
stft Documentation

Release 0.4.7

International AudioLabs Erlangen

March 31, 2016

1	Installation	3
2	Examples	5
2.1	Simple Example	5
2.2	Back and Forth Example	5
2.3	Passing multiple transfer functions	5
2.4	Saving Settings Example	6
3	Modules	7
3.1	stft package	7
3.2	stft types module	12
4	Indices and tables	15
	Python Module Index	17

This is a package for calculating short time fourier transforms with NumPy.

Contents:

Installation

This package can be installed using PIP:

```
pip install stft
```

Examples

2.1 Simple Example

Loading a file and calculating the spectrogram.

```
import stft
import scipy.io.wavfile as wav

fs, audio = wav.read('input.wav')
specgram = stft.spectrogram(audio)
```

See also:

module *stft.spectrogram()*

2.2 Back and Forth Example

Loading a file and calculating the spectrogram, its inverse and saving the result.

```
import stft
import scipy.io.wavfile as wav

fs, audio = wav.read('input.wav')
specgram = stft.spectrogram(audio)
output = stft.ispectrogram(specgram)
wav.write('output.wav', fs, output)
```

See also:

modules *stft.spectrogram()* *stft.ispectrogram()*

2.3 Passing multiple transfer functions

stft.spectrogram() and *stft.ispectrogram()* allow passing multiple transform functions as a list.

STFT will pick each transform for each frame it processes, the list of transforms will be extended indefinitely for as long as many frames need to be processed.

```
import stft
import scipy.io.wavfile as wav

fs, audio = wav.read('input.wav')
specgram = stft.spectrogram(audio, transform=[scipy.fftpack.fft, numpy.fft.fft])
output = stft.ispectrogram(specgram, transform=[scipy.fftpack.ifft, numpy.fft.ifft])
wav.write('output.wav', fs, output)
```

In this case, each frame will be processed using `scipy.fftpack.fft`, then `numpy.fft.fft`, then `scipy.fftpack.ifft` again etc.

2.4 Saving Settings Example

You do not need to pass the same settings to `stft.spectrogram()` and `stft.ispectrogram()` twice as the settings are saved in the array itself.

```
import stft
import scipy.io.wavfile as wav

fs, audio = wav.read('input.wav')
specgram = stft.spectrogram(audio, framelength=512, overlap=4)
output = stft.ispectrogram(specgram)
wav.write('output.wav', fs, output)
```

See also:

modules `stft.spectrogram()` `stft.ispectrogram()` `stft.types.SpectrogramArray`

3.1 stft package

3.1.1 Module contents

```
stft.spectrogram(data, framelength=1024, hopsize=None, overlap=None, centered=True, window=None, halved=True, transform=None, padding=0, save_settings=True)
```

Calculate the spectrogram of a signal

Parameters

- **data** (*array_like*) – The signal to be transformed. May be a 1D vector for single channel or a 2D matrix for multi channel data. In case of a mono signal, the data must be a 1D vector of length `samples`. In case of a multi channel signal, the data must be in the shape of `samples x channels`.
- **framelength** (*int*) – The signal frame length. Defaults to 1024.
- **hopsize** (*int*) – The signal frame hopsize. Defaults to `None`. Setting this value will override `overlap`.
- **overlap** (*int*) – The signal frame overlap coefficient. Value `x` means `1/x` overlap. Defaults to 2.
- **centered** (*boolean*) – Pad input signal so that the first and last window are centered around the beginning of the signal. Defaults to `true`.
- **window** (*callable, array_like*) – Window to be used for deringing. Can be `False` to disable windowing. Defaults to `scipy.signal.cosine`.
- **halved** (*boolean*) – Switch for turning on signal truncation. For real signals, the fourier transform of real signals returns a symmetrically mirrored spectrum. This additional data is not needed and can be removed. Defaults to `True`.
- **transform** (*callable*) – The transform to be used. Defaults to `scipy.fftpack.fft`.
- **padding** (*int*) – Zero-pad signal with `x` times the number of samples.
- **save_settings** (*boolean*) – Save settings used here in attribute `out.stft_settings` so that `ispectrogram()` can infer these settings without the developer having to pass them again.

