

---

# **thelper Documentation**

***Release 0.6.2***

**Pierre-Luc St-Charles**

**Oct 01, 2020**



# CONTENTS

<b>1</b>	<b>Overview</b>	<b>1</b>
1.1	Notes . . . . .	1
<b>2</b>	<b>FAQ</b>	<b>3</b>
2.1	What it is. . . . .	3
2.2	What it is <b>NOT</b> . . . . .	3
2.3	What it supports. . . . .	3
2.4	How do I. . . . .	4
<b>3</b>	<b>Installation</b>	<b>5</b>
3.1	Docker . . . . .	5
3.2	Installing from source . . . . .	5
3.3	Anaconda . . . . .	6
3.4	Testing the installation . . . . .	6
3.5	Documentation . . . . .	7
<b>4</b>	<b>User Guide</b>	<b>9</b>
4.1	Command-Line Interface . . . . .	9
4.2	Configuration Files . . . . .	11
4.3	Session Directories . . . . .	17
<b>5</b>	<b>Use Cases</b>	<b>21</b>
5.1	Image classification . . . . .	21
5.2	Image segmentation . . . . .	24
5.3	Object Detection . . . . .	24
5.4	Super-resolution . . . . .	24
5.5	Creating a new dataset interface . . . . .	24
5.6	Dataset/Loader visualization . . . . .	24
5.7	Dataset annotation . . . . .	24
5.8	Rebalancing a dataset . . . . .	25
5.9	Exporting a dataset . . . . .	25
5.10	Defining a data augmentation pipeline . . . . .	25
5.11	Supporting a custom trainer . . . . .	25
5.12	Supporting a custom task . . . . .	25
5.13	Supporting a custom model . . . . .	25
5.14	Visualizing metrics using <code>tensorboardX</code> . . . . .	25
5.15	Manually reloading a model . . . . .	25
5.16	Exporting a model . . . . .	26
<b>6</b>	<b>thelper package</b>	<b>29</b>
6.1	Subpackages . . . . .	29

6.2	Submodules	110
6.3	thelper.cli module	110
6.4	thelper.concepts module	114
6.5	thelper.draw module	116
6.6	thelper.ifaces module	118
6.7	thelper.typedefs module	120
6.8	thelper.utils module	120
<b>7</b>	<b>Contributing</b>	<b>127</b>
7.1	Bug reports	127
7.2	Documentation improvements	127
7.3	Feature requests and feedback	127
7.4	Development	128
<b>8</b>	<b>Authors</b>	<b>129</b>
<b>9</b>	<b>Maintainer Guide</b>	<b>131</b>
9.1	Testing	131
9.2	Releases	132
9.3	Documentation	132
9.4	Deployment	132
<b>10</b>	<b>Changelog</b>	<b>135</b>
10.1	Unreleased (latest)	135
10.2	0.6.2 (2020/10/01)	135
10.3	0.6.1 (2020/07/29)	135
10.4	0.6.0 (2020/07/28)	135
10.5	0.5.0 (2020/07/21)	135
10.6	0.5.0-rc2 (2020/07/08)	136
10.7	0.5.0-rc1 (2020/07/07)	136
10.8	0.5.0-rc0 (2020/04/25)	136
10.9	0.4.7 (2019/11/20)	136
10.10	0.4.6 (2019/11/20)	136
10.11	0.4.5 (2019/11/18)	136
10.12	0.4.4 (2019/11/18)	137
10.13	0.4.3 (2019/11/06)	137
10.14	0.4.2 (2019/11/06)	137
10.15	0.4.1 (2019/10/15)	138
10.16	0.4.0 (2019/10/11)	138
10.17	0.3.14 (2019/09/30)	138
10.18	0.3.13 (2019/09/26)	138
10.19	0.3.12 (2019/09/13)	138
10.20	0.3.11 (2019/09/09)	139
10.21	0.3.10 (2019/09/05)	139
10.22	0.3.9 (2019/08/20)	139
10.23	0.3.8 (2019/08/08)	139
10.24	0.3.7 (2019/07/31)	140
10.25	0.3.6 (2019/07/26)	140
10.26	0.3.5 (2019/07/23)	140
10.27	0.3.4 (2019/07/12)	140
10.28	0.3.3 (2019/07/09)	141
10.29	0.3.2 (2019/07/05)	141
10.30	0.3.1 (2019/06/17)	141
10.31	0.3.0 (2019/06/12)	141
10.32	0.2.8 (2019/03/17)	141

10.33 0.2.7 (2019/02/04) . . . . .	142
10.34 0.2.6 (2019/01/31) . . . . .	142
10.35 0.2.5 (2019/01/29) . . . . .	142
10.36 0.2.4 (2019/01/29) . . . . .	142
10.37 0.2.3 (2019/01/29) . . . . .	142
10.38 0.2.2 (2019/01/29) . . . . .	142
10.39 0.2.1 (2019/01/24) . . . . .	143
10.40 0.2.0 (2019/01/15) . . . . .	143
10.41 0.1.1 (2019/01/14) . . . . .	143
10.42 0.1.0 (2018/11/28) . . . . .	143
10.43 0.0.2 (2018/10/18) . . . . .	144
10.44 0.0.1 (2018/10/03) . . . . .	144
<b>11 Indices and tables</b>	<b>145</b>
<b>Python Module Index</b>	<b>147</b>
<b>Index</b>	<b>149</b>



**OVERVIEW**

dependencies	
ci-status	
releases	
packages	

This package provides a training framework and CLI for PyTorch-based machine learning projects. This is free software distributed under the [Apache Software License version 2.0](#) built by researchers and developers from the Centre de Recherche Informatique de Montréal / Computer Research Institute of Montreal (CRIM).

To get a general idea of what this framework can be used for, visit the [FAQ page](#). For installation instructions, refer to the [installation guide](#). For usage instructions, refer to the [user guide](#). The auto-generated documentation is available via [readthedocs.io](#).

## 1.1 Notes

Development is still on-going — the API and internal classes may change in the future.

The project's structure was originally generated by [cookiecutter](#) via [ionelmc's template](#).





We answer some of the simple and frequently asked questions about the framework below. If you think of any other question that should be in this list, send a mail to one of the maintainers, and it will be added here.

## 2.1 What it is...

- This framework is used to simplify the exploration, development, and testing of models that you create yourself, or that you import from other libraries or frameworks.
- This framework is used to enforce good reproducibility standards for your experiments via the use of global configuration files, checkpoints, and logs.
- This framework is used to easily swap, split, scale, combine, and augment datasets used in your experiments.
- This framework can help you fine-tune, debug, measure, visualize, and understand the behavior of your models more easily.

## 2.2 What it is NOT...

- This framework is **NOT** used to obtain off-the-shelf solutions. In most cases, you will have to put in some work by at least modifying a pre-existing configuration file to get a solution for a new task (e.g. image classification, segmentation, ...).
- This framework is **NOT** a model (or model architecture) zoo. By design, importing models for training from 3rd-party packages is easy, and so is importing a model architecture from a local Python file. However, due to the fast-paced nature of deep learning, we do not plan to keep a repository of “state-of-the-art” models in the framework.

## 2.3 What it supports...

- **PyTorch.** For now, the entire backend is based on the design patterns, interfaces, and concepts of the PyTorch library ([\[more info\]](#)).
- Image classification, segmentation, object detection, super-resolution, and generic regression tasks. More types of tasks are planned in the future. Users can also implement their own task interfaces and trainers to support custom scenarios if needed (e.g. for multi-task learning).
- Live evaluation and monitoring of predefined metrics. The framework implements *[several types of metrics]*, but custom metrics can also be defined and evaluated at run time.

- Data augmentation. The framework implements basic *[transformation operations and wrappers]* for large augmentation libraries such as `albumentations` (*[more info]*).
- Model fine-tuning and exportation. Models obtained from the `torchvision` package (*[more info]*) or pre-trained using the framework can be loaded and fine-tuned directly for any compatible task. They can also be exported in PyTorch-JIT/ONNX format for external inference.
- Tensorboard. Event logs are generated using `tensorboardX` (*[more info]*) and may contain plots, visualizations, histograms, graph module trees and more.

## 2.4 How do I...

This section is still a work in progress; see the use case examples *[here]* for a list of code snippets and tutorials on how to use the framework. For high-level documentation, refer to the *[user guide]*.

## INSTALLATION

### 3.1 Docker

Starting with v0.3.2, the latest stable version of the framework is pre-built and available from the [docker hub](#). To get a copy, simply pull it via:

```
$ docker pull plstcharles/thelper
```

You should then be able to launch sessions in containers as such:

```
$ docker run -it plstcharles/thelper thelper <CLI_ARGS_HERE>
```

The image is built from `nvidia/cuda`, meaning that it is compatible with `nvidia-docker` and supports CUDA-enabled GPUs. To run a GPU-enabled container, install [the runtime using these instructions](#), and add `--runtime=nvidia` to the arguments given to `docker run`.

### 3.2 Installing from source

If you wish to modify the framework's source code or develop new modules within the framework itself, follow the installation instructions below.

#### 3.2.1 Linux

You can use the provided Makefile to automatically create a conda environment on your system that will contain the thelper framework and all its dependencies. In your terminal, simply enter:

```
$ cd <THELPER_ROOT>
$ make install
```

If you already have conda installed somewhere, you can force the Makefile to use it for the installation of the new environment by setting the `CONDA_HOME` variable before calling make:

```
$ export CONDA_HOME=/some/path/to/miniconda3
$ cd <THELPER_ROOT>
$ make install
```

The newly created conda environment will be called 'thelper', and can then be activated using:

```
$ conda activate thelper
```

Or, assuming conda is not already in your path:

```
$ source /some/path/to/miniconda3/bin/activate thelper
```

### 3.2.2 Other systems

If you cannot use the Makefile, you will have to install the dependencies yourself. These dependencies are listed in the [requirements file](#), and can also be installed using the conda environment configuration file provided [here](#). For the latter case, call the following from your terminal:

```
$ conda env create --file <THELPER_ROOT>/conda-env.yml -n thelper
```

Then, simply activate your environment and install the thelper package within it:

```
$ conda activate thelper
$ pip install -e <THELPER_ROOT> --no-deps
```

On the other hand, although it is *not* recommended since it tends to break PyTorch, you can install the dependencies directly through pip:

```
$ pip install -r <THELPER_ROOT>/requirements.txt
$ pip install -e <THELPER_ROOT> --no-deps
```

## 3.3 Anaconda

Starting with v0.2.5, a stable version of the framework can be installed directly (with its dependencies) via [Anaconda](#). In a conda environment, simply enter:

```
$ conda config --env --add channels plstcharles
$ conda config --env --add channels conda-forge
$ conda config --env --add channels pytorch
$ conda install thelper
```

This should install a stable version of the framework on Windows and Linux for Python 3.6 or 3.7. You can check the release notes [on GitHub](#), and pre-built packages [here](#).

Note that due to Travis build limitations (as of November 2019), conda package builds and deployments have been stalling and have required manual uploads. This means that the conda packages are fairly likely to be out-of-date compared to those on Docker Hub and PyPI. As such, we now recommend users to install the framework through the “Install from source” method above.

## 3.4 Testing the installation

You should now be able to print the thelper package version number to see if the package is properly installed and that all dependencies can be loaded at runtime:

```
(conda-env:thelper) username@hostname:~/devel/thelper$ python
Python X.Y.Z |Anaconda, Inc.| (default, YYYY XX ZZZ, AA:BB:CC)
[GCC X.Y.Z] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>> import thelper
```

(continues on next page)

(continued from previous page)

```
>>> print(thelper.__version__)  
x.y.z
```

You can now refer to the [\[user guide\]](#) for more information on how to use the framework.

## 3.5 Documentation

The sphinx documentation is generated automatically via [readthedocs.io](#), but it might still be incomplete due to buggy apidoc usage/platform limitations. To build it yourself, use the makefile:

```
$ cd <THELPER_ROOT>  
$ make docs
```

The HTML documentation should then be generated inside `<THELPER_ROOT>/docs/build/html`. To browse it, simply open the `index.html` file there.



## USER GUIDE

This guide provides an overview of the basic functionalities and typical use cases of the *thelper* framework. For installation instructions, refer to the installation guide [\[here\]](#).

Currently, the framework can be used to tackle image classification, image segmentation, object detection, image super-resolution, and generic regression tasks. Models for all of these tasks can be trained out-of-the-box using PyTorch. More task types are expected to follow in the future. The goal of the framework is not to solve those problems for you; its goal is to facilitate your model exploration and development process. This is achieved by providing a centralized interface for the control of all your experiment settings, by offering a simple solution for model checkpointing and fine-tuning, and by providing debugging tools and visualizations to help you understand your model's behavior. It can also help users working with GPU clusters by keeping track of their jobs more easily. This framework will not directly give you the perfect solution for your particular problem, but it will help you discover a solution while enforcing good reproducibility standards.

If your problem is related to one of the aforementioned tasks, and if you can solve this problem using a standard model architecture already included in PyTorch or in the framework itself, then you might be able to train and export a solution without writing a single line of code. It is however typical to work with a custom model, a custom trainer, or even a custom task/objective. This is also supported by the framework, as most classes can be either imported as-is, or they can derive from and replace the internal classes of the framework.

In the sections below, we introduce the framework's *Command-Line Interface (CLI)* used to launch jobs, the *session configuration files* used to define the settings of these jobs, and the *session directories* that contain job outputs. Use cases that show how to use different functionalities of the framework are available in [\[a different section\]](#).

---

### 4.1 Command-Line Interface

The Command-Line Interface (CLI) of the framework offers the main entrypoint from which jobs are executed. A number of different operations are supported; these are detailed in the following subsections, and listed [\[in the documentation\]](#). For now, note that these operations all rely on a configuration dictionary which is typically parsed from a JSON file. The fields of this dictionary that are required by each operation are detailed [in the next section](#).

Note that using the framework's CLI is not mandatory. If you prefer bypassing it and creating your own high-level job dispatcher, you can do so by deconstructing one of the already-existing CLI entrypoints, and by calling the same high-level functions it uses to load the components you need. These might include for example `thelper.data.utils.create_loaders()` and `thelper.nn.utils.create_model()`. Calling those functions directly may also be necessary if you intend on embedding the framework inside another application.

### 4.1.1 Creating a training session

Usage from the terminal:

```
$ thelper new <PATH_TO_CONFIG_FILE.json> <PATH_TO_ROOT_SAVE_DIR>
```

To create a training session, the `new` operation of the CLI is used. This redirects the execution flow to `thelper.cli.create_session()`. The configuration dictionary that is provided must contain all sections required to train a model, namely `datasets`, `loaders`, `model`, and `trainer`. It is also mandatory to provide a `name` field in the global space for the training session to be properly identified later on.

No distinction is made at this stage regarding the task that the training session is tackling. The nature of this task (e.g. image classification) will be deduced from the `datasets` section of the configuration later in the process. This CLI entrypoint can therefore be used to start training sessions for any task.

Finally, note that since starting a training session produces logs and data, the path to a directory where the output can be created must be provided as the second argument.

### 4.1.2 Resuming a training session

Usage from the terminal:

```
$ thelper resume <PATH_TO_SESSION_DIR_OR_CHECKPT> [-m MAP_LOCATION] [-c OVERRIDE_CFG]   
→ [...]
```

If a previously created training session was halted for any reason, it is possible to resume it with the `resume` operation of the CLI. To do so, you must provide either the path to the session directory or to a checkpoint created by the framework. If a directory path is given, it will be searched for checkpoints and the latest one will be loaded. The training session will then be resumed using the loaded model and optimizer state, and subsequent outputs will be saved in the original session directory.

A session can be resumed with an overriding configuration dictionary adding e.g. new metrics. If no configuration is provided at all, the original one contained in the loaded checkpoint will be used. Compatibility between an overriding configuration dictionary and the original one must be ensured by the user. A session can also be resumed only to evaluate the (best) trained model performance on the testing set. This is done by adding the `--eval-only` flag at the end of the command line. For more information on the parameters, see the documentation of `thelper.cli.resume_session()`.

### 4.1.3 Visualizing data

Usage from the terminal:

```
$ thelper viz <PATH_TO_CONFIG_FILE.json>
```

Visualizing the images that will be forwarded to the model during training after applying data augmentation operations can be useful to determine whether they still look natural or not. The `viz` operation of the CLI allows you to do just this. It relies on the dataset parsers or data loaders defined in a configuration dictionary that would normally be given to the CLI under the `new` or `resume` operation modes. For more information on this mode, see the documentation of `thelper.cli.visualize_data()`.



### 4.1.4 Annotating data

Usage from the terminal:

```
$ thelper annot <PATH_TO_CONFIG_FILE.json> <PATH_TO_ROOT_SAVE_DIR>
```

The `annot` CLI operation allows the user to browse a dataset and annotate individual samples from it using a specialized GUI tool. The configuration dictionary that is provided must contain a `datasets` section to define the parsers that load the data, and an `annotator` section that defines the GUI tool settings used to create annotations. During an annotation session, all annotations that are created by the user will be saved into the session directory. For more information on the parameters, refer to the documentation of `thelper.cli.annotate_data()`.

### 4.1.5 Split data

Usage from the terminal:

```
$ thelper split <PATH_TO_CONFIG_FILE.json> <PATH_TO_ROOT_SAVE_DIR>
```

When training a model, the framework will typically split the datasets into non-overlapping data loaders. This split must be performed every time a training session is created or resumed. This can be a lengthy process based on the amount of preprocessing and parsing required by the dataset constructors. Using the `split` CLI operation allows the user to pre-compute this split and archive the training, validation, and test sets into a HDF5 archive. This archive can then be parsed by an interface provided in the framework to speed up the creation/resuming of training sessions, or simply for external tests. See `thelper.data.parsers.HDF5Dataset` for more information on the dataset interface, or `thelper.cli.split_data()` on the operation itself.

### 4.1.6 Export model

Usage from the terminal:

```
$ thelper export <PATH_TO_CONFIG_FILE.json> <PATH_TO_ROOT_SAVE_DIR>
```

The `export` CLI operation allows the user to export a trained model for external use as defined in a configuration file. The export format is a new checkpoint that may optionally contain an optimized version of the model compiled using PyTorch's JIT engine. This is still an experimental feature. See the documentation of `thelper.cli.export_model()` or the [\[example here\]](#) for more information.

[\[to top\]](#)

## 4.2 Configuration Files

Configuration files are at the heart of the framework. These essentially contain all the settings that might affect the behavior of a training session, and therefore of a trained model. The framework itself does not enforce that all parameters must be passed through the configuration file, but it is good to follow this principle, as it helps enforce reproducibility. Configuration files also essentially always contain a dictionary so that parameters can be split into sections. We thus often refer to them as 'configuration dictionaries'.

The framework will automatically skip sections of the configuration file that it does not need to use or that it does not understand. This is useful when sections or subsections are added for custom needs, or when only a portion of the configuration is relevant to some use case (for example, the 'visualization' mode of the CLI will only look at the datasets and data loaders sections).

For now, all configuration files are expected to be in JSON or YAML format. Future versions of the framework should support raw python modules (.py files) that define each subsection as a dictionary. Examples of complete configuration files used for various purposes are available in the `config` directory located with the code ([\[see them here\]](#)).

## 4.2.1 Datasets section

The `datasets` section of the configuration defines the dataset “parsers” that will be instantiated by the framework and passed to the data loaders. These are responsible for parsing the structure of a dataset and providing the total number of samples that it contains. Dataset parsers should expose a `__getitem__` function that returns an individual data sample when queried by index. The dataset parsers provided in the `torchvision.datasets` package are all fully compatible with these requirements.

The configuration section itself should be built like a dictionary of objects to instantiate. The key associated with each parser is the name that will be used to refer to it internally as well as in the `loaders` section. If a dataset parser that does not derive from `thelper.data.parsers.Dataset` is needed, you will have to specify a task object inside its definition. An example configuration based on the CIFAR10 class provided by `torchvision` ([\[more info here\]](#)) is shown below:

```
"datasets": {
  "cifar10_train": { # name of the first dataset parser
    "type": "torchvision.datasets.CIFAR10", # class to instantiate
    "params": { # parameters forwarded to the class constructor
      "root": "data/cifar/train",
      "train": true,
      "download": true
    },
    "task": { # task defined explicitly due to external type
      "type": "thelper.tasks.Classification",
      "params": { # by default, we just need to know the class names
        "class_names": [
          "airplane", "car", "bird", "cat", "deer",
          "dog", "frog", "horse", "ship", "truck"
        ],
        # torchvision loads samples as tuples; we map the indices
        "input_key": "0", # input = element at index#0 in tuple
        "label_key": "1" # label = element at index#1 in tuple
      }
    }
  },
  "cifar10_test": { # name of the second dataset parser
    "type": "torchvision.datasets.CIFAR10", # class to instantiate
    "params": { # parameters forwarded to the class constructor
      "root": "data/cifar/test",
      "train": false, # here, fetch test data instead of train data
      "download": true
    },
    "task": {
      # we use the same task info as above, both will be merged
      "type": "thelper.tasks.Classification",
      "params": {
        "class_names": [
          "airplane", "car", "bird", "cat", "deer",
          "dog", "frog", "horse", "ship", "truck"
        ],
        "input_key": "0",
        "label_key": "1"
      }
    }
  }
}
```

(continues on next page)

(continued from previous page)

```

    }
  }
}

```

The example above defines two dataset parsers, `cifar10_train` and `cifar10_test`, that can now be referred to in the `loaders` section of a configuration file (*described next*). For more information on the instantiation of dataset parsers, refer to `thelper.data.utils.create_parsers()`.

## 4.2.2 Loaders section

The `loaders` section of the configuration defines all data loader-related settings including split ratios, samplers, batch sizes, base transforms and augmentations, seeds, memory pinning, and async worker count. The first important concept to understand here is that multiple data parsers (*defined earlier*) can be combined or split into one or more data loaders. Moreover, there are exactly three data loaders defined for all experiments: the training data loader, the validation data loader, and the test data loader. For more information on the fundamental role of each loader, see [\[this link\]](#). In short, data loaders deal with parsers to load and transform data samples efficiently before packing them into batches that we can feed our models.

Some of the settings defined in this section apply to all three data loaders (e.g. memory pinning, base data transforms), while others can be specified for each loader individually (e.g. augmentations, batch size). The meta-settings that should always be set however are the split ratios that define the fraction of samples from each parser to use in a data loader. As shown in the example below, these ratios allow us to split a dataset into different loaders automatically, and without any possibility of data leakage between them. If all RNG seeds are set in this section, then the split will be reproducible between experiments. The split can also be precomputed using the `split` operation of the CLI ([click here for more information](#)).

Besides, base transformations defined in this section are used to ensure that all samples loaded by parsers are compatible with the input format expected by the model during training. For example, typical image classification pipelines expect images to have a resolution of 224x224 pixels, with each color channel normalized to either the `[-1, 1]` range, or using pre-computed mean and standard deviation values. We can define such operations directly using the classes available in the `thelper.transforms` module. This is also demonstrated in the example configuration below:

```

# note: this example is tied with the "datasets" example given earlier
"loaders": {
    "batch_size": 32,      # pack 32 images per minibatch (for all loaders)
    "test_seed": 0,        # fix the test set splitting seed
    "valid_seed": 0,       # fix the validation set splitting seed
    "torch_seed": 0,       # fix the PyTorch RNG seed for transforms/augments
    "numpy_seed": 0,       # fix the numpy RNG seed for transforms/augments
    "random_seed": 0,      # fix the random package RNG seed for transforms/augments
    # note: non-fixed seeds will be initialized randomly and printed in logs
    "workers": 4,          # each loader will be loading 4 minibatches in parallel
    "base_transforms": [   # defines the operations to apply to all loaded samples
        {
            # first, normalize 8-bit images to the [-1, 1] range
            "operation": "thelper.transforms.NormalizeMinMax",
            "params": {
                "min": [127, 127, 127],
                "max": [255, 255, 255]
            }
        },
        {
            # next, resize the CIFAR10 images to 224x224 for the model

```

(continues on next page)

(continued from previous page)

```

        "operation": "thelper.transforms.Resize",
        "params": {
            "dsize": [224, 224]
        }
    },
    {
        # finally, transform the opencv/numpy arrays to torch.Tensor arrays
        "operation": "torchvision.transforms.ToTensor"
    }
],
# we reserve 20% of the samples from the training parser for validation
"train_split": {
    "cifar10_train": 0.8
},
"valid_split": {
    "cifar10_train": 0.2
},
# we use 100% of the samples from the test parser for testing
"test_split": {
    "cifar10_test": 1.0
}
}

```

The example above prepares the CIFAR10 data using a 80%-20% training-validation split, and keeps all the original CIFAR10 testing data for actual testing. All loaded samples will be normalized and resized to fit the expected input resolution of a typical model, as shown in the next subsection. This example however contains no data augmentation pipelines; refer to the [\[relevant sections here\]](#) for actual usage examples. Similarly, no sampler is used above to rebalance the classes; [\[see here\]](#) for a use case. Finally, for more information on other parameters that are not discussed here, refer to the documentation of `thelper.data.utils.create_loaders()`.

### 4.2.3 Model section

The `model` section of the configuration defines the model that will be trained, fine-tuned, evaluated, or exported during the session. The model can be defined in several ways. If you are creating a new model from scratch (i.e. using randomly initialized weights), you simply have to specify the type of the class that implements the model's architecture along with its constructor's parameters. This is shown in the example below for an instance of MobileNet:

```

"model": {
    "type": "thelper.nn.mobilenet.MobileNetV2",
    "params": {
        "input_size": 224
    }
}

```

In this case, the constructor of `thelper.nn.mobilenet.MobileNetV2` will only receive a single argument, `input_size`, i.e. the size of the tensors it should expect as input. Some implementations of model architectures such as those in `torchvision.models` ([\[see them here\]](#)) might allow you to specify a pretrained parameter. Setting this parameter to `True` will let you automatically download the weights of that model and thus allow you to fine-tune it directly:

```

"model": {
    "type": "torchvision.models.resnet.resnet18",
    "params": {
        "pretrained": true
    }
}

```

(continues on next page)

(continued from previous page)

```
}
}
```

The second option to fine-tune a model that is not available via `torchvision` is to specify the path to a checkpoint produced by the framework as such:

```
"model": {
    "ckptdata" : "<PATH_TO_ANY_THELPER_CHECKPOINT.pth>"
}
```

When using this approach, the framework will first open the checkpoint and reinstantiate the model using its original fully qualified class name and the parameters originally passed to its constructor. Then, that model will be checked for task compatibility, and its weights will finally be loaded in. For more information on the checkpoints produced by the framework, see the [\[relevant section below\]](#). For more information on the model creation/loading process, refer to `thelper.nn.utils.create_model()`.

## 4.2.4 Trainer section

The `trainer` section of the configuration defines trainer, optimization, and metric-related settings used in a session. These settings include the type of trainer to use, the number of epochs to train for, the list of metrics to compute during training, the name of the metric to continuously monitor for improvements, the loss function to use, the optimizer, the scheduler, and the device (CUDA or CPU) that the session should be executed on.

First, note here that the type of trainer that is picked must be compatible with the task(s) exposed by the dataset parser(s) listed earlier in the configuration. If no trainer type is provided, the framework will automatically deduce which one to use for the current task. This deduction might fail for custom trainers/task combinations. If you are using a custom task, or if your model relies on multiple loss functions (or any other similar exotic thing), you might have to create your own trainer implementation derived from `thelper.train.base.Trainer`. Otherwise, see the trainers module (`thelper.train`) for a list of all available trainers.

All optimization settings are grouped into the `optimization` subsection of the `trainer` section. While specifying a scheduler is optional, an optimizer and a loss function must always be specified. The loss function can be provided via the typical type/params setup (as shown below), or obtained from the model via a getter function. For more information on the latter option, see `thelper.optim.utils.create_loss_fn()`. On the other hand, the nature of the optimizer and scheduler can only be specified via a type/param setup (as also shown below). The weights of the model specified in the last section will always be passed as the first argument of the optimizer's constructor at runtime. This behavior is compatible with all optimizers defined by PyTorch ([\[more info here\]](#)).

The `trainer` section finally contains another subsection titled `metrics`. This subsection defines a dictionary of named metrics that should be continuously updated during training, and evaluated at the end of each epoch. Numerous types of metrics are already implemented in `thelper.optim.metrics`, and many more will be added in the future. Metrics typically measure the performance of the model based on a specific criteria, but they can also do things like save model predictions and create graphs. A special “monitored” metric can also be defined in the `trainer` section, and it will be used to determine whether the model is improving or not during the training session. This is used to keep track of the “best” model weights while creating checkpoints, and it might also be used for scheduling.

A complete example of a trainer configuration is shown below:

```
"trainer": {
    # this example is in line with the earlier examples; we create a classifier
    "type": "thelper.train.ImageClassifTrainer", # type could be deduced_
    ↪ automatically
    "device": "cuda:all", # by default, run the session on all GPUs in parallel
    "epochs": 50, # run the session for a maximum of 50 epochs
```

(continues on next page)

(continued from previous page)

```

    "save_freq": 1,          # save the model in a checkpoint every epoch
    "monitor": "accuracy",  # monitor the 'accuracy' metric defined below for_
↪improvements
    "use_tbx": true,        # activate tensorboardX metric logging in output directory
    "optimization": {
        "loss": {
            "type": "torch.nn.CrossEntropyLoss",
            "params": {}     # empty sections like these can be omitted
        },
        "optimizer": {
            "type": "torch.optim.RMSprop",
            "params": {
                "lr": 0.01, # default learning rate used at the first epoch
                "weight_decay": 0.00004
            }
        },
        "scheduler": {
            # here, we create a fancy scheduler that will check a metric for its steps
            "type": "torch.optim.lr_scheduler.ReduceLROnPlateau",
            "params": {
                "mode": "max", # since we will monitor accuracy, we want to_
↪maximize it
                "factor": 0.1, # when a plateau is detected, decrease lr by 90%
                "patience": 3  # wait three epochs with no improvement before_
↪stepping
            },
            # now, we just name the metric defined below for the scheduler to use
            "step_metric": "accuracy"
        }
    },
    "metrics": { # this is the list of all metrics we will be evaluating
        "accuracy": { # the name of each metric should be unique
            "type": "thelper.optim.Accuracy",
            "params": {
                "top_k": 1
            }
        },
        "confmat": {
            # this is a special consumer used to create confusion matrices
            # (we can't monitor this one, as it is not an actual "metric")
            "type": "thelper.train.ConfusionMatrix"
        }
    },
    "test_metrics": { # metrics in this section will only be used for testing
        "logger": {
            # (can't monitor this one either, as it is not an actual "metric")
            "type": "thelper.train.ClassifLogger",
            "params": {
                "top_k": 3
            }
        }
    }
}

```

For more information on the metrics available in the framework, see [thelper.optim.metrics](#).

### 4.2.5 Annotator section

The `annotator` section of the configuration is used solely to define GUI-related settings during annotation sessions. For now, it should only contain the type and constructor parameters of the GUI tool that will be instantiated to create the annotations. An example is shown below:

```
"annotator": {
  "type": "thelper.gui.ImageSegmentAnnotator", # type of annotator to instantiate
  "params": {
    "sample_input_key": "image", # this key is tied to the data parser's output
    "labels": [
      # for this example, we only use one brush type that draws using solid red
      {"id": 255, "name": "foreground", "color": [0, 0, 255]}
    ]
  }
}
```

In this case, an image segmentation GUI is created that will allow the “image” loaded in each sample to be annotated by user with a brush tool. This section (as well as all GUI tools) are still experimental. For more information on annotators, refer to `thelper.gui.annotators`.

### 4.2.6 Global parameters

Finally, session configurations can also contain global parameters located outside the main sections detailed so far. For example, the session name is a global flag which is often mandatory as it is used to identify the session and create its output directory. Other global parameters are used to change the behavior of imported package, or are just hacky solutions to problems that should be fixed otherwise.

For now, the global parameters considered “of interest” are the following:

- `name` : specifies the name of the session (mandatory in most operation modes).
- `cudnn_benchmark` : specifies whether to activate/deactivate cuDNN benchmarking mode.
- `cudnn_deterministic` : specifies whether to activate/deactivate cuDNN deterministic mode.

Future global parameters will most likely be handled via `thelper.utils.setup_globals()`.

[\[to top\]](#)

## 4.3 Session Directories

If the framework is used in a way that requires it to produce outputs, they will always be located somewhere in the “session directory”. This directory is created in the root output directory provided to the CLI (also often called the “save” directory), and it is named after the session itself. The session directory contains three main folders that hold checkpoints, logs, and outputs. These are discussed in the following subsections. The general structure of a session directory is shown below:

```
<session_directory_name>
|
|-- checkpoints
|   |-- ckpt.0000.<platform>--<date>--<time>.pth
|   |-- ckpt.0001.<platform>--<date>--<time>.pth
|   |-- ckpt.0002.<platform>--<date>--<time>.pth
```

(continues on next page)

(continued from previous page)

```

|      |-- ...
|      \-- ckpt.best.pth
|
|-- logs
|      |-- <dataset1-name>.log
|      |-- <dataset2-name>.log
|      |-- ...
|      |-- config.<platform>--<date>--<time>.json
|      |-- data.log
|      |-- modules.log
|      |-- packages.log
|      |-- task.log
|      \-- trainer.log
|
|-- output
|      \-- <session_directory_name>
|          |-- train-<platform>--<date>--<time>
|              |-- events.out.tfevents.<something>.<platform>
|              \-- ...
|          |-- valid-<platform>--<date>--<time>
|              |-- events.out.tfevents.<something>.<platform>
|              \-- ...
|          |-- test-<platform>--<date>--<time>
|              |-- events.out.tfevents.<something>.<platform>
|              \-- ...
|          \-- ...
|
\-- config.latest.json

```

### 4.3.1 Checkpoints

The `checkpoints` folder contains the binary files pickled by PyTorch that store all training data required to resume a session. These files are automatically saved at the end of each epoch during a training session. The checkpoints are named using the `ckpt.XXXX.YYYYYY-ZZZZZZ-ZZZZZZ.pth` convention, where `XXXX` is the epoch index (0-based), `YYYYYY` is the platform or hostname, and `ZZZZZZ-ZZZZZZ` defines the date and time of their creation (in `YYYYMMDD-HHMMSS` format). All checkpoints created by the framework during training will use this naming convention except for the `best` checkpoint that might be created in monitored training sessions (as part of early stopping and for final model evaluation). In this case, it will simply be named `ckpt.best.pth`. Its content is the same as other checkpoints however, and it is actually just a copy of the corresponding “best” checkpoint in the same directory.

Checkpoints can be opened directly using `torch.load()`. They contain a dictionary with the following fields:

- `name` : the name of the session. Used as a unique identifier for many types of output.
- `epoch` : the epoch index (0-based) at the end of which the checkpoint was saved. This value is optional, and may only be saved in training sessions.
- `iter` : the total number of iterations computed so far in the training session. This value is optional, and may only be saved in training sessions.
- `source` : the name of the host that created the checkpoint and its time of creation.
- `sha1` : the sha1 signature of the framework’s latest git commit. Used for debugging purposed only.



- `version` : the version of the framework that created this checkpoint. Will be used for data and configuration file migration if necessary when reloading the checkpoint.
- `task` : a copy or string representation of the task the model was being trained for. Used to keep track of expected model input/output mappings (e.g. class names).
- `outputs` : the outputs (e.g. metrics) generated by the trainer for all epochs thus far. This object is optional, and may only be saved in training sessions.
- `model` : the weights (or “state dictionary”) of the model, or a path to where these weights may be found. This field can be used to hold a link to an external JIT trace or ONNX version of the model.
- `model_type` : the type (or class name) of the model that may be used to reinstantiate it.
- `model_params` : the constructor parameters of the model that may be used to reinstantiate it.
- `optimizer` : the state of the optimizer at the end of the latest training epoch. This value is optional, and may only be saved in training sessions.
- `scheduler` : the state of the scheduler at the end of the latest training epoch. This value is optional, and may only be saved in training sessions.
- `monitor_best` : the “best” value for the metric being monitored so far. This value is optional, and may only be saved in training sessions.
- `config` : the full session configuration dictionary originally passed to the CLI endpoint.

By default, these fields do not contain pickled objects directly tied to the framework, meaning any PyTorch installation should be able to open a checkpoint without crashing. This also means that a model trained with this framework can be opened and reused in any other framework, as long as you are willing to extract its weights from the checkpoint yourself. An example of this procedure is given [\[here\]](#).

Experimental support for checkpoint creation outside a training session is available through the CLI’s `export` operation. *See the section above for more information.*

### 4.3.2 Session logs

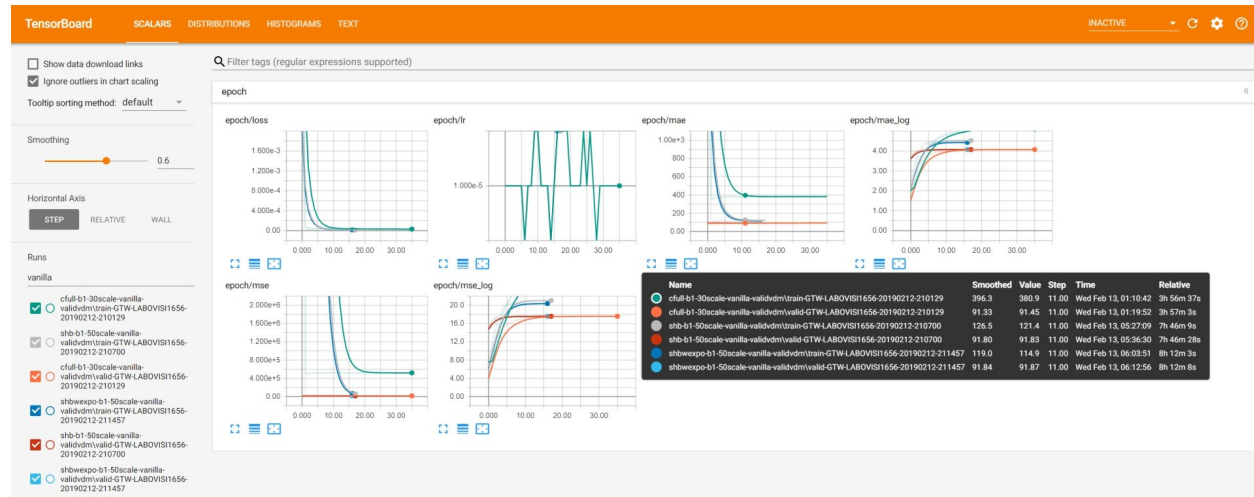
All information printed to the terminal during a session will also automatically be printed to files located in the `logs` folder of the session directory. Moreover, useful information about the training environment and datasets will be printed in other files in the same location. A brief description of these files is provided below:

- `<dataset_name>.log` : contains metadata (in JSON format) of the named dataset, its loaded sample count, and the separation of its sample indices across the train/valid/test sets. Can be used to validate the data split and keep track of which sample is used in which set.
- `config.<PLATFORM>-<DATE>-<TIME>.json` : backup of the (JSON) configuration file of the session that created or modified the current session directory.
- `data.log` : logger output that provides high-level information about the loaded dataset parsers including given names, sizes, task interfaces, and base transforms.
- `modules.log` : logger output that provides details regarding the instantiated model type (class name), the parameters passed to its constructor, and a full list of its layers once constructed.
- `packages.log` : lists all packages installed in the runtime environment as well as their version.
- `task.log` : provides the full string representation of the task object used during the session.
- `trainer.log` : logger output that details the training progress during the session. This file can become very large for long sessions; and might be rotated past a certain size in the future.

Specialized CLI operations and trainers as well as custom implementations might create additional logs in this directory. In all cases, logs are provided as nice-to-have for debugging purposes only, and their content/structure might change in future versions.

### 4.3.3 Outputs

Finally, the session directory contains an `output` folder that is used to store all the evaluation results generated by the metrics as well as the `tensorboard` event files. The first level of the output directory is named after the session itself so that it may easily be copied elsewhere without creating conflicts. This also allows `tensorboard` to display the session name in its UI. That folder then contains the training, validation, and testing outputs generated for each session. These outputs are separated so that individual curves can be turned on and off in `tensorboard`. A typical output directory loaded in `tensorboard` is shown below.



