



SPECFNU: SPECTRAL FUNCTIONS UNDER FORMANTS

(with Command Line Usage)

Modify spectral shape in relation to formant peaks, or show formant data.

SPECFNU is in effect a group of 23 individual processes, all of which manipulate frequency analysis files to modify the spectral shape while retaining the original formants. This is particularly important for processing vocal sounds, which have well-defined formant frequencies, but also applies to any sound where you wish to retain the original formant peaks.

For vocal sounds, some understanding of formants and the key formant frequencies is useful. See for example: the article [Harmonics vs formants](#) or the [Wikipedia](#) article.

The command line usage in each case is:

specfnu specfnu 1-23 inanalfile outanalfile [parameters]

SPECFNU MODES:

1 NARROW FORMANTS

Steepen skirts of formant peaks by power factor

2 SQUEEZE SPECTRUM

Squeeze spectrum around specified formant

3 INVERT FORMANTS

Formant peaks become troughs, and troughs peaks

4 ROTATE FORMANTS

Formant peaks & freqs move up (or down) spectrum, re-appearing at foot (or top) on reaching formants' edge

5 SPECTRAL NEGATIVE

Spectral values inverted for each channel

6 SUPPRESS FORMANTS

Suppress the selected formant(s)

7 GENERATE FILTER(S) FROM FORMANTS

Output Varibank filter data based on formant peaks

8 MOVE FORMANTS BY

Displace individual formants by a Hz value

9 MOVE FORMANTS TO

Displace individual formants to specified frequencies

10 ARPEGGIATE

Arpeggiate partials of sound, under formants

11 OCTAVE-SHIFT

Octave-shift pitch of sound, under formants

12 TRANSPOSE

Transpose pitch of sound, under formants

13 FREQ-SHIFT

Frequency shift partials of source sound, under formants

14 RESPACE PARTIALS

Respace partials in source spectrum, retaining formants.

15 PITCH-INVERT

Invert pitch of sound, under formants

16 PITCH-EXAGG/SMOOTH

Exaggerate/Smooth pitch line, under formants

17 PITCH-QUANTISE

Force pitch onto pitch field, under formants

18 PITCH-RANDOMISE

Randomise pitch of source, under formants

19 RANDOMISE PARTIALS

Random-shift partials of sound, under formants

20 SEE SPEC ENVELOPES

Outputs viewable (not playable) soundfile

21 SEE SPEC PEAKS/TROUGHS

Print textfile of frequencies of peaks & troughs per window

22 GET LOUDNESS TROUGHS

Print textfile of times-of-troughs between syllables

23 SINE SPEECH

Convert formant frequencies to sine tones. A single sine wave represents each formant.

Return to Main Index for the CDP System.

Return to Spectral Index

SPECFNU: Usage and Parameters for each Mode

SPECFNU MODE 1

NARROW FORMANTS – Steepen skirts of formant peaks by power factor

Usage

specfnu specfnu 1 *inanalfil outanalfil narrow* [-**ggain**] [-**ooff**] [-**t**] [-**f**] [-**s**] [-**x** |**-k**] [-**r**]

Parameters

inanalfil – input analysis file.

outanalfil – output analysis file.

narrow – narrowing of individual formant peaks; may vary through time. (Range: 1 to 1000).

-**ggain** – output gain or attenuation. (Range: 0.01 to 10)

-**ooff** – text list of which formants to suppress (up to 4),

e.g. "1" means suppress 1st, while "134" means suppress 1st, 3rd and 4th.

-**t** – zero top of spectrum (above fourth formant).

-**f** – force lowest formant to use fundamental frequency as peak.

-**s** – use a short window for extracting spectral envelope (default is large window).

-**x** – exclude non-harmonic partials.

-**k** – kill harmonic partials.

-**r** – replace unpitched (or extremely quiet) windows by silence.

Understanding SPECFNU Mode 1: NARROW FORMANTS

Narrowing the formants focuses on the spectral peaks and therefore acts like an increasingly tight filter: the material between the peaks is diminished and eventually disappears. *Narrow* can normally be set quite high (in 100s) and is also time-variable, which will alter the timbre as the sound progresses. You can also choose to omit some of the first four formants altogether (*off*). As the overall level is reduced, it is often useful to set some extra *gain* to compensate.

Spectral window-size: if the "force fundamental" option (-f) is used, selecting a large window will find the fundamental, whereas a short window (option -s) may not.

The process may be compared with [FOCUS EXAG](#).

Musical Applications

Related Function

[FOCUS EXAG](#)

End of SPECFNU Mode 1: NARROW FORMANTS

SPECFNU MODE 2

SQUEEZE SPECTRUM AROUND FORMANT – Squeeze around specified formant

Usage

specfnu specfnu 2 *inanal outanal squeeze centre* [-**ggain**] [-**t**] [-**f**] [-**s**] [-**x** | -**k**] [-**r**]

Parameters

inanal – input analysis file.

outanal – output analysis file.

squeeze – squeeze factor; may vary through time. (Range: 1 to 10)

centre – formant peak at centre of squeeze. (Range: 1 to 4)

-ggain – output gain or attenuation. (Range: 0.01 to 10)

-t – squeeze around trough above specified peak.

-f – force lowest formant to use fundamental frequency as peak.

-s – use a short window for extracting spectral envelope (default is large window).

-x – exclude non-harmonic partials.

-k – kill harmonic partials.

-r – replace unpitched (or extremely quiet) windows by silence.

