
asteroid Documentation

Release 0.4.0rc1

Manuel Pariente et al.

Feb 18, 2021

Start here

1	What is Asteroid?	3
2	Installation	5
3	What is a recipe?	7
4	Datasets and tasks	11
5	Training and Evaluation	17
6	Pretrained models	19
7	Command-line interface	21
8	FAQ	23
9	PyTorch Datasets	27
10	Filterbank API	29
11	DNN building blocks	43
12	Models	45
13	Losses & Metrics	47
14	Lightning Wrapper	49
15	Optimizers & Schedulers	51
16	DSP Modules	53
17	Utils	55
18	Asteroid High-Level Contribution Guide	57
19	How to contribute	63
20	Indices and tables	65

Python Module Index	67
Index	69



Asteroid is a Pytorch-based audio source separation toolkit that enables fast experimentation on common datasets. It comes with a source code that supports a large range of datasets and architectures, and a set of recipes to reproduce some important papers.

CHAPTER 1

What is Asteroid?

Asteroid is a PyTorch-based audio source separation toolkit.

The main goals of Asteroid are:

- Gather a wider **community** around audio source separation by lowering the barriers to entry.
- **Promote reproducibility** by replicating important research papers.
- Automatize most engineering and **make way for research**.
- Simplify **model sharing** to reduce compute costs and carbon footprint.

So, how do we do that? We aim to provide

- PyTorch `Dataset` for **common datasets**.
- Ready-to-use state-of-the art source separation architectures in **native PyTorch**.
- **Configurable recipes** from data preparation to evaluation.
- **Pretrained models** for a wide variety of tasks and architectures.

1.1 Who is it for?

Asteroid has several target usage:

- Use asteroid in your own code, as a package.
- Use available recipes to build your own separation model.
- Use pretrained models to process your files.
- Hit the ground running with your research ideas!

1.2 Want to know more?

- Visit our webpage
- Read our paper
- Watch the presentation video
- **‘Check how we won the PyTorch Hackathon 2020 !<<https://devpost.com/software/asteroid-the-pytorch-based-source-separation-toolkit>>‘__**

CHAPTER 2

Installation

By following the instructions below, first install PyTorch and then Asteroid (using either pip/dev install). We recommend the development installation for users likely to modify the source code.

2.1 CUDA and PyTorch

Asteroid is based on PyTorch. To run Asteroid on GPU, you will need a CUDA-enabled PyTorch installation. Visit this site for the instructions: <https://pytorch.org/get-started/locally/>.

2.2 Pip

Asteroid is regularly updated on PyPI, install the latest stable version with:

```
pip install asteroid
```

2.3 Development installation

For development installation, you can fork/clone the GitHub repo and locally install it with pip:

```
git clone https://github.com/asteroid-team/asteroid
cd asteroid
pip install -e .
```

This is an editable install (`-e` flag), it means that source code changes (or branch switching) are automatically taken into account when importing asteroid.

You can also use `conda env create -f environment.yml` to create a Conda env directly.

CHAPTER 3

What is a recipe?

A recipe is a set of scripts that use Asteroid to build a source separation system. Each directory corresponds to a dataset and each subdirectory corresponds to a system build on this dataset. You can start by reading [this recipe](#) to get familiar with them.

3.1 How is it organized?

Most recipes are organized as follows. When you clone the repo, `data`, `exp` and `logs` won't be there yet, it's normal.

```
├── data/
├── exp/
├── logs/
├── local/
│   ├── convert_sphere2wav.sh
│   ├── prepare_data.sh
│   ├── conf.yml
│   └── preprocess_wham.py
├── utils/
│   ├── parse_options.sh
│   └── prepare_python_env.sh
├── run.sh
├── train.py
├── model.py
└── eval.py
```

A small graph might help.



```
readmes/../../docs/source/_static/images/code_example_cropped.png
```

3.2 How does it work?

Let's try to summarize how recipes work :

- There is a master file, `run.sh`, from which all the steps are ran (install dependencies, download data, create dataset, train a model evaluate it and so on..). This recipe style is borrowed from [Kaldi](#) and [ESPnet](#).
 - You usually have to change some variables in the top of the file (comments are around it to help you) like data directory, python path etc..
 - This script is controlled by several arguments. Among them, `stage` controls from where do you start the script. You already generated the data? No need to do it again, set `stage=3`!
 - All steps until training are dataset-specific and the corresponding scripts are stored in `./local`
- The training and evaluation scripts are then called from `run.sh`
 - There is a script, `model.py`, where the model should be defined along with the `System` subclass used for training (if needed).
 - We wrap the model definition in one function (`make_model_and_optimizer`). The function receives a dictionary which is also saved in the experiment folder. This make checkpoint restoring easy without any additional constraints.
 - We also write a function to load the best model (`load_best_model`) after training. This is useful to load the model several time (evaluation, separation of new examples...).
- The arguments flow through bash/python/yaml in a specific way, which was designed by us and suits our use-cases until now:
 - The very first step is the `local/conf.yml` file where is a hierarchical configuration file,
 - **On the python side** : This file is parsed as a dictionary of dictionary in `training.py`. From this dict, we create an argument parser which can accept all the second-level keys from the dictionary (so second-level keys should be unique) as arguments and has the default values from the `conf.yml` file.
 - **On the bash side**: we also parse arguments from the command line (using `utils/parse_options.sh`). The arguments above the line `. utils/parse_options.sh` can be parsed, the rest are fixed. Most arguments will be passed to the training script. Others control the data preparation, GPU usage etc...
 - In light of all this the config file should have sensible default values that shouldn't be modified by hand much. The quickly configurable part of the recipe are added to `run.sh` (you want to experiment with the batch size, add an argument in and pass it to python. If you want it fixed, no need to put it in bash, the `conf.yml` file keeps it for you.) This makes it possible to directly identify the important parts of the experiment, without reading lots of lines of argparser or bash arguments.
- Some more notes :
 - After the first execution, you can go and change `stage` in `run.sh` to avoid redoing all the steps every-time.
 - To use GPUs for training, run `run.sh --id 0,1` where 0 and 1 are the GPUs you want to use, training should automatically take advantage of both GPUs.
 - By default, a random id is generated for each run, you can also add a `tag` to name the experiments how you want. For example `run.sh --tag with_cool_loss` will save all results to `exp/train_{arch_name}_with_cool_loss`. You'll also find the corresponding log file in `logs/train_{arch_name}_with_cool_loss.log`.
 - Model loading methods suppose that the model architecture is the same as when training was performed. Be careful when you change it.

Again, you have a doubt, a question, a suggestion or a request, [open an issue](#) or [join the slack](#), we'll be happy to help you.

Datasets and tasks

The following is a list of supported datasets, sorted by task. If you're more interested in the corresponding PyTorch Dataset, see [this page](#)

4.1 Speech separation

4.1.1 wsj0-2mix dataset

wsj0-2mix is a single channel speech separation dataset base on WSJ0. Three speaker extension (wsj0-3mix) is also considered here.

Reference

```
@article{Hershey_2016,
  title={Deep clustering: Discriminative embeddings for segmentation and separation},
  ISBN={9781479999880},
  url={http://dx.doi.org/10.1109/ICASSP.2016.7471631},
  DOI={10.1109/icassp.2016.7471631},
  journal={2016 IEEE International Conference on Acoustics, Speech and Signal_
↪Processing (ICASSP)},
  publisher={IEEE},
  author={Hershey, John R. and Chen, Zhuo and Le Roux, Jonathan and Watanabe, Shinji}
↪,
  year={2016},
}
```

4.1.2 WHAM dataset

WHAM! is a noisy single-channel speech separation dataset based on WSJ0. It is a noisy extension of [wsj0-2mix](#).

More info [here](#).

References

```
@inproceedings{WHAMWichern2019,  
  author={Gordon Wicher and Joe Antognini and Michael Flynn and Licheng Richard Zhu,  
↪ and Emmett McQuinn and Dwight Crow and Ethan Manilow and Jonathan Le Roux},  
  title={{WHAM!}: extending speech separation to noisy environments},  
  year=2019,  
  booktitle={Proc. Interspeech},  
  pages={1368--1372},  
  doi={10.21437/Interspeech.2019-2821},  
  url={http://dx.doi.org/10.21437/Interspeech.2019-2821}  
}
```

4.1.3 WHAMR dataset

WHAMR! is a noisy and reverberant single-channel speech separation dataset based on WSJ0. It is a reverberant extension of [WHAM!](#).

Note that WHAMR! can synthesize binaural recordings, but we only consider the single channel for now.