In this example, the training and validation outputs of several sessions are combined. The metrics of each session that produced scalar values were used to generate plots. The scalars are evaluated once every epoch, and are grouped automatically in a section named `epoch`. The loss and learning rates are also automatically plotted in this section. Additional tabs holding model weight histograms and text outputs are also available. If a metric had been used that generated images, those would also be available in another tab.

For more information on available metrics, see [thelper.optim.metrics](#). For more information about `tensorboard`, visit [\[the official site\]](#).

[\[to top\]](#)

## USE CASES

This section is still under construction. Some example configuration files are available in the `config` directory of the repository ([\[see them here\]](#)). For high-level information on generic parts of the framework, refer to the [\[user guide\]](#).

### 5.1 Image classification

Building an image classifier is probably the simplest thing you can do with the framework. Here, we provide an in-depth look at a JSON configuration used to build a 10-class object classification model based on the CIFAR-10 dataset.

As usual, we must define four different fields in our configuration for everything to work: the datasets, the data loaders, the model, and the trainer. First off, the dataset. As mentionned before, here we will work with the CIFAR-10 dataset provided by `torchvision`. We do so simply to have a configuration that can run “off-the-shelf”, but in reality, you will probably be using your own data. To learn how to create your own dataset interface and load your own data, see [\[this section\]](#). So, for our simple task, we define the `datasets` field as follows:

```
"datasets": {
  "cifar10": {
    "type": "torchvision.datasets.CIFAR10",
    "params": {"root": "data/cifar/train", "download": true},
    "task": {
      "type": "thelper.tasks.Classification",
      "params": {
        "class_names": [
          "airplane", "car", "bird", "cat", "deer",
          "dog", "frog", "horse", "ship", "truck"
        ],
        "input_key": "0", "label_key": "1"
      }
    }
  }
}
```

In summary, we will instantiate a single dataset named “cifar10” based on the `torchvision.datasets.CIFAR10` class. The constructor of that class will receive two arguments, namely the (relative) path where to save the data, and a boolean indicating that the dataset should be downloaded if not already available. More info on these two arguments can actually be found in the [\[PyTorch documentation\]](#). Here, since this dataset interface does not explicitly define a “task”, we need to provide one ourselves. Therefore, we add a “task” field in which we specify that the task type is related to classification, and provide the task’s construction arguments. In this case, this is the list of class names that correspond to the indices that the dataset interface will be associating to each loaded sample, and the key strings. Note that since `torchvision.datasets.CIFAR10` will be loading samples as tuples, so the key strings merely correspond to tuple indices.

Next, we will define how this cute little dataset will be used to load training and/or evaluation data. This is accomplished through the `loaders` fields as follows:

```
"loaders": {
  "train_split": {"cifar10": 0.9},
  "valid_split": {"cifar10": 0.1},
  "batch_size": 32,
  "base_transforms": [
    {
      "operation": "thelper.transforms.NormalizeMinMax",
      "params": {
        "min": [127, 127, 127], "max": [255, 255, 255]
      }
    },
    {
      "operation": "thelper.transforms.Resize",
      "params": {"dsize": [224, 224]}
    },
    {
      "operation": "torchvision.transforms.ToTensor"
    }
  ]
}
```

In this example, we ask the training and validation data loaders to split the “cifar10” dataset we defined above using a 90%-10% ratio. By default, this split will be randomized, but repeatable across experiments. Seeds used to shuffle the data samples will be printed in the logs, and can even be set manually in this section if needed. Next, we set the batch size for all data loaders to 32, and define base transforms to apply to all samples. In this case, the transforms will normalize each 8-bit image by min-maxing it, resize it 224x224 (as required by our model), and finally transform it into a tensor. Many transformation operations can be defined in the `loaders` section, and it may be interesting to visualize the output before trying to feed it to a model. For more information on how to do so, see [\[this use case\]](#).

Next, we will define the model architecture to train. Again, you might want to use your own architecture here. If so, refer to the [\[relevant use case\]](#). Here, we keep things extremely simple and rely on the pre-existing ResNet implementation located inside the framework:

```
"model": {
  "type": "thelper.nn.resnet.ResNet"
}
```

Since this class can be instantiated as-is without having to provide any arguments (it defaults to ResNet-34), we do not even need to specify a “params” field. Once created, this model will be adapted to our classification problem by providing it the “task” object we defined in our dataset interface. This means that its default 1000-class output layer will be updated to produce 10 outputs (since we have 10 classes).

Finally, we define the `trainer` field to provide all remaining training parameters:

```
"trainer": {
  "epochs": 5,
  "monitor": "accuracy",
  "optimization": {
    "loss": {"type": "torch.nn.CrossEntropyLoss"},
    "optimizer": {"type": "torch.optim.Adam"}
  },
  "metrics": {
    "accuracy": {"type": "thelper.optim.Accuracy"}
  }
}
```

(continues on next page)

(continued from previous page)

```
}
}
```

Here, we limit the training to 5 epochs. The loss is a traditional cross-entropy, and we use Adam to update model weights via backprop. The loss function and optimizer could both receive extra parameters (using a “params” field once more), but we keep the defaults everywhere. Finally, we define a single metric to be evaluated during training (accuracy), and set it as the “monitoring” metric to use for early stopping.

The complete configuration is shown below:

```
{
  "name": "classif-cifar10",
  "datasets": {
    "cifar10": {
      "type": "torchvision.datasets.CIFAR10",
      "params": {"root": "data/cifar/train", "download": true},
      "task": {
        "type": "thelper.tasks.Classification",
        "params": {
          "class_names": [
            "airplane", "car", "bird", "cat", "deer",
            "dog", "frog", "horse", "ship", "truck"
          ],
          "input_key": "0", "label_key": "1"
        }
      }
    }
  },
  "loaders": {
    "batch_size": 32,
    "base_transforms": [
      {
        "operation": "thelper.transforms.NormalizeMinMax",
        "params": {
          "min": [127, 127, 127], "max": [255, 255, 255]
        }
      },
      {
        "operation": "thelper.transforms.Resize",
        "params": {"dsize": [224, 224]}
      },
      {
        "operation": "torchvision.transforms.ToTensor"
      }
    ],
    "train_split": {"cifar10": 0.9},
    "valid_split": {"cifar10": 0.1}
  },
  "model": {"type": "thelper.nn.resnet.ResNet"},
  "trainer": {
    "epochs": 5,
    "monitor": "accuracy",
    "optimization": {
      "loss": {"type": "torch.nn.CrossEntropyLoss"},
      "optimizer": {"type": "torch.optim.Adam"}
    },
    "metrics": {
```

(continues on next page)

(continued from previous page)

```
        "accuracy": {"type": "thelper.optim.Accuracy"}
    }
}
```

Once saved to a json file, we will be able to launch the training session via:

```
$ thelper new <PATH_TO_CLASSIF_CIFAR10_CONFIG>.json <PATH_TO_OUTPUT_DIR>
```

The dataset will first be downloaded, split, and passed to data loaders. Then, the model will be instantiated, and all objects will be given to the trainer to start the session. Right away, some log files will be created in a new folder named “classif-cifar10” in the directory provided as the second argument on the command line. Once the training is complete, that folder will contain the model checkpoints as well as the final evaluation results.

## 5.2 Image segmentation

Section statement here @ @ @ @ @

## 5.3 Object Detection

Section statement here @ @ @ @ @

## 5.4 Super-resolution

Section statement here @ @ @ @ @

## 5.5 Creating a new dataset interface

Section statement here @ @ @ @ @

## 5.6 Dataset/Loader visualization

Section statement here @ @ @ @ @

## 5.7 Dataset annotation

Section statement here @ @ @ @ @

## 5.8 Rebalancing a dataset

Section statement here @ @ @ @ @

## 5.9 Exporting a dataset

Section statement here @ @ @ @ @

## 5.10 Defining a data augmentation pipeline

Section statement here @ @ @ @ @

## 5.11 Supporting a custom trainer

Section statement here @ @ @ @ @

## 5.12 Supporting a custom task

Section statement here @ @ @ @ @

## 5.13 Supporting a custom model

Section statement here @ @ @ @ @

## 5.14 Visualizing metrics using `tensorboardX`

Section statement here @ @ @ @ @

## 5.15 Manually reloading a model

Section statement here @ @ @ @ @

## 5.16 Exporting a model

Once you have trained a model (using the framework or otherwise), you might want to share it with others. Models are typically exported in two parts: architecture and weights. However, metadata related to the task the model was built for would be missing with only those two components. Here, we show a solution for exporting a classification model trained using the framework under ONNX, TraceScript, or pickle format along with its corresponding index-to-class-name mapping. Further down, we also give tips on similarly exporting a model trained in another framework.

The advantage of ONNX and TraceScript exports is that whoever reloads your model does not need to have the class that you used to define the model's architecture at hand. However, this approach might make fine-tuning or retraining your model more complicated (you should consider it a 'read-only' export).

Models/checkpoints exported this way can be easily reloaded using the framework, and may also be opened manually by others to extract only the information they need.

So, first off, let's start by training a classification model using the following configuration:

```
{
  "name": "classif-cifar10",
  "datasets": {
    "cifar10": {
      "type": "torchvision.datasets.CIFAR10",
      "params": {"root": "data/cifar/train"},
      "task": {
        "type": "thelper.tasks.Classification",
        "params": {
          "class_names": [
            "airplane", "car", "bird", "cat", "deer",
            "dog", "frog", "horse", "ship", "truck"
          ],
          "input_key": "0", "label_key": "1"
        }
      }
    }
  },
  "loaders": {
    "batch_size": 32,
    "base_transforms": [
      {
        "operation": "thelper.transforms.NormalizeMinMax",
        "params": {
          "min": [127, 127, 127], "max": [255, 255, 255]
        }
      },
      {
        "operation": "thelper.transforms.Resize",
        "params": {"dsize": [224, 224]}
      },
      {
        "operation": "torchvision.transforms.ToTensor"
      }
    ],
    "train_split": {"cifar10": 0.9},
    "valid_split": {"cifar10": 0.1}
  },
  "model": {"type": "thelper.nn.resnet.ResNet"},
  "trainer": {
    "epochs": 5,
```

(continues on next page)



(continued from previous page)

```

    "monitor": "accuracy",
    "optimization": {
        "loss": {"type": "torch.nn.CrossEntropyLoss"},
        "optimizer": {"type": "torch.optim.Adam"}
    },
    "metrics": {
        "accuracy": {"type": "thelper.optim.Accuracy"}
    }
}

```

The above configuration essentially means that we will be training a ResNet model with default settings on CIFAR10 using all 10 classes. You can launch the training process via:

```
$ thelper new <PATH_TO_CLASSIF_CIFAR10_CONFIG>.json <PATH_TO_OUTPUT_DIR>
```

See the [\[user guide\]](#) for more information on creating training sessions. Once that's done, you should obtain a folder named `classif-cifar10` in your output directory that contains training logs as well as checkpoints. To export this model in a new checkpoint, we will use the following session configuration:

```

{
    "name": "export-classif-cifar10",
    "model": {
        "ckptdata": "<PATH_TO_OUTPUT_DIR>/classif-cifar10/checkpoints/ckpt.best.pth"
    },
    "export": {
        "ckpt_name": "test-export.pth",
        "trace_name": "test-export.zip",
        "save_raw": true,
        "trace_input": "torch.rand(1, 3, 224, 224)"
    }
}

```

This configuration essentially specifies where to find the ‘best’ checkpoint for the model we just trained, and how to export a trace of it. For more information on the export operation, refer to [\[the user guide\]](#). We now provide the configuration as a JSON to the CLI one more:

```
$ thelper export <PATH_TO_EXPORT_CONFIG>.json <PATH_TO_OUTPUT_DIR>
```

If everything goes well, `<PATH_TO_OUTPUT_DIR>/export-classif-cifar10` should now contain a checkpoint with the exported model trace and all metadata required to reinstantiate it. Note that as of 2019/06, PyTorch exports model traces as zip files, meaning you will have to copy two files from the output session folder. In this case, that would be `test-export.pth` and `test-export.zip`.

Finally, note that if you are attempting to export a model that was trained outside the framework, you will have to specify which task this model was trained for as well as the type of the model to instantiate and possibly the path to its weights in the `model` field of the configuration above. An example configuration is given below:

```

{
    "name": "export-classif-custom",
    "model": {
        "type": "fully.qualified.name.to.model",
        "params": {
            # here, provide all model constructor parameters
        },
        "weights": "path_to_model_state_dictionary.pth"
    }
}

```

(continues on next page)

(continued from previous page)

```
    },
    "export": {
        "ckpt_name": "test-export.pth",
        "trace_name": "test-export.zip",
        "save_raw": true,
        "trace_input": "torch.rand(1, 3, 224, 224)"
    }
}
```

For more information on model importation, refer to the documentation of `thelper.nn.utils.create_model()`.

[\[to top\]](#)

## HELPER PACKAGE

Top-level package for the ‘helper’ framework.

Running `import helper` will recursively import all important subpackages and modules.

### 6.1 Subpackages

#### 6.1.1 `helper.data` package

Dataset parsing/loading package.

This package contains classes and functions whose role is to fetch the data required to train, validate, and test a model. The `helper.data.utils.create_loaders()` function contained herein is responsible for preparing the task and data loaders for a training session. This package also contains the base interfaces for dataset parsers.

#### Subpackages

##### `helper.data.geo` package

Geospatial dataset parsing/loading package.

This package contains classes and functions whose role is to fetch the data required to train, validate, and test a model on geospatial data. Importing the modules inside this package requires GDAL.

#### Submodules

##### `helper.data.geo.agravis` module

Agricultural Semantic Segmentation Challenge Dataset Interface

Original author: David Landry ([david.landry@crim.ca](mailto:david.landry@crim.ca)) Updated by Pierre-Luc St-Charles (April 2020)

## thelper.data.geo.bigearthnet module

## thelper.data.geo.gdl module

Data parsers & utilities for cross-framework compatibility with Geo Deep Learning (GDL).

Geo Deep Learning (GDL) is a machine learning framework initiative for geospatial projects lead by the wonderful folks at NRCan's CCMEQ. See <https://github.com/NRCan/geo-deep-learning> for more information.

The classes and functions defined here were used for the exploration of research topics and for the validation and testing of new software components.

## thelper.data.geo.infer module

```
class thelper.data.geo.infer.SlidingWindowTester(session_name, session_dir, model,  
                                                task, loaders, config, ckpt-  
                                                data=None)
```

Bases: *thelper.infer.base.Tester*

Tester that satisfies the requirements of the `Tester` in order to run classification inference

```
__init__(session_name, session_dir, model, task, loaders, config, ckptdata=None)
```

Receives the trainer configuration dictionary, parses it, and sets up the session.

```
eval_epoch(model, epoch, dev, loader, metrics, output_path)
```

Computes the pixelwise prediction on an image.

It does the prediction per batch size of `N` pixels. It returns the class predicted and its probability. The results are saved into two images created with the same size and projection info as the input rasters.

The `class` image gives the class id, a number between 1 and the number of classes for corresponding pixels. Class id 0 is reserved for `nodata`.

The `probs` image contains `N`-class channels with the probability values of the pixels for each class. The probabilities by default are normalised.

Also, a `config-classes.json` file is created listing the `name-to-class-id` mapping that was used to generate the values in the `class` image (i.e.: class names defined by the pre-trained model).

### Parameters

- **model** – the model with which to run inference that is already uploaded to the target device(s).
- **epoch** – the epoch index we are training for (0-based, and should normally only be 0 for single test pass).
- **dev** – the target device that tensors should be uploaded to (corresponding to model's device(s)).
- **loader** – the data loader used to get transformed test samples.
- **metrics** – the dictionary of metrics/consumers to report inference results (mostly loggers and basic report generator in this case since there shouldn't be ground truth labels to validate against).
- **output\_path** – directory where output files should be written, if necessary.

```
supports_classification = True
```

## thelper.data.geo.ogc module

Data parsers & utilities module for OGC-related projects.

**class** `thelper.data.geo.ogc.TB15D104`

Bases: `object`

Wrapper class for OGC Testbed-15 (D104) identifiers.

**BACKGROUND\_ID** = 0

**LAKE\_ID** = 1

**TYPECE\_LAKE** = '21'

**TYPECE\_RIVER** = '10'

**class** `thelper.data.geo.ogc.TB15D104DetectLogger` (*conf\_threshold=0.5*)

Bases: `thelper.train.utils.DetectLogger`

**\_\_init\_\_** (*conf\_threshold=0.5*)

Receives the logging parameters & the optional class label names used to decorate the log.

**report\_geojson** ()

`thelper.data.geo.ogc.postproc_features` (*input\_file*, *bboxes\_srs*, *orig\_geoms\_path*,  
*output\_file*, *final\_srs=None*,  
*write\_shapefile\_copy=False*)

Post-processes bounding box detections produced during an evaluation session into a GeoJSON file.

## thelper.data.geo.parsers module

Geospatial data parser & utilities module.

## thelper.data.geo.utils module

`thelper.data.geo.utils.export_geojson_with_crs` (*features*, *srs\_target*)

Exports a list of features along with their SRS into a GeoJSON-compatible string.

`thelper.data.geo.utils.export_geotiff` (*filepath*, *crop*, *srs*, *geotransform*)

`thelper.data.geo.utils.get_feature_bbox` (*geom*, *offsets=None*)

`thelper.data.geo.utils.get_feature_roi` (*geom*, *px\_size*, *skew*, *roi\_buffer=None*,  
*crop\_img\_size=None*, *crop\_real\_size=None*)

`thelper.data.geo.utils.get_geocoord` (*geotransform*, *x*, *y*)

`thelper.data.geo.utils.get_geoextent` (*geotransform*, *x*, *y*, *cols*, *rows*)

`thelper.data.geo.utils.get_pxcoord` (*geotransform*, *x*, *y*)

`thelper.data.geo.utils.open_rasterfile` (*raster\_data*, *keep\_rasters\_open=False*)

`thelper.data.geo.utils.parse_geojson` (*geojson*, *srs\_target=None*, *roi=None*, *al-*  
*low\_outlying=False*, *clip\_outlying=False*)

`thelper.data.geo.utils.parse_geojson_crs` (*body*)

Imports a coordinate reference system (CRS) from a GeoJSON tree.

```
thelper.data.geo.utils.parse_raster_metadata (raster_metadata, raster_dataset=None)
```

Parses the provided raster metadata and updates it by adding extra details required for later use.

The provided raster metadata is updated directly. Metadata is validated against the matching data storage. If any important, required or requested (bands) metadata is missing, the function raises the issue immediately.

#### Parameters

- **raster\_metadata** (*dict*) – raster metadata dictionary with minimally a file ‘path’ and list of ‘bands’ indices to process.
- **raster\_dataset** (*gdal.Dataset*) – (optional) preloaded dataset object corresponding to the raster metadata.

#### Raises

- **ValueError** – at least one input raster was missing a required metadata parameter or a parameter is erroneous.
- **IOError** – the raster path could not be found or reading it did not generate a valid raster using GDAL.

```
thelper.data.geo.utils.parse_rasters (raster_paths, srs_target=None, reproj=False)
```

```
thelper.data.geo.utils.parse_roi (roi_path, srs_target=None)
```

```
thelper.data.geo.utils.parse_shapefile (shapefile_path, srs_target=None, roi=None,  
                                         allow_outlying=False, clip_outlying=False,  
                                         layer_id=0)
```

```
thelper.data.geo.utils.reproject_coords (coords, src_srs, tgt_srs)
```

```
thelper.data.geo.utils.reproject_crop (raster, crop_raster, crop_size, crop_datatype,  
                                         crop_nodataval=None, reproj_opt=None,  
                                         fill_nodata=False)
```

## Submodules

### thelper.data.loaders module

Dataset loaders module.

This module contains a dataset loader specialization used to properly seed samplers and workers.

```
class thelper.data.loaders.LoaderFactory (config)
```

Bases: `object`

Factory used for preparing and splitting dataset parsers into usable data loader objects.

This class is responsible for parsing the parameters contained in the ‘loaders’ field of a configuration dictionary, instantiating the data loaders, and shuffling/splitting the samples. An example configuration is presented in `thelper.data.utils.create_loaders()`.

**See also:**

```
thelper.data.utils.create_loaders()  
thelper.transforms.utils.load_augments()  
thelper.transforms.utils.load_transforms()
```

**\_\_init\_\_** (*config*)

Receives and parses the data configuration dictionary.

**create\_loaders** (*datasets, train\_idx, valid\_idx, test\_idx*)

Returns the data loaders for the train/valid/test sets based on a prior split.

This function essentially takes the dataset parser interfaces and indices maps, and instantiates data loaders that are ready to produce samples for training or evaluation. Note that the dataset parsers will be deep-copied in each data loader, meaning that they should ideally not contain a persistent loading state or a large buffer.

#### Parameters

- **datasets** – the map of dataset parsers, where each has a name (key) and a parser (value).
- **train\_idx** – training data samples indices map.
- **valid\_idx** – validation data samples indices map.
- **test\_idx** – test data samples indices map.

**Returns** A three-element tuple containing the training, validation, and test data loaders, respectively.

**get\_base\_transforms** ()

Returns the (global) sample transformation operations parsed in the data configuration.

**get\_split** (*datasets, task*)

Returns the train/valid/test sample indices split for a given dataset (name-parser) map.

Note that the returned indices are unique, possibly shuffled, and never duplicated between sets. If the samples have a class attribute (i.e. the task is related to classification), the split will respect the initial distribution and apply the ratios within the classes themselves. For example, consider a dataset of three classes (*A*, *B*, and *C*) that contains 100 samples such as:

$$|A| = 50, |B| = 30, |C| = 20$$

If we require a 80%-10%-10% ratio distribution for the training, validation, and test loaders respectively, the resulting split will contain the following sample counts:

$$\text{training loader} = 40A + 24B + 16C$$

$$\text{validation loader} = 5A + 3B + 2C$$

$$\text{test loader} = 5A + 3B + 2C$$

In the case of multi-label classification datasets, there is no guarantee that the classes will be balanced across the training/validation/test sets. Instead, for a given class list, the classes with fewer samples will be split first.

#### Parameters

- **datasets** – the map of datasets to split, where each has a name (key) and a parser (value).
- **task** – a task object that should be compatible with all provided datasets (can be `None`).

**Returns** A three-element tuple containing the maps of the training, validation, and test sets respectively. These maps associate dataset names to a list of sample indices.

`thelper.data.loaders.default_collate` (*batch, force\_tensor=True*)

Puts each data field into a tensor with outer dimension batch size.

This function is copied from PyTorch's `torch.utils.data._utils.collate.default_collate`, but additionally supports custom objects from the framework (such as bounding boxes). These will not be converted to tensors, and it will be up to the trainer to handle them accordingly.

See `torch.utils.data.DataLoader` for more information.

### thelper.data.parsers module

Dataset parsers module.

This module contains dataset parser interfaces and base classes that define basic i/o operations so that the framework can automatically interact with training data.

### thelper.data.pascalvoc module

PASCAL VOC dataset parser module.

This module contains a dataset parser used to load the PASCAL Visual Object Classes (VOC) dataset for semantic segmentation or object detection. See <http://host.robots.ox.ac.uk/pascal/VOC/> for more info.

### thelper.data.samplers module

Samplers module.

This module contains classes used for raw dataset rebalancing or augmentation.

All samplers here should aim to be compatible with PyTorch's sampling interface (`torch.utils.data.sampler.Sampler`) so that they can be instantiated at runtime through a configuration file and used as the input of a data loader.

### thelper.data.utils module

Dataset utility functions and tools.

This module contains utility functions and tools used to instantiate data loaders and parsers.

`thelper.data.utils.create_hdf5` (*archive\_path*, *task*, *train\_loader*, *valid\_loader*, *test\_loader*, *compression=None*, *config\_backup=None*)

Saves the samples loaded from train/valid/test data loaders into an HDF5 archive.

The loaded minibatches are decomposed into individual samples. The keys provided via the task interface are used to fetch elements (input, groundtruth, ...) from the samples, and save them in the archive. The archive will contain three groups (*train*, *valid*, and *test*), and each group will contain a dataset for each element originally found in the samples.

Note that the compression operates at the sample level, not at the dataset level. This means that elements of each sample will be compressed individually, not as an array. Therefore, if you are trying to compress very correlated samples (e.g. frames in a video sequence), this approach will be pretty bad.

#### Parameters

- **archive\_path** – path pointing where the HDF5 archive should be created.
- **task** – task object that defines the input, groundtruth, and meta keys tied to elements that should be parsed from loaded samples and saved in the HDF5 archive.
- **train\_loader** – training data loader (can be *None*).
- **valid\_loader** – validation data loader (can be *None*).



- **test\_loader** – testing data loader (can be *None*).
- **compression** – the compression configuration dictionary that will be parsed to determine how sample elements should be compressed. If a mapping is missing, that element will not be compressed.
- **config\_backup** – optional session configuration file that should be saved in the HDF5 archive.

Example compression configuration:

```
# the config is given as a dictionary
{
    # each field is a key that corresponds to an element in each sample
    "key1": {
        # the 'type' identifies the compression approach to use
        # (see thelper.utils.encode_data for more information)
        "type": "jpg",
        # extra parameters might be needed to encode the data
        # (see thelper.utils.encode_data for more information)
        "encode_params": {},
        # these parameters are packed and kept for decoding
        # (see thelper.utils.decode_data for more information)
        "decode_params": {"flags": "cv.IMREAD_COLOR"}
    },
    "key2": {
        # this explicitly means that no encoding should be performed
        "type": "none"
    },
    ...
    # if a key is missing, its elements will not be compressed
}
```

See also:

```
thelper.cli.split_data()
thelper.data.parsers.HDF5Dataset
thelper.utils.encode_data()
thelper.utils.decode_data()
```

`thelper.data.utils.create_loaders` (*config*, *save\_dir=None*)

Prepares the task and data loaders for a model trainer based on a provided data configuration.

This function will parse a configuration dictionary and extract all the information required to instantiate the requested dataset parsers. Then, combining the task metadata of all these parsers, it will evenly split the available samples into three sets (training, validation, test) to be handled by different data loaders. These will finally be returned along with the (global) task object.

The configuration dictionary is expected to contain two fields: `loaders`, which specifies all parameters required for establishing the dataset split, shuffling seeds, and batch size (these are listed and detailed below); and `datasets`, which lists the dataset parser interfaces to instantiate as well as their parameters. For more information on the `datasets` field, refer to `thelper.data.utils.create_parsers()`.

The parameters expected in the ‘loaders’ configuration field are the following:

- `<train_/valid_/test_>batch_size` (mandatory): specifies the (mini)batch size to use in data loaders. If you get an ‘out of memory’ error at runtime, try reducing it.

- `<train_/valid_/test_>collate_fn` (optional): specifies the collate function to use in data loaders. The default one is typically fine, but some datasets might require a custom function.
- `shuffle` (optional, default=True): specifies whether the data loaders should shuffle their samples or not.
- `test_seed` (optional): specifies the RNG seed to use when splitting test data. If no seed is specified, the RNG will be initialized with a device-specific or time-related seed.
- `valid_seed` (optional): specifies the RNG seed to use when splitting validation data. If no seed is specified, the RNG will be initialized with a device-specific or time-related seed.
- `torch_seed` (optional): specifies the RNG seed to use for torch-related stochastic operations (e.g. for data augmentation). If no seed is specified, the RNG will be initialized with a device-specific or time-related seed.
- `numpy_seed` (optional): specifies the RNG seed to use for numpy-related stochastic operations (e.g. for data augmentation). If no seed is specified, the RNG will be initialized with a device-specific or time-related seed.
- `random_seed` (optional): specifies the RNG seed to use for stochastic operations with python's 'random' package. If no seed is specified, the RNG will be initialized with a device-specific or time-related seed.
- `workers` (optional, default=1): specifies the number of threads to use to preload batches in parallel; can be 0 (loading will be on main thread), or an integer  $\geq 1$ .
- `pin_memory` (optional, default=False): specifies whether the data loaders will copy tensors into CUDA-pinned memory before returning them.
- `drop_last` (optional, default=False): specifies whether to drop the last incomplete batch or not if the dataset size is not a multiple of the batch size.
- `sampler` (optional): specifies a type of sampler and its constructor parameters to be used in the data loaders. This can be used for example to help rebalance a dataset based on its class distribution. See [thelper.data.samplers](#) for more information.
- `augments` (optional): provides a list of transformation operations used to augment all samples of a dataset. See [thelper.transforms.utils.load\\_augments\(\)](#) for more info.
- `train_augments` (optional): provides a list of transformation operations used to augment the training samples of a dataset. See [thelper.transforms.utils.load\\_augments\(\)](#) for more info.
- `valid_augments` (optional): provides a list of transformation operations used to augment the validation samples of a dataset. See [thelper.transforms.utils.load\\_augments\(\)](#) for more info.
- `test_augments` (optional): provides a list of transformation operations used to augment the test samples of a dataset. See [thelper.transforms.utils.load\\_augments\(\)](#) for more info.
- `eval_augments` (optional): provides a list of transformation operations used to augment the validation and test samples of a dataset. See [thelper.transforms.utils.load\\_augments\(\)](#) for more info.
- `base_transforms` (optional): provides a list of transformation operations to apply to all loaded samples. This list will be passed to the constructor of all instantiated dataset parsers. See [thelper.transforms.utils.load\\_transforms\(\)](#) for more info.
- `train_split` (optional): provides the proportion of samples of each dataset to hand off to the training data loader. These proportions are given in a dictionary format (`name: ratio`).
- `valid_split` (optional): provides the proportion of samples of each dataset to hand off to the validation data loader. These proportions are given in a dictionary format (`name: ratio`).
- `test_split` (optional): provides the proportion of samples of each dataset to hand off to the test data loader. These proportions are given in a dictionary format (`name: ratio`).

- `skip_verif` (optional, default=True): specifies whether the dataset split should be verified if resuming a session by parsing the log files generated earlier.
- `skip_split_norm` (optional, default=False): specifies whether the question about normalizing the split ratios should be skipped or not.
- `skip_class_balancing` (optional, default=False): specifies whether the balancing of class labels should be skipped in case the task is classification-related.

Example configuration file:

```
# ...
"loaders": {
    "batch_size": 128, # batch size to use in data loaders
    "shuffle": true, # specifies that the data should be shuffled
    "workers": 4, # number of threads to pre-fetch data batches with
    "train_sampler": { # we can use a data sampler to rebalance classes_
↪ (optional)
        # see e.g. 'thelper.data.samplers.WeightedSubsetRandomSampler'
        # ...
    },
    "train_augments": { # training data augmentation operations
        # see 'thelper.transforms.utils.load_augments'
        # ...
    },
    "eval_augments": { # evaluation (valid/test) data augmentation operations
        # see 'thelper.transforms.utils.load_augments'
        # ...
    },
    "base_transforms": { # global sample transformation operations
        # see 'thelper.transforms.utils.load_transforms'
        # ...
    },
    # optionally indicate how to resolve dataset loader task vs model task_
↪ incompatibility if any
    # leave blank to get more details about each case during runtime if this_
↪ situation happens
    "task_compat_mode": "old|new|compat",
    # finally, we define a 80%-10%-10% split for our data
    # (we could instead use one dataset for training and one for testing)
    "train_split": {
        "dataset_A": 0.8
        "dataset_B": 0.8
    },
    "valid_split": {
        "dataset_A": 0.1
        "dataset_B": 0.1
    },
    "test_split": {
        "dataset_A": 0.1
        "dataset_B": 0.1
    }
    # (note that the dataset names above are defined in the field below)
},
"datasets": {
    "dataset_A": {
        # type of dataset interface to instantiate
        "type": "...",
        "params": {
```

(continues on next page)

(continued from previous page)

```

        # ...
    },
    "dataset_B": {
        # type of dataset interface to instantiate
        "type": "...",
        "params": {
            # ...
        },
        # if it does not derive from 'thelper.data.parsers.Dataset', a task is_
        ↪needed:
        "task": {
            # this type must derive from 'thelper.tasks.Task'
            "type": "...",
            "params": {
                # ...
            }
        }
    },
    # ...
},
# ...

```

### Parameters

- **config** – a dictionary that provides all required data configuration information under two fields, namely ‘datasets’ and ‘loaders’.
- **save\_dir** – the path to the root directory where the session directory should be saved. Note that this is not the path to the session directory itself, but its parent, which may also contain other session directories.

**Returns** A 4-element tuple that contains – 1) the global task object to specialize models and trainers with; 2) the training data loader; 3) the validation data loader; and 4) the test data loader.

### See also:

```

thelper.data.utils.create_parsers()
thelper.transforms.utils.load_augments()
thelper.transforms.utils.load_transforms()

```

`thelper.data.utils.create_parsers` (*config*, *base\_transforms=None*)

Instantiates dataset parsers based on a provided dictionary.

This function will instantiate dataset parsers as defined in a name-type-param dictionary. If multiple datasets are instantiated, this function will also verify their task compatibility and return the global task. The dataset interfaces themselves should be derived from `thelper.data.parsers.Dataset`, be compatible with `thelper.data.parsers.ExternalDataset`, or should provide a ‘task’ field specifying all the information related to sample dictionary keys and model i/o.

The provided configuration will be parsed for a ‘datasets’ dictionary entry. The keys in this dictionary are treated as unique dataset names and are used for lookups. The value associated to each key (or dataset name) should be a type-params dictionary that can be parsed to instantiate the dataset interface.

An example configuration dictionary is given in `thelper.data.utils.create_loaders()`.

**Parameters**

- **config** – a dictionary that provides unique dataset names and parameters needed for instantiation under the ‘datasets’ field.
- **base\_transforms** – the transform operation that should be applied to all loaded samples, and that will be provided to the constructor of all instantiated dataset parsers.

**Returns** *A 2-element tuple that contains* – 1) the list of dataset interfaces/parsers that were instantiated; and 2) a task object compatible with all of those (see [\*thelper.tasks.utils.Task\*](#) for more information).

**See also:**

```
thelper.data.utils.create_loaders()
thelper.data.parsers.Dataset
thelper.data.parsers.ExternalDataset
thelper.tasks.utils.Task
```

```
thelper.data.utils.get_class_weights(label_map, stype='linear', maxw=inf, minw=0.0,
                                     norm=True, invmax=False)
```

Returns a map of label weights that may be adjusted based on a given rebalancing strategy.

**Parameters**

- **label\_map** – map of index lists or sample counts tied to class labels.
- **stype** – weighting strategy (‘uniform’, ‘linear’, or ‘rootX’). Using ‘uniform’ will provide a uniform map of weights. Using ‘linear’ will return the actual weights, unmodified. Using ‘rootX’ will rebalance the weights according to factor ‘X’. See [\*thelper.data.samplers.WeightedSubsetRandomSampler\*](#) for more information on these strategies.
- **maxw** – maximum allowed weight value (applied after invmax, if required).
- **minw** – minimum allowed weight value (applied after invmax, if required).
- **norm** – specifies whether the returned weights should be normalized (default=True, i.e. normalized).
- **invmax** – specifies whether to max-invert the weight vector (thus creating cost factors) or not. Not compatible with `norm` (it would return weights again instead of factors).

**Returns** Map of weights tied to class labels.

**See also:**

```
thelper.data.samplers.WeightedSubsetRandomSampler
```

## 6.1.2 thelper.gui package

Graphical User Interface (GUI) package.

This package contains various tools and annotators used to simplify user interactions with data or models. Since the training framework is CLI-based, these tools are not used to train models, but can be helpful when debugging them. They can also be used to annotate and explore datasets.

### Submodules

#### thelper.gui.annotators module

Graphical User Interface (GUI) annotator module.

This module contains various annotators that define interactive ways to visualize and annotate data loaded via a dataset parser.

**class** `thelper.gui.annotators.Annotator` (*session\_name, config, save\_dir, datasets*)

Bases: `object`

Abstract annotation tool used to define common functions for all GUI-based I/O.

Example configuration file:

```
# ...
"annotator": {
    # type of annotator to instantiate
    "type": "thelper.gui.ImageSegmentAnnotator",
    # ...
    # provide all extra parameters to the specialized anntator here
    "params": {
        # ...
    }
},
# ...
```

See also:

`thelper.gui.annotators.ImageSegmentAnnotator`

`thelper.gui.utils.create_annotator()`

**\_\_init\_\_** (*session\_name, config, save\_dir, datasets*)

Receives the annotator configuration dictionary, parses it, and sets up the basic session attributes.

**abstract run** ()

Displays the GUI tool and blocks until it is closed by the user.

**class** `thelper.gui.annotators.ImageSegmentAnnotator` (*session\_name, config, save\_dir, datasets*)

Bases: `thelper.gui.annotators.Annotator`

Annotator interface specialized for image segmentation annotation generation.

This interface will create a GUI tool with a brush and a zoomed tooltip window that allows images to be painted over with predetermined class labels. The generated masks will then be saved in the session directory as PNG images.

The configuration is expected to provide values for at least the following parameters:

- `sample_input_key`: specifies the key to use when extracting images from loaded samples. This is typically a string defined by the dataset parser.
- `labels`: provides a list of labels that will be available to use in the GUI. These labels are expected to be given as dictionaries that each define an `id` (uint8 value used in output masks), a `name` (string used for display/lookup purposes), and a `color` (3-element integer tuple).

Other parameters can also be provided to alter the GUI's default behavior:

- `default_brush_radius`: the size of the brush at startup (default=10).
- `background_id`: the integer id to use for the background label (default=0).
- `window_zoom_crop_size`: the crop size displayed in the zoom tooltip (default=250x250).
- `window_zoom_size`: the size of the zoom tooltip window (default=500x500).
- `zoom_interp_type`: the interpolation type to use when zooming (default=cv2.INTER\_NEAREST).
- `start_sample_idx`: the index of the first sample to display (default=0).
- `window_name`: the name of the main display window (default=image-segm-annotator).
- `window_size`: the size of the main display window (default=1000).
- `brush_thickness`: the size of the GUI brush tooltip border display (default=2).
- `gui_bar_size`: the width of the GUI bar displayed on top of the main window (default=50).
- `default_mask_opacity`: the default opacity of the segmentation mask (default=0.3).
- `default_fill_id`: the label id to fill all new masks with (default=0).

See also:

*thelper.gui.annotators.Annotator*

**BRUSH\_SIZE = None**

**class Brush** (*config*)

Bases: object

Brush manager used to refresh/draw mask contents based on mouse input.

**\_\_init\_\_** (*config*)

Parses the input config and extracts brush-related parameters.

**static draw\_stroke** (*mask, label\_id, start, end*)

Draws a brush stroke on the mask with a given label id between two points.

**refresh** (*mask, label*)

Fetches the latest mouse state and updates the mask if necessary.

**CURRENT\_KEY = -1**

**GUI\_BAR\_SIZE = None**

**GUI\_DIRTY = True**

**LATEST\_PT = (-1, -1)**

**LATEST\_RAW\_PT = (-1, -1)**

```
MASK_DIRTY = True
MOUSE_FLAGS = 0
WINDOW_SIZE = None

class ZoomTooltip (config)
    Bases: object

    Zoom tooltip manager used to visualize image details based on mouse location.

    __init__ (config)
        Parses the input config and extracts zoom-related parameters.

    refresh (image, mask, mask_colormap, mask_opacity, coords)
        Fetches the latest mouse position and updates the zoom window tooltip.

__init__ (session_name, config, save_dir, datasets)
    Parses the input samples and initializes the annotator GUI elements.

get_mask_path (index)
    Returns the path where the mask of a specific sample should be located.

handle_keys ()
    Fetches the latest keyboard press and updates the annotator state accordingly.

load (index)
    Loads the image and mask associated to a specific sample.

static on_mouse (event, x, y, flags, param)
    Callback entrypoint for opencv to register mouse movement/clicks.

static on_press (key)
    Callback entrypoint for pynput to register keyboard presses.

refresh_gui ()
    Updates and displays the main window based on the latest changes.

refresh_layers ()
    Updates the image, mask, and tool display layers based on the latest changes.

run ()
    Displays the main window and other GUI elements in a loop until it is closed by the user.
```

## thelper.gui.utils module

Graphical User Interface (GUI) utility module.

This module contains various tools and utilities used to instantiate annotators and GUI elements.

`thelper.gui.utils.create_annotator` (*session\_name, save\_dir, config, datasets*)  
Instantiates a GUI annotation tool based on the type contained in the config dictionary.