Understanding SPECFNU Mode 2: SQUEEZE SPECTRUM

This mode squeezes the spectrum around a specified formant. It is not yet clear what this means technically, but the result typically turns a female speaker into a deep hoarse voice.

Spectral window-size: if the "force fundamental" option (-f) is used, selecting a large window will find the fundamental, whereas a short window (option -s) may not.

Musical Applications

Related Function

SPECFNU SQUEEZE

End of SPECFNU Mode 2: SQUEEZE SPECTRUM

SPECFNU MODE 3 INVERT FORMANTS – Formant peaks become troughs, and troughs peaks

Usage

specfnu specfnu 3 *inanal outanal vibrate* [-**ggain**] [-**s**] [-**x** | -**k**] [-**r**]

Parameters

inanal – input analysis file.

outanal – output analysis file.

vibrate – if not zero, cycle between original and inverted values at this frequency; may vary through time. (Range: 0 to 300 Hz)

-**ggain** – output gain or attenuation. (Range: 0.01 to 10)

-**s** – use a short window for extracting spectral envelope (default is large window).

-**x** – exclude non-harmonic partials.

-**k** – kill harmonic partials.

-**r** – replace unpitched (or extremely quiet) windows by silence.

Understanding SPECFNU Mode 3: INVERT FORMANTS

A spectral peak is an area of higher amplitude among the frequencies of the spectrum. A trough is an area of low amplitude between spectral peaks. The INVERT mode swaps these, so that the peaks become troughs and vice versa. This should make the sound noisier, with buzzy high frequencies.

Vibrate allow you to oscillate between the original and inverted states.

Musical Applications

End of SPECFNU Mode 3: INVERT FORMANTS

SPECFNU MODE 4 ROTATE FORMANTS – Formant peaks & frequencies move up (or down) spectrum

Usage

specfnu specfnu 4 *inanal outanal rspeed* [-**ggain**] [-**s**] [-**x** | -**k**] [-**r**]

Parameters

inanal – input analysis file.

outanal – output analysis file.

rspeed – How quickly the spectrum rotates every second; may vary through time. (Range: -300 to +300 rotations per second)

-**ggain** – output gain or attenuation. (Range: 0.01 to 10)

-**s** – use a short window for extracting spectral envelope (default is large window).

-**x** – exclude non-harmonic partials.

-**k** – kill harmonic partials.

-**r** – replace unpitched (or extremely quiet) windows by silence.

Understanding SPECFNU Mode 4: ROTATE FORMANTS

In this ROTATE mode, the formants move up (or down) in frequency, reappearing at the foot (or top) on reaching formants' edge. The result is a rather wobbly sound.

To move the formants down instead of up, set a negative value for *rspeed*. Note that this is time-variable, so the formants could rotate both ways through the sound.

Musical Applications

End of SPECFNU Mode 4: ROTATE FORMANTS

SPECFNU MODE 5 SPECTRAL NEGATIVE – Spectral values are inverted for each channel

Usage

specfnu specfnu 5 *inanal outanal* [-**ggain**] [-**f**]

Parameters

inanal – input analysis file.

outanal – output analysis file.

-**ggain** – output gain or attenuation. (Range: 0.01 to 10)

-**f** – flat: does not re-envelope the output spectrum.

Understanding SPECFNU Mode 5: SPECTRAL NEGATIVE

This mode inverts the spectral values for each channel, usually producing a highly distorted output. It appears to flip the spectral envelope in some way (still to be clarified). Users are recommended to set the -f ('flat') flag, at least at first, particularly as the only other control (*gain*) often produces this obscure error message: SUPPRESS NON-HARMONICS WITH SUPPRESS-HARMONICS WILL PRODUCE ZERO SIGNAL LEVEL.

Musical Applications

End of SPECFNU Mode 5: SPECTRAL NEGATIVE

SPECFNU MODE 6

SUPPRESS FORMANTS – Suppress the selected formant(s)

Usage

specfnu specfnu 6 *inanalfile outanalfile formantlist* [-**ggain**] [-**s**] [-**x**]

Parameters

inanalfile – input analysis file.

outanalfile – output analysis file.

formantlist – text list of which formants to suppress (up to 4),

e.g. "1" means suppress 1st, while "134" means suppress 1st, 3rd and 4th.

-ggain – amplitude gain attenuation. (Range: 0.1 to 10)

-s – use a short window for extracting spectral envelope (default is large window).

-x – exclude non-harmonic partials.

Understanding SPECFNU Mode 6: SUPPRESS FORMANTS

This mode suppresses the selected formant(s), which is a selective type of filtering. Suppressing formant 1 makes the sound shriller, whereas suppressing formant 4 may have only a slightly dulling effect on the upper frequencies.

Musical Applications

Related Functions

BLUR SUPPRESS

SUPPRESS PARTIALS

End of SPECFNU Mode 6: SUPPRESS FORMANTS

SPECFNU MODE 7

GENERATE FILTER(S) FROM FORMANTS – Output textfile of Varibank filter data based on formant peaks

Usage

specfnu specfnu 7 *inanalfile outfiltfile datafile filtcnt* [-b*below*] [-k | -i] [-f] [-s]

Parameters

inanalfile – input analysis file.

outfiltfile – output text file in VARIBANK format.

datafile – textfile containing:

- (1) List of times at which source is to be divided into blocks for analysis. Resulting filter step-changes between values created from each block. Value **ZERO** means use the entire source to make one fixed filter.
- (2) Data about grid on which pitches you are searching for must lie. There must first be a marker indicating the grid type. These are:

#HS followed by listed MIDI pitches to search for.

