
MMClassification Documentation

Release 0.18.0

MMClassification Contributors

Dec 01, 2021

GET STARTED

1	Installation	3
2	Getting Started	7
3	Tutorial 1: Learn about Configs	13
4	Tutorial 2: Fine-tune Models	25
5	Tutorial 3: Adding New Dataset	31
6	Tutorial 4: Custom Data Pipelines	35
7	Tutorial 5: Adding New Modules	39
8	Tutorial 6: Customize Schedule	45
9	Tutorial 7: Customize Runtime Settings	53
10	Model Zoo Summary	59
11	Model Zoo	61
12	MLP-Mixer: An all-MLP Architecture for Vision	63
13	MobileNetV2: Inverted Residuals and Linear Bottlenecks	65
14	Searching for MobileNetV3	67
15	Designing Network Design Spaces	69
16	Repvgg: Making vgg-style convnets great again	71
17	Res2Net: A New Multi-scale Backbone Architecture	73
18	Deep Residual Learning for Image Recognition	75
19	Aggregated Residual Transformations for Deep Neural Networks	77
20	Squeeze-and-Excitation Networks	79
21	ShuffleNet: An Extremely Efficient Convolutional Neural Network for Mobile Devices	81
22	Shufflenet v2: Practical guidelines for efficient cnn architecture design	83

23 Swin Transformer: Hierarchical Vision Transformer using Shifted Windows	85
24 Tokens-to-Token ViT: Training Vision Transformers from Scratch on ImageNet	87
25 Transformer in Transformer	89
26 Very Deep Convolutional Networks for Large-Scale Image Recognition	91
27 An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale	93
28 Pytorch to ONNX (Experimental)	95
29 ONNX to TensorRT (Experimental)	99
30 Pytorch to TorchScript (Experimental)	101
31 Model Serving	103
32 Visualization	105
33 Contributing to OpenMMLab	109
34 mmcls.apis	111
35 mmcls.core	113
36 mmcls.models	117
37 mmcls.datasets	153
38 mmcls.utils	169
39 Indices and tables	171
Python Module Index	173
Index	175

You can switch between Chinese and English documentation in the lower-left corner of the layout.
您可以在页面左下角切换中英文文档。

INSTALLATION

1.1 Requirements

- Python 3.6+
- PyTorch 1.5+
- [MMCV](#)

The compatible MMClassification and MMCV versions are as below. Please install the correct version of MMCV to avoid installation issues.

Note: Since the `master` branch is under frequent development, the `mmcv` version dependency may be inaccurate. If you encounter problems when using the `master` branch, please try to update `mmcv` to the latest version.

1.2 Install MMClassification

a. Create a conda virtual environment and activate it.

```
conda create -n open-mmlab python=3.8 -y
conda activate open-mmlab
```

b. Install PyTorch and torchvision following the [official instructions](#), e.g.,

```
conda install pytorch torchvision -c pytorch
```

Note: Make sure that your compilation CUDA version and runtime CUDA version match. You can check the supported CUDA version for precompiled packages on the [PyTorch website](#).

E.g. 1 If you have CUDA 10.1 installed under `/usr/local/cuda` and would like to install PyTorch 1.5.1, you need to install the prebuilt PyTorch with CUDA 10.1.

```
conda install pytorch==1.5.1 torchvision==0.6.1 cudatoolkit=10.1 -c pytorch
```

E.g. 2 If you have CUDA 11.3 installed under `/usr/local/cuda` and would like to install PyTorch 1.10.0., you need to install the prebuilt PyTorch with CUDA 11.3.

```
conda install pytorch==1.10.0 torchvision==0.11.1 cudatoolkit=11.3 -c pytorch
```

If you build PyTorch from source instead of installing the prebuilt package, you can use more CUDA versions such as 9.0.

c. Install MMClassification repository.

1.2.1 Release version

We recommend you to install MMClassification with [MIM](#).

```
pip install git+https://github.com/open-mmlab/mim.git
mim install mmcls
```

MIM can automatically install OpenMMLab projects and their requirements, and it can also help us to train, parameter search and pretrain model download.

Or, you can install MMClassification with pip:

```
pip install mmcls
```

1.2.2 Develop version

First, clone the MMClassification repository.

```
git clone https://github.com/open-mmlab/mclassification.git
cd mclassification
```

And then, install build requirements and install MMClassification.

```
pip install -e . # or "python setup.py develop"
```

Note: Following above instructions, MMClassification is installed on dev mode, any local modifications made to the code will take effect without the need to reinstall it (unless you submit some commits and want to update the version number).

1.2.3 Another option: Docker Image

We provide a [Dockerfile](#) to build an image.

```
# build an image with PyTorch 1.6.0, CUDA 10.1, CUDNN 7.
docker build -f ./docker/Dockerfile --rm -t mmcls:torch1.6.0-cuda10.1-cudnn7 .
```

Important: Make sure you've installed the [nvidia-container-toolkit](#).

Run a container built from mmcls image with command:

```
docker run --gpus all --shm-size=8g -it -v {DATA_DIR}:/workspace/mclassification/data_
↪mmcls:torch1.6.0-cuda10.1-cudnn7 /bin/bash
```


1.3 Using multiple MMClassification versions

The train and test scripts already modify the PYTHONPATH to ensure the script use the MMClassification in the current directory.

To use the default MMClassification installed in the environment rather than that you are working with, you can remove the following line in those scripts

```
PYTHONPATH="$(dirname $0)/..":$PYTHONPATH
```


GETTING STARTED

This page provides basic tutorials about the usage of MMClassification.

2.1 Prepare datasets

It is recommended to symlink the dataset root to `$MMCLASSIFICATION/data`. If your folder structure is different, you may need to change the corresponding paths in config files.

```
mmclassification
├── mmcls
├── tools
├── configs
├── docs
├── data
│   ├── imagenet
│   │   ├── meta
│   │   ├── train
│   │   └── val
│   ├── cifar
│   │   └── cifar-10-batches-py
│   └── mnist
│       ├── train-images-idx3-ubyte
│       ├── train-labels-idx1-ubyte
│       ├── t10k-images-idx3-ubyte
│       └── t10k-labels-idx1-ubyte
```

For ImageNet, it has multiple versions, but the most commonly used one is [ILSVRC 2012](#). It can be accessed with the following steps.

1. Register an account and login to the [download page](#).
2. Find download links for ILSVRC2012 and download the following two files
 - ILSVRC2012_img_train.tar (~138GB)
 - ILSVRC2012_img_val.tar (~6.3GB)
3. Untar the downloaded files
4. Download meta data using this [script](#)

For MNIST, CIFAR10 and CIFAR100, the datasets will be downloaded and unzipped automatically if they are not found.

For using custom datasets, please refer to *Tutorials 2: Adding New Dataset*.

2.2 Inference with pretrained models

We provide scripts to inference a single image, inference a dataset and test a dataset (e.g., ImageNet).

2.2.1 Inference a single image

```
python demo/image_demo.py ${IMAGE_FILE} ${CONFIG_FILE} ${CHECKPOINT_FILE}

# Example
python demo/image_demo.py demo/demo.JPEG configs/resnet/resnet50_8xb32_in1k.py \
  https://download.openmmlab.com/mmcclassification/v0/resnet/resnet50_8xb32_in1k_20210831-
  ↪ea4938fc.pth
```

2.2.2 Inference and test a dataset

- single GPU
- single node multiple GPU
- multiple node

You can use the following commands to infer a dataset.

```
# single-gpu
python tools/test.py ${CONFIG_FILE} ${CHECKPOINT_FILE} [--metrics ${METRICS}] [--out $
  ↪${RESULT_FILE}]

# multi-gpu
./tools/dist_test.sh ${CONFIG_FILE} ${CHECKPOINT_FILE} ${GPU_NUM} [--metrics ${METRICS}] \
  ↪[--out ${RESULT_FILE}]

# multi-node in slurm environment
python tools/test.py ${CONFIG_FILE} ${CHECKPOINT_FILE} [--metrics ${METRICS}] [--out $
  ↪${RESULT_FILE}] --launcher slurm
```

Optional arguments:

- RESULT_FILE: Filename of the output results. If not specified, the results will not be saved to a file. Support formats include json, yaml and pickle.
- METRICS: Items to be evaluated on the results, like accuracy, precision, recall, etc.

Examples:

Assume that you have already downloaded the checkpoints to the directory checkpoints/. Infer ResNet-50 on ImageNet validation set to get predicted labels and their corresponding predicted scores.

```
python tools/test.py configs/resnet/resnet50_8xb16_cifar10.py \
  https://download.openmmlab.com/mmcclassification/v0/resnet/resnet50_b16x8_cifar10_
  ↪20210528-f54bfad9.pth \
  --out result.pkl
```

2.3 Train a model

MMClassification implements distributed training and non-distributed training, which uses `MMDistributedDataParallel` and `MMDDataParallel` respectively.

All outputs (log files and checkpoints) will be saved to the working directory, which is specified by `work_dir` in the config file.

By default we evaluate the model on the validation set after each epoch, you can change the evaluation interval by adding the `interval` argument in the training config.

```
evaluation = dict(interval=12) # Evaluate the model per 12 epochs.
```

2.3.1 Train with a single GPU

```
python tools/train.py ${CONFIG_FILE} [optional arguments]
```

If you want to specify the working directory in the command, you can add an argument `--work_dir` `YOUR_WORK_DIR`.

2.3.2 Train with multiple GPUs

```
./tools/dist_train.sh ${CONFIG_FILE} ${GPU_NUM} [optional arguments]
```

Optional arguments are:

- `--no-validate` (**not suggested**): By default, the codebase will perform evaluation at every `k` (default value is 1) epochs during the training. To disable this behavior, use `--no-validate`.
- `--work-dir` `WORK_DIR`: Override the working directory specified in the config file.
- `--resume-from` `CHECKPOINT_FILE`: Resume from a previous checkpoint file.

Difference between `resume-from` and `load-from`: `resume-from` loads both the model weights and optimizer status, and the epoch is also inherited from the specified checkpoint. It is usually used for resuming the training process that is interrupted accidentally. `load-from` only loads the model weights and the training epoch starts from 0. It is usually used for finetuning.

2.3.3 Train with multiple machines

If you run MMClassification on a cluster managed with `slurm`, you can use the script `slurm_train.sh`. (This script also supports single machine training.)

```
[GPUS=${GPUS}] ./tools/slurm_train.sh ${PARTITION} ${JOB_NAME} ${CONFIG_FILE} ${WORK_DIR}
```

You can check `slurm_train.sh` for full arguments and environment variables.

If you have just multiple machines connected with ethernet, you can refer to PyTorch [launch utility](#). Usually it is slow if you do not have high speed networking like InfiniBand.

2.3.4 Launch multiple jobs on a single machine

If you launch multiple jobs on a single machine, e.g., 2 jobs of 4-GPU training on a machine with 8 GPUs, you need to specify different ports (29500 by default) for each job to avoid communication conflict.

If you use `dist_train.sh` to launch training jobs, you can set the port in commands.

```
CUDA_VISIBLE_DEVICES=0,1,2,3 PORT=29500 ./tools/dist_train.sh ${CONFIG_FILE} 4
CUDA_VISIBLE_DEVICES=4,5,6,7 PORT=29501 ./tools/dist_train.sh ${CONFIG_FILE} 4
```

If you use launch training jobs with Slurm, you need to modify the config files (usually the 6th line from the bottom in config files) to set different communication ports.

In `config1.py`,

```
dist_params = dict(backend='nccl', port=29500)
```

In `config2.py`,

```
dist_params = dict(backend='nccl', port=29501)
```

Then you can launch two jobs with `config1.py` and `config2.py`.

```
CUDA_VISIBLE_DEVICES=0,1,2,3 GPUS=4 ./tools/slurm_train.sh ${PARTITION} ${JOB_NAME} \
↪ config1.py ${WORK_DIR}
CUDA_VISIBLE_DEVICES=4,5,6,7 GPUS=4 ./tools/slurm_train.sh ${PARTITION} ${JOB_NAME} \
↪ config2.py ${WORK_DIR}
```

2.4 Useful tools

We provide lots of useful tools under `tools/` directory.

2.4.1 Get the FLOPs and params (experimental)

We provide a script adapted from `flops-counter.pytorch` to compute the FLOPs and params of a given model.

```
python tools/get_flops.py ${CONFIG_FILE} [--shape ${INPUT_SHAPE}]
```

You will get the result like this.

```
=====
Input shape: (3, 224, 224)
Flops: 4.12 GFLOPs
Params: 25.56 M
=====
```

Warning: This tool is still experimental and we do not guarantee that the number is correct. You may well use the result for simple comparisons, but double check it before you adopt it in technical reports or papers.

- FLOPs are related to the input shape while parameters are not. The default input shape is (1, 3, 224, 224).

- Some operators are not counted into FLOPs like GN and custom operators. Refer to `mmcv.cnn.get_model_complexity_info()` for details.

2.4.2 Publish a model

Before you publish a model, you may want to

1. Convert model weights to CPU tensors.
2. Delete the optimizer states.
3. Compute the hash of the checkpoint file and append the hash id to the filename.

```
python tools/convert_models/publish_model.py  ${INPUT_FILENAME}  ${OUTPUT_FILENAME}
```

E.g.,

```
python tools/convert_models/publish_model.py work_dirs/resnet50/latest.pth imagenet_
↪resnet50.pth
```

The final output filename will be `imagenet_resnet50_{date}-{hash id}.pth`.

2.5 Tutorials

Currently, we provide five tutorials for users.

- *learn about config*
- *finetune models*
- *add new dataset*
- *design data pipeline*
- *add new modules*
- *customize schedule*
- *customize runtime settings.*

TUTORIAL 1: LEARN ABOUT CONFIGS

MMClassification mainly uses python files as configs. The design of our configuration file system integrates modularity and inheritance, facilitating users to conduct various experiments. All configuration files are placed in the `configs` folder, which mainly contains the primitive configuration folder of `_base_` and many algorithm folders such as `resnet`, `swin_transformer`, `vision_transformer`, etc.

If you wish to inspect the config file, you may run `python tools/analysis/print_config.py /PATH/TO/CONFIG` to see the complete config.

- *Config File and Checkpoint Naming Convention*
- *Config File Structure*
- *Inherit and Modify Config File*
 - *Use intermediate variables in configs*
 - *Ignore some fields in the base configs*
 - *Use some fields in the base configs*
- *Modify config through script arguments*
- *Import user-defined modules*
- *FAQ*

3.1 Config File and Checkpoint Naming Convention

We follow the below convention to name config files. Contributors are advised to follow the same style. The config file names are divided into four parts: algorithm info, module information, training information and data information. Logically, different parts are concatenated by underscores `'_'`, and words in the same part are concatenated by dashes `'-'`.

```
{algorithm info}_{module info}_{training info}_{data info}.py
```

- `algorithm info`: algorithm information, model name and neural network architecture, such as `resnet`, etc.;
- `module info`: module information is used to represent some special neck, head and pretrain information;
- `training info`: Training information, some training schedule, including batch size, lr schedule, data augment and the like;
- `data info`: Data information, dataset name, input size and so on, such as `imagenet`, `cifar`, etc.;

3.1.1 Algorithm information

The main algorithm name and the corresponding branch architecture information. E.g:

- `resnet50`
- `mobilenet-v3-large`
- `vit-small-patch32` : `patch32` represents the size of the partition in ViT algorithm;
- `seresnext101-32x4d` : SeResNet101 network structure, `32x4d` means that `groups` and `width_per_group` are 32 and 4 respectively in Bottleneck;

3.1.2 Module information

Some special `neck`, `head` and `pretrain` information. In classification tasks, `pretrain` information is the most commonly used:

- `in21k-pre` : pre-trained on ImageNet21k;
- `in21k-pre-3rd-party` : pre-trained on ImageNet21k and the checkpoint is converted from a third-party repository;

3.1.3 Training information

Training schedule, including `training type`, `batch size`, `lr schedule`, `data augment`, `special loss functions` and so on:

- `format {gpu x batch_per_gpu}`, such as `8xb32`

Training type (mainly seen in the transformer network, such as the ViT algorithm, which is usually divided into two training type: pre-training and fine-tuning):

- `ft` : configuration file for fine-tuning
- `pt` : configuration file for pretraining

Training recipe. Usually, only the part that is different from the original paper will be marked. These methods will be arranged in the order `{pipeline aug}`-`{train aug}`-`{loss trick}`-`{scheduler}`-`{epochs}`.