The tool type is expected to be in the configuration dictionary's *annotator* field, under the *type* key. For more information on the configuration, refer to `thelper.gui.annotators.Annotator`. The instantiated type must be compatible with the constructor signature of `thelper.gui.annotators.Annotator`. The object's constructor will be given the full config dictionary.

### Parameters

- **session\_name** – name of the annotation session used for printing and to create output directories.
- **save\_dir** – path to the session directory where annotations and other outputs will be saved.



- **config** – full configuration dictionary that will be parsed for annotator parameters.
- **datasets** – map of named dataset parsers that will provide the data to annotate.

**Returns** The fully-constructed annotator object, ready to begin annotation via its `run()` function.

**See also:**

*thelper.gui.annotators.Annotator*

`thelper.gui.utils.create_key_listener(callback)`

Returns a key press listener based on `pynput.keyboard` (used for mocking).

### 6.1.3 thelper.infer package

Infer package.

This package contains classes specialized for inference of models on various tasks.

#### Submodules

##### thelper.infer.base module

**class** `thelper.infer.base.Tester`(*session\_name, session\_dir, model, task, loaders, config, ckpt-data=None*)

Bases: *thelper.train.base.Trainer*

Base interface of a session runner for testing.

This call mostly delegates calls to existing Trainer implementation, but limiting their use to ‘eval’ methods to make sure that ‘train’ operations are not called by mistake.

**See also:**

*thelper.train.base.Trainer*

**\_\_init\_\_**(*session\_name, session\_dir, model, task, loaders, config, ckptdata=None*)

Receives the trainer configuration dictionary, parses it, and sets up the session.

**abstract eval\_epoch**(*model, epoch, dev, loader, metrics, output\_path*)

Evaluates the model using the provided objects.

#### Parameters

- **model** – the model with which to run inference that is already uploaded to the target device(s).
- **epoch** – the epoch index we are training for (0-based, and should normally only be 0 for single test pass).
- **dev** – the target device that tensors should be uploaded to (corresponding to model’s device(s)).
- **loader** – the data loader used to get transformed test samples.

- **metrics** – the dictionary of metrics/consumers to report inference results (mostly loggers and basic report generator in this case since there shouldn't be ground truth labels to validate against).
- **output\_path** – directory where output files should be written, if necessary.

**test()**

**test\_epoch** (\*args, \*\*kwargs)

**train()**

Starts the training process.

This function will train the model until the required number of epochs is reached, and then evaluate it on the test data. The setup of loggers, tensorboard writers is done here, so is model improvement tracking via monitored metrics. However, the code related to loss computation and back propagation is implemented in a derived class via `thelper.train.base.Trainer.train_epoch()`.

**train\_epoch** (model, epoch, dev, loss, optimizer, loader, metrics, output\_path)

Trains the model for a single epoch using the provided objects.

#### Parameters

- **model** – the model to train that is already uploaded to the target device(s).
- **epoch** – the epoch index we are training for (0-based).
- **dev** – the target device that tensors should be uploaded to.
- **loss** – the loss function used to evaluate model fidelity.
- **optimizer** – the optimizer used for back propagation.
- **loader** – the data loader used to get transformed training samples.
- **metrics** – the dictionary of metrics/consumers to update every iteration.
- **output\_path** – directory where output files should be written, if necessary.

`thelper.infer.base.make_tester_from_trainer(trainer)`

## thelper.infer.impl module

Explicit Tester definitions from existing Trainers.

## thelper.infer.utils module

`thelper.infer.utils.create_tester(session_name, save_dir, config, model, task, loaders, ckpt_data=None)`

Instantiates the tester object based on the type contained in the config dictionary.

The tester type is expected to be in the configuration dictionary's `tester` field, under the `type` key. For backward compatibility, the fields `runner` and `trainer` will also be looked for. For more information on the configuration, refer to `thelper.train.base.Trainer`. The instantiated type must be compatible with the constructor signature of `thelper.train.base.Trainer`. The object's constructor will be given the full config dictionary and the checkpoint data for resuming the session (if available).

If the trainer type is missing, it will be automatically deduced based on the task object.

#### Parameters

- **session\_name** – name of the training session used for printing and to create internal tensorboardX directories.

- **save\_dir** – path to the session directory where logs and checkpoints will be saved.
- **config** – full configuration dictionary that will be parsed for trainer parameters and saved in checkpoints.
- **model** – model to train/evaluate; should be compatible with `thelper.nn.utils.Module`.
- **task** – global task interface defining the type of model and training goal for the session.
- **loaders** – a tuple containing the training/validation/test data loaders (a loader can be `None` if empty).
- **ckptdata** – raw checkpoint to parse data from when resuming a session (if `None`, will start from scratch).

**Returns** The fully-constructed trainer object, ready to begin model training/evaluation.

**See also:**

`thelper.infer.base.Tester`

### 6.1.4 thelper.nn package

Neural network and model package.

This package contains classes that define blocks and modules used in various neural network architectures. Most of these classes have been adapted from external sources; see their individual headers for more information.

#### Subpackages

##### thelper.nn.segmentation package

Neural network and model package for segmentation.

This package contains classes that define blocks and modules used in various neural network for segmentation

#### Submodules

##### thelper.nn.segmentation.base module

##### thelper.nn.segmentation.deeplabv3 module

##### thelper.nn.segmentation.fcn module

##### thelper.nn.sr package

Neural network and model package for super resolution.

This package contains classes that define blocks and modules used in various neural network for super resolution architectures. Most of these classes have been adapted from external sources; see their individual headers for more information.

## Submodules

### thelper.nn.sr.srcnn module

### thelper.nn.sr.vdsr module

## Submodules

### thelper.nn.common module

```
thelper.nn.common.shave (imgs, border_size=0)
thelper.nn.common.weights_init_kaiming (m)
thelper.nn.common.weights_init_xavier (m)
```

### thelper.nn.coordconv module

```
thelper.nn.coordconv.get_coords_map (height, width, centered=True, normalized=True,
                                     noise=None, dtype=torch.float32)
    Returns a HxW intrinsic coordinates map tensor (shape=2xHxW).

thelper.nn.coordconv.make_conv2d (*args, coordconv=False, centered=True, normalized=True,
                                  noise=None, radius_channel=False, scale=None, **kwargs)
    Creates a 2D convolution layer with optional CoordConv support.

thelper.nn.coordconv.swap_coordconv_layers (module, centered=True, normalized=True,
                                             noise=None, radius_channel=False,
                                             scale=None)
    Modifies the provided module by swapping Conv2d layers for CoordConv-equivalent layers.
```

### thelper.nn.densenet module

```
thelper.nn.densenet.densenet121 (pretrained=False, **kwargs)
    Densenet-121 model from “Densely Connected Convolutional Networks”

    Parameters pretrained (bool) – If True, returns a model pre-trained on ImageNet

thelper.nn.densenet.densenet161 (pretrained=False, **kwargs)
    Densenet-161 model from “Densely Connected Convolutional Networks”

    Parameters pretrained (bool) – If True, returns a model pre-trained on ImageNet

thelper.nn.densenet.densenet169 (pretrained=False, **kwargs)
    Densenet-169 model from “Densely Connected Convolutional Networks”

    Parameters pretrained (bool) – If True, returns a model pre-trained on ImageNet

thelper.nn.densenet.densenet201 (pretrained=False, **kwargs)
    Densenet-201 model from “Densely Connected Convolutional Networks”

    Parameters pretrained (bool) – If True, returns a model pre-trained on ImageNet
```

**thelper.nn.efficientnet module****thelper.nn.fcn module**

`thelper.nn.fcn.get_upsampling_weight` (*in\_channels, out\_channels, kernel\_size*)

**thelper.nn.inceptionresnetv2 module****thelper.nn.lenet module****thelper.nn.mobilenet module**

`thelper.nn.mobilenet.conv_1x1_bn` (*inp, oup*)

`thelper.nn.mobilenet.conv_bn` (*inp, oup, stride*)

**thelper.nn.resnet module**

`thelper.nn.resnet.get_activation_layer` (*name: AnyStr*) → `torch.nn.Module`

**thelper.nn.srm module**

`thelper.nn.srm.setup_srm_layer` (*input\_channels: int = 3*) → `torch.nn.Module`

Creates a SRM convolution layer for noise analysis.

`thelper.nn.srm.setup_srm_weights` (*input\_channels: int = 3*) → `torch.Tensor`

Creates the SRM kernels for noise analysis.

**thelper.nn.unet module****thelper.nn.utils module**

Neural network utility functions and classes.

This module contains base interfaces and utility functions used to define and instantiate neural network models.

`thelper.nn.utils.create_model` (*config, task, save\_dir=None, ckptdata=None*)

Instantiates a model based on a provided task object.

The configuration must be given as a dictionary object. This dictionary will be parsed for a ‘model’ field. This field is expected to be a dictionary itself. It may then specify a type to instantiate as well as the parameters to provide to that class constructor, or a path to a checkpoint from which a model should be loaded.

All models must derive from `thelper.nn.utils.Module`, or they must be instantiable through `thelper.nn.utils.ExternalModule` (or one of its specialized classes). The provided task object will be used to make sure that the model has the required input/output layers for the requested objective.

If checkpoint data is provided by the caller, the weights it contains will be loaded into the returned model.

Usage examples inside a session configuration file:

```
# ...
# the function will look for a 'model' field in the provided config dict
"model": {
    # the type provides the class name to instantiate an object from
    "type": "thelper.nn.mobilenet.MobileNetV2",
    # the parameters listed below are passed to the model's constructor
    "params": {
        # ...
    }
}
# ...
```

### Parameters

- **config** – a session dictionary that provides a ‘model’ field containing a dictionary.
- **task** – a task object that will be passed to the model’s constructor in order to specialize it. Can be `None` if a checkpoint is provided, and if the previous task is wanted instead of a new one.
- **save\_dir** – if not `None`, a log file containing model information will be created there.
- **ckptdata** – raw checkpoint data loaded via `torch.load()`; the model will be given its previous state.

**Returns** The instantiated model, compatible with the interface of both `thelper.nn.utils.Module` and `torch.nn.Module`.

### See also:

```
thelper.nn.utils.Module
thelper.nn.utils.ExternalModule
thelper.tasks.utils.Task
thelper.utils.load_checkpoint()
```

`thelper.nn.utils.get_learnable_param_count(model: torch.nn.Module) → int`  
Returns the learnable (grad-enabled) parameter count in a module.

## 6.1.5 thelper.optim package

Optimization/metrics package.

This package contains metrics implementations used to monitor training sessions and evaluate models, and optimization methods used to control the learning behavior of these models.

## Submodules

### thelper.optim.eval module

Evaluation classes/funcs module.

This module contains procedures used to evaluate models and prediction results on specific tasks or datasets. These procedures may be used as part of metric classes (defined in `thelper.optim.metrics`) or high-level debug/drawing utilities.

`thelper.optim.eval.compute_average_precision` (*precision, recall, method='all-points'*)

Computes the average precision given an array of precision and recall values.

This function is inspired from the ‘Object Detection Metrics’ repository of Rafael Padilla. See <https://github.com/rafaelpadilla/Object-Detection-Metrics> for more information. The original code is distributed under the MIT License, Copyright (c) 2018 Rafael Padilla.

#### Parameters

- **precision** – list of precision values for the evaluated predictions of a class.
- **recall** – list of recall values for the evaluated predictions of a class.
- **method** – the evaluation method to use; can be the the latest & official PASCAL VOC toolkit approach (“all-points”), or the 11-point approach (“11-points”) described in the original paper (“The PASCAL Visual Object Classes(VOC) Challenge”).

**Returns** A 4-element tuple containing the average precision, rectified precision/recall arrays, and the indices used for the integral.

`thelper.optim.eval.compute_bbox_iou` (*bbox1, bbox2*)

Computes and returns the Intersection over Union (IoU) of two bounding boxes.

`thelper.optim.eval.compute_mask_iou` (*mask1, mask2, class\_indices=None, dontcare=None*)

Computes and returns a map of Intersection over Union (IoU) scores for two segmentation masks.

`thelper.optim.eval.compute_pascalvoc_metrics` (*pred\_bboxes, gt\_bboxes, task, iou\_threshold=0.5, method='all-points'*)

Computes the metrics used by the VOC Pascal 2012 challenge.

This function is inspired from the ‘Object Detection Metrics’ repository of Rafael Padilla. See <https://github.com/rafaelpadilla/Object-Detection-Metrics> for more information. The original code is distributed under the MIT License, Copyright (c) 2018 Rafael Padilla.

#### Parameters

- **pred\_bboxes** – list of bbox predictions generated by the model under evaluation.
- **gt\_bboxes** – list of groundtruth bounding boxes defined by the dataset.
- **task** – task definition object that holds a vector of all class names.
- **iou\_threshold** – Intersection Over Union (IOU) threshold for true/false positive classification.
- **method** – the evaluation method to use; can be the the latest & official PASCAL VOC toolkit approach (“all-points”), or the 11-point approach (“11-points”) described in the original paper (“The PASCAL Visual Object Classes(VOC) Challenge”).

#### Returns

*A dictionary containing evaluation information and metrics for each class. Each entry contains*

–

- `precision`: array with the precision values;
- `recall`: array with the recall values;
- `AP`: average precision;
- `interpolated precision`: interpolated precision values;
- `interpolated recall`: interpolated recall values;
- `total positives`: total number of ground truth positives;
- `total TP`: total number of True Positive detections;
- `total FP`: total number of False Negative detections.

## thelper.optim.losses module

## thelper.optim.metrics module

Metrics module.

This module contains classes that implement metrics used to monitor training sessions and evaluate models. These metrics should all inherit from `thelper.optim.metrics.Metric` to allow them to be dynamically instantiated by the framework from a configuration file, and evaluated automatically inside a training session. For more information on this, refer to `thelper.train.base.Trainer`.

**class** `thelper.optim.metrics.Accuracy` (*top\_k=1, max\_win\_size=None*)  
Bases: `thelper.optim.metrics.Metric`

Classification accuracy metric interface.

This is a scalar metric used to monitor the label prediction accuracy of a model. By default, it works in `top-k` mode, meaning that the evaluation result is given by:

$$\text{accuracy} = \frac{\text{nb. correct predictions}}{\text{nb. total predictions}} \cdot 100$$

When  $k > 1$ , a ‘correct’ prediction is obtained if any of the model’s top  $k$  predictions (i.e. the  $k$  predictions with the highest score) match the groundtruth label. Otherwise, if  $k = 1$ , then only the top prediction is compared to the groundtruth label. Note that for binary classification problems,  $k$  should always be set to 1.

This metric’s goal is to maximize its value  $\in [0, 100]$  (a percentage is returned).

Usage example inside a session configuration file:

```
# ...
# lists all metrics to instantiate as a dictionary
"metrics": {
    # ...
    # this is the name of the example metric; it is used for lookup/printing only
    "top_5_accuracy": {
        # this type is used to instantiate the accuracy metric
        "type": "thelper.optim.metrics.Accuracy",
        # these parameters are passed to the wrapper's constructor
        "params": {
            # the top prediction count to check for a match with the groundtruth
            "top_k": 5
        }
    },
    # ...
}
```

(continues on next page)



(continued from previous page)

```
}
# ...
```

Todo: add support for ‘dont care’ target value?

### Variables

- **top\_k** – number of top predictions to consider when matching with the groundtruth (default=1).
- **max\_win\_size** – maximum moving average window size to use (default=None, which equals dataset size).
- **correct** – total number of correct predictions stored using an array for window-based averaging.
- **total** – total number of predictions stored using an array for window-based averaging.
- **warned\_eval\_bad** – toggles whether the division-by-zero warning has been flagged or not.

**\_\_init\_\_** (*top\_k=1, max\_win\_size=None*)

Receives the number of predictions to consider for matches (*top\_k*) and the moving average window size (*window\_size*).

Note that by default, if *max\_win\_size* is not provided here, the value given to *max\_iters* on the first update call will be used instead to fix the sliding window length. In any case, the smallest of *max\_iters* and *max\_win\_size* will be used to determine the actual window size.

**eval** ()

Returns the current accuracy (in percentage) based on the accumulated prediction counts.

Will issue a warning if no predictions have been accumulated yet.

**property goal**

Returns the scalar optimization goal of this metric (maximization).

**reset** ()

Toggles a reset of the metric’s internal state, deallocating count arrays.

**supports\_classification = True**

**supports\_segmentation = True**

**update** (*task, input, pred, target, sample, loss, iter\_idx, max\_iters, epoch\_idx, max\_epochs, output\_path, \*\*kwargs*)

Receives the latest class prediction and groundtruth labels from the training session.

This function computes and accumulate the number of correct and total predictions in the internal arrays, cycling over the iteration index if the maximum window length is reached.

The exact signature of this function should match the one of the callbacks defined in *thelper.train.base.Trainer* and specified by *thelper.typedefs.IterCallbackParams*.

**class** *thelper.optim.metrics.AveragePrecision* (*target\_class=None, iou\_threshold=0.5, method='all-points', max\_win\_size=None*)

Bases: *thelper.optim.metrics.Metric*

Object detection average precision score from PascalVOC.

This metric is computed based on the evaluator function implemented in *thelper.optim.eval*. It can target a single class at a time, or produce the mean average precision for all classes.

Usage example inside a session configuration file:

```
# ...
# lists all metrics to instantiate as a dictionary
"metrics": {
    # ...
    # this is the name of the example metric; it is used for lookup/printing only
    "mAP": {
        # this type is used to instantiate the AP metric
        "type": "thelper.optim.metrics.AveragePrecision",
        # these parameters are passed to the wrapper's constructor
        "params": {
            # no parameters means we will compute the mAP
        }
    },
    # ...
}
```

### Variables

- **target\_class** – name of the class to target; if ‘None’, will compute mAP instead of AP.
- **iou\_threshold** – Intersection Over Union (IOU) threshold for true/false positive classification.
- **method** – the evaluation method to use; can be the the latest & official PASCAL VOC toolkit approach (“all-points”), or the 11-point approach (“11-points”) described in the original paper (“The PASCAL Visual Object Classes(VOC) Challenge”).
- **max\_win\_size** – maximum moving average window size to use (default=None, which equals dataset size).
- **preds** – array holding the predicted bounding boxes for all input samples.
- **targets** – array holding the target bounding boxes for all input samples.

**\_\_init\_\_** (*target\_class=None, iou\_threshold=0.5, method='all-points', max\_win\_size=None*)

Initializes metric attributes.

Note that by default, if `max_win_size` is not provided here, the value given to `max_iters` on the first update call will be used instead to fix the sliding window length. In any case, the smallest of `max_iters` and `max_win_size` will be used to determine the actual window size.

**eval** ()

Returns the current accuracy (in percentage) based on the accumulated prediction counts.

Will issue a warning if no predictions have been accumulated yet.

**property goal**

Returns the scalar optimization goal of this metric (maximization).

**property live\_eval**

Returns whether this metric can/should be evaluated at every backprop iteration or not.

**reset** ()

Toggles a reset of the metric’s internal state, deallocating bbox arrays.

**supports\_detection = True**

**update** (*task, input, pred, target, sample, loss, iter\_idx, max\_iters, epoch\_idx, max\_epochs, output\_path, \*\*kwargs*)

Receives the latest bbox predictions and targets from the training session.

The exact signature of this function should match the one of the callbacks defined in `thelper.train.base.Trainer` and specified by `thelper.typedefs.IterCallbackParams`.

```
class thelper.optim.metrics.ExternalMetric(metric_name, metric_type, metric_goal,
                                           metric_params=None, target_name=None,
                                           class_names=None, max_win_size=None,
                                           force_softmax=True, live_eval=True)
```

Bases: `thelper.optim.metrics.Metric`, `thelper.ifaces.ClassNamesHandler`

External metric wrapping interface.

This interface is used to wrap external metrics and use them in the training framework. The metrics of `sklearn.metrics` are good candidates that have been used extensively with this interface in the past, but those of other libraries might also be compatible.

Along with the name of the class to import and its constructor's parameters, the user must provide a handling mode that specifies how prediction and groundtruth data should be handled in this wrapper. Also, extra arguments such as target label names, goal information, and window sizes can be provided for specific use cases related to the selected handling mode.

For now, two metric handling modes (both related to classification) are supported:

- `classif_best`: the wrapper will accumulate the predicted and groundtruth classification labels forwarded by the trainer and provide them to the external metric for evaluation. If a target label name is specified, then only classifications related to that label will be accumulated. This is the handling mode required for count-based classification metrics such as accuracy, F-Measure, precision, recall, etc.
- `classif_score`: the wrapper will accumulate the prediction score of the targeted label along with a boolean that indicates whether this label was the groundtruth label or not. This is the handling mode required for score-based classification metrics such as when computing the area under the ROC curve (AUC).

Usage examples inside a session configuration file:

```
# ...
# lists all metrics to instantiate as a dictionary
"metrics": {
    # ...
    # this is the name of the first example metric; it is used for lookup/
    ↪printing only
    "fl_score_reject": {
        # this type is used to instantiate the wrapper
        "type": "thelper.optim.metrics.ExternalMetric",
        # these parameters are passed to the wrapper's constructor
        "params": {
            # the external class to import
            "metric_name": "sklearn.metrics.fl_score",
            # the parameters passed to the external class's constructor
            "metric_params": {},
            # the wrapper metric handling mode
            "metric_type": "classif_best",
            # the target class name (note: dataset-specific)
            "target_name": "reject",
            # the goal type of the external metric
            "metric_goal": "max"
        }
    },
    # this is the name of the second example metric; it is used for lookup/
    ↪printing only
    "roc_auc_accept": {
```

(continues on next page)

(continued from previous page)

```

# this type is used to instantiate the wrapper
"type": "thelper.optim.metrics.ExternalMetric",
# these parameters are passed to the wrapper's constructor
"params": {
    # the external class to import
    "metric_name": "sklearn.metrics.roc_auc_score",
    # the parameters passed to the external class's constructor
    "metric_params": {},
    # the wrapper metric handling mode
    "metric_type": "classif_score",
    # the target class name (note: dataset-specific)
    "target_name": "accept",
    # the goal type of the external metric
    "metric_goal": "max"
}
},
# ...
}
# ...

```

### Variables

- **metric\_goal** – goal of the external metric, used for monitoring. Can be `min` or `max`.
- **metric\_type** – handling mode of the external metric. Can only be one of the predetermined values.
- **metric** – type of the external metric that will be instantiated when `eval` is called.
- **metric\_params** – dictionary of parameters passed to the external metric on instantiation.
- **target\_name** – name of the targeted label. Used only in handling modes related to classification.
- **target\_idx** – index of the targeted label. Used only in handling modes related to classification.
- **class\_names** – holds the list of class label names provided by the dataset parser. If it is not provided when the constructor is called, it will be set by the trainer at runtime.
- **force\_softmax** – specifies whether a softmax operation should be applied to the prediction scores obtained from the trainer. Only used with the “`classif_score`” handling mode.
- **max\_win\_size** – maximum moving average window size to use (default=`None`, which equals dataset size).
- **pred** – queue used to store predictions-related values for window-based averaging.
- **target** – queue used to store groundtruth-related values for window-based averaging.

**\_\_init\_\_**(*metric\_name*, *metric\_type*, *metric\_goal*, *metric\_params*=`None`, *target\_name*=`None`, *class\_names*=`None`, *max\_win\_size*=`None`, *force\_softmax*=`True`, *live\_eval*=`True`)  
 Receives all necessary arguments for wrapper initialization and external metric instantiation.

See [thelper.optim.metrics.ExternalMetric](#) for information on arguments.

### property **class\_names**

Returns the list of class names considered “of interest” by the derived class.

### **eval**()

Returns the external metric’s evaluation result.

**property goal**

Returns the scalar optimization goal of this metric (user-defined).

**property live\_eval**

Returns whether this metric can/should be evaluated at every backprop iteration or not.

By default, this returns `True`, but implementations that are quite slow may return `False`.

**reset()**

Toggles a reset of the metric's internal state, emptying pred/target queues.

**supports\_classification = True****supports\_segmentation = True****update** (*task, input, pred, target, sample, loss, iter\_idx, max\_iters, epoch\_idx, max\_epochs, output\_path, \*\*kwargs*)

Receives the latest predictions and target values from the training session.

The handling of the data received here will depend on the current metric's handling mode.

The exact signature of this function should match the one of the callbacks defined in `thelper.train.base.Trainer` and specified by `thelper.typedefs.IterCallbackParams`.

```
class thelper.optim.metrics.IntersectionOverUnion (target_names=None,  
                                                    global_score=True,  
                                                    max_win_size=None)
```

Bases: `thelper.optim.metrics.Metric`

Computes the intersection over union over image classes.

It can target a single class at a time, or produce the mean IoU (mIoU) for a number of classes. It can also average IoU scores from each images, or sum up all intersection and union areas and compute a global score.

Usage example inside a session configuration file:

```
# ...
# lists all metrics to instantiate as a dictionary
"metrics": {
    # ...
    # this is the name of the example metric; it is used for lookup/printing only
    "mIoU": {
        # this type is used to instantiate the IoU metric
        "type": "thelper.optim.metrics.IntersectionOverUnion",
        # these parameters are passed to the wrapper's constructor
        "params": {
            # no parameters means we will compute the mIoU with global scoring
        }
    },
    # ...
}
```

**Variables**

- **target\_names** – name(s) of the class(es) to target; if 'None' or list, will compute mIoU instead of IoU.
- **max\_win\_size** – maximum moving average window size to use (default=None, which equals dataset size).
- **inters** – array holding the intersection areas or IoU scores for all input samples.

- **unions** – array holding the union areas for all input samples.

**\_\_init\_\_** (*target\_names=None, global\_score=True, max\_win\_size=None*)  
Initializes metric attributes.

Note that by default, if `max_win_size` is not provided here, the value given to `max_iters` on the first update call will be used instead to fix the sliding window length. In any case, the smallest of `max_iters` and `max_win_size` will be used to determine the actual window size.

**eval** ()  
Returns the current IoU ratio based on the accumulated counts.  
  
Will issue a warning if no predictions have been accumulated yet.

**property goal**  
Returns the scalar optimization goal of this metric (maximization).

**reset** ()  
Toggles a reset of the metric's internal state, deallocating bbox arrays.

**supports\_segmentation = True**

**update** (*task, input, pred, target, sample, loss, iter\_idx, max\_iters, epoch\_idx, max\_epochs, output\_path, \*\*kwargs*)  
Receives the latest bbox predictions and targets from the training session.  
  
The exact signature of this function should match the one of the callbacks defined in `thelper.train.base.Trainer` and specified by `thelper.typedefs.IterCallbackParams`.

**class** `thelper.optim.metrics.MeanAbsoluteError` (*reduction='mean', max\_win\_size=None*)  
Bases: `thelper.optim.metrics.Metric`

Mean absolute error metric interface.

This is a scalar metric used to monitor the mean absolute deviation (or error) for a model's predictions. This regression metric can be described as:

$$e(x, y) = E = \{e_1, \dots, e_N\}^\top, \quad e_n = |x_n - y_n|,$$

where  $N$  is the batch size. If `reduction` is not 'none', then:

$$\text{MAE}(x, y) = \begin{cases} \text{mean}(E), & \text{if } \text{reduction} = \text{mean.} \\ \text{sum}(E), & \text{if } \text{reduction} = \text{sum.} \end{cases}$$

$x$  and  $y$  are tensors of arbitrary shapes with a total of  $n$  elements each.

Usage example inside a session configuration file:

```
# ...
# lists all metrics to instantiate as a dictionary
"metrics": {
    # ...
    # this is the name of the example metric; it is used for lookup/printing only
    "mae": {
        # this type is used to instantiate the error metric
        "type": "thelper.optim.metrics.MeanAbsoluteError",
        "params": {
            "reduction": "mean"
        }
    },
    # ...
}
```

(continues on next page)

(continued from previous page)

```
}
# ...
```

Todo: add support for 'dont care' target value?

### Variables

- **max\_win\_size** – maximum moving average window size to use (default=None, which equals dataset size).
- **reduction** – string representing the tensor reduction strategy to use.
- **errors** – array of error values stored for window-based averaging.
- **warned\_eval\_bad** – toggles whether the division-by-zero warning has been flagged or not.

**\_\_init\_\_** (*reduction='mean', max\_win\_size=None*)

Receives the reduction strategy and the moving average window size (*window\_size*).

Note that by default, if *max\_win\_size* is not provided here, the value given to *max\_iters* on the first update call will be used instead to fix the sliding window length. In any case, the smallest of *max\_iters* and *max\_win\_size* will be used to determine the actual window size.

**eval** ()

Returns the current (average) mean absolute error based on the accumulated values.

Will issue a warning if no predictions have been accumulated yet.

**property goal**

Returns the scalar optimization goal of this metric (minimization).

**reset** ()

Toggles a reset of the metric's internal state, deallocating the errors array.

**supports\_regression = True**

**update** (*task, input, pred, target, sample, loss, iter\_idx, max\_iters, epoch\_idx, max\_epochs, output\_path, \*\*kwargs*)

Receives the latest predictions and target values from the training session.

This function computes and accumulates the L1 distance between predictions and targets in the internal array, cycling over the iteration index if the maximum window length is reached.

The exact signature of this function should match the one of the callbacks defined in *thelper.train.base.Trainer* and specified by *thelper.typedefs.IterCallbackParams*.

**class** *thelper.optim.metrics.MeanSquaredError* (*reduction='mean', max\_win\_size=None*)

Bases: *thelper.optim.metrics.Metric*

Mean squared error metric interface.

This is a scalar metric used to monitor the mean squared deviation (or error) for a model's predictions. This regression metric can be described as:

$$e(x, y) = E = \{e_1, \dots, e_N\}^\top, \quad e_n = (x_n - y_n)^2,$$

where  $N$  is the batch size. If *reduction* is not 'none', then:

$$\text{MSE}(x, y) = \begin{cases} \text{mean}(E), & \text{if } \text{reduction} = \text{mean.} \\ \text{sum}(E), & \text{if } \text{reduction} = \text{sum.} \end{cases}$$

$x$  and  $y$  are tensors of arbitrary shapes with a total of  $n$  elements each.

Usage example inside a session configuration file:

```
# ...
# lists all metrics to instantiate as a dictionary
"metrics": {
    # ...
    # this is the name of the example metric; it is used for lookup/printing only
    "mse": {
        # this type is used to instantiate the error metric
        "type": "thelper.optim.metrics.MeanSquaredError",
        "params": {
            "reduction": "mean"
        }
    },
    # ...
}
```

Todo: add support for 'dont care' target value?

### Variables

- **max\_win\_size** – maximum moving average window size to use (default=None, which equals dataset size).
- **reduction** – string representing the tensor reduction strategy to use.
- **errors** – array of error values stored for window-based averaging.
- **warned\_eval\_bad** – toggles whether the division-by-zero warning has been flagged or not.

**\_\_init\_\_** (*reduction='mean', max\_win\_size=None*)

Receives the reduction strategy and the moving average window size (*window\_size*).

Note that by default, if *max\_win\_size* is not provided here, the value given to *max\_iters* on the first update call will be used instead to fix the sliding window length. In any case, the smallest of *max\_iters* and *max\_win\_size* will be used to determine the actual window size.

**eval** ()

Returns the current (average) mean squared error based on the accumulated values.

Will issue a warning if no predictions have been accumulated yet.

**property goal**

Returns the scalar optimization goal of this metric (minimization).

**reset** ()

Toggles a reset of the metric's internal state, deallocating the errors array.

**supports\_regression = True**

**update** (*task, input, pred, target, sample, loss, iter\_idx, max\_iters, epoch\_idx, max\_epochs, output\_path, \*\*kwargs*)

Receives the latest predictions and target values from the training session.

This function computes and accumulates the mean squared error between predictions and targets in the internal array, cycling over the iteration index if the maximum window length is reached.

The exact signature of this function should match the one of the callbacks defined in *thelper.train.base.Trainer* and specified by *thelper.typedefs.IterCallbackParams*.

**class** *thelper.optim.metrics.Metric*

Bases: *thelper.ifaces.PredictionConsumer*



Abstract metric interface.

This interface defines basic functions required so that `thelper.train.base.Trainer` can figure out how to instantiate, update, and optimize a given metric while training/evaluating a model.

All metrics, by definition, must be ‘optimizable’. This means that they should return a scalar value when ‘evaluated’ and define an optimal goal (-inf or +inf). If this is not possible, then the class should probably be derived using the more generic `thelper.ifaces.PredictionConsumer` instead.

**abstract eval()**

Returns the metric’s evaluation result.

The returned value should be a scalar. As a model improves, this scalar should get closer to the optimization goal (defined through the ‘goal’ attribute). This value will be queried at the end of each training epoch by the trainer.

**property goal**

Returns the scalar optimization goal of the metric.

The returned goal can be the `minimize` or `maximize` members of `thelper.optim.metrics.Metric` if the class’s evaluation returns a scalar value, and `None` otherwise. The trainer will check this value to see if monitoring the metric’s evaluation result progression is possible.

**property live\_eval**

Returns whether this metric can/should be evaluated at every backprop iteration or not.

By default, this returns `True`, but implementations that are quite slow may return `False`.

**maximize = inf**

Possible value of the `goal` attribute of this metric.

**minimize = -inf**

Possible value of the `goal` attribute of this metric.

**abstract update** (*task, input, pred, target, sample, loss, iter\_idx, max\_iters, epoch\_idx, max\_epochs, output\_path, \*\*kwargs*)

Receives the latest prediction and groundtruth tensors from the training session.

The data given here will be “consumed” internally, but it should NOT be modified. For example, a classification accuracy metric might accumulate the correct number of predictions in comparison to groundtruth labels, but never alter those predictions. The iteration/epoch indices may be used to ‘reset’ the internal state of this object when needed (for example, at the start of each new epoch).

Remember that input, prediction, and target tensors received here will all have a batch dimension!

The exact signature of this function should match the one of the callbacks defined in `thelper.train.base.Trainer` and specified by `thelper.typedefs.IterCallbackParams`.

**class** `thelper.optim.metrics.PSNR` (*data\_range=1.0, max\_win\_size=None*)

Bases: `thelper.optim.metrics.Metric`

Peak Signal-to-Noise Ratio (PSNR) metric interface.

This is a scalar metric used to monitor the change in quality of a signal (or image) following a transformation. For more information, see its definition on [\[Wikipedia\]](#).

The PSNR (in decibels, dB) between a modified signal  $x$  and its original version  $y$  is defined as:

$$\text{PSNR}(x, y) = 10 * \log_{10} \left( \frac{R^2}{\text{MSE}(x, y)} \right)$$

where  $\text{MSE}(x, y)$  returns the mean squared error (see `thelper.optim.metrics.MeanSquaredError` for more information), and  $R$  is the maximum possible value for a single element in the input signal (i.e. its maximum “range”).

Usage example inside a session configuration file:

```
# ...
# lists all metrics to instantiate as a dictionary
"metrics": {
    # ...
    # this is the name of the example metric; it is used for lookup/printing only
    "psnr": {
        # this type is used to instantiate the metric
        "type": "thelper.optim.metrics.PSNR",
        "params": {
            "data_range": "255"
        }
    },
    # ...
}
```

### Variables

- **max\_win\_size** – maximum moving average window size to use (default=None, which equals dataset size).
- **data\_range** – maximum value of an element in the target signal.
- **psnrs** – array of psnr values stored for window-based averaging.
- **warned\_eval\_bad** – toggles whether the division-by-zero warning has been flagged or not.

**\_\_init\_\_** (*data\_range=1.0, max\_win\_size=None*)

Receives all necessary initialization arguments to compute signal PSNRs,

See [thelper.optim.metrics.PSNR](#) for information on arguments.

**eval** ()

Returns the current (average) PSNR based on the accumulated values.

Will issue a warning if no predictions have been accumulated yet.

**property goal**

Returns the scalar optimization goal of this metric (maximization).

**reset** ()

Toggles a reset of the metric's internal state, deallocating the psnrs array.

**supports\_regression = True**

**update** (*task, input, pred, target, sample, loss, iter\_idx, max\_iters, epoch\_idx, max\_epochs, output\_path, \*\*kwargs*)

Receives the latest predictions and target values from the training session.

The exact signature of this function should match the one of the callbacks defined in [thelper.train.base.Trainer](#) and specified by [thelper.typedefs.IterCallbackParams](#).

**class** [thelper.optim.metrics.ROCCurve](#) (*target\_name, target\_tpr=None, target\_fpr=None, class\_names=None, force\_softmax=True, sample\_weight=None, drop\_intermediate=True*)

Bases: [thelper.optim.metrics.Metric](#), [thelper.ifaces.ClassNamesHandler](#)

Receiver operating characteristic (ROC) computation interface.

This class provides an interface to `sklearn.metrics.roc_curve` and `sklearn.metrics.roc_auc_score` that can produce various types of ROC-related information including the area under the curve (AUC), the false positive and negative rates for various operating points, and the ROC curve itself as an image (also compatible with tensorboardX).

By default, evaluating this metric returns the Area Under the Curve (AUC). If a target operating point is set, it will instead return the false positive/negative prediction rate of the model at that point.

Usage examples inside a session configuration file:

```
# ...
# lists all metrics to instantiate as a dictionary
"metrics": {
  # ...
  # this is the name of the first example; it will output the AUC of the "reject
→ " class
  "roc_reject_auc": {
    # this type is used to instantiate the ROC metric
    "type": "thelper.optim.metrics.ROCCurve",
    # these parameters are passed to the constructor
    "params": {
      # the name of the class to evaluate
      "target_name": "reject"
    }
  },
  # this is the name of the second example; it will output the FPR at TPR=0.99
  "roc_reject_0.99tpr": {
    # this type is used to instantiate the ROC metric
    "type": "thelper.optim.metrics.ROCCurve",
    # these parameters are passed to the constructor
    "params": {
      # the name of the class to evaluate
      "target_name": "reject",
      # the target true positive rate (TPR) operating point
      "target_tpr": 0.99
    }
  },
  # ...
}
```

## Variables

- **target\_inv** – used to target all classes except the named one(s); experimental!
- **target\_name** – name of targeted class to generate the roc curve/auc information for.
- **target\_tpr** – target operating point in terms of true positive rate (provided in constructor).
- **target\_fpr** – target operating point in terms of false positive rate (provided in constructor).
- **target\_idx** – index of the targeted class, mapped from target\_name using the class\_names list.
- **class\_names** – holds the list of class label names provided by the dataset parser. If it is not provided when the constructor is called, it will be set by the trainer at runtime.
- **force\_softmax** – specifies whether a softmax operation should be applied to the prediction scores obtained from the trainer.

- **curve** – roc curve generator function, called at evaluation time to generate the output string.
- **auc** – auc score generator function, called at evaluation time to generate the output string.
- **score** – queue used to store prediction score values for window-based averaging.
- **true** – queue used to store groundtruth label values for window-based averaging.

**\_\_init\_\_** (*target\_name, target\_tpr=None, target\_fpr=None, class\_names=None, force\_softmax=True, sample\_weight=None, drop\_intermediate=True*)

Receives the target class/operating point info, log parameters, and roc computation arguments.

#### Parameters

- **target\_name** – name of targeted class to generate the roc curve/auc information for.
- **target\_tpr** – target operating point in terms of true positive rate (provided in constructor).
- **target\_fpr** – target operating point in terms of false positive rate (provided in constructor).
- **class\_names** – holds the list of class label names provided by the dataset parser. If it is not provided when the constructor is called, it will be set by the trainer at runtime.
- **force\_softmax** – specifies whether a softmax operation should be applied to the prediction scores obtained from the trainer.
- **sample\_weight** – passed to `sklearn.metrics.roc_curve` and `sklearn.metrics.roc_auc_score`.
- **drop\_intermediate** – passed to `sklearn.metrics.roc_curve`.

#### **property class\_names**

Returns the list of class names considered “of interest” by the derived class.

#### **eval()**

Returns the evaluation result (AUC/TPR/FPR).

If no target operating point is set, the returned value is the AUC for the target class. If a target TPR is set, the returned value is the FPR for that operating point. If a target FPR is set, the returned value is the TPR for that operating point.

#### **property goal**

Returns the scalar optimization goal of this metric (variable based on target op point).

#### **property live\_eval**

Returns whether this metric can/should be evaluated at every backprop iteration or not.