More info [here](#). **References**

```
@misc{maciejewski2019whamr,  
  title={WHAMR!: Noisy and Reverberant Single-Channel Speech Separation},  
  author={Matthew Maciejewski and Gordon Wicher and Emmett McQuinn and Jonathan Le,  
↪ Roux},  
  year={2019},  
  eprint={1910.10279},  
  archivePrefix={arXiv},  
  primaryClass={cs.SD}  
}
```

4.1.4 LibriMix dataset

The LibriMix dataset is an open source dataset derived from LibriSpeech dataset. It's meant as an alternative and complement to [WHAM](#).

More info [here](#).

References

```
@misc{cosentino2020librimix,  
  title={LibriMix: An Open-Source Dataset for Generalizable Speech Separation},  
  author={Joris Cosentino and Manuel Pariente and Samuele Cornell and Antoine,  
↪ Deleforge and Emmanuel Vincent},  
  year={2020},  
  eprint={2005.11262},  
  archivePrefix={arXiv},  
  primaryClass={eess.AS}  
}
```

4.1.5 Kinect-WSJ dataset

Kinect-WSJ is a reverberated, noisy version of the WSJ0-2MIX dataset. Microphones are placed on a linear array with spacing between the devices resembling that of Microsoft Kinect™, the device used to record the CHiME-5 dataset.

This was done so that we could use the real ambient noise captured as part of CHiME-5 dataset. The room impulse responses (RIR) were simulated for a sampling rate of 16,000 Hz.

Requirements

- `wsj_path` : Path to precomputed wsj-2mix dataset. Should contain the folder 2speakers/wav16k/. If you don't have wsj_mix dataset, please create it using the scripts in `egs/wsj0_mix`
- `chime_path` : Path to chime-5 dataset. Should contain the folders train, dev and eval
- `dihard_path` : Path to dihard labels. Should contain `*.lab` files for the train and dev set

References [Original repo](#)

```
@inproceedings{sivasankaran2020,
  booktitle = {2020 28th {{European Signal Processing Conference}} ({{EUSIPCO}})},
  title={Analyzing the impact of speaker localization errors on speech separation for_
↪automatic speech recognition},
  author={Sunit Sivasankaran and Emmanuel Vincent and Dominique Fohr},
  year={2021},
  month = Jan,
}
```

4.1.6 SMS_WSJ dataset

SMS_WSJ (stands for Spatialized Multi-Speaker Wall Street Journal) is a multichannel source separation dataset, based on WSJ0 and WSJ1.

All the information regarding the dataset can be found in [this repo](#).

References If you use this dataset, please cite the corresponding paper as follows :

```
@Article{SmsWsj19,
  author   = {Drude, Lukas and Heitkaemper, Jens and Boeddeker, Christoph and Haeb-
↪Umbach, Reinhold},
  title    = {{SMS-WSJ}: Database, performance measures, and baseline recipe for_
↪multi-channel source separation and recognition},
  journal  = {arXiv preprint arXiv:1910.13934},
  year     = {2019},
}
```

4.2 Speech enhancement

4.2.1 DNS Challenge's dataset

The Deep Noise Suppression (DNS) Challenge is a single-channel speech enhancement challenge organized by Microsoft, with a focus on real-time applications. More info can be found on the [official page](#).

References The challenge paper, [here](#).

```
@misc{DNSChallenge2020,
  title={The INTERSPEECH 2020 Deep Noise Suppression Challenge: Datasets, Subjective_
↪Speech Quality and Testing Framework},
  author={Chandan K. A. Reddy and Ebrahim Beyrami and Harishchandra Dubey and Vishak_
↪Gopal and Roger Cheng and Ross Cutler and Sergiy Matuskevych and Robert Aichner and_
↪Ashkan Aazami and Sebastian Braun and Puneet Rana and Sriram Srinivasan and_
↪Johannes Gehrke}, year={2020},
```

(continues on next page)

(continued from previous page)

```
eprint={2001.08662},  
}
```

The baseline paper, [here](#).

```
@misc{xia2020weighted,  
title={Weighted Speech Distortion Losses for Neural-network-based Real-time Speech_  
↪Enhancement},  
author={Yangyang Xia and Sebastian Braun and Chandan K. A. Reddy and Harishchandra_  
↪Dubey and Ross Cutler and Ivan Tashev},  
year={2020},  
eprint={2001.10601},  
}
```

4.3 Music source separation

4.3.1 MUSDB18 Dataset

The musdb18 is a dataset of 150 full lengths music tracks (~10h duration) of different genres along with their isolated drums, bass, vocals and others stems.

More info [here](#).

4.3.2 DAMP-VSEP dataset

All the information regarding the dataset can be found in [zenodo](#).

References If you use this dataset, please cite as follows :

```
@dataset{smule_inc_2019_3553059,  
  author      = {Smule, Inc},  
  title       = {{DAMP-VSEP: Smule Digital Archive of Mobile  
                  Performances - Vocal Separation}},  
  month       = oct,  
  year        = 2019,  
  publisher   = {Zenodo},  
  version     = {1.0.1},  
  doi         = {10.5281/zenodo.3553059},  
  url         = {https://doi.org/10.5281/zenodo.3553059}  
}
```

4.4 Environmental sound separation

4.4.1 FUSS dataset

The Free Universal Sound Separation (FUSS) dataset comprises audio mixtures of arbitrary sounds with source references for use in experiments on arbitrary sound separation.

All the information related to this dataset can be found in [this repo](#).

References If you use this dataset, please cite the corresponding paper as follows:

```
@Article{Wisdom2020,
  author    = {Scott Wisdom and Hakan Erdogan and Daniel P. W. Ellis and Romain_
↪Serizel and Nicolas Turpault and Eduardo Fonseca and Justin Salamon and Prem_
↪Seetharaman and John R. Hershey},
  title     = {What's All the FUSS About Free Universal Sound Separation Data?},
  journal   = {in preparation},
  year      = {2020},
}
```

4.5 Audio-visual source separation

4.5.1 AVSpeech dataset

AVSpeech is an audio-visual speech separation dataset which was introduced by Google in this article [Looking to Listen at the Cocktail Party: A Speaker-Independent Audio-Visual Model for Speech Separation](#).

More info [here](#).

References

```
@article{Ephrat_2018,
  title={Looking to listen at the cocktail party},
  volume={37},
  url={http://dx.doi.org/10.1145/3197517.3201357},
  DOI={10.1145/3197517.3201357},
  journal={ACM Transactions on Graphics},
  publisher={Association for Computing Machinery (ACM)},
  author={Ephrat, Ariel and Mosseri, Inbar and Lang, Oran and Dekel, Tali and Wilson,
↪Kevin and Hassidim, Avinatan and Freeman, William T. and Rubinstein, Michael},
  year={2018},
  pages={1-11}
}
```

4.6 Speaker extraction

Training and Evaluation

Training and evaluation are the two essential parts of the recipes. For training, we offer a thin wrapper around `PyTorchLightning` that seamlessly enables distributed training, experiment logging and more, without sacrificing flexibility. For evaluation we released `pb_bss_eval` on PyPI, which is the evaluation part of `pb_bss`. All the credit goes to the original authors from the Paderborn University.

5.1 Training with PyTorchLightning

First, have a look [here](#) for an overview of PyTorchLightning. As you saw, the `LightningModule` is a central class of PyTorchLightning where a large part of the research-related logic lives. Instead of subclassing it everytime, we use `System`, a thin wrapper that separately gathers the essential parts of every deep learning project:

1. A model
2. Optimizer
3. Loss function
4. Train/val data

```
class System(pl.LightningModule):
    def __init__(self, model, optimizer, loss_func, train_loader, val_loader):
        ...

    def common_step(self, batch):
        """ common_step is the method that'll be called at both train/val time. """
        inputs, targets = batch
        est_targets = self(inputs)
        loss = self.loss_func(est_targets, targets)
        return loss
```

Only overwriting `common_step` will often be enough to obtain a desired behavior, while avoiding boilerplate code. Then, we can use the native PyTorchLightning `Trainer` to train the models.

5.2 Evaluation

Asteroid's function `compute_metrics` that calls `pb_bss_eval` is used to compute the following common source separation metrics:

- SDR / SIR / SAR
- STOI
- PESQ
- SI-SDR

Pretrained models

Asteroid provides pretrained models through the [Asteroid community](#) in Zenodo. Have a look at the Zenodo page to choose which model you want to use.

Enjoy having pretrained models? **Please share your models** if you train some, we made it simple with the `asteroid-upload` CLI, check the next sections.

6.1 Using them

Loading a pretrained model is super simple!

```
from asteroid.models import ConvTasNet
model = ConvTasNet.from_pretrained('mpariente/ConvTasNet_WHAM!_sepclean')
```

Use the [search page](#) if you want to narrow your search.