- `coslr-200e` : use cosine scheduler to train 200 epochs
- `autoaug-mixup-lbs-coslr-50e` : use autoaug, mixup, label smooth, cosine scheduler to train 50 epochs

3.1.4 Data information

- `in1k` : ImageNet1k dataset, default to use the input image size of 224x224;
- `in21k` : ImageNet21k dataset, also called ImageNet22k dataset, default to use the input image size of 224x224;
- `in1k-384px` : Indicates that the input image size is 384x384;
- `cifar100`

3.1.5 Config File Name Example

```
repvgg-D2se_deploy_4xb64-autoaug-lbs-mixup-coslr-200e_in1k.py
```

- `repvgg-D2se`: Algorithm information
 - `repvgg`: The main algorithm.
 - `D2se`: The architecture.
- `deploy`: Module information, means the backbone is in the deploy state.
- `4xb64-autoaug-lbs-mixup-coslr-200e`: Training information.
 - `4xb64`: Use 4 GPUs and the size of batches per GPU is 64.
 - `autoaug`: Use AutoAugment in training pipeline.
 - `lbs`: Use label smoothing loss.
 - `mixup`: Use mixup training augment method.
 - `coslr`: Use cosine learning rate scheduler.
 - `200e`: Train the model for 200 epochs.
- `in1k`: Dataset information. The config is for ImageNet 1k dataset and the input size is 224x224.

Note: Some configuration files currently do not follow this naming convention, and related files will be updated in the near future.

3.1.6 Checkpoint Naming Convention

The naming of the weight mainly includes the configuration file name, date and hash value.

```
{config_name}_{date}-{hash}.pth
```

3.2 Config File Structure

There are four kinds of basic component file in the `configs/_base_` folders, namely:

- `models`
- `datasets`
- `schedules`
- `runtime`

You can easily build your own training config file by inherit some base config files. And the configs that are composed by components from `_base_` are called *primitive*.

For easy understanding, we use `ResNet50 primitive config` as a example and comment the meaning of each line. For more detaile, please refer to the API documentation.

```

_base_ = [
    './_base_/models/resnet50.py',          # model
    './_base_/datasets/imagenet_bs32.py',  # data
    './_base_/schedules/imagenet_bs256.py', # training schedule
    './_base_/default_runtime.py'         # runtime setting
]

```

The four parts are explained separately below, and the above-mentioned ResNet50 primitive config are also used as an example.

3.2.1 model

The parameter "model" is a python dictionary in the configuration file, which mainly includes information such as network structure and loss function:

- `type` : Classifier name, MMCIs supports `ImageClassifier`, refer to [API documentation](#).
- `backbone` : Backbone configs, refer to [API documentation](#) for available options.
- `neck` : Neck network name, MMCIs supports `GlobalAveragePooling`, please refer to [API documentation](#).
- `head`: Head network name, MMCIs supports single-label and multi-label classification head networks, available options refer to [API documentation](#).
 - `loss`: Loss function type, supports `CrossEntropyLoss`, `LabelSmoothLoss` etc., For available options, refer to [API documentation](#).
- `train_cfg` : Training augment config, MMCIs supports `mixup`, `cutmix` and other augments.

Note: The 'type' in the configuration file is not a constructed parameter, but a class name.

```

model = dict(
    type='ImageClassifier',          # Classifier name
    backbone=dict(
        type='ResNet',              # Backbones name
        depth=50,                   # depth of backbone, ResNet has options of 18, 34, 50,
↪101, 152.
        num_stages=4,               # number of stages. The feature maps generated by these
↪states are used as the input for the subsequent neck and head.
        out_indices=(3, ),          # The output index of the output feature maps.
        frozen_stages=-1,           # the stage to be frozen, '-1' means not be frozen
        style='pytorch'),           # The style of backbone, 'pytorch' means that stride 2
↪layers are in 3x3 conv, 'caffe' means stride 2 layers are in 1x1 convs.
    neck=dict(type='GlobalAveragePooling'), # neck network name
    head=dict(
        type='LinearClsHead',        # linear classification head,
        num_classes=1000,           # The number of output categories, consistent with the
↪number of categories in the dataset
        in_channels=2048,           # The number of input channels, consistent with the
↪output channel of the neck
        loss=dict(type='CrossEntropyLoss', loss_weight=1.0), # Loss function
↪configuration information
        topk=(1, 5),                # Evaluation index, Top-k accuracy rate, here is the
↪accuracy rate of top1 and top5

```

(continues on next page)

(continued from previous page)

))

3.2.2 data

The parameter "data" is a python dictionary in the configuration file, which mainly includes information to construct dataloader:

- `samples_per_gpu` : the BatchSize of each GPU when building the dataloader
- `workers_per_gpu` : the number of threads per GPU when building dataloader
- `train | val | test` : config to construct dataset
 - `type`: Dataset name, MMCLs supports ImageNet, Cifar etc., refer to [API documentation](#)
 - `data_prefix` : Dataset root directory
 - `pipeline` : Data processing pipeline, refer to related tutorial [CUSTOM DATA PIPELINES](#)

The parameter evaluation is also a dictionary, which is the configuration information of evaluation hook, mainly including evaluation interval, evaluation index, etc..

```
# dataset settings
dataset_type = 'ImageNet' # dataset name,
img_norm_cfg = dict(      # Image normalization config to normalize the input images
    mean=[123.675, 116.28, 103.53], # Mean values used to pre-training the pre-trained
    ↪backbone models
    std=[58.395, 57.12, 57.375],   # Standard variance used to pre-training the pre-
    ↪trained backbone models
    to_rgb=True)                  # Whether to invert the color channel, rgb2bgr or
    ↪bgr2rgb.
# train data pipeline
train_pipeline = [
    dict(type='LoadImageFromFile'), # First pipeline to load images from
    ↪file path
    dict(type='RandomResizedCrop', size=224), # RandomResizedCrop
    dict(type='RandomFlip', flip_prob=0.5, direction='horizontal'), # Randomly flip the
    ↪picture horizontally with a probability of 0.5
    dict(type='Normalize', **img_norm_cfg), # normalization
    dict(type='ImageToTensor', keys=['img']), # convert image from numpy into torch.
    ↪Tensor
    dict(type='ToTensor', keys=['gt_label']), # convert gt_label into torch.Tensor
    dict(type='Collect', keys=['img', 'gt_label']) # Pipeline that decides which keys in
    ↪the data should be passed to the detector
]
# test data pipeline
test_pipeline = [
    dict(type='LoadImageFromFile'),
    dict(type='Resize', size=(256, -1)),
    dict(type='CenterCrop', crop_size=224),
    dict(type='Normalize', **img_norm_cfg),
    dict(type='ImageToTensor', keys=['img']),
    dict(type='Collect', keys=['img']) # do not pass gt_label while testing
]
```

(continues on next page)

(continued from previous page)

```

data = dict(
    samples_per_gpu=32,      # Batch size of a single GPU
    workers_per_gpu=2,      # Worker to pre-fetch data for each single GPU
    train=dict( # Train dataset config
        train=dict( # train data config
            type=dataset_type, # dataset name
            data_prefix='data/imagenet/train', # Dataset root, when ann_file does not exist,
            ↪ the category information is automatically obtained from the root folder
            pipeline=train_pipeline), # train data pipeline
        val=dict( # val data config
            type=dataset_type,
            data_prefix='data/imagenet/val',
            ann_file='data/imagenet/meta/val.txt', # ann_file existes, the category_
            ↪ information is obtained from file
            pipeline=test_pipeline),
        test=dict( # test data config
            type=dataset_type,
            data_prefix='data/imagenet/val',
            ann_file='data/imagenet/meta/val.txt',
            pipeline=test_pipeline))
    evaluation = dict( # The config to build the evaluation hook, refer to https://
            ↪ github.com/open-mmlab/mmdetection/blob/master/mmdet/core/evaluation/eval_hooks.py#L7_
            ↪ for more details.
            interval=1, # Evaluation interval
            metric='accuracy') # Metrics used during evaluation

```

3.2.3 training schedule

Mainly include optimizer settings, optimizer hook settings, learning rate schedule and runner settings:

- `optimizer`: optimizer setting , support all optimizers in pytorch, refer to related `mmcv` documentation.
- `optimizer_config`: optimizer hook configuration file, such as setting gradient limit, refer to related `mmcv` code.
- `lr_config`: Learning rate scheduler, supports “CosineAnnealing”, “Step”, “Cyclic”, etc. refer to related `mmcv` documentation for more options.
- `runner`: For runner, please refer to `mmcv` for `runner` introduction document.

```

# he configuration file used to build the optimizer, support all optimizers in PyTorch.
optimizer = dict(type='SGD', # Optimizer type
                 lr=0.1, # Learning rate of optimizers, see detail usages of_
                 ↪ the parameters in the documentation of PyTorch
                 momentum=0.9, # Momentum
                 weight_decay=0.0001) # Weight decay of SGD
# Config used to build the optimizer hook, refer to https://github.com/open-mmlab/mmcv/
↪ blob/master/mmcv/runner/hooks/optimizer.py#L8 for implementation details.
optimizer_config = dict(grad_clip=None) # Most of the methods do not use gradient clip
# Learning rate scheduler config used to register LrUpdater hook
lr_config = dict(policy='step', # The policy of scheduler, also support_
                 ↪ CosineAnnealing, Cyclic, etc. Refer to details of supported LrUpdater from https://
                 ↪ github.com/open-mmlab/mmcv/blob/master/mmcv/runner/hooks/lr_updater.py#L9.

```

(continues on next page)

(continued from previous page)

```

        step=[30, 60, 90])    # Steps to decay the learning rate
runner = dict(type='EpochBasedRunner', # Type of runner to use (i.e. IterBasedRunner_
↳or EpochBasedRunner)
        max_epochs=100)    # Runner that runs the workflow in total max_epochs. For_
↳IterBasedRunner use `max_iters`

```

3.2.4 runtime setting

This part mainly includes saving the checkpoint strategy, log configuration, training parameters, breakpoint weight path, working directory, etc..

```

# Config to set the checkpoint hook, Refer to https://github.com/open-mmlab/mmcv/blob/
↳master/mmcv/runner/hooks/checkpoint.py for implementation.
checkpoint_config = dict(interval=1)    # The save interval is 1
# config to register logger hook
log_config = dict(
    interval=100,                # Interval to print the log
    hooks=[
        dict(type='TextLoggerHook'),    # The Tensorboard logger is also supported
        # dict(type='TensorboardLoggerHook')
    ])
dist_params = dict(backend='nccl')    # Parameters to setup distributed training, the_
↳port can also be set.
log_level = 'INFO'                # The output level of the log.
resume_from = None                # Resume checkpoints from a given path, the training will_
↳be resumed from the epoch when the checkpoint's is saved.
workflow = [('train', 1)]        # Workflow for runner. [('train', 1)] means there is only_
↳one workflow and the workflow named 'train' is executed once.
work_dir = 'work_dir'            # Directory to save the model checkpoints and logs for_
↳the current experiments.

```

3.3 Inherit and Modify Config File

For easy understanding, we recommend contributors to inherit from existing methods.

For all configs under the same folder, it is recommended to have only **one primitive** config. All other configs should inherit from the *primitive* config. In this way, the maximum of inheritance level is 3.

For example, if your config file is based on ResNet with some other modification, you can first inherit the basic ResNet structure, dataset and other training setting by specifying `_base_ = './resnet50_8xb32_in1k.py'` (The path relative to your config file), and then modify the necessary parameters in the config file. A more specific example, now we want to use almost all configs in `configs/resnet/resnet50_8xb32_in1k.py`, but change the number of training epochs from 100 to 300, modify when to decay the learning rate, and modify the dataset path, you can create a new config file `configs/resnet/resnet50_8xb32-300e_in1k.py` with content as below:

```

_base_ = './resnet50_8xb32_in1k.py'

runner = dict(max_epochs=300)
lr_config = dict(step=[150, 200, 250])

```

(continues on next page)

(continued from previous page)

```

data = dict(
    train=dict(data_prefix='mydata/imagenet/train'),
    val=dict(data_prefix='mydata/imagenet/train', ),
    test=dict(data_prefix='mydata/imagenet/train', )
)

```

3.3.1 Use intermediate variables in configs

Some intermediate variables are used in the configuration file. The intermediate variables make the configuration file clearer and easier to modify.

For example, `train_pipeline / test_pipeline` is the intermediate variable of the data pipeline. We first need to define `train_pipeline / test_pipeline`, and then pass them to `data`. If you want to modify the size of the input image during training and testing, you need to modify the intermediate variables of `train_pipeline / test_pipeline`.

```

img_norm_cfg = dict(
    mean=[123.675, 116.28, 103.53], std=[58.395, 57.12, 57.375], to_rgb=True)
train_pipeline = [
    dict(type='LoadImageFromFile'),
    dict(type='RandomResizedCrop', size=384, backend='pillow'),
    dict(type='RandomFlip', flip_prob=0.5, direction='horizontal'),
    dict(type='Normalize', **img_norm_cfg),
    dict(type='ImageToTensor', keys=['img']),
    dict(type='ToTensor', keys=['gt_label']),
    dict(type='Collect', keys=['img', 'gt_label'])
]
test_pipeline = [
    dict(type='LoadImageFromFile'),
    dict(type='Resize', size=384, backend='pillow'),
    dict(type='Normalize', **img_norm_cfg),
    dict(type='ImageToTensor', keys=['img']),
    dict(type='Collect', keys=['img'])
]
data = dict(
    train=dict(pipeline=train_pipeline),
    val=dict(pipeline=test_pipeline),
    test=dict(pipeline=test_pipeline))

```

3.3.2 Ignore some fields in the base configs

Sometimes, you need to set `_delete_=True` to ignore some domain content in the basic configuration file. You can refer to `mmcv` for more instructions.

The following is an example. If you want to use cosine schedule in the above ResNet50 case, just using inheritance and directly modify it will report `get_unexpected keyword 'step' error`, because the `'step'` field of the basic config in `lr_config` domain information is reserved, and you need to add `_delete_=True` to ignore the content of `lr_config` related fields in the basic configuration file:

```

_base_ = '../..//configs/resnet/resnet50_8xb32_in1k.py'

```

(continues on next page)

(continued from previous page)

```
lr_config = dict(
    _delete_=True,
    policy='CosineAnnealing',
    min_lr=0,
    warmup='linear',
    by_epoch=True,
    warmup_iters=5,
    warmup_ratio=0.1
)
```

3.3.3 Use some fields in the base configs

Sometimes, you may refer to some fields in the `_base_` config, so as to avoid duplication of definitions. You can refer to `mmcv` for some more instructions.

The following is an example of using auto augment in the training data preprocessing pipeline, refer to `configs/_base_/datasets/imagenet_bs64_autoaug.py`. When defining `train_pipeline`, just add the definition file name of auto augment to `_base_`, and then use `{{_base_.auto_increasing_policies}}` to reference the variables:

```
_base_ = ['./pipelines/auto_aug.py']

# dataset settings
dataset_type = 'ImageNet'
img_norm_cfg = dict(
    mean=[123.675, 116.28, 103.53], std=[58.395, 57.12, 57.375], to_rgb=True)
train_pipeline = [
    dict(type='LoadImageFromFile'),
    dict(type='RandomResizedCrop', size=224),
    dict(type='RandomFlip', flip_prob=0.5, direction='horizontal'),
    dict(type='AutoAugment', policies={{_base_.auto_increasing_policies}}),
    dict(type='Normalize', **img_norm_cfg),
    dict(type='ImageToTensor', keys=['img']),
    dict(type='ToTensor', keys=['gt_label']),
    dict(type='Collect', keys=['img', 'gt_label'])
]
test_pipeline = [...]
data = dict(
    samples_per_gpu=64,
    workers_per_gpu=2,
    train=dict(..., pipeline=train_pipeline),
    val=dict(..., pipeline=test_pipeline))
evaluation = dict(interval=1, metric='accuracy')
```

3.4 Modify config through script arguments

When users use the script “tools/train.py” or “tools/test.py” to submit tasks or use some other tools, they can directly modify the content of the configuration file used by specifying the `--cfg-options` parameter.

- Update config keys of dict chains.

The config options can be specified following the order of the dict keys in the original config. For example, `--cfg-options model.backbone.norm_eval=False` changes the all BN modules in model backbones to train mode.

- Update keys inside a list of configs.

Some config dicts are composed as a list in your config. For example, the training pipeline `data.train.pipeline` is normally a list e.g. `[dict(type='LoadImageFromFile'), dict(type='TopDownRandomFlip', flip_prob=0.5), ...]`. If you want to change `'flip_prob=0.5'` to `'flip_prob=0.0'` in the pipeline, you may specify `--cfg-options data.train.pipeline.1.flip_prob=0.0`.

- Update values of list/tuples.

If the value to be updated is a list or a tuple. For example, the config file normally sets `workflow=[('train', 1)]`. If you want to change this key, you may specify `--cfg-options workflow="[(train,1), (val,1)]"`. Note that the quotation mark “” is necessary to support list/tuple data types, and that **NO** white space is allowed inside the quotation marks in the specified value.

