

- ***nv*** = *mirnovelty(sm)*
- ***p*** = *mirpeaks(nv)*
- ***sg*** = *mirsegment('mysong', p)*
- *mirplay(sg)*
- ***s*** = *mirmfcc(sg)*
- ***sm*** = *mirsimatrix(s)*


sm



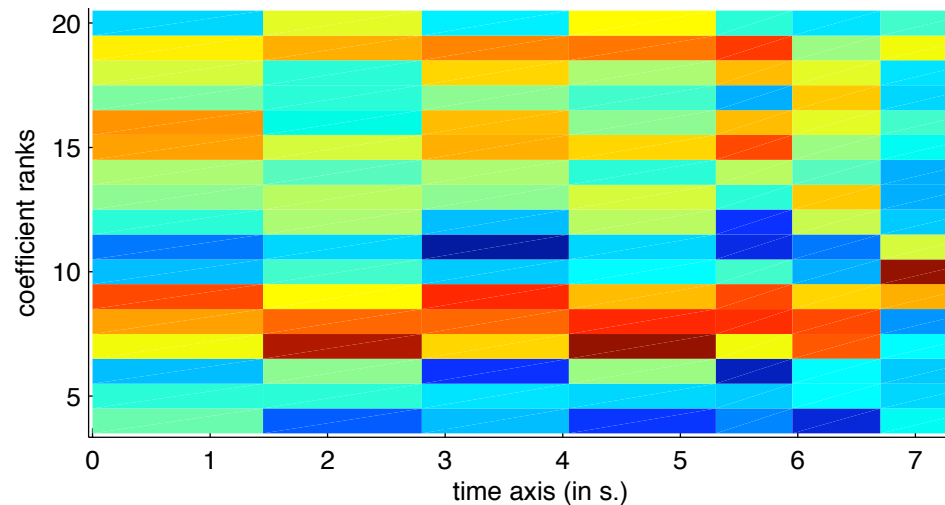
nv



sg





s



Statistics

- ***mirstat***
 - mean
 - standard deviation
 - slope
 - periodicity
- ***mirhisto***
 - distribution histograms



- moments
 - ***mircentroid***
 - ***mirspread***
 - ***mirskewness***
 - ***mirkurtosis***



mirfeatures

batch of features


- *mirzerocross*
- *mircentroid*
- *mirbrightness*
- *mirspread*
- *mirskewness*
- *mirkurtosis*
- *mirrolloff*
- *mirentropy*
- *mirflatness*
- *mirroughness*
- *mirregularity*
- *mirinharmonicity*
- *mirmfcc*
- *mirfluctuation*
- *mirattacktime*
- *mirattackslope*
- *mirlowenergy*
- *mirflux*
- *mirpitch*
- *mirchromagram*
- *mirkeystrength*
- *mirkey*
- *mirmode*
- *mirhcdf*
- *mirtempo*
- *mirpulseclarity*

mirfeatures(‘Folder’, ‘**Stat**’)


mirexport

exportation of statistical data to files

- *mirexport(filename, ...)* adding one or several data from *MIRtoolbox* operators.
- *mirexport('result.**txt**', ...)* saved in a text file.
- *mirexport('result.**arff**', ...)* exported to WEKA for data-mining.
- *mirexport('Workspace', ...)* saved in a Matlab variable.



Eerola & Vuoskoski. A comparison of the discrete and dimensional models of emotion in music. Psychology of Music.



	3D				2D		
	R^2	β			R^2	β	
		<i>v</i>	<i>a</i>	<i>t</i>		<i>v</i>	<i>a</i>
happy	.89	.93	.79	-.35	.89	.85	.49
sad	.63	-.20	-.84	-.22	.63	-.05	-.69
tender	.77	.33	-.45	-.58	.77	.50	-.51
fear	.87	-.83	.07	.63	.87	-.90	.24
anger	.64	-.52	.32	.35	.64	-.55	.35
mean	.76				.76		

Eerola & Vuoskoski. A comparison of the discrete and dimensional models of emotion in music. Psychology of Music.

miremotion

Box-Cox

R^2	h	s	t	f	a
MLR	.55	.56	.42	.51	.54
PCA	.34	.40	.23	.13	.27
PLS	.59	.64	.52	.60	.57
MLR	.62	.58	.44	.60	.63
PCA	.38	.49	.47	.45	.62
PLS	.65	.61	.55	.64	.67

happy	
Feature	β
fluctuation	-.40
sp. centroid	.13
sp. spread	-.19
chrom. peaks	-.05
majorness	.03

sad	
Feature	β
roughness	.12
register	-.08
register var	.09
majorness	.02
harm. change	-.03

tender	
Feature	β
RMS var	-.42
sp. centroid	.14
key clarity	.11
harm. change	-.10
tonal novelty	-.01

fear	
Feature	β
RMS var	-.79
fluctuation	-.21
key clarity	-.09
harm. change	.08
tonal novelty	-.02

anger	
Feature	β
RMS var	.44
pulse clarity	-.13
sp. centroid	-.13
key clarity	-.04
tonal novelty	.02

Part 2

- Rhythm, metrical structure
- Tonal analysis
- Segmentation, structure
- Statistics
- Music & emotion
- **Advanced use**

mirgetdata

return data in Matlab format

s = *mirspectrum*('file'); →

Encapsulated data **s**
numerical data,
related sampling rates,
related file name,
etc.

mirgetdata(s) ↙

vector
|

mirgetdata



return data in Matlab format

$s = \text{mirspectrum}(\text{'file'},$
Frame');



Encapsulated data **s**
grouped numerical data,
related sampling rates,
related file name,
etc.

mirgetdata(s)



mirgetdata

return data in Matlab format


f = *mirfilterbank*('file',
'Frame');



Encapsulated data
grouped numerical data,
related sampling rates,
related file name,
etc.

f

mirgetdata(*f*)



3D-matrix




mirgetdata

return data in Matlab format

```
sg = mirsegment('file')  
s = mirspectrum(sg, 'Frame');
```