#### **render()**

Returns the ROC curve as a numpy-compatible RGBA image drawn by pyplot.

#### **reset()**

Toggles a reset of the metric’s internal state, emptying queues.

**supports\_classification = True**

**supports\_segmentation = True**

**update** (*task, input, pred, target, sample, loss, iter\_idx, max\_iters, epoch\_idx, max\_epochs, output\_path, \*\*kwargs*)

Receives the latest predictions and target values from the training session.

The exact signature of this function should match the one of the callbacks defined in `thelper.train.base.Trainer` and specified by `thelper.typedefs.IterCallbackParams`.

## thelper.optim.schedulers module

Schedulers module.

This module contains classes used for scheduling learning rate changes while training a model. All classes defined here should derive from `torch.optim.lr_scheduler._LRScheduler` to remain torch-compatible.

## thelper.optim.utils module

Optimization/metrics utility module.

This module contains utility functions and tools used by the other modules of this package.

`thelper.optim.utils.create_loss_fn` (*config*, *model*, *loader=None*, *uploader=None*)

Instantiates and returns the loss function to use for training.

The default way to specify the loss function to use is to provide a callable type to instantiate as well as its initialization parameters. For example:

```
# ...
"loss": {
    # if we want to use PyTorch's cross entropy loss:
    # >>> loss_fn = torch.nn.CrossEntropyLoss(**params)
    # >>> ...
    # >>> loss = loss_fn(pred, target)
    "type": "torch.nn.CrossEntropyLoss"
    "params": {
        "weight": [ ... ],
        "reduction": "mean",
        # ...
    }
},
# ...
```

The loss function can also be queried from a member function of the model class, as such:

```
# ...
"loss": {
    # to query the model for the loss function:
    # >>> loss_fn = model.get_loss_fn(**params)
    # >>> ...
    # >>> loss = loss_fn(pred, target)
    "model_getter": "get_loss_fn"
    "params": {
        # ...
    }
},
# ...
```

If the model is supposed to compute its own loss, we suggest creating a specialized trainer class. In that case only, the 'loss' field can be omitted from the session configuration file.

If the task is related to image classification or semantic segmentation, the classes can be weighted based on extra parameters in the loss configuration. The strategy used to compute the weights is related to the one in `thelper.data.samplers.WeightedSubsetRandomSampler`. The exact parameters that are expected for class reweighting are the following:

- `weight_distribution` (mandatory, toggle): the dictionary of weights assigned to each class, or the rebalancing strategy to use. If omitted entirely, no class weighting will be performed.

- `weight_param_name` (optional, default="weight"): name of the constructor parameter that expects the weight list.
- `weight_max` (optional, default=inf): the maximum weight that can be assigned to a class.
- `weight_min` (optional, default=0): the minimum weight that can be assigned to a class.
- `weight_norm` (optional, default=True): specifies whether the weights should be normalized or not.

This function also supports an extra special parameter if the task is related to semantic segmentation: `ignore_index`. If this parameter is found and not `None` (integer), then the loss function will ignore the given value when computing the loss of a sample. The exact parameters that are expected in this case are the following:

- `ignore_index_param_name` (optional, default="ignore\_index"): name of the constructor parameter that expects the ignore value.
- `ignore_index_label_name` (optional, default="dontcare"): name of the label to pass the ignore value from.

`thelper.optim.utils.create_metrics` (*config*)

Instantiates and returns the metrics defined in the configuration dictionary.

All arguments are expected to be handed in through the configuration via a dictionary named 'params'.

`thelper.optim.utils.create_optimizer` (*config, model*)

Instantiates and returns the optimizer to use for training.

By default, the optimizer will be instantiated with the model parameters given as the first argument of its constructor. All supplementary arguments are expected to be handed in through the configuration via a dictionary named 'params'.

`thelper.optim.utils.create_scheduler` (*config, optimizer*)

Instantiates and returns the learning rate scheduler to use for training.

All arguments are expected to be handed in through the configuration via a dictionary named 'params'.

`thelper.optim.utils.get_lr` (*optimizer*)

Returns the optimizer's learning rate, or 0 if not found.

## 6.1.6 thelper.session package

### Submodules

#### thelper.session.base module

**class** `thelper.session.base.SessionRunner` (*session\_name, session\_dir, model, task, loaders, config, ckptdata=None*)

Bases: object

Abstract session runner interface that defines basic session i/o and setup operations.

This class offers the most basic methods that can be employed by more specialized training or inference sessions. By itself, it doesn't actually run anything.

#### Variables

- **checkpoint\_dir** – session checkpoint output directory (located within the 'session directory').
- **config** – session configuration dictionary holding all original settings, including trainer configuration.

- **devices** – list of (cuda) device IDs to upload the model/tensors to; can be empty if only the CPU is available.
- **epochs** – number of epochs to train the model for.
- **logger** – used to output debug/warning/error messages to session log.
- **model** – reference to the model being trained or used for evaluation/prediction.
- **monitor** – name of the training/validation metric that should be monitored for model improvement.
- **name** – name of the session, used for printing and creating log folders.
- **optimization\_config** – dictionary of optim-related parameters, parsed at training time.
- **output\_paths** – map of session output paths where training/evaluation results should be saved.
- **save\_freq** – frequency of checkpoint saves while training (i.e. save every X epochs).
- **save\_raw** – specifies whether to save raw types or thelper objects in checkpoints.
- **skip\_eval\_iter** – number of evaluation iterations to skip (useful for resuming a session).
- **skip\_tbx\_histograms** – flag used to skip the generation of graph histograms in tbx (useful for large models).
- **task** – reference to the object used to specialize the model and that holds task metainformation.
- **tbx\_histogram\_freq** – frequency of tbx histogram saves while training (i.e. save every X epochs).
- **use\_tbx** – defines whether to use tensorboardX writers for logging or not.
- **writers** – map of tbx writers used to save training/evaluation events.

See also:

*thelper.train.base.Trainer*

`__init__` (*session\_name, session\_dir, model, task, loaders, config, ckptdata=None*)  
Receives the trainer configuration dictionary, parses it, and sets up the session.

### 6.1.7 thelper.tasks package

Task definition package.

This package contains classes and functions whose role is to define the input/output formats and operations expected from a trained model. This essentially defines the ‘goal’ of the model, and is used to specialize and automate its training.

## Submodules

### thelper.tasks.classif module

Classification task interface module.

This module contains a class that defines the objectives of models/trainers for classification tasks.

```
class thelper.tasks.classif.Classification (class_names: Iterable[AnyStr], input_key:
                                         collections.abc.Hashable, label_key: col-
                                         lections.abc.Hashable, meta_keys: Op-
                                         tional[Iterable[collections.abc.Hashable]] =
                                         None, multi_label: bool = False)
```

Bases: *thelper.tasks.utils.Task, thelper.ifaces.ClassNamesHandler*

Interface for input labeling/classification tasks.

This specialization requests that when given an input tensor, the trained model should provide prediction scores for each predefined label (or class). The label names are used here to help categorize samples, and to assure that two tasks are only identical when their label counts and ordering match.

#### Variables

- **class\_names** – list of label (class) names to predict (each name should be a string).
- **input\_key** – the key used to fetch input tensors from a sample dictionary.
- **label\_key** – the key used to fetch label (class) names/indices from a sample dictionary.
- **meta\_keys** – the list of extra keys provided by the data parser inside each sample.

See also:

*thelper.tasks.utils.Task*  
*thelper.train.classif.ImageClassifTrainer*

```
__init__ (class_names: Iterable[AnyStr], input_key: collections.abc.Hashable, label_key: col-
         lections.abc.Hashable, meta_keys: Optional[Iterable[collections.abc.Hashable]] = None,
         multi_label: bool = False)
```

Receives and stores the class (or label) names to predict, the input tensor key, the groundtruth label (class) key, and the extra (meta) keys produced by the dataset parser(s).

The class names can be provided as a list of strings, or as a path to a json file that contains such a list. The list must contain at least two items. All other arguments are used as-is to index dictionaries, and must therefore be key-compatible types.

If the *multi\_label* is activated, samples with non-scalar class labels will be allowed in the *get\_class\_sizes* and *get\_class\_sample\_map* functions.

```
check_compat (task: thelper.tasks.utils.Task, exact: bool = False) → bool
```

Returns whether the current task is compatible with the provided one or not.

This is useful for sanity-checking, and to see if the inputs/outputs of two models are compatible. If *exact* = *True*, all fields will be checked for exact (perfect) compatibility (in this case, matching meta keys and class name order).

```
get_class_sample_map (samples: Iterable, unset_key: collections.abc.Hashable = None) →
                     Dict[AnyStr, List[int]]
```

Splits a list of samples based on their labels into a map of sample lists.



This function is useful if we need to split a dataset based on its label categories in order to sort it, augment it, or re-balance it. The samples do not need to be fully loaded for this to work, as only their label (gt) value will be queried. If a sample is missing its label, it will be ignored and left out of the generated dictionary unless a value is given for `unset_key`.

#### Parameters

- **samples** – the samples to split, where each sample is provided as a dictionary.
- **unset\_key** – a key under which all unlabeled samples should be kept (`None` = ignore).

**Returns** A dictionary that maps each class label to its corresponding list of samples.

**get\_class\_sizes** (*samples: Iterable*) → Dict[AnyStr, int]

Given a list of samples, returns a map of sample counts for each class label.

**get\_compat** (*task: [thelper.tasks.utils.Task](#)*) → *[thelper.tasks.utils.Task](#)*

Returns a task instance compatible with the current task and the given one.

**supports\_classification** = True

## thelper.tasks.detect module

Detection task interface module.

This module contains classes that define object detection utilities and task interfaces.

```
class thelper.tasks.detect.BoundingBox (class_id, bbox, include_margin=True, difficult=False, occluded=False, truncated=False, iscrowd=False, confidence=None, image_id=None, task=None)
```

Bases: `object`

Interface used to hold instance metadata for object detection tasks.

Object detection trainers and display utilities in the framework will expect this interface to be used when parsing a predicted detection or an annotation. The base contents are based on the PASCALVOC metadata structure, and this class can be derived if necessary to contain more metadata.

#### Variables

- **class\_id** – type identifier for the underlying object instance.
- **bbox** – four-element tuple holding the (xmin,ymin,xmax,ymax) bounding box parameters.
- **include\_margin** – defines whether xmax/ymax is included in the bounding box area or not.
- **difficult** – defines whether this instance is considered “difficult” (false by default).
- **occluded** – defines whether this instance is considered “occluded” (false by default).
- **truncated** – defines whether this instance is considered “truncated” (false by default).
- **iscrowd** – defines whether this instance covers a “crowd” of objects or not (false by default).
- **confidence** – scalar or array of prediction confidence values tied to class types (empty by default).
- **image\_id** – string used to identify the image containing this bounding box (i.e. file path or uuid).

- **task** – reference to the task object that holds extra metadata regarding the content of the bbox (None by default).

See also:

`thelper.tasks.utils.Task`

`thelper.tasks.detect.Detection`

**\_\_init\_\_** (*class\_id, bbox, include\_margin=True, difficult=False, occluded=False, truncated=False, is-crowd=False, confidence=None, image\_id=None, task=None*)

Receives and stores low-level input detection metadata for later access.

**property area**

Returns a scalar indicating the total surface of the annotation (may be None if unknown/unspecified).

**property bbox**

Returns the bounding box tuple ( $x_{min}, y_{min}, x_{max}, y_{max}$ ).

**property bottom**

Returns the bottom bounding box edge origin offset value.

**property bottom\_right**

Returns the bottom right bounding box corner coordinates ( $x, y$ ).

**property centroid**

Returns the bounding box centroid coordinates ( $x, y$ ).

**property class\_id**

Returns the object class type identifier.

**property confidence**

Returns the confidence value (or array of confidence values) associated to the predicted class types.

**static decode** (*vec, format=None*)

Returns a BoundingBox object from a vectorized representation in a specified format.

---

**Note:** The input bbox is expected to be a 4 element array ( $x_{min}, y_{min}, x_{max}, y_{max}$ ).

---

**property difficult**

Returns whether this bounding box is considered *difficult* by the dataset (false by default).

**encode** (*format=None*)

Returns a vectorizable representation of this bounding box in a specified format.

WARNING: Encoding might cause information loss (e.g. task reference is discarded).

**property height**

Returns the height of the bounding box.

**property image\_id**

Returns the image string identifier.

**property include\_margin**

Returns whether  $x_{max}$  and  $y_{max}$  are included in the bounding box area or not

**intersects** (*geom*)

Returns whether the bounding box intersects a geometry (i.e. a 2D point or another bbox).

**property iscrowd**

Returns whether this instance covers a *crowd* of objects or not (false by default).

**json()**

Gets a JSON-serializable representation of the bounding box parameters.

**property left**

Returns the left bounding box edge origin offset value.

**property occluded**

Returns whether this bounding box is considered *occluded* by the dataset (false by default).

**property right**

Returns the right bounding box edge origin offset value.

**supports\_detection = True****property task**

Returns the reference to the task object that holds extra metadata regarding the content of the bbox.

**tolist()**

Gets a `list` representation of the underlying bounding box tuple  $(x_{min}, y_{min}, x_{max}, y_{max})$ .

This ensures that `Tensor` objects are converted to native *Python* types.

**property top**

Returns the top bounding box edge origin offset value.

**property top\_left**

Returns the top left bounding box corner coordinates  $(x, y)$ .

**totuple()**

Gets a `tuple` representation of the underlying bounding box tuple  $(x_{min}, y_{min}, x_{max}, y_{max})$ .

This ensures that `Tensor` objects are converted to native *Python* types.

**property truncated**

Returns whether this bounding box is considered *truncated* by the dataset (false by default).

**property width**

Returns the width of the bounding box.

```
class thelper.tasks.detect.Detection(class_names, input_key, bboxes_key, meta_keys=None,
                                     input_shape=None, target_shape=None, target_min=None, target_max=None, background=None,
                                     color_map=None)
```

Bases: `thelper.tasks.regr.Regression`, `thelper.ifaces.ClassNamesHandler`

Interface for object detection tasks.

This specialization requests that when given an input image, the trained model should provide a list of bounding box (bbox) proposals that correspond to probable objects detected in the image.

This specialized regression interface is currently used to help display functions.

**Variables**

- **class\_names** – map of class name-value pairs for object types to detect.
- **input\_key** – the key used to fetch input tensors from a sample dictionary.
- **bboxes\_key** – the key used to fetch target (groundtruth) bboxes from a sample dictionary.
- **meta\_keys** – the list of extra keys provided by the data parser inside each sample.
- **input\_shape** – a numpy-compatible shape to expect input images to possess.

- **target\_shape** – a numpy-compatible shape to expect the predictions to be in.
- **target\_min** – a 2-dim tensor containing minimum (x,y) bounding box corner values.
- **target\_max** – a 2-dim tensor containing maximum (x,y) bounding box corner values.
- **background** – value of the ‘background’ label (if any) used in the class map.
- **color\_map** – map of class name-color pairs to use when displaying results.

See also:

*thelper.tasks.utils.Task*  
*thelper.tasks.regr.Regression*  
*thelper.train.detect.ObjDetectTrainer*

**\_\_init\_\_** (*class\_names*, *input\_key*, *bboxes\_key*, *meta\_keys=None*, *input\_shape=None*, *target\_shape=None*, *target\_min=None*, *target\_max=None*, *background=None*, *color\_map=None*)

Receives and stores the bbox types to detect, the input tensor key, the groundtruth bboxes list key, the extra (meta) keys produced by the dataset parser(s), and the color map used to color bboxes when displaying results.

The class names can be provided as a list of strings, as a path to a json file that contains such a list, or as a map of predefined name-value pairs to use in gt maps. This list/map must contain at least two elements (background and one class). All other arguments are used as-is to index dictionaries, and must therefore be key-compatible types.

**property background**

Returns the ‘background’ label value used in loss functions (can be None).

**check\_compat** (*task*, *exact=False*)

Returns whether the current task is compatible with the provided one or not.

This is useful for sanity-checking, and to see if the inputs/outputs of two models are compatible. If *exact* = *True*, all fields will be checked for exact (perfect) compatibility (in this case, matching meta keys and class maps).

**property color\_map**

Returns the color map used to swap label indices for colors when displaying results.

**get\_class\_sizes** (*samples*, *bbox\_format=None*)

Given a list of samples, returns a map of element counts for each object type.

**get\_compat** (*task*)

Returns a task instance compatible with the current task and the given one.

**supports\_detection** = *True*

## thelper.tasks.regr module

Regression task interface module.

This module contains a class that defines the objectives of models/trainers for regression tasks.

```
class thelper.tasks.regr.Regression(input_key, target_key, meta_keys=None, input_shape=None, target_shape=None, target_type=None, target_min=None, target_max=None)
```

Bases: *thelper.tasks.utils.Task*

Interface for n-dimension regression tasks.

This specialization requests that when given an input tensor, the trained model should provide an n-dimensional target prediction. This is a fairly generic task that (unlike image classification and semantic segmentation) is not linked to a pre-existing set of possible solutions. The task interface is used to carry useful metadata for this task, e.g. input/output shapes, types, and min/max values for rounding/saturation.

### Variables

- *input\_shape* – a numpy-compatible shape to expect model inputs to be in.
- *target\_shape* – a numpy-compatible shape to expect the predictions to be in.
- *target\_type* – a numpy-compatible type to cast the predictions to (if needed).
- *target\_min* – an n-dim tensor containing minimum target values (if applicable).
- *target\_max* – an n-dim tensor containing maximum target values (if applicable).
- *input\_key* – the key used to fetch input tensors from a sample dictionary.
- *target\_key* – the key used to fetch target (groundtruth) values from a sample dictionary.
- *meta\_keys* – the list of extra keys provided by the data parser inside each sample.

See also:

```
thelper.tasks.utils.Task
thelper.train.regr.RegressionTrainer
thelper.tasks.regr.SuperResolution
thelper.tasks.detect.Detection
```

```
__init__(input_key, target_key, meta_keys=None, input_shape=None, target_shape=None, target_type=None, target_min=None, target_max=None)
    Receives and stores the keys produced by the dataset parser(s).
```

```
check_compat(task, exact=False)
    Returns whether the current task is compatible with the provided one or not.
```

This is useful for sanity-checking, and to see if the inputs/outputs of two models are compatible. If `exact = True`, all fields will be checked for exact (perfect) compatibility (in this case, matching meta keys).

```
get_compat(task)
    Returns a task instance compatible with the current task and the given one.
```

```
property input_shape
    Returns the shape of the input tensors to be processed by the model.
```

```
supports_regression = True
```

**property target\_max**

Returns the maximum target value(s) to be generated by the model.

**property target\_min**

Returns the minimum target value(s) to be generated by the model.

**property target\_shape**

Returns the shape of the output tensors to be generated by the model.

**property target\_type**

Returns the type of the output tensors to be generated by the model.

```
class thelper.tasks.regr.SuperResolution(input_key, target_key, meta_keys=None, input_shape=None, target_type=None, target_min=None, target_max=None)
```

Bases: *thelper.tasks.regr.Regression*

Interface for super-resolution tasks.

This specialization requests that when given an input tensor, the trained model should provide an identically-shape target prediction that essentially contains more (or more adequate) high-frequency spatial components.

This specialized regression interface is currently used to help display functions.

**Variables**

- **input\_shape** – a numpy-compatible shape to expect model inputs/outputs to be in.
- **target\_type** – a numpy-compatible type to cast the predictions to (if needed).
- **target\_min** – an n-dim tensor containing minimum target values (if applicable).
- **target\_max** – an n-dim tensor containing maximum target values (if applicable).
- **input\_key** – the key used to fetch input tensors from a sample dictionary.
- **target\_key** – the key used to fetch target (groundtruth) values from a sample dictionary.
- **meta\_keys** – the list of extra keys provided by the data parser inside each sample.

See also:

*thelper.tasks.utils.Task*

*thelper.tasks.regr.Regression*

*thelper.train.regr.RegressionTrainer*

```
__init__(input_key, target_key, meta_keys=None, input_shape=None, target_type=None, target_min=None, target_max=None)
```

Receives and stores the keys produced by the dataset parser(s).

```
supports_regression = True
```

## thelper.tasks.segm module

Segmentation task interface module.

This module contains a class that defines the objectives of models/trainers for segmentation tasks.

```
class thelper.tasks.segm.Segmentation (class_names,      input_key,      label_map_key,
                                         meta_keys=None,    dontcare=None,
                                         color_map=None)
```

Bases: *thelper.tasks.utils.Task, thelper.ifaces.ClassNamesHandler*

Interface for pixel-level labeling/classification (segmentation) tasks.

This specialization requests that when given an input tensor, the trained model should provide prediction scores for each predefined label (or class), for each element of the input tensor. The label names are used here to help categorize samples, and to assure that two tasks are only identical when their label counts and ordering match.

### Variables

- **class\_names** – map of class name-value pairs to predict for each pixel.
- **input\_key** – the key used to fetch input tensors from a sample dictionary.
- **label\_map\_key** – the key used to fetch label (class) maps from a sample dictionary.
- **meta\_keys** – the list of extra keys provided by the data parser inside each sample.
- **dontcare** – value of the ‘dontcare’ label (if any) used in the class map.
- **color\_map** – map of class name-color pairs to use when displaying results.

See also:

*thelper.tasks.utils.Task*

*thelper.train.segm.ImageSegmTrainer*

```
__init__ (class_names,  input_key,  label_map_key,  meta_keys=None,  dontcare=None,
          color_map=None)
```

Receives and stores the class (or label) names to predict, the input tensor key, the groundtruth label (class) map key, the extra (meta) keys produced by the dataset parser(s), the ‘dontcare’ label value that might be present in gt maps (if any), and the color map used to swap label indices for colors when displaying results.

The class names can be provided as a list of strings, as a path to a json file that contains such a list, or as a map of predefined name-value pairs to use in gt maps. This list/map must contain at least two elements. All other arguments are used as-is to index dictionaries, and must therefore be key-compatible types.

```
check_compat (task, exact=False)
```

Returns whether the current task is compatible with the provided one or not.

This is useful for sanity-checking, and to see if the inputs/outputs of two models are compatible. If `exact = True`, all fields will be checked for exact (perfect) compatibility (in this case, matching meta keys and class maps).

```
property color_map
```

Returns the color map used to swap label indices for colors when displaying results.

```
property dontcare
```

Returns the ‘dontcare’ label value used in loss functions (can be `None`).

```
get_class_sizes (samples)
```

Given a list of samples, returns a map of element counts for each class label.

**get\_compat** (*task*)

Returns a task instance compatible with the current task and the given one.

**supports\_segmentation** = True

## thelper.tasks.utils module

Task utility functions & base interface module.

This module contains utility functions used to instantiate tasks and check their compatibility, and the base interface used to define new tasks.

```
class thelper.tasks.utils.Task (input_key: collections.abc.Hashable, gt_key: collections.abc.Hashable = None, meta_keys: Optional[Iterable[collections.abc.Hashable]] = None)
```

Bases: object

Basic task interface that defines a training objective and that holds sample i/o keys.

Since the framework's data loaders expect samples to be passed in as dictionaries, keys are required to obtain the input that should be forwarded to a model, and to obtain the groundtruth required for the evaluation of model predictions. Other keys might also be kept by this interface for reference (these are considered meta keys).

Note that while this interface can be instantiated directly, trainers and models might not be provided enough information about their goal to be correctly instantiated. Thus, specialized task objects derived from this base class should be used if possible.

### Variables

- **input\_key** – the key used to fetch input tensors from a sample dictionary.
- **gt\_key** – the key used to fetch gt tensors from a sample dictionary.
- **meta\_keys** – the list of extra keys provided by the data parser inside each sample.

See also:

*thelper.tasks.classif.Classification*

*thelper.tasks.segm.Segmentation*

*thelper.tasks.regr.Regression*

*thelper.tasks.detect.Detection*

```
__init__ (input_key: collections.abc.Hashable, gt_key: collections.abc.Hashable = None, meta_keys: Optional[Iterable[collections.abc.Hashable]] = None)
```

Receives and stores the keys used to index dataset sample contents.

**check\_compat** (*task*: thelper.tasks.utils.Task, *exact*: bool = False) → bool

Returns whether the current task is compatible with the provided one or not.

This is useful for sanity-checking, and to see if the inputs/outputs of two models are compatible. It should be overridden in derived classes to specialize the compatibility verification. If `exact = True`, all fields will be checked for exact compatibility.

**get\_compat** (*task*: thelper.tasks.utils.Task) → thelper.tasks.utils.Task

Returns a task instance compatible with the current task and the given one.

**property gt\_key**

Returns the key used to fetch groundtruth data tensors from a sample dictionary.



**property input\_key**

Returns the key used to fetch input data tensors from a sample dictionary.

**property keys**

Returns a list of all keys used to carry tensors and metadata in samples.

**property meta\_keys**

Returns the list of keys used to carry meta/auxiliary data in samples.

`thelper.tasks.utils.create_global_task` (*tasks: Optional[Iterable[Task]]*) → Optional[*thelper.tasks.utils.Task*]

Returns a new task object that is compatible with a list of subtasks.

When different datasets must be combined in a session, the tasks they define must also be merged. This functions allows us to do so as long as the tasks all share a common objective. If creating a globally-compatible task is impossible, this function will raise an exception. Otherwise, the returned task object can be used to replace the subtasks of all used datasets.

**See also:**

*thelper.tasks.utils.Task*  
*thelper.tasks.utils.create\_task()*  
*thelper.data.utils.create\_parsers()*

`thelper.tasks.utils.create_task` (*config: Union[Dict[AnyStr, Union[AnyStr, float, int, List[Any], Dict[Any, Any], ConfigDict]], AnyStr]*) → *Task*

Parses a configuration dictionary or repr string and instantiates a task from it.

If a string is provided, it will first be parsed to get the task type, and then the object will be instantiated by forwarding the parameters contained in the string to the constructor of that type. Note that it is important for this function to work that the constructor argument names match the names of parameters printed in the task's `__repr__` function.

If a dict is provided, it should contain a 'type' and a 'params' field with the values required for direct instantiation.

If a *Task* instance was specified, it is directly returned.

**See also:**

*thelper.tasks.utils.Task*

## 6.1.8 thelper.train package

Trainer package.

This package contains classes specialized for training models on various tasks.

## Submodules

### thelper.train.ae module

**class** `thelper.train.ae.AutoEncoderTrainer` (*session\_name, session\_dir, model, task, loaders, config, ckptdata=None*)

Bases: `thelper.train.base.Trainer`

**\_\_init\_\_** (*session\_name, session\_dir, model, task, loaders, config, ckptdata=None*)  
Receives session parameters, parses image/label keys from task object, and sets up metrics.

**eval\_epoch** (*model, epoch, dev, loader, metrics, output\_path*)  
Evaluates the model using the provided objects.

#### Parameters

- **model** – the model to evaluate that is already uploaded to the target device(s).
- **epoch** – the epoch index we are training for (0-based).
- **dev** – the target device that tensors should be uploaded to.
- **loader** – the data loader used to get transformed valid/test samples.
- **metrics** – the dictionary of metrics/consumers to update every iteration.
- **output\_path** – directory where output files should be written, if necessary.

**supports\_classification** = **True**

**supports\_segmentation** = **True**

**train\_epoch** (*model, epoch, dev, classif\_loss, optimizer, loader, metrics, output\_path*)  
Trains the model for a single epoch using the provided objects.

#### Parameters

- **model** – the model to train that is already uploaded to the target device(s).
- **epoch** – the epoch index we are training for (0-based).
- **dev** – the target device that tensors should be uploaded to.
- **loss** – the loss function used to evaluate model fidelity.
- **optimizer** – the optimizer used for back propagation.
- **loader** – the data loader used to get transformed training samples.
- **metrics** – the dictionary of metrics/consumers to update every iteration.
- **output\_path** – directory where output files should be written, if necessary.

### thelper.train.base module

Model training/evaluation base interface module.

This module contains the interface required to train and/or evaluate a model based on different tasks. The trainers based on this interface are instantiated in launched sessions based on configuration dictionaries.

**class** `thelper.train.base.Trainer` (*session\_name, session\_dir, model, task, loaders, config, ckptdata=None*)

Bases: `thelper.session.base.SessionRunner`

Abstract trainer interface that defines basic session i/o and setup operations.

This interface defines the general behavior of a training session which includes configuration parsing, tensorboard setup, metrics and goal setup, and loss/optimizer setup. It also provides utilities for uploading models and tensors on specific devices, and for saving the state of a session. This interface should be specialized for every task by implementing the `train_epoch` and `eval_epoch` functions in a derived class. For better support from visualization utilities, the derived class should also implement `to_tensor`. See `thelper.train.classif.ImageClassifTrainer` for a complete example.

Any trainer derived from this class will alternate between training and validation epochs. It will also support post-training (final) evaluation using a separate test set. If requested, visualizations can be computed after the validation epoch during training (e.g. sample activation maps, or t-SNE plots). See `thelper.viz` for more information on these.

The main parameters that will be parsed by this interface from a configuration dictionary are the following:

- `epochs` (mandatory if training): number of epochs to train for; one epoch is one iteration over all mini-batches.
- `optimization` (mandatory if training): sub-dictionary containing types and extra parameters required for instantiating the loss, optimizer, and scheduler objects. See the code of each related loading function for more information on special parameters.
- `save_freq` (optional, default=1): checkpoint save frequency (will save every epoch multiple of given number).
- `save_raw` (optional, default=True): specifies whether to save raw types or helper objects in checkpoints.
- `use_tbx` (optional, default=False): defines whether to use tensorboardX writers for logging or not.
- `device` (optional): specifies which device to train/evaluate the model on (default=all available).
- `metrics`: list of metrics to instantiate and update during training/evaluation; see related loading function for more information.
- `monitor`: specifies the name of the metric that should be monitored on the validation set for model improvement.

Example configuration file:

```
# ...
"trainer": {
    # type of trainer to instantiate (linked to task type)
    "type": "thelper.train.ImageClassifTrainer",
    # train for 40 epochs
    "epochs": 40,
    # save every 5 epochs
    "save_freq": 5,
    # monitor validation accuracy and save best model based on that
    "monitor": "accuracy",
    # optimization parameters block
    "optimization": {
        # all types & params below provided by PyTorch
        "loss": {
            "type": "torch.nn.CrossEntropyLoss"
        },
        "optimizer": {
            "type": "torch.optim.SGD",
            "params": {
                "lr": 0.1,
                "momentum": 0.9,
                "weight_decay": 1e-06,
                "nesterov": true
            }
        }
    }
}
```

(continues on next page)

(continued from previous page)

```

        }
    },
    "scheduler": {
        "type": "torch.optim.lr_scheduler.StepLR",
        "params": {
            "step_size": 10,
            "step_size": 0.1
        }
    }
},
# visualization block (optional)
"viz": {
    # multiple visualization techniques can be toggled by name
    "tsne": {
        # visualization parameters would be provided here
    }
},
# in this example, we use two consumers in total
# (one metric for monitoring, and one for logging)
"metrics": {
    "accuracy": {
        "type": "thelper.optim.Accuracy"
    },
    "fullreport": {
        "type": "thelper.train.ClassifReport"
    }
}
}
# ...

```

### Variables

- **config** – session configuration dictionary holding all original settings, including trainer configuration.
- **model** – reference to the model being trained or used for evaluation/prediction.
- **task** – reference to the object used to specialize the model and that holds task meta-information.

TODO: move static utils to their related modules

See also:

```

thelper.session.base.SessionRunner
thelper.train.classif.ImageClassifTrainer
thelper.train.segm.ImageSegmTrainer
thelper.train.detect.ObjDetectTrainer
thelper.train.regr.RegressionTrainer
thelper.train.utils.create_trainer()

```

**\_\_init\_\_** (session\_name, session\_dir, model, task, loaders, config, ckptdata=None)  
 Receives the trainer configuration dictionary, parses it, and sets up the session.

**eval()**

Starts the evaluation process.

This function will evaluate the model using the test data (or the validation data, if no test data is available), and return the results. Note that the code related to the forwarding of samples inside the model itself is implemented in a derived class via `thelper.train.base.Trainer.eval_epoch()`.

**abstract eval\_epoch(model, epoch, device, loader, metrics, output\_path)**

Evaluates the model using the provided objects.

**Parameters**

- **model** – the model to evaluate that is already uploaded to the target device(s).
- **epoch** – the epoch index we are training for (0-based).
- **device** – the target device that tensors should be uploaded to.
- **loader** – the data loader used to get transformed valid/test samples.
- **metrics** – the dictionary of metrics/consumers to update every iteration.
- **output\_path** – directory where output files should be written, if necessary.

**train()**

Starts the training process.

This function will train the model until the required number of epochs is reached, and then evaluate it on the test data. The setup of loggers, tensorboard writers is done here, so is model improvement tracking via monitored metrics. However, the code related to loss computation and back propagation is implemented in a derived class via `thelper.train.base.Trainer.train_epoch()`.

**abstract train\_epoch(model, epoch, dev, loss, optimizer, loader, metrics, output\_path)**

Trains the model for a single epoch using the provided objects.

**Parameters**

- **model** – the model to train that is already uploaded to the target device(s).
- **epoch** – the epoch index we are training for (0-based).
- **dev** – the target device that tensors should be uploaded to.
- **loss** – the loss function used to evaluate model fidelity.
- **optimizer** – the optimizer used for back propagation.
- **loader** – the data loader used to get transformed training samples.
- **metrics** – the dictionary of metrics/consumers to update every iteration.
- **output\_path** – directory where output files should be written, if necessary.

**thelper.train.classif module**

Classification trainer/evaluator implementation module.

**class** `thelper.train.classif.ImageClassifTrainer`(*session\_name, session\_dir, model, task, loaders, config, ckptdata=None*)

Bases: `thelper.train.base.Trainer`

Trainer interface specialized for image classification.

This class implements the abstract functions of `thelper.train.base.Trainer` required to train/evaluate a model for image classification or recognition. It also provides a utility function for fetching i/o packets (images, class labels) from a sample, and that converts those into tensors for forwarding and loss estimation.

See also:

`thelper.train.base.Trainer`

**\_\_init\_\_** (*session\_name, session\_dir, model, task, loaders, config, ckptdata=None*)

Receives session parameters, parses image/label keys from task object, and sets up metrics.

**eval\_epoch** (*model, epoch, dev, loader, metrics, output\_path*)

Evaluates the model using the provided objects.

#### Parameters

- **model** – the model to evaluate that is already uploaded to the target device(s).
- **epoch** – the epoch index we are training for (0-based).
- **dev** – the target device that tensors should be uploaded to.
- **loader** – the data loader used to get transformed valid/test samples.
- **metrics** – the dictionary of metrics/consumers to update every iteration.
- **output\_path** – directory where output files should be written, if necessary.

**supports\_classification = True**

**train\_epoch** (*model, epoch, dev, loss, optimizer, loader, metrics, output\_path*)

Trains the model for a single epoch using the provided objects.

#### Parameters

- **model** – the model to train that is already uploaded to the target device(s).
- **epoch** – the epoch index we are training for (0-based).
- **dev** – the target device that tensors should be uploaded to.
- **loss** – the loss function used to evaluate model fidelity.
- **optimizer** – the optimizer used for back propagation.
- **loader** – the data loader used to get transformed training samples.
- **metrics** – the dictionary of metrics/consumers to update every iteration.
- **output\_path** – directory where output files should be written, if necessary.

## thelper.train.detect module

Object detection trainer/evaluator implementation module.

**class** `thelper.train.detect.ObjDetectTrainer` (*session\_name, session\_dir, model, task, loaders, config, ckptdata=None*)

Bases: `thelper.train.base.Trainer`

Trainer interface specialized for object detection.

This class implements the abstract functions of `thelper.train.base.Trainer` required to train/evaluate a model for object detection (i.e. 2D bounding box regression). It also provides a utility function for fetching

i/o packets (input images, bounding boxes) from a sample, and that converts those into tensors for forwarding and loss estimation.

**See also:**

*thelper.train.base.Trainer*

**\_\_init\_\_** (*session\_name, session\_dir, model, task, loaders, config, ckptdata=None*)  
 Receives session parameters, parses tensor/target keys from task object, and sets up metrics.

**eval\_epoch** (*model, epoch, dev, loader, metrics, output\_path*)  
 Evaluates the model using the provided objects.

#### Parameters

- **model** – the model to evaluate that is already uploaded to the target device(s).
- **epoch** – the epoch index we are training for (0-based).
- **dev** – the target device that tensors should be uploaded to.
- **loader** – the data loader used to get transformed valid/test samples.
- **metrics** – the dictionary of metrics/consumers to update every iteration.
- **output\_path** – directory where output files should be written, if necessary.

**supports\_detection = True**

**train\_epoch** (*model, epoch, dev, loss, optimizer, loader, metrics, output\_path*)  
 Trains the model for a single epoch using the provided objects.

#### Parameters

- **model** – the model to train that is already uploaded to the target device(s).
- **epoch** – the epoch index we are training for (0-based).
- **dev** – the target device that tensors should be uploaded to.
- **loss** – the loss function used to evaluate model fidelity.
- **optimizer** – the optimizer used for back propagation.
- **loader** – the data loader used to get transformed training samples.
- **metrics** – the dictionary of metrics/consumers to update every iteration.
- **output\_path** – directory where output files should be written, if necessary.

## thelper.train.regr module

Regression trainer/evaluator implementation module.

**class** *thelper.train.regr.RegressionTrainer* (*session\_name, session\_dir, model, task, loaders, config, ckptdata=None*)

Bases: *thelper.train.base.Trainer*

Trainer interface specialized for generic (n-dim) regression.

This class implements the abstract functions of *thelper.train.base.Trainer* required to train/evaluate a model for generic regression (i.e. n-dim target value prediction). It also provides a utility function for fetching

i/o packets (input tensors, target values) from a sample, and that converts those into tensors for forwarding and loss estimation.

**See also:**

*thelper.train.base.Trainer*

**\_\_init\_\_** (*session\_name, session\_dir, model, task, loaders, config, ckptdata=None*)  
Receives session parameters, parses tensor/target keys from task object, and sets up metrics.

**eval\_epoch** (*model, epoch, dev, loader, metrics, output\_path*)  
Evaluates the model using the provided objects.

#### Parameters

- **model** – the model to evaluate that is already uploaded to the target device(s).
- **epoch** – the epoch index we are training for (0-based).
- **dev** – the target device that tensors should be uploaded to.
- **loader** – the data loader used to get transformed valid/test samples.
- **metrics** – the dictionary of metrics/consumers to update every iteration.
- **output\_path** – directory where output files should be written, if necessary.

**supports\_regression = True**

**train\_epoch** (*model, epoch, dev, loss, optimizer, loader, metrics, output\_path*)  
Trains the model for a single epoch using the provided objects.

#### Parameters

- **model** – the model to train that is already uploaded to the target device(s).
- **epoch** – the epoch index we are training for (0-based).
- **dev** – the target device that tensors should be uploaded to.
- **loss** – the loss function used to evaluate model fidelity.
- **optimizer** – the optimizer used for back propagation.
- **loader** – the data loader used to get transformed training samples.
- **metrics** – the dictionary of metrics/consumers to update every iteration.
- **output\_path** – directory where output files should be written, if necessary.

## **thelper.train.segm module**

Segmentation trainer/evaluator implementation module.

**class** *thelper.train.segm.ImageSegmTrainer* (*session\_name, session\_dir, model, task, loaders, config, ckptdata=None*)

Bases: *thelper.train.base.Trainer*

Trainer interface specialized for image segmentation.

This class implements the abstract functions of *thelper.train.base.Trainer* required to train/evaluate a model for image segmentation (i.e. pixel-level classification/labeling). It also provides a utility function for



fetching i/o packets (images, class labels) from a sample, and that converts those into tensors for forwarding and loss estimation.

**See also:**

*thelper.train.base.Trainer*

**\_\_init\_\_** (*session\_name, session\_dir, model, task, loaders, config, ckptdata=None*)

Receives session parameters, parses image/label keys from task object, and sets up metrics.

**eval\_epoch** (*model, epoch, dev, loader, metrics, output\_path*)

Evaluates the model using the provided objects.

#### Parameters

- **model** – the model to evaluate that is already uploaded to the target device(s).
- **epoch** – the epoch index we are training for (0-based).
- **dev** – the target device that tensors should be uploaded to.
- **loader** – the data loader used to get transformed valid/test samples.
- **metrics** – the dictionary of metrics/consumers to update every iteration.
- **output\_path** – directory where output files should be written, if necessary.