3.5 Import user-defined modules

Note: This part may only be used when using MMClassification as a third party library to build your own project, and beginners can skip it.

After studying the follow-up tutorials [ADDING NEW DATASET](#), [CUSTOM DATA PIPELINES](#), [ADDING NEW MODULES](#). You may use MMClassification to complete your project and create new classes of datasets, models, data enhancements, etc. in the project. In order to streamline the code, you can use MMClassification as a third-party library, you just need to keep your own extra code and import your own custom module in the configuration files. For examples, you may refer to [OpenMMLab Algorithm Competition Project](#) .

Add the following code to your own configuration files:

```
custom_imports = dict(
    imports=['your_dataset_class',
            'your_transforme_class',
            'your_model_class',
            'your_module_class'],
    allow_failed_imports=False)
```

3.6 FAQ

- None

TUTORIAL 2: FINE-TUNE MODELS

Classification models pre-trained on the ImageNet dataset have been demonstrated to be effective for other datasets and other downstream tasks. This tutorial provides instructions for users to use the models provided in the *Model Zoo* for other datasets to obtain better performance.

There are two steps to fine-tune a model on a new dataset.

- Add support for the new dataset following *Tutorial 2: Adding New Dataset*.
- Modify the configs as will be discussed in this tutorial.

Assume we have a ResNet-50 model pre-trained on the ImageNet-2012 dataset and want to take the fine-tuning on the CIFAR-10 dataset, we need to modify five parts in the config.

4.1 Inherit base configs

At first, create a new config file `configs/tutorial/resnet50_finetune_cifar.py` to store our configs. Of course, the path can be customized by yourself.

To reuse the common parts among different configs, we support inheriting configs from multiple existing configs. To fine-tune a ResNet-50 model, the new config needs to inherit `configs/_base_/models/resnet50.py` to build the basic structure of the model. To use the CIFAR-10 dataset, the new config can also simply inherit `configs/_base_/datasets/cifar10_bs16.py`. For runtime settings such as training schedules, the new config needs to inherit `configs/_base_/default_runtime.py`.

To inherit all above configs, put the following code at the config file.

```
_base_ = [  
    '../_base_/models/resnet50.py',  
    '../_base_/datasets/cifar10_bs16.py', '../_base_/default_runtime.py'  
]
```

Besides, you can also choose to write the whole contents rather than use inheritance, like `configs/lenet/lenet5_mnist.py`.

4.2 Modify model

When fine-tuning a model, usually we want to load the pre-trained backbone weights and train a new classification head.

To load the pre-trained backbone, we need to change the initialization config of the backbone and use `Pretrained` initialization function. Besides, in the `init_cfg`, we use `prefix='backbone'` to tell the initialization function to remove the prefix of keys in the checkpoint, for example, it will change `backbone.conv1` to `conv1`. And here we use an online checkpoint, it will be downloaded during training, you can also download the model manually and use a local path.

And then we need to modify the head according to the class numbers of the new datasets by just changing `num_classes` in the head.

```
model = dict(
    backbone=dict(
        init_cfg=dict(
            type='Pretrained',
            checkpoint='https://download.openmmlab.com/mmcclassification/v0/resnet/
↪resnet50_8xb32_in1k_20210831-ea4938fc.pth',
            prefix='backbone',
        )),
    head=dict(num_classes=10),
)
```

Tip: Here we only need to set the part of configs we want to modify, because the inherited configs will be merged and get the entire configs.

Sometimes, we want to freeze the first several layers' parameters of the backbone, that will help the network to keep ability to extract low-level information learnt from pre-trained model. In MMClassification, you can simply specify how many layers to freeze by `frozen_stages` argument. For example, to freeze the first two layers' parameters, just use the following config:

```
model = dict(
    backbone=dict(
        frozen_stages=2,
        init_cfg=dict(
            type='Pretrained',
            checkpoint='https://download.openmmlab.com/mmcclassification/v0/resnet/
↪resnet50_8xb32_in1k_20210831-ea4938fc.pth',
            prefix='backbone',
        )),
    head=dict(num_classes=10),
)
```

Note: Not all backbones support the `frozen_stages` argument by now. Please check [the docs](#) to confirm if your backbone supports it.

4.3 Modify dataset

When fine-tuning on a new dataset, usually we need to modify some dataset configs. Here, we need to modify the pipeline to resize the image from 32 to 224 to fit the input size of the model pre-trained on ImageNet, and some other configs.

```
img_norm_cfg = dict(
    mean=[125.307, 122.961, 113.8575],
    std=[51.5865, 50.847, 51.255],
    to_rgb=False,
)
train_pipeline = [
    dict(type='RandomCrop', size=32, padding=4),
    dict(type='RandomFlip', flip_prob=0.5, direction='horizontal'),
    dict(type='Resize', size=224),
    dict(type='Normalize', **img_norm_cfg),
    dict(type='ImageToTensor', keys=['img']),
    dict(type='ToTensor', keys=['gt_label']),
    dict(type='Collect', keys=['img', 'gt_label']),
]
test_pipeline = [
    dict(type='Resize', size=224),
    dict(type='Normalize', **img_norm_cfg),
    dict(type='ImageToTensor', keys=['img']),
    dict(type='Collect', keys=['img']),
]
data = dict(
    train=dict(pipeline=train_pipeline),
    val=dict(pipeline=test_pipeline),
    test=dict(pipeline=test_pipeline),
)
```

4.4 Modify training schedule

The fine-tuning hyper parameters vary from the default schedule. It usually requires smaller learning rate and less training epochs.

```
# lr is set for a batch size of 128
optimizer = dict(type='SGD', lr=0.01, momentum=0.9, weight_decay=0.0001)
optimizer_config = dict(grad_clip=None)
# learning policy
lr_config = dict(policy='step', step=[15])
runner = dict(type='EpochBasedRunner', max_epochs=200)
log_config = dict(interval=100)
```

4.5 Start Training

Now, we have finished the fine-tuning config file as following:

```

_base_ = [
    './_base_/models/resnet50.py',
    './_base_/datasets/cifar10_bs16.py', './_base_/default_runtime.py'
]

# Model config
model = dict(
    backbone=dict(
        frozen_stages=2,
        init_cfg=dict(
            type='Pretrained',
            checkpoint='https://download.openmmlab.com/mmlclassification/v0/resnet/
↪resnet50_8xb32_in1k_20210831-ea4938fc.pth',
            prefix='backbone',
        )),
    head=dict(num_classes=10),
)

# Dataset config
img_norm_cfg = dict(
    mean=[125.307, 122.961, 113.8575],
    std=[51.5865, 50.847, 51.255],
    to_rgb=False,
)

train_pipeline = [
    dict(type='RandomCrop', size=32, padding=4),
    dict(type='RandomFlip', flip_prob=0.5, direction='horizontal'),
    dict(type='Resize', size=224),
    dict(type='Normalize', **img_norm_cfg),
    dict(type='ImageToTensor', keys=['img']),
    dict(type='ToTensor', keys=['gt_label']),
    dict(type='Collect', keys=['img', 'gt_label']),
]

test_pipeline = [
    dict(type='Resize', size=224),
    dict(type='Normalize', **img_norm_cfg),
    dict(type='ImageToTensor', keys=['img']),
    dict(type='Collect', keys=['img']),
]

data = dict(
    train=dict(pipeline=train_pipeline),
    val=dict(pipeline=test_pipeline),
    test=dict(pipeline=test_pipeline),
)

# Training schedule config
# lr is set for a batch size of 128
optimizer = dict(type='SGD', lr=0.01, momentum=0.9, weight_decay=0.0001)
optimizer_config = dict(grad_clip=None)

```

(continues on next page)

(continued from previous page)

```
# learning policy
lr_config = dict(policy='step', step=[15])
runner = dict(type='EpochBasedRunner', max_epochs=200)
log_config = dict(interval=100)
```

Here we use 8 GPUs on your computer to train the model with the following command:

```
bash tools/dist_train.sh configs/tutorial/resnet50_finetune_cifar.py 8
```

Also, you can use only one GPU to train the model with the following command:

```
python tools/train.py configs/tutorial/resnet50_finetune_cifar.py
```

But wait, an important config need to be changed if using one GPU. We need to change the dataset config as following:

```
data = dict(
    samples_per_gpu=128,
    train=dict(pipeline=train_pipeline),
    val=dict(pipeline=test_pipeline),
    test=dict(pipeline=test_pipeline),
)
```

It's because our training schedule is for a batch size of 128. If using 8 GPUs, just use `samples_per_gpu=16` config in the base config file, and the total batch size will be 128. But if using one GPU, you need to change it to 128 manually to match the training schedule.

TUTORIAL 3: ADDING NEW DATASET

5.1 Customize datasets by reorganizing data

5.1.1 Reorganize dataset to existing format

The simplest way is to convert your dataset to existing dataset formats (ImageNet).

For training, it differentiates classes by folders. The directory of training data is as follows:

```
imagenet
├── ...
├── train
│   ├── n01440764
│   │   ├── n01440764_10026.JPEG
│   │   ├── n01440764_10027.JPEG
│   │   └── ...
│   ├── ...
│   └── n15075141
│       ├── n15075141_999.JPEG
│       ├── n15075141_9993.JPEG
│       └── ...
```

For validation, we provide an annotation list. Each line of the list contains a filename and its corresponding ground-truth labels. The format is as follows:

```
ILSVRC2012_val_00000001.JPEG 65
ILSVRC2012_val_00000002.JPEG 970
ILSVRC2012_val_00000003.JPEG 230
ILSVRC2012_val_00000004.JPEG 809
ILSVRC2012_val_00000005.JPEG 516
```

Note: The value of ground-truth labels should fall in range $[0, \text{num_classes} - 1]$.

5.1.2 An example of customized dataset

You can write a new Dataset class inherited from BaseDataset, and overwrite `load_annotations(self)`, like CIFAR10 and ImageNet. Typically, this function returns a list, where each sample is a dict, containing necessary data information, e.g., `img` and `gt_label`.

Assume we are going to implement a Filelist dataset, which takes filelists for both training and testing. The format of annotation list is as follows:

```
000001.jpg 0
000002.jpg 1
```

We can create a new dataset in `mmcls/datasets/filelist.py` to load the data.

```
import mmcv
import numpy as np

from .builder import DATASETS
from .base_dataset import BaseDataset

@DATASETS.register_module()
class Filelist(BaseDataset):

    def load_annotations(self):
        assert isinstance(self.ann_file, str)

        data_infos = []
        with open(self.ann_file) as f:
            samples = [x.strip().split(' ') for x in f.readlines()]
            for filename, gt_label in samples:
                info = {'img_prefix': self.data_prefix}
                info['img_info'] = {'filename': filename}
                info['gt_label'] = np.array(gt_label, dtype=np.int64)
                data_infos.append(info)
        return data_infos
```

And add this dataset class in `mmcls/datasets/__init__.py`

```
from .base_dataset import BaseDataset
...
from .filelist import Filelist

__all__ = [
    'BaseDataset', ... , 'Filelist'
]
```

Then in the config, to use Filelist you can modify the config as the following

```
train = dict(
    type='Filelist',
    ann_file = 'image_list.txt',
    pipeline=train_pipeline
)
```

5.2 Customize datasets by mixing dataset

MMClassification also supports to mix dataset for training. Currently it supports to concat and repeat datasets.

5.2.1 Repeat dataset

We use RepeatDataset as wrapper to repeat the dataset. For example, suppose the original dataset is Dataset_A, to repeat it, the config looks like the following

```
dataset_A_train = dict(  
    type='RepeatDataset',  
    times=N,  
    dataset=dict( # This is the original config of Dataset_A  
        type='Dataset_A',  
        ...  
        pipeline=train_pipeline  
    )  
)
```

5.2.2 Class balanced dataset

We use ClassBalancedDataset as wrapper to repeat the dataset based on category frequency. The dataset to repeat needs to instantiate function `self.get_cat_ids(idx)` to support ClassBalancedDataset. For example, to repeat Dataset_A with `oversample_thr=1e-3`, the config looks like the following

```
dataset_A_train = dict(  
    type='ClassBalancedDataset',  
    oversample_thr=1e-3,  
    dataset=dict( # This is the original config of Dataset_A  
        type='Dataset_A',  
        ...  
        pipeline=train_pipeline  
    )  
)
```

You may refer to [source code](#) for details.

TUTORIAL 4: CUSTOM DATA PIPELINES

6.1 Design of Data pipelines

Following typical conventions, we use `Dataset` and `DataLoader` for data loading with multiple workers. Indexing `Dataset` returns a dict of data items corresponding to the arguments of models forward method.

The data preparation pipeline and the dataset is decomposed. Usually a dataset defines how to process the annotations and a data pipeline defines all the steps to prepare a data dict. A pipeline consists of a sequence of operations. Each operation takes a dict as input and also output a dict for the next transform.

The operations are categorized into data loading, pre-processing and formatting.

Here is an pipeline example for ResNet-50 training on ImageNet.

```
img_norm_cfg = dict(
    mean=[123.675, 116.28, 103.53], std=[58.395, 57.12, 57.375], to_rgb=True)
train_pipeline = [
    dict(type='LoadImageFromFile'),
    dict(type='RandomResizedCrop', size=224),
    dict(type='RandomFlip', flip_prob=0.5, direction='horizontal'),
    dict(type='Normalize', **img_norm_cfg),
    dict(type='ImageToTensor', keys=['img']),
    dict(type='ToTensor', keys=['gt_label']),
    dict(type='Collect', keys=['img', 'gt_label'])
]
test_pipeline = [
    dict(type='LoadImageFromFile'),
    dict(type='Resize', size=256),
    dict(type='CenterCrop', crop_size=224),
    dict(type='Normalize', **img_norm_cfg),
    dict(type='ImageToTensor', keys=['img']),
    dict(type='Collect', keys=['img'])
]
```

For each operation, we list the related dict fields that are added/updated/removed. At the end of the pipeline, we use `Collect` to only retain the necessary items for forward computation.

6.1.1 Data loading

LoadImageFromFile

- add: img, img_shape, ori_shape

By default, `LoadImageFromFile` loads images from disk but it may lead to IO bottleneck for efficient small models. Various backends are supported by `mmcv` to accelerate this process. For example, if the training machines have setup `memcached`, we can revise the config as follows.

```
memcached_root = '/mnt/xxx/memcached_client/'
train_pipeline = [
    dict(
        type='LoadImageFromFile',
        file_client_args=dict(
            backend='memcached',
            server_list_cfg=osp.join(memcached_root, 'server_list.conf'),
            client_cfg=osp.join(memcached_root, 'client.conf')),
    ]
```

More supported backends can be found in `mmcv.fileio.FileClient`.

6.1.2 Pre-processing

Resize

- add: scale, scale_idx, pad_shape, scale_factor, keep_ratio
- update: img, img_shape

RandomFlip

- add: flip, flip_direction
- update: img

RandomCrop

- update: img, pad_shape

Normalize

- add: img_norm_cfg
- update: img

6.1.3 Formatting

ToTensor

- update: specified by keys.

ImageToTensor

- update: specified by keys.

Collect

- remove: all other keys except for those specified by keys

6.2 Extend and use custom pipelines

1. Write a new pipeline in any file, e.g., `my_pipeline.py`, and place it in the folder `mmcls/datasets/pipelines/`. The pipeline class needs to override the `__call__` method which takes a dict as input and returns a dict.

```
from mmcls.datasets import PIPELINES

@PIPELINES.register_module()
class MyTransform(object):

    def __call__(self, results):
        # apply transforms on results['img']
        return results
```

2. Import the new class in `mmcls/datasets/pipelines/__init__.py`.

```
...
from .my_pipeline import MyTransform

__all__ = [
    ..., 'MyTransform'
]
```

3. Use it in config files.

```
img_norm_cfg = dict(
    mean=[123.675, 116.28, 103.53], std=[58.395, 57.12, 57.375], to_rgb=True)
train_pipeline = [
    dict(type='LoadImageFromFile'),
    dict(type='RandomResizedCrop', size=224),
    dict(type='RandomFlip', flip_prob=0.5, direction='horizontal'),
    dict(type='MyTransform'),
    dict(type='Normalize', **img_norm_cfg),
    dict(type='ImageToTensor', keys=['img']),
    dict(type='ToTensor', keys=['gt_label']),
    dict(type='Collect', keys=['img', 'gt_label'])
]
```

6.3 Pipeline visualization

After designing data pipelines, you can use the *visualization tools* to view the performance.

TUTORIAL 5: ADDING NEW MODULES

7.1 Develop new components

We basically categorize model components into 3 types.

- backbone: usually an feature extraction network, e.g., ResNet, MobileNet.
- neck: the component between backbones and heads, e.g., GlobalAveragePooling.
- head: the component for specific tasks, e.g., classification or regression.