You can also load it with Hub

```
from torch import hub
model = hub.load('mpariente/asteroid', 'conv_tasnet', 'mpariente/ConvTasNet_WHAM!_
↪sepclean')
```

6.2 Model caching

When using a `from_pretrained` method, the model is downloaded and cached. The cache directory is either the value in the `$ASTEROID_CACHE` environment variable, or `~/ .cache/torch/asteroid`.

6.3 Share your models

At the end of each sharing-enabled recipe, all the necessary infos are gathered into a file, the only thing that's left to do is to run

```
asteroid-upload exp/your_exp_dir/publish_dir --uploader "Name Here"
```

Ok, not really. First you need to register to [Zenodo](#) (Sign in with GitHub: ok), [create a token](#) and use it with the `--token` option of the CLI, or by setting the `ACCESS_TOKEN` environment variable. If you plan to upload more models (and you should :innocent:), you can fill in your infos in `uploader_info.yml` at the root, like this.

```
uploader: Manuel Pariente
affiliation: INRIA
git_username: mpariente
token: TOKEN_HERE
```

6.4 Note about licenses

All Asteroid's pretrained models are shared under the [Attribution-ShareAlike 3.0 \(CC BY-SA 3.0\)](#) license. This means that models are released under the same license as the original training data. **If any non-commercial data is used during training (wsj0, WHAM's noises etc.), the models are non-commercial use only.** This is indicated in the bottom of the corresponding Zenodo page (ex: [here](#)).

Command-line interface

7.1 Inference

7.1.1 asteroid-infer

Example

```
asteroid-infer "mpariente/ConvTasNet_WHAM!_sepclean" --files myaudio.wav --resample --  
→ola-window 8000 --ola-hop 4000
```

Reference

Command 'asteroid-infer -help' failed: [Errno 2] No such file or directory: 'asteroid-infer': 'asteroid-infer'

7.2 Publishing models

7.2.1 asteroid-upload

Reference

Command 'asteroid-upload -help' failed: [Errno 2] No such file or directory: 'asteroid-upload': 'asteroid-upload'

7.2.2 asteroid-register-sr


Reference

Command 'asteroid-register-sr -help' failed: [Errno 2] No such file or directory: 'asteroid-register-sr': 'asteroid-register-sr'

8.1 My results are worse than the ones reported in the README, why?

There are few possibilities here:

1. Your data is wrong. We had this examples with wsj0-mix, WHAM etc.. where wv2 was used instead of wv1 to generate the data. This was fixed in [#166](#). Chances are there is a pretrained model available for the given dataset, run the evaluation with it. If your results are different, it's a data problem. Refs: [#164](#), [#165](#) and [#188](#).
2. You stopped training too early. We've seen this happen, specially with DPRNN. Be sure that your training/validation losses are completely flat at the end of training.



3. If it's not both, there is a real bug and we're happy you caught it! Please, open an issue with your torch/pytorch_lightning/asteroid versions to let us know.

8.2 How long does it take to train a model?

Need a log here.

8.3 Can I use the pretrained models for commercial purposes?

Not always. See the note on pretrained models Licenses *Note about licenses*

8.4 Separated audio is really bad, what is happening?

There are several possible cause to this, a common one is clipping.

1. When training with scale invariant losses (e.g. SI-SNR) the audio output can be unbounded. However, waveform values should be normalized to $[-1, 1]$ range before saving, otherwise they will be clipped. See [Clipping on Wikipedia](#) and [issue #250](#)

2. As all supervised learning approaches, source separation can suffer from generalization error when evaluated on unseen data. If your model works well on data similar to your training data but doesn't work on real data, that's probably why. More about this [on Wikipedia](#).

CHAPTER 9

PyTorch Datasets

This page lists the supported datasets and their corresponding PyTorch's `Dataset` class. If you're interested in the datasets more than in the code, see [this page](#).

9.1 LibriMix

9.2 Wsj0mix

9.3 WHAM!

9.4 WHAMR!

9.5 SMS-WSJ

9.6 KinectWSJMix

9.7 DNSDataset

9.8 MUSDB18

9.9 DAMP-VSEP

9.10 FUSS

9.11 AVSpeech

10.1 Filterbank, Encoder and Decoder

```
class asteroid_filterbanks.Filterbank (n_filters, kernel_size, stride=None, sample_rate=8000.0)
```

Bases: sphinx.ext.autodoc.importer._MockObject

Base Filterbank class. Each subclass has to implement a `filters` method.

Parameters

- **n_filters** (*int*) – Number of filters.
- **kernel_size** (*int*) – Length of the filters.
- **stride** (*int*, *optional*) – Stride of the conv or transposed conv. (Hop size). If None (default), set to `kernel_size // 2`.
- **sample_rate** (*float*) – Sample rate of the expected audio. Defaults to 8000.

Variables **n_feats_out** (*int*) – Number of output filters.

filters ()

Abstract method for filters.

pre_analysis (*wav*: <sphinx.ext.autodoc.importer._MockObject object at 0x7ff682d0fd90>)

Apply transform before encoder convolution.

post_analysis (*spec*: <sphinx.ext.autodoc.importer._MockObject object at 0x7ff682d0fe50>)

Apply transform to encoder convolution.

pre_synthesis (*spec*: <sphinx.ext.autodoc.importer._MockObject object at 0x7ff690802bd0>)

Apply transform before decoder transposed convolution.

post_synthesis (*wav*: <sphinx.ext.autodoc.importer._MockObject object at 0x7ff690802d50>)

Apply transform after decoder transposed convolution.

get_config()

Returns dictionary of arguments to re-instantiate the class. Needs to be subclassed if the filterbanks takes additional arguments than `n_filters` `kernel_size` `stride` and `sample_rate`.

class `asteroid_filterbanks.Encoder` (*filterbank*, *is_pinv=False*, *as_conv1d=True*, *padding=0*)

Bases: `asteroid_filterbanks.enc_dec._EncDec`

Encoder class.

Add encoding methods to Filterbank classes. Not intended to be subclassed.

Parameters

- **filterbank** (*Filterbank*) – The filterbank to use as an encoder.
- **is_pinv** (*bool*) – Whether to be the pseudo inverse of filterbank.
- **as_conv1d** (*bool*) – Whether to behave like `nn.Conv1d`. If True (default), forwarding input with shape *(batch, 1, time)* will output a tensor of shape *(batch, freq, conv_time)*. If False, will output a tensor of shape *(batch, 1, freq, conv_time)*.
- **padding** (*int*) – Zero-padding added to both sides of the input.

classmethod `pinv_of` (*filterbank*, ***kwargs*)

Returns an *Encoder*, pseudo inverse of a *Filterbank* or *Decoder*.

forward (*waveform*)

Convolve input waveform with the filters from a filterbank.

Parameters *waveform* (`torch.Tensor`) – any tensor with samples along the last dimension. The waveform representation with and batch/channel etc.. dimension.

Returns `torch.Tensor` – The corresponding TF domain signal.

Shapes

```
>>> (time, ) -> (freq, conv_time)
>>> (batch, time) -> (batch, freq, conv_time) # Avoid
>>> if as_conv1d:
>>>     (batch, 1, time) -> (batch, freq, conv_time)
>>>     (batch, chan, time) -> (batch, chan, freq, conv_time)
>>> else:
>>>     (batch, chan, time) -> (batch, chan, freq, conv_time)
>>> (batch, any, dim, time) -> (batch, any, dim, freq, conv_time)
```

class `asteroid_filterbanks.Decoder` (*filterbank*, *is_pinv=False*, *padding=0*, *output_padding=0*)

Bases: `asteroid_filterbanks.enc_dec._EncDec`

Decoder class.

Add decoding methods to Filterbank classes. Not intended to be subclassed.

Parameters

- **filterbank** (*Filterbank*) – The filterbank to use as a decoder.
- **is_pinv** (*bool*) – Whether to be the pseudo inverse of filterbank.
- **padding** (*int*) – Zero-padding added to both sides of the input.
- **output_padding** (*int*) – Additional size added to one side of the output shape.

Note: `padding` and `output_padding` arguments are directly passed to `F.conv_transpose1d`.

classmethod `pinv_of` (*filterbank*)

Returns an Decoder, pseudo inverse of a filterbank or Encoder.

forward (*spec*, *length*: *Optional[int] = None*) → <sphinx.ext.autodoc.importer._MockObject object at 0x7ff690802810>

Applies transposed convolution to a TF representation.

This is equivalent to overlap-add.

Parameters

- **spec** (*torch.Tensor*) – 3D or 4D Tensor. The TF representation. (Output of *Encoder.forward()*).
- **length** – desired output length.

Returns *torch.Tensor* – The corresponding time domain signal.

class `asteroid_filterbanks.make_enc_dec`

Creates congruent encoder and decoder from the same filterbank family.