**supports\_segmentation = True**

**train\_epoch** (*model, epoch, dev, loss, optimizer, loader, metrics, output\_path*)

Trains the model for a single epoch using the provided objects.

#### Parameters

- **model** – the model to train that is already uploaded to the target device(s).
- **epoch** – the epoch index we are training for (0-based).
- **dev** – the target device that tensors should be uploaded to.
- **loss** – the loss function used to evaluate model fidelity.
- **optimizer** – the optimizer used for back propagation.
- **loader** – the data loader used to get transformed training samples.
- **metrics** – the dictionary of metrics/consumers to update every iteration.
- **output\_path** – directory where output files should be written, if necessary.

## thelper.train.utils module

Training/evaluation utilities module.

This module contains utilities and tools used to instantiate training sessions. It also contains the prediction consumer interface used by metrics and loggers to receive iteration data during training. See *thelper.optim.metrics* for more information on metrics.

```
class thelper.train.utils.ClassifLogger(top_k=1,                                conf_threshold=None,
                                         class_names=None,                 target_name=None,
                                         viz_count=0, report_count=None, log_keys=None,
                                         force_softmax=True, format=None)
```

Bases: `thelper.ifaces.PredictionConsumer`, `thelper.ifaces.ClassNamesHandler`,  
`thelper.ifaces.FormatHandler`

Classification output logger.

This class provides a simple logging interface for accumulating and saving the predictions of a classifier. By default, all predictions will be logged. However, a confidence threshold can be set to focus on “hard” samples if necessary. It also optionally offers tensorboardX-compatible output images that can be saved locally or posted to tensorboard for browser-based visualization.

Usage examples inside a session configuration file:

```
# ...
# lists all metrics to instantiate as a dictionary
"metrics": {
    # ...
    # this is the name of the example consumer; it is used for lookup/printing_
    ↪only
    "logger": {
        # this type is used to instantiate the confusion matrix report object
        "type": "thelper.train.utils.ClassifLogger",
        "params": {
            # log the three 'best' predictions for each sample
            "top_k": 3,
            # keep updating a set of 10 samples for visualization via tensorboardX
            "viz_count": 10
        }
    },
    # ...
}
```

### Variables

- **top\_k** – number of ‘best’ predictions to keep for each sample (along with the gt label).
- **conf\_threshold** – threshold used to eliminate uncertain predictions.
- **class\_names** – holds the list of class label names provided by the dataset parser. If it is not provided when the constructor is called, it will be set by the trainer at runtime.
- **target\_name** – name of the targeted label (may be ‘None’ if all classes are used).
- **target\_idx** – index of the targeted label (may be ‘None’ if all classes are used).
- **viz\_count** – number of tensorboardX images to generate and update at each epoch.
- **report\_count** – number of samples to print in reports (use ‘None’ if all samples must be printed).
- **log\_keys** – list of metadata field keys to copy from samples into the log for each prediction.
- **force\_softmax** – specifies whether a softmax operation should be applied to the prediction scores obtained from the trainer.
- **score** – array used to store prediction scores for logging.

- **true** – array used to store groundtruth labels for logging.
- **meta** – array used to store metadata pulled from samples for logging.
- **format** – output format of the produced log (supports: text, CSV)

**\_\_init\_\_** (*top\_k=1, conf\_threshold=None, class\_names=None, target\_name=None, viz\_count=0, report\_count=None, log\_keys=None, force\_softmax=True, format=None*)  
 Receives the logging parameters & the optional class label names used to decorate the log.

**property class\_names**

Returns the list of class names considered “of interest” by the derived class.

**render** ()

Returns an image of predicted outputs as a numpy-compatible RGBA image drawn by pyplot.

**report\_csv** ()

Returns the logged metadata of predicted samples.

The returned object is a print-friendly CSV string that can be consumed directly by tensorboardX. Note that this string might be very long if the dataset is large (i.e. it will contain one line per sample).

**report\_json** ()

**report\_text** ()

Must be implemented by inheriting classes. Default report text representation.

**reset** ()

Toggles a reset of the internal state, emptying storage arrays.

**supports\_classification = True**

**update** (*task, input, pred, target, sample, loss, iter\_idx, max\_iters, epoch\_idx, max\_epochs, output\_path, \*\*kwargs*)

Receives the latest predictions and target values from the training session.

The exact signature of this function should match the one of the callbacks defined in *thelper.train.base.Trainer* and specified by *thelper.typedefs.IterCallbackParams*.

**class** *thelper.train.utils.ClassifReport* (*class\_names=None, sample\_weight=None, digits=4, format=None*)

Bases: *thelper.ifaces.PredictionConsumer, thelper.ifaces.ClassNamesHandler, thelper.ifaces.FormatHandler*

Classification report interface.

This class provides a simple interface to *sklearn.metrics.classification\_report* so that all count-based metrics can be reported at once under a string-based representation.

Usage example inside a session configuration file:

```
# ...
# lists all metrics to instantiate as a dictionary
"metrics": {
    # ...
    # this is the name of the example consumer; it is used for lookup/printing
    ↪only
    "report": {
        # this type is used to instantiate the classification report object
        "type": "thelper.train.utils.ClassifReport",
        # we do not need to provide any parameters to the constructor, defaults
        ↪are fine
        "params": {
```

(continues on next page)

(continued from previous page)

```

        # optional parameter that will indicate output as JSON is desired,
        ↪ plain 'text' otherwise
        "format": "json"
    }
    },
    # ...
}
# ...

```

### Variables

- **class\_names** – holds the list of class label names provided by the dataset parser. If it is not provided when the constructor is called, it will be set by the trainer at runtime.
- **pred** – queue used to store the top-1 (best) predicted class indices at each iteration.
- **format** – output format of the produced log (supports: text, JSON)

**\_\_init\_\_** (*class\_names=None, sample\_weight=None, digits=4, format=None*)

Receives the optional class names and arguments passed to the report generator function.

### Parameters

- **class\_names** – holds the list of class label names provided by the dataset parser. If it is not provided when the constructor is called, it will be set by the trainer at runtime.
- **sample\_weight** – sample weights, forwarded to `sklearn.metrics.classification_report`.
- **digits** – metrics output digit count, forwarded to `sklearn.metrics.classification_report`.
- **format** – output format of the produced log.

**gen\_report** (*as\_dict=False*)

**report\_json** ()

Returns the classification report as a JSON formatted string.

**report\_text** ()

Returns the classification report as a multi-line print-friendly string.

**reset** ()

Toggles a reset of the metric's internal state, emptying queues.

**supports\_classification = True**

**update** (*task, input, pred, target, sample, loss, iter\_idx, max\_iters, epoch\_idx, max\_epochs, output\_path, \*\*kwargs*)

Receives the latest predictions and target values from the training session.

The exact signature of this function should match the one of the callbacks defined in `thelper.train.base.Trainer` and specified by `thelper.typedefs.IterCallbackParams`.

**class** `thelper.train.utils.ConfusionMatrix` (*class\_names=None, draw\_normalized=True*)

Bases: `thelper.ifaces.PredictionConsumer`, `thelper.ifaces.ClassNamesHandler`

Confusion matrix report interface.

This class provides a simple interface to `sklearn.metrics.confusion_matrix` so that a full confusion matrix can be easily reported under a string-based representation. It also offers a tensorboardX-compatible output image that can be saved locally or posted to tensorboard for browser-based visualization.

Usage example inside a session configuration file:

```
# ...
# lists all metrics to instantiate as a dictionary
"metrics": {
    # ...
    # this is the name of the example consumer; it is used for lookup/printing
    ↪only
    "confmat": {
        # this type is used to instantiate the confusion matrix report object
        "type": "thelper.train.utils.ConfusionMatrix",
        # we do not need to provide any parameters to the constructor, defaults
    ↪are fine
        "params": {
            # optional parameter that will indicate output as JSON is desired,
    ↪plain 'text' otherwise
            "format": "json"
        }
    },
    # ...
}
# ...
```

### Variables

- **matrix** – report generator function, called at evaluation time to generate the output string.
- **class\_names** – holds the list of class label names provided by the dataset parser. If it is not provided when the constructor is called, it will be set by the trainer at runtime.
- **draw\_normalized** – defines whether rendered confusion matrices should be normalized or not.
- **pred** – queue used to store the top-1 (best) predicted class indices at each iteration.

**\_\_init\_\_** (*class\_names=None, draw\_normalized=True*)

Receives the optional class label names used to decorate the output string.

### Parameters

- **class\_names** – holds the list of class label names provided by the dataset parser. If it is not provided when the constructor is called, it will be set by the trainer at runtime.
- **draw\_normalized** – defines whether rendered confusion matrices should be normalized or not.

**render** ()

Returns the confusion matrix as a numpy-compatible RGBA image drawn by pyplot.

**report** ()

Returns the confusion matrix as a multi-line print-friendly string.

**reset** ()

Toggles a reset of the metric's internal state, emptying queues.

**supports\_classification = True**

**supports\_segmentation = True**

**update** (*task, input, pred, target, sample, loss, iter\_idx, max\_iters, epoch\_idx, max\_epochs, output\_path, \*\*kwargs*)

Receives the latest predictions and target values from the training session.

The exact signature of this function should match the one of the callbacks defined in `thelper.train.base.Trainer` and specified by `thelper.typedefs.IterCallbackParams`.

```
class thelper.train.utils.DetectLogger (top_k=None,                conf_threshold=None,
                                         iou_threshold=None,  class_names=None,  tar-
                                         get_name=None, viz_count=0, report_count=None,
                                         log_keys=None, format=None)
```

Bases: `thelper.ifaces.PredictionConsumer`, `thelper.ifaces.ClassNamesHandler`,  
`thelper.ifaces.FormatHandler`

Detection output logger.

This class provides a simple logging interface for accumulating and saving the bounding boxes of an object detector. By default, all detections will be logged. However, a confidence threshold can be set to focus on strong predictions if necessary.

Usage examples inside a session configuration file:

```
# ...
# lists all metrics to instantiate as a dictionary
"metrics": {
    # ...
    # this is the name of the example consumer; it is used for lookup/printing_
    ↪only
    "logger": {
        # this type is used to instantiate the confusion matrix report object
        "type": "thelper.train.utils.DetectLogger",
        "params": {
            # (optional) log the three 'best' detections for each target
            "top_k": 3
        }
    },
    # ...
}
```

### Variables

- **top\_k** – number of ‘best’ detections to keep for each target bbox (along with the target label). If omitted, lists all bounding box predictions by the model after applying IoU and confidence thresholds.
- **conf\_threshold** – threshold used to eliminate uncertain predictions (if they support confidence). If confidence is not supported by the model bbox predictions, this parameter is ignored.
- **iou\_threshold** – threshold used to eliminate predictions too far from target (regardless of confidence). If omitted, will ignore only completely non-overlapping predicted bounding boxes ( $IoU = 0$ ). If no target bounding boxes are provided (prediction-only), this parameter is ignored.
- **class\_names** – holds the list of class label names provided by the dataset parser. If it is not provided when the constructor is called, it will be set by the trainer at runtime.
- **target\_name** – name of the targeted label (may be ‘None’ if all classes are used).
- **target\_idx** – index of the targeted label (may be ‘None’ if all classes are used).
- **viz\_count** – number of tensorboardX images to generate and update at each epoch.

- **report\_count** – number of samples to print in reports (use ‘None’ if all samples must be printed).
- **log\_keys** – list of metadata field keys to copy from samples into the log for each prediction.
- **bbox** – array used to store prediction bounding boxes for logging.
- **true** – array used to store groundtruth labels for logging.
- **meta** – array used to store metadata pulled from samples for logging.
- **format** – output format of the produced log (supports: text, CSV, JSON)

**\_\_init\_\_**(*top\_k=None, conf\_threshold=None, iou\_threshold=None, class\_names=None, target\_name=None, viz\_count=0, report\_count=None, log\_keys=None, format=None*)  
 Receives the logging parameters & the optional class label names used to decorate the log.

#### **property class\_names**

Returns the list of class names considered “of interest” by the derived class.

#### **gen\_report**( )

Returns the logged metadata of predicted bounding boxes per sample target in a JSON-like structure.

For every target bounding box, the corresponding *best*-sorted detections are returned. Sample metadata is appended to every corresponding sub-target if any where requested.

If `report_count` was specified, the returned report will be limited to that requested amount of targets.

**See also:**

[\*DetectLogger.group\\_bbox\(\)\*](#) for formatting, sorting and filtering details.

#### **group\_bbox**(*target\_bboxes, detect\_bboxes*)

Groups a sample’s detected bounding boxes with target bounding boxes according to configuration parameters.

Returns a list of detections grouped by target(s) with following format:

```
[
  {
    "target": <associated-target-bbox>,
    "detect": [
      {
        "bbox": <detection-bbox>,
        "iou": <IoU(detect-bbox, target-bbox)>
      },
      ...
    ]
  },
  ...
]
```

The associated target bounding box and *IoU* can be `None` if no target was provided (ie: during inference). In this case, the returned list will have only a single element with all detections associated to it. A single element list can also be returned if only one target was specified for this sample.

When multiple ground truth targets were provided, the returned list will have the same length and ordering as these targets. The associated detected bounding boxes will depend on IoU between target/detection combinations.

All filtering thresholds specified as configuration parameter will be applied for the returned list. Detected bounding boxes will also be sorted by highest confidence (if available) or by highest IoU as fallback.

**render ()**

Returns an image of predicted outputs as a numpy-compatible RGBA image drawn by pyplot.

**report\_csv ()**

Returns the logged metadata of predicted bounding boxes.

The returned object is a print-friendly CSV string.

Note that this string might be very long if the dataset is large or if the model tends to generate a lot of detections. The string will contain at least  $N_{sample} \cdot N_{target}$  lines and each line will have up to  $N_{box}$  detections, unless limited by configuration parameters.

**report\_json ()**

Returns the logged metadata of predicted bounding boxes as a JSON formatted string.

**report\_text ()**

Must be implemented by inheriting classes. Default report text representation.

**reset ()**

Toggles a reset of the internal state, emptying storage arrays.

**supports\_detection = True**

**update (task, input, pred, target, sample, loss, iter\_idx, max\_iters, epoch\_idx, max\_epochs, output\_path, \*\*kwargs)**

Receives the latest predictions and target values from the training session.

The exact signature of this function should match the one of the callbacks defined in `thelper.train.base.Trainer` and specified by `thelper.typedefs.IterCallbackParams`.

**class** `thelper.train.utils.PredictionCallback (callback_func, callback_kwargs=None)`

Bases: `thelper.ifaces.PredictionConsumer`

Callback function wrapper compatible with the consumer interface.

This interface is used to hide user-defined callbacks into the list of prediction consumers given to trainer implementations. The callbacks must always be compatible with the list of arguments defined by `thelper.typedefs.IterCallbackParams`, but may also receive extra arguments defined in advance and passed to the constructor of this class.

#### Variables

- **callback\_func** – user-defined function to call on every update from the trainer.
- **callback\_kwargs** – user-defined extra arguments to provide to the callback function.

**\_\_init\_\_ (callback\_func, callback\_kwargs=None)**

Initialize self. See `help(type(self))` for accurate signature.

**update (\*args, \*\*kwargs)**

Forwards the latest prediction data from the training session to the user callback.

`thelper.train.utils.create_consumers (config)`

Instantiates and returns the prediction consumers defined in the configuration dictionary.

All arguments are expected to be handed in through the configuration via a dictionary named ‘params’.



`thelper.train.utils.create_trainer` (*session\_name*, *save\_dir*, *config*, *model*, *task*, *loaders*, *ckpt-data=None*)

Instantiates the trainer object based on the type contained in the config dictionary.

The trainer type is expected to be in the configuration dictionary's *trainer* field, under the *type* key. For more information on the configuration, refer to `thelper.train.base.Trainer`. The instantiated type must be compatible with the constructor signature of `thelper.train.base.Trainer`. The object's constructor will be given the full config dictionary and the checkpoint data for resuming the session (if available).

If the trainer type is missing, it will be automatically deduced based on the task object.

#### Parameters

- **session\_name** – name of the training session used for printing and to create internal tensor-boardX directories.
- **save\_dir** – path to the session directory where logs and checkpoints will be saved.
- **config** – full configuration dictionary that will be parsed for trainer parameters and saved in checkpoints.
- **model** – model to train/evaluate; should be compatible with `thelper.nn.utils.Module`.
- **task** – global task interface defining the type of model and training goal for the session.
- **loaders** – a tuple containing the training/validation/test data loaders (a loader can be `None` if empty).
- **ckptdata** – raw checkpoint to parse data from when resuming a session (if `None`, will start from scratch).

**Returns** The fully-constructed trainer object, ready to begin model training/evaluation.

See also:

`thelper.train.base.Trainer`

## 6.1.9 thelper.transforms package

Transformation operations package.

This package contains data transformation classes and wrappers for preprocessing, augmentation, and normalization of data samples.

### Submodules

#### thelper.transforms.composers module

Transformation composers module.

All transforms should aim to be compatible with both numpy arrays and PyTorch tensors. By default, images are processed using `__call__`, meaning that for a given transformation `t`, we apply it via:

```
image_transformed = t(image)
```

All important parameters for an operation should also be passed in the constructor and exposed in the operation's `__repr__` function so that external parsers can discover exactly how to reproduce their behavior. For now, these representations are used for debugging more than anything else.

## thelper.transforms.operations module

Transformation operations module.

All transforms should aim to be compatible with both numpy arrays and PyTorch tensors. By default, images are processed using `__call__`, meaning that for a given transformation `t`, we apply it via:

```
image_transformed = t(image)
```

All important parameters for an operation should also be passed in the constructor and exposed in the operation's `__repr__` function so that external parsers can discover exactly how to reproduce their behavior. For now, these representations are used for debugging more than anything else.

**class** `thelper.transforms.operations.Affine` (*transf*, *out\_size=None*, *flags=None*, *border\_mode=None*, *border\_val=None*)

Bases: `object`

Warpes a given image using an affine matrix via OpenCV and numpy.

This operation is deterministic. The code relies on OpenCV, meaning the border arguments must be compatible with `cv2.warpAffine`.

### Variables

- **transf** – the 2x3 transformation matrix passed to `cv2.warpAffine`.
- **out\_size** – target image size (tuple of width, height). If `None`, same as original.
- **flags** – extra warp flags forwarded to `cv2.warpAffine`.
- **border\_mode** – border extrapolation mode forwarded to `cv2.warpAffine`.
- **border\_val** – border constant extrapolation value forwarded to `cv2.warpAffine`.

`__call__` (*sample*)

Warpes a given image using an affine matrix.

**Parameters** *sample* – the image to warp; should be a 2d or 3d numpy array.

**Returns** The warped image.

`__init__` (*transf*, *out\_size=None*, *flags=None*, *border\_mode=None*, *border\_val=None*)

Validates and initializes affine warp parameters.

### Parameters

- **transf** – the 2x3 transformation matrix passed to `cv2.warpAffine`.
- **out\_size** – target image size (tuple of width, height). If `None`, same as original.
- **flags** – extra warp flags forwarded to `cv2.warpAffine`.
- **border\_mode** – border extrapolation mode forwarded to `cv2.warpAffine`.
- **border\_val** – border constant extrapolation value forwarded to `cv2.warpAffine`.

**invert** (*sample*)

Inverts the warp transformation, but only if the output image has not been cropped before.

**class** `thelper.transforms.operations.CenterCrop` (*size*, *border\_type=cv2.BORDER\_CONSTANT*, *bordervalue=0*)

Bases: `object`

Returns a center crop from a given image via OpenCV and numpy.

This operation is deterministic. The code relies on OpenCV, meaning the border arguments must be compatible with `cv2.copyMakeBorder`.

#### Variables

- **size** – the size of the target crop (tuple of width, height).
- **relative** – specifies whether the target crop size is relative to the image size or not.
- **bordertype** – argument forwarded to `cv2.copyMakeBorder`.
- **borderval** – argument forwarded to `cv2.copyMakeBorder`.

`__call__` (*sample*)

Extracts and returns a central crop from the provided image.

**Parameters** **sample** – the image to generate the crop from; should be a 2d or 3d numpy array.

**Returns** The center crop.

`__init__` (*size, bordertype=cv2.BORDER\_CONSTANT, borderval=0*)

Validates and initializes center crop parameters.

#### Parameters

- **size** – size of the target crop, provided as tuple or list. If integer values are used, the size is assumed to be absolute. If floating point values are used, the size is assumed to be relative, and will be determined dynamically for each sample. If a tuple is used, it is assumed to be (width, height).
- **bordertype** – border copy type to use when the image is too small for the required crop size. See `cv2.copyMakeBorder` for more information.
- **borderval** – border value to use when the image is too small for the required crop size. See `cv2.copyMakeBorder` for more information.

**invert** (*sample*)

Specifies that this operation cannot be inverted, as data loss is incurred during image transformation.

**class** `thelper.transforms.operations.Duplicator` (*count, deepcopy=False*)

Bases: `thelper.transforms.operations.NoTransform`

Duplicates and returns a list of copies of the input sample.

This operation is used in data augmentation pipelines that rely on probabilistic or preset transformations. It can produce a fixed number of simple copies or deep copies of the input samples as required.

**Warning:** Since the duplicates will be given directly to the data loader as part of the same minibatch, using too many copies can adversely affect gradient descent for that minibatch. To simply increase the total size of the training set while still allowing a proper shuffling of samples and/or to keep the minibatch size intact, we instead recommend setting the `train_scale` configuration value in the data loader. See `thelper.data.utils.create_loaders()` for more information.

#### Variables

- **count** – number of copies to generate.
- **deepcopy** – specifies whether to deep-copy samples or not.

`__call__` (*sample*)

Generates and returns duplicates of the sample/object.

If a dictionary is provided, its values will be expanded into lists that will contain all duplicates. Otherwise, the duplicates will be returned directly as a list.

**Parameters** **sample** – the sample/object to duplicate.

**Returns** A list of duplicated samples, or a dictionary of duplicate lists.

`__init__` (*count*, *deepcopy=False*)

Validates and initializes duplication parameters.

**Parameters**

- **count** – number of copies to generate.
- **deepcopy** – specifies whether to deep-copy samples or not.

`invert` (*sample*)

Returns the first instance of the list of duplicates.

**class** `thelper.transforms.operations.NoTransform`

Bases: `object`

Used to flag some ops that should not be externally wrapped for sample/key handling.

`__call__` (*sample*)

Identity transform.

`invert` (*sample*)

Identity transform.

**class** `thelper.transforms.operations.NormalizeMinMax` (*min*, *max*,  
*out\_type=numpy.float32*)

Bases: `object`

Normalizes a given image using a set of minimum and maximum values.

The samples will be transformed such that  $s = (s - \min) / (\max - \min)$ .

Note that this operation is also not restricted to images.

**Variables**

- **min** – an array of minimum values to subtract with.
- **max** – an array of maximum values to divide with.
- **out\_type** – the output data type to cast the normalization result to.

`__call__` (*sample*)

Normalizes a given sample.

**Parameters**

- **sample** – the sample to normalize. If given as a PIL image, it
- **will be converted to a numpy array first.**

**Returns** The warped sample, in a numpy array of type `self.out_type`.

`__init__` (*min*, *max*, *out\_type=numpy.float32*)

Validates and initializes normalization parameters.

**Parameters**

- **min** – an array of minimum values to subtract with.
- **max** – an array of maximum values to divide with.
- **out\_type** – the output data type to cast the normalization result to.

**invert** (*sample*)

Inverts the normalization.

**class** `thelper.transforms.operations.NormalizeZeroMeanUnitVar` (*mean*, *std*,  
*out\_type=numpy.float32*)

Bases: `object`

Normalizes a given image using a set of mean and standard deviation parameters.

The samples will be transformed such that  $s = (s - \text{mean}) / \text{std}$ .

This can be used for whitening; see [https://en.wikipedia.org/wiki/Whitening\\_transformation](https://en.wikipedia.org/wiki/Whitening_transformation) for more information. Note that this operation is also not restricted to images.

#### Variables

- **mean** – an array of mean values to subtract from data samples.
- **std** – an array of standard deviation values to divide with.
- **out\_type** – the output data type to cast the normalization result to.

**\_\_call\_\_** (*sample*)

Normalizes a given sample.

#### Parameters

- **sample** – the sample to normalize. If given as a PIL image, it
- **will be converted to a numpy array first.**

**Returns** The warped sample, in a numpy array of type `self.out_type`.

**\_\_init\_\_** (*mean*, *std*, *out\_type=numpy.float32*)

Validates and initializes normalization parameters.

#### Parameters

- **mean** – an array of mean values to subtract from data samples.
- **std** – an array of standard deviation values to divide with.
- **out\_type** – the output data type to cast the normalization result to.

**invert** (*sample*)

Inverts the normalization.

**class** `thelper.transforms.operations.RandomResizedCrop` (*output\_size*, *input\_size=0.08*,  
*1.0*, *ratio=0.75*, *1.33*,  
*probability=1.0*, *ran-*  
*dom\_attempts=10*,  
*min\_roi\_iou=1.0*,  
*flags=cv2.INTER\_LINEAR*)

Bases: `object`

Returns a resized crop of a randomly selected image region.

This operation is stochastic, and thus cannot be inverted. Each time the operation is called, a random check will determine whether a transformation is applied or not. The code relies on OpenCV, meaning the interpolation arguments must be compatible with `cv2.resize`.

#### Variables

- **output\_size** – size of the output crop, provided as a single element (`edge_size`) or as a two-element tuple or list (`[width, height]`). If integer values are used, the size is assumed to be absolute. If floating point values are used (i.e. in `[0,1]`), the output size is

assumed to be relative to the original image size, and will be determined at execution time for each sample. If set to `None`, the crop will not be resized.

- **input\_size** – range of the input region sizes, provided as a pair of elements (`[min_edge_size, max_edge_size]`) or as a pair of tuples or lists (`[[min_width, min_height], [max_width, max_height]]`). If the pair-of-pairs format is used, the `ratio` argument cannot be used. If integer values are used, the ranges are assumed to be absolute. If floating point values are used (i.e. in `[0,1]`), the ranges are assumed to be relative to the original image size, and will be determined at execution time for each sample.
- **ratio** – range of minimum/maximum input region aspect ratios to use. This argument cannot be used if the pair-of-pairs format is used for the `input_size` argument.
- **probability** – the probability that the transformation will be applied when called; if not applied, the returned image will be the original.
- **random\_attempts** – the number of random sampling attempts to try before reverting to center or most-probably-valid crop generation.
- **min\_roi\_iou** – minimum roi intersection over union (IoU) required for accepting a tile (in `[0,1]`).
- **flags** – interpolation flag forwarded to `cv2.resize`.

**\_\_call\_\_** (*image, roi=None, mask=None, bboxes=None*)

Extracts and returns a random (resized) crop from the provided image.

#### Parameters

- **image** – the image to generate the crop from. If given as a 2-element list, it is assumed to contain both the image and the roi (passed through a composer).
- **roi** – the roi to check tile intersections with (may be `None`).
- **mask** – a mask to crop simultaneously with the input image (may be `None`).
- **bboxes** – a list or array of bounding boxes to crop with the input image (may be `None`).

**Returns** The randomly selected and resized crop. If mask and/or bboxes is given, the output will be a dictionary containing the results under the `image`, `mask`, and `bboxes` keys.

**\_\_init\_\_** (*output\_size, input\_size=0.08, 1.0, ratio=0.75, 1.33, probability=1.0, random\_attempts=10, min\_roi\_iou=1.0, flags=cv2.INTER\_LINEAR*)

Validates and initializes center crop parameters.

#### Parameters

- **output\_size** – size of the output crop, provided as a single element (`edge_size`) or as a two-element tuple or list (`[width, height]`). If integer values are used, the size is assumed to be absolute. If floating point values are used (i.e. in `[0,1]`), the output size is assumed to be relative to the original image size, and will be determined at execution time for each sample. If set to `None`, the crop will not be resized.
- **input\_size** – range of the input region sizes, provided as a pair of elements (`[min_edge_size, max_edge_size]`) or as a pair of tuples or lists (`[[min_width, min_height], [max_width, max_height]]`). If the pair-of-pairs format is used, the `ratio` argument cannot be used. If integer values are used, the ranges are assumed to be absolute. If floating point values are used (i.e. in `[0,1]`), the ranges are assumed to be relative to the original image size, and will be determined at execution time for each sample.

- **ratio** – range of minimum/maximum input region aspect ratios to use. This argument cannot be used if the pair-of-pairs format is used for the `input_size` argument.
- **probability** – the probability that the transformation will be applied when called; if not applied, the returned image will be the original.
- **random\_attempts** – the number of random sampling attempts to try before reverting to center or most-probably-valid crop generation.
- **min\_roi\_iou** – minimum roi intersection over union (IoU) required for producing a tile.
- **flags** – interpolation flag forwarded to `cv2.resize`.

**invert** (*image*)

Specifies that this operation cannot be inverted, as data loss is incurred during image transformation.

**set\_seed** (*seed*)

Sets the internal seed to use for stochastic ops.

```
class thelper.transforms.operations.RandomShift (min, max, probability=1.0,  
                                                flags=None, border_mode=None,  
                                                border_val=None)
```

Bases: object

Randomly translates an image in a provided range via OpenCV and numpy.

This operation is stochastic, and thus cannot be inverted. Each time the operation is called, a random check will determine whether a transformation is applied or not. The code relies on OpenCV, meaning the border arguments must be compatible with `cv2.warpAffine`.

#### Variables

- **min** – the minimum pixel shift that can be applied stochastically.
- **max** – the maximum pixel shift that can be applied stochastically.
- **probability** – the probability that the transformation will be applied when called.
- **flags** – extra warp flags forwarded to `cv2.warpAffine`.
- **border\_mode** – border extrapolation mode forwarded to `cv2.warpAffine`.
- **border\_val** – border constant extrapolation value forwarded to `cv2.warpAffine`.

**\_\_call\_\_** (*sample*)

Translates a given image using a predetermined min/max range.

**Parameters** *sample* – the image to translate; should be a 2d or 3d numpy array.

**Returns** The translated image.

**\_\_init\_\_** (*min, max, probability=1.0, flags=None, border\_mode=None, border\_val=None*)

Validates and initializes shift parameters.

#### Parameters

- **min** – the minimum pixel shift that can be applied stochastically.
- **max** – the maximum pixel shift that can be applied stochastically.
- **probability** – the probability that the transformation will be applied when called.
- **flags** – extra warp flags forwarded to `cv2.warpAffine`.
- **border\_mode** – border extrapolation mode forwarded to `cv2.warpAffine`.
- **border\_val** – border constant extrapolation value forwarded to `cv2.warpAffine`.

**invert** (*sample*)

Specifies that this operation cannot be inverted, as it is stochastic, and data loss occurs during transformation.

**set\_seed** (*seed*)

Sets the internal seed to use for stochastic ops.

**class** `thelper.transforms.operations.Resize` (*dsize*, *fx=0*, *fy=0*, *interp=cv2.INTER\_LINEAR*, *buffer=False*)

Bases: `object`

Resizes a given image using OpenCV and numpy.

This operation is deterministic. The code relies on OpenCV, meaning the interpolation arguments must be compatible with `cv2.resize`.

#### Variables

- **interp** – interpolation type to use (forwarded to `cv2.resize`)
- **buffer** – specifies whether a destination buffer should be used to avoid allocations
- **dsize** – target image size (tuple of width, height).
- **dst** – buffer used to avoid reallocations if `self.buffer == True`
- **fx** – argument forwarded to `cv2.resize`.
- **fy** – argument forwarded to `cv2.resize`.

**\_\_call\_\_** (*sample*)

Returns a resized copy of the provided image.

**Parameters** **sample** – the image to resize; should be a 2d or 3d numpy array.

**Returns** The resized image. May be allocated on the spot, or be a pointer to a local buffer.

**\_\_init\_\_** (*dsize*, *fx=0*, *fy=0*, *interp=cv2.INTER\_LINEAR*, *buffer=False*)

Validates and initializes resize parameters.

#### Parameters

- **dsize** – size of the target image, forwarded to `cv2.resize`.
- **fx** – x-scaling factor, forwarded to `cv2.resize`.
- **fy** – y-scaling factor, forwarded to `cv2.resize`.
- **interp** – resize interpolation type, forwarded to `cv2.resize`.
- **buffer** – specifies whether a destination buffer should be used to avoid allocations

**invert** (*sample*)

Specifies that this operation cannot be inverted, as data loss is incurred during image transformation.

**class** `thelper.transforms.operations.SelectChannels` (*channels*)

Bases: `object`

Returns selected channels by indices from an array in `numpy`, `torch.Tensor` or `PIL.Image` format.

This operation does not attempt to interpret the meaning of the channel's content. It only moves them around by indices. It is up to the user to make sure indices make sense for desired result. Input is expected to be encoded as HxWxC and will be returned as numpy array of same format.

Behavior according to provided `channels`:

- single index (`int`): only that channel is extracted to form a single channel array as result



- multiple indices (`list`, `tuple`, `set`): all specified unique channel indices are extracted and placed in the specified order as result
- indices map (`dict` with `int` values): channels at key index are moved to value index values

**Examples::** `out = SelectChannels(0)(img)` # only the first channel is kept, out is HxWx1  
`out = SelectChannels([2,1,0])(img3C)` # out image will have channels in reversed order  
`out = SelectChannels([0,1,3])(img4C)` # out image will drop the channel at index #2  
`out = SelectChannels({3:0, 0:1, 1:2})(img4C)` # out image is HxWx3 with remapped channels using

# <from:to> definitions and drops channel #2

# all of the following are equivalent, some implicit and other explicit with `None`  
`out = SelectChannels([0, 1, 3])(img4C)`  
`out = SelectChannels({0:0, 1:1, 3:3})(img4C)`  
`out = SelectChannels({0:0, 1:1, 2:None, 3:3})(img4C)`

**Variables** `channels` – indices or map of the channels to select from the original array.

**Returns** numpy array of selected channels

`__call__` (*sample*)

Converts and returns an array in numpy format with selected channels.

**Parameters** `sample` – the array to convert; should be a tensor, numpy array, or PIL image.

**Returns** The numpy-converted array with selected channels.

`__init__` (*channels*)

Initializes transformation parameters.

`invert` (*sample*)

Specifies that this operation cannot be inverted. Original data type is unknown and channels can be dropped.

```
class thelper.transforms.operations.Tile (tile_size, tile_overlap=0.0, min_mask_iou=1.0,
                                         offset_overlap=False, border-
                                         type=cv2.BORDER_CONSTANT, borderval=0)
```

Bases: `object`

Returns a list of tiles cut out from a given image.

This operation can perform tiling given an optional mask with a target intersection over union (IoU) score, and with an optional overlap between tiles. The tiling is deterministic and can thus be inverted, but only if a mask is not used, as some image regions may be lost otherwise.

If a mask is used, the first tile position is tested exhaustively by iterating over all input coordinates starting from the top-left corner of the image. Otherwise, the first tile position is set as (0,0). Then, all other tiles are found by offsetting from these coordinates, and testing for IoU with the mask (if needed).

**Variables**

- **tile\_size** – size of the output tiles, provided as a single element (`edge_size`) or as a two-element tuple or list (`[width, height]`). If integer values are used, the size is assumed to be absolute. If floating point values are used (i.e. in `[0,1]`), the output size is assumed to be relative to the original image size, and will be determined at execution time for each image.
- **tile\_overlap** – overlap allowed between two neighboring tiles; should be a ratio in `[0,1]`.
- **min\_mask\_iou** – minimum mask intersection over union (IoU) required for accepting a tile (in `[0,1]`).

- **offset\_overlap** – specifies whether the overlap tiling should be offset outside the image or not.
- **bordertype** – border copy type to use when the image is too small for the required crop size. See `cv2.copyMakeBorder` for more information.
- **borderval** – border value to use when the image is too small for the required crop size. See `cv2.copyMakeBorder` for more information.

**\_\_call\_\_** (*image*, *mask=None*)

Extracts and returns a list of tiles cut out from the given image.

#### Parameters

- **image** – the image to cut into tiles. If given as a 2-element list, it is assumed to contain both the image and the mask (passed through a composer).
- **mask** – the mask to check tile intersections with (may be `None`).

**Returns** A list of tiles (numpy-compatible images).

**\_\_init\_\_** (*tile\_size*, *tile\_overlap=0.0*, *min\_mask\_iou=1.0*, *offset\_overlap=False*, *bordertype=cv2.BORDER\_CONSTANT*, *borderval=0*)

Validates and initializes tiling parameters.

#### Parameters

- **tile\_size** – size of the output tiles, provided as a single element (*edge\_size*) or as a two-element tuple or list (*[width, height]*). If integer values are used, the size is assumed to be absolute. If floating point values are used (i.e. in *[0,1]*), the output size is assumed to be relative to the original image size, and will be determined at execution time for each image.
- **tile\_overlap** – overlap ratio between two consecutive (neighboring) tiles; should be in *[0,1]*.
- **min\_mask\_iou** – minimum mask intersection over union (IoU) required for producing a tile.
- **offset\_overlap** – specifies whether the overlap tiling should be offset outside the image or not.
- **bordertype** – border copy type to use when the image is too small for the required crop size. See `cv2.copyMakeBorder` for more information.
- **borderval** – border value to use when the image is too small for the required crop size. See `cv2.copyMakeBorder` for more information.

**count\_tiles** (*image*, *mask=None*)

Returns the number of tiles that would be cut out from the given image.

#### Parameters

- **image** – the image to cut into tiles. If given as a 2-element list, it is assumed to contain both the image and the mask (passed through a composer).
- **mask** – the mask to check tile intersections with (may be `None`).

**Returns** The number of tiles that would be cut with `thelper.transforms.operations.Tile.__call__()`.

**invert** (*image*, *mask=None*)

Returns the reconstituted image from a list of tiles, or throws if a mask was used.

**class** `thelper.transforms.operations.ToColor`

Bases: `object`

Converts a single-channel image to color (RGB).

This operation is deterministic and reversible. It CANNOT be applied to images with more than one channel (HxWx1). The byte ordering (BGR or RGB) does not matter.

**\_\_call\_\_** (*sample*)

Call self as a function.

**\_\_init\_\_** ()

Does nothing, there's no attribute to store for this operation.

**invert** (*sample*)

Inverts the operation by calling the 'ToGray' operation.

Note that this operation is probably lossy due to OpenCV's grayscale conversion code which uses "0.21 R + 0.72 G + 0.07 B" to compute the luminosity of a pixel.

**class** `thelper.transforms.operations.ToGray`

Bases: `object`

Converts a multi-channel image to grayscale.

This operation is deterministic, but not reversible. It can be applied to images with more than three channels (RGB) — in that case, it will compute their per-pixel mean value. Note that in any case, the last dimension (that corresponds to the channels) will remain and be of size 1.

**\_\_call\_\_** (*sample*)

Call self as a function.

**\_\_init\_\_** ()

Does nothing, there's no attribute to store for this operation.

**invert** (*sample*)

Specifies that this operation cannot be inverted, as data loss occurs during transformation.

**class** `thelper.transforms.operations.ToNumpy` (*reorder\_bgr=False*)

Bases: `object`

Converts and returns an image in numpy format from a `torch.Tensor` or `PIL.Image` format.

This operation is deterministic. The returned image will always be encoded as HxWxC, where if the input has three channels, the ordering might be optionally changed.

**Variables** `reorder_bgr` – specifies whether the channels should be reordered in OpenCV format.

**\_\_call\_\_** (*sample*)

Converts and returns an image in numpy format.

**Parameters** `sample` – the image to convert; should be a tensor, numpy array, or PIL image.

**Returns** The numpy-converted image.

**\_\_init\_\_** (*reorder\_bgr=False*)

Initializes transformation parameters.

**invert** (*sample*)

Specifies that this operation cannot be inverted, as the original data type is unknown.

**class** `thelper.transforms.operations.Transpose` (*axes*)

Bases: `object`

Transposes an image via numpy.

This operation is deterministic.

**Variables**

- **axes** – the axes on which to apply the transpose; forwarded to `numpy.transpose`.
- **axes\_inv** – used to invert the tranpose; also forwarded to `numpy.transpose`.