7.1.1 Add new backbones

Here we show how to develop new components with an example of ResNet_CIFAR. As the input size of CIFAR is 32x32, this backbone replaces the `kernel_size=7, stride=2` to `kernel_size=3, stride=1` and remove the MaxPooling after stem, to avoid forwarding small feature maps to residual blocks. It inherits from ResNet and only modifies the stem layers.

1. Create a new file `mmcls/models/backbones/resnet_cifar.py`.

```
import torch.nn as nn

from ..builder import BACKBONES
from .resnet import ResNet

@BACKBONES.register_module()
class ResNet_CIFAR(ResNet):

    """ResNet backbone for CIFAR.

    short description of the backbone

    Args:
        depth(int): Network depth, from {18, 34, 50, 101, 152}.
        ...
    """

    def __init__(self, depth, deep_stem, **kwargs):
        # call ResNet init
        super(ResNet_CIFAR, self).__init__(depth, deep_stem=deep_stem, **kwargs)
        # other specific initialization
```

(continues on next page)

(continued from previous page)

```

assert not self.deep_stem, 'ResNet_CIFAR do not support deep_stem'

def _make_stem_layer(self, in_channels, base_channels):
    # override ResNet method to modify the network structure
    self.conv1 = build_conv_layer(
        self.conv_cfg,
        in_channels,
        base_channels,
        kernel_size=3,
        stride=1,
        padding=1,
        bias=False)
    self.norm1_name, norm1 = build_norm_layer(
        self.norm_cfg, base_channels, postfix=1)
    self.add_module(self.norm1_name, norm1)
    self.relu = nn.ReLU(inplace=True)

def forward(self, x): # should return a tuple
    pass # implementation is ignored

def init_weights(self, pretrained=None):
    pass # override ResNet init_weights if necessary

def train(self, mode=True):
    pass # override ResNet train if necessary

```

1. Import the module in `mmcls/models/backbones/__init__.py`.

```

...
from .resnet_cifar import ResNet_CIFAR

__all__ = [
    ..., 'ResNet_CIFAR'
]

```

1. Use it in your config file.

```

model = dict(
    ...
    backbone=dict(
        type='ResNet_CIFAR',
        depth=18,
        other_arg=xxx),
    ...

```

7.1.2 Add new necks

Here we take `GlobalAveragePooling` as an example. It is a very simple neck without any arguments. To add a new neck, we mainly implement the `forward` function, which applies some operation on the output from backbone and forward the results to head.

1. Create a new file in `mmcls/models/necks/gap.py`.

```
import torch.nn as nn

from ..builder import NECKS

@NECKS.register_module()
class GlobalAveragePooling(nn.Module):

    def __init__(self):
        self.gap = nn.AdaptiveAvgPool2d((1, 1))

    def forward(self, inputs):
        # we regard inputs as tensor for simplicity
        outs = self.gap(inputs)
        outs = outs.view(inputs.size(0), -1)
        return outs
```

2. Import the module in `mmcls/models/necks/__init__.py`.

```
...
from .gap import GlobalAveragePooling

__all__ = [
    ..., 'GlobalAveragePooling'
]
```

3. Modify the config file.

```
model = dict(
    neck=dict(type='GlobalAveragePooling'),
)
```

7.1.3 Add new heads

Here we show how to develop a new head with the example of `LinearClsHead` as the following. To implement a new head, basically we need to implement `forward_train`, which takes the feature maps from necks or backbones as input and compute loss based on ground-truth labels.

1. Create a new file in `mmcls/models/heads/linear_head.py`.

```
from ..builder import HEADS
from .cls_head import ClsHead

@HEADS.register_module()
class LinearClsHead(ClsHead):
```

(continues on next page)

(continued from previous page)

```

def __init__(self,
              num_classes,
              in_channels,
              loss=dict(type='CrossEntropyLoss', loss_weight=1.0),
              topk=(1, )):
    super(LinearClsHead, self).__init__(loss=loss, topk=topk)
    self.in_channels = in_channels
    self.num_classes = num_classes

    if self.num_classes <= 0:
        raise ValueError(
            f'num_classes={num_classes} must be a positive integer')

    self._init_layers()

def _init_layers(self):
    self.fc = nn.Linear(self.in_channels, self.num_classes)

def init_weights(self):
    normal_init(self.fc, mean=0, std=0.01, bias=0)

def forward_train(self, x, gt_label):
    cls_score = self.fc(x)
    losses = self.loss(cls_score, gt_label)
    return losses

```

2. Import the module in `mmcls/models/heads/__init__.py`.

```

...
from .linear_head import LinearClsHead

__all__ = [
    ..., 'LinearClsHead'
]

```

3. Modify the config file.

Together with the added GlobalAveragePooling neck, an entire config for a model is as follows.

```

model = dict(
    type='ImageClassifier',
    backbone=dict(
        type='ResNet',
        depth=50,
        num_stages=4,
        out_indices=(3, ),
        style='pytorch'),
    neck=dict(type='GlobalAveragePooling'),
    head=dict(
        type='LinearClsHead',
        num_classes=1000,
        in_channels=2048,

```

(continues on next page)

(continued from previous page)

```

    loss=dict(type='CrossEntropyLoss', loss_weight=1.0),
    topk=(1, 5),
))

```

7.1.4 Add new loss

To add a new loss function, we mainly implement the forward function in the loss module. In addition, it is helpful to leverage the decorator `weighted_loss` to weight the loss for each element. Assuming that we want to mimic a probabilistic distribution generated from another classification model, we implement a `L1Loss` to fulfil the purpose as below.

1. Create a new file in `mmcls/models/losses/l1_loss.py`.

```

import torch
import torch.nn as nn

from ..builder import LOSSES
from .utils import weighted_loss

@weighted_loss
def l1_loss(pred, target):
    assert pred.size() == target.size() and target.numel() > 0
    loss = torch.abs(pred - target)
    return loss

@LOSSES.register_module()
class L1Loss(nn.Module):

    def __init__(self, reduction='mean', loss_weight=1.0):
        super(L1Loss, self).__init__()
        self.reduction = reduction
        self.loss_weight = loss_weight

    def forward(self,
                pred,
                target,
                weight=None,
                avg_factor=None,
                reduction_override=None):
        assert reduction_override in (None, 'none', 'mean', 'sum')
        reduction = (
            reduction_override if reduction_override else self.reduction)
        loss = self.loss_weight * l1_loss(
            pred, target, weight, reduction=reduction, avg_factor=avg_factor)
        return loss

```

2. Import the module in `mmcls/models/losses/__init__.py`.

```

...
from .l1_loss import L1Loss, l1_loss

```

(continues on next page)

(continued from previous page)

```
__all__ = [  
    ..., 'L1Loss', 'l1_loss'  
]
```

3. Modify loss field in the config.

```
loss=dict(type='L1Loss', loss_weight=1.0))
```


TUTORIAL 6: CUSTOMIZE SCHEDULE

In this tutorial, we will introduce some methods about how to construct optimizers, customize learning rate and momentum schedules, parameter-wise finely configuration, gradient clipping, gradient accumulation, and customize self-implemented methods for the project.

- *Customize optimizer supported by PyTorch*
- *Customize learning rate schedules*
 - *Learning rate decay*
 - *Warmup strategy*
- *Customize momentum schedules*
- *Parameter-wise finely configuration*
- *Gradient clipping and gradient accumulation*
 - *Gradient clipping*
 - *Gradient accumulation*
- *Customize self-implemented methods*
 - *Customize self-implemented optimizer*
 - *Customize optimizer constructor*

8.1 Customize optimizer supported by PyTorch

We already support to use all the optimizers implemented by PyTorch, and to use and modify them, please change the `optimizer` field of config files.

For example, if you want to use SGD, the modification could be as the following.

```
optimizer = dict(type='SGD', lr=0.0003, weight_decay=0.0001)
```

To modify the learning rate of the model, just modify the `lr` in the config of optimizer. You can also directly set other arguments according to the [API doc](#) of PyTorch.

For example, if you want to use Adam with the setting like `torch.optim.Adam(params, lr=0.001, betas=(0.9, 0.999), eps=1e-08, weight_decay=0, amsgrad=False)` in PyTorch, the config should looks like.

```
optimizer = dict(type='Adam', lr=0.001, betas=(0.9, 0.999), eps=1e-08, weight_decay=0, ↵  
↵amsgrad=False)
```

8.2 Customize learning rate schedules

8.2.1 Learning rate decay

Learning rate decay is widely used to improve performance. And to use learning rate decay, please set the `lr_config` field in config files.

For example, we use step policy as the default learning rate decay policy of ResNet, and the config is:

```
lr_config = dict(policy='step', step=[100, 150])
```

Then during training, the program will call `StepLRHook` periodically to update the learning rate.

We also support many other learning rate schedules [here](#), such as `CosineAnnealing` and `Poly` schedule. Here are some examples

- `CosineAnnealing` schedule:

```
lr_config = dict(
    policy='CosineAnnealing',
    warmup='linear',
    warmup_iters=1000,
    warmup_ratio=1.0 / 10,
    min_lr_ratio=1e-5)
```

- `Poly` schedule:

```
lr_config = dict(policy='poly', power=0.9, min_lr=1e-4, by_epoch=False)
```

8.2.2 Warmup strategy

In the early stage, training is easy to be volatile, and warmup is a technique to reduce volatility. With warmup, the learning rate will increase gradually from a minor value to the expected value.

In MMClassification, we use `lr_config` to configure the warmup strategy, the main parameters are as follows:

- `warmup`: The warmup curve type. Please choose one from 'constant', 'linear', 'exp' and None, and None means disable warmup.
- `warmup_by_epoch`: if warmup by epoch or not, default to be True, if set to be False, warmup by iter.
- `warmup_iters`: the number of warm-up iterations, when `warmup_by_epoch=True`, the unit is epoch; when `warmup_by_epoch=False`, the unit is the number of iterations (iter).
- `warmup_ratio`: warm-up initial learning rate will calculate as `lr = lr * warmup_ratio`.

Here are some examples

1. linear & warmup by iter

```
lr_config = dict(
    policy='CosineAnnealing',
    by_epoch=False,
    min_lr_ratio=1e-2,
    warmup='linear',
    warmup_ratio=1e-3,
```

(continues on next page)

(continued from previous page)

```
warmup_iters=20 * 1252,  
warmup_by_epoch=False)
```

2. exp & warmup by epoch

```
lr_config = dict(  
    policy='CosineAnnealing',  
    min_lr=0,  
    warmup='exp',  
    warmup_iters=5,  
    warmup_ratio=0.1,  
    warmup_by_epoch=True)
```

Tip: After completing your configuration file, you could use [learning rate visualization tool](#) to draw the corresponding learning rate adjustment curve.

8.3 Customize momentum schedules

We support the momentum scheduler to modify the model's momentum according to learning rate, which could make the model converge in a faster way.

Momentum scheduler is usually used with LR scheduler, for example, the following config is used to accelerate convergence. For more details, please refer to the implementation of [CyclicLrUpdater](#) and [CyclicMomentumUpdater](#).

Here is an example

```
lr_config = dict(  
    policy='cyclic',  
    target_ratio=(10, 1e-4),  
    cyclic_times=1,  
    step_ratio_up=0.4,  
)  
momentum_config = dict(  
    policy='cyclic',  
    target_ratio=(0.85 / 0.95, 1),  
    cyclic_times=1,  
    step_ratio_up=0.4,  
)
```

8.4 Parameter-wise finely configuration

Some models may have some parameter-specific settings for optimization, for example, no weight decay to the Batch-Norm layer or using different learning rates for different network layers. To finely configuration them, we can use the `paramwise_cfg` option in optimizer.

We provide some examples here and more usages refer to [DefaultOptimizerConstructor](#).

- Using specified options

The `DefaultOptimizerConstructor` provides options including `bias_lr_mult`, `bias_decay_mult`, `norm_decay_mult`, `dwconv_decay_mult`, `dcn_offset_lr_mult` and `bypass_duplicate` to configure special optimizer behaviors of bias, normalization, depth-wise convolution, deformable convolution and duplicated parameter. E.g:

1. No weight decay to the BatchNorm layer

```
optimizer = dict(  
    type='SGD',  
    lr=0.8,  
    weight_decay=1e-4,  
    paramwise_cfg=dict(norm_decay_mult=0.))
```

- Using `custom_keys` dict

MMClassification can use `custom_keys` to specify different parameters to use different learning rates or weight decays, for example:

1. No weight decay for specific parameters

```
paramwise_cfg = dict(  
    custom_keys={  
        'backbone.cls_token': dict(decay_mult=0.0),  
        'backbone.pos_embed': dict(decay_mult=0.0)  
    })  
  
optimizer = dict(  
    type='SGD',  
    lr=0.8,  
    weight_decay=1e-4,  
    paramwise_cfg=paramwise_cfg)
```

1. Using a smaller learning rate and a weight decay for the backbone layers

```
optimizer = dict(  
    type='SGD',  
    lr=0.8,  
    weight_decay=1e-4,  
    # 'lr' for backbone and 'weight_decay' are 0.1 * lr and 0.9 * weight_decay  
    paramwise_cfg=dict(  
        custom_keys={'backbone': dict(lr_mult=0.1, decay_mult=0.9)}))
```

8.5 Gradient clipping and gradient accumulation

Besides the basic function of PyTorch optimizers, we also provide some enhancement functions, such as gradient clipping, gradient accumulation, etc., refer to [MMCV](#).

8.5.1 Gradient clipping

During the training process, the loss function may get close to a cliffy region and cause gradient explosion. And gradient clipping is helpful to stabilize the training process. More introduction can be found in [this page](#).

Currently we support `grad_clip` option in `optimizer_config`, and the arguments refer to [PyTorch Documentation](#).

Here is an example:

```
optimizer_config = dict(grad_clip=dict(max_norm=35, norm_type=2))
# norm_type: type of the used p-norm, here norm_type is 2.
```

When inheriting from base and modifying configs, if `grad_clip=None` in base, `_delete_=True` is needed. For more details about `_delete_` you can refer to [TUTORIAL 1: LEARN ABOUT CONFIGS](#). For example,

```
_base_ = [./_base_/schedules/imagenet_bs256_coslr.py]

optimizer_config = dict(grad_clip=dict(max_norm=35, norm_type=2), _delete_=True, type=
↪ 'OptimizerHook')
# you can ignore type if type is 'OptimizerHook', otherwise you must add "type=
↪ 'xxxxxOptimizerHook'" here
```

8.5.2 Gradient accumulation

When computing resources are lacking, the batch size can only be set to a small value, which may affect the performance of models. Gradient accumulation can be used to solve this problem.

Here is an example:

```
data = dict(samples_per_gpu=64)
optimizer_config = dict(type="GradientCumulativeOptimizerHook", cumulative_iters=4)
```

Indicates that during training, back-propagation is performed every 4 iters. And the above is equivalent to:

```
data = dict(samples_per_gpu=256)
optimizer_config = dict(type="OptimizerHook")
```

Note: When the optimizer hook type is not specified in `optimizer_config`, `OptimizerHook` is used by default.

8.6 Customize self-implemented methods

In academic research and industrial practice, it may be necessary to use optimization methods not implemented by MMClassification, and you can add them through the following methods.

Note: This part will modify the MMClassification source code or add code to the MMClassification framework, beginners can skip it.

8.6.1 Customize self-implemented optimizer

1. Define a new optimizer

A customized optimizer could be defined as below.

Assume you want to add an optimizer named MyOptimizer, which has arguments a, b, and c. You need to create a new directory named `mmcls/core/optimizer`. And then implement the new optimizer in a file, e.g., in `mmcls/core/optimizer/my_optimizer.py`:

```
from mmcv.runner import OPTIMIZERS
from torch.optim import Optimizer

@OPTIMIZERS.register_module()
class MyOptimizer(Optimizer):

    def __init__(self, a, b, c):
```

2. Add the optimizer to registry

To find the above module defined above, this module should be imported into the main namespace at first. There are two ways to achieve it.

- Modify `mmcls/core/optimizer/__init__.py` to import it into optimizer package, and then modify `mmcls/core/__init__.py` to import the new optimizer package.

Create the `mmcls/core/optimizer` folder and the `mmcls/core/optimizer/__init__.py` file if they don't exist. The newly defined module should be imported in `mmcls/core/optimizer/__init__.py` and `mmcls/core/__init__.py` so that the registry will find the new module and add it:

```
# In mmcls/core/optimizer/__init__.py
from .my_optimizer import MyOptimizer # MyOptimizer maybe other class name

__all__ = ['MyOptimizer']
```

```
# In mmcls/core/__init__.py
...
from .optimizer import * # noqa: F401, F403
```

- Use `custom_imports` in the config to manually import it

```
custom_imports = dict(imports=['mmcls.core.optimizer.my_optimizer'], allow_failed_
↳ imports=False)
```

The module `mmcls.core.optimizer.my_optimizer` will be imported at the beginning of the program and the class `MyOptimizer` is then automatically registered. Note that only the package containing the class `MyOptimizer` should be imported. `mmcls.core.optimizer.my_optimizer.MyOptimizer` **cannot** be imported directly.