Returns data – The spectrogram (or tensor of spectrograms) In case of a mono signal, the data is formatted as bins x frames. In case of a multi channel signal, the data is formatted as bins x frames x channels.

Return type array_like

Notes

The data will be padded to be a multiple of the desired FFT length.

See also:

`stft.stft.process()` The function used to transform the data

`stft.ispectrogram(data, framelength=None, hopsize=None, overlap=None, centered=None, window=None, halved=None, transform=None, padding=None, outlength=None)`

Calculate the inverse spectrogram of a signal

Parameters

- **data** (*array_like*) – The spectrogram to be inverted. May be a 2D matrix for single channel or a 3D tensor for multi channel data. In case of a mono signal, the data must be in the shape of bins x frames. In case of a multi channel signal, the data must be in the shape of bins x frames x channels.
- **framelength** (*int*) – The signal frame length. Defaults to infer from data.
- **hopsize** (*int*) – The signal frame hopsize. Defaults to infer from data. Setting this value will override overlap.
- **overlap** (*int*) – The signal frame overlap coefficient. Value x means 1/x overlap. Defaults to infer from data.
- **centered** (*boolean*) – Pad input signal so that the first and last window are centered around the beginning of the signal. Defaults to to infer from data.
- **window** (*callable, array_like*) – Window to be used for deringing. Can be False to disable windowing. Defaults to to infer from data.
- **halved** (*boolean*) – Switch to reconstruct the other half of the spectrum if the forward transform has been truncated. Defaults to to infer from data.
- **transform** (*callable*) – The transform to be used. Defaults to infer from data.
- **padding** (*int*) – Zero-pad signal with x times the number of samples. Defaults to infer from data.
- **outlength** (*int*) – Crop output signal to length. Useful when input length of spectrogram did not fit into framelength and input data had to be padded. Not setting this value will disable cropping, the output data may be longer than expected.

Returns data – The signal (or matrix of signals). In case of a mono output signal, the data is formatted as a 1D vector of length samples. In case of a multi channel output signal, the data is formatted as samples x channels.

Return type array_like

Notes

By default `spectrogram()` saves its transformation parameters in the output array. This data is used to infer the transform parameters here. Any aspect of the settings can be overridden by passing the according parameter to this function.

During transform the data will be padded to be a multiple of the desired FFT length. Hence, the result of the inverse transform might be longer than the input signal. However it is safe to remove the additional data, e.g. by using

```
output.resize(input.shape)
```

where `input` is the input of `stft.spectrogram()` and `output` is the output of `stft.ispectrogram()`

See also:

`stft.stft.iprocess()` The function used to transform the data

Module to transform signals

`stft.stft.cosine(M)`

Generate a halfcosine window of given length

Uses `scipy.signal.cosine` by default. However since this window function has only recently been merged into mainline SciPy, a fallback calculation is in place.

Parameters `M(int)` – Length of the window.

Returns `data` – The window function

Return type `array_like`

`stft.stft.iprocess(data, window, halved, transform, padding)`

Calculate the inverse short time fourier transform of a spectrum

Parameters

- **data** (`array_like`) – The spectrum to be calculated. Must be a 1D array.
- **window** (`array_like`) – Tapering window
- **halved** (`boolean`) – Switch for turning on signal truncation. For real output signals, the inverse fourier transform consumes a symmetrically mirrored spectrum. This additional data is not needed and can be removed. Setting this value to `True` will automatically create a mirrored spectrum.
- **transform** (`callable`) – The transform to be used.
- **padding** (`int`) – Signal before FFT transform was padded with x zeros.