Encapsulated data s
grouped numerical data,
related sampling rates,
related file name,
etc.

mirgetdata(s)



mirgetdata


return data in Matlab format

```
sg = mirsegment('Folder')  
s = mirspectrum(sg, 'Frame');
```



Encapsulated data *s*
grouped numerical data,
related sampling rates,
related file name,
etc.

```
mirgetdata(s)
```



mirgetdata

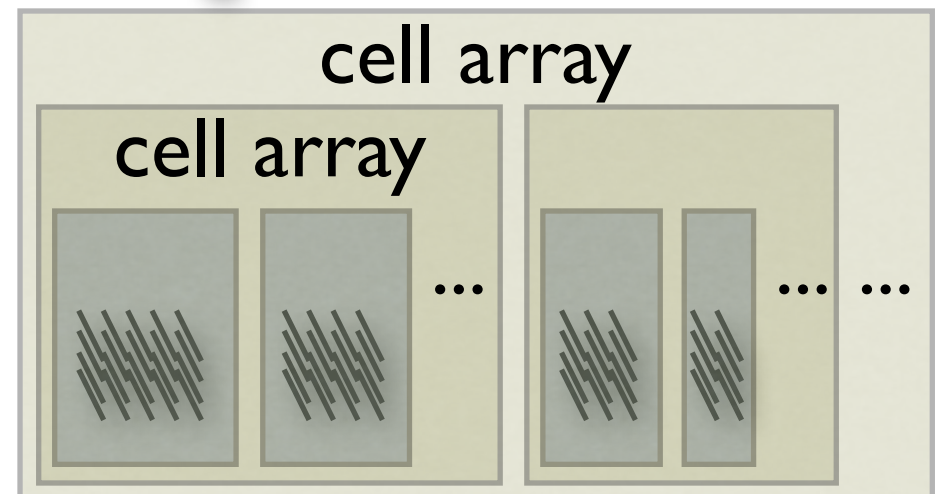

return data in Matlab format

$p = \text{mirpeaks}(a) \longrightarrow$

Encapsulated data p
grouped numerical data,
related sampling rates,
related file name,
etc.

$\text{mirgetdata}(p)$
peak positions

peak values



get

returns fields of encapsulated data

`a = miraudio('ragtime');` →

Encapsulated data

- `get(a, 'xName')`
- `get(a, 'xData')`
- `get(a, 'yName')`
- `get(a, 'yData')`
- `get(a, 'yUnit')`
- `get(a, 'FramePos')`
- `get(a, 'Sampling')`
- `get(a, 'NBits')`
- `get(a, 'Title')`
- `get(a, 'FileName')`
- `get(a, 'Label')`
- `get(a, 'Channels')`
- `get(a, 'xPeakSample')`
- `get(a, 'xPeakUnit')`
- `get(a, 'xPeakInterpol')`
- `get(a, 'yPeak')`
- `get(a, 'yPeakInterpol')`