**\_\_call\_\_** (*sample*)

Transposes a given image.

**Parameters** **sample** – the image to transpose; should be a numpy array.

**Returns** The transposed image.

**\_\_init\_\_** (*axes*)

Validates and initializes tranpose parameters.

**Parameters** **axes** – the axes on which to apply the transpose; forwarded to `numpy.transpose`.

**invert** (*sample*)

Invert-transposes a given image.

**Parameters** **sample** – the image to invert-transpose; should be a numpy array.

**Returns** The invert-transposed image.

**class** `thelper.transforms.operations.Unsqueeze` (*axis*)

Bases: `object`

Expands a dimension in the input array via numpy/PyTorch.

This operation is deterministic.

**Variables** **axis** – the axis on which to apply the expansion.

**\_\_call\_\_** (*sample*)

Expands a dimension in the input array via numpy/PyTorch.

**Parameters** **sample** – the array to expand.

**Returns** The array with an extra dimension.

**\_\_init\_\_** (*axis*)

Validates and initializes tranpose parameters.

**Parameters** **axis** – the axis on which to apply the expansion.

**invert** (*sample*)

Squeezes a dimension in the input array via numpy/PyTorch.

**Parameters** **sample** – the array to squeeze.

**Returns** The array with one less dimension.

## thelper.transforms.utils module

Transformations utilities module.

This module contains utility functions used to instantiate transformation/augmentation ops.

`thelper.transforms.utils.load_augments` (*config*)

Loads a data augmentation pipeline.

An augmentation pipeline is essentially a specialized transformation pipeline that can be appended or prefixed to the base transforms defined for all samples. Augmentations are typically used to diversify the samples within the training set in order to help model generalization. They can also be applied to validation and test samples in order to get multiple responses for the same input so that they can be averaged/concatenated into a single output.

Usage examples inside a session configuration file:

```
# ...
# the 'loaders' field can contain several augmentation pipelines
# (see 'thelper.data.utils.create_loaders' for more information on these_
↳pipelines)
"loaders": {
  # ...
  # the 'train_augments' operations are applied to training samples only
  "train_augments": {
    # specifies whether to apply the augmentations before or after the base_
↳transforms
    "append": false,
    "transforms": [
      {
        # here, we use a single stage, which is actually an augmentor sub-
↳pipeline
        # that is purely probabilistic (i.e. it does not increase input_
↳sample count)
        "operation": "Augmentor.Pipeline",
        "params": {
          # the augmentor pipeline defines two operations: rotations_
↳and flips
          "rotate_random_90": {"probability": 0.75},
          "flip_random": {"probability": 0.75}
        }
      }
    ]
  },
  # ...
}
```

**Parameters** `config` – the configuration dictionary defining the meta parameters as well as the list of transformation operations of the augmentation pipeline.

**Returns** A tuple that consists of a pipeline compatible with the `torchvision.transforms` interfaces, and a bool specifying whether this pipeline should be appended or prefixed to the base transforms.

See also:

`thelper.transforms.wrappers.AugmentorWrapper`

```
thelper.transforms.utils.load_transforms()
thelper.data.utils.create_loaders()
```

`thelper.transforms.utils.load_transforms(stages, avoid_transform_wrapper=False)`

Loads a transformation pipeline from a list of stages.

Each entry in the provided list will be considered a stage in the pipeline. The ordering of the stages is important, as some transformations might not be compatible if taken out of order. The entries must each be dictionaries that define an operation, its parameters, and some meta-parameters (detailed below).

The `operation` field of each stage will be used to dynamically import a specific type of operation to apply. The `params` field of each stage will then be used to pass parameters to the constructor of this operation.

If an operation is identified as `"Augmentor.Pipeline"` or `"albumentations.Compose"`, it will be specially handled. In both case, the `params` field becomes mandatory in the stage dictionary, and it must specify the Augmentor or albumentations pipeline operation names and parameters (as a dictionary). Two additional optional config fields can then be set for Augmentor pipelines: `input_tensor` (bool) which specifies whether the previous stage provides a `torch.Tensor` to the pipeline (default=False); and `output_tensor` (bool) which specifies whether the output of the pipeline should be converted into a tensor (default=False). For albumentations pipelines, two additional fields are also available, namely `bbox_params` (dict) and `keypoint_params` (dict). For more information on these, refer to the documentation of `albumentations.core.composition.Compose`. Finally, when unpacking dictionaries for albumentations pipelines, the keys associated to bounding boxes/masks/keypoints that must be forwarded to the composer can be specified via the `bboxes_key`, `mask_key`, and `keypoints_key` fields.

All operations can also specify which sample components they should be applied to via the `target_key` field. This field can contain a single key (typically a string), or a list of keys. The operation will be applied at runtime to all values which are found in the samples with one of those keys. If no key is provided for an operation, it will be applied to all array-like components of the sample. Finally, all operations can specify a `linked_fate` field (bool) to specify whether the samples provided in lists should all have the same fate or not (default=True).

Usage examples inside a session configuration file:

```
# ...
# the 'loaders' field may contain several transformation pipelines
# (see 'thelper.data.utils.create_loaders' for more information on these_
↪pipelines)
"loaders": {
    # ...
    # the 'base_transforms' operations are applied to all loaded samples
    "base_transforms": [
        {
            "operation": "...",
            "params": {
                ...
            },
            "target_key": [ ... ],
            "linked_fate": ...
        },
        {
            "operation": "...",
            "params": {
                ...
            },
            "target_key": [ ... ],
            "linked_fate": ...
        },
    ],
}
```

(continues on next page)

(continued from previous page)

```

    ...
],
# ...

```

**Parameters** `stages` – a list defining a series of transformations to apply as a single pipeline.

**Returns** A transformation pipeline object compatible with the `torchvision.transforms` interface.

**See also:**

```

thelper.transforms.wrappers.AlbumentationsWrapper
thelper.transforms.wrappers.AugmentorWrapper
thelper.transforms.wrappers.TransformWrapper
thelper.transforms.utils.load_augments()
thelper.data.utils.create_loaders()

```

## thelper.transforms.wrappers module

Transformations wrappers module.

The wrapper classes herein are used to either support inline operations on odd sample types (e.g. lists of images) or for external libraries (e.g. Augmentor).

```

class thelper.transforms.wrappers.AlbumentationsWrapper (transforms,
                                                         bbox_params=None,
                                                         add_targets=None,
                                                         image_key='image',
                                                         bboxes_key='bboxes',
                                                         mask_key='mask',      key-
                                                         points_key='keypoints',
                                                         probability=1.0,
                                                         cvt_kpts_to_bboxes=False,
                                                         linked_fate=False)

```

Bases: `object`

Albumentations pipeline wrapper that allows dictionary unpacking.

See <https://github.com/albu/albumentations> for more information.

### Variables

- **pipeline** – the augmentor pipeline instance to apply to images.
- **image\_key** – the key to fetch images from (when dictionaries are passed in).
- **bboxes\_key** – the key to fetch bounding boxes from (when dictionaries are passed in).
- **mask\_key** – the key to fetch masks from (when dictionaries are passed in).
- **keypoints\_key** – the key to fetch keypoints from (when dictionaries are passed in).
- **cvt\_kpts\_to\_bboxes** – specifies whether keypoints should be converted to bboxes for compatibility.
- **linked\_fate** – specifies whether input list samples should all have the same fate or not.

See also:

`thelper.transforms.utils.load_transforms()`

`__call__` (*sample*, *force\_linked\_fate=False*, *op\_seed=None*)

Transforms a (dict) sample, a single image, or a list of images using the augmentor pipeline.

#### Parameters

- **sample** – the sample or image(s) to transform (can also contain embedded lists/tuples of images).
- **force\_linked\_fate** – override flag for recursive use allowing forced linking of arrays.
- **op\_seed** – seed to set before calling the wrapped operation.

**Returns** The transformed image(s), with the same list/tuple formatting as the input.

`__init__` (*transforms*, *bbox\_params=None*, *add\_targets=None*, *image\_key='image'*,  
*bboxes\_key='bboxes'*, *mask\_key='mask'*, *keypoints\_key='keypoints'*, *probability=1.0*,  
*cvt\_kpts\_to\_bboxes=False*, *linked\_fate=False*)

Receives and stores an augmentor pipeline for later use.

The pipeline itself is instantiated in `thelper.transforms.utils.load_transforms()`.

`set_epoch` (*epoch=0*)

Sets the current epoch number in order to change the behavior of some suboperations.

`set_seed` (*seed*)

Sets the internal seed to use for stochastic ops.

**class** `thelper.transforms.wrappers.AugmentorWrapper` (*pipeline*, *target\_keys=None*,  
*linked\_fate=True*)

Bases: `object`

Augmentor pipeline wrapper that allows pickling and multi-threading.

See <https://github.com/mdbloice/Augmentor> for more information. This wrapper was last updated to work with version 0.2.2 — more recent versions introduced yet unfixed (as of 2018/08) issues on some platforms.

All original transforms are supported here. This wrapper also fixes the list output bug for single-image samples when using operations individually.

#### Variables

- **pipeline** – the augmentor pipeline instance to apply to images.
- **target\_keys** – the sample keys to apply the pipeline to (when dictionaries are passed in).
- **linked\_fate** – specifies whether input list samples should all have the same fate or not.

See also:

`thelper.transforms.utils.load_transforms()`

`__call__` (*sample*, *force\_linked\_fate=False*, *op\_seed=None*, *in\_cvts=None*)

Transforms a (dict) sample, a single image, or a list of images using the augmentor pipeline.

#### Parameters



- **sample** – the sample or image(s) to transform (can also contain embedded lists/tuples of images).
- **force\_linked\_fate** – override flag for recursive use allowing forced linking of arrays.
- **op\_seed** – seed to set before calling the wrapped operation.
- **in\_cvts** – holds the input conversion flag array (for recursive usage).

**Returns** The transformed image(s), with the same list/tuple formatting as the input.

**\_\_init\_\_** (*pipeline*, *target\_keys=None*, *linked\_fate=True*)

Receives and stores an augmentor pipeline for later use.

The pipeline itself is instantiated in `thelper.transforms.utils.load_transforms()`.

**set\_epoch** (*epoch=0*)

Sets the current epoch number in order to change the behavior of some suboperations.

**set\_seed** (*seed*)

Sets the internal seed to use for stochastic ops.

```
class helper.transforms.wrappers.TransformWrapper(operation, params=None, probability=1, convert_pil=False,
target_keys=None, linked_fate=True)
```

Bases: object

Transform wrapper that allows operations on samples, lists, tuples, and single elements.

Can be used to wrap the operations in `thelper.transforms` or in `torchvision.transforms` that only accept array-like objects as input. Will optionally force-convert content to PIL images.

Can also be used to transform a list/tuple of images uniformly based on a shared dice roll, or to ensure that each image is transformed independently.

**Warning:** Stochastic transforms (e.g. `torchvision.transforms.RandomCrop`) will always treat each image in a list differently. If the same operations are to be applied to all images, you should consider using a series non-stochastic operations wrapped inside an instance of `torchvision.transforms.RandomApply`, or simply provide the probability of applying the transforms to this wrapper's constructor.

### Variables

- **operation** – the wrapped operation (callable object or class name string to import).
- **params** – the parameters that are passed to the operation when init'd or called.
- **probability** – the probability that the wrapped operation will be applied.
- **convert\_pil** – specifies whether images should be converted into PIL format or not.
- **target\_keys** – the sample keys to apply the transform to (when dictionaries are passed in).
- **linked\_fate** – specifies whether images given in a list/tuple should have the same fate or not.

**\_\_call\_\_** (*sample*, *force\_linked\_fate=False*, *op\_seed=None*, *in\_cvts=None*)

Transforms a (dict) sample, a single image, or a list of images using a wrapped operation.

### Parameters

- **sample** – the sample or image(s) to transform (can also contain embedded lists/tuples of images).
- **force\_linked\_fate** – override flag for recursive use allowing forced linking of arrays.
- **op\_seed** – seed to set before calling the wrapped operation.
- **in\_cvts** – holds the input conversion flag array (for recursive usage).

**Returns** The transformed image(s), with the same list/tuple formatting as the input.

**\_\_init\_\_**(*operation*, *params=None*, *probability=1*, *convert\_pil=False*, *target\_keys=None*, *linked\_fate=True*)

Receives and stores a torchvision transform operation for later use.

If the operation is given as a string, it is assumed to be a class name and it will be imported. The parameters (if any) will then be given to the constructor of that class. Otherwise, the operation is assumed to be a callable object, and its parameters (if any) will be provided at call-time.

#### Parameters

- **operation** – the wrapped operation (callable object or class name string to import).
- **params** – the parameters that are passed to the operation when init'd or called.
- **probability** – the probability that the wrapped operation will be applied.
- **convert\_pil** – specifies whether images should be forced into PIL format or not.
- **target\_keys** – the sample keys to apply the pipeline to (when dictionaries are passed in).
- **linked\_fate** – specifies whether images given in a list/tuple should have the same fate or not.

**set\_epoch**(*epoch=0*)

Sets the current epoch number in order to change the behavior of some suboperations.

**set\_seed**(*seed*)

Sets the internal seed to use for stochastic ops.

### 6.1.10 thelper.viz package

Visualization package.

In contrast with `thelper.draw`, this package regroups utilities and tools used to create visualizations that can be used to debug and understand the behavior of models. For example, it contains a t-SNE module that can create a projection of high-dimensional embeddings created by a model, and different modules to visualize the image regions that cause certain activations inside a model. All these techniques are used to create logs and/or images which in turn can be displayed using the `thelper.draw` module.

`thelper.viz.visualize`(*model*, *task*, *loader*, *viz\_type*, *\*\*kwargs*)

Dispatches a visualization call to the proper package module.

## Submodules

### thelper.viz.tsne module

Tools related to the t-Distributed Stochastic Neighbor Embedding (t-SNE, or TSNE).

For more information on t-SNE, see <https://lvdmaaten.github.io/tsne/> for the original author's repository, or <https://scikit-learn.org/stable/modules/generated/sklearn.manifold.TSNE.html> for the `scikit-learn` tools.

`thelper.viz.tsne.plot` (*projs, targets, preds, color\_map=None, task=None, \*\*kwargs*)

Returns a matplotlib figure of a set of projected embeddings with optional target/predicted labels.

`thelper.viz.tsne.visualize` (*model, task, loader, draw=False, color\_map=None, max\_samples=None, return\_meta=False, \*\*kwargs*)

Creates (and optionally displays) a 2D t-SNE visualization of sample embeddings.

By default, all samples from the data loader will be projected using the model and used for the visualization. If the task is related to classification, the prediction and groundtruth labels will be highlighting using various colors.

If the model does not possess a `get_embedding` attribute, its raw output will be used for projections. Otherwise, `get_embedding` will be called.

#### Parameters

- **model** – the model which will be used to produce embeddings.
- **task** – the task object used to decode predictions and color samples (if possible).
- **loader** – the data loader used to get data samples to project.
- **draw** – boolean flag used to toggle internal display call on or off.
- **color\_map** – map of RGB triplets used to color predictions (for classification only).
- **max\_samples** – maximum number of samples to draw from the data loader.
- **return\_meta** – toggles whether sample metadata should be provided as output or not.

**Returns** A dictionary of the visualization result (an RGB image in numpy format), a list of projected embedding coordinates, the labels of the samples, and the predictions of the samples.

### thelper.viz.umap module

Tools related to the Uniform Manifold Approximation and Projection (UMAP).

For more information on UMAP, see <https://github.com/lmcinnes/umap> for the original author's repository.

`thelper.viz.umap.visualize` (*model, task, loader, draw=False, color\_map=None, max\_samples=None, return\_meta=False, \*\*kwargs*)

Creates (and optionally displays) a 2D UMAP visualization of sample embeddings.

By default, all samples from the data loader will be projected using the model and used for the visualization. If the task is related to classification, the prediction and groundtruth labels will be highlighting using various colors.

If the model does not possess a `get_embedding` attribute, its raw output will be used for projections. Otherwise, `get_embedding` will be called.

#### Parameters

- **model** – the model which will be used to produce embeddings.
- **task** – the task object used to decode predictions and color samples (if possible).

- **loader** – the data loader used to get data samples to project.
- **draw** – boolean flag used to toggle internal display call on or off.
- **color\_map** – map of RGB triplets used to color predictions (for classification only).
- **max\_samples** – maximum number of samples to draw from the data loader.
- **return\_meta** – toggles whether sample metadata should be provided as output or not.

**Returns** A dictionary of the visualization result (an RGB image in numpy format), a list of projected embedding coordinates, the labels of the samples, and the predictions of the samples.

## 6.2 Submodules

### 6.3 `thelper.cli` module

Command-line module, for use with a `__main__` entrypoint.

This module contains the primary functions used to create or resume a training session, to start a visualization session, or to start an annotation session. The basic argument that needs to be provided by the user to create any kind of session is a configuration dictionary. For sessions that produce outputs, the path to a directory where to save the data is also needed.

`thelper.cli.annotate_data(config, save_dir)`

Launches an annotation session for a dataset using a specialized GUI tool.

Note that the annotation type must be supported by the GUI tool. The annotations created by the user during the session will be saved in the session directory.

#### Parameters

- **config** – a dictionary that provides all required dataset and GUI tool configuration parameters; see `thelper.data.utils.create_parsers()` and `thelper.gui.utils.create_annotator()` for more information.
- **save\_dir** – the path to the root directory where the session directory should be saved. Note that this is not the path to the session directory itself, but its parent, which may also contain other session directories.

See also:

```
thelper.gui.annotators.Annotator()
thelper.gui.annotators.ImageSegmentAnnotator()
```

`thelper.cli.create_session(config, save_dir)`

Creates a session to train a model.

All generated outputs (model checkpoints and logs) will be saved in a directory named after the session (the name itself is specified in `config`), and located in `save_dir`.

#### Parameters

- **config** – a dictionary that provides all required data configuration and trainer parameters; see `thelper.train.base.Trainer` and `thelper.data.utils.create_loaders()` for more information. Here, it is only expected to contain a name field that specifies the name of the session.

- **save\_dir** – the path to the root directory where the session directory should be saved. Note that this is not the path to the session directory itself, but its parent, which may also contain other session directories.

See also:

`thelper.train.base.Trainer`

`thelper.cli.export_model(config, save_dir)`

Launches a model exportation session.

This function will export a model defined via a configuration file into a new checkpoint that can be loaded elsewhere. The model can be built using the framework, or provided via its type, construction parameters, and weights. Its exported format will be compatible with the framework, and may also be an optimized/compiled version obtained using PyTorch's JIT tracer.

The configuration dictionary must minimally contain a 'model' section that provides details on the model to be exported. A section named 'export' can be used to provide settings regarding the exportation approaches to use, and the task interface to save with the model. If a task is not explicitly defined in the 'export' section, the session configuration will be parsed for a 'datasets' section that can be used to define it. Otherwise, it must be provided through the model.

The exported checkpoint containing the model will be saved in the session's output directory.

#### Parameters

- **config** – a dictionary that provides all required data configuration parameters; see `thelper.nn.utils.create_model()` for more information.
- **save\_dir** – the path to the root directory where the session directory should be saved. Note that this is not the path to the session directory itself, but its parent, which may also contain other session directories.

See also:

`thelper.nn.utils.create_model()`

`thelper.cli.inference_session(config, save_dir=None, ckpt_path=None)`

Executes an inference session on samples with a trained model checkpoint.

In order to run inference, a model is mandatory and therefore expected to be provided in the configuration. Similarly, a list of input sample file paths are expected for which to run inference on. These inputs can provide additional data according to how they are being parsed by lower level operations.

The session will save the current configuration as `config-infer.json` and the employed model's training configuration as `config-train.json`. Other outputs depend on the specific implementation of the session runner.

#### Parameters

- **config** – a dictionary that provides all required data configuration.
- **save\_dir** – the path to the root directory where the session directory should be saved. Note that this is not the path to the session directory itself, but its parent, which may also contain other session directories. If not provided, will use the best value extracted from either the configuration path or the configuration dictionary itself.

- **ckpt\_path** – explicit checkpoint path to use for loading a model to execute inference. Otherwise look for a model definition in the configuration.

See also:

```
thelper.data.geo.utils.prepare_raster_metadata()  
thelper.data.geo.utils.sliding_window_inference()
```

`thelper.cli.main` (*args=None, argparser=None*)

Main entrypoint to use with console applications.

This function parses command line arguments and dispatches the execution based on the selected operating mode. Run with `--help` for information on the available arguments.

**Warning:** If you are trying to resume a session that was previously executed using a now unavailable GPU, you will have to force the checkpoint data to be loaded on CPU using `--map-location=cpu` (or using `-m=cpu`).

See also:

```
thelper.cli.create_session()  
thelper.cli.resume_session()  
thelper.cli.visualize_data()  
thelper.cli.annotate_data()  
thelper.cli.split_data()  
thelper.cli.inference_session()
```

`thelper.cli.make_argparser` ()

Creates the (default) argument parser to use for the main entrypoint.

The argument parser will contain different “operating modes” that dictate the high-level behavior of the CLI. This function may be modified in branches of the framework to add project-specific features.

`thelper.cli.resume_session` (*ckptdata, save\_dir, config=None, eval\_only=False, task\_compat=None*)

Resumes a previously created training session.

Since the saved checkpoints contain the original session’s configuration, the `config` argument can be set to `None` if the session should simply pick up where it was interrupted. Otherwise, the `config` argument can be set to a new configuration that will override the older one. This is useful when fine-tuning a model, or when testing on a new dataset.

**Warning:** If a session is resumed with an overriding configuration, the user must make sure that the inputs/outputs of the older model are compatible with the new parameters. For example, with classifiers, this means that the number of classes parsed by the dataset (and thus to be predicted by the model) should remain the same. This is a limitation of the framework that should be addressed in a future update.

**Warning:** A resumed session will not be compatible with its original RNG states if the number of workers used is changed. To get 100% reproducible results, make sure you run with the same worker count.

### Parameters

- **ckptdata** – raw checkpoint data loaded via `torch.load()`; it will be parsed by the various parts of the framework that need to reload their previous state.
- **save\_dir** – the path to the root directory where the session directory should be saved. Note that this is not the path to the session directory itself, but its parent, which may also contain other session directories.
- **config** – a dictionary that provides all required data configuration and trainer parameters; see `thelper.train.base.Trainer` and `thelper.data.utils.create_loaders()` for more information. Here, it is only expected to contain a `name` field that specifies the name of the session.
- **eval\_only** – specifies whether training should be resumed or the model should only be evaluated.
- **task\_compat** – specifies how to handle discrepancy between old task from checkpoint and new task from config

See also:

`thelper.train.base.Trainer`

`thelper.cli.setup(args=None, argparser=None)`

Sets up the argument parser (if not already done externally) and parses the input CLI arguments.

This function may return an error code (integer) if the program should exit immediately. Otherwise, it will return the parsed arguments to use in order to redirect the execution flow of the entrypoint.

`thelper.cli.split_data(config, save_dir)`

Launches a dataset splitting session.

This mode will generate an HDF5 archive that contains the split datasets defined in the session configuration file. This archive can then be reused in a new training session to guarantee a fixed distribution of training, validation, and testing samples. It can also be used outside the framework in order to reproduce an experiment.

The configuration dictionary must minimally contain two sections: ‘datasets’ and ‘loaders’. A third section, ‘split’, can be used to provide settings regarding the archive packing and compression approaches to use.

The HDF5 archive will be saved in the session’s output directory.

### Parameters

- **config** – a dictionary that provides all required data configuration parameters; see `thelper.data.utils.create_loaders()` for more information.
- **save\_dir** – the path to the root directory where the session directory should be saved. Note that this is not the path to the session directory itself, but its parent, which may also contain other session directories.

See also:

```
thelper.data.utils.create_loaders()
thelper.data.utils.create_hdf5()
thelper.data.parsers.HDF5Dataset
```

`thelper.cli.visualize_data` (*config*)

Displays the images used in a training session.

This mode does not generate any output, and is only used to visualize the (transformed) images used in a training session. This is useful to debug the data augmentation and base transformation pipelines and make sure the modified images are valid. It does not attempt to load a model or instantiate a trainer, meaning the related fields are not required inside `config`.

If the configuration dictionary includes a ‘loaders’ field, it will be parsed and used. Otherwise, if only a ‘datasets’ field is available, basic loaders will be instantiated to load the data. The ‘loaders’ field can also be ignored if ‘ignore\_loaders’ is found within the ‘viz’ section of the config and set to `True`. Each minibatch will be displayed via `pyplot` or `OpenCV`. The display will block and wait for user input, unless ‘block’ is set within the ‘viz’ section’s ‘kwargs’ config as `False`.

**Parameters** `config` – a dictionary that provides all required data configuration parameters; see

`thelper.data.utils.create_loaders()` for more information.

**See also:**

```
thelper.data.utils.create_loaders()
thelper.data.utils.create_parsers()
```

## 6.4 thelper.concepts module

Framework concepts module for high-level interface documenting.

The decorators and type checkers defined in this module help highlight the purpose of classes and functions with respect to high-level ML tasks.

`thelper.concepts.SUPPORT_PREFIX = 'supports_'`

Prefix that is applied before any ‘concept’ decorator.

`thelper.concepts.apply_support` (*func\_or\_cls=None, concept=None*)

Utility decorator that allows marking a function or a class as *supporting* a certain concept.

### Notes

concept support by the function or class is marked only as documentation reference, no strict validation is accomplished to ensure that further underlying requirements are met for the *concept*.

**See also:**

```
thelper.concepts.classification()
thelper.concepts.detection()
thelper.concepts.segmentation()
thelper.concepts.regression()
```



`thelper.concepts.classification` (*func\_or\_cls=None*)

Decorator that allows marking a function or class as *supporting* the image classification task.

Example:

```
@thelper.concepts.classification
class ClassifObject():
    pass

c = ClassifObject()
c.supports_classification
> True

thelper.concepts.supports(c, "classification")
> True
```

`thelper.concepts.detection` (*func\_or\_cls=None*)

Decorator that allows marking a function or class as *supporting* the object detection task.

Example:

```
@thelper.concepts.detection
class DetectObject():
    pass

d = DetectObject()
d.supports_detection
> True

thelper.concepts.supports(d, "detection")
> True
```

`thelper.concepts.regression` (*func\_or\_cls=None*)

Decorator that allows marking a function or class as *supporting* the generic regression task.

Example:

```
@thelper.concepts.regression
class RegrObject():
    pass

r = RegrObject()
r.supports_regression
> True

thelper.concepts.supports(r, "regression")
> True
```

`thelper.concepts.segmentation` (*func\_or\_cls=None*)

Decorator that allows marking a function or class as *supporting* the image segmentation task.

Example:

```
@thelper.concepts.segmentation
class SegmentObject():
    pass

s = SegmentObject()
s.supports_segmentation
```

(continues on next page)

(continued from previous page)

```
> True

thelper.concepts.supports(s, "segmentation")
> True
```

`thelper.concepts.supports` (*thing*, *concept*)

Utility method to evaluate if *thing supports* a given *concept* as defined by decorators.

#### Parameters

- **thing** – any type, function, method, class or object instance to evaluate if it is marked by the concept
- **concept** – concept to check

See also:

```
thelper.concepts.classification()
thelper.concepts.detection()
thelper.concepts.segmentation()
thelper.concepts.regression()
```

## 6.5 thelper.draw module

Drawing/display utilities module.

These functions currently rely on OpenCV and/or matplotlib.

`thelper.draw.apply_color_map` (*image*, *colormap*, *dst=None*)

Applies a color map to an image of 8-bit color indices; works similarly to `cv2.applyColorMap` (v3.3.1).

`thelper.draw.draw` (*task*, *input*, *pred=None*, *target=None*, *block=False*, *ch\_transpose=True*, *flip\_bgr=False*, *redraw=None*, *\*\*kwargs*)

Draws and returns a figure of a model input/predictions/targets using pyplot or OpenCV.

`thelper.draw.draw_bbox` (*image*, *tl*, *br*, *text*, *color*, *box\_thickness=2*, *font\_thickness=1*, *font\_scale=0.4*, *show=False*, *block=False*, *win\_name='bbox'*)

Draws a single bounding box on a given image (used in `thelper.draw.draw_bboxes()`).

`thelper.draw.draw_bboxes` (*images*, *preds=None*, *bboxes=None*, *color\_map=None*, *redraw=None*, *block=False*, *min\_confidence=0.5*, *class\_map=None*, *\*\*kwargs*)

Draws a set of bounding box prediction results on images.

#### Parameters

- **images** – images with first dimension as list index, and other dimensions are each image's content
- **preds** – predicted bounding boxes per image to be displayed, must match images count if provided
- **bboxes** – ground truth (targets) bounding boxes per image to be displayed, must match images count if provided
- **color\_map** – mapping of class-id to color to be applied to drawn bounding boxes on the image
- **redraw** – existing figure and axes to reuse for drawing the new images and bounding boxes

- **block** – indicate whether to block execution until all figures have been closed or not
- **min\_confidence** – ignore display of bounding boxes that have a confidence below this value, if available
- **class\_map** – alternative class-id to class-name mapping to employ for display. This overrides the default class names retrieved from each bounding box’s attributed task. Useful for displaying generic bounding boxes obtained from raw input values without a specific task.
- **kwargs** – other arguments to be passed down to further drawing functions or drawing settings (amongst other settings, box\_thickness, font\_thickness and font\_scale can be provided)

```
thelper.draw.draw_classifs(images, preds=None, labels=None, is_multi_label=False,
                           class_names_map=None, redraw=None, block=False, **kwargs)
```

Draws and returns a set of classification results.

```
thelper.draw.draw_confmat(confmat, class_list, size_inch=5, 5, dpi=160, normalize=False,
                           keep_unset=False, show=False, block=False)
```

Draws and returns an a confusion matrix figure using pyplot.

```
thelper.draw.draw_errbars(labels, min_values, max_values, stddev_values, mean_values, xlabel="",
                           ylabel='Raw Value', show=False, block=False)
```

Draws and returns an error bar histogram figure using pyplot.

```
thelper.draw.draw_histogram(data, bins=50, xlabel="", ylabel='Proportion', show=False,
                             block=False)
```

Draws and returns a histogram figure using pyplot.

```
thelper.draw.draw_images(images, captions=None, redraw=None, show=True, block=False,
                           use_cv2=True, cv2_flip_bgr=True, img_shape=None,
                           max_img_size=None, grid_size_x=None, grid_size_y=None, cap-
                           tion_opts=None, window_name=None)
```

Draws a set of images with optional captions.

```
thelper.draw.draw_pascalvoc_curve(metrics, size_inch=5, 5, dpi=160, show=False,
                                   block=False)
```

Draws and returns a precision-recall curve according to pascalvoc metrics.

```
thelper.draw.draw_popbars(labels, counts, xlabel="", ylabel='Pop. Count', show=False,
                           block=False)
```

Draws and returns a bar histogram figure using pyplot.

```
thelper.draw.draw_predicts(images, preds=None, targets=None, swap_channels=False, re-
                             draw=None, block=False, **kwargs)
```

Draws and returns a set of generic prediction results.

```
thelper.draw.draw_roc_curve(fpr, tpr, labels=None, size_inch=5, 5, dpi=160, show=False,
                             block=False)
```

Draws and returns an ROC curve figure using pyplot.

```
thelper.draw.draw_segments(images, preds=None, masks=None, color_map=None, redraw=None,
                            block=False, segm_threshold=None, target_class=None, tar-
                            get_threshold=None, **kwargs)
```

Draws and returns a set of segmentation results.

```
thelper.draw.fig2array(fig)
```

Transforms a pyplot figure into a numpy-compatible RGB array.

```
thelper.draw.get_bgr_from_hsl(hue, sat, light)
```

Converts a single HSL triplet (0-360 hue, 0-1 sat & lightness) into an 8-bit RGB triplet.

```
thelper.draw.get_displayable_heatmap(array, convert_rgb=True)
```

Returns a ‘displayable’ array that has been min-maxed and mapped to color triplets.

`thelper.draw.get_displayable_image(image, grayscale=False)`

Returns a ‘displayable’ image that has been normalized and padded to three channels.

`thelper.draw.get_label_color_mapping(idx)`

Returns the PASCAL VOC color triplet for a given label index.

`thelper.draw.get_label_html_color_code(idx)`

Returns the PASCAL VOC HTML color code for a given label index.

`thelper.draw.safe_crop(image, tl, br, bordertype=cv2.BORDER_CONSTANT, borderval=0, force_copy=False)`

Safely crops a region from within an image, padding borders if needed.

#### Parameters

- **image** – the image to crop (provided as a numpy array).
- **tl** – a tuple or list specifying the (x,y) coordinates of the top-left crop corner.
- **br** – a tuple or list specifying the (x,y) coordinates of the bottom-right crop corner.
- **bordertype** – border copy type to use when the image is too small for the required crop size. See `cv2.copyMakeBorder` for more information.
- **borderval** – border value to use when the image is too small for the required crop size. See `cv2.copyMakeBorder` for more information.
- **force\_copy** – defines whether to force a copy of the target image region even when it can be avoided.

**Returns** The cropped image.

## 6.6 thelper.ifaces module

Common object interfaces module.

The interfaces defined here are fairly generic and used to eliminate issues related to circular module importation.

**class** `thelper.ifaces.ClassNamesHandler` (*class\_names: Optional[Iterable[AnyStr]] = None, \*args, \*\*kwargs*)

Bases: `abc.ABC`

Generic interface to handle class names operations for inheriting classes.

#### Variables

- **class\_names** – holds the list of class label names.
- **class\_indices** – holds a mapping (dict) of class-names-to-label-indices.

**\_\_init\_\_** (*class\_names: Optional[Iterable[AnyStr]] = None, \*args, \*\*kwargs*) → None

Initializes the class names array, if an object is provided.

**property class\_indices**

Returns the class-name-to-index map used for encoding labels as integers.

**property class\_names**

Returns the list of class names considered “of interest” by the derived class.

**class** `thelper.ifaces.FormatHandler` (*format: Optional[AnyStr] = 'text', \*args, \*\*kwargs*)

Bases: `abc.ABC`

Generic interface to handle format output operations for inheriting classes.

If `format` is specified and matches a supported one (with a matching `report_<format>` method), this method is used to generate the output. Defaults to "text" if not specified or provided value is not found within supported formatting methods.

#### Variables

- **format** – format to be used for producing the report (default: "text")
- **ext** – extension associated with generated format (default: "txt")

**\_\_init\_\_** (*format: Optional[AnyStr] = 'text', \*args, \*\*kwargs*) → None  
Initialize self. See help(type(self)) for accurate signature.

**report** (*format: Optional[AnyStr] = None*) → Optional[AnyStr]  
Returns the report as a print-friendly string, matching the specified format if specified in configuration.

**Parameters** **format** – format to be used for producing the report (default: initialization attribute or "text" if invalid)

**abstract report\_text** () → Optional[AnyStr]  
Must be implemented by inheriting classes. Default report text representation.

**solve\_format** (*format: Optional[AnyStr] = None*) → None

**class** `thelper.ifaces.PredictionConsumer`  
Bases: `abc.ABC`

Abstract model prediction consumer class.

This interface defines basic functions required so that `thelper.train.base.Trainer` can figure out how to instantiate and update a model prediction consumer. The most notable class derived from this interface is `thelper.optim.metrics.Metric` which is used to monitor the improvement of a model during a training session. Other prediction consumers defined in `thelper.train.utils` will instead log predictions to local files, create graphs, etc.

**reset** () → None  
Resets the internal state of the consumer.

May be called for example by the trainer between two evaluation epochs. The default implementation does nothing, and if a reset behavior is needed, it should be implemented by the derived class.

**abstract update** (*task: `thelper.tasks.utils.Task`, input: `torch.Tensor`, pred: `Union[torch.Tensor, torch.Tensor, List[List[thelper.tasks.detect.BoundingBox]]], torch.Tensor`], target: `Union[torch.Tensor, torch.Tensor, List[List[thelper.tasks.detect.BoundingBox]]], torch.Tensor`, sample: `Dict[Union[AnyStr, int], Any]`, loss: `Optional[float]`, iter\_idx: `int`, max\_iters: `int`, epoch\_idx: `int`, max\_epochs: `int`, output\_path: `AnyStr`, \*\*kwargs*) → None

Receives the latest prediction and groundtruth tensors from the training session.

The data given here will be "consumed" internally, but it should NOT be modified. For example, a classification accuracy metric would accumulate the correct number of predictions in comparison to groundtruth labels, while a plotting logger would add new corresponding dots to a curve.

Remember that input, prediction, and target tensors received here will all have a batch dimension!

The exact signature of this function should match the one of the callbacks defined in `thelper.train.base.Trainer` and specified by `thelper.typedefs.IterCallbackParams`.

## 6.7 thelper.typedefs module

Typing definitions for thelper.

## 6.8 thelper.utils module

General utilities module.

This module only contains non-ML specific functions, i/o helpers, and matplotlib/pyplot drawing calls.

**class** thelper.utils.**Struct** (\*\*kwargs)

Bases: object

Generic runtime-defined C-like data structure (maps constructor elements to fields).

**\_\_init\_\_** (\*\*kwargs)

Initialize self. See help(type(self)) for accurate signature.

thelper.utils.**check\_func\_signature** (func, params)

Checks whether the signature of a function matches the expected parameter list.

thelper.utils.**check\_installed** (package\_name)

Attempts to import a specified package by name, returning a boolean indicating success.

thelper.utils.**check\_version** (version\_check, version\_required)

Verifies that the checked version is not greater than the required one (ie: not a future version).

Version format is MAJOR [ .MINOR [ .PATCH [ [-] <RELEASE> ] ] ] .

Note that for RELEASE part, comparison depends on alphabetical order if all other previous parts were equal (i.e.: alpha will be lower than beta, which in turn is lower than rc and so on). The - is optional and will be removed for comparison (i.e.: 0.5.0-rc is exactly the same as 0.5.0rc and the additional - will not result in evaluating 0.5.0a0 as a greater version because of - being lower ascii than a).

### Parameters

- **version\_check** – the version string that needs to be verified and compared for lower than the required version.
- **version\_required** – the control version against which the check is done.

**Returns** Tuple of the validated check, and lists of both parsed version parts as [MAJOR, MINOR, PATCH, 'RELEASE']. The returned lists are *guaranteed* to be formed of 4 elements, adding 0 or '' as applicable for missing parts.

thelper.utils.**clipstr** (s, size, fill='')

Clips a string to a specific length, with an optional fill character.

thelper.utils.**create\_hdf5\_dataset** (fd, name, max\_len, batch\_like, compression='chunk\_lz4',  
chunk\_size=None, flatten=True)

Creates an HDF5 dataset inside the provided HDF5.File object descriptor.

thelper.utils.**decode\_data** (data, approach='lz4', \*\*kwargs)

Decodes a binary array using a given coding approach.

### Parameters

- **data** – the binary array to decode.
- **approach** – the encoding; supports *none*, *lz4*, *jpg*, *png*.

See also:

```
thelper.utils.encode_data()
```

```
thelper.utils.download_file(url, root, filename, md5=None)
```

Downloads a file from a given URL to a local destination.

#### Parameters

- **url** – path to query for the file (query will be based on urllib).
- **root** – destination folder where the file should be saved.
- **filename** – destination name for the file.
- **md5** – optional, for md5 integrity check.

**Returns** The path to the downloaded file.

```
thelper.utils.encode_data(data, approach='lz4', **kwargs)
```

Encodes a numpy array using a given coding approach.

#### Parameters

- **data** – the numpy array to encode.
- **approach** – the encoding; supports *none*, *lz4*, *jpg*, *png*.

See also:

```
thelper.utils.decode_data()
```

```
thelper.utils.extract_tar(filepath, root, flags='r:gz')
```

Extracts the content of a tar file to a specific location.

#### Parameters

- **filepath** – location of the tar archive.
- **root** – where to extract the archive's content.
- **flags** – extra flags passed to `tarfile.open`.

```
thelper.utils.fetch_hdf5_sample(dset, idx, dtype='auto', shape='auto', compression='auto',
                                **decompr_kwargs)
```

Returns a sample from the specified HDF5 dataset object.

```
thelper.utils.fill_hdf5_sample(dset, dset_idx, array_idx, array, compression='chunk_lz4',
                                **compr_kwargs)
```

Fills a sample inside the specified HDF5 dataset object.

```
thelper.utils.get_available_cuda_devices(attempts_per_device=5)
```

Tests all visible cuda devices and returns a list of available ones.