#HF followed by listed MIDI pitches to search for in all octaves.

#SCALE followed by just TWO values:

- (a) the number of equal pitch divisions in an octave.
- (b) MIDI pitch of any pitch to tune the scales to.

#ELACS followed by just THREE values:

- (a) size of "octave" in (possibly fractional) semitones.
- (b) the number of equal pitch divisions in the "octave".
- (c) MIDI-pitch of any pitch to tune the scales to.

filtcnt – (max) no. of peaks from each formant to use. (Default: 1)

-bbelow – try to ensure 1 (or more) pitches below MIDI pitch value "below".

-k – keep relative (summed) loudnesses of peaks as part of filter design.

-i – keep inverse of loudnesses as part of filter design. (Default: all filter amplitudes set to 1.0)

-f – force fundamental as pitch in lowest formant.

-s – use a short window for extracting spectral envelope (default is large window).

Understanding SPECFNU Mode 7: GENERATE FILTERS

There is a clear link between formant peaks and filtering, as the peaks carry the most energy in the sound. This mode produces time-varying filter data in [VARIBANK](#) format, which can then be used by FILTER VARIBANK itself, or other processes which employ this format, such as TUNEVARY and SYNFLT.

Trevor Wishart adds: "with a time-varying filter, the transitions between time blocks can be extreme. It's wise to cut the input into segments and produce fixed filters for each, or pre-envelope the input file to zero the level around transitions, deriving appropriate envelope times from the filter file times."

Spectral window-size: if the "force fundamental" option (-f) is used, selecting a large window will find the fundamental, whereas a short window (option -s) may not.

Musical Applications

End of SPECFNU Mode 7: GENERATE FILTERS

SPECFNU MODE 8 MOVE FORMANTS BY – Displace individual formants by a given frequency value

Usage

specfnu specfnu 8 *inanal outanal mov1 mov2 mov3 mov4* [-**ggain**] [-**t**] [-**s**] [-**x** | -**k**] [-**r**]

Parameters

inanal – input analysis file.

outanal – output analysis file.

mov1 – frequency displacement, up or down, of formant 1 (in Hz).

mov2, mov3, mov4 – similarly for the other formants. *Mov1, 2, 3* and 4 may vary through time.

Note that formants moving below zero or above nyquist/2 will disappear.

-**ggain** – output gain or attenuation. (Range: 0.01 to 10)

-**t** – zero top of spectrum.

-**s** – use a short window for extracting spectral envelope (default is large window).

-**x** – exclude non-harmonic partials.

-**k** – kill harmonic partials.

-**r** – replace unpitched (or extremely quiet) windows by silence.

Understanding SPECFNU Mode 8: MOVE FORMANTS BY

This mode moves the formant frequencies by a specified amount up or down. There is a control for each of four formants (*mov1, mov2, mov3, mov4*).

Musical Applications

End of SPECFNU Mode 8: MOVE FORMANTS BY

SPECFNU MODE 9

MOVE FORMANTS TO – Displace individual formants to specified frequencies

Usage

specfnu specfnu 9 *inanal outanal frq1 frq2 frq3 frq4* [-**ggain**] [-**t**] [-**s**] [-**n**] [-**x** | -**k**] [-**r**]

Parameters

inanal – input analysis file.

outanal – output analysis file.

frq1 – new frequency of formant 1 (Hz).

frq2, frq3, frq4 – similarly for the other formants. Frq1, 2, 3 and 4 may vary through time.

-ggain – output gain or attenuation. (Range: 0.01 to 10)

-t – zero top of spectrum.

-n – use narrow formant bands.

Using narrow bands can give counter-intuitive results: it doesn't capture so many harmonic peaks, so sounds duller.

-s – use a short window for extracting spectral envelope (default is large window).

-x – exclude non-harmonic partials.

-k – kill harmonic partials.

-r – replace unpitched (or extremely quiet) windows by silence.

Understanding SPECFNU Mode 9: MOVE FORMANTS TO

This mode complement Mode 8 ("Move formants by"): here the formants are moved to specific frequencies. As with Mode 8, there is a control for each of four formants (*frq1, frq2, frq3, frq4*).

Musical Applications

End of SPECFNU Mode 9: MOVE FORMANTS TO

SPECFNU MODE 10

ARPEGGIATE – Arpeggiate spectrum of source, under formants

Usage

specfnu specfnu 10 *inanal outanal arprate* [-**ggain**] [-**s**] [-**x**] [-**r**] [-**d** | -**c**]

Parameters

inanal – input analysis file.

outanal – output analysis file.

arprate – rate of arpeggiation of spectrum. (Range: -50 to +50 Hz)

-ggain – output gain or attenuation. (Range: 0.01 to 10)

-s – use a short window for extracting spectral envelope (default is large window).

-x – exclude non-harmonic partials.

-r – replace unpitched (or extremely quiet) windows by silence.

-d – arpeggiate downwards (default upwards).

-c – arpeggiate up and down (default upwards).

Understanding SPECFNU Mode 10: ARPEGGIATE

This mode creates a sweeping-filter effect, based on formants. (Use STRETCH TIME to slow this down and hear the process in detail.)

The arpeggiation here is not like that of HILTE ARPEG, which causes a waveform at a specified frequency to sweep through the spectrum of a sound, emphasizing individual partials in turn.

NB: On some PCs (e.g., Windows 7) this operation may give a silent output. This is currently under investigation.