3. Specify the optimizer in the config file

Then you can use `MyOptimizer` in `optimizer` field of config files. In the configs, the optimizers are defined by the field `optimizer` like the following:

```
optimizer = dict(type='SGD', lr=0.02, momentum=0.9, weight_decay=0.0001)
```

To use your own optimizer, the field can be changed to

```
optimizer = dict(type='MyOptimizer', a=a_value, b=b_value, c=c_value)
```

8.6.2 Customize optimizer constructor

Some models may have some parameter-specific settings for optimization, e.g. weight decay for BatchNorm layers.

Although our `DefaultOptimizerConstructor` is powerful, it may still not cover your need. If that, you can do those fine-grained parameter tuning through customizing optimizer constructor.

```
from mmcv.runner.optimizer import OPTIMIZER_BUILDERS

@OPTIMIZER_BUILDERS.register_module()
class MyOptimizerConstructor:

    def __init__(self, optimizer_cfg, paramwise_cfg=None):
        pass

    def __call__(self, model):
        ... # Construct your optimizer here.
        return my_optimizer
```

The default optimizer constructor is implemented [here](#), which could also serve as a template for new optimizer constructor.

TUTORIAL 7: CUSTOMIZE RUNTIME SETTINGS

In this tutorial, we will introduce some methods about how to customize workflow and hooks when running your own settings for the project.

- *Customize Workflow*
- *Hooks*
 - *Default training hooks*
 - * *CheckpointHook*
 - * *LoggerHooks*
 - * *EvalHook*
 - *Use other implemented hooks*
 - *Customize self-implemented hooks*
 - * *1. Implement a new hook*
 - * *2. Register the new hook*
 - * *3. Modify the config*
- *FAQ*

9.1 Customize Workflow

Workflow is a list of (phase, duration) to specify the running order and duration. The meaning of “duration” depends on the runner’s type.

For example, we use epoch-based runner by default, and the “duration” means how many epochs the phase to be executed in a cycle. Usually, we only want to execute training phase, just use the following config.

```
workflow = [('train', 1)]
```

Sometimes we may want to check some metrics (e.g. loss, accuracy) about the model on the validate set. In such case, we can set the workflow as

```
[('train', 1), ('val', 1)]
```

so that 1 epoch for training and 1 epoch for validation will be run iteratively.

By default, we recommend using **EvalHook** to do evaluation after the training epoch, but you can still use `val` workflow as an alternative.

Note:

1. The parameters of model will not be updated during the val epoch.
 2. Keyword `max_epochs` in the config only controls the number of training epochs and will not affect the validation workflow.
 3. Workflows `[('train', 1), ('val', 1)]` and `[('train', 1)]` will not change the behavior of `EvalHook` because `EvalHook` is called by `after_train_epoch` and validation workflow only affect hooks that are called through `after_val_epoch`. Therefore, the only difference between `[('train', 1), ('val', 1)]` and `[('train', 1)]` is that the runner will calculate losses on the validation set after each training epoch.
-

9.2 Hooks

The hook mechanism is widely used in the OpenMMLab open-source algorithm library. Combined with the `Runner`, the entire life cycle of the training process can be managed easily. You can learn more about the hook through [related article](#).

Hooks only work after being registered into the runner. At present, hooks are mainly divided into two categories:

- default training hooks

The default training hooks are registered by the runner by default. Generally, they are hooks for some basic functions, and have a certain priority, you don't need to modify the priority.

- custom hooks

The custom hooks are registered through `custom_hooks`. Generally, they are hooks with enhanced functions. The priority needs to be specified in the configuration file. If you do not specify the priority of the hook, it will be set to 'NORMAL' by default.

Priority list

The priority determines the execution order of the hooks. Before training, the log will print out the execution order of the hooks at each stage to facilitate debugging.

9.2.1 default training hooks

Some common hooks are not registered through `custom_hooks`, they are

`OptimizerHook`, `MomentumUpdaterHook` and `LrUpdaterHook` have been introduced in *schedule strategy*. `IterTimerHook` is used to record elapsed time and does not support modification.

Here we reveal how to customize `CheckpointHook`, `LoggerHooks`, and `EvalHook`.

CheckpointHook

The MMCV runner will use `checkpoint_config` to initialize `CheckpointHook`.

```
checkpoint_config = dict(interval=1)
```

We could set `max_keep_ckpts` to save only a small number of checkpoints or decide whether to store state dict of optimizer by `save_optimizer`. More details of the arguments are [here](#)

LoggerHooks

The `log_config` wraps multiple logger hooks and enables to set intervals. Now MMCV supports `TextLoggerHook`, `WandbLoggerHook`, `MlflowLoggerHook`, `NeptuneLoggerHook`, `DvcLiveLoggerHook` and `TensorboardLoggerHook`. The detailed usages can be found in the [doc](#).

```
log_config = dict(
    interval=50,
    hooks=[
        dict(type='TextLoggerHook'),
        dict(type='TensorboardLoggerHook')
    ])

```

EvalHook

The config of evaluation will be used to initialize the `EvalHook`.

The `EvalHook` has some reserved keys, such as `interval`, `save_best` and `start`, and the other arguments such as `metrics` will be passed to the `dataset.evaluate()`

```
evaluation = dict(interval=1, metric='accuracy', metric_options={'topk': (1, )})
```

You can save the model weight when the best verification result is obtained by modifying the parameter `save_best`:

```
# "auto" means automatically select the metrics to compare.
# You can also use a specific key like "accuracy_top-1".
evaluation = dict(interval=1, save_best="auto", metric='accuracy', metric_options={'topk
↪': (1, )})
```

When running some large experiments, you can skip the validation step at the beginning of training by modifying the parameter `start` as below:

```
evaluation = dict(interval=1, start=200, metric='accuracy', metric_options={'topk': (1,
↪)})
```

This indicates that, before the 200th epoch, evaluations would not be executed. Since the 200th epoch, evaluations would be executed after the training process.

Note: In the default configuration files of MMClassification, the evaluation field is generally placed in the datasets configs.

9.2.2 Use other implemented hooks

Some hooks have been already implemented in MMCV and MMClassification, they are:

- EMAHook
- SyncBuffersHook
- EmptyCacheHook
- ProfilerHook
-

If the hook is already implemented in MMCV, you can directly modify the config to use the hook as below

```
mmcv_hooks = [  
    dict(type='MMCVHook', a=a_value, b=b_value, priority='NORMAL')  
]
```

such as using EMAHook, interval is 100 iters:

```
custom_hooks = [  
    dict(type='EMAHook', interval=100, priority='HIGH')  
]
```

9.3 Customize self-implemented hooks

9.3.1 1. Implement a new hook

Here we give an example of creating a new hook in MMClassification and using it in training.

```
from mmcv.runner import HOOKS, Hook  
  
@HOOKS.register_module()  
class MyHook(Hook):  
  
    def __init__(self, a, b):  
        pass  
  
    def before_run(self, runner):  
        pass  
  
    def after_run(self, runner):  
        pass  
  
    def before_epoch(self, runner):  
        pass  
  
    def after_epoch(self, runner):  
        pass  
  
    def before_iter(self, runner):
```

(continues on next page)

(continued from previous page)

```

pass

def after_iter(self, runner):
    pass

```

Depending on the functionality of the hook, the users need to specify what the hook will do at each stage of the training in `before_run`, `after_run`, `before_epoch`, `after_epoch`, `before_iter`, and `after_iter`.

9.3.2 2. Register the new hook

Then we need to make `MyHook` imported. Assuming the file is in `mmcls/core/utils/my_hook.py` there are two ways to do that:

- Modify `mmcls/core/utils/__init__.py` to import it.

The newly defined module should be imported in `mmcls/core/utils/__init__.py` so that the registry will find the new module and add it:

```
from .my_hook import MyHook
```

- Use `custom_imports` in the config to manually import it

```
custom_imports = dict(imports=['mmcls.core.utils.my_hook'], allow_failed_imports=False)
```

9.3.3 3. Modify the config

```

custom_hooks = [
    dict(type='MyHook', a=a_value, b=b_value)
]

```

You can also set the priority of the hook as below:

```

custom_hooks = [
    dict(type='MyHook', a=a_value, b=b_value, priority='ABOVE_NORMAL')
]

```

By default, the hook's priority is set as `NORMAL` during registration.

9.4 FAQ

9.4.1 1. `resume_from` and `load_from` and `init_cfg.Pretrained`

- `load_from`: only imports model weights, which is mainly used to load pre-trained or trained models;
- `resume_from`: not only import model weights, but also optimizer information, current epoch information, mainly used to continue training from the checkpoint.
- `init_cfg.Pretrained`: Load weights during weight initialization, and you can specify which module to load. This is usually used when fine-tuning a model, refer to *Tutorial 2: Fine-tune Models*.

MODEL ZOO SUMMARY

- Number of papers: 16
 - ALGORITHM: 16
- Number of checkpoints: 80
 - [ALGORITHM] *MLP-Mixer: An all-MLP Architecture for Vision* (2 ckpts)
 - [ALGORITHM] *MobileNetV2: Inverted Residuals and Linear Bottlenecks* (1 ckpts)
 - [ALGORITHM] *Searching for MobileNetV3* (2 ckpts)
 - [ALGORITHM] *Designing Network Design Spaces* (8 ckpts)
 - [ALGORITHM] *Repvgg: Making vgg-style convnets great again* (12 ckpts)
 - [ALGORITHM] *Res2Net: A New Multi-scale Backbone Architecture* (3 ckpts)
 - [ALGORITHM] *Deep Residual Learning for Image Recognition* (15 ckpts)
 - [ALGORITHM] *Aggregated Residual Transformations for Deep Neural Networks* (4 ckpts)
 - [ALGORITHM] *Squeeze-and-Excitation Networks* (2 ckpts)
 - [ALGORITHM] *ShuffleNet: An Extremely Efficient Convolutional Neural Network for Mobile Devices* (1 ckpts)
 - [ALGORITHM] *Shufflenet v2: Practical guidelines for efficient cnn architecture design* (1 ckpts)
 - [ALGORITHM] *Swin Transformer: Hierarchical Vision Transformer using Shifted Windows* (11 ckpts)
 - [ALGORITHM] *Tokens-to-Token ViT: Training Vision Transformers from Scratch on ImageNet* (3 ckpts)
 - [ALGORITHM] *Transformer in Transformer* (1 ckpts)
 - [ALGORITHM] *Very Deep Convolutional Networks for Large-Scale Image Recognition* (8 ckpts)
 - [ALGORITHM] *An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale* (6 ckpts)

11.1 ImageNet

ImageNet has multiple versions, but the most commonly used one is [ILSVRC 2012](#). The ResNet family models below are trained by standard data augmentations, i.e., RandomResizedCrop, RandomHorizontalFlip and Normalize.

Models with * are converted from other repos, others are trained by ourselves.

11.2 CIFAR10

MLP-MIXER: AN ALL-MLP ARCHITECTURE FOR VISION

12.1 Abstract

Convolutional Neural Networks (CNNs) are the go-to model for computer vision. Recently, attention-based networks, such as the Vision Transformer, have also become popular. In this paper we show that while convolutions and attention are both sufficient for good performance, neither of them are necessary. We present MLP-Mixer, an architecture based exclusively on multi-layer perceptrons (MLPs). MLP-Mixer contains two types of layers: one with MLPs applied independently to image patches (i.e. “mixing” the per-location features), and one with MLPs applied across patches (i.e. “mixing” spatial information). When trained on large datasets, or with modern regularization schemes, MLP-Mixer attains competitive scores on image classification benchmarks, with pre-training and inference cost comparable to state-of-the-art models. We hope that these results spark further research beyond the realms of well established CNNs and Transformers.

12.2 Citation

```
@misc{tolstikhin2021mlpmixer,  
  title={MLP-Mixer: An all-MLP Architecture for Vision},  
  author={Ilya Tolstikhin and Neil Houlsby and Alexander Kolesnikov and Lucas Beyer_  
↪and Xiaohua Zhai and Thomas Unterthiner and Jessica Yung and Andreas Steiner and_  
↪Daniel Keysers and Jakob Uszkoreit and Mario Lucic and Alexey Dosovitskiy},  
  year={2021},  
  eprint={2105.01601},  
  archivePrefix={arXiv},  
  primaryClass={cs.CV}  
}
```

12.3 Pretrain model

The pre-trained models are converted from `timm`.

12.3.1 ImageNet-1k

*Models with * are converted from other repos.*

MOBILENETV2: INVERTED RESIDUALS AND LINEAR BOTTLENECKS

13.1 Abstract

In this paper we describe a new mobile architecture, MobileNetV2, that improves the state of the art performance of mobile models on multiple tasks and benchmarks as well as across a spectrum of different model sizes. We also describe efficient ways of applying these mobile models to object detection in a novel framework we call SSDLite. Additionally, we demonstrate how to build mobile semantic segmentation models through a reduced form of DeepLabv3 which we call Mobile DeepLabv3.

The MobileNetV2 architecture is based on an inverted residual structure where the input and output of the residual block are thin bottleneck layers opposite to traditional residual models which use expanded representations in the input. MobileNetV2 uses lightweight depthwise convolutions to filter features in the intermediate expansion layer. Additionally, we find that it is important to remove non-linearities in the narrow layers in order to maintain representational power. We demonstrate that this improves performance and provide an intuition that led to this design. Finally, our approach allows decoupling of the input/output domains from the expressiveness of the transformation, which provides a convenient framework for further analysis. We measure our performance on Imagenet classification, COCO object detection, VOC image segmentation. We evaluate the trade-offs between accuracy, and number of operations measured by multiply-adds (MAdd), as well as the number of parameters

13.2 Citation

```
@INPROCEEDINGS{8578572,
  author={M. {Sandler} and A. {Howard} and M. {Zhu} and A. {Zhmoginov} and L. {Chen}},
  booktitle={2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition},
  title={MobileNetV2: Inverted Residuals and Linear Bottlenecks},
  year={2018},
  volume={},
  number={},
  pages={4510-4520},
  doi={10.1109/CVPR.2018.00474}}
}
```

13.3 Results and models

13.3.1 ImageNet

SEARCHING FOR MOBILENETV3

14.1 Abstract

We present the next generation of MobileNets based on a combination of complementary search techniques as well as a novel architecture design. MobileNetV3 is tuned to mobile phone CPUs through a combination of hardware-aware network architecture search (NAS) complemented by the NetAdapt algorithm and then subsequently improved through novel architecture advances. This paper starts the exploration of how automated search algorithms and network design can work together to harness complementary approaches improving the overall state of the art. Through this process we create two new MobileNet models for release: MobileNetV3-Large and MobileNetV3-Small which are targeted for high and low resource use cases. These models are then adapted and applied to the tasks of object detection and semantic segmentation. For the task of semantic segmentation (or any dense pixel prediction), we propose a new efficient segmentation decoder Lite Reduced Atrous Spatial Pyramid Pooling (LR-ASPP). We achieve new state of the art results for mobile classification, detection and segmentation. MobileNetV3-Large is 3.2% more accurate on ImageNet classification while reducing latency by 15% compared to MobileNetV2. MobileNetV3-Small is 4.6% more accurate while reducing latency by 5% compared to MobileNetV2. MobileNetV3-Large detection is 25% faster at roughly the same accuracy as MobileNetV2 on COCO detection. MobileNetV3-Large LR-ASPP is 30% faster than MobileNetV2 R-ASPP at similar accuracy for Cityscapes segmentation.

14.2 Citation

```
@inproceedings{Howard_2019_ICCV,  
  author = {Howard, Andrew and Sandler, Mark and Chu, Grace and Chen, Liang-Chieh and  
↪Chen, Bo and Tan, Mingxing and Wang, Weijun and Zhu, Yukun and Pang, Ruoming and  
↪Vasudevan, Vijay and Le, Quoc V. and Adam, Hartwig},  
  title = {Searching for MobileNetV3},  
  booktitle = {Proceedings of the IEEE/CVF International Conference on Computer Vision  
↪(ICCV)},  
  month = {October},  
  year = {2019}  
}
```

14.3 Pretrain model

The pre-trained modles are converted from `torchvision`.

14.3.1 ImageNet

14.4 Results and models

Waiting for adding.

