



CDP Functions to FOCUS Spectral Data

(with Command Line Usage)

Functions which FOCUS Spectral Data

Quick Musical Overview

Musical Effect	FOCUS Module(s)
internal glissandos	ACCUMULATE, SUPERACCU
tonal changes	EXAG, FOCUS, FOLD
sample-hold	FREEZE, HOLD, STEP

FOCUS ACCU

Sustain each spectral band, until louder data appears in that band

FOCUS EXAG

Exaggerate the spectral (formant) contour

FOCUS FOCUS

Focus spectral energy onto the peaks in the spectrum

FOCUS FOLD

Octave-transpose spectral components into a specified frequency range

FOCUS FREEZE

Freeze the spectral characteristics in a sound, at given times, for specified durations

FOCUS HOLD

Hold sound spectrum, at given times

SPECFOLD

Fold, invert or randomise the spectrum

FOCUS STEP

Step-frame through a sound by freezing the spectrum at regular time intervals

SUPERACCU

Sustain each spectral band until louder data appears in that band

The 'typical command lines' shown below use parameters described in the text and should produce good results, although exactly what happens is dependent on the nature of the input sound. The complex sound used was that of a strum on a 3 ft spring, 1 cm diameter. The pitched sound was a sustained French horn tone.

FOCUS ACCU – Sustain (accumulate) each spectral band, until louder data appears in that band

*This function has been superseded and extended by **SUPERACCU**.*

Usage

focus accu *inanalysisfile outanalysisfile -ddecay -gglis*

Typical command line:

focus accu raspdt.ana raspdtacc.ana -d0.75 -g-0.9

Parameters

inanalysisfile – input analysis file

outanalysisfile – output analysis file

-ddecay sustained channel data decays by a factor of *decay* each second

(Possible Range: 0.001000 to 1.0. Default is 1.0. – Suggested effective Range: 0.000001 to 0.5)

-gglis sustained channel data glissandos at *glis* octaves per second

(Approximate Range: -11.7 to 11.7. Default is 0)

Understanding the FOCUS ACCU Process

Frequencies are sustained into subsequent windows. The overall effect is one of sustaining, but one which also makes the spectrum more complex.

The *gliss* parameter appears to change the result the most. It produces glissandos within the spectrum of the sound. Very effective slow glissandos are produced when *gliss* is near 0, e.g., -0.9 or 0.1. At 0.5, there are several glissandos, at 1, they are fairly fast, and at 10 it becomes a wash.

The range allowed for *gliss* is really very extreme, especially in the negative range below zero. Values e.g., of -5.0 or less lead to the sound settling down to a steady, sustained pitch after an initial short glissando. **Recommended effective range is from -0.5 to 10.**

The *delay* parameter is important for making the glissandos more perceptible. A low value of e.g., 0.1 is really very dry, whereas something like 0.75 brings out the glissandos nicely, with a kind of reverberant aura.

Musical Applications

The main type of output from this function creates a supple, internally modulating sound, moving slowly at the low end of the *gliss* range and increasingly quickly as one moves towards the high end.

With a clearly pitched input, the results are aurally clearer, but otherwise similar. **TIMESTRETCH** applied to these results produces nicely slow-moving timbral and pitch changes.

SEE ALSO: HILITE ARPEG. Churning, bubbly textures can also be achieved with **BLUR SCATTER** (high values for *keep*) and **HILITE TRACE** (low value for *channels*).

End of FOCUS ACCU

FOCUS EXAG – Exaggerate spectral contour

Usage

focus exag *inanalysisfile outanalysisfile exaggeration*

Typical command line:

focus exag raspdt.ana raspdtexg.ana 0.5

Parameters

inanalysisfile – input analysis file

outanalysisfile – output analysis file

exaggeration – **< 1** widens troughs and narrows formants (focuses on the peaks); **> 1** narrows troughs and widens formants (diffusing the peaks). Range: 0.00100 to 1000.

exaggeration may vary over time

Understanding the FOCUS EXAG Process

FOCUS EXAG plays with the frequency focus of regions of high amplitude among the partials of a spectrum. A spectral peak is an area of high amplitude among the frequencies of the spectrum. A trough is an area of low amplitude between spectral peaks. Formants are slightly more than spectral peaks: i.e., they are the area around a spectral peak, and might be narrow, wide etc.).