Parameters

- **fb_name** (*str*, *className*) – Filterbank family from which to make encoder and decoder. To choose among ['free', 'analytic_free', 'param_sinc', 'stft']. Can also be a class defined in a submodule in this subpackage (e.g. *FreeFB*).
- **n_filters** (*int*) – Number of filters.
- **kernel_size** (*int*) – Length of the filters.
- **stride** (*int*, *optional*) – Stride of the convolution. If *None* (default), set to `kernel_size // 2`.
- **sample_rate** (*float*) – Sample rate of the expected audio. Defaults to 8000.0.
- **who_is_pinv** (*str*, *optional*) – If *None*, no pseudo-inverse filters will be used. If string (among ['encoder', 'decoder']), decides which of *Encoder* or *Decoder* will be the pseudo inverse of the other one.
- **padding** (*int*) – Zero-padding added to both sides of the input. Passed to *Encoder* and *Decoder*.
- **output_padding** (*int*) – Additional size added to one side of the output shape. Passed to *Decoder*.
- ****kwargs** – Arguments which will be passed to the filterbank class additionally to the usual *n_filters*, *kernel_size* and *stride*. Depends on the filterbank family.

Returns *Encoder*, *Decoder*

class `asteroid_filterbanks.get`

Returns a filterbank class from a string. Returns its input if it is callable (already a *Filterbank* for example).

Parameters *identifier* (*str* or *Callable* or *None*) – the filterbank identifier.

Returns *Filterbank* or *None*

10.2 Learnable filterbanks

10.2.1 Free

```
class asteroid_filterbanks.free_fb.FreeFB(n_filters, kernel_size, stride=None, sample_rate=8000.0, **kwargs)
```

Bases: `asteroid_filterbanks.enc_dec.Filterbank`

Free filterbank without any constraints. Equivalent to `nn.Conv1d`.

Parameters

- **n_filters** (*int*) – Number of filters.
- **kernel_size** (*int*) – Length of the filters.
- **stride** (*int*, *optional*) – Stride of the convolution. If None (default), set to `kernel_size // 2`.
- **sample_rate** (*float*) – Sample rate of the expected audio. Defaults to 8000.

Variables **n_feats_out** (*int*) – Number of output filters.

References [1] : “Filterbank design for end-to-end speech separation”. ICASSP 2020. Manuel Pariente, Samuele Cornell, Antoine Deleforge, Emmanuel Vincent.

filters ()

Abstract method for filters.

10.2.2 Analytic Free

```
class asteroid_filterbanks.analytic_free_fb.AnalyticFreeFB(n_filters, kernel_size,  
                                                         stride=None, sample_rate=8000.0,  
                                                         **kwargs)
```

Bases: `asteroid_filterbanks.enc_dec.Filterbank`

Free analytic (fully learned with analyticity constraints) filterbank. For more details, see [1].

Parameters

- **n_filters** (*int*) – Number of filters. Half of *n_filters* will have parameters, the other half will be the hilbert transforms. *n_filters* should be even.
- **kernel_size** (*int*) – Length of the filters.
- **stride** (*int*, *optional*) – Stride of the convolution. If None (default), set to `kernel_size // 2`.
- **sample_rate** (*float*) – Sample rate of the expected audio. Defaults to 8000.

Variables **n_feats_out** (*int*) – Number of output filters.

References [1] : “Filterbank design for end-to-end speech separation”. ICASSP 2020. Manuel Pariente, Samuele Cornell, Antoine Deleforge, Emmanuel Vincent.

filters ()

Abstract method for filters.

10.2.3 Parameterized Sinc

```
class asteroid_filterbanks.param_sinc_fb.ParamSincFB(n_filters, kernel_size,
                                                    stride=None, sample_rate=16000.0,
                                                    min_low_hz=50,
                                                    min_band_hz=50, **kwargs)
```

Bases: `asteroid_filterbanks.enc_dec.Filterbank`

Extension of the parameterized filterbank from [1] proposed in [2]. Modified and extended from <https://github.com/mravanelli/SincNet>

Parameters

- **n_filters** (*int*) – Number of filters. Half of *n_filters* (the real parts) will have parameters, the other half will correspond to the imaginary parts. *n_filters* should be even.
- **kernel_size** (*int*) – Length of the filters.
- **stride** (*int*, *optional*) – Stride of the convolution. If None (default), set to `kernel_size // 2`.
- **sample_rate** (*float*, *optional*) – The sample rate (used for initialization).
- **min_low_hz** (*int*, *optional*) – Lowest low frequency allowed (Hz).
- **min_band_hz** (*int*, *optional*) – Lowest band frequency allowed (Hz).

Variables **n_feats_out** (*int*) – Number of output filters.

References [1]: “Speaker Recognition from raw waveform with SincNet”. SLT 2018. Mirco Ravanelli, Yoshua Bengio. <https://arxiv.org/abs/1808.00158>

[2]: “Filterbank design for end-to-end speech separation”. ICASSP 2020. Manuel Pariente, Samuele Cornell, Antoine Deleforge, Emmanuel Vincent. <https://arxiv.org/abs/1910.10400>

filters ()

Compute filters from parameters

get_config ()

Returns dictionary of arguments to re-instantiate the class.

10.3 Fixed filterbanks

10.3.1 STFT

```
class asteroid_filterbanks.stft_fb.STFTFB(n_filters, kernel_size, stride=None, window=None,
                                           sample_rate=8000.0, **kwargs)
```

Bases: `asteroid_filterbanks.enc_dec.Filterbank`

STFT filterbank.

Parameters

- **n_filters** (*int*) – Number of filters. Determines the length of the STFT filters before windowing.
- **kernel_size** (*int*) – Length of the filters (i.e the window).

- **stride** (*int*, *optional*) – Stride of the convolution (hop size). If None (default), set to `kernel_size // 2`.
- **window** (`numpy.ndarray`, *optional*) – If None, defaults to `np.sqrt(np.hanning())`.
- **sample_rate** (*float*) – Sample rate of the expected audio. Defaults to 8000.

Variables **n_feats_out** (*int*) – Number of output filters.

filters ()

Abstract method for filters.

`asteroid_filterbanks.stft_fb.perfect_synthesis_window` (*analysis_window*, *hop_size*)

Computes a window for perfect synthesis given an analysis window and a hop size.

Parameters

- **analysis_window** (`np.array`) – Analysis window of the transform.
- **hop_size** (*int*) – Hop size in number of samples.

Returns `np.array` – the synthesis window to use for perfectly inverting the STFT.

10.3.2 MelGram

```
class asteroid_filterbanks.melgram_fb.MelGramFB(n_filters, kernel_size, stride=None,
                                                window=None, sample_rate=8000.0,
                                                n_mels=40, fmin=0.0, fmax=None,
                                                norm='slaney', **kwargs)
```

Bases: `asteroid_filterbanks.stft_fb.STFTFB`

Mel magnitude spectrogram filterbank.

Parameters

- **n_filters** (*int*) – Number of filters. Determines the length of the STFT filters before windowing.
- **kernel_size** (*int*) – Length of the filters (i.e the window).
- **stride** (*int*, *optional*) – Stride of the convolution (hop size). If None (default), set to `kernel_size // 2`.
- **window** (`numpy.ndarray`, *optional*) – If None, defaults to `np.sqrt(np.hanning())`.
- **sample_rate** (*float*) – Sample rate of the expected audio. Defaults to 8000.
- **n_mels** (*int*) – Number of mel bands.
- **fmin** (*float*) – Minimum frequency of the mel filters.
- **fmax** (*float*) – Maximum frequency of the mel filters. Defaults to `sample_rate//2`.
- **norm** (*str*) – Mel normalization {None, 'slaney', or number}. See `librosa.filters.mel`
- ****kwargs** –

post_analysis (*spec*: `<sphinx.ext.autodoc.importer._MockObject object at 0x7ff682c93f90>`)

Apply transform to encoder convolution.

get_config()

Returns dictionary of arguments to re-instantiate the class. Needs to be subclassed if the filterbanks takes additional arguments than `n_filters` `kernel_size` `stride` and `sample_rate`.

```
class asteroid_filterbanks.melgram_fb.MelScale (n_filters, sample_rate=8000.0,
                                              n_mels=40, fmin=0.0, fmax=None,
                                              norm='slaney')
```

Bases: `sphinx.ext.autodoc.importer._MockObject`

Mel-scale filterbank matrix.