Returns `data` – The signal

Return type `array_like`

`stft.stft.ispectrogram(data, framelength=None, hopsize=None, overlap=None, centered=None, window=None, halved=None, transform=None, padding=None, outlength=None)`

Calculate the inverse spectrogram of a signal

Parameters

- **data** (*array_like*) – The spectrogram to be inverted. May be a 2D matrix for single channel or a 3D tensor for multi channel data. In case of a mono signal, the data must be in the shape of `bins x frames`. In case of a multi channel signal, the data must be in the shape of `bins x frames x channels`.
- **framelength** (*int*) – The signal frame length. Defaults to infer from data.
- **hopsize** (*int*) – The signal frame hopsize. Defaults to infer from data. Setting this value will override `overlap`.
- **overlap** (*int*) – The signal frame overlap coefficient. Value `x` means `1/x` overlap. Defaults to infer from data.
- **centered** (*boolean*) – Pad input signal so that the first and last window are centered around the beginning of the signal. Defaults to to infer from data.
- **window** (*callable, array_like*) – Window to be used for deringing. Can be `False` to disable windowing. Defaults to to infer from data.
- **halved** (*boolean*) – Switch to reconstruct the other halve of the spectrum if the forward transform has been truncated. Defaults to to infer from data.
- **transform** (*callable*) – The transform to be used. Defaults to infer from data.
- **padding** (*int*) – Zero-pad signal with `x` times the number of samples. Defaults to infer from data.
- **outlength** (*int*) – Crop output signal to length. Useful when input length of spectrogram did not fit into `framelength` and input data had to be padded. Not setting this value will disable cropping, the output data may be longer than expected.

Returns data – The signal (or matrix of signals). In case of a mono output signal, the data is formatted as a 1D vector of length `samples`. In case of a multi channel output signal, the data is formatted as `samples x channels`.

Return type `array_like`

Notes

By default `spectrogram()` saves its transformation parameters in the output array. This data is used to infer the transform parameters here. Any aspect of the settings can be overridden by passing the according parameter to this function.

During transform the data will be padded to be a multiple of the desired FFT length. Hence, the result of the inverse transform might be longer than the input signal. However it is safe to remove the additional data, e.g. by using

```
output.resize(input.shape)
```

where `input` is the input of `stft.spectrogram()` and `output` is the output of `stft.ispectrogram()`

See also:

`stft.stft.iprocess()` The function used to transform the data

`stft.stft.process(data, window, halved, transform, padding)`

Calculate a windowed transform of a signal

Parameters

- **data** (*array_like*) – The signal to be calculated. Must be a 1D array.
- **window** (*array_like*) – Tapering window
- **halved** (*boolean*) – Switch for turning on signal truncation. For real signals, the fourier transform of real signals returns a symmetrically mirrored spectrum. This additional data is not needed and can be removed.
- **transform** (*callable*) – The transform to be used.
- **padding** (*int*) – Zero-pad signal with x times the number of samples.

Returns **data** – The spectrum

Return type *array_like*

`stft.stft.spectrogram(data, framelength=1024, hopsize=None, overlap=None, centered=True, window=None, halved=True, transform=None, padding=0, save_settings=True)`
Calculate the spectrogram of a signal

Parameters

- **data** (*array_like*) – The signal to be transformed. May be a 1D vector for single channel or a 2D matrix for multi channel data. In case of a mono signal, the data is must be a 1D vector of length `samples`. In case of a multi channel signal, the data must be in the shape of `samples x channels`.
- **framelength** (*int*) – The signal frame length. Defaults to 1024.
- **hopsize** (*int*) – The signal frame hopsize. Defaults to `None`. Setting this value will override `overlap`.
- **overlap** (*int*) – The signal frame overlap coefficient. Value x means 1/x overlap. Defaults to 2.
- **centered** (*boolean*) – Pad input signal so that the first and last window are centered around the beginning of the signal. Defaults to `true`.
- **window** (*callable, array_like*) – Window to be used for deringing. Can be `False` to disable windowing. Defaults to `scipy.signal.cosine`.
- **halved** (*boolean*) – Switch for turning on signal truncation. For real signals, the fourier transform of real signals returns a symmetrically mirrored spectrum. This additional data is not needed and can be removed. Defaults to `True`.
- **transform** (*callable*) – The transform to be used. Defaults to `scipy.fftpack.fft`.
- **padding** (*int*) – Zero-pad signal with x times the number of samples.
- **save_settings** (*boolean*) – Save settings used here in attribute `out.stft_settings` so that `ispectrogram()` can infer these settings without the developer having to pass them again.