- Narrowing the formants (**< 1**) increases the frequency focus of each peak area, and therefore acts like an increasingly tight filter.
- Widening the formants (**> 1**) reduces the frequency focus of the formants, and therefore makes the sound increasingly fuzzy.

Musical Applications

FOCUS EXAG can be used to make a sound thinner and pitched, just as a bandpass filter can create pitches within a complex sound. Working in the spectral domain, it does this by manipulating frequency regions of high amplitude, otherwise known as 'formants'. These can be visualised as little mountain peaks. 'Narrowing' the formants means making the 'mountains' narrower and its sides steeper; furthermore, the material inbetween virtually disappears. 'Widening' the formants means making the mountainsides slope more gradually, therefore encompassing more of the surrounding region. Because this is actually moving the position of partials, this results in timbral change and the addition of noise elements, hence the fuzziness referred to above.

I have found that a value of *exag* = 0.5 maintains a nice balance between the original and the treated sound. Lower values start to produce results reminiscent of BLUR TRACE, while higher values, e.g., 10, produce quite a fuzzy, noisy result, at least with a complex sound as input. 5 with a pitched sound produced a light buzz sheen on it.

It can also be used in connection with the extraction of a spectral envelope in order to alter the contour shape.

Also see: [FORMANTS GET](#), [FORMANTS PUT](#) and [HILITE TRACE](#).

End of FOCUS EXAG

FOCUS FOCUS – Focus spectral energy onto the peaks in the spectrum

Usage

focus focus *inanalysisfile outanalysisfile* **-fN** | **-pN** [-i] *pk bw* [-b*bt*] [-t*tp*] [-s*val*]

Typical command line:

```
focus focus raspdt.ana raspdtfoc.ana -f1 16 0.09 -b20 -t20000
```

Parameters

inanalysisfile – input analysis file

outanalysisfile – output analysis file

-f extract formant envelope linear frequency-wise, using 1 point for every *N* equally-spaced frequency-channels

-p extract formant envelope linear pitchwise, using *N* equally-spaced pitch-bands per octave

-i quicksearch for formants (less accurate)

pk (maximum) number of peaks to find Range 1 – 16

bw bandwidth of peak-centred filters, in octaves

-b*bt* bottom frequency at which to start the search for peaks

-t*tp* top frequency at which to end the search for peaks

-s*val* the number of windows over which the peaks are averaged. This is an attempt to retain only peaks which are STABLE over time. Range: 2 – 4097. Default is 9.

bw (bandwidth), *bt* (bottom frequency) & *tp* (top frequency) may vary over time

Understanding the FOCUS FOCUS Process

The program searches out the most prominent and persistent spectral peaks in a sound, creates a set of filters centred on those peaks, and then filters the sound with those filters. This focuses the spectral energy of the sound around those persistent spectral peaks.

The *channels* and *bandwidth* parameters are both the key sensitive areas affecting the outcome. Increasing the number of channels moves the output towards a BLUR TRACE result, while wider bandwidths, as one would expect, 'let through' more of the original sound.

I found that there was a good balance between original and processed sound with these parameter settings: *Channels* = 1, *Peaks* = 16, *Bandwidth* = 0.09, *Lofrq* = 20 and *HiFrq* = 20000.

Musical Applications

The results depend very much on the spectral configuration of the source, so this function needs to be explored on a 'try it and see' basis. On the whole, it appears to be an alternative approach to BLUR TRACE type effects, i.e., data reduction. However, trials with a pitched sound did produce fairly similar results.

Also see: [HILITE TRACE](#).

End of FOCUS FOCUS

FOCUS FOLD – Octave-transpose spectral components into a specified frequency range

Usage

focus fold *inanalysisfile outanalysisfile lofrq hifrq -x*

Typical command line:

focus fold raspdt.ana raspdtfold.ana 100 900 -x

Parameters

inanalysisfile – input analysis file

outanalysisfile – output analysis file

lofrq lowest frequency of range into which the spectrum is folded

hifrq highest frequency of range into which the spectrum is folded

-x fuller spectrum

lofrq and *hifrq* may vary over time

Understanding the FOCUS FOLD Process

This process octave transposes all the analysis data of a sound into a specified frequency range. Anything below the range is octave-transposed up until it is in (the bottom of) the range. Anything above gets octave-transposed down until it is in (the top of) the range.