get

returns fields of encapsulated data

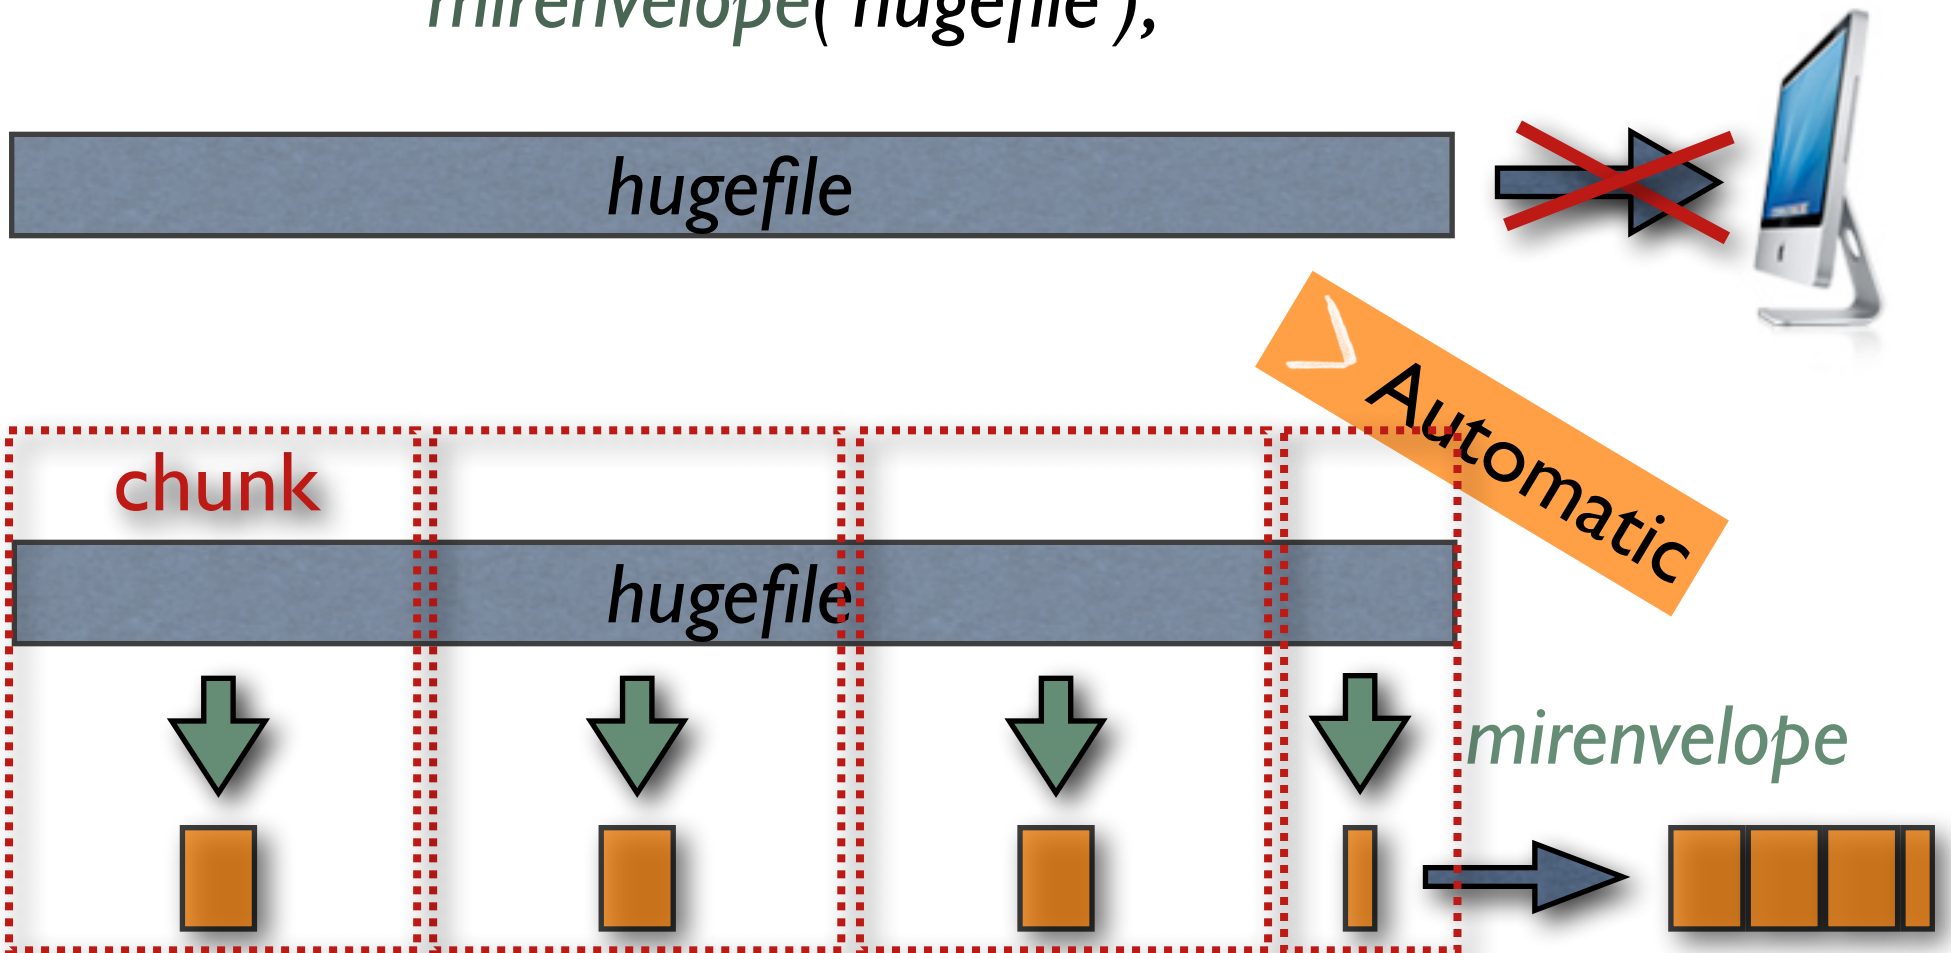
`s = mirspectrum('ragtime');` \longrightarrow Encapsulated data

- `get(s, Frequency) = get(s, 'xData')`
- `get(s, Magnitude) = get(s, 'yData')`
- `get(s, Phase)`
- `get(s, xScale) (= 'Freq', 'Mel', 'Bark')`
- `get(s, Power)`
- `get(s, dB)`

etc.

memory management

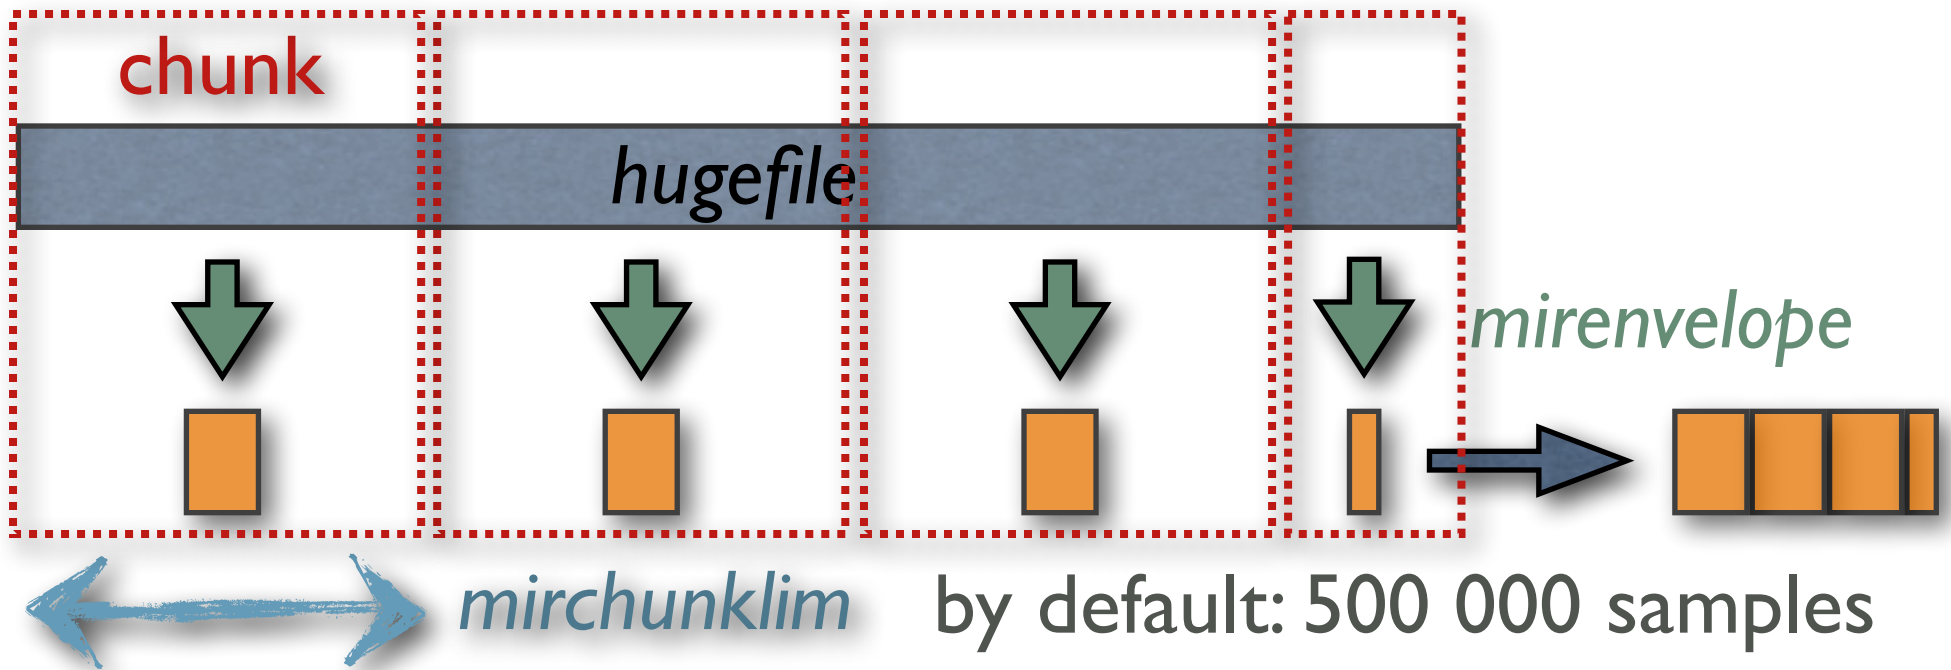
mirenvelope('hugefile');