DESIGNING NETWORK DESIGN SPACES

15.1 Abstract

In this work, we present a new network design paradigm. Our goal is to help advance the understanding of network design and discover design principles that generalize across settings. Instead of focusing on designing individual network instances, we design network design spaces that parametrize populations of networks. The overall process is analogous to classic manual design of networks, but elevated to the design space level. Using our methodology we explore the structure aspect of network design and arrive at a low-dimensional design space consisting of simple, regular networks that we call RegNet. The core insight of the RegNet parametrization is surprisingly simple: widths and depths of good networks can be explained by a quantized linear function. We analyze the RegNet design space and arrive at interesting findings that do not match the current practice of network design. The RegNet design space provides simple and fast networks that work well across a wide range of flop regimes. Under comparable training settings and flops, the RegNet models outperform the popular EfficientNet models while being up to 5x faster on GPUs.

15.2 Citation

```
@article{radosavovic2020designing,  
  title={Designing Network Design Spaces},  
  author={Ilija Radosavovic and Raj Prateek Kosaraju and Ross Girshick and Kaiming He_  
↪and Piotr Dollár},  
  year={2020},  
  eprint={2003.13678},  
  archivePrefix={arXiv},  
  primaryClass={cs.CV}  
}
```

15.3 Pretrain model

The pre-trained modles are converted from [model zoo](#) of [pycls](#).

15.3.1 ImageNet

15.4 Results and models

Waiting for adding.

REPVGG: MAKING VGG-STYLE CONVNETS GREAT AGAIN

16.1 Abstract

We present a simple but powerful architecture of convolutional neural network, which has a VGG-like inference-time body composed of nothing but a stack of 3x3 convolution and ReLU, while the training-time model has a multi-branch topology. Such decoupling of the training-time and inference-time architecture is realized by a structural re-parameterization technique so that the model is named RepVGG. On ImageNet, RepVGG reaches over 80% top-1 accuracy, which is the first time for a plain model, to the best of our knowledge. On NVIDIA 1080Ti GPU, RepVGG models run 83% faster than ResNet-50 or 101% faster than ResNet-101 with higher accuracy and show favorable accuracy-speed trade-off compared to the state-of-the-art models like EfficientNet and RegNet.

16.2 Citation

```
@inproceedings{ding2021repvgg,  
  title={Repvgg: Making vgg-style convnets great again},  
  author={Ding, Xiaohan and Zhang, Xiangyu and Ma, Ningning and Han, Jungong and Ding, Guiguang and Sun, Jian},  
  booktitle={Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition},  
  pages={13733--13742},  
  year={2021}  
}
```

16.3 Pretrain model

*Models with * are converted from other repos.*

16.4 Reparameterize RepVGG

The checkpoints provided are all in train form. Use the reparameterize tool to switch them to more efficient deploy form, which not only has fewer parameters but also less calculations.

```
python ./tools/convert_models/reparameterize_repvgg.py ${CFG_PATH} ${SRC_CKPT_PATH} $
↪ ${TARGET_CKPT_PATH}
```

`${CFG_PATH}` is the config file, `${SRC_CKPT_PATH}` is the source checkpoint file, `${TARGET_CKPT_PATH}` is the target deploy weight file path.

To use reparameterized repvgg weight, the config file must switch to the deploy config files as below:

```
python ./tools/test.py ${RapVGG_Deploy_CFG} ${CHECK_POINT}
```

RES2NET: A NEW MULTI-SCALE BACKBONE ARCHITECTURE

17.1 Abstract

Representing features at multiple scales is of great importance for numerous vision tasks. Recent advances in backbone convolutional neural networks (CNNs) continually demonstrate stronger multi-scale representation ability, leading to consistent performance gains on a wide range of applications. However, most existing methods represent the multi-scale features in a layer-wise manner. In this paper, we propose a novel building block for CNNs, namely Res2Net, by constructing hierarchical residual-like connections within one single residual block. The Res2Net represents multi-scale features at a granular level and increases the range of receptive fields for each network layer. The proposed Res2Net block can be plugged into the state-of-the-art backbone CNN models, e.g., ResNet, ResNeXt, and DLA. We evaluate the Res2Net block on all these models and demonstrate consistent performance gains over baseline models on widely-used datasets, e.g., CIFAR-100 and ImageNet. Further ablation studies and experimental results on representative computer vision tasks, i.e., object detection, class activation mapping, and salient object detection, further verify the superiority of the Res2Net over the state-of-the-art baseline methods.

17.2 Citation

```
@article{gao2019res2net,  
  title={Res2Net: A New Multi-scale Backbone Architecture},  
  author={Gao, Shang-Hua and Cheng, Ming-Ming and Zhao, Kai and Zhang, Xin-Yu and Yang,   
↪Ming-Hsuan and Torr, Philip},  
  journal={IEEE TPAMI},  
  year={2021},  
  doi={10.1109/TPAMI.2019.2938758},  
}
```

17.3 Pretrain model

The pre-trained models are converted from [official repo](#).

17.3.1 ImageNet 1k

*Models with * are converted from other repos.*

DEEP RESIDUAL LEARNING FOR IMAGE RECOGNITION

18.1 Abstract

Deeper neural networks are more difficult to train. We present a residual learning framework to ease the training of networks that are substantially deeper than those used previously. We explicitly reformulate the layers as learning residual functions with reference to the layer inputs, instead of learning unreferenced functions. We provide comprehensive empirical evidence showing that these residual networks are easier to optimize, and can gain accuracy from considerably increased depth. On the ImageNet dataset we evaluate residual nets with a depth of up to 152 layers—8x deeper than VGG nets but still having lower complexity. An ensemble of these residual nets achieves 3.57% error on the ImageNet test set. This result won the 1st place on the ILSVRC 2015 classification task. We also present analysis on CIFAR-10 with 100 and 1000 layers.

The depth of representations is of central importance for many visual recognition tasks. Solely due to our extremely deep representations, we obtain a 28% relative improvement on the COCO object detection dataset. Deep residual nets are foundations of our submissions to ILSVRC & COCO 2015 competitions, where we also won the 1st places on the tasks of ImageNet detection, ImageNet localization, COCO detection, and COCO segmentation.

18.2 Citation

```
@inproceedings{he2016deep,  
  title={Deep residual learning for image recognition},  
  author={He, Kaiming and Zhang, Xiangyu and Ren, Shaoqing and Sun, Jian},  
  booktitle={Proceedings of the IEEE conference on computer vision and pattern  
↪recognition},  
  pages={770--778},  
  year={2016}  
}
```

18.3 Results and models

18.4 Cifar10

18.5 Cifar100

18.5.1 ImageNet

AGGREGATED RESIDUAL TRANSFORMATIONS FOR DEEP NEURAL NETWORKS

19.1 Abstract

We present a simple, highly modularized network architecture for image classification. Our network is constructed by repeating a building block that aggregates a set of transformations with the same topology. Our simple design results in a homogeneous, multi-branch architecture that has only a few hyper-parameters to set. This strategy exposes a new dimension, which we call “cardinality” (the size of the set of transformations), as an essential factor in addition to the dimensions of depth and width. On the ImageNet-1K dataset, we empirically show that even under the restricted condition of maintaining complexity, increasing cardinality is able to improve classification accuracy. Moreover, increasing cardinality is more effective than going deeper or wider when we increase the capacity. Our models, named ResNeXt, are the foundations of our entry to the ILSVRC 2016 classification task in which we secured 2nd place. We further investigate ResNeXt on an ImageNet-5K set and the COCO detection set, also showing better results than its ResNet counterpart. The code and models are publicly available online.

19.2 Citation

```
@inproceedings{xie2017aggregated,  
  title={Aggregated residual transformations for deep neural networks},  
  author={Xie, Saining and Girshick, Ross and Doll{\`a}r, Piotr and Tu, Zhuowen and He,  
↪Kaiming},  
  booktitle={Proceedings of the IEEE conference on computer vision and pattern  
↪recognition},  
  pages={1492--1500},  
  year={2017}  
}
```

19.3 Results and models

19.3.1 ImageNet

SQUEEZE-AND-EXCITATION NETWORKS

20.1 Abstract

The central building block of convolutional neural networks (CNNs) is the convolution operator, which enables networks to construct informative features by fusing both spatial and channel-wise information within local receptive fields at each layer. A broad range of prior research has investigated the spatial component of this relationship, seeking to strengthen the representational power of a CNN by enhancing the quality of spatial encodings throughout its feature hierarchy. In this work, we focus instead on the channel relationship and propose a novel architectural unit, which we term the “Squeeze-and-Excitation” (SE) block, that adaptively recalibrates channel-wise feature responses by explicitly modelling interdependencies between channels. We show that these blocks can be stacked together to form SENet architectures that generalise extremely effectively across different datasets. We further demonstrate that SE blocks bring significant improvements in performance for existing state-of-the-art CNNs at slight additional computational cost. Squeeze-and-Excitation Networks formed the foundation of our ILSVRC 2017 classification submission which won first place and reduced the top-5 error to 2.251%, surpassing the winning entry of 2016 by a relative improvement of ~25%.

20.2 Citation

```
@inproceedings{hu2018squeeze,  
  title={Squeeze-and-excitation networks},  
  author={Hu, Jie and Shen, Li and Sun, Gang},  
  booktitle={Proceedings of the IEEE conference on computer vision and pattern  
↪recognition},  
  pages={7132--7141},  
  year={2018}  
}
```

20.3 Results and models

20.3.1 ImageNet

SHUFFLENET: AN EXTREMELY EFFICIENT CONVOLUTIONAL NEURAL NETWORK FOR MOBILE DEVICES

21.1 Abstract

We introduce an extremely computation-efficient CNN architecture named ShuffleNet, which is designed specially for mobile devices with very limited computing power (e.g., 10-150 MFLOPs). The new architecture utilizes two new operations, pointwise group convolution and channel shuffle, to greatly reduce computation cost while maintaining accuracy. Experiments on ImageNet classification and MS COCO object detection demonstrate the superior performance of ShuffleNet over other structures, e.g. lower top-1 error (absolute 7.8%) than recent MobileNet on ImageNet classification task, under the computation budget of 40 MFLOPs. On an ARM-based mobile device, ShuffleNet achieves ~13x actual speedup over AlexNet while maintaining comparable accuracy.

21.2 Citation

```
@inproceedings{zhang2018shufflenet,  
  title={Shufflenet: An extremely efficient convolutional neural network for mobile  
↪ devices},  
  author={Zhang, Xiangyu and Zhou, Xinyu and Lin, Mengxiao and Sun, Jian},  
  booktitle={Proceedings of the IEEE conference on computer vision and pattern  
↪ recognition},  
  pages={6848--6856},  
  year={2018}  
}
```

21.3 Results and models

21.3.1 ImageNet

SHUFFLENET V2: PRACTICAL GUIDELINES FOR EFFICIENT CNN ARCHITECTURE DESIGN

22.1 Abstract

Currently, the neural network architecture design is mostly guided by the *indirect* metric of computation complexity, i.e., FLOPs. However, the *direct* metric, e.g., speed, also depends on the other factors such as memory access cost and platform characteristics. Thus, this work proposes to evaluate the direct metric on the target platform, beyond only considering FLOPs. Based on a series of controlled experiments, this work derives several practical *guidelines* for efficient network design. Accordingly, a new architecture is presented, called *ShuffleNet V2*. Comprehensive ablation experiments verify that our model is the state-of-the-art in terms of speed and accuracy tradeoff.

22.2 Citation

```
@inproceedings{ma2018shufflenet,  
  title={Shufflenet v2: Practical guidelines for efficient cnn architecture design},  
  author={Ma, Ningning and Zhang, Xiangyu and Zheng, Hai-Tao and Sun, Jian},  
  booktitle={Proceedings of the European conference on computer vision (ECCV)},  
  pages={116--131},  
  year={2018}  
}
```

22.3 Results and models

22.3.1 ImageNet

SWIN TRANSFORMER: HIERARCHICAL VISION TRANSFORMER USING SHIFTED WINDOWS

23.1 Abstract

This paper presents a new vision Transformer, called Swin Transformer, that capably serves as a general-purpose backbone for computer vision. Challenges in adapting Transformer from language to vision arise from differences between the two domains, such as large variations in the scale of visual entities and the high resolution of pixels in images compared to words in text. To address these differences, we propose a hierarchical Transformer whose representation is computed with **Shifted windows**. The shifted windowing scheme brings greater efficiency by limiting self-attention computation to non-overlapping local windows while also allowing for cross-window connection. This hierarchical architecture has the flexibility to model at various scales and has linear computational complexity with respect to image size. These qualities of Swin Transformer make it compatible with a broad range of vision tasks, including image classification (87.3 top-1 accuracy on ImageNet-1K) and dense prediction tasks such as object detection (58.7 box AP and 51.1 mask AP on COCO test-dev) and semantic segmentation (53.5 mIoU on ADE20K val). Its performance surpasses the previous state-of-the-art by a large margin of +2.7 box AP and +2.6 mask AP on COCO, and +3.2 mIoU on ADE20K, demonstrating the potential of Transformer-based models as vision backbones. The hierarchical design and the shifted window approach also prove beneficial for all-MLP architectures.

23.2 Citation

```
@article{liu2021Swin,  
  title={Swin Transformer: Hierarchical Vision Transformer using Shifted Windows},  
  author={Liu, Ze and Lin, Yutong and Cao, Yue and Hu, Han and Wei, Yixuan and Zhang,  
↪Zheng and Lin, Stephen and Guo, Baining},  
  journal={arXiv preprint arXiv:2103.14030},  
  year={2021}  
}
```

23.3 Pretrain model

The pre-trained models are converted from model zoo of Swin Transformer.

23.3.1 ImageNet 1k

23.4 Results and models

23.4.1 ImageNet 1k

Model	Pretrain	resolution	Params(M)	Flops(G)	Top-1 (%)	Top-5 (%)	Config	Download
Swin-T	ImageNet-1k	224x224	28.29	4.36	81.18	95.61	config	model log
Swin-S	ImageNet-1k	224x224	49.61	8.52	83.02	96.29	config	model log
Swin-B	ImageNet-1k	224x224	87.77	15.14	83.36	96.44	config	model log

TOKENS-TO-TOKEN ViT: TRAINING VISION TRANSFORMERS FROM SCRATCH ON IMAGENET

24.1 Abstract

Transformers, which are popular for language modeling, have been explored for solving vision tasks recently, \eg, the Vision Transformer (ViT) for image classification. The ViT model splits each image into a sequence of tokens with fixed length and then applies multiple Transformer layers to model their global relation for classification. However, ViT achieves inferior performance to CNNs when trained from scratch on a midsize dataset like ImageNet. We find it is because: 1) the simple tokenization of input images fails to model the important local structure such as edges and lines among neighboring pixels, leading to low training sample efficiency; 2) the redundant attention backbone design of ViT leads to limited feature richness for fixed computation budgets and limited training samples. To overcome such limitations, we propose a new Tokens-to-Token Vision Transformer (T2T-ViT), which incorporates 1) a layer-wise Tokens-to-Token (T2T) transformation to progressively structurize the image to tokens by recursively aggregating neighboring Tokens into one Token (Tokens-to-Token), such that local structure represented by surrounding tokens can be modeled and tokens length can be reduced; 2) an efficient backbone with a deep-narrow structure for vision transformer motivated by CNN architecture design after empirical study. Notably, T2T-ViT reduces the parameter count and MACs of vanilla ViT by half, while achieving more than 3.0% improvement when trained from scratch on ImageNet. It also outperforms ResNets and achieves comparable performance with MobileNets by directly training on ImageNet. For example, T2T-ViT with comparable size to ResNet50 (21.5M parameters) can achieve 83.3% top1 accuracy in image resolution 384×384 on ImageNet.

24.2 Citation

```
@article{yuan2021tokens,  
  title={Tokens-to-token vit: Training vision transformers from scratch on imagenet},  
  author={Yuan, Li and Chen, Yunpeng and Wang, Tao and Yu, Weihao and Shi, Yujun and Tay,  
↪ Francis EH and Feng, Jiashi and Yan, Shuicheng},  
  journal={arXiv preprint arXiv:2101.11986},  
  year={2021}  
}
```

24.3 Pretrain model

The pre-trained modles are converted from [official repo](#).

24.3.1 ImageNet-1k

*Models with * are converted from other repos.*

24.4 Results and models

Waiting for adding.