Parameters

- **n_filters** (*int*) – Number of filters. Determines the length of the STFT filters before windowing.
- **sample_rate** (*float*) – Sample rate of the expected audio. Defaults to 8000.
- **n_mels** (*int*) – Number of mel bands.
- **fmin** (*float*) – Minimum frequency of the mel filters.
- **fmax** (*float*) – Maximum frequency of the mel filters. Defaults to `sample_rate//2`.
- **norm** (*str*) – Mel normalization {None, 'slaney', or number}. See `librosa.filters.mel`

10.3.3 MPGT

```
class asteroid_filterbanks.multiphase_gammatone_fb.MultiphaseGammatoneFB (n_filters=128,
                                                                    kernel_size=16,
                                                                    sample_rate=8000.0,
                                                                    stride=None,
                                                                    **kwargs)
```

Bases: `asteroid_filterbanks.enc_dec.Filterbank`

Multi-Phase Gammatone Filterbank as described in [1].

Please cite [1] whenever using this.

Original code repository:

Parameters

- **n_filters** (*int*) – Number of filters.
- **kernel_size** (*int*) – Length of the filters.
- **sample_rate** (*float*, *optional*) – The sample rate (used for initialization).
- **stride** (*int*, *optional*) – Stride of the convolution. If None (default), set to `kernel_size // 2`.

References [1] David Ditter, Timo Gerkmann, “A Multi-Phase Gammatone Filterbank for Speech Separation via TasNet”, ICASSP 2020 Available: <https://ieeexplore.ieee.org/document/9053602/>

filters()

Abstract method for filters.

10.4 Transforms

10.4.1 Griffin-Lim and MISI

`asteroid_filterbanks.griffin_lim.griffin_lim`(*mag_specgram*, *stft_enc*, *angles=None*,
istft_dec=None, *n_iter=6*, *momentum=0.9*)

Estimates matching phase from magnitude spectrogram using the ‘fast’ Griffin Lim algorithm [1].

Parameters

- **mag_specgram** (*torch.Tensor*) – (any, dim, ension, freq, frames) as returned by *Encoder(STFTFB)*, the magnitude spectrogram to be inverted.
- **stft_enc** (*Encoder[STFTFB]*) – The *Encoder(STFTFB())* object that was used to compute the input *mag_spec*.
- **angles** (*None* or *Tensor*) – Angles to use to initialize the algorithm. If *None* (default), angles are init with uniform ditribution.
- **istft_dec** (*None* or *Decoder[STFTFB]*) – Optional Decoder to use to get back to the time domain. If *None* (default), a perfect reconstruction Decoder is built from *stft_enc*.
- **n_iter** (*int*) – Number of griffin-lim iterations to run.
- **momentum** (*float*) – The momentum of fast Griffin-Lim. Original Griffin-Lim is obtained for momentum=0.

Returns *torch.Tensor* – estimated waveforms of shape (any, dim, ension, time).

Examples

```
>>> stft = Encoder(STFTFB(n_filters=256, kernel_size=256, stride=128))
>>> wav = torch.randn(2, 1, 8000)
>>> spec = stft(wav)
>>> masked_spec = spec * torch.sigmoid(torch.randn_like(spec))
>>> mag = transforms.mag(masked_spec, -2)
>>> est_wav = griffin_lim(mag, stft, n_iter=32)
```

References [1] Perraudin et al. “A fast Griffin-Lim algorithm,” WASPAA 2013.

[2] D. W. Griffin and J. S. Lim: “Signal estimation from modified short-time Fourier transform,” ASSP 1984.

`asteroid_filterbanks.griffin_lim.misi`(*mixture_wav*, *mag_specgrams*, *stft_enc*, *angles=None*,
istft_dec=None, *n_iter=6*, *momentum=0.0*, *src_weights=None*, *dim=1*)

Jointly estimates matching phase from magnitude spectrograms using the Multiple Input Spectrogram Inversion (MISI) algorithm [1].

Parameters

- **mixture_wav** (*torch.Tensor*) – (batch, time)
- **mag_specgrams** (*torch.Tensor*) – (batch, n_src, freq, frames) as returned by *Encoder(STFTFB)*, the magnitude spectrograms to be jointly inverted using MISI (modified or not).
- **stft_enc** (*Encoder[STFTFB]*) – The *Encoder(STFTFB())* object that was used to compute the input *mag_spec*.

- **angles** (*None* or *Tensor*) – Angles to use to initialize the algorithm. If *None* (default), angles are init with uniform ditribution.
- **istft_dec** (*None* or *Decoder[STFTFB]*) – Optional Decoder to use to get back to the time domain. If *None* (default), a perfect reconstruction Decoder is built from *stft_enc*.
- **n_iter** (*int*) – Number of MISI iterations to run.
- **momentum** (*float*) – Momentum on updates (this argument comes from GriffinLim). Defaults to 0 as it was never proposed anywhere.
- **src_weights** (*None* or *torch.Tensor*) – Consistency weight for each source. Shape needs to be broadcastable to *istft_dec(mag_specgrams)*. We make sure that the weights sum up to 1 along dim *dim*. If *src_weights* is *None*, compute them based on relative power.
- **dim** (*int*) – Axis which contains the sources in *mag_specgrams*. Used for consistency constraint.

Returns *torch.Tensor* – estimated waveforms of shape (batch, n_src, time).

Examples

```
>>> stft = Encoder(STFTFB(n_filters=256, kernel_size=256, stride=128))
>>> wav = torch.randn(2, 3, 8000)
>>> specs = stft(wav)
>>> masked_specs = specs * torch.sigmoid(torch.randn_like(specs))
>>> mag = transforms.mag(masked_specs, -2)
>>> est_wav = misi(wav.sum(1), mag, stft, n_iter=32)
```

References [1] Gunawan and Sen, “Iterative Phase Estimation for the Synthesis of Separated Sources From Single-Channel Mixtures,” in IEEE Signal Processing Letters, 2010.

[2] Wang, LeRoux et al. “End-to-End Speech Separation with Unfolded Iterative Phase Reconstruction.” Interspeech 2018 (2018)

10.4.2 Complex transforms

`asteroid_filterbanks.transforms.mul_c(inp, other, dim: int = -2)`

Entrywise product for complex valued tensors.

Operands are assumed to have the real parts of each entry followed by the imaginary parts of each entry along dimension *dim*, e.g. for, *dim* = 1, the matrix

is interpreted as

where *j* is such that $j * j = -I$.

Parameters

- **inp** (*torch.Tensor*) – The first operand with real and imaginary parts concatenated on the *dim* axis.
- **other** (*torch.Tensor*) – The second operand.
- **dim** (*int*, *optional*) – frequency (or equivalent) dimension along which real and imaginary values are concatenated.

Returns

torch.Tensor – The complex multiplication between *inp* and *other*

For now, it assumes that *other* has the same shape as *inp* along *dim*.

```
asteroid_filterbanks.transforms.reim(x, dim: int = -2) → Tuple[<sphinx.ext.autodoc.importer._MockObject
object at 0x7ff682cbc6d0>,
<sphinx.ext.autodoc.importer._MockObject object
at 0x7ff682cbc850>]
```

Returns a tuple (re, im).

Parameters

- **x** (`torch.Tensor`) – Complex valued tensor.
- **dim** (`int`) – frequency (or equivalent) dimension along which real and imaginary values are concatenated.

```
asteroid_filterbanks.transforms.mag(x, dim: int = -2, EPS: float = 1e-08)
```

Takes the magnitude of a complex tensor.

The operands is assumed to have the real parts of each entry followed by the imaginary parts of each entry along dimension *dim*, e.g. for, *dim* = 1, the matrix

is interpreted as

where *j* is such that $j * j = -I$.

Parameters

- **x** (`torch.Tensor`) – Complex valued tensor.
- **dim** (`int`) – frequency (or equivalent) dimension along which real and imaginary values are concatenated.

Returns `torch.Tensor` – The magnitude of x.

```
asteroid_filterbanks.transforms.magreim(x, dim: int = -2)
```

Returns a concatenation of (mag, re, im).

Parameters

- **x** (`torch.Tensor`) – Complex valued tensor.
- **dim** (`int`) – frequency (or equivalent) dimension along which real and imaginary values are concatenated.

```
asteroid_filterbanks.transforms.apply_real_mask(tf_rep, mask, dim: int = -2)
```

Applies a real-valued mask to a real-valued representation.

It corresponds to ReIm mask in [1].

Parameters

- **tf_rep** (`torch.Tensor`) – The time frequency representation to apply the mask to.
- **mask** (`torch.Tensor`) – The real-valued mask to be applied.
- **dim** (`int`) – Kept to have the same interface with the other ones.

Returns `torch.Tensor` – *tf_rep* multiplied by the *mask*.

```
asteroid_filterbanks.transforms.apply_mag_mask(tf_rep, mask, dim: int = -2)
```

Applies a real-valued mask to a complex-valued representation.

If *tf_rep* has 2N elements along *dim*, *mask* has N elements, *mask* is duplicated along *dim* to apply the same mask to both the Re and Im.

tf_rep is assumed to have the real parts of each entry followed by the imaginary parts of each entry along dimension dim , e.g. for, $dim = 1$, the matrix

is interpreted as

where j is such that $j * j = -I$.