mirchunklim

chunk size limitation

```
mirenvelope('hugefile');
```




► If memory overflow problems, decrease *mirchunklim*:

```
mirchunklim(50000) set to 50 000 samples
```


avoid useless call to *miraudio*

- $a = \text{miraudio}(\text{'hugefile'})$;
- $c = \text{mirmfcc}(a)$

a
↓
 c




- $c = \text{mirmfcc}(\text{'hugefile'})$




avoid useless call to *miraudio*

- $a = \text{miraudio}(\text{Folder});$
- $c = \text{mirmfcc}(a)$

a
↓
 c




- $c = \text{mirmfcc}(\text{Folder});$




avoid useless call to *mirframe*

- $f = \text{mirframe}(\text{'hugefile'});$
- $c = \text{mirmfcc}(f)$

f
↓
 c



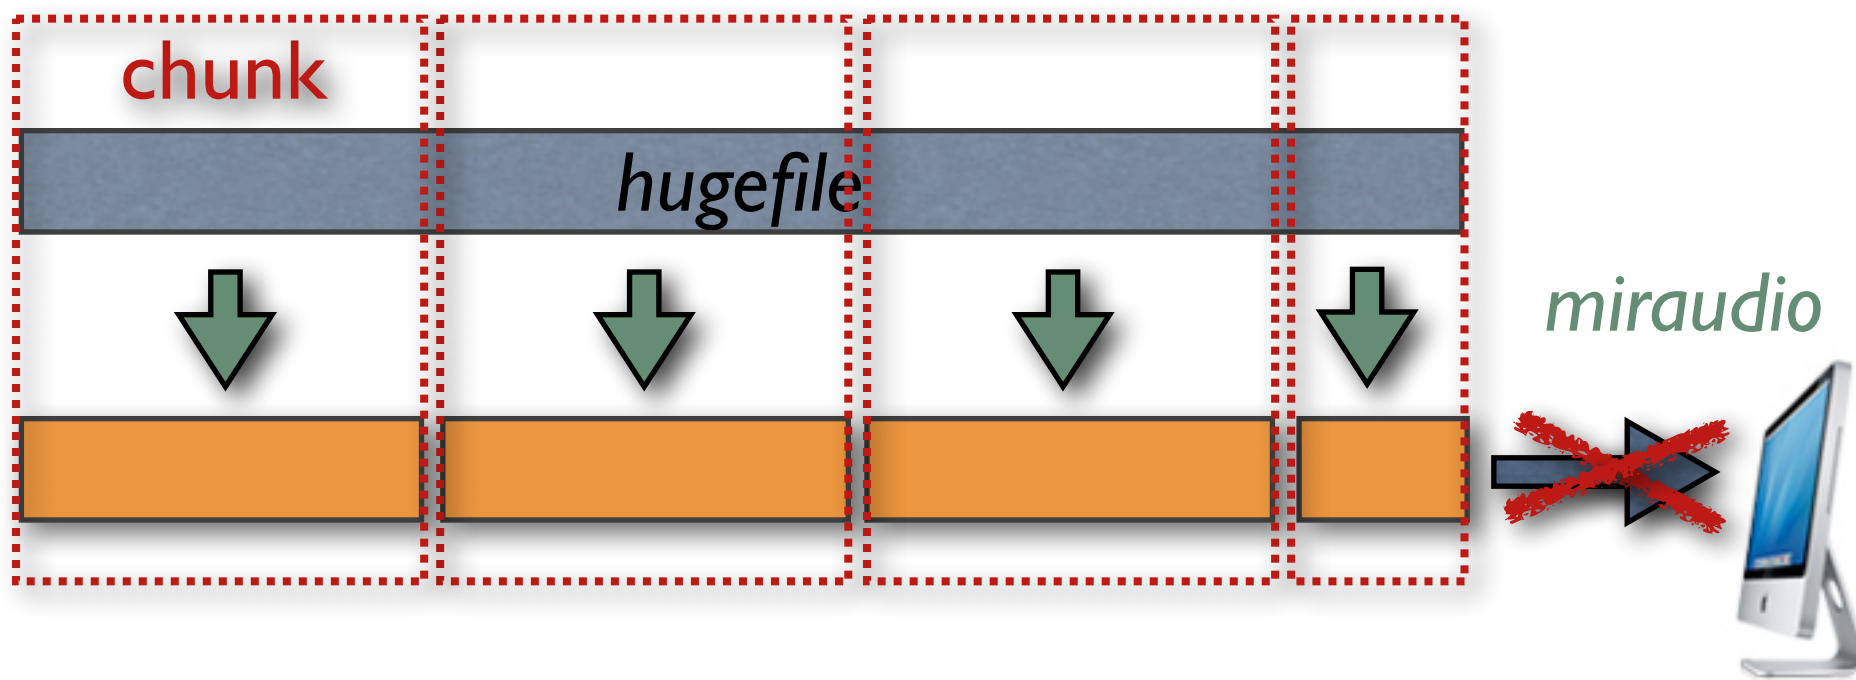
- $c = \text{mirmfcc}(\text{'hugefile'}, \text{'Frame'})$