TRANSFORMER IN TRANSFORMER

25.1 Abstract

Transformer is a new kind of neural architecture which encodes the input data as powerful features via the attention mechanism. Basically, the visual transformers first divide the input images into several local patches and then calculate both representations and their relationship. Since natural images are of high complexity with abundant detail and color information, the granularity of the patch dividing is not fine enough for excavating features of objects in different scales and locations. In this paper, we point out that the attention inside these local patches are also essential for building visual transformers with high performance and we explore a new architecture, namely, Transformer in Transformer (TNT). Specifically, we regard the local patches (e.g., 16×16) as “visual sentences” and present to further divide them into smaller patches (e.g., 4×4) as “visual words”. The attention of each word will be calculated with other words in the given visual sentence with negligible computational costs. Features of both words and sentences will be aggregated to enhance the representation ability. Experiments on several benchmarks demonstrate the effectiveness of the proposed TNT architecture, e.g., we achieve an 81.5% top-1 accuracy on the ImageNet, which is about 1.7% higher than that of the state-of-the-art visual transformer with similar computational cost.

25.2 Citation

```
@misc{han2021transformer,
  title={Transformer in Transformer},
  author={Kai Han and An Xiao and Enhua Wu and Jianyuan Guo and Chunjing Xu and
↪ Yunhe Wang},
  year={2021},
  eprint={2103.00112},
  archivePrefix={arXiv},
  primaryClass={cs.CV}
}
```

25.3 Pretrain model

The pre-trained models are converted from `timm`.

25.3.1 ImageNet

*Models with * are converted from other repos.*

25.4 Results and models

Waiting for adding.

VERY DEEP CONVOLUTIONAL NETWORKS FOR LARGE-SCALE IMAGE RECOGNITION

26.1 Abstract

In this work we investigate the effect of the convolutional network depth on its accuracy in the large-scale image recognition setting. Our main contribution is a thorough evaluation of networks of increasing depth using an architecture with very small (3x3) convolution filters, which shows that a significant improvement on the prior-art configurations can be achieved by pushing the depth to 16-19 weight layers. These findings were the basis of our ImageNet Challenge 2014 submission, where our team secured the first and the second places in the localisation and classification tracks respectively. We also show that our representations generalise well to other datasets, where they achieve state-of-the-art results. We have made our two best-performing ConvNet models publicly available to facilitate further research on the use of deep visual representations in computer vision.

26.2 Citation

```
@article{simonyan2014very,  
  title={Very deep convolutional networks for large-scale image recognition},  
  author={Simonyan, Karen and Zisserman, Andrew},  
  journal={arXiv preprint arXiv:1409.1556},  
  year={2014}  
}
```

26.3 Results and models

26.3.1 ImageNet

AN IMAGE IS WORTH 16X16 WORDS: TRANSFORMERS FOR IMAGE RECOGNITION AT SCALE

27.1 Abstract

While the Transformer architecture has become the de-facto standard for natural language processing tasks, its applications to computer vision remain limited. In vision, attention is either applied in conjunction with convolutional networks, or used to replace certain components of convolutional networks while keeping their overall structure in place. We show that this reliance on CNNs is not necessary and a pure transformer applied directly to sequences of image patches can perform very well on image classification tasks. When pre-trained on large amounts of data and transferred to multiple mid-sized or small image recognition benchmarks (ImageNet, CIFAR-100, VTAB, etc.), Vision Transformer (ViT) attains excellent results compared to state-of-the-art convolutional networks while requiring substantially fewer computational resources to train.

27.2 Citation

```
@inproceedings{
  dosovitskiy2021an,
  title={An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale},
  author={Alexey Dosovitskiy and Lucas Beyer and Alexander Kolesnikov and Dirk
↪Weissenborn and Xiaohua Zhai and Thomas Unterthiner and Mostafa Dehghani and Matthias
↪Minderer and Georg Heigold and Sylvain Gelly and Jakob Uszkoreit and Neil Houlsby},
  booktitle={International Conference on Learning Representations},
  year={2021},
  url={https://openreview.net/forum?id=YicbFdNTTy}
}
```

The training step of Vision Transformers is divided into two steps. The first step is training the model on a large dataset, like ImageNet-21k, and get the pretrain model. And the second step is training the model on the target dataset, like ImageNet-1k, and get the finetune model. Here, we provide both pretrain models and finetune models.

27.3 Pretrain model

The pre-trained models are converted from [model zoo of Google Research](#).

27.3.1 ImageNet 21k

*Models with * are converted from other repos.*

27.4 Finetune model

The finetune models are converted from [model zoo of Google Research](#).

27.4.1 ImageNet 1k

Model	Pretrain	resolu- tion	Params(M)	Flops(G)	Top-1 (%)	Top-5 (%)	Con- fig	Down- load
ViT- B16*	ImageNet- 21k	384x384	86.86	33.03	85.43	97.77	config	model
ViT- B32*	ImageNet- 21k	384x384	88.30	8.56	84.01	97.08	config	model
ViT- L16*	ImageNet- 21k	384x384	304.72	116.68	85.63	97.63	config	model

*Models with * are converted from other repos.*

PYTORCH TO ONNX (EXPERIMENTAL)

- *Pytorch to ONNX (Experimental)*
 - *How to convert models from Pytorch to ONNX*
 - * *Prerequisite*
 - * *Usage*
 - * *Description of all arguments:*
 - *How to evaluate ONNX models with ONNX Runtime*
 - * *Prerequisite*
 - * *Usage*
 - * *Description of all arguments*
 - * *Results and Models*
 - *List of supported models exportable to ONNX*
 - *Reminders*
 - *FAQs*

28.1 How to convert models from Pytorch to ONNX

28.1.1 Prerequisite

1. Please refer to [install](#) for installation of MMClassification.
2. Install onnx and onnxruntime

```
pip install onnx onnxruntime==1.5.1
```

28.1.2 Usage

```
python tools/deployment/pytorch2onnx.py \  
    ${CONFIG_FILE} \  
    --checkpoint ${CHECKPOINT_FILE} \  
    --output-file ${OUTPUT_FILE} \  
    --shape ${IMAGE_SHAPE} \  
    --opset-version ${OPSET_VERSION} \  
    --dynamic-export \  
    --show \  
    --simplify \  
    --verify
```

28.1.3 Description of all arguments:

- `config` : The path of a model config file.
- `--checkpoint` : The path of a model checkpoint file.
- `--output-file`: The path of output ONNX model. If not specified, it will be set to `tmp.onnx`.
- `--shape`: The height and width of input tensor to the model. If not specified, it will be set to 224 224.
- `--opset-version` : The opset version of ONNX. If not specified, it will be set to 11.
- `--dynamic-export` : Determines whether to export ONNX with dynamic input shape and output shapes. If not specified, it will be set to `False`.
- `--show`: Determines whether to print the architecture of the exported model. If not specified, it will be set to `False`.
- `--simplify`: Determines whether to simplify the exported ONNX model. If not specified, it will be set to `False`.
- `--verify`: Determines whether to verify the correctness of an exported model. If not specified, it will be set to `False`.

Example:

```
python tools/deployment/pytorch2onnx.py \  
    configs/resnet/resnet18_8xb16_cifar10.py \  
    --checkpoint checkpoints/resnet/resnet18_8xb16_cifar10.pth \  
    --output-file checkpoints/resnet/resnet18_8xb16_cifar10.onnx \  
    --dynamic-export \  
    --show \  
    --simplify \  
    --verify
```

28.2 How to evaluate ONNX models with ONNX Runtime

We prepare a tool `tools/deployment/test.py` to evaluate ONNX models with ONNXRuntime or TensorRT.

28.2.1 Prerequisite

- Install `onnx` and `onnxruntime-gpu`

```
pip install onnx onnxruntime-gpu
```

28.2.2 Usage

```
python tools/deployment/test.py \
    ${CONFIG_FILE} \
    ${ONNX_FILE} \
    --backend ${BACKEND} \
    --out ${OUTPUT_FILE} \
    --metrics ${EVALUATION_METRICS} \
    --metric-options ${EVALUATION_OPTIONS} \
    --show
    --show-dir ${SHOW_DIRECTORY} \
    --cfg-options ${CFG_OPTIONS} \
```

28.2.3 Description of all arguments

- `config`: The path of a model config file.
- `model`: The path of a ONNX model file.
- `--backend`: Backend for input model to run and should be `onnxruntime` or `tensorrt`.
- `--out`: The path of output result file in pickle format.
- `--metrics`: Evaluation metrics, which depends on the dataset, e.g., “accuracy”, “precision”, “recall”, “f1_score”, “support” for single label dataset, and “mAP”, “CP”, “CR”, “CF1”, “OP”, “OR”, “OF1” for multi-label dataset.
- `--show`: Determines whether to show classifier outputs. If not specified, it will be set to `False`.
- `--show-dir`: Directory where painted images will be saved
- `--metrics-options`: Custom options for evaluation, the key-value pair in `xxx=yyy` format will be kwargs for `dataset.evaluate()` function
- `--cfg-options`: Override some settings in the used config file, the key-value pair in `xxx=yyy` format will be merged into config file.

28.2.4 Results and Models

This part selects ImageNet for onnxruntime verification. ImageNet has multiple versions, but the most commonly used one is [ILSVRC 2012](#).

28.3 List of supported models exportable to ONNX

The table below lists the models that are guaranteed to be exportable to ONNX and runnable in ONNX Runtime.

Notes:

- *All models above are tested with Pytorch==1.6.0*

28.4 Reminders

- If you meet any problem with the listed models above, please create an issue and it would be taken care of soon. For models not included in the list, please try to dig a little deeper and debug a little bit more and hopefully solve them by yourself.

28.5 FAQs

- None

ONNX TO TENSORRT (EXPERIMENTAL)

- *ONNX to TensorRT (Experimental)*
 - *How to convert models from ONNX to TensorRT*
 - * *Prerequisite*
 - * *Usage*
 - *List of supported models convertible to TensorRT*
 - *Reminders*
 - *FAQs*

29.1 How to convert models from ONNX to TensorRT

29.1.1 Prerequisite

1. Please refer to [install.md](#) for installation of MMClassification from source.
2. Use our tool [pytorch2onnx.md](#) to convert the model from PyTorch to ONNX.

29.1.2 Usage

```
python tools/deployment/onnx2tensorrt.py \  
    ${MODEL} \  
    --trt-file ${TRT_FILE} \  
    --shape ${IMAGE_SHAPE} \  
    --max-batch-size ${MAX_BATCH_SIZE} \  
    --workspace-size ${WORKSPACE_SIZE} \  
    --fp16 \  
    --show \  
    --verify \  
    \
```

Description of all arguments:

- `model` : The path of an ONNX model file.
- `--trt-file`: The Path of output TensorRT engine file. If not specified, it will be set to `tmp.trt`.
- `--shape`: The height and width of model input. If not specified, it will be set to 224 224.
- `--max-batch-size`: The max batch size of TensorRT model, should not be less than 1.

- `--fp16`: Enable fp16 mode.
- `--workspace-size` : The required GPU workspace size in GiB to build TensorRT engine. If not specified, it will be set to 1 GiB.
- `--show`: Determines whether to show the outputs of the model. If not specified, it will be set to `False`.
- `--verify`: Determines whether to verify the correctness of models between ONNXRuntime and TensorRT. If not specified, it will be set to `False`.

Example:

```
python tools/deployment/onnx2tensorrt.py \  
  checkpoints/resnet/resnet18_b16x8_cifar10.onnx \  
  --trt-file checkpoints/resnet/resnet18_b16x8_cifar10.trt \  
  --shape 224 224 \  
  --show \  
  --verify \  

```

29.2 List of supported models convertible to TensorRT

The table below lists the models that are guaranteed to be convertible to TensorRT.

Notes:

- *All models above are tested with Pytorch==1.6.0 and TensorRT-7.2.1.6.Ubuntu-16.04.x86_64-gnu.cuda-10.2.cudnn8.0*

29.3 Reminders

- If you meet any problem with the listed models above, please create an issue and it would be taken care of soon. For models not included in the list, we may not provide much help here due to the limited resources. Please try to dig a little deeper and debug by yourself.

29.4 FAQs

- None

PYTORCH TO TORCHSCRIPT (EXPERIMENTAL)

- *Pytorch to TorchScript (Experimental)*
 - *How to convert models from Pytorch to TorchScript*
 - * *Usage*
 - * *Description of all arguments*
 - *Reminders*
 - *FAQs*

30.1 How to convert models from Pytorch to TorchScript

30.1.1 Usage

```
python tools/deployment/pytorch2torchscript.py \  
    ${CONFIG_FILE} \  
    --checkpoint ${CHECKPOINT_FILE} \  
    --output-file ${OUTPUT_FILE} \  
    --shape ${IMAGE_SHAPE} \  
    --verify \  
    \
```

30.1.2 Description of all arguments

- `config` : The path of a model config file.
- `--checkpoint` : The path of a model checkpoint file.
- `--output-file`: The path of output TorchScript model. If not specified, it will be set to `tmp.pt`.
- `--shape`: The height and width of input tensor to the model. If not specified, it will be set to `224 224`.
- `--verify`: Determines whether to verify the correctness of an exported model. If not specified, it will be set to `False`.

Example:

```
python tools/deployment/pytorch2onnx.py \  
    configs/resnet/resnet18_8xb16_cifar10.py \  
    --checkpoint checkpoints/resnet/resnet18_8xb16_cifar10.pth \  
    \
```

(continues on next page)

(continued from previous page)

```
--output-file checkpoints/resnet/resnet18_8xb16_cifar10.pt \  
--verify \  

```

Notes:

- *All models are tested with Pytorch==1.8.1*

30.2 Reminders

- For `torch.jit.is_tracing()` is only supported after v1.6. For users with pytorch v1.3-v1.5, we suggest early returning tensors manually.
- If you meet any problem with the models in this repo, please create an issue and it would be taken care of soon.

30.3 FAQs

- None

MODEL SERVING

In order to serve an MMLClassification model with `TorchServe`, you can follow the steps:

31.1 1. Convert model from MMLClassification to TorchServe

```
python tools/deployment/mmls2torchserve.py ${CONFIG_FILE} ${CHECKPOINT_FILE} \  
--output-folder ${MODEL_STORE} \  
--model-name ${MODEL_NAME}
```

Note: `MODEL_STORE` needs to be an absolute path to a folder.

Example:

```
python tools/deployment/mmls2torchserve.py \  
  configs/resnet/resnet18_8xb32_in1k.py \  
  checkpoints/resnet18_8xb32_in1k_20210831-fbbb1da6.pth \  
  --output-folder ./checkpoints \  
  --model-name resnet18_in1k
```

31.2 2. Build `mmls-serve` docker image

```
docker build -t mmls-serve:latest docker/serve/
```

31.3 3. Run `mmls-serve`

Check the official docs for [running TorchServe with docker](#).

In order to run in GPU, you need to install `nvidia-docker`. You can omit the `--gpus` argument in order to run in GPU.

Example:

```
docker run --rm \  
--cpus 8 \  
--gpus device=0 \  
-p8080:8080 -p8081:8081 -p8082:8082 \  
/
```

(continues on next page)

(continued from previous page)

```
--mount type=bind,source=`realpath ./checkpoints`,target=/home/model-server/model-store \  
mmcls-serve:latest
```

Note: `realpath ./checkpoints` points to the absolute path of “./checkpoints”, and you can replace it with the absolute path where you store torchserve models.

Read the docs about the Inference (8080), Management (8081) and Metrics (8082) APIs

31.4 4. Test deployment

```
curl http://127.0.0.1:8080/predictions/${MODEL_NAME} -T demo/demo.JPEG
```

You should obtain a response similar to:

```
{  
  "pred_label": 58,  
  "pred_score": 0.38102269172668457,  
  "pred_class": "water snake"  
}
```

And you can use `test_torchserver.py` to compare result of TorchServe and PyTorch, and visualize them.

```
python tools/deployment/test_torchserver.py ${IMAGE_FILE} ${CONFIG_FILE} ${CHECKPOINT_  
↪FILE} ${MODEL_NAME}  
[--inference-addr ${INFERENCE_ADDR}] [--device ${DEVICE}]
```

Example:

```
python tools/deployment/test_torchserver.py \  
demo/demo.JPEG \  
configs/resnet/resnet18_8xb32_in1k.py \  
checkpoints/resnet18_8xb32_in1k_20210831-fbbb1da6.pth \  
resnet18_in1k
```

VISUALIZATION

- *Visualization*
 - *Pipeline Visualization*
 - *Learning Rate Schedule Visualization*
 - *FAQs*

32.1 Pipeline Visualization

```
python tools/visualizations/vis_pipeline.py \  
    ${CONFIG_FILE} \  
    --output-dir ${OUTPUT_DIR} \  
    --phase ${DATASET_PHASE} \  
    --number ${BUNBER_IMAGES_DISPLAY} \  
    --skip-type ${SKIP_TRANSFORM_TYPE} \  
    --mode ${DISPLAY_MODE} \  
    --show \  
    --adaptive \  
    --min-edge-length ${MIN_EDGE_LENGTH} \  
    --max-edge-length ${MAX_EDGE_LENGTH} \  
    --bgr2rgb \  
    --window-size ${WINDOW_SIZE}
```

Description of all arguments:

- `config` : The path of a model config file.
- `--output-dir`: The output path for visualized images. If not specified, it will be set to '', which means not to save.
- `--phase`: Phase of visualizing dataset, must be one of [train, val, test]. If not specified, it will be set to train.
- `--number`: The number of samples to visualize. If not specified, display all images in the dataset.
- `--skip-type`: The pipelines to be skipped. If not specified, it will be set to ['ToTensor', 'Normalize', 'ImageToTensor', 'Collect'].
- `--mode`: The display mode, can be one of [original, pipeline, concat]. If not specified, it will be set to concat.
- `--show`: If set, display pictures in pop-up windows.
- `--adaptive`: If set, automatically adjust the size of the visualization images.