If more than two pieces of transposed information compete for the same channel in the resultant file, the data of the lowest amplitude is thrown away.

Musical Applications

The key here is the frequency band: high wide it is and where in the frequency range it is placed. Fairly narrow bands seem the most useful for this function, otherwise too much of the original comes through, so there is not very much change to the sound.

A low-placed range such as 100 to 900 (Hz) produces a deep, muffled sound, with the amount of internal modulation dependent on how much the spectrum of the original input changes. With a rasping spring sound, this produced an 'eerie wind', such as one might imagine a passing spaceship to sound like. On a pitched sound, these parameters resulted in a double pitch: the original and the folded down.

A high-placed range such as 1800 to 3200 yields a fairly high and thin, slightly ringing sound. With the rasping spring sound, this produced a somewhat ethereal sound. On a pitched sound, these parameters were rather high, and not very much of the input remained.

The *fuller spectrum* option does seem to strengthen the effect, perhaps bringing through a little more of the original sound.

This process will focus the energy in a tight band while retaining the pitch characteristics (because of the octave-transpose). It is therefore a way to intensify a sound. However there may also be a radical and somewhat unpredictable transformation of the sound, depending on the input. It seems most effective with sounds containing many and complex partials.

SEE ALSO: [SPECFOLD](#), [STRANGE INVERT](#) (NB: seems to overload easily).

End of FOCUS FOLD

FOCUS FREEZE – Freeze spectral characteristics in a sound, at given times

Usage

focus freeze mode *inanalysisfile outanalysisfile datafile*

Typical command line:

focus freeze 3 rasprd.ana rasprdfz.ana freeze.txt

(See example datafiles in the text below.)

Modes

- 1** freeze channel amplitudes
- 2** freeze channel frequencies
- 3** freeze channel amplitudes & frequencies

Parameters

inanalysisfile – input analysis file

outanalysisfile – output analysis file

datafile – text file containing times at which the spectrum is frozen. These times may be preceded by character **markers**:

- a** use window here as freezefreeze window for spectrum AFTER this time
- b** use window here as freezefreeze window for spectrum BEFORE this time

Otherwise, times are *end/start* of freeze established at one of these markers.

Understanding the FOCUS FREEZE Process

FOCUS FREEZE freezes the spectrum on one window's frequencies or amplitudes. If the frequencies are frozen, the amplitudes continue to move, and v.v.s.: if the amplitudes are frozen, the frequencies continue to move.

The following *datafiles* illustrate how **a** (AFTER) and **b** (BEFORE) work.

marker&times	(action)
a1.0	(freeze AFTER this time)
3.0	(unfreeze)
a4.0	(freeze AFTER this time)
5.0	(unfreeze)

Here we take the values of the window at 1.0 sec. and freeze them until the 3 sec. mark. Then the window is unfrozen, and the sound continues as normal until 4.0, whereupon it is frozen again with the values at 4.0 sec., until 5.0. (The action comments are of course not included in the file!). Similarly:

marker&times	(action)
2.0	(no freezing action here)
b3.0	(value taken here and applied backwards until 2.0)
4.0	(unfrozen between 3.0 and 4.0)
b5.0	(value taken here and applied backwards until 4.0)

One can also mix **a** and **b** as long as the two operations don't overlap. E.g.,

marker&times	(action)
a2.0	(freezing starts here, value taken from 2.0)
3.0	(unfreeze)
4.0	(freezing starts here, value taken from 5.0)
b5.0	(value taken from here and applied backwards until 4.0)

Musical Applications

This process prolongs specific timbral configurations in the sound, so it is a way of stopping the continuous spectral evolution which normally takes place in a sound. The results will vary enormously depending on the source and of course the contents of the spectrum at the freeze points. Static and stepped qualities are introduced into the sound by this process.