what if *miraudio* (or *mirframe*) really necessary?

?

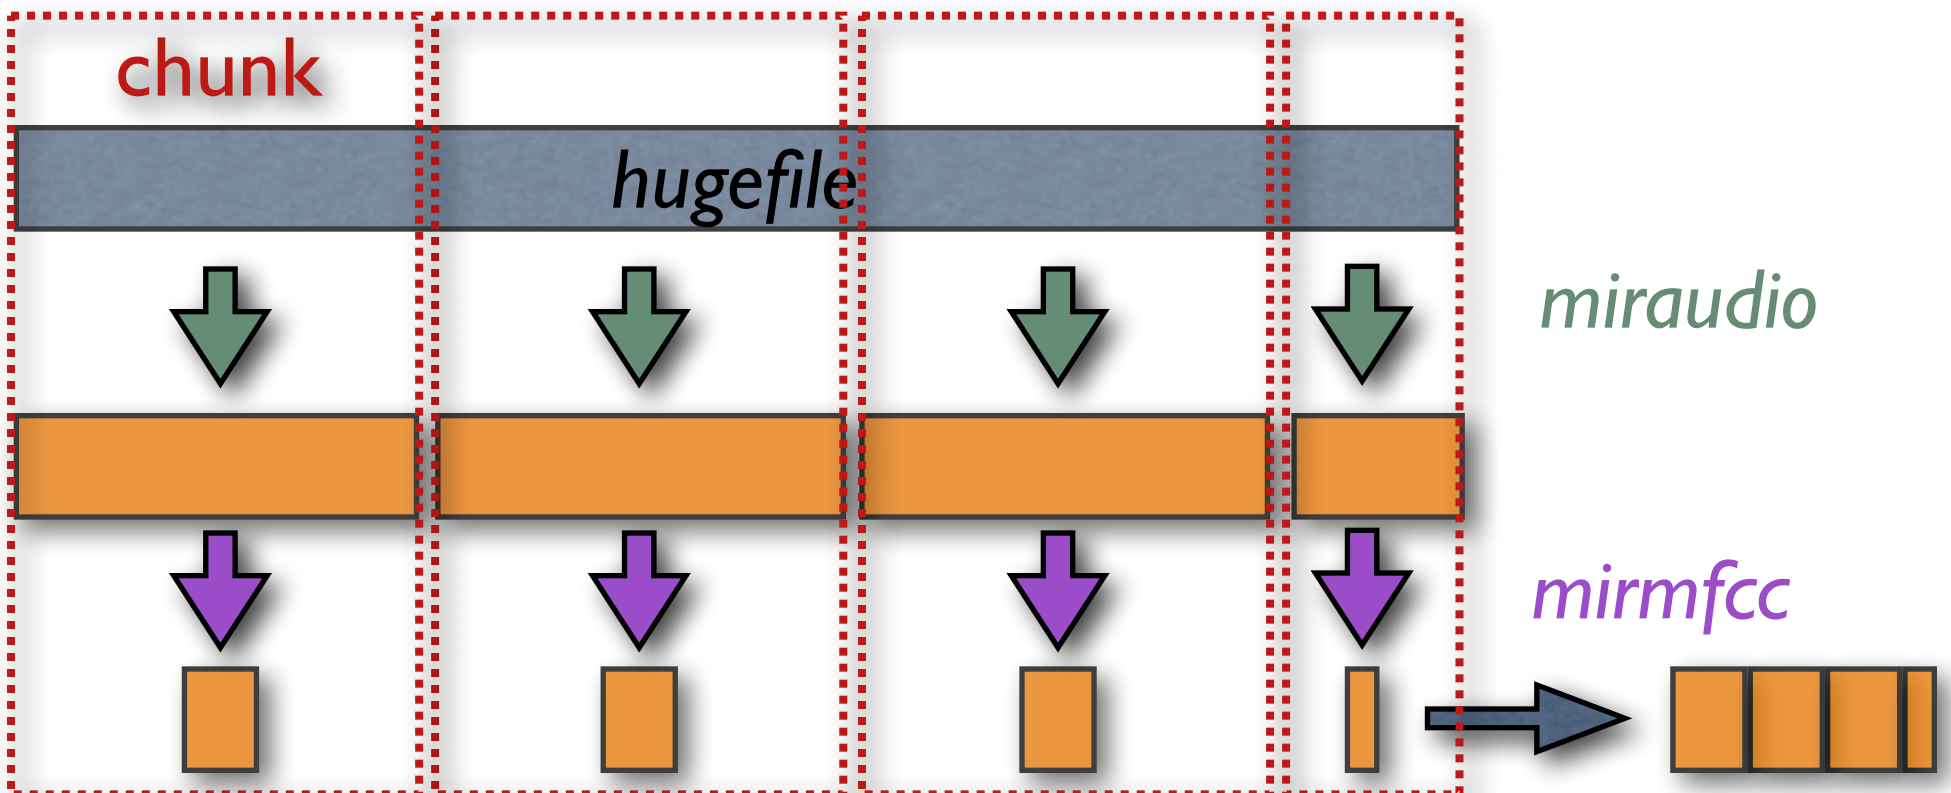
- $a = \text{miraudio}(\text{'hugefile'}, \text{'Sampling'}, 11025);$ a
- $c = \text{mirmfcc}(a)$ \downarrow
 c



mireval

flowchart design and evaluation

- $a = \text{miraudio}(\text{“Design”}, \text{“Sampling”}, 11025);$ a
- $c = \text{mirmfcc}(a);$ \downarrow
 c
- **$\text{mireval}(c, \text{“hugefile”})$**