Musical Applications

End of SPECFNU Mode 10: ARPEGGIATE

SPECFNU MODE 11

OCTAVE-SHIFT – Octave-shift pitch of sound, under formants

Usage

specfnu specfnu 11 *inanal outanal octshift* [-**ggain**] [-**parprate**] [-**llocut**] [-**hhicut**] [-**s**] [-**x**] [-**r**] [-**d** | -**c**] [-**f**]

Parameters

inanal – input analysis file.

outanal – output analysis file.

octshift – octave-shift; may vary through time. (Range: -4 to + 4 8ves)

-**ggain** – output gain or attenuation. (Range: 0.01 to 10)

-**parprate** – rate of arpeggiation of spectrum. (Range: -50 to +50 Hz)

-**llocut** – cut off frequencies below this. (Range: 0 to 10000 Hz)

-**hhicut** – cut off frequencies above this. (Range: 50 to 10000 Hz)

-**s** – use a short window for extracting spectral envelope (default is large window).

-**x** – exclude non-harmonic partials.

-**r** – replace unpitched (or extremely quiet) windows by silence.

-**d** – arpeggiate downwards (default upwards).

-**c** – arpeggiate up and down (default upwards).

-**f** – if transposing downwards, fill spectrum top with extra harmonics.

Understanding SPECFNU Mode 11: OCTAVE-SHIFT

This mode transposes the pitch of the sound, while retaining the original formants. It may be compared with REPITCH PITCHSHIFT/TRANSPPOSEF or PITCH OCTMOVE. In tests, the SPECFNU version produced more noticeable artefacts, but these could be suppressed by setting the "silence unpitched" flag (-**r**) and boosting the output level (with *gain*).

You can optionally filter out frequencies below *locut* or above *hicut*, as required. It is also possible to arpeggiate the spectrum. Set *arprate* if you want arpeggiation; otherwise, do NOT set it. Note also the two flags for the direction of arpeggiation (-**d** and -**c**.)

See also: SPECFNU Mode 12.

Musical Applications

Related Functions

[REPITCH TRANSPPOSEF](#)

[PITCH OCTMOVE](#)

[REPITCH PCHSHIFT](#)

End of SPECFNU Mode 11: OCTAVE-SHIFT

SPECFNU MODE 12

TRANSPOSE – Transpose pitch of sound, under formants

Usage

specfnu specfnu 12 *inanal outanal transpos* [-**gain**] [-**parprate**] [-**locut**] [-**hicut**] [-**s**] [-**x**] [-**r**] [-**d**] [-**c**] [-**f**]

Parameters

inanal – input analysis file.

outanal – output analysis file.

transpos – semitone shift; may vary through time. (Range: -48 to + 48 semitones)

-gain – output gain or attenuation. (Range: 0.01 to 10)

-parprate – rate of arpeggiation of spectrum. (Range: -50 to +50 Hz)

-locut – cut off frequencies below this. (Range: 0 to 10000 Hz)

-hicut – cut off frequencies above this. (Range: 50 to 10000 Hz)

-s – use a short window for extracting spectral envelope (default is large window).

-x – exclude non-harmonic partials.

-r – replace unpitched (or extremely quiet) windows by silence.

-d – arpeggiate downwards (default upwards).

-c – arpeggiate up and down (default upwards).

-f – if transposing downwards, fill spectrum top with extra harmonics.

Understanding SPECFNU Mode 12: TRANSPOSE

This mode transposes the pitch of the sound, while retaining the original formants. It may be compared with REPITCH PITCHSHIFT/TRANSPPOSEF or PITCH OCTMOVE. In tests, the SPECFNU version produced more noticeable artefacts, but these could be suppressed by setting the "silence unpitched" flag (**-r**) and boosting the output level (with *gain*).

You can optionally filter out frequencies below *locut* or above *hicut*, as required. It is also possible to arpeggiate the spectrum. Set *arprate* if you want arpeggiation; otherwise, do NOT set it. Note also the two flags for the direction of arpeggiation (**-d** and **-c**.)

See also **SPECFNU Mode 11**.

Musical Applications

Related Functions

REPITCH TRANSPPOSEF

PITCH OCTMOVE

REPITCH PCHSHIFT

End of SPECFNU Mode 12: TRANSPOSE

SPECFNU MODE 13

FREQUENCY SHIFT – Frequency-shift partials of source sound, under formants

Usage

specfnu specfnu 13 *inanal outanal frqshift* [-**ggain**] [-**parprate**] [-**llocut**] [-**hhicut**] [-**s**] [-**x**] [-**r**] [-**d**] [-**c**] [-**f**]

Parameters

inanal – input analysis file.

outanal – output analysis file.

frqshift – frequency shift; may vary through time. (Range: -1000 to +1000 Hz)

-**ggain** – output gain or attenuation. (Range: 0.01 to 10)

-**parprate** – rate of arpeggiation of spectrum. (Range: -50 to +50 Hz)

-**llocut** – cut off frequencies below this. (Range: 0 to 10000 Hz)

-**hhicut** – cut off frequencies above this. (Range: 50 to 10000 Hz)

-**s** – use a short window for extracting spectral envelope (default is large window).

-**x** – exclude non-harmonic partials.

-**r** – replace unpitched (or extremely quiet) windows by silence.

-**d** – arpeggiate downwards (default upwards). (Only set when *arprate* > 0.)

-**c** – arpeggiate up and down (default upwards). (Only set when *arprate* > 0.)

-**f** – if transposing downwards, fill spectrum top with extra harmonics.