**Returns** List of available cuda device IDs (integers). An empty list means no cuda device is available, and the app should fallback to cpu.

```
thelper.utils.get_caller_name(skip=2, base_class=False)
```

Returns the name of a caller in the format `module.class.method`.

#### Parameters

- **skip** – specifies how many levels of stack to skip while getting the caller.

- **base\_class** – specified if the base class should be returned or the top-most class in case of inheritance. If the caller is not a class, this doesn't do anything.

**Returns** An empty string is returned if skipped levels exceed stack height; otherwise, returns the requested caller name.

`thelper.utils.get_checkpoint_session_root(ckpt_path)`

Returns the session root directory associated with a checkpoint path.

The given path can point to a checkpoint file or to a directory that contains checkpoints. The returned output directory will be the top-level of the session that created the checkpoint, or None if it cannot be deduced.

**Parameters** **ckpt\_path** – the path to a checkpoint or to an existing directory that contains checkpoints.

**Returns** The path to the session root directory. Will always point to an existing directory, or be None.

`thelper.utils.get_class_logger(skip=0, base=False)`

Shorthand to get logger for current class frame.

`thelper.utils.get_config_output_root(config)`

Returns the output root directory as defined inside a configuration dictionary.

The current implementation will scan for multiple keywords and return the first value found. If no keyword is matched, the function will return None.

**Parameters** **config** – the configuration dictionary to parse for a root output directory.

**Returns** The path to the output root directory. Can point to a non-existing directory, or be None.

`thelper.utils.get_config_session_name(config)`

Returns the 'name' of a session as defined inside a configuration dictionary.

The current implementation will scan for multiple keywords and return the first value found. If no keyword is matched, the function will return None.

**Parameters** **config** – the configuration dictionary to parse for a name.

**Returns** The name that should be given to the session (or 'None' if unknown/unavailable).

`thelper.utils.get_env_list()`

Returns a list of all packages installed in the current environment.

If the required packages cannot be imported, the returned list will be empty. Note that some packages may not be properly detected by this approach, and it is pretty hacky, so use it with a grain of salt (i.e. logging is fine).

`thelper.utils.get_file_paths(input_path, data_root, allow_glob=False, can_be_dir=False)`

Parse a wildcard-compatible file name pattern at a given root level for valid file paths.

`thelper.utils.get_func_logger(skip=0)`

Shorthand to get logger for current function frame.

`thelper.utils.get_git_stamp()`

Returns a print-friendly SHA signature for the framework's underlying git repository (if found).

`thelper.utils.get_glob_paths(input_glob_pattern, can_be_dir=False)`

Parse a wildcard-compatible file name pattern for valid file paths.

`thelper.utils.get_key(key, config, msg=None, delete=False)`

Returns a value given a dictionary key, throwing if not available.

`thelper.utils.get_key_def(key, config, default=None, msg=None, delete=False)`

Returns a value given a dictionary key, or the default value if it cannot be found.



`thelper.utils.get_log_stamp()`

Returns a print-friendly and filename-friendly identification string containing platform and time.

`thelper.utils.get_params_hash(*args, **kwargs)`

Returns a sha1 hash for the given list of parameters (useful for caching).

`thelper.utils.get_save_dir(out_root, dir_name, config=None, resume=False, backup_ext='.json')`

Returns a directory path in which the app can save its data.

If a folder with name `dir_name` already exists in the directory `out_root`, then the user will be asked to pick a new name. If the user refuses, `sys.exit(1)` is called. If `config` is not `None`, it will be saved to the output directory as a json file. Finally, a `logs` directory will also be created in the output directory for writing logger files.

#### Parameters

- **out\_root** – path to the directory root where the save directory should be created.
- **dir\_name** – name of the save directory to create. If it already exists, a new one will be requested.
- **config** – dictionary of app configuration parameters. Used to overwrite i/o queries, and will be written to the save directory in json format to test writing. Default is `None`.
- **resume** – specifies whether this session is new, or resumed from an older one (in the latter case, overwriting is allowed, and the user will never have to choose a new folder)
- **backup\_ext** – extension to use when creating configuration file backups.

**Returns** The path to the created save directory for this session.

`thelper.utils.get_slurm_tmpdir() → str`

Returns the local SLURM\_TMPDIR path if available, or `None`.

`thelper.utils.import_class(fullname)`

General-purpose runtime class importer.

#### Supported syntax:

1. `module.package.Class` will import the fully qualified `Class` located in `package` from the *installed* module
2. `/some/path/mod.pkg.Cls` will import `Cls` as fully qualified `mod.pkg.Cls` from `/some/path` directory

**Parameters** `fullname` – the fully qualified class name to be imported.

**Returns** The imported class.

`thelper.utils.import_function(func, params=None)`

General-purpose runtime function importer, with support for parameter binding.

#### Parameters

- **func** – the fully qualified function name to be imported, or a dictionary with two members (a `type` and optional `params`), or a list of any of these.
- **params** – optional params dictionary to bind to the function call via `functools`. If a dictionary of parameters is also provided in `func`, both will be merged.

**Returns** The imported function, with optionally bound parameters.

`thelper.utils.init_logger(log_level=0, filename=None, force_stdout=False)`

Initializes the framework logger with a specific filter level, and optional file output.

`thelper.utils.is_scalar(val)`

Returns whether the input value is a scalar according to numpy and PyTorch.

`thelper.utils.load_checkpoint(ckpt, map_location=None, always_load_latest=False, check_version=True)`

Loads a session checkpoint via PyTorch, check its compatibility, and returns its data.

If the `ckpt` parameter is a path to a valid directory, then that directly will be searched for a checkpoint. If multiple checkpoints are found, the latest will be returned (based on the epoch index in its name). If `always_load_latest` is set to `False` and if a checkpoint named `ckpt.best.pth` is found, it will be returned instead.

#### Parameters

- **ckpt** – a file-like object or a path to the checkpoint file or session directory.
- **map\_location** – a function, string or a dict specifying how to remap storage locations. See `torch.load` for more information.
- **always\_load\_latest** – toggles whether to always try to load the latest checkpoint if a session directory is provided (instead of loading the ‘best’ checkpoint).
- **check\_version** – toggles whether the checkpoint’s version should be checked for compatibility issues, and query the user for how to proceed.

**Returns** Content of the checkpoint (a dictionary).

`thelper.utils.load_config(path, as_json=False, add_name_if_missing=True, **kwargs)`

Loads the configuration dictionary from the provided path.

The type of file that is loaded is based on the extension in the path.

If the loaded configuration dictionary does not contain a ‘name’ field, the name of the file itself will be inserted as a value.

#### Parameters

- **path** – the path specifying which configuration to be loaded. only supported types are loaded unless `as_json` is `True`.
- **as\_json** – specifies if an alternate extension should be considered as JSON format.
- **add\_name\_if\_missing** – specifies whether the file name should be added to the config dictionary if it is missing a ‘name’ field.
- **kwargs** – forwarded to the extension-specific importer.

`thelper.utils.lreplace(string, old_prefix, new_prefix)`

Replaces a single occurrence of `old_prefix` in the given string by `new_prefix`.

`thelper.utils.migrate_checkpoint(ckptdata)`

Migrates the content of an incompatible or outdated checkpoint to the current version of the framework.

This function might not be able to fix all backward compatibility issues (e.g. it cannot fix class interfaces that were changed). Perfect reproducibility of tests cannot be guaranteed either if this migration tool is used.

**Parameters** `ckptdata` – checkpoint data in dictionary form obtained via `thelper.utils.load_checkpoint`. Note that the data contained in this dictionary will be modified in-place.

**Returns** An updated checkpoint dictionary that should be compatible with the current version of the framework.

`thelper.utils.migrate_config(config, cfg_ver_str)`

Migrates the content of an incompatible or outdated configuration to the current version of the framework.

This function might not be able to fix all backward compatibility issues (e.g. it cannot fix class interfaces that were changed). Perfect reproducibility of tests cannot be guaranteed either if this migration tool is used.

#### Parameters

- **config** – session configuration dictionary obtained e.g. by parsing a JSON file. Note that the data contained in this dictionary will be modified in-place.
- **cfg\_ver\_str** – string representing the version for which the configuration was created (e.g. “0.2.0”).

**Returns** An updated configuration dictionary that should be compatible with the current version of the framework.

`thelper.utils.query_string(question, choices=None, default=None, allow_empty=False, bypass=None)`

Asks the user a question and returns the answer (a generic string).

#### Parameters

- **question** – the string that is presented to the user.
- **choices** – a list of predefined choices that the user can pick from. If `None`, then whatever the user types will be accepted.
- **default** – the presumed answer if the user just hits <Enter>. If `None`, then an answer is required to continue.
- **allow\_empty** – defines whether an empty answer should be accepted.
- **bypass** – the returned value if the `bypass_queries` global variable is set to `True`. Can be `None`, in which case the function will throw an exception.

**Returns** The string entered by the user.

`thelper.utils.query_yes_no(question, default=None, bypass=None)`

Asks the user a yes/no question and returns the answer.

#### Parameters

- **question** – the string that is presented to the user.
- **default** – the presumed answer if the user just hits <Enter>. It must be ‘yes’, ‘no’, or `None` (meaning an answer is required).
- **bypass** – the option to select if the `bypass_queries` global variable is set to `True`. Can be `None`, in which case the function will throw an exception.

**Returns** `True` for ‘yes’, or `False` for ‘no’ (or their respective variations).

`thelper.utils.report_orion_results(session_runner: SessionRunner) → None`

Reports the results of a session runner, but only if the config allows it (true by default).

`thelper.utils.reporhook(count, block_size, total_size)`

Report hook used to display a download progression bar when using urllib requests.

`thelper.utils.resolve_import(fullname)`

Class name resolver.

Takes a string corresponding to a module and class fullname to be imported with `thelper.utils.import_class()` and resolves any back compatibility issues related to renamed or moved classes.

**Parameters** **fullname** – the fully qualified class name to be resolved.

**Returns** The resolved class fullname.

`thelper.utils.save_config(config, path, force_convert=True, as_json=False, **kwargs)`

Saves the given session/object configuration dictionary to the provided path.

The type of file that is created is based on the extension specified in the path. If the file cannot hold some of the objects within the configuration, they will be converted to strings before serialization, unless *force\_convert* is set to *False* (in which case the function will raise an exception).

#### Parameters

- **config** – the session/object configuration dictionary to save.
- **path** – the path specifying where to create the output file. The extension used will determine what type of backup to create (e.g. Pickle = .pkl, JSON = .json, YAML = .yaml/.yml). if *as\_json* is *True*, then any specified extension will be preserved bump dumped as JSON.
- **force\_convert** – specifies whether non-serializable types should be converted if necessary.
- **as\_json** – specifies if an alternate extension should be considered as JSON format.
- **kwargs** – forwarded to the extension-specific exporter.

`thelper.utils.save_env_list(path)`

Saves a list of all packages installed in the current environment to a log file.

**Parameters** **path** – the path where the log file should be created.

`thelper.utils.set_matplotlib_agg()`

Sets the matplotlib backend to Agg.

`thelper.utils.setup_cudnn(config)`

Parses the provided config for CUDNN flags and sets up PyTorch accordingly.

`thelper.utils.setup_cv2(config)`

Parses the provided config for OpenCV flags and sets up its global state accordingly.

`thelper.utils.setup_gdal(config)`

Parses the provided config for GDAL flags and sets up its global state accordingly.

`thelper.utils.setup_globals(config)`

Parses the provided config for global flags and sets up the global state accordingly.

`thelper.utils.setup_plt(config)`

Parses the provided config for matplotlib flags and sets up its global state accordingly.

`thelper.utils.setup_sys(config)`

Parses the provided config for PYTHON sys paths and sets up its global state accordingly.

`thelper.utils.str2bool(s)`

Converts a string to a boolean.

If the lower case version of the provided string matches any of 'true', '1', or 'yes', then the function returns True.

`thelper.utils.str2size(input_str)`

Returns a (WIDTH, HEIGHT) integer size tuple from a string formatted as 'WxH'.

`thelper.utils.stringify_confmat(confmat, class_list, hide_zeroes=False, hide_diagonal=False, hide_threshold=None)`

Transforms a confusion matrix array obtained in list or numpy format into a printable string.

`thelper.utils.test_cuda_device_availability(device_idx)`

Tests the availability of a single cuda device and returns its status.

`thelper.utils.to_numpy(array)`

Converts a list or PyTorch tensor to numpy. Does nothing if already a numpy array.

## CONTRIBUTING

Contributions are welcome, and they are greatly appreciated! Every little bit helps, and credit will always be given.

### 7.1 Bug reports

When [reporting a bug](#) please include:

- Your operating system name and version.
- Any details about your local setup that might be helpful in troubleshooting.
- Detailed steps to reproduce the bug.

### 7.2 Documentation improvements

The framework could always use more documentation, whether as part of the official docs, in docstrings, or even on the web in blog posts, articles, and such.

### 7.3 Feature requests and feedback

The best way to send feedback is to file an issue at <https://github.com/plstcharles/thelper/issues>.

If you are proposing a feature:

- Explain in detail how it would work.
- Keep the scope as narrow as possible, to make it easier to implement.
- Remember that this is a volunteer-driven project, and that code contributions are welcome :)

## 7.4 Development

To set up *thelper* for local development:

1. Fork [thelper](#) (look for the “Fork” button).
2. Clone your fork locally:

```
$ git clone git@github.com:your_name_here/thelper.git
```

3. Create a branch for local development:

```
$ git checkout -b name-of-your-bugfix-or-feature
```

Now you can make your changes locally.

4. When you’re done making changes, run all the tests via:

```
$ make test-all
$ make docs
```

5. Commit your changes and push your branch to GitHub:

```
$ git add .
$ git commit -m "Your detailed description of your changes."
$ git push origin name-of-your-bugfix-or-feature
```

6. Submit a pull request through the GitHub website.

### 7.4.1 Pull Request Guidelines

If you need some code review or feedback while you’re developing the code just make the pull request.

For merging, you should:

1. Make sure all tests are passing (run `make test-all`)<sup>1</sup>.
2. Update documentation when there’s new API, functionality etc.
3. Add a note to `CHANGELOG.rst` about the changes.
4. Add yourself to `AUTHORS.rst`.

---

<sup>1</sup> If you don’t have all the necessary python versions available locally you can rely on Travis - it will [run the tests](#) for each change you add in the pull request.  
It will be slower though ...

## AUTHORS

- Pierre-Luc St-Charles - stcharpl<at>crim.ca
- Francis Charette Migneault - francis.charette-migneault<at>crim.ca
- Mario Beaulieu - mario.beaulieu<at>crim.ca
- Justine Boulent - justine.boulent<at>usherbrooke.ca
- Cyril Pecoraro - cyril.pecoraro<at>crim.ca





## MAINTAINER GUIDE

This guide provides an overview of the basic steps required for testing, creating new releases, and deploying new builds of the framework. For installation instructions, refer to the installation guide [\[here\]](#).

As of November 2019 (v0.4.3), the framework's Continuous Integration (CI) pipeline is based primarily on [Travis CI](#) and [Docker Hub](#). Automated conda builds have slowly been getting harder and harder to fix on Travis, and have gradually been phased out in favor of installation from source.

Maintainers are expected to be using an installation from source and running on a platform that supports Makefiles. All commands below are also provided under the assumption that they will be executed from the root directory of the framework.

### 9.1 Testing

The simplest way to test whether local changes have broken something in the framework is to use make:

```
$ make test-all
```

This command will start by running linters (flake8, isort, twine, check-manifest), and then execute all tests and produce a coverage report. To only run the linters, you can use:

```
$ make check
```

The framework tests are based on pytest, and can be executed independently via:

```
$ make test
```

Since Travis CI runs on GPU-less platforms, unit tests and integration tests with mocked device-aware components are preferred. Future regression tests should also keep this limitation in consideration.

Tests can leave some logs and artifacts in the working directory, especially if they are cancelled in the middle of a run. To get rid of these, you can use:

```
$ make clean-test
```

## 9.2 Releases

To tag a commit for a release, you should use `bumpversion`. It is pre-configured to update all framework version references everywhere and automatically create a new commit with the required tag. Bumping the version works by incrementing the patch (v0.0.X), minor (v0.X.0) or major (vX.0.0) integer:

```
$ bumpversion patch
    or
$ bumpversion minor
    or
$ bumpversion major
```

Creating and pushing any tag on GitHub will trigger the deployment phase on Travis CI.

## 9.3 Documentation

Building the documentation can be accomplished by simply calling:

```
$ make docs
```

This will create the documentation pages in HTML format and display them in your browser. The same documentation will be built and deployed on [readthedocs.io](https://readthedocs.io).

## 9.4 Deployment

Travis CI will automatically attempt to deploy the framework after successfully testing a tagged version. The deployment will target PyPI, Docker Hub, and Anaconda. If any of these steps fail, the deployment can be completed manually as specified below.

### 9.4.1 Python Package Index (PyPI)

A source distribution (sdist) and be prepared and uploaded using the following commands:

```
$ python setup.py sdist
$ twine upload dist/* --skip-existing
```

This will allow you to upload to your own project page, or to the [origin](#) (if you have collaborator access).

### 9.4.2 Docker Hub

A docker image can be prepared and uploaded using the following commands:

```
$ docker build -t ${DOCKER_REPO}:${TAG} -t ${DOCKER_REPO}:latest .
$ docker push ${DOCKER_REPO}
```

Again, uploading to the [original project page](#) will require collaborator access, but you can also upload the image to your own private repository.

### 9.4.3 Anaconda

The Anaconda package build process is very long, and requires a lot of disk space. It tends to fail on Travis CI, and must often be completed manually. To do so, you must first configure your conda environment to use custom channels to find proper project dependencies:

```
$ conda config --prepend channels conda-forge
$ conda config --prepend channels albumentations
$ conda config --prepend channels pytorch
```

Then, updating conda itself is never a bad idea:

```
$ conda update -q conda
```

To build and upload a package, you must also install the correct CLI tools:

```
$ conda install conda-build conda-verify anaconda-client
```

Finally, using the meta-config already available inside the project, you can build the package via:

```
$ conda build ci/
```

Instructions will be printed in the terminal regarding how to upload the built packages.



## CHANGELOG

### 10.1 Unreleased (latest)

### 10.2 0.6.2 (2020/10/01)

- Added `SelectChannels` transform operator for inplace sample modifications that might need it.

### 10.3 0.6.1 (2020/07/29)

- Fixed conda package builds for tagged deployments

### 10.4 0.6.0 (2020/07/28)

- Refactored and cleaned up HDF5 data extraction/parsing classes
- Added dataset interfaces for BigEarthNet, Agri-Vis challenge
- Update classification task to allow multi-label classification
- Added activation layer customization for in-framework ResNet archs
- Updated default `move_tensor` behavior to be non-blocking
- Added trainer implementation for auto-encoder-type models
- Added Orion reporting support for hyperparameter explorations
- Added SLURM cluster utilities (tmpdir getter, launch scripts)

### 10.5 0.5.0 (2020/07/21)

- Skip image save call during metric rendering if the provided value is `None` as employed by basic logger/reporter.
- Add JSON implementation for `thelper.train.utils.ClassifLogger`.
- Fix `concepts` to handle any variation of upper/lower concept name.

## 10.6 0.5.0-rc2 (2020/07/08)

- Employ `requirements.txt` within `conda-env.yml` to kept dependencies in sync.
- Fixes built docker image not using appropriate dependencies enforced through `requirements.txt`.

## 10.7 0.5.0-rc1 (2020/07/07)

- Fix version comparison check when validating configuration and/or checkpoint against package version. Version can now have a release part which was not considered.
- Fix incorrect calculation of sample coordinates in `thelper.data.geo.parsers.SlidingWindowDataset`.
- Remove `not_skip = __init__.py` config option for `isort` since `__init__.py` is included since 4.3.5. Also force `isort<5` since many import checks break suddenly (e.g.: direct import with *as* alias break).

## 10.8 0.5.0-rc0 (2020/04/25)

- Update this changelog to use rst links (renders on github and readthedocs)
- Add `infer` mode for classification of geo-referenced rasters
- Add `Dockerfile-geo` to build `thelper` with pre-installed geo packages
- Add geo-related build instructions to travis-ci build steps
- Add auto-documentation of makefile targets and docker related targets

## 10.9 0.4.7 (2019/11/20)

- Removed optional dependencies from conda build env

## 10.10 0.4.6 (2019/11/20)

- Travis deploy test w/ split conda/docker stages

## 10.11 0.4.5 (2019/11/18)

- Split travis deploy stage into two phases
- Fixed `draw_segment` threshold usage & params lookup
- Fixed FCResNet embedding getter wrt latest pooling update
- Update all matplotlib plots to use 160 dpi by default
- Refactor trainer data/metric writer to save all viz data

## 10.12 0.4.4 (2019/11/18)

- Added viz pkg w/ t-SNE & UMAP support for in-trainer usage
- Fixed geo pkg documentation build issue related to mocking
- Fixed type and output format checks in numerous metrics
- Updated all callback readers to rely on new utility function
- Cleaned and optimize coordconv implementation
- Added U-Net architecture implementation to nn package
- Added IoU metric implementation
- Added support for SRM kernels and SRM convolutions
- Updated documentation (install, faq, maintenance)
- Added fixed weight sampler to data package
- Added lots of extra unit tests
- Added efficientnet 3rd-party module wrapper
- Fixed potential conflicts in task class names ordering

## 10.13 0.4.3 (2019/11/06)

- Fixed pytest-mock scope usage in metrics utests

## 10.14 0.4.2 (2019/11/06)

- Updated common resnet impl to support segmentation heads
- Fixed samples usage for auto-weighting of loss functions
- Cleaned up samples usage in loader factory data splitter
- Add GDL compatibility module to geo package
- Fix segmentation task dontcare default color mapping
- Cleaned up and simplified coordconv implementation
- Update segmentation trainer to use long-typed label maps
- Cleaned up augmentor/albumentations demo configurations

## 10.15 0.4.1 (2019/10/15)

- Removed travis check in deploy stage for master branch

## 10.16 0.4.0 (2019/10/11)

- Added geo subpackage
- Added geo vector/raster parsing classes
- Added ogc module for testbed15-specific utilities
- Added testbed15 train/viz configuration files
- Cleaned up makefile targets & coverage usage
- Replaced tox build system with makefile completely
- Merged 3rdparty configs into setup.cfg
- Updated travis to rely on makefile directly

## 10.17 0.3.14 (2019/09/30)

- Added extra logging calls in trainer and framework utils
- Cleaned up data configuration parsing logger calls
- Bypassed full device check when specific one is requested

## 10.18 0.3.13 (2019/09/26)

- Moved drawing utilities to new module
- Cleaned up output root/save directory parsing
- Cleaned up potential circular imports
- Moved optional dependency imports inside relevant functions
- Added support for root directory specification via config
- Updated config load/save to make naming optional

## 10.19 0.3.12 (2019/09/13)

- Fixed potential issue when reinstantiating custom ResNet
- Fixed ClassifLogger prediction logger w/o groundtruth



## 10.20 0.3.11 (2019/09/09)

- Add cli/config override for task compatibility mode setting

## 10.21 0.3.10 (2019/09/05)

- Cleaned up dependency lists, docstrings
- Fixed bbox iou computation with mixed int/float
- Fixed dontcare label deletion in segmentation task
- Cleaned up training session output directory localization
- Fixed object detection trainer empty bbox lists
- Fixed exponential parsing with pyyaml
- Fixed bbox display when using integer coords values

## 10.22 0.3.9 (2019/08/20)

- Fixed collate issues for pytorch  $\geq 1.2$
- Fixed null-size batch issues
- Cleaned up params#kwargs parsing in trainer
- Added pickled hashed param support utils
- Added support for yaml-based session configuration
- Added concept decorators for metrics/consumer classes
- Cleaned up shared interfaces to fix circular dependencies
- Added detection (bbox) logger class

## 10.23 0.3.8 (2019/08/08)

- Fixed nn modules constructor args forwarding
- Updated class importer to allow parsing of non-package dirs
- Fixed file-based logging from submodules (e.g. for all data)
- Cleaned and API-fied the CLI entrypoints for external use

## 10.24 0.3.7 (2019/07/31)

- Fixed travis timeouts on long deploy operations
- Added output path to trainer callback impls
- Added new draw-and-save display callback
- Added togray/tocolor transformation operations
- Cleaned up matplotlib use and show/block across draw functions
- Fixed various dependency and logging issues

## 10.25 0.3.6 (2019/07/26)

- Fixed torch version checks in custom default collate impl
- Fixed bbox predictions forwarding and evaluation in objdetect
- Refactored metrics/callbacks to clean up trainer impls
- Added pretrained opt to default resnet impl
- Fixed objdetect trainer display and prediction callbacks

## 10.26 0.3.5 (2019/07/23)

- Refactored metrics/consumers into separate interfaces
- Added unit tests for all metrics/prediction consumers
- Updated trainer callback signatures to include more data
- Updated install doc with links to anaconda/docker hubs
- Cleaned drawing functions args wrt callback refactoring
- Added eval module to optim w/ pascalvoc evaluation funcs

## 10.27 0.3.4 (2019/07/12)

- Fixed issues when reloading objdet model checkpoints
- Fixed issues when trying to use missing color maps
- Fixed backward compat issues when reloading old tasks
- Cleaned up object detection drawing utilities

## 10.28 0.3.3 (2019/07/09)

- Fixed travis conda build dependencies & channels

## 10.29 0.3.2 (2019/07/05)

- Update documentation use cases (model export) & faq
- Cleanup module base class config backup
- Fixed docker build and automated it via travis

## 10.30 0.3.1 (2019/06/17)

- Fix metrics RawPredictions not returning predictions during eval
- Fix parsing of checkpoint base path

## 10.31 0.3.0 (2019/06/12)

- Added dockerfile for containerized builds
- Added object detection task & trainer implementations
- Added CLI model/checkpoint export support
- Added CLI dataset splitting/HDF5 support
- Added baseline superresolution implementations
- Added lots of new unit tests & docstrings
- Cleaned up transform & display operations

## 10.32 0.2.8 (2019/03/17)

- Cleaned up build tools & docstrings throughout api
- Added user guide in documentation build
- Update tasks to allow dataset interface override
- Cleaned up trainer output logs
- Added fully convolutional resnet implementation
- Fixup various issues related to fine-tuning via 'resume'

### 10.33 0.2.7 (2019/02/04)

- Updated conda build recipe for python variants w/ auto upload

### 10.34 0.2.6 (2019/01/31)

- Added framework checkpoint/configuration migration utilities
- Fixed minor config parsing backward compatibility issues
- Fixed minor bugs related to query & drawing utilities

### 10.35 0.2.5 (2019/01/29)

- Fix travis-ci conda build/env path

### 10.36 0.2.4 (2019/01/29)

- Fix travis-ci conda channel setup

### 10.37 0.2.3 (2019/01/29)

- Fix openssl dependency

### 10.38 0.2.2 (2019/01/29)

- Fixed travis-ci matrix configuration
- Added travis-ci deployment step for pypi
- Fixed readthedocs documentation building
- Updated readme shields & front page look
- Cleaned up cli module entrypt
- Fixed openssl dependency issues for travis tox check jobs
- Updated travis post-deploy to try to fix conda packaging (wip)

## 10.39 0.2.1 (2019/01/24)

- Added typedef module & cleaned up parameter inspections
- Cleaned up all drawing utils & added callback support to trainers
- Added support for albumentation pipelines via wrapper
- Updated all trainers/schedulers to rely on 0-based indexing
- Updated travis/rtd configs for auto-deploy & 3.6 support

## 10.40 0.2.0 (2019/01/15)

- Added regression/segmentation tasks and trainers
- Added interface for pascalvoc dataset
- Refactored data loaders/parsers and cleaned up data package
- Added lots of new utilities in base trainer implementation
- Added new unit tests for transformations
- Refactored transformations to use wrappers for augments/lists
- Added new samplers with dataset scaling support
- Added baseline implementation for FCN32s
- Added mae/mse metrics implementations
- Added trainer support for loss computation via external members
- Added utils to download/verify/extract files

## 10.41 0.1.1 (2019/01/14)

- Minor fixups and updates for CCFB02 compatibility
- Added RawPredictions metric to fetch data from trainers

## 10.42 0.1.0 (2018/11/28)

- Fixed readthedocs sphinx auto-build w/ mocking.
- Refactored package structure to avoid env issues.
- Rewrote seeding to allow 100% reproducible sessions.
- Cleaned up config file parameter lists.
- Cleaned up session output vars/logs/images.
- Add support for eval-time augmentation.
- Update transform wrappers for multi-channels & lists.
- Add gui module w/ basic segmentation annotation tool.

- Refactored task interfaces to allow merging.
- Simplified model fine-tuning via checkpoints.

### **10.43 0.0.2 (2018/10/18)**

- Completed first documentation pass.
- Fixed travis/rtfd builds.
- Fixed device mapping/loading issues.

### **10.44 0.0.1 (2018/10/03)**

- Initial release (work in progress).

## INDICES AND TABLES

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





## PYTHON MODULE INDEX

### t

- `thelper`, 29
- `thelper.cli`, 110
- `thelper.concepts`, 114
- `thelper.data`, 29
- `thelper.data.geo`, 29
  - `thelper.data.geo.agravis`, 29
  - `thelper.data.geo.gdl`, 30
  - `thelper.data.geo.infer`, 30
  - `thelper.data.geo.ogc`, 31
  - `thelper.data.geo.parsers`, 31
  - `thelper.data.geo.utils`, 31
  - `thelper.data.loaders`, 32
  - `thelper.data.parsers`, 34
  - `thelper.data.pascalvoc`, 34
  - `thelper.data.samplers`, 34
  - `thelper.data.utils`, 34
- `thelper.draw`, 116
- `thelper.gui`, 40
  - `thelper.gui.annotators`, 40
  - `thelper.gui.utils`, 42
- `thelper.ifaces`, 118
- `thelper.infer`, 43
  - `thelper.infer.base`, 43
  - `thelper.infer.impl`, 44
  - `thelper.infer.utils`, 44
- `thelper.nn`, 45
  - `thelper.nn.common`, 46
  - `thelper.nn.coordconv`, 46
  - `thelper.nn.densenet`, 46
  - `thelper.nn.efficientnet`, 47
  - `thelper.nn.fcn`, 47
  - `thelper.nn.inceptionresnetv2`, 47
  - `thelper.nn.lenet`, 47
  - `thelper.nn.mobilenet`, 47
  - `thelper.nn.resnet`, 47
  - `thelper.nn.segmentation`, 45
    - `thelper.nn.segmentation.base`, 45
    - `thelper.nn.segmentation.deeplabv3`, 45
    - `thelper.nn.segmentation.fcn`, 45
  - `thelper.nn.sr`, 45
    - `thelper.nn.sr.srcnn`, 46
    - `thelper.nn.sr.vdsr`, 46
  - `thelper.nn.srm`, 47
  - `thelper.nn.unet`, 47
  - `thelper.nn.utils`, 47
- `thelper.optim`, 48
  - `thelper.optim.eval`, 49
  - `thelper.optim.losses`, 50
  - `thelper.optim.metrics`, 50
  - `thelper.optim.schedulers`, 63
  - `thelper.optim.utils`, 63
- `thelper.session`, 64
  - `thelper.session.base`, 64
- `thelper.tasks`, 65
  - `thelper.tasks.classif`, 66
  - `thelper.tasks.detect`, 67
  - `thelper.tasks.regr`, 71
  - `thelper.tasks.segm`, 73
  - `thelper.tasks.utils`, 74
- `thelper.train`, 75
  - `thelper.train.ae`, 76
  - `thelper.train.base`, 76
  - `thelper.train.classif`, 79
  - `thelper.train.detect`, 80
  - `thelper.train.regr`, 81
  - `thelper.train.segm`, 82
  - `thelper.train.utils`, 83
- `thelper.transforms`, 91
  - `thelper.transforms.composers`, 91
  - `thelper.transforms.operations`, 92
  - `thelper.transforms.utils`, 103
  - `thelper.transforms.wrappers`, 105
- `thelper.typedefs`, 120
- `thelper.utils`, 120
- `thelper.viz`, 108
  - `thelper.viz.tsne`, 109
  - `thelper.viz.umap`, 109



## Symbols

<code>__call__()</code>	( <i>thelper.transforms.operations.Affine</i> method), 92	<code>__init__()</code>	( <i>thelper.data.loaders.LoaderFactory</i> method), 32
<code>__call__()</code>	( <i>thelper.transforms.operations.CenterCrop</i> method), 93	<code>__init__()</code>	( <i>thelper.gui.annotators.Annotator</i> method), 40
<code>__call__()</code>	( <i>thelper.transforms.operations.Duplicator</i> method), 93	<code>__init__()</code>	( <i>thelper.gui.annotators.ImageSegmentAnnotator</i> method), 42
<code>__call__()</code>	( <i>thelper.transforms.operations.NoTransform</i> method), 94	<code>__init__()</code>	( <i>thelper.gui.annotators.ImageSegmentAnnotator.Brush</i> method), 41
<code>__call__()</code>	( <i>thelper.transforms.operations.NormalizeMinMax</i> method), 94	<code>__init__()</code>	( <i>thelper.gui.annotators.ImageSegmentAnnotator.ZoomTool</i> method), 42
<code>__call__()</code>	( <i>thelper.transforms.operations.NormalizeZeroMeanUnitVar</i> method), 95	<code>__init__()</code>	( <i>thelper.ifaces.ClassNamesHandler</i> method), 118
<code>__call__()</code>	( <i>thelper.transforms.operations.RandomResizedCrop</i> method), 96	<code>__init__()</code>	( <i>thelper.ifaces.FormatHandler</i> method), 119
<code>__call__()</code>	( <i>thelper.transforms.operations.RandomShift</i> method), 97	<code>__init__()</code>	( <i>thelper.infer.base.Tester</i> method), 43
<code>__call__()</code>	( <i>thelper.transforms.operations.Resize</i> method), 98	<code>__init__()</code>	( <i>thelper.optim.metrics.Accuracy</i> method), 51
<code>__call__()</code>	( <i>thelper.transforms.operations.SelectChannels</i> method), 99	<code>__init__()</code>	( <i>thelper.optim.metrics.AveragePrecision</i> method), 52
<code>__call__()</code>	( <i>thelper.transforms.operations.Tile</i> method), 100	<code>__init__()</code>	( <i>thelper.optim.metrics.ExternalMetric</i> method), 54
<code>__call__()</code>	( <i>thelper.transforms.operations.ToColor</i> method), 101	<code>__init__()</code>	( <i>thelper.optim.metrics.IntersectionOverUnion</i> method), 56
<code>__call__()</code>	( <i>thelper.transforms.operations.ToGray</i> method), 101	<code>__init__()</code>	( <i>thelper.optim.metrics.MeanAbsoluteError</i> method), 57
<code>__call__()</code>	( <i>thelper.transforms.operations.ToNumpy</i> method), 101	<code>__init__()</code>	( <i>thelper.optim.metrics.MeanSquaredError</i> method), 58
<code>__call__()</code>	( <i>thelper.transforms.operations.Transpose</i> method), 102	<code>__init__()</code>	( <i>thelper.optim.metrics.PSNR</i> method), 60
<code>__call__()</code>	( <i>thelper.transforms.operations.Unsqueeze</i> method), 102	<code>__init__()</code>	( <i>thelper.optim.metrics.ROCCurve</i> method), 62
<code>__call__()</code>	( <i>thelper.transforms.wrappers.AlbumentationsWrapper</i> method), 106	<code>__init__()</code>	( <i>thelper.session.base.SessionRunner</i> method), 65
<code>__call__()</code>	( <i>thelper.transforms.wrappers.AugmentorWrapper</i> method), 106	<code>__init__()</code>	( <i>thelper.tasks.classif.Classification</i> method), 66
<code>__call__()</code>	( <i>thelper.transforms.wrappers.TransformWrapper</i> method), 107	<code>__init__()</code>	( <i>thelper.tasks.detect.BoundingBox</i> method), 68
<code>__init__()</code>	( <i>thelper.data.geo.infer.SlidingWindowTester</i> method), 30	<code>__init__()</code>	( <i>thelper.tasks.detect.Detection</i> method), 70
<code>__init__()</code>	( <i>thelper.data.geo.ogc.TB15D104DetectLogger</i> method), 31	<code>__init__()</code>	( <i>thelper.tasks.regr.Regression</i> method), 71
		<code>__init__()</code>	( <i>thelper.tasks.regr.SuperResolution</i> method), 72

`__init__()` (*thelper.tasks.segm.Segmentation method*), 73

`__init__()` (*thelper.tasks.utils.Task method*), 74

`__init__()` (*thelper.train.ae.AutoEncoderTrainer method*), 76

`__init__()` (*thelper.train.base.Trainer method*), 78

`__init__()` (*thelper.train.classif.ImageClassifTrainer method*), 80

`__init__()` (*thelper.train.detect.ObjDetectTrainer method*), 81

`__init__()` (*thelper.train.regr.RegressionTrainer method*), 82

`__init__()` (*thelper.train.segm.ImageSegmTrainer method*), 83

`__init__()` (*thelper.train.utils.ClassifLogger method*), 85

`__init__()` (*thelper.train.utils.ClassifReport method*), 86

`__init__()` (*thelper.train.utils.ConfusionMatrix method*), 87

`__init__()` (*thelper.train.utils.DetectLogger method*), 89

`__init__()` (*thelper.train.utils.PredictionCallback method*), 90

`__init__()` (*thelper.transforms.operations.Affine method*), 92

`__init__()` (*thelper.transforms.operations.CenterCrop method*), 93

`__init__()` (*thelper.transforms.operations.Duplicator method*), 94

`__init__()` (*thelper.transforms.operations.NormalizeMinMax method*), 94

`__init__()` (*thelper.transforms.operations.NormalizeZeroMeanUnitVar method*), 95