End of FOCUS FREEZE



FOCUS HOLD – Hold sound spectrum, at given times

Usage

focus hold *inanalysisfile outanalysisfile datafile*

Typical command line:

```
focus hold raspdt.ana raspdthold.ana holds.txt
```

(See example datafile in the text below.)

Parameters

inanalysisfile – input analysis file

outanalysisfile – output analysis file

datafile – contains a list of paired *times* at which the spectrum is held and a *hold-duration* for each time. These data items must be paired correctly.

Understanding the FOCUS HOLD Process

The process expands each hold window to the duration given, before proceeding to the next window. The output file is therefore longer than the input file. The timbral quality of the held sections depends on what is happening at the point(s), the window(s), where the holds occur.

Musical Applications

This is like SPEC MAGNIFY except that a list of times and hold-durations can be given. The key is to identify the times when the most interesting spectral data is present. Alternatively, you may just want to introduce some additional sustaining. Note that the soundfile plays normally at the times when the holds finish. This means that you can create outputs which mix normal and held sections in various ways.

For example, a useful objective is to alternate the original sound with the held spectrum. This is done in the following file (for a 5 sec. source), where the length of time held (0.6 sec.) is shorter than the time-interval to the next hold point (1 sec.). Thus the sound is 'released' after 1 second and the original sounds for 0.4 second until the next hold point is reached:

Hold_start_time	Length_of_time_held
0.5	0.6
1.5	0.6
2.5	0.6
3.5	0.6
4.5	0.4

The above alternation of original and held was very noticeable with a complex input such as the spring strum (raspdt.wav). However, with a clear, steady pitch such as a horn sound, there were some light surface changes, but overall a smooth flow between original and held – this could therefore be a useful way to extend, i.e., lengthen, a pitched sound, as it adds the held sections to the duration of the soundfile.

End of FOCUS HOLD

FOCUS STEP – Step-frame through sound by freezing spectrum at regular time intervals

Usage

focus step *inanalysisfile outanalysisfile timestep*

Typical command line:

focus step raspdt.ana raspdtstep.ana 0.1

Parameters

inanalysisfile – input analysis file

outanalysisfile – output analysis file

timestep duration of the steps. Must be \geq the duration of two analysis frames. The value here is rounded internally to a multiple of analysis frame time.

Understanding the FOCUS STEP Process

This is an automatic FREEZE of the spectrum at intervals of *timestep*, similar to the familiar 'sample-hold' effect. The main thing here is that the time intervals are *regular*. Thus the spectrum of the sound, which normally is constantly changing, is frozen with the values at the beginning of each *timestep*.

The effect with fairly large *timesteps* can be a relatively slow-moving sample-hold effect. But as the duration of the *timestep* was shortened, the complex source developed regular 'jangling' effect, while the clearly pitched source only changed its tone.

The 'jangling' effect seems to occur with *timestep* at about 0.1 or even 0.05. At 0.25, the sample-hold is quite noticeable, rather slow at 0.5 and probably too slow at 1.0 – but this setting could be used to produce sections of steady-state material that can then be CUT and used elsewhere.

Timestep can be very small. The Phase Vocoder reference manual explains that the value shown by DIRSF for the sample rate of an analysis file is actually frame rate: the number of analysis frames per second, e.g. 172 or 344 etc. (The sample rate and value for **-N** both alter the frame rate.) So if you divide 1 sec by the frame rate, you will learn the duration of one frame: e.g., $1/344 = 0.0029$ (2.9ms). As *timestep* gets larger, timbral alteration increases.

Musical Applications

It is recommended that you start exploring this function by using very small values for *timestep*. This should produce a granulation effect. But note that it is doing this without creating enveloped 'grains' as does MODIFY BRASSAGE. It is working directly on the spectrum of the sound.

End of FOCUS STEP

SPECFOLD – Fold, invert or randomise the spectrum

NB: This is a newer version of **FOCUS FOLD**.

Usage

specfold specfold 1 *inanalfile outanalfile stt len cnt* [-a]
specfold specfold 2 *inanalfile outanalfile stt len* [-a]
specfold specfold 3 *inanalfile outanalfile stt len seed* [-a]

Modes

- 1 Fold the spectrum.
- 2 Invert the spectrum.
- 3 Randomise the spectrum.