Understanding SPECFNU Mode 13: FREQUENCY SHIFT

This mode shifts the frequencies in the sound by a specified amount (producing, for example, inharmonic sounds out of harmonic ones). This process may be compared with STRANGE SHIFT, which does not preserve the formants.

As with Modes 11 and 12, there many artefacts present. To cure this problem, set the "exclude non-harmonic partials" flag (-**x**). (Unlike Modes 11 and 12, do not set the "silence unpitched" flag (-**r**), or the output will be seriously quiet.)

You can optionally filter out frequencies below *locut* or above *hicut*, as required. It is also possible to arpeggiate the spectrum. Set *arprate* if you want arpeggiation; otherwise, do NOT set it. Note also the two flags for the direction of arpeggiation (-**d** and -**c**.)

SEE ALSO: SPECFNU MODE 14 – Respace

Musical Applications

Related Function

FREQ. SHIFT not reatining partials

End of SPECFNU Mode 13: FREQUENCY SHIFT

SPECFNU MODE 14 RESPACE PARTIALS – Respace partials in source spectrum, retaining formants

Usage

specfnu specfnu 14 *inanal outanal respace* [-**ggain**] [-**parprate**] [-**locut**] [-**hicut**] [-**s**] [-**x**] [-**r**] [-**d** | -**c**] [-**f**]

Parameters

inanal – input analysis file.

outanal – output analysis file.

respace – new frequency spacing of partials; may vary through time. (Range: 1 to 1000 Hz)

-ggain – output gain or attenuation. (Range: 0.01 to 10)

-parprate – rate of arpeggiation of spectrum. (Range: -50 to +50 Hz)

-locut – cut off frequencies below this. (Range: 0 to 10000 Hz)

-hicut – cut off frequencies above this. (Range: 50 to 10000 Hz)

-s – use a short window for extracting spectral envelope (default is large window).

-x – exclude non-harmonic partials.

-r – replace unpitched (or extremely quiet) windows by silence.

-d – arpeggiate downwards (default upwards). (Only set when *arprate* > 0.)

-c – arpeggiate up and down (default upwards). (Only set when *arprate* > 0.)

-f – if narrowed spacing, fill spectrum top with extra shifted partials.

Understanding SPECFNU Mode 14: RESPACE PARTIALS

In this mode, *respace* sets a new fundamental for the sound, as well as re-spacing the partials as harmonics of this fundamental, while retaining the existing formants.

With a vocal source, the "exclude non-harmonic partials" flag (**-x**) seemed generally redundant, while the "silence unpitched" flag (**-r**) produced a high-pass filter effect. However, with a non-harmonic source (handbell), the process produced many artefacts which, though suppressed by the "silence unpitched" flag (**-r**), left gaps in the sound.

You can optionally filter out frequencies below *locut* or above *hicut*, as required. It is also possible to arpeggiate the spectrum. Set *arprate* if you want arpeggiation; otherwise, do NOT set it. Note also the two flags for the direction of arpeggiation (**-d** and **-c**.)

SEE ALSO: SPECFNU MODE 13 – Frequency Shift and MODE 19 – Randomise Partials

Musical Applications

End of SPECFNU Mode 14: RESPACE PARTIALS

SPECFNU MODE 15

PITCH-INVERT – Invert pitch of sound, under formants

Usage

specfnu specfnu 15 *inanal outanal map about* [-**ggain**] [-**parprate**] [-**llocut**] [-**hhipcut**] [-**blpch**] [-**thipch**] [-**s**] [-**x**] [-**r**] [-**d**] [-**c**]

Parameters

inanal – input analysis file.

outanal – output analysis file.

map – set to zero if no mapping is required; otherwise, map is a textfile of paired values showing how the intervals (in semitones) are to be mapped onto their inversions. (Range: +/-96.000000 semitones, possibly fractional)

about – pitch about which to invert the pitch of the source; may vary through time. (Range: 4 to 127, as MIDI, possibly fractional)

If the value zero is entered, the mean pitch of the input is used.

-ggain – output gain or attenuation. (Range: 0.01 to 10)

-parprate – rate of arpeggiation of spectrum. (Range: -50 to +50 Hz)

-llocut – cut off frequencies below this. (Range: 0 to 10000 Hz)

-hhipcut – cut off frequencies above this. (Range: 50 to 10000 Hz)

-blpch – minimum acceptable pitch. (Range: 1-127, as MIDI, possibly fractional)

-thipch – maximum acceptable pitch. (Range: 1-127, as MIDI, possibly fractional)

-s – use a short window for extracting spectral envelope (default is large window).

-x – exclude non-harmonic partials.

-r – replace unpitched (or extremely quiet) windows by silence.

-d – arpeggiate downwards (default upwards). (Only set when *arprate* > 0.)

-c – arpeggiate up and down (default upwards). (Only set when *arprate* > 0.)

Understanding SPECFNU Mode 15: PITCH-INVERT

This mode inverts the pitch around a given MIDI pitch, so that frequencies above this frequency are mirrored below it and vice versa. It may be compared with REPITCH INVERT, which also inverts the pitch, or with SPECNU SLICE 5 (invert frequency).

In tests, a MIDI pitch of 60 (middle C) or lower gave coherent results, whereas a high wrap-around pitch (MIDI 80) produced many artefacts.

You can optionally restrict the pitch range with *lopch* and *hipch* and filter out frequencies below *llocut* or above *hhipcut*, as required. It is also possible to arpeggiate the spectrum. Set *arprate* if you want arpeggiation; otherwise, do NOT set it. Note also the two flags for the direction of arpeggiation (**-d** and **-c**.)