`__init__()` (*thelper.transforms.operations.RandomResizeCrop method*), 96

`__init__()` (*thelper.transforms.operations.RandomShift method*), 97

`__init__()` (*thelper.transforms.operations.Resize method*), 98

`__init__()` (*thelper.transforms.operations.SelectChannels method*), 99

`__init__()` (*thelper.transforms.operations.Tile method*), 100

`__init__()` (*thelper.transforms.operations.ToColor method*), 101

`__init__()` (*thelper.transforms.operations.ToGray method*), 101

`__init__()` (*thelper.transforms.operations.ToNumpy method*), 101

`__init__()` (*thelper.transforms.operations.Transpose method*), 102

`__init__()` (*thelper.transforms.operations.Unsqueeze method*), 102

`__init__()` (*thelper.transforms.wrappers.AlbumentationsWrapper method*), 106

`__init__()` (*thelper.transforms.wrappers.AugmentorWrapper method*), 107

`__init__()` (*thelper.transforms.wrappers.TransformWrapper method*), 108

`__init__()` (*thelper.utils.Struct method*), 120

## A

Accuracy (class in *thelper.optim.metrics*), 50

Affine (class in *thelper.transforms.operations*), 92

AlbumentationsWrapper (class in *thelper.transforms.wrappers*), 105

annotate\_data() (in module *thelper.cli*), 110

Annotator (class in *thelper.gui.annotators*), 40

apply\_color\_map() (in module *thelper.draw*), 116

apply\_support() (in module *thelper.concepts*), 114

area() (*thelper.tasks.detect.BoundingBox* property), 68

AugmentorWrapper (class in *thelper.transforms.wrappers*), 106

AutoEncoderTrainer (class in *thelper.train.ae*), 76

AveragePrecision (class in *thelper.optim.metrics*), 51

## B

background() (*thelper.tasks.detect.Detection* property), 70

BACKGROUND\_ID (*thelper.data.geo.ogc.TB15D104* attribute), 31

bbox() (*thelper.tasks.detect.BoundingBox* property), 68

bottom() (*thelper.tasks.detect.BoundingBox* property), 68

bottom\_right() (*thelper.tasks.detect.BoundingBox* property), 68

BoundingBox (class in *thelper.tasks.detect*), 67

BRUSH\_SIZE (*thelper.gui.annotators.ImageSegmentAnnotator* attribute), 41

## C

CenterCrop (class in *thelper.transforms.operations*), 92

centroid() (*thelper.tasks.detect.BoundingBox* property), 68

check\_compat() (*thelper.tasks.classif.Classification* method), 66

check\_compat() (*thelper.tasks.detect.Detection* method), 70

check\_compat() (*thelper.tasks.regr.Regression* method), 71

check\_compat() (*thelper.tasks.segm.Segmentation* method), 73

check\_compat() (*thelper.tasks.utils.Task* method), 74

check\_func\_signature() (in module *thelper.utils*), 120

check\_installed() (in module *thelper.utils*), 120  
 check\_version() (in module *thelper.utils*), 120  
 class\_id() (*thelper.tasks.detect.BoundingBox* property), 68  
 class\_indices() (*thelper.ifaces.ClassNamesHandler* property), 118  
 class\_names() (*thelper.ifaces.ClassNamesHandler* property), 118  
 class\_names() (*thelper.optim.metrics.ExternalMetric* property), 54  
 class\_names() (*thelper.optim.metrics.ROCCurve* property), 62  
 class\_names() (*thelper.train.utils.ClassifLogger* property), 85  
 class\_names() (*thelper.train.utils.DetectLogger* property), 89  
 Classification (class in *thelper.tasks.classif*), 66  
 classification() (in module *thelper.concepts*), 114  
 ClassifLogger (class in *thelper.train.utils*), 83  
 ClassifReport (class in *thelper.train.utils*), 85  
 ClassNamesHandler (class in *thelper.ifaces*), 118  
 clipstr() (in module *thelper.utils*), 120  
 color\_map() (*thelper.tasks.detect.Detection* property), 70  
 color\_map() (*thelper.tasks.segm.Segmentation* property), 73  
 compute\_average\_precision() (in module *thelper.optim.eval*), 49  
 compute\_bbox\_iou() (in module *thelper.optim.eval*), 49  
 compute\_mask\_iou() (in module *thelper.optim.eval*), 49  
 compute\_pascalvoc\_metrics() (in module *thelper.optim.eval*), 49  
 confidence() (*thelper.tasks.detect.BoundingBox* property), 68  
 ConfusionMatrix (class in *thelper.train.utils*), 86  
 conv\_1x1\_bn() (in module *thelper.nn.mobilenet*), 47  
 conv\_bn() (in module *thelper.nn.mobilenet*), 47  
 count\_tiles() (*thelper.transforms.operations.Tile* method), 100  
 create\_annotator() (in module *thelper.gui.utils*), 42  
 create\_consumers() (in module *thelper.train.utils*), 90  
 create\_global\_task() (in module *thelper.tasks.utils*), 75  
 create\_hdf5() (in module *thelper.data.utils*), 34  
 create\_hdf5\_dataset() (in module *thelper.utils*), 120  
 create\_key\_listener() (in module *thelper.gui.utils*), 43  
 create\_loaders() (in module *thelper.data.utils*), 35  
 create\_loaders() (*thelper.data.loaders.LoaderFactory* method), 33  
 create\_loss\_fn() (in module *thelper.optim.utils*), 63  
 create\_metrics() (in module *thelper.optim.utils*), 64  
 create\_model() (in module *thelper.nn.utils*), 47  
 create\_optimizer() (in module *thelper.optim.utils*), 64  
 create\_parsers() (in module *thelper.data.utils*), 38  
 create\_scheduler() (in module *thelper.optim.utils*), 64  
 create\_session() (in module *thelper.cli*), 110  
 create\_task() (in module *thelper.tasks.utils*), 75  
 create\_tester() (in module *thelper.infer.utils*), 44  
 create\_trainer() (in module *thelper.train.utils*), 90  
 CURRENT\_KEY (*thelper.gui.annotators.ImageSegmentAnnotator* attribute), 41

## D

decode() (*thelper.tasks.detect.BoundingBox* static method), 68  
 decode\_data() (in module *thelper.utils*), 120  
 default\_collate() (in module *thelper.data.loaders*), 33  
 densenet121() (in module *thelper.nn.densenet*), 46  
 densenet161() (in module *thelper.nn.densenet*), 46  
 densenet169() (in module *thelper.nn.densenet*), 46  
 densenet201() (in module *thelper.nn.densenet*), 46  
 Detection (class in *thelper.tasks.detect*), 69  
 detection() (in module *thelper.concepts*), 115  
 DetectLogger (class in *thelper.train.utils*), 88  
 difficult() (*thelper.tasks.detect.BoundingBox* property), 68  
 dontcare() (*thelper.tasks.segm.Segmentation* property), 73  
 download\_file() (in module *thelper.utils*), 121  
 draw() (in module *thelper.draw*), 116  
 draw\_bbox() (in module *thelper.draw*), 116  
 draw\_bboxes() (in module *thelper.draw*), 116  
 draw\_classifs() (in module *thelper.draw*), 117  
 draw\_confmat() (in module *thelper.draw*), 117  
 draw\_errbars() (in module *thelper.draw*), 117  
 draw\_histogram() (in module *thelper.draw*), 117  
 draw\_images() (in module *thelper.draw*), 117  
 draw\_pascalvoc\_curve() (in module *thelper.draw*), 117  
 draw\_popbars() (in module *thelper.draw*), 117  
 draw\_predicts() (in module *thelper.draw*), 117  
 draw\_roc\_curve() (in module *thelper.draw*), 117  
 draw\_segments() (in module *thelper.draw*), 117  
 draw\_stroke() (*thelper.gui.annotators.ImageSegmentAnnotator.Brush* static method), 41  
 Duplicator (class in *thelper.transforms.operations*), 93

## E

`encode()` (*thelper.tasks.detect.BoundingBox* method), 68

`encode_data()` (in module *thelper.utils*), 121

`eval()` (*thelper.optim.metrics.Accuracy* method), 51

`eval()` (*thelper.optim.metrics.AveragePrecision* method), 52

`eval()` (*thelper.optim.metrics.ExternalMetric* method), 54

`eval()` (*thelper.optim.metrics.IntersectionOverUnion* method), 56

`eval()` (*thelper.optim.metrics.MeanAbsoluteError* method), 57

`eval()` (*thelper.optim.metrics.MeanSquaredError* method), 58

`eval()` (*thelper.optim.metrics.Metric* method), 59

`eval()` (*thelper.optim.metrics.PSNR* method), 60

`eval()` (*thelper.optim.metrics.ROCCurve* method), 62

`eval()` (*thelper.train.base.Trainer* method), 78

`eval_epoch()` (*thelper.data.geo.infer.SlidingWindowTester* method), 30

`eval_epoch()` (*thelper.infer.base.Tester* method), 43

`eval_epoch()` (*thelper.train.ae.AutoEncoderTrainer* method), 76

`eval_epoch()` (*thelper.train.base.Trainer* method), 79

`eval_epoch()` (*thelper.train.classif.ImageClassifTrainer* method), 80

`eval_epoch()` (*thelper.train.detect.ObjDetectTrainer* method), 81

`eval_epoch()` (*thelper.train.regr.RegressionTrainer* method), 82

`eval_epoch()` (*thelper.train.segm.ImageSegmTrainer* method), 83

`export_geojson_with_crs()` (in module *thelper.data.geo.utils*), 31

`export_geotiff()` (in module *thelper.data.geo.utils*), 31

`export_model()` (in module *thelper.cli*), 111

`ExternalMetric` (class in *thelper.optim.metrics*), 53

`extract_tar()` (in module *thelper.utils*), 121

## F

`fetch_hdf5_sample()` (in module *thelper.utils*), 121

`fig2array()` (in module *thelper.draw*), 117

`fill_hdf5_sample()` (in module *thelper.utils*), 121

`FormatHandler` (class in *thelper.ifaces*), 118

## G

`gen_report()` (*thelper.train.utils.ClassifReport* method), 86

`gen_report()` (*thelper.train.utils.DetectLogger* method), 89

`get_activation_layer()` (in module *thelper.nn.resnet*), 47

`get_available_cuda_devices()` (in module *thelper.utils*), 121

`get_base_transforms()` (*thelper.data.loaders.LoaderFactory* method), 33

`get_bgr_from_hsl()` (in module *thelper.draw*), 117

`get_caller_name()` (in module *thelper.utils*), 121

`get_checkpoint_session_root()` (in module *thelper.utils*), 122

`get_class_logger()` (in module *thelper.utils*), 122

`get_class_sample_map()` (*thelper.tasks.classif.Classification* method), 66

`get_class_sizes()` (*thelper.tasks.classif.Classification* method), 67

`get_class_sizes()` (*thelper.tasks.detect.Detection* method), 70

`get_class_sizes()` (*thelper.tasks.segm.Segmentation* method), 73

`get_class_weights()` (in module *thelper.data.utils*), 39

`get_compat()` (*thelper.tasks.classif.Classification* method), 67

`get_compat()` (*thelper.tasks.detect.Detection* method), 70

`get_compat()` (*thelper.tasks.regr.Regression* method), 71

`get_compat()` (*thelper.tasks.segm.Segmentation* method), 73

`get_compat()` (*thelper.tasks.utils.Task* method), 74

`get_config_output_root()` (in module *thelper.utils*), 122

`get_config_session_name()` (in module *thelper.utils*), 122

`get_coords_map()` (in module *thelper.nn.coordconv*), 46

`get_displayable_heatmap()` (in module *thelper.draw*), 117

`get_displayable_image()` (in module *thelper.draw*), 117

`get_env_list()` (in module *thelper.utils*), 122

`get_feature_bbox()` (in module *thelper.data.geo.utils*), 31

`get_feature_roi()` (in module *thelper.data.geo.utils*), 31

`get_file_paths()` (in module *thelper.utils*), 122

`get_func_logger()` (in module *thelper.utils*), 122

`get_geocoord()` (in module *thelper.data.geo.utils*), 31

`get_geoextent()` (in module *thelper.data.geo.utils*), 31

`get_git_stamp()` (in module *thelper.utils*), 122



- [get\\_glob\\_paths\(\)](#) (in module *thelper.utils*), 122  
[get\\_key\(\)](#) (in module *thelper.utils*), 122  
[get\\_key\\_def\(\)](#) (in module *thelper.utils*), 122  
[get\\_label\\_color\\_mapping\(\)](#) (in module *thelper.draw*), 118  
[get\\_label\\_html\\_color\\_code\(\)](#) (in module *thelper.draw*), 118  
[get\\_learnable\\_param\\_count\(\)](#) (in module *thelper.nn.utils*), 48  
[get\\_log\\_stamp\(\)](#) (in module *thelper.utils*), 122  
[get\\_lr\(\)](#) (in module *thelper.optim.utils*), 64  
[get\\_mask\\_path\(\)](#) (*thelper.gui.annotators.ImageSegmentAnnotator* method), 42  
[get\\_params\\_hash\(\)](#) (in module *thelper.utils*), 123  
[get\\_pxcoord\(\)](#) (in module *thelper.data.geo.utils*), 31  
[get\\_save\\_dir\(\)](#) (in module *thelper.utils*), 123  
[get\\_slurm\\_tmpdir\(\)](#) (in module *thelper.utils*), 123  
[get\\_split\(\)](#) (*thelper.data.loaders.LoaderFactory* method), 33  
[get\\_upsampling\\_weight\(\)](#) (in module *thelper.nn.fcn*), 47  
[goal\(\)](#) (*thelper.optim.metrics.Accuracy* property), 51  
[goal\(\)](#) (*thelper.optim.metrics.AveragePrecision* property), 52  
[goal\(\)](#) (*thelper.optim.metrics.ExternalMetric* property), 54  
[goal\(\)](#) (*thelper.optim.metrics.IntersectionOverUnion* property), 56  
[goal\(\)](#) (*thelper.optim.metrics.MeanAbsoluteError* property), 57  
[goal\(\)](#) (*thelper.optim.metrics.MeanSquaredError* property), 58  
[goal\(\)](#) (*thelper.optim.metrics.Metric* property), 59  
[goal\(\)](#) (*thelper.optim.metrics.PSNR* property), 60  
[goal\(\)](#) (*thelper.optim.metrics.ROCCurve* property), 62  
[group\\_bbox\(\)](#) (*thelper.train.utils.DetectLogger* method), 89  
[gt\\_key\(\)](#) (*thelper.tasks.utils.Task* property), 74  
[GUI\\_BAR\\_SIZE](#) (*thelper.gui.annotators.ImageSegmentAnnotator* attribute), 41  
[GUI\\_DIRTY](#) (*thelper.gui.annotators.ImageSegmentAnnotator* attribute), 41
- ## H
- [handle\\_keys\(\)](#) (*thelper.gui.annotators.ImageSegmentAnnotator* method), 42  
[height\(\)](#) (*thelper.tasks.detect.BoundingBox* property), 68
- ## I
- [image\\_id\(\)](#) (*thelper.tasks.detect.BoundingBox* property), 68  
[ImageClassifTrainer](#) (class in *thelper.train.classif*), 79  
[ImageSegmentAnnotator](#) (class in *thelper.gui.annotators*), 40  
[ImageSegmentAnnotator.Brush](#) (class in *thelper.gui.annotators*), 41  
[ImageSegmentAnnotator.ZoomTooltip](#) (class in *thelper.gui.annotators*), 42  
[ImageSegmTrainer](#) (class in *thelper.train.segm*), 82  
[import\\_class\(\)](#) (in module *thelper.utils*), 123  
[import\\_function\(\)](#) (in module *thelper.utils*), 123  
[include\\_margin\(\)](#) (*thelper.tasks.detect.BoundingBox* property), 68  
[init\\_force\\_session\(\)](#) (in module *thelper.cli*), 111  
[init\\_logger\(\)](#) (in module *thelper.utils*), 123  
[input\\_key\(\)](#) (*thelper.tasks.utils.Task* property), 74  
[input\\_shape\(\)](#) (*thelper.tasks.regr.Regression* property), 71  
[IntersectionOverUnion](#) (class in *thelper.optim.metrics*), 55  
[intersects\(\)](#) (*thelper.tasks.detect.BoundingBox* method), 68  
[invert\(\)](#) (*thelper.transforms.operations.Affine* method), 92  
[invert\(\)](#) (*thelper.transforms.operations.CenterCrop* method), 93  
[invert\(\)](#) (*thelper.transforms.operations.Duplicator* method), 94  
[invert\(\)](#) (*thelper.transforms.operations.NormalizeMinMax* method), 95  
[invert\(\)](#) (*thelper.transforms.operations.NormalizeZeroMeanUnitVar* method), 95  
[invert\(\)](#) (*thelper.transforms.operations.NoTransform* method), 94  
[invert\(\)](#) (*thelper.transforms.operations.RandomResizedCrop* method), 97  
[invert\(\)](#) (*thelper.transforms.operations.RandomShift* method), 97  
[invert\(\)](#) (*thelper.transforms.operations.Resize* method), 98  
[invert\(\)](#) (*thelper.transforms.operations.SelectChannels* method), 99  
[invert\(\)](#) (*thelper.transforms.operations.Tile* method), 100  
[invert\(\)](#) (*thelper.transforms.operations.ToColor* method), 101  
[invert\(\)](#) (*thelper.transforms.operations.ToGray* method), 101  
[invert\(\)](#) (*thelper.transforms.operations.ToNumpy* method), 101  
[invert\(\)](#) (*thelper.transforms.operations.Transpose* method), 102  
[invert\(\)](#) (*thelper.transforms.operations.Unsqueeze* method), 102  
[is\\_scalar\(\)](#) (in module *thelper.utils*), 123  
[iscrowd\(\)](#) (*thelper.tasks.detect.BoundingBox* prop-

erty), 68

## J

`json()` (*thelper.tasks.detect.BoundingBox method*), 69

## K

`keys()` (*thelper.tasks.utils.Task property*), 75

## L

`LAKE_ID` (*thelper.data.geo.ogc.TB15D104 attribute*), 31

`LATEST_PT` (*thelper.gui.annotators.ImageSegmentAnnotator attribute*), 41

`LATEST_RAW_PT` (*thelper.gui.annotators.ImageSegmentAnnotator attribute*), 41

`left()` (*thelper.tasks.detect.BoundingBox property*), 69

`live_eval()` (*thelper.optim.metrics.AveragePrecision property*), 52

`live_eval()` (*thelper.optim.metrics.ExternalMetric property*), 55

`live_eval()` (*thelper.optim.metrics.Metric property*), 59

`live_eval()` (*thelper.optim.metrics.ROCCurve property*), 62

`load()` (*thelper.gui.annotators.ImageSegmentAnnotator method*), 42

`load_augments()` (*in module thelper.transforms.utils*), 103

`load_checkpoint()` (*in module thelper.utils*), 124

`load_config()` (*in module thelper.utils*), 124

`load_transforms()` (*in module thelper.transforms.utils*), 104

`LoaderFactory` (*class in thelper.data.loaders*), 32

`lreplace()` (*in module thelper.utils*), 124

## M

`main()` (*in module thelper.cli*), 112

`make_argparser()` (*in module thelper.cli*), 112

`make_conv2d()` (*in module thelper.nn.coordconv*), 46

`make_tester_from_trainer()` (*in module thelper.infer.base*), 44

`MASK_DIRTY` (*thelper.gui.annotators.ImageSegmentAnnotator attribute*), 41

`maximize` (*thelper.optim.metrics.Metric attribute*), 59

`MeanAbsoluteError` (*class in thelper.optim.metrics*), 56

`MeanSquaredError` (*class in thelper.optim.metrics*), 57

`meta_keys()` (*thelper.tasks.utils.Task property*), 75

`Metric` (*class in thelper.optim.metrics*), 58

`migrate_checkpoint()` (*in module thelper.utils*), 124

`migrate_config()` (*in module thelper.utils*), 124

`minimize` (*thelper.optim.metrics.Metric attribute*), 59

## module

`thelper`, 29

`thelper.cli`, 110

`thelper.concepts`, 114

`thelper.data`, 29

`thelper.data.geo`, 29

`thelper.data.geo.agravis`, 29

`thelper.data.geo.gdl`, 30

`thelper.data.geo.infer`, 30

`thelper.data.geo.ogc`, 31

`thelper.data.geo.parsers`, 31

`thelper.data.geo.utils`, 31

`thelper.data.loaders`, 32

`thelper.data.parsers`, 34

`thelper.data.pascalvoc`, 34

`thelper.data.samplers`, 34

`thelper.data.utils`, 34

`thelper.draw`, 116

`thelper.gui`, 40

`thelper.gui.annotators`, 40

`thelper.gui.utils`, 42

`thelper.ifaces`, 118

`thelper.infer`, 43

`thelper.infer.base`, 43

`thelper.infer.impl`, 44

`thelper.infer.utils`, 44

`thelper.nn`, 45

`thelper.nn.common`, 46

`thelper.nn.coordconv`, 46

`thelper.nn.densenet`, 46

`thelper.nn.efficientnet`, 47

`thelper.nn.fcn`, 47

`thelper.nn.inceptionresnetv2`, 47

`thelper.nn.lenet`, 47

`thelper.nn.mobilenet`, 47

`thelper.nn.resnet`, 47

`thelper.nn.segmentation`, 45

`thelper.nn.segmentation.base`, 45

`thelper.nn.segmentation.deeplabv3`, 45

`thelper.nn.segmentation.fcn`, 45

`thelper.nn.sr`, 45

`thelper.nn.sr.srcnn`, 46

`thelper.nn.sr.vdsr`, 46

`thelper.nn.srm`, 47

`thelper.nn.unet`, 47

`thelper.nn.utils`, 47

`thelper.optim`, 48

`thelper.optim.eval`, 49

`thelper.optim.losses`, 50

`thelper.optim.metrics`, 50

`thelper.optim.schedulers`, 63

`thelper.optim.utils`, 63

`thelper.session`, 64



[thelper.session.base](#), 64  
[thelper.tasks](#), 65  
[thelper.tasks.classif](#), 66  
[thelper.tasks.detect](#), 67  
[thelper.tasks.regr](#), 71  
[thelper.tasks.segm](#), 73  
[thelper.tasks.utils](#), 74  
[thelper.train](#), 75  
[thelper.train.ae](#), 76  
[thelper.train.base](#), 76  
[thelper.train.classif](#), 79  
[thelper.train.detect](#), 80  
[thelper.train.regr](#), 81  
[thelper.train.segm](#), 82  
[thelper.train.utils](#), 83  
[thelper.transforms](#), 91  
[thelper.transforms.composers](#), 91  
[thelper.transforms.operations](#), 92  
[thelper.transforms.utils](#), 103  
[thelper.transforms.wrappers](#), 105  
[thelper.typedefs](#), 120  
[thelper.utils](#), 120  
[thelper.viz](#), 108  
[thelper.viz.tsne](#), 109  
[thelper.viz.umap](#), 109

[MOUSE\\_FLAGS](#) (*thelper.gui.annotators.ImageSegmentAnnotator* attribute), 42

## N

[NormalizeMinMax](#) (class in *thelper.transforms.operations*), 94  
[NormalizeZeroMeanUnitVar](#) (class in *thelper.transforms.operations*), 95  
[NoTransform](#) (class in *thelper.transforms.operations*), 94

## O

[ObjDetectTrainer](#) (class in *thelper.train.detect*), 80  
[occluded\(\)](#) (*thelper.tasks.detect.BoundingBox* property), 69  
[on\\_mouse\(\)](#) (*thelper.gui.annotators.ImageSegmentAnnotator* static method), 42  
[on\\_press\(\)](#) (*thelper.gui.annotators.ImageSegmentAnnotator* static method), 42  
[open\\_rasterfile\(\)](#) (in module *thelper.data.geo.utils*), 31

## P

[parse\\_geojson\(\)](#) (in module *thelper.data.geo.utils*), 31  
[parse\\_geojson\\_crs\(\)](#) (in module *thelper.data.geo.utils*), 31  
[parse\\_raster\\_metadata\(\)](#) (in module *thelper.data.geo.utils*), 31

[parse\\_rasters\(\)](#) (in module *thelper.data.geo.utils*), 32  
[parse\\_roi\(\)](#) (in module *thelper.data.geo.utils*), 32  
[parse\\_shapefile\(\)](#) (in module *thelper.data.geo.utils*), 32  
[plot\(\)](#) (in module *thelper.viz.tsne*), 109  
[postproc\\_features\(\)](#) (in module *thelper.data.geo.ogc*), 31  
[PredictionCallback](#) (class in *thelper.train.utils*), 90  
[PredictionConsumer](#) (class in *thelper.ifaces*), 119  
[PSNR](#) (class in *thelper.optim.metrics*), 59

## Q

[query\\_string\(\)](#) (in module *thelper.utils*), 125  
[query\\_yes\\_no\(\)](#) (in module *thelper.utils*), 125

## R

[RandomResizedCrop](#) (class in *thelper.transforms.operations*), 95  
[RandomShift](#) (class in *thelper.transforms.operations*), 97  
[refresh\(\)](#) (*thelper.gui.annotators.ImageSegmentAnnotator.Brush* method), 41  
[refresh\(\)](#) (*thelper.gui.annotators.ImageSegmentAnnotator.ZoomTooltip* method), 42  
[refresh\\_gui\(\)](#) (*thelper.gui.annotators.ImageSegmentAnnotator* method), 42  
[refresh\\_layers\(\)](#) (*thelper.gui.annotators.ImageSegmentAnnotator* method), 42  
[Regression](#) (class in *thelper.tasks.regr*), 71  
[regression\(\)](#) (in module *thelper.concepts*), 115  
[RegressionTrainer](#) (class in *thelper.train.regr*), 81  
[render\(\)](#) (*thelper.optim.metrics.ROCCurve* method), 62  
[render\(\)](#) (*thelper.train.utils.ClassifLogger* method), 85  
[render\(\)](#) (*thelper.train.utils.ConfusionMatrix* method), 87  
[render\(\)](#) (*thelper.train.utils.DetectLogger* method), 90  
[report\(\)](#) (*thelper.ifaces.FormatHandler* method), 119  
[report\(\)](#) (*thelper.train.utils.ConfusionMatrix* method), 87  
[report\\_csv\(\)](#) (*thelper.train.utils.ClassifLogger* method), 85  
[report\\_csv\(\)](#) (*thelper.train.utils.DetectLogger* method), 90  
[report\\_geojson\(\)](#) (*thelper.data.geo.ogc.TB15D104DetectLogger* method), 31  
[report\\_json\(\)](#) (*thelper.train.utils.ClassifLogger* method), 85  
[report\\_json\(\)](#) (*thelper.train.utils.ClassifReport* method), 86

`report_json()` (*thelper.train.utils.DetectLogger method*), 90  
`report_orion_results()` (*in module `thelper.utils`*), 125  
`report_text()` (*thelper.ifaces.FormatHandler method*), 119  
`report_text()` (*thelper.train.utils.ClassifLogger method*), 85  
`report_text()` (*thelper.train.utils.ClassifReport method*), 86  
`report_text()` (*thelper.train.utils.DetectLogger method*), 90  
`reprothook()` (*in module `thelper.utils`*), 125  
`reproject_coords()` (*in module `thelper.data.geo.utils`*), 32  
`reproject_crop()` (*in module `thelper.data.geo.utils`*), 32  
`reset()` (*thelper.ifaces.PredictionConsumer method*), 119  
`reset()` (*thelper.optim.metrics.Accuracy method*), 51  
`reset()` (*thelper.optim.metrics.AveragePrecision method*), 52  
`reset()` (*thelper.optim.metrics.ExternalMetric method*), 55  
`reset()` (*thelper.optim.metrics.IntersectionOverUnion method*), 56  
`reset()` (*thelper.optim.metrics.MeanAbsoluteError method*), 57  
`reset()` (*thelper.optim.metrics.MeanSquaredError method*), 58  
`reset()` (*thelper.optim.metrics.PSNR method*), 60  
`reset()` (*thelper.optim.metrics.ROCCurve method*), 62  
`reset()` (*thelper.train.utils.ClassifLogger method*), 85  
`reset()` (*thelper.train.utils.ClassifReport method*), 86  
`reset()` (*thelper.train.utils.ConfusionMatrix method*), 87  
`reset()` (*thelper.train.utils.DetectLogger method*), 90  
`Resize` (*class in `thelper.transforms.operations`*), 98  
`resolve_import()` (*in module `thelper.utils`*), 125  
`resume_session()` (*in module `thelper.cli`*), 112  
`right()` (*thelper.tasks.detect.BoundingBox property*), 69  
`ROCCurve` (*class in `thelper.optim.metrics`*), 60  
`run()` (*thelper.gui.annotators.Annotator method*), 40  
`run()` (*thelper.gui.annotators.ImageSegmentAnnotator method*), 42

**S**

`safe_crop()` (*in module `thelper.draw`*), 118  
`save_config()` (*in module `thelper.utils`*), 125  
`save_env_list()` (*in module `thelper.utils`*), 126  
`Segmentation` (*class in `thelper.tasks.segm`*), 73  
`segmentation()` (*in module `thelper.concepts`*), 115  
`SelectChannels` (*class in `thelper.transforms.operations`*), 98  
`SessionRunner` (*class in `thelper.session.base`*), 64  
`set_epoch()` (*thelper.transforms.wrappers.AlbumentationsWrapper method*), 106  
`set_epoch()` (*thelper.transforms.wrappers.AugmentorWrapper method*), 107  
`set_epoch()` (*thelper.transforms.wrappers.TransformWrapper method*), 108  
`set_matplotlib_agg()` (*in module `thelper.utils`*), 126  
`set_seed()` (*thelper.transforms.operations.RandomResizedCrop method*), 97  
`set_seed()` (*thelper.transforms.operations.RandomShift method*), 98  
`set_seed()` (*thelper.transforms.wrappers.AlbumentationsWrapper method*), 106  
`set_seed()` (*thelper.transforms.wrappers.AugmentorWrapper method*), 107  
`set_seed()` (*thelper.transforms.wrappers.TransformWrapper method*), 108  
`setup()` (*in module `thelper.cli`*), 113  
`setup_cudnn()` (*in module `thelper.utils`*), 126  
`setup_cv2()` (*in module `thelper.utils`*), 126  
`setup_gdal()` (*in module `thelper.utils`*), 126  
`setup_globals()` (*in module `thelper.utils`*), 126  
`setup_plt()` (*in module `thelper.utils`*), 126  
`setup_srm_layer()` (*in module `thelper.nn.srm`*), 47  
`setup_srm_weights()` (*in module `thelper.nn.srm`*), 47  
`setup_sys()` (*in module `thelper.utils`*), 126  
`shave()` (*in module `thelper.nn.common`*), 46  
`SlidingWindowTester` (*class in `thelper.data.geo.infer`*), 30  
`solve_format()` (*thelper.ifaces.FormatHandler method*), 119  
`split_data()` (*in module `thelper.cli`*), 113  
`str2bool()` (*in module `thelper.utils`*), 126  
`str2size()` (*in module `thelper.utils`*), 126  
`stringify_confmat()` (*in module `thelper.utils`*), 126  
`Struct` (*class in `thelper.utils`*), 120  
`SuperResolution` (*class in `thelper.tasks.regr`*), 72  
`SUPPORT_PREFIX` (*in module `thelper.concepts`*), 114  
`supports()` (*in module `thelper.concepts`*), 116  
`supports_classification` (*thelper.data.geo.infer.SlidingWindowTester attribute*), 30  
`supports_classification` (*thelper.optim.metrics.Accuracy attribute*), 51  
`supports_classification` (*thelper.optim.metrics.ExternalMetric attribute*), 55

[supports\\_classification](#)  
 ([thelper.optim.metrics.ROCCurve](#) attribute), [62](#)  
[supports\\_classification](#)  
 ([thelper.tasks.classif.Classification](#) attribute),  
[67](#)  
[supports\\_classification](#)  
 ([thelper.train.ae.AutoEncoderTrainer](#) at-  
 tribute), [76](#)  
[supports\\_classification](#)  
 ([thelper.train.classif.ImageClassifTrainer](#)  
 attribute), [80](#)  
[supports\\_classification](#)  
 ([thelper.train.utils.ClassifLogger](#) attribute),  
[85](#)  
[supports\\_classification](#)  
 ([thelper.train.utils.ClassifReport](#) attribute),  
[86](#)  
[supports\\_classification](#)  
 ([thelper.train.utils.ConfusionMatrix](#) attribute),  
[87](#)  
[supports\\_detection](#)  
 ([thelper.optim.metrics.AveragePrecision](#)  
 attribute), [52](#)  
[supports\\_detection](#)  
 ([thelper.tasks.detect.BoundingBox](#) attribute),  
[69](#)  
[supports\\_detection](#)  
 ([thelper.tasks.detect.Detection](#) attribute),  
[70](#)  
[supports\\_detection](#)  
 ([thelper.train.detect.ObjDetectTrainer](#) at-  
 tribute), [81](#)  
[supports\\_detection](#)  
 ([thelper.train.utils.DetectLogger](#) attribute),  
[90](#)  
[supports\\_regression](#)  
 ([thelper.optim.metrics.MeanAbsoluteError](#)  
 attribute), [57](#)  
[supports\\_regression](#)  
 ([thelper.optim.metrics.MeanSquaredError](#)  
 attribute), [58](#)  
[supports\\_regression](#)  
 ([thelper.optim.metrics.PSNR](#) attribute), [60](#)  
[supports\\_regression](#)  
 ([thelper.tasks.regr.Regression](#) attribute),  
[71](#)  
[supports\\_regression](#)  
 ([thelper.tasks.regr.SuperResolution](#) attribute),  
[72](#)  
[supports\\_regression](#)  
 ([thelper.train.regr.RegressionTrainer](#) attribute),  
[82](#)  
[supports\\_segmentation](#)  
 ([thelper.optim.metrics.Accuracy](#) attribute),  
[51](#)  
[supports\\_segmentation](#)  
 ([thelper.optim.metrics.ExternalMetric](#) at-  
 tribute), [55](#)  
[supports\\_segmentation](#)  
 ([thelper.optim.metrics.IntersectionOverUnion](#)  
 attribute), [56](#)  
[supports\\_segmentation](#)  
 ([thelper.optim.metrics.ROCCurve](#) attribute), [62](#)  
[supports\\_segmentation](#)  
 ([thelper.tasks.segm.Segmentation](#) attribute), [74](#)  
[supports\\_segmentation](#)  
 ([thelper.train.ae.AutoEncoderTrainer](#) at-  
 tribute), [76](#)  
[supports\\_segmentation](#)  
 ([thelper.train.segm.ImageSegmTrainer](#) at-  
 tribute), [83](#)  
[supports\\_segmentation](#)  
 ([thelper.train.utils.ConfusionMatrix](#) attribute),  
[87](#)  
[swap\\_coordconv\\_layers\(\)](#) (in module  
[thelper.nn.coordconv](#)), [46](#)

## T

[target\\_max\(\)](#) ([thelper.tasks.regr.Regression](#) prop-  
 erty), [71](#)  
[target\\_min\(\)](#) ([thelper.tasks.regr.Regression](#) prop-  
 erty), [72](#)  
[target\\_shape\(\)](#) ([thelper.tasks.regr.Regression](#) prop-  
 erty), [72](#)  
[target\\_type\(\)](#) ([thelper.tasks.regr.Regression](#) prop-  
 erty), [72](#)  
[Task](#) (class in [thelper.tasks.utils](#)), [74](#)  
[task\(\)](#) ([thelper.tasks.detect.BoundingBox](#) property), [69](#)  
[TB15D104](#) (class in [thelper.data.geo.ogc](#)), [31](#)  
[TB15D104DetectLogger](#) (class in  
[thelper.data.geo.ogc](#)), [31](#)  
[test\(\)](#) ([thelper.infer.base.Tester](#) method), [44](#)  
[test\\_cuda\\_device\\_availability\(\)](#) (in mod-  
 ule [thelper.utils](#)), [126](#)  
[test\\_epoch\(\)](#) ([thelper.infer.base.Tester](#) method), [44](#)  
[Tester](#) (class in [thelper.infer.base](#)), [43](#)  
[thelper](#)  
 module, [29](#)  
[thelper.cli](#)  
 module, [110](#)  
[thelper.concepts](#)  
 module, [114](#)  
[thelper.data](#)  
 module, [29](#)  
[thelper.data.geo](#)  
 module, [29](#)  
[thelper.data.geo.agravis](#)  
 module, [29](#)

thelper.data.geo.gdl	thelper.nn.mobilenet
module, 30	module, 47
thelper.data.geo.infer	thelper.nn.resnet
module, 30	module, 47
thelper.data.geo.ogc	thelper.nn.segmentation
module, 31	module, 45
thelper.data.geo.parsers	thelper.nn.segmentation.base
module, 31	module, 45
thelper.data.geo.utils	thelper.nn.segmentation.deeplabv3
module, 31	module, 45
thelper.data.loaders	thelper.nn.segmentation.fcn
module, 32	module, 45
thelper.data.parsers	thelper.nn.sr
module, 34	module, 45
thelper.data.pascalvoc	thelper.nn.sr.srcnn
module, 34	module, 46
thelper.data.samplers	thelper.nn.sr.vdsr
module, 34	module, 46
thelper.data.utils	thelper.nn.srm
module, 34	module, 47
thelper.draw	thelper.nn.unet
module, 116	module, 47
thelper.gui	thelper.nn.utils
module, 40	module, 47
thelper.gui.annotators	thelper.optim
module, 40	module, 48
thelper.gui.utils	thelper.optim.eval
module, 42	module, 49
thelper.ifaces	thelper.optim.losses
module, 118	module, 50
thelper.infer	thelper.optim.metrics
module, 43	module, 50
thelper.infer.base	thelper.optim.schedulers
module, 43	module, 63
thelper.infer.impl	thelper.optim.utils
module, 44	module, 63
thelper.infer.utils	thelper.session
module, 44	module, 64
thelper.nn	thelper.session.base
module, 45	module, 64
thelper.nn.common	thelper.tasks
module, 46	module, 65
thelper.nn.coordconv	thelper.tasks.classif
module, 46	module, 66
thelper.nn.densenet	thelper.tasks.detect
module, 46	module, 67
thelper.nn.efficientnet	thelper.tasks.regr
module, 47	module, 71
thelper.nn.fcn	thelper.tasks.segm
module, 47	module, 73
thelper.nn.inceptionresnetv2	thelper.tasks.utils
module, 47	module, 74
thelper.nn.lenet	thelper.train
module, 47	module, 75

thelper.train.ae  
     module, 76  
 thelper.train.base  
     module, 76  
 thelper.train.classif  
     module, 79  
 thelper.train.detect  
     module, 80  
 thelper.train.regr  
     module, 81  
 thelper.train.segm  
     module, 82  
 thelper.train.utils  
     module, 83  
 thelper.transforms  
     module, 91  
 thelper.transforms.composers  
     module, 91  
 thelper.transforms.operations  
     module, 92  
 thelper.transforms.utils  
     module, 103  
 thelper.transforms.wrappers  
     module, 105  
 thelper.typedefs  
     module, 120  
 thelper.utils  
     module, 120  
 thelper.viz  
     module, 108  
 thelper.viz.tsne  
     module, 109  
 thelper.viz.umap  
     module, 109  
 Tile (class in thelper.transforms.operations), 99  
 to\_numpy() (in module thelper.utils), 126  
 ToColor (class in thelper.transforms.operations), 100  
 ToGray (class in thelper.transforms.operations), 101  
 tolist() (thelper.tasks.detect.BoundingBox method), 69  
 ToNumpy (class in thelper.transforms.operations), 101  
 top() (thelper.tasks.detect.BoundingBox property), 69  
 top\_left() (thelper.tasks.detect.BoundingBox property), 69  
 totuple() (thelper.tasks.detect.BoundingBox method), 69  
 train() (thelper.infer.base.Tester method), 44  
 train() (thelper.train.base.Trainer method), 79  
 train\_epoch() (thelper.infer.base.Tester method), 44  
 train\_epoch() (thelper.train.ae.AutoEncoderTrainer method), 76  
 train\_epoch() (thelper.train.base.Trainer method), 79  
 train\_epoch() (thelper.train.classif.ImageClassifTrainer method), 80  
 train\_epoch() (thelper.train.detect.ObjDetectTrainer method), 81  
 train\_epoch() (thelper.train.regr.RegressionTrainer method), 82  
 train\_epoch() (thelper.train.segm.ImageSegmTrainer method), 83  
 Trainer (class in thelper.train.base), 76  
 TransformWrapper (class in thelper.transforms.wrappers), 107  
 Transpose (class in thelper.transforms.operations), 101  
 truncated() (thelper.tasks.detect.BoundingBox property), 69  
 TYPECE\_LAKE (thelper.data.geo.ogc.TB15D104 attribute), 31  
 TYPECE\_RIVER (thelper.data.geo.ogc.TB15D104 attribute), 31

## U

Unsqueeze (class in thelper.transforms.operations), 102  
 update() (thelper.ifaces.PredictionConsumer method), 119  
 update() (thelper.optim.metrics.Accuracy method), 51  
 update() (thelper.optim.metrics.AveragePrecision method), 52  
 update() (thelper.optim.metrics.ExternalMetric method), 55  
 update() (thelper.optim.metrics.IntersectionOverUnion method), 56  
 update() (thelper.optim.metrics.MeanAbsoluteError method), 57  
 update() (thelper.optim.metrics.MeanSquaredError method), 58  
 update() (thelper.optim.metrics.Metric method), 59  
 update() (thelper.optim.metrics.PSNR method), 60  
 update() (thelper.optim.metrics.ROCCurve method), 62  
 update() (thelper.train.utils.ClassifLogger method), 85  
 update() (thelper.train.utils.ClassifReport method), 86  
 update() (thelper.train.utils.ConfusionMatrix method), 87  
 update() (thelper.train.utils.DetectLogger method), 90  
 update() (thelper.train.utils.PredictionCallback method), 90

## V

visualize() (in module thelper.viz), 108  
 visualize() (in module thelper.viz.tsne), 109  
 visualize() (in module thelper.viz.umap), 109  
 visualize\_data() (in module thelper.cli), 114

## W

`weights_init_kaiming()` (in module `thelper.nn.common`), [46](#)  
`weights_init_xavier()` (in module `thelper.nn.common`), [46](#)  
`width()` (`thelper.tasks.detect.BoundingBox` property), [69](#)  
`WINDOW_SIZE` (`thelper.gui.annotators.ImageSegmentAnnotator` attribute), [42](#)