Parameters

- **tf_rep** (`torch.Tensor`) – The time frequency representation to apply the mask to. Re and Im are concatenated along dim .
- **mask** (`torch.Tensor`) – The real-valued mask to be applied.
- **dim** (`int`) – The frequency (or equivalent) dimension of both tf_rep and $mask$ along which real and imaginary values are concatenated.

Returns `torch.Tensor` – tf_rep multiplied by the $mask$.

`asteroid_filterbanks.transforms.apply_complex_mask(tf_rep, mask, dim: int = -2)`

Applies a complex-valued mask to a complex-valued representation.

Operands are assumed to have the real parts of each entry followed by the imaginary parts of each entry along dimension dim , e.g. for, $dim = 1$, the matrix

is interpreted as

where j is such that $j * j = -I$.

Parameters

- **tf_rep** (`torch.Tensor`) – The time frequency representation to apply the mask to.
- **(class (mask) – torch.Tensor)**: The complex-valued mask to be applied.
- **dim** (`int`) – The frequency (or equivalent) dimension of both tf_rep and $mask$ along which real and imaginary values are concatenated.

Returns `torch.Tensor` – tf_rep multiplied by the $mask$ in the complex sense.

`asteroid_filterbanks.transforms.is_asteroid_complex(tensor, dim: int = -2)`

Check if tensor is complex-like in a given dimension.

Parameters

- **tensor** (`torch.Tensor`) – tensor to be checked.
- **dim** (`int`) – the frequency (or equivalent) dimension along which real and imaginary values are concatenated.

Returns True if dimension is even in the specified dimension, otherwise False

`asteroid_filterbanks.transforms.check_complex(tensor, dim: int = -2)`

Assert that tensor is an Asteroid-style complex in a given dimension.

Parameters

- **tensor** (`torch.Tensor`) – tensor to be checked.
- **dim** (`int`) – the frequency (or equivalent) dimension along which real and imaginary values are concatenated.

Raises `AssertionError` if dimension is not even in the specified dimension

`asteroid_filterbanks.transforms.to_numpy(tensor, dim: int = -2)`

Convert complex-like torch tensor to numpy complex array

Parameters

- **tensor** (*torch.Tensor*) – Complex tensor to convert to numpy.
- **dim** (*int*, *optional*) – the frequency (or equivalent) dimension along which real and imaginary values are concatenated.

Returns `numpy.array` – Corresponding complex array.

`asteroid_filterbanks.transforms.from_numpy(array, dim: int = -2)`

Convert complex numpy array to complex-like torch tensor.

Parameters

- **array** (*np.array*) – array to be converted.
- **dim** (*int*, *optional*) – the frequency (or equivalent) dimension along which real and imaginary values are concatenated.

Returns `torch.Tensor` – Corresponding torch.Tensor (complex axis in dim ‘dim’=

`asteroid_filterbanks.transforms.is_torchaudio_complex(x)`

Check if tensor is Torchaudio-style complex-like (last dimension is 2).

Parameters **x** (*torch.Tensor*) – tensor to be checked.

Returns True if last dimension is 2, else False.

`asteroid_filterbanks.transforms.check_torchaudio_complex(tensor)`

Assert that tensor is Torchaudio-style complex-like (last dimension is 2).

Parameters **tensor** (*torch.Tensor*) – tensor to be checked.

Raises `AssertionError` if last dimension is != 2.

`asteroid_filterbanks.transforms.to_torchaudio(tensor, dim: int = -2)`

Converts complex-like torch tensor to torchaudio style complex tensor.

Parameters

- **tensor** (*torch.tensor*) – asteroid-style complex-like torch tensor.
- **dim** (*int*, *optional*) – the frequency (or equivalent) dimension along which real and imaginary values are concatenated.

Returns `torch.Tensor` – torchaudio-style complex-like torch tensor.

`asteroid_filterbanks.transforms.from_torchaudio(tensor, dim: int = -2)`

Converts torchaudio style complex tensor to complex-like torch tensor.

Parameters

- **tensor** (*torch.tensor*) – torchaudio-style complex-like torch tensor.
- **dim** (*int*, *optional*) – the frequency (or equivalent) dimension along which real and imaginary values are concatenated.

Returns `torch.Tensor` – asteroid-style complex-like torch tensor.

`asteroid_filterbanks.transforms.to_torch_complex(tensor, dim: int = -2)`

Converts complex-like torch tensor to native PyTorch complex tensor.

Parameters

- **tensor** (*torch.tensor*) – asteroid-style complex-like torch tensor.
- **dim** (*int*, *optional*) – the frequency (or equivalent) dimension along which real and imaginary values are concatenated.

Returns `torch.Tensor` – Pytorch native complex-like torch tensor.

`asteroid_filterbanks.transforms.from_torch_complex` (*tensor*, *dim*: *int* = -2)

Converts Pytorch native complex tensor to complex-like torch tensor.

Parameters

- **tensor** (*torch.tensor*) – PyTorch native complex-like torch tensor.
- **dim** (*int*, *optional*) – the frequency (or equivalent) dimension along which real and imaginary values are concatenated.

Returns `torch.Tensor` – asteroid-style complex-like torch tensor.

`asteroid_filterbanks.transforms.angle` (*tensor*, *dim*: *int* = -2)

Return the angle of the complex-like torch tensor.

Parameters

- **tensor** (*torch.Tensor*) – the complex tensor from which to extract the phase.
- **dim** (*int*, *optional*) – the frequency (or equivalent) dimension along which real and imaginary values are concatenated.

Returns `torch.Tensor` – The counterclockwise angle from the positive real axis on the complex plane in radians.

`asteroid_filterbanks.transforms.from_magphase` (*mag_spec*, *phase*, *dim*: *int* = -2)

Return a complex-like torch tensor from magnitude and phase components.

Parameters

- **mag_spec** (*torch.tensor*) – magnitude of the tensor.
- **phase** (*torch.tensor*) – angle of the tensor
- **dim** (*int*, *optional*) – the frequency (or equivalent) dimension along which real and imaginary values are concatenated.

Returns `torch.Tensor` – The corresponding complex-like torch tensor.

`asteroid_filterbanks.transforms.magphase` (*spec*: `<sphinx.ext.autodoc.importer._MockObject object at 0x7ff682cbca50>`, *dim*: *int* = -2) → `Tuple[<sphinx.ext.autodoc.importer._MockObject object at 0x7ff682cbce10>, <sphinx.ext.autodoc.importer._MockObject object at 0x7ff682cbcd0>]`

Splits Asteroid complex-like tensor into magnitude and phase.

`asteroid_filterbanks.transforms.centerfreq_correction` (*spec*: `<sphinx.ext.autodoc.importer._MockObject object at 0x7ff682cbcd90>`, *kernel_size*: *int*, *stride*: *int* = *None*, *dim*: *int* = -2) → `<sphinx.ext.autodoc.importer._MockObject object at 0x7ff682cbcf50>`

Corrects phase from the input spectrogram so that a sinusoid in the middle of a bin keeps the same phase from one frame to the next.

Parameters

- **spec** – Spectrogram tensor of shape (batch, *n_freq* + 2, frames).
- **kernel_size** (*int*) – Kernel size of the STFT.
- **stride** (*int*) – Stride of the STFT.

- **dim** (*int*) – Only works of dim=-2.

Returns *Tensor* – the input spec with corrected phase.

```
asteroid_filterbanks.transforms.phase_centerfreq_correction (phase:
    <sphinx.ext.autodoc.importer._MockObject
    object          at
    0x7ff682cbcf90>,
    kernel_size:
    int,   stride:   int
    =     None)      →
    <sphinx.ext.autodoc.importer._MockObject
    object          at
    0x7ff682cbce90>
```

Corrects phase so that a sinusoid in the middle of a bin keeps the same phase from one frame to the next.

Parameters

- **phase** – tensor of shape (batch, n_freq//2 + 1, frames)
- **kernel_size** (*int*) – Kernel size of the STFT.
- **stride** (*int*) – Stride of the STFT.

Returns *Tensor* – corrected phase.

11.1 Convolutional blocks

11.2 Recurrent blocks

11.3 Attention blocks

11.4 Norms

11.5 Complex number support

12.1 Base classes

12.2 Ready-to-use models

12.3 Publishing models

13.1 Permutation invariant training (PIT) made easy

Asteroid supports regular Permutation Invariant Training (PIT), it's extension using Sinkhorn algorithm (SinkPIT) as well as Mixture Invariant Training (MixIT).