Returns **data** – The spectrogram (or tensor of spectrograms) In case of a mono signal, the data is formatted as `bins x frames`. In case of a multi channel signal, the data is formatted as `bins x frames x channels`.

Return type *array_like*

Notes

The data will be padded to be a multiple of the desired FFT length.

See also:

`stft.stft.process()` The function used to transform the data

3.2 stft types module

3.2.1 Module contents

class `stft.types.SpectrogramArray`

Bases: `numpy.ndarray`

NumpyArray with additional `stft_settings` attribute for saving stft-specific settings.

Attributes

<code>T</code>	Same as <code>self.transpose()</code> , except that <code>self</code> is returned if <code>self.ndim < 2</code> .
<code>base</code>	Base object if memory is from some other object.
<code>ctypes</code>	An object to simplify the interaction of the array with the <code>ctypes</code> module.
<code>data</code>	Python buffer object pointing to the start of the array's data.
<code>dtype</code>	Data-type of the array's elements.
<code>flags</code>	Information about the memory layout of the array.
<code>flat</code>	A 1-D iterator over the array.
<code>imag</code>	The imaginary part of the array.
<code>itemsize</code>	Length of one array element in bytes.
<code>nbytes</code>	Total bytes consumed by the elements of the array.
<code>ndim</code>	Number of array dimensions.
<code>real</code>	The real part of the array.
<code>shape</code>	Tuple of array dimensions.
<code>size</code>	Number of elements in the array.
<code>strides</code>	Tuple of bytes to step in each dimension when traversing an array.

Methods

<code>all([axis, out])</code>	Returns True if all elements evaluate to True.
<code>any([axis, out])</code>	Returns True if any of the elements of <i>a</i> evaluate to True.
<code>argmax([axis, out])</code>	Return indices of the maximum values along the given axis.
<code>argmin([axis, out])</code>	Return indices of the minimum values along the given axis of <i>a</i> .
<code>argpartition(kth[, axis, kind, order])</code>	Returns the indices that would partition this array.
<code>argsort([axis, kind, order])</code>	Returns the indices that would sort this array.
<code>astype(dtype[, order, casting, subok, copy])</code>	Copy of the array, cast to a specified type.
<code>byteswap(inplace)</code>	Swap the bytes of the array elements
<code>choose(choices[, out, mode])</code>	Use an index array to construct a new array from a set of choices.
<code>clip(a_min, a_max[, out])</code>	Return an array whose values are limited to <code>[a_min, a_max]</code> .
<code>compress(condition[, axis, out])</code>	Return selected slices of this array along given axis.
<code>conj()</code>	Complex-conjugate all elements.
<code>conjugate()</code>	Return the complex conjugate, element-wise.
<code>copy([order])</code>	Return a copy of the array.
<code>cumprod([axis, dtype, out])</code>	Return the cumulative product of the elements along the given axis.