Musical Applications

Related Function

REPITCH INVERT

End of SPECFNU Mode 15: PITCH-INVERT

SPECFNU MODE 16

PITCH-EXAGGERATE/SMOOTH UNDER FORMANTS – Exaggerate or smooth the pitch-line, retaining the formants.

Usage

specfnu specfnu 16 *inanal outanal about rang* [-**ggain**] [-**parprate**] [-**llocut**] [-**hhipcut**] [-**blpoch**] [-**thipch**] [-**T**] [-**F**] [-**M**] [-**A**] [-**B**] [-**s**] [-**x**] [-**r**] [-**d**] [-**c**]

Parameters

inanal – input analysis file.

outanal – output analysis file.

about – pitch about which to exaggerate or smooth pitch line of the source; may vary through time.

(Range: 4 to 127, as MIDI, possibly fractional)

If the value zero is entered, the mean pitch of the input is used.

rang – expand or contract the range of the pitch line. (Range: 0 to 1)

-**ggain** – output gain or attenuation. (Range: 0.01 to 10)

-**parprate** – rate of arpeggiation of spectrum. (Range: -50 to +50 Hz)

-**llocut** – cut off frequencies below this. (Range: 0 to 10000 Hz)

-**hhipcut** – cut off frequencies above this. (Range: 50 to 10000 Hz)

-**blpoch** – minimum acceptable pitch. (Range: 1-127, as MIDI, possibly fractional)

-**thipch** – maximum acceptable pitch. (Range: 1-127, as MIDI, possibly fractional)

-**T** – tie to top of pitch range.

-**F** – tie to foot of pitch range.

-**M** – tie to range middle: only if "-T" and "-F" flags are also set.

-**A** – do range change above mean only.

-**B** – do range change below mean only.

-**s** – use a short window for extracting spectral envelope (default is large window).

-**x** – exclude non-harmonic partials.

-**r** – replace unpitched (or extremely quiet) windows by silence.

-**d** – arpeggiate downwards (default upwards). (Only set when *arprate* > 0.)

-**c** – arpeggiate up and down (default upwards). (Only set when *arprate* > 0.)

Understanding SPECFNU Mode 16: PITCH-EXAGG/SMOOTH

This mode aims to exaggerate the pitch line, under the existing formants. The range parameter (*rang*) sets the degree of possible change; some settings may produce a consistent transposition.

The various flags set options for restricting the pitch range that might be exaggerated.

You can optionally restrict the pitch range with *lopch* and *hipch* and filter out frequencies below *locut* or above *hipcut*, as required. It is also possible to arpeggiate the spectrum. Set *arprate* if you want arpeggiation; otherwise, do NOT set it. Note also the two flags for the direction of arpeggiation (-**d** and -**c**.)

Musical Applications

Related Function

REPITCH EXAG

SPECFNU MODE 17

PITCH-QUANTISE – Force pitch onto a specified pitch field, under formants

Usage

specfnu specfnu 17 *inanal outanal datafile* [-**ggain**] [-**parprate**] [-**llocut**] [-**hhicut**] [-**blotch**] [-**thipch**] [-**s**] [-**x**] [-**r**] [-**d**] [-**c**] [-**o**] [-**n**]

Parameters

inanal – input analysis file.

outanal – output analysis file.

datafile – textfile containing data about pitches to quantise to, as follows:

There must first be a marker indicating the grid type. These are:

#HS followed by listed MIDI pitches to quantise to (harmonic set).

#HF followed by listed MIDI pitches to quantise to in all octaves (harmonic field).

#THF followed by lines each with time plus a list of MIDI pitches (timed harmonic field).

First time must be zero & times must increase.

Each MIDI list must be the same length.

To change the number of pitches from line to line, duplicate values.

#SCALE followed by just TWO values:

(a) the number of equal pitch divisions in an octave.

(b) MIDI pitch of any pitch to tune the scales to.

#ELACS followed by just THREE values:

(a) size of "pseudo-octave" in (possibly fractional) semitones.

(b) the number of equal pitch divisions in the pseudo-octave.

(c) MIDI-pitch of any pitch to tune the scales to.

-ggain – output gain or attenuation. (Range: 0.01 to 10)

-parprate – rate of arpeggiation of spectrum. (Range: -50 to +50 Hz)

-llocut – cut off frequencies below this. (Range: 0 to 10000 Hz)

-hhicut – cut off frequencies above this. (Range: 50 to 10000 Hz)

-blotch – minimum acceptable pitch. (Range: 1-127, as MIDI, possibly fractional)

-thipch – maximum acceptable pitch. (Range: 1-127, as MIDI, possibly fractional)

-s – use a short window for extracting spectral envelope (default is large window).

-x – exclude non-harmonic partials.

-r – replace unpitched (or extremely quiet) windows by silence.

-d – arpeggiate downwards (default upwards). (Only set when *arprate* > 0.)

-c – arpeggiate up and down (default upwards). (Only set when *arprate* > 0.)

-o – allow ornaments in the quantised pitch-line.

-n – no smoothing of transitions between pitches.

Understanding SPECFNU Mode 17: PITCH-QUANTISE

Mode 17 quantizes the pitch content in a sound while preserving the original formants. The actual pitches are forced onto the nearest listed pitch value. The process needs a text *datafile* about pitches to quantize to.