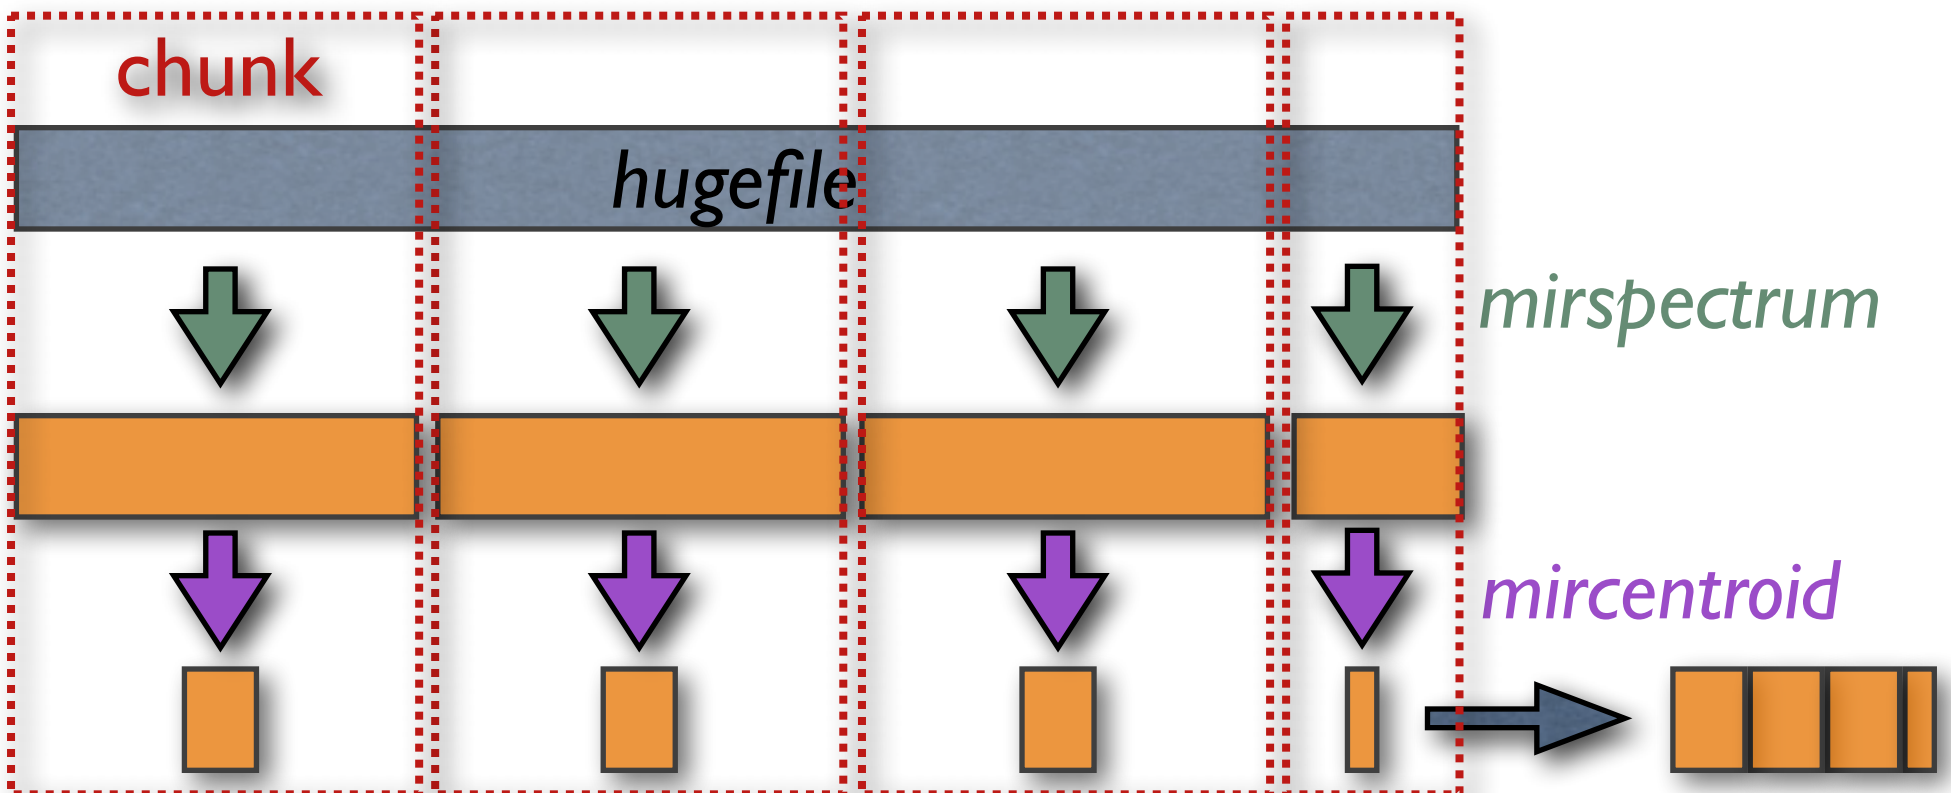


mireval

flowchart design and evaluation

- $s = \text{mirspectrum}(\text{Design}, \text{Frame})$;
- $c = \text{mircentroid}(s)$;
- **$\text{mireval}(c, \text{hugefile})$**