Table 3.2 – continued from previous page

<code>cumsum([axis, dtype, out])</code>	Return the cumulative sum of the elements along the given axis.
<code>diagonal([offset, axis1, axis2])</code>	Return specified diagonals.
<code>dot(b[, out])</code>	Dot product of two arrays.
<code>dump(file)</code>	Dump a pickle of the array to the specified file.
<code>dumps()</code>	Returns the pickle of the array as a string.
<code>fill(value)</code>	Fill the array with a scalar value.
<code>flatten([order])</code>	Return a copy of the array collapsed into one dimension.
<code>getfield(dtype[, offset])</code>	Returns a field of the given array as a certain type.
<code>item(*args)</code>	Copy an element of an array to a standard Python scalar and return it.
<code>itemset(*args)</code>	Insert scalar into an array (scalar is cast to array's dtype, if possible)
<code>max([axis, out])</code>	Return the maximum along a given axis.
<code>mean([axis, dtype, out])</code>	Returns the average of the array elements along given axis.
<code>min([axis, out])</code>	Return the minimum along a given axis.
<code>newbyteorder([new_order])</code>	Return the array with the same data viewed with a different byte order.
<code>nonzero()</code>	Return the indices of the elements that are non-zero.
<code>partition(kth[, axis, kind, order])</code>	Rearranges the elements in the array in such a way that value of the element in kth position is in the middle of the sorted order.
<code>prod([axis, dtype, out])</code>	Return the product of the array elements over the given axis
<code>ptp([axis, out])</code>	Peak to peak (maximum - minimum) value along a given axis.
<code>put(indices, values[, mode])</code>	Set <code>a.flat[n] = values[n]</code> for all <code>n</code> in indices.
<code>ravel([order])</code>	Return a flattened array.
<code>repeat(repeats[, axis])</code>	Repeat elements of an array.
<code>reshape(shape[, order])</code>	Returns an array containing the same data with a new shape.
<code>resize(new_shape[, refcheck])</code>	Change shape and size of array in-place.
<code>round([decimals, out])</code>	Return <code>a</code> with each element rounded to the given number of decimals.
<code>searchsorted(v[, side, sorter])</code>	Find indices where elements of <code>v</code> should be inserted in <code>a</code> to maintain order.
<code>setfield(val, dtype[, offset])</code>	Put a value into a specified place in a field defined by a data-type.
<code>setflags([write, align, uic])</code>	Set array flags WRITEABLE, ALIGNED, and UPDATEIFCOPY, respectively.
<code>sort([axis, kind, order])</code>	Sort an array, in-place.
<code>squeeze([axis])</code>	Remove single-dimensional entries from the shape of <code>a</code> .
<code>std([axis, dtype, out, ddof])</code>	Returns the standard deviation of the array elements along given axis.
<code>sum([axis, dtype, out])</code>	Return the sum of the array elements over the given axis.
<code>swapaxes(axis1, axis2)</code>	Return a view of the array with <code>axis1</code> and <code>axis2</code> interchanged.
<code>take(indices[, axis, out, mode])</code>	Return an array formed from the elements of <code>a</code> at the given indices.
<code>tofile(fid[, sep, format])</code>	Write array to a file as text or binary (default).
<code>tolist()</code>	Return the array as a (possibly nested) list.
<code>tostring([order])</code>	Construct a Python string containing the raw data bytes in the array.
<code>trace([offset, axis1, axis2, dtype, out])</code>	Return the sum along diagonals of the array.
<code>transpose(*axes)</code>	Returns a view of the array with axes transposed.
<code>var([axis, dtype, out, ddof])</code>	Returns the variance of the array elements, along given axis.
<code>view([dtype, type])</code>	New view of array with the same data.

Indices and tables

- `genindex`
- `modindex`
- `search`

S

`stft`, [7](#)
`stft.stft`, [9](#)
`stft.types`, [12](#)

C

`cosine()` (in module `stft.stft`), 9

I

`iprocess()` (in module `stft.stft`), 9

`ispectrogram()` (in module `stft`), 8

`ispectrogram()` (in module `stft.stft`), 9

P

`process()` (in module `stft.stft`), 10

S

`spectrogram()` (in module `stft`), 7

`spectrogram()` (in module `stft.stft`), 11

`SpectrogramArray` (class in `stft.types`), 12

`stft` (module), 7

`stft.stft` (module), 9

`stft.types` (module), 12