The format of the pitch data file is much more flexible than that of REPITCH QUANTISE. There are five possible types, each preceded by a different initial text marker: 1. #HS (harmonic set) + a list of pitches; 2. #HF (harmonic field) + a list of pitches to be duplicated in all 8ves; 3. #THF (timed harmonic field) + a time-changing list of pitches; 4. #SCALE: an equal-tempered scale dividing the octave, plus a pitch to tune the scale to; e) #ELACS: as d) but using a pseudo-octave of any size.

You can optionally restrict the pitch range with *lopch* and *hipch* and filter out frequencies below *locut* or above *hicut*, as required. There is also an arpeggiation option, triggered by the optional *arprate*, with options to arpeggiate downwards and/or upwards instead of up.

Musical Applications

Related Function

REPITCH QUANTISE

End of SPECFNU Mode 17: PITCH-QUANTISE

SPECFNU MODE 18

PITCH RANDOMISE – Randomise pitch of source, under formants

Usage

specfnu specfnu 18 *inanal outanal datafile range slew* [-**ggain**] [-**parprate**] [-**llocut**] [-**hhicut**] [-**blotch**] [-**thipch**] [-**s**] [-**x**] [-**r**] [-**d** | -**c**] [-**o**] [-**n**] [-**k**]

Parameters

inanal – input analysis file.

outanal – output analysis file.

datafile – if set to zero, no pre-quantisation of pitch takes place;

Otherwise, this a textfile containing data about pitches to quantise to, as follows:

There must first be a marker indicating the grid type. These are:

#HS followed by listed MIDI pitches to quantise to.

#HF followed by listed MIDI pitches to quantise to in all octaves.

#THF followed by lines each with time plus a list of MIDI pitches.

First time must be zero & times must increase.

Each MIDI list must be the same length.

To change the number of pitches from line to line, duplicate values.

#SCALE followed by just TWO values:

(a) the number of equal pitch divisions in an octave.

(b) MIDI pitch of any pitch to tune the scales to.

#ELACS followed by just THREE values:

(a) size of "pseudo-octave" in (possibly fractional) semitones.

(b) the number of equal pitch divisions in the pseudo-octave.

(c) MIDI-pitch of any pitch to tune the scales to.

range – semitone range within which random offsets are generated. (Range: 0 to 96 semitones)

slew – relationship between (possible) upward & downward random variation,

e.g. 2: uprange = 2 * downrange

0.5: uprange = 0.5 * downrange.

(Ranges 0.1 to 10, where 1 has no effect)

-**ggain** – output gain or attenuation. (Range: 0.01 to 10)

-**parprate** – rate of arpeggiation of spectrum. (Range: -50 to +50 Hz)

-**llocut** – cut off frequencies below this. (Range: 0 to 10000 Hz)

-**hhicut** – cut off frequencies above this. (Range: 50 to 10000 Hz)

-**blotch** – minimum acceptable pitch. (Range: 1-127, as MIDI, possibly fractional)

-**thipch** – maximum acceptable pitch. (Range: 1-127, as MIDI, possibly fractional)

-**s** – use a short window for extracting spectral envelope (default is large window).

-**x** – exclude non-harmonic partials.

-**r** – replace unpitched (or extremely quiet) windows by silence.

-**d** – arpeggiate downwards (default upwards). (Only set when *arprate* > 0.)

-**c** – arpeggiate up and down (default upwards). (Only set when *arprate* > 0.)

-**o** – open up fast movement in the quantised pitch-line.

-**n** – no smoothing of transitions between pitches.

-**k** – kill-off formant reshaping.

Understanding SPECFNU Mode 18: PITCH-RANDOMISE

Mode 18 randomises pitch under the existing formants. It may be compared with REPITCH RANDOMISE.

Before randomising, the pitch may optionally be quantised as in Mode 17, with the same type of datafile.

As with Mode 17, you can optionally restrict the pitch range with *lopch* and *hipch* and filter out frequencies below *locut* or above *hicut*, as required. There is also an arpeggiation option, triggered by the optional *arprate*, with options to arpeggiate downwards and/or upwards instead of up.

Musical Applications

Related Function

REPITCH RANDOMISE

End of SPECFNU Mode 18: PITCH-RANDOMISE

SPECFNU MODE 19

RANDOMISE PARTIALS – Random-shift partials, under formants

Usage

specfnu specfnu 19 *inanal outanal rand* [-**ggain**] [-**parprate**] [-**llocut**] [-**hhicut**] [-**s**] [-**x**] [-**r**] [-**d**] [-**c**] [-**o**]

Parameters

- inanal* – input analysis file.
- outanal* – output analysis file.
- rand* – randomisation of partial frequencies; may vary through time. (Range: 0 to 1)
- ggain** – output gain or attenuation. (Range: 0.01 to 10)
- parprate** – rate of arpeggiation of spectrum. (Range: -50 to +50 Hz)
- llocut** – cut off frequencies below this. (Range: 0 to 10000 Hz)
- hhicut** – cut off frequencies above this. (Range: 50 to 10000 Hz)
- s** – use a short window for extracting spectral envelope (default is large window).
- x** – exclude non-harmonic partials.
- r** – replace unpitched (or extremely quiet) windows by silence.
- d** – arpeggiate downwards (default upwards). (Only set when *arprate* > 0.)
- c** – arpeggiate up and down (default upwards). (Only set when *arprate* > 0.)

Understanding SPECFNU Mode 19: RANDOMISE PARTIALS

In Mode 19, the partials are randomised (as opposed to the pitch in Mode 18). The degree of randomisation is set by *rand*.