s
↓
 c

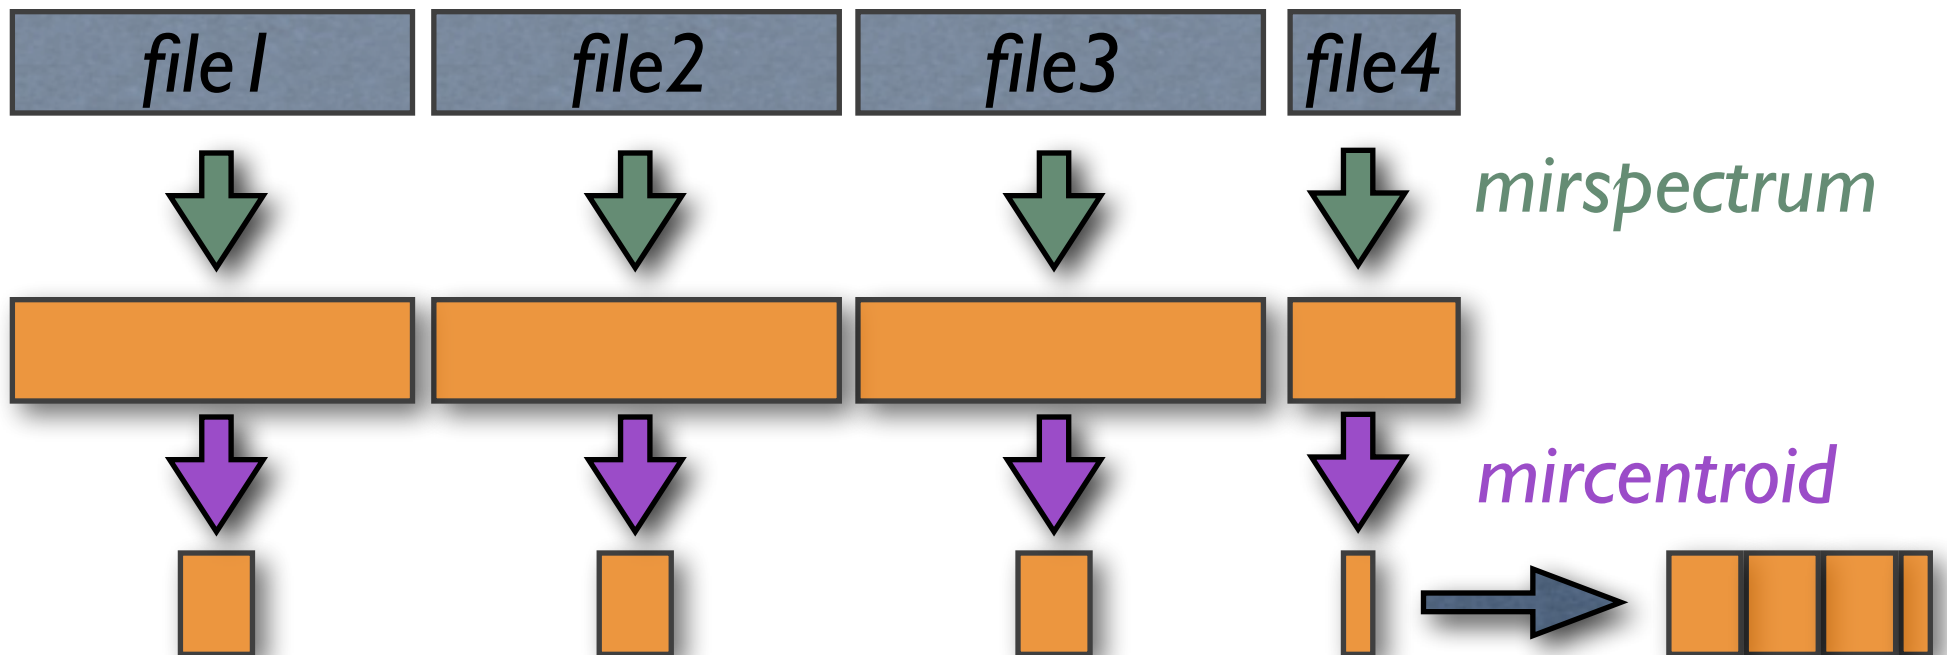


mireval

flowchart design and evaluation

- $s = \text{mirspectrum}(\text{Design}, \text{Frame});$
- $c = \text{mircentroid}(s);$
- **$\text{mireval}(c, \text{Folder})$**

s
↓
 c

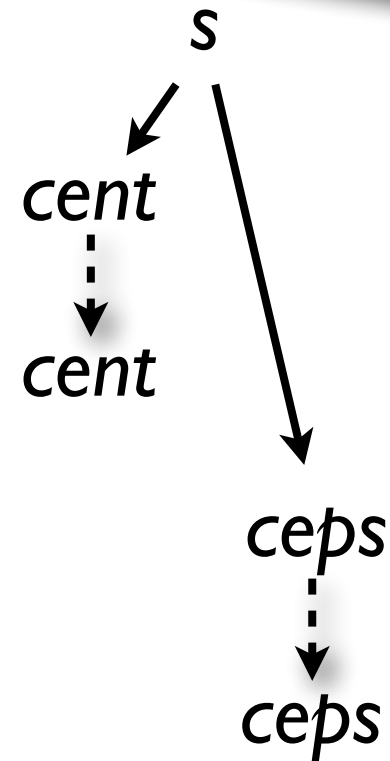


mireval

flowchart evaluation?

s is evaluated twice!