13.1.1 PIT

13.1.2 MixIT

13.1.3 SinkPIT

13.2 Available loss functions

`PITLossWrapper` supports three types of loss function. For “easy” losses, we implement the three types (pairwise point, single-source loss and multi-source loss). For others, we only implement the single-source loss which can be aggregated into both PIT and nonPIT training.

13.2.1 MSE

13.2.2 SDR

13.2.3 PMSQE

13.2.4 STOI

13.2.5 MultiScale Spectral Loss

13.2.6 Deep clustering (Affinity) loss

13.3 Computing metrics

CHAPTER 14

Lightning Wrapper

As explained in *Training and Evaluation*, Asteroid provides a thin wrapper on the top of `PyTorchLightning` for training your models.

15.1 Optimizers

Asteroid relies on `torch_optimizer` and `torch` for optimizers. We provide a simple `get` method that retrieves optimizers from string, which makes it easy to specify optimizers from the command line.

Here is a list of supported optimizers, retrievable from string:

- AccSGD
- AdaBound
- AdaMod
- DiffGrad
- Lamb
- NovoGrad
- PID
- QHAdam
- QHM
- RAdam
- SGDW
- Yogi
- Ranger
- RangerQH
- RangerVA
- Adam
- RMSprop

- SGD
- Adadelta
- Adagrad
- Adamax
- AdamW
- ASG

15.2 Schedulers

Asteroid provides step-wise learning schedulers, integrable to `pytorch-lightning` via `System`.

CHAPTER 16

DSP Modules

16.1 LambdaOverlapAdd

16.2 DualPath Processing

16.3 Mixture Consistency

16.4 VAD

16.5 Delta Features

17.1 Parser utils

Asteroid has its own argument parser (built on `argparse`) that handles dict-like structure, created from a config YAML file.

17.2 Torch utils

17.3 Hub utils

17.4 Generic utils

Asteroid High-Level Contribution Guide

Asteroid is a Pytorch-based audio source separation toolkit that enables fast experimentation on common datasets.

18.1 The Asteroid Contribution Process

The Asteroid development process involves a healthy amount of open discussions between the core development team and the community.

Asteroid operates similar to most open source projects on GitHub. However, if you've never contributed to an open source project before, here is the basic process.

- **Figure out what you're going to work on.** The majority of open source contributions come from people scratching their own itches. However, if you don't know what you want to work on, or are just looking to get more acquainted with the project, here are some tips for how to find appropriate tasks:
 - Look through the [issue tracker](#) and see if there are any issues you know how to fix. Issues that are confirmed by other contributors tend to be better to investigate.
 - Join us on Slack and let us know you're interested in getting to know Asteroid. We're very happy to help out researchers and partners get up to speed with the codebase.
- **Figure out the scope of your change and reach out for design comments on a GitHub issue if it's large.** The majority of pull requests are small; in that case, no need to let us know about what you want to do, just get cracking. But if the change is going to be large, it's usually a good idea to get some design comments about it first.
 - If you don't know how big a change is going to be, we can help you figure it out! Just post about it on issues or Slack.
 - Some feature additions are very standardized; for example, lots of people add new datasets or architectures to Asteroid. Design discussion in these cases boils down mostly to, "Do we want this dataset/architecture?" Giving evidence for its utility, e.g., usage in peer reviewed papers, or existence in other frameworks, helps a bit when making this case.

- Core changes and refactors can be quite difficult to coordinate, as the pace of development on Asteroid master is quite fast. Definitely reach out about fundamental or cross-cutting changes; we can often give guidance about how to stage such changes into more easily reviewable pieces.
- **Code it out!**
 - See the technical guide and read the code for advice for working with Asteroid in a technical form.
- **Open a pull request.**
 - If you are not ready for the pull request to be reviewed, tag it with [WIP]. We will ignore it when doing review passes. If you are working on a complex change, it's good to start things off as WIP, because you will need to spend time looking at CI results to see if things worked out or not.
 - Find an appropriate reviewer for your change. We have some folks who regularly go through the PR queue and try to review everything, but if you happen to know who the maintainer for a given subsystem affected by your patch is, feel free to include them directly on the pull request.
- **Iterate on the pull request until it's accepted!**
 - We'll try our best to minimize the number of review roundtrips and block PRs only when there are major issues. For the most common issues in pull requests, take a look at [Common Mistakes](#).
 - Once a pull request is accepted and CI is passing, there is nothing else you need to do; we will merge the PR for you.

18.2 Getting Started

18.2.1 Proposing new features

New feature ideas are best discussed on a specific issue. Please include as much information as you can, any accompanying data, and your proposed solution. The Asteroid team and community frequently reviews new issues and comments where they think they can help. If you feel confident in your solution, go ahead and implement it.

18.2.2 Reporting Issues

If you've identified an issue, first search through the [list of existing issues](#) on the repo. If you are unable to find a similar issue, then create a new one. Supply as much information you can to reproduce the problematic behavior. Also, include any additional insights like the behavior you expect.

18.2.3 Implementing Features or Fixing Bugs

If you want to fix a specific issue, it's best to comment on the individual issue with your intent. However, we do not lock or assign issues except in cases where we have worked with the developer before. It's best to strike up a conversation on the issue and discuss your proposed solution. We can provide guidance that saves you time.

18.2.4 Adding Tutorials

Most our tutorials come from our team but we are very open to additional contributions. Have a notebook leveraging Asteroid? Open a PR to let us know!

18.2.5 Improving Documentation & Tutorials

We aim to produce high quality documentation and tutorials. On some occasions that content includes typos or bugs. If you find something you can fix, send us a pull request for consideration.

Take a look at the *Documentation* section to learn how our system works.

18.2.6 Participating in online discussions

You can find active discussions happening on our [slack workspace](#).

18.2.7 Submitting pull requests to fix open issues

You can view a list of all open issues [here](#). Commenting on an issue is a great way to get the attention of the team. From here you can share your ideas and how you plan to resolve the issue.

For more challenging issues, the team will provide feedback and direction for how to best solve the issue.

If you're not able to fix the issue itself, commenting and sharing whether you can reproduce the issue can be useful for helping the team identify problem areas.

18.2.8 Reviewing open pull requests

We appreciate your help reviewing and commenting on pull requests. Our team strives to keep the number of open pull requests at a manageable size, we respond quickly for more information if we need it, and we merge PRs that we think are useful. However, additional eyes on pull requests is always appreciated.

18.2.9 Improving code readability

Improve code readability helps everyone. We plan to integrate `black/DeepSource` in the CI process, but readability issues can still persist and we'll welcome your corrections.

18.2.10 Adding test cases to make the codebase more robust

Additional test coverage is **always** appreciated.

18.2.11 Promoting Asteroid

Your use of Asteroid in your projects, research papers, write ups, blogs, or general discussions around the internet helps to raise awareness for Asteroid and our growing community. Please reach out to [us](#) for support.

18.2.12 Triaging issues

If you feel that an issue could benefit from a particular tag or level of complexity comment on the issue and share your opinion. If an you feel an issue isn't categorized properly comment and let the team know.

18.3 About open source development

If this is your first time contributing to an open source project, some aspects of the development process may seem unusual to you.

- **There is no way to “claim” issues.** People often want to “claim” an issue when they decide to work on it, to ensure that there isn’t wasted work when someone else ends up working on it. This doesn’t really work too well in open source, since someone may decide to work on something, and end up not having time to do it. Feel free to give information in an advisory fashion, but at the end of the day, we will take running code and rough consensus.
- **There is a high bar for new functionality that is added.** Unlike in a corporate environment, where the person who wrote code implicitly “owns” it and can be expected to take care of it in the beginning of its lifetime, once a pull request is merged into an open source project, it immediately becomes the collective responsibility of all maintainers on the project. When we merge code, we are saying that we, the maintainers, are able to review subsequent changes and make a bugfix to the code. This naturally leads to a higher standard of contribution.

18.4 Common Mistakes To Avoid

- **Did you add tests?** (Or if the change is hard to test, did you describe how you tested your change?)
 - We have a few motivations for why we ask for tests:
 1. to help us tell if we break it later
 2. to help us tell if the patch is correct in the first place (yes, we did review it, but as Knuth says, “beware of the following code, for I have not run it, merely proven it correct”)
 - When is it OK not to add a test? Sometimes a change can’t be conveniently tested, or the change is so obviously correct (and unlikely to be broken) that it’s OK not to test it. On the contrary, if a change is seems likely (or is known to be likely) to be accidentally broken, it’s important to put in the time to work out a testing strategy.
- **Is your PR too long?** It’s easier for us to review and merge small PRs. Difficulty of reviewing a PR scales nonlinearly with its size. You can try to split it up if possible, else it helps if there is a complete description of the contents of the PR: it’s easier to review code if we know what’s inside!
- **Comments for subtle things?** In cases where behavior of your code is nuanced, please include extra comments and documentation to allow us to better understand the intention of your code.
- **Did you add a hack?** Sometimes a hack is the right answer. But usually we will have to discuss it.
- **Do you want to touch a very core component?** In order to prevent major regressions, pull requests that touch core components receive extra scrutiny. Make sure you’ve discussed your changes with the team before undertaking major changes.
- **Want to add a new feature?** If you want to add new features, comment your intention on the related issue. Our team tries to comment on and provide feedback to the community. It’s better to have an open discussion with the team and the rest of the community prior to building new features. This helps us stay aware of what you’re working on and increases the chance that it’ll be merged.
- **Did you touch unrelated code to the PR?** To aid in code review, please only include files in your pull request that are directly related to your changes.