Parameters

inanalfile – input analysis file.
outanalfile – output analysis file.
stt – lowest spectral channel to process (Range: ≥ 1)
len – number of channels to process (≥ 4); for folding (Mode 1) this must be an even number.
cnt – (Mode 1) number of folds (less than *len*).
seed – (Mode 3) seed value sets a specific random permutation of channels. (Range: 1 to 64)
-a – process the amplitudes only. Default: process frequencies and amplitudes.

Understanding the SPECFOLD Process

SPECFOLD is a newer alternative to the processes **FOCUS FOLD** and **SPECNU SLICE** (Mode 5: Invert spectrum), offering also a randomising of spectral frequencies. For all three modes, the frequency range to work within is set by the lowest channel (*stt*) and the number of channels *len* to process above that. (The utility SPECINFO FREQUENCY reports the centre frequency for any given channel, while SPECINFO CHANNEL gives the channel for a given frequency.)

The folding process in Mode 1 is as follows:

- Channels in the fold group are split into two equal sets of consecutive channels.
- Values in the first set are moved from channel *N* to channel *2N*.
- Values in the second set are assigned to remaining channels, in descending order.
- *cnt* sets the number of fold groups to fit into the frequency range defined by *stt* and *len*.

Mode 2 inverts the spectrum, putting high frequency channels into low ones and vice versa, while Mode 3 randomly permutes the channels within the channel range.

Setting the **-a** flag processes just the amplitudes, not the frequencies; in a typical sound with harmonics reducing in amplitude as frequency rises, this should make the sound brighter. Overall, the channel settings tend to brighten the sound considerably, above the lowest channel (set this higher to preserve the bass).

Related Functions

SEPCNU SLICE Mode 5 (invert)
STRANGE INVERT
SPECFNU: RANDOMISE PARTIALS (Mode 19)

End of SPECFOLD

SUPERACCU – Sustain each spectral band until louder data appears in that band

Usage

superaccu superaccu 1 *inanalysisfile outanalysisfile* [-**ddecay**] [-**gglis**] [-**r**]
superaccu superaccu 2 *inanalysisfile outanalysisfile* [-**ddecay**] [-**gglis**] [-**r**]
superaccu superaccu 3 *inanalysisfile outanalysisfile tuningtextfile* [-**ddecay**] [-**gglis**] [-**r**]
superaccu superaccu 4 *inanalysisfile outanalysisfile tuningtextfile* [-**ddecay**] [-**gglis**] [-**r**]

Example command line to sustain until louder data is present:

```
superaccu superaccu 3 rasp.ana raspsuperaccu.ana supaccutune.txt -d0.5 -g0.9  
supaccutune.txt: 60 63 67 70 73
```

(This example seemed to produce rather too much 'resonant tail' – the last 10 seconds were inaudible; the input file was 10 seconds long. So you may need to CUT some of the latter part of the output analysis file with SPEC CUT.)

Modes

- 1 Operates like [FOCUS ACCU](#)
- 2 Forces the (start of) resonances to the tempered scale
- 3 The frequencies are specified in a 'tuning' file (a 'harmonic set')
- 4 The frequencies and their octaves are specified in a 'tuning' file (a harmonic field)

Parameters

inanalysisfile – input analysis file

outanalysisfile – output analysis file

tuningtextfile – text file containing the set of frequencies in (possibly fractional) MIDI values (1-128) for tuning the onsets of the resonance

[-**ddecay**] – sustained channel data attenuates by the factor *decay* per second. Possible Range: >0.0 to 0.9 Default: 1.0 (no attenuation)

[-**gglis**] – sustained channel data glissandos at *glis* octaves per second. Approximate Range: -11.7 to 11.7 Default: 0

[-**r**] – reassign the glissandoing pitches to appropriate channels

Understanding the SUPERACCU Process

SUPERACCU is an extension of [FOCUS ACCU](#) with additional modes to tune the sustained spectral channels to a resonance template. The tuning has a remarkable effect on speech, especially in Mode 2. Warning: High values of *decay* can generate very long soundfiles.

End of SUPERACCU