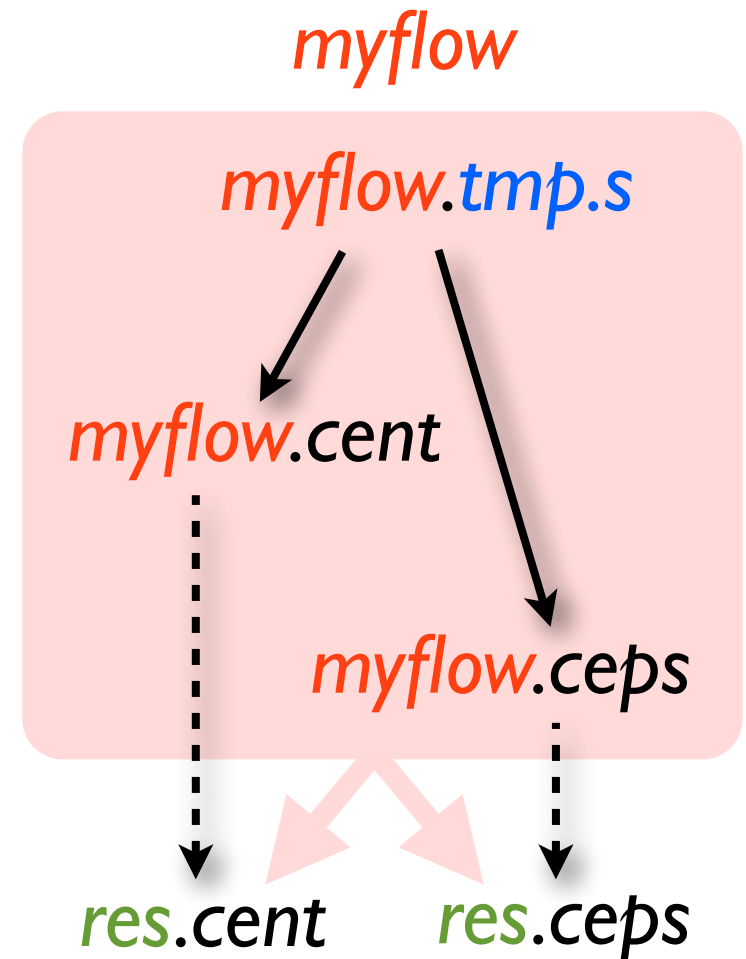
- $s = \text{mirspectrum}(\text{Design}, \text{Frame});$
- $\text{cent} = \text{mircentroid}(s);$
- $\text{cent} = \text{mireval}(\text{cent}, \text{Folder});$
- $\text{ceps} = \text{mircepstrum}(s);$
- $\text{ceps} = \text{mireval}(\text{ceps}, \text{Folder});$



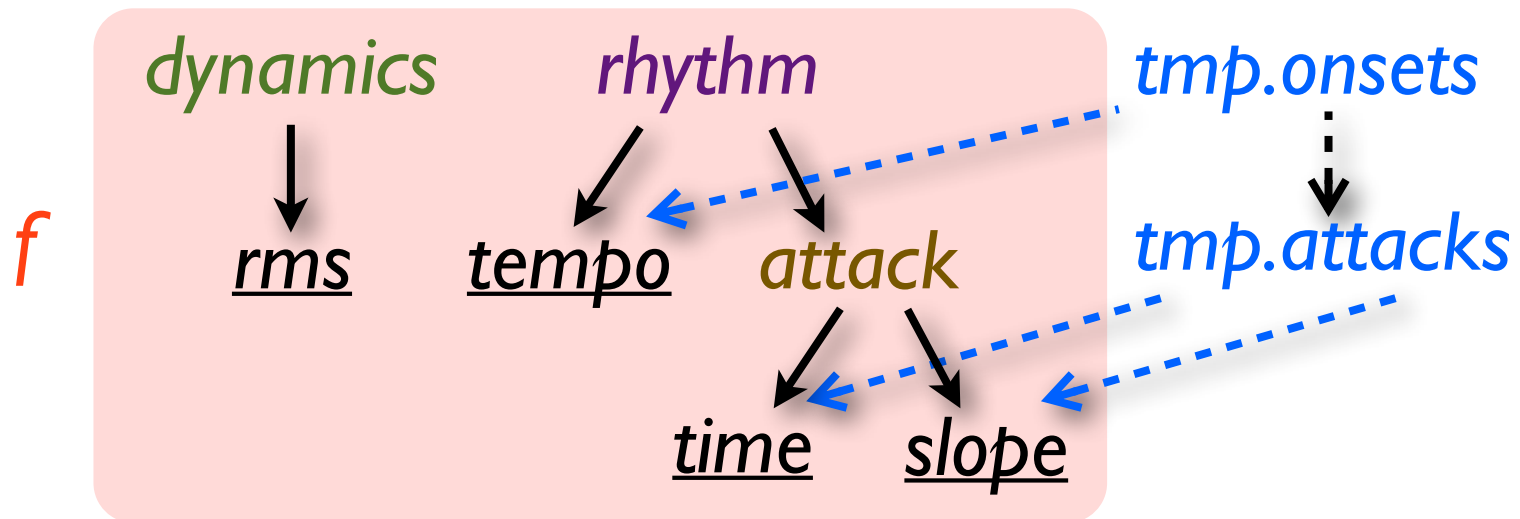
mirstruct

complex flowchart

- *myflow* = **mirstruct**;
- *myflow.tmp.s* =
mirspectrum('Design', 'Frame');
- *myflow.cent* =
mircentroid(*myflow.tmp.s*);
- *myflow.ceps* =
mircepstrum(*myflow.tmp.s*);
- *res* = **mireval**(*myflow*, 'Folder');



complex flowchart



- f = mirstruct;
- $f.dynamics.rms$ = mirrms("Design", 'Frame')
- $f.tmp.onsets$ = mironsets("Design");
- $f.rhythm.tempo$ = mirtempo($f.tmp.onsets$, 'Frame');
- $f.tmp.attacks$ = mironsets($f.tmp.onsets$, 'Attacks');
- $f.rhythm.attack.time$ = mirattacktime($f.tmp.attacks$);
- $f.rhythm.attack.slope$ = mirattackslope($f.tmp.attacks$);