18.5 Frequently asked questions

- **How can I contribute as a reviewer?** There is lots of value if community developer reproduce issues, try out new functionality, or otherwise help us identify or troubleshoot issues. Commenting on tasks or pull requests with your environment details is helpful and appreciated.
- **CI tests failed, what does it mean?** Maybe you need to merge with master or rebase with latest changes. Pushing your changes should re-trigger CI tests. If the tests persist, you'll want to trace through the error messages and resolve the related issues.

18.6 Attribution

This Contribution Guide is adapted from PyTorch's Contribution Guide available [here](#).

CHAPTER 19

How to contribute

The general way to contribute to Asteroid is to fork the main repository on GitHub:

1. Fork the [main repo](#) and `git clone` it.
2. Make your changes, test them, commit them and push them to your fork.
3. You can open a pull request on GitHub when you're satisfied.

Things don't need to be perfect for PRs to be opened.

If you made changes to the source code, you'll want to try them out without installing asteroid everytime you change something. To do that, install asteroid in develop mode either with `pip pip install -e .[tests]` or with `python python setup.py develop`.

To avoid formatting roundtrips in PRs, Asteroid relies on `'black'` [\(<https://github.com/psf/black>'](https://github.com/psf/black) and `'pre-commit-hooks'` [\(<https://github.com/pre-commit/pre-commit-hooks>'](https://github.com/pre-commit/pre-commit-hooks) to handle formatting for us. You'll need to install `requirements/dev.txt` and install git hooks with `pre-commit install`.

Here is a summary:

```
### Install
git clone your_fork_url
cd asteroid
pip install -r requirements/dev.txt
pip install -e .
pre-commit install # To run black before commit

# Make your changes
# Test them locally
# Commit your changes
# Push your changes
# Open a PR!
```

19.1 Source code contributions

All contributions to the source code of asteroid should be documented and unit-tested. See [here](#) to run the tests with coverage reports. Docstrings follow the [Google format](#), have a look at other docstrings in the codebase for examples. Examples in docstrings can be very useful, don't hesitate to add some!

19.2 Writing new recipes.

Most new recipes should follow the standard format that is described [here](#). We are not dogmatic about it, but another organization should be explained and motivated. We welcome any recipe on standard or new datasets, with standard or new architectures. You can even link a paper submission with a PR number if you'd like!

19.3 Improving the docs.

If you found a typo, think something could be more explicit etc... Improving the documentation is always welcome. The instructions to install dependencies and build the docs can be found [here](#). Docstrings follow the [Google format](#), have a look at other docstrings in the codebase for examples.

19.4 Coding style

We use [pre-commit hooks](#) to format the code using `black`. The code is checked for `black`- and `flake8`- compliance on every commit with GitHub actions. Remember, continuous integration is not here to be all green, be to help us see where to improve !

If you have any question, [open an issue](#) or [join the slack](#), we'll be happy to help you.

CHAPTER 20

Indices and tables

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

a

`asteroid_filterbanks.analytic_free_fb`,
32
`asteroid_filterbanks.free_fb`, 32
`asteroid_filterbanks.griffin_lim`, 36
`asteroid_filterbanks.melgram_fb`, 34
`asteroid_filterbanks.param_sinc_fb`, 33
`asteroid_filterbanks.stft_fb`, 33
`asteroid_filterbanks.transforms`, 37

A

AnalyticFreeFB (class in asteroid_filterbanks.analytic_free_fb), 32
 angle() (in module asteroid_filterbanks.transforms), 41
 apply_complex_mask() (in module asteroid_filterbanks.transforms), 39
 apply_mag_mask() (in module asteroid_filterbanks.transforms), 38
 apply_real_mask() (in module asteroid_filterbanks.transforms), 38
 asteroid_filterbanks.analytic_free_fb (module), 32
 asteroid_filterbanks.free_fb (module), 32
 asteroid_filterbanks.griffin_lim (module), 36
 asteroid_filterbanks.melgram_fb (module), 34
 asteroid_filterbanks.param_sinc_fb (module), 33
 asteroid_filterbanks.stft_fb (module), 33
 asteroid_filterbanks.transforms (module), 37

C

centerfreq_correction() (in module asteroid_filterbanks.transforms), 41
 check_complex() (in module asteroid_filterbanks.transforms), 39
 check_torchaudio_complex() (in module asteroid_filterbanks.transforms), 40

D

Decoder (class in asteroid_filterbanks), 30

E

Encoder (class in asteroid_filterbanks), 30

F

Filterbank (class in asteroid_filterbanks), 29

filters() (asteroid_filterbanks.analytic_free_fb.AnalyticFreeFB method), 32
 filters() (asteroid_filterbanks.Filterbank method), 29
 filters() (asteroid_filterbanks.free_fb.FreeFB method), 32
 filters() (asteroid_filterbanks.multiphase_gammatone_fb.MultiphaseG method), 35
 filters() (asteroid_filterbanks.param_sinc_fb.ParamSincFB method), 33
 filters() (asteroid_filterbanks.stft_fb.STFTFB method), 34
 forward() (asteroid_filterbanks.Decoder method), 31
 forward() (asteroid_filterbanks.Encoder method), 30
 FreeFB (class in asteroid_filterbanks.free_fb), 32
 from_magphase() (in module asteroid_filterbanks.transforms), 41
 from_numpy() (in module asteroid_filterbanks.transforms), 40
 from_torch_complex() (in module asteroid_filterbanks.transforms), 40
 from_torchaudio() (in module asteroid_filterbanks.transforms), 40

G

get (class in asteroid_filterbanks), 31
 get_config() (asteroid_filterbanks.Filterbank method), 29
 get_config() (asteroid_filterbanks.melgram_fb.MelGramFB method), 34
 get_config() (asteroid_filterbanks.param_sinc_fb.ParamSincFB method), 33
 griffin_lim() (in module asteroid_filterbanks.griffin_lim), 36

I

is_asteroid_complex() (in module asteroid_filterbanks.transforms), 39

`is_torchaudio_complex()` (in module `asteroid_filterbanks.transforms`), 40

M

`mag()` (in module `asteroid_filterbanks.transforms`), 38

`magphase()` (in module `asteroid_filterbanks.transforms`), 41

`magreim()` (in module `asteroid_filterbanks.transforms`), 38

`make_enc_dec` (class in `asteroid_filterbanks`), 31

`MelGramFB` (class in `asteroid_filterbanks.melgram_fb`), 34

`MelScale` (class in `asteroid_filterbanks.melgram_fb`), 35

`mis()` (in module `asteroid_filterbanks.griffin_lim`), 36

`mul_c()` (in module `asteroid_filterbanks.transforms`), 37

`MultiphaseGammatoneFB` (class in `asteroid_filterbanks.multiphase_gammatone_fb`), 35

P

`ParamSincFB` (class in `asteroid_filterbanks.param_sinc_fb`), 33

`perfect_synthesis_window()` (in module `asteroid_filterbanks.stft_fb`), 34

`phase_centerfreq_correction()` (in module `asteroid_filterbanks.transforms`), 42

`pinv_of()` (`asteroid_filterbanks.Decoder` class method), 31

`pinv_of()` (`asteroid_filterbanks.Encoder` class method), 30

`post_analysis()` (`asteroid_filterbanks.Filterbank` method), 29

`post_analysis()` (`asteroid_filterbanks.melgram_fb.MelGramFB` method), 34

`post_synthesis()` (`asteroid_filterbanks.Filterbank` method), 29

`pre_analysis()` (`asteroid_filterbanks.Filterbank` method), 29

`pre_synthesis()` (`asteroid_filterbanks.Filterbank` method), 29

R

`reim()` (in module `asteroid_filterbanks.transforms`), 38

S

`STFTFB` (class in `asteroid_filterbanks.stft_fb`), 33

T

`to_numpy()` (in module `asteroid_filterbanks.transforms`), 39

`to_torch_complex()` (in module `asteroid_filterbanks.transforms`), 40

`to_torchaudio()` (in module `asteroid_filterbanks.transforms`), 40