You can optionally filter out frequencies below *locut* or above *hicut*, as required. There is also an arpeggiation option, triggered by the optional *arprate*, with options to arpeggiate downwards and/or upwards instead of up.

Musical Applications

...

Related Function

SPECFOLD Mode 3 (randomise spectrum)

End of SPECFNU Mode 19: RANDOMISE PARTIALS

SPECFNU MODE 20 SEE SPECTRAL ENVELOPE – output viewable (not playable) soundfile

Usage

specfnu specfnu 20 *inanalfile outpseudosndfile* [-s]

Parameters

inanalfile – input analysis file.

outpseudosndfile – output psuedo-soundfile (DO NOT PLAY THIS!)

-s – use a short window for extracting spectral envelope (default is large window).

Understanding SPECFNU Mode 20: SEE SPECTRAL ENVELOPE

The function creates a viewable (but not playable) soundfile showing the spectral envelope at each window as a block of positive samples.

Musical Applications

End of SPECFNU Mode 20: SEE SPECTRAL ENVELOPE

SPECFNU MODE 21 SEE SPECTRAL PEAKS & TROUGHS – Print textfile of frequencies of peaks & troughs per window

Usage

specfnu specfnu 21 *inanalfile outtextfile* [-s]

Parameters

inanalfile – input analysis file.

outtextfile –output text file of peaks and troughs per window. Output format:

Trough Peak-1 Trough Peak-2 Trough Peak-3 Trough Peak-4 Trough

-s – use a short window for extracting spectral envelope (default is large window).

Understanding SPECFNU Mode 21: SEE SPECTRAL PEAKS/TROUGHS

Musical Applications

End of SPECFNU Mode 21: SEE SPECTRAL PEAKS/TROUGHS

SPECFNU MODE 22

GET LOUDNESS TROUGHS – Print textfile of times-of-troughs between syllables

Usage

specfnu specfnu 22 *inanalfile outtextfile* [-**ssyldur**] [-**ppktrof**] [-**P** | -**B**]

Parameters

inanalfile – input analysis file.

outtextfile – output text file of trough times between syllables.

-**ssyldur** – minimum acceptable duration to define a syllable (default 0.0800).

-**ppktrof** – minimum height of peak above the surrounding troughs (default 0.0800).
The maximum possible height is 1.0.

-**P** – get peaks.

-**B** – get both troughs and peaks.

Understanding SPECFNU Mode 22: GET LOUDNESS TROUGHS

The output text file will usually need post-correction "by hand".

Musical Applications

End of SPECFNU Mode 22: GET LOUDNESS TROUGHS

SPECFNU MODE 23

SINE SPEECH – Convert formant frequencies to sine tones.

Usage

specfnu specfnu 23 *inanalfile outanalfile hffile sining*
 [-**again**] [-**bamp1**] [-**camp2**] [-**damp3**] [-**eamp4**] [-**nqdep1**] [-**oqdep2**] [-**pqdep3**] [-**qqdep4**] [-**s**] [-**f**] [-**r**] [-**S**]

Parameters

inanalfile – input analysis file.

outanalfile – output analysis file.

hffile – if set to zero, no pitch quantisation takes place.

Otherwise, this a textfile containing data about pitches to quantise to, as follows:

There must first be a marker indicating the grid type. These are:

#HS followed by listed MIDI pitches to quantise to.

#HF followed by listed MIDI pitches to quantise to in all octaves.

#THF followed by lines each with time plus a list of MIDI pitches.

First time must be zero & times must increase.

Each MIDI list must be the same length.

To change the number of pitches from line to line, duplicate values.

#SCALE followed by just TWO values:

(a) the number of equal pitch divisions in an octave.

(b) MIDI pitch of any pitch to tune the scales to.

#ELACS followed by just THREE values:

(a) size of "pseudo-octave" in (possibly fractional) semitones.

(b) the number of equal pitch divisions in the pseudo-octave.

(c) MIDI-pitch of any pitch to tune the scales to.

sining – degree of sinusoidisation. (Range: 0 to 1)

-again – output gain or attenuation. (Range: 0.01 to 10)

-bamp1 – relative gain or attenuation of formant 1. (Range: 0.01 to 10)

-camp2, -damp3, -eamp4 – similarly for the other formants.

-nqdep1 – if HF is applied (*hffile* non-zero), how strongly to force pitches' formant 1 to hf. (Range: 0-1)

-oqdep2, -pqdep3, -qqdep4 – similarly for the other formant bands.

-s – use a short window for extracting spectral envelope (default is large window).

-f – force lowest formant to use fundamental frequency as peak.

-r – replace unpitched (or extremely quiet) windows by silence.

-S – smoothing between pitches in formant traces.

Understanding SPECFNU Mode 23: SINE SPEECH

The final mode of SPECFNU replaces formant frequencies with sine tones. A single sine wave represents each formant. The original vocal sound is usually preserved in "shadow" form. Each of the main four formants has an optional level control (*amp1*, *amp2*, *amp3*, *amp4*).

Before converting, the pitch may optionally be quantised as in Modes 17 and 18, with the same type of datafile. If quantisation is employed, then each of the main four formants has a control for the degree of forcing the formant onto the pitch set (*dep1*, *dep2*, *dep3*, *dep4*).

Spectral window-size: if the "force fundamental" option (-f) is used, selecting a large window will find the fundamental, whereas a short window (option -s) may not.

Musical Applications

End of SPECFNU Mode 23: SINE SPEECH