- `--min-edge-length`: The minimum edge length, used when `--adaptive` is set. When any side of the picture is smaller than `${MIN_EDGE_LENGTH}`, the picture will be enlarged while keeping the aspect ratio unchanged, and the short side will be aligned to `${MIN_EDGE_LENGTH}`. If not specified, it will be set to 200.
- `--max-edge-length`: The maximum edge length, used when `--adaptive` is set. When any side of the picture is larger than `${MAX_EDGE_LENGTH}`, the picture will be reduced while keeping the aspect ratio unchanged, and the long side will be aligned to `${MAX_EDGE_LENGTH}`. If not specified, it will be set to 1000.
- `--bgr2rgb`: If set, flip the color channel order of images.
- `--window-size`: The shape of the display window. If not specified, it will be set to 12*7. If used, it must be in the format 'W*H'.

Note:

1. If the `--mode` is not specified, it will be set to `concat` as default, get the pictures stitched together by original pictures and transformed pictures; if the `--mode` is set to `original`, get the original pictures; if the `--mode` is set to `pipeline`, get the transformed pictures.
2. When `--adaptive` option is set, images that are too large or too small will be automatically adjusted, you can use `--min-edge-length` and `--max-edge-length` to set the adjust size.

Examples:

1. Visualize all the transformed pictures of the ImageNet training set and display them in pop-up windows:

```
python ./tools/visualizations/vis_pipeline.py ./configs/resnet/resnet50_8xb32_in1k.py --
↳ show --mode pipeline
```

1. Visualize 10 comparison pictures in the ImageNet train set and save them in the `./tmp` folder:

```
python ./tools/visualizations/vis_pipeline.py configs/swin_transformer/swin_base_224_
↳ b16x64_300e_imagenet.py --phase train --output-dir tmp --number 10 --adaptive
```

1. Visualize 100 original pictures in the CIFAR100 validation set, then display and save them in the `./tmp` folder:

```
python ./tools/visualizations/vis_pipeline.py configs/resnet/resnet50_8xb16_cifar100.py -
↳ -phase val --output-dir tmp --mode original --number 100 --show --adaptive --bgr2rgb
```

32.2 Learning Rate Schedule Visualization

```
python tools/visualizations/vis_lr.py \
  ${CONFIG_FILE} \
  --dataset-size ${DATASET_SIZE} \
  --ngpus ${NUM_GPUS}
  --save-path ${SAVE_PATH} \
  --title ${TITLE} \
  --style ${STYLE} \
  --window-size ${WINDOW_SIZE}
  --cfg-options
```

Description of all arguments:

- `config` : The path of a model config file.

- `dataset-size` : The size of the datasets. If set, `build_dataset` will be skipped and `#{DATASET_SIZE}` will be used as the size. Default to use the function `build_dataset`.
- `ngpus` : The number of GPUs used in training, default to be 1.
- `save-path` : The learning rate curve plot save path, default not to save.
- `title` : Title of figure. If not set, default to be config file name.
- `style` : Style of plt. If not set, default to be `whitegrid`.
- `window-size` : The shape of the display window. If not specified, it will be set to `12*7`. If used, it must be in the format `'W*H'`.
- `cfg-options` : Modifications to the configuration file, refer to [Tutorial 1: Learn about Configs](#).

Note: Loading annotations maybe consume much time, you can directly specify the size of the dataset with `dataset-size` to save time.

Examples:

```
python tools/visualizations/vis_lr.py configs/resnet/resnet50_b16x8_cifar100.py
```

When using ImageNet, directly specify the size of ImageNet, as below:

```
python tools/visualizations/vis_lr.py configs/repvgg/repvgg-B3g4_4xb64-autoaug-lbs-mixup-  
↪coslr-200e_in1k.py --dataset-size 1281167 --ngpus 4 --save-path ./repvgg-B3g4_4xb64-lr.  
↪jpg
```

32.3 FAQs

- None

CONTRIBUTING TO OPENMMLAB

All kinds of contributions are welcome, including but not limited to the following.

- Fixes (typo, bugs)
- New features and components

33.1 Workflow

1. fork and pull the latest OpenMMLab repository (mmclassification)
2. checkout a new branch (do not use master branch for PRs)
3. commit your changes
4. create a PR

Note: If you plan to add some new features that involve large changes, it is encouraged to open an issue for discussion first.

33.2 Code style

33.2.1 Python

We adopt [PEP8](#) as the preferred code style.

We use the following tools for linting and formatting:

- [flake8](#): A wrapper around some linter tools.
- [yapf](#): A formatter for Python files.
- [isort](#): A Python utility to sort imports.
- [markdownlint](#): A linter to check markdown files and flag style issues.
- [docformatter](#): A formatter to format docstring.

Style configurations of yapf and isort can be found in [setup.cfg](#).

We use [pre-commit hook](#) that checks and formats for [flake8](#), [yapf](#), [isort](#), [trailing whitespaces](#), [markdown files](#), [fixes end-of-files](#), [double-quoted-strings](#), [python-encoding-pragma](#), [mixed-line-ending](#), [sorts requirements.txt](#) automatically on every commit. The config for a pre-commit hook is stored in [.pre-commit-config](#).

After you clone the repository, you will need to install initialize pre-commit hook.

```
pip install -U pre-commit
```

From the repository folder

```
pre-commit install
```

Try the following steps to install ruby when you encounter an issue on installing markdownlint

```
# install rvm
curl -L https://get.rvm.io | bash -s -- --autolibs=read-fail
[[ -s "$HOME/.rvm/scripts/rvm" ]] && source "$HOME/.rvm/scripts/rvm"
rvm autolibs disable

# install ruby
rvm install 2.7.1
```

Or refer to [this repo](#) and take `zzruby.sh` according its instruction.

After this on every commit check code linters and formatter will be enforced.

Important: Before you create a PR, make sure that your code lints and is formatted by yapf.

33.2.2 C++ and CUDA

We follow the [Google C++ Style Guide](#).

`mmcls.apis.inference_model(model, img)`
Inference image(s) with the classifier.

Parameters

- **model** (*nn.Module*) – The loaded classifier.
- **img** (*str/ndarray*) – The image filename or loaded image.

Returns

The classification results that contains *class_name*, *pred_label* and *pred_score*.

Return type result (dict)

`mmcls.apis.init_model(config, checkpoint=None, device='cuda:0', options=None)`
Initialize a classifier from config file.

Parameters

- **config** (*str* or *mmcv.Config*) – Config file path or the config object.
- **checkpoint** (*str*, *optional*) – Checkpoint path. If left as *None*, the model will not load any weights.
- **options** (*dict*) – Options to override some settings in the used config.

Returns The constructed classifier.

Return type *nn.Module*

`mmcls.apis.multi_gpu_test(model, data_loader, tmpdir=None, gpu_collect=False)`
Test model with multiple gpus.

This method tests model with multiple gpus and collects the results under two different modes: *gpu* and *cpu* modes. By setting '*gpu_collect=True*' it encodes results to *gpu* tensors and use *gpu* communication for results collection. On *cpu* mode it saves the results on different gpus to '*tmpdir*' and collects them by the rank 0 worker.

Parameters

- **model** (*nn.Module*) – Model to be tested.
- **data_loader** (*nn.DataLoader*) – Pytorch data loader.
- **tmpdir** (*str*) – Path of directory to save the temporary results from different gpus under *cpu* mode.
- **gpu_collect** (*bool*) – Option to use either *gpu* or *cpu* to collect results.

Returns The prediction results.

Return type list

`mmcls.apis.set_random_seed(seed, deterministic=False)`

Set random seed.

Parameters

- **seed** (*int*) – Seed to be used.
- **deterministic** (*bool*) – Whether to set the deterministic option for CUDNN backend, i.e., set `torch.backends.cudnn.deterministic` to True and `torch.backends.cudnn.benchmark` to False. Default: False.

`mmcls.apis.show_result_pyplot(model, img, result, fig_size=(15, 10), title='result', wait_time=0)`

Visualize the classification results on the image.

Parameters

- **model** (*nn.Module*) – The loaded classifier.
- **img** (*str or np.ndarray*) – Image filename or loaded image.
- **result** (*list*) – The classification result.
- **fig_size** (*tuple*) – Figure size of the pyplot figure. Defaults to (15, 10).
- **title** (*str*) – Title of the pyplot figure. Defaults to 'result'.
- **wait_time** (*int*) – How many seconds to display the image. Defaults to 0.

35.1 evaluation

```
class mmcls.core.evaluation.DistEvalHook(dataloader, interval=1, gpu_collect=False, by_epoch=True,  
                                         **eval_kwargs)
```

Distributed evaluation hook.

Parameters

- **dataloader** (*DataLoader*) – A PyTorch dataloader.
- **interval** (*int*) – Evaluation interval (by epochs). Default: 1.
- **tmpdir** (*str*, *optional*) – Temporary directory to save the results of all processes. Default: None.
- **gpu_collect** (*bool*) – Whether to use gpu or cpu to collect results. Default: False.

```
class mmcls.core.evaluation.EvalHook(dataloader, interval=1, by_epoch=True, **eval_kwargs)
```

Evaluation hook.

Parameters

- **dataloader** (*DataLoader*) – A PyTorch dataloader.
- **interval** (*int*) – Evaluation interval (by epochs). Default: 1.

```
mmcls.core.evaluation.average_performance(pred, target, thr=None, k=None)
```

Calculate CP, CR, CF1, OP, OR, OF1, where C stands for per-class average, O stands for overall average, P stands for precision, R stands for recall and F1 stands for F1-score.

Parameters

- **pred** (*torch.Tensor* / *np.ndarray*) – The model prediction with shape (N, C), where C is the number of classes.
- **target** (*torch.Tensor* / *np.ndarray*) – The target of each prediction with shape (N, C), where C is the number of classes. 1 stands for positive examples, 0 stands for negative examples and -1 stands for difficult examples.
- **thr** (*float*) – The confidence threshold. Defaults to None.
- **k** (*int*) – Top-k performance. Note that if thr and k are both given, k will be ignored. Defaults to None.

Returns (CP, CR, CF1, OP, OR, OF1)

Return type tuple

`mmcls.core.evaluation.average_precision(pred, target)`

Calculate the average precision for a single class.

AP summarizes a precision-recall curve as the weighted mean of maximum precisions obtained for any $r' > r$, where r is the recall:

$$AP = \sum_n (R_n - R_{n-1}) P_n$$

Note that no approximation is involved since the curve is piecewise constant.

Parameters

- **pred** (*np.ndarray*) – The model prediction with shape (N,).
- **target** (*np.ndarray*) – The target of each prediction with shape (N,).

Returns a single float as average precision value.

Return type float

`mmcls.core.evaluation.calculate_confusion_matrix(pred, target)`

Calculate confusion matrix according to the prediction and target.

Parameters

- **pred** (*torch.Tensor* | *np.array*) – The model prediction with shape (N, C).
- **target** (*torch.Tensor* | *np.array*) – The target of each prediction with shape (N, 1) or (N,).

Returns

Confusion matrix The shape is (C, C), where C is the number of classes.

Return type torch.Tensor

`mmcls.core.evaluation.f1_score(pred, target, average_mode='macro', thrs=0.0)`

Calculate F1 score according to the prediction and target.

Parameters

- **pred** (*torch.Tensor* | *np.array*) – The model prediction with shape (N, C).
- **target** (*torch.Tensor* | *np.array*) – The target of each prediction with shape (N, 1) or (N,).
- **average_mode** (*str*) – The type of averaging performed on the result. Options are ‘macro’ and ‘none’. If ‘none’, the scores for each class are returned. If ‘macro’, calculate metrics for each class, and find their unweighted mean. Defaults to ‘macro’.
- **thrs** (*Number* | *tuple[Number]*, *optional*) – Predictions with scores under the thresholds are considered negative. Default to 0.

Returns F1 score.

Return type

float | np.array | list[float | np.array]

Args	thrs is number	thrs is tuple
average_mode = “macro”	float	list[float]
average_mode = “none”	np.array	list[np.array]

`mmcls.core.evaluation.mAP(pred, target)`

Calculate the mean average precision with respect of classes.

Parameters

- **pred** (*torch.Tensor* | *np.ndarray*) – The model prediction with shape (N, C), where C is the number of classes.
- **target** (*torch.Tensor* | *np.ndarray*) – The target of each prediction with shape (N, C), where C is the number of classes. 1 stands for positive examples, 0 stands for negative examples and -1 stands for difficult examples.

Returns A single float as mAP value.

Return type float

`mmcls.core.evaluation.precision(pred, target, average_mode='macro', thrs=0.0)`

Calculate precision according to the prediction and target.

Parameters

- **pred** (*torch.Tensor* | *np.array*) – The model prediction with shape (N, C).
- **target** (*torch.Tensor* | *np.array*) – The target of each prediction with shape (N, 1) or (N,).
- **average_mode** (*str*) – The type of averaging performed on the result. Options are ‘macro’ and ‘none’. If ‘none’, the scores for each class are returned. If ‘macro’, calculate metrics for each class, and find their unweighted mean. Defaults to ‘macro’.
- **thrs** (*Number* | *tuple[Number]*, *optional*) – Predictions with scores under the thresholds are considered negative. Default to 0.

Returns Precision.

Return type

float | np.array | list[float | np.array]

Args	thrs is number	thrs is tuple
average_mode = “macro”	float	list[float]
average_mode = “none”	np.array	list[np.array]

`mmcls.core.evaluation.precision_recall_f1(pred, target, average_mode='macro', thrs=0.0)`

Calculate precision, recall and f1 score according to the prediction and target.

Parameters

- **pred** (*torch.Tensor* | *np.array*) – The model prediction with shape (N, C).
- **target** (*torch.Tensor* | *np.array*) – The target of each prediction with shape (N, 1) or (N,).
- **average_mode** (*str*) – The type of averaging performed on the result. Options are ‘macro’ and ‘none’. If ‘none’, the scores for each class are returned. If ‘macro’, calculate metrics for each class, and find their unweighted mean. Defaults to ‘macro’.
- **thrs** (*Number* | *tuple[Number]*, *optional*) – Predictions with scores under the thresholds are considered negative. Default to 0.

Returns

tuple containing precision, recall, f1 score.

The type of precision, recall, f1 score is one of the following:

Args	thrs is number	thrs is tuple
average_mode = "macro"	float	list[float]
average_mode = "none"	np.array	list[np.array]

Return type tuple

`mmcls.core.evaluation.recall(pred, target, average_mode='macro', thrs=0.0)`

Calculate recall according to the prediction and target.

Parameters

- **pred** (*torch.Tensor* / *np.array*) – The model prediction with shape (N, C).
- **target** (*torch.Tensor* / *np.array*) – The target of each prediction with shape (N, 1) or (N,).
- **average_mode** (*str*) – The type of averaging performed on the result. Options are ‘macro’ and ‘none’. If ‘none’, the scores for each class are returned. If ‘macro’, calculate metrics for each class, and find their unweighted mean. Defaults to ‘macro’.
- **thrs** (*Number* / *tuple[Number]*, *optional*) – Predictions with scores under the thresholds are considered negative. Default to 0.

Returns Recall.

Return type

float | np.array | list[float | np.array]

Args	thrs is number	thrs is tuple
average_mode = "macro"	float	list[float]
average_mode = "none"	np.array	list[np.array]

`mmcls.core.evaluation.support(pred, target, average_mode='macro')`

Calculate the total number of occurrences of each label according to the prediction and target.

Parameters

- **pred** (*torch.Tensor* / *np.array*) – The model prediction with shape (N, C).
- **target** (*torch.Tensor* / *np.array*) – The target of each prediction with shape (N, 1) or (N,).
- **average_mode** (*str*) – The type of averaging performed on the result. Options are ‘macro’ and ‘none’. If ‘none’, the scores for each class are returned. If ‘macro’, calculate metrics for each class, and find their unweighted sum. Defaults to ‘macro’.

Returns

Support.

- If the `average_mode` is set to `macro`, the function returns a single float.
- If the `average_mode` is set to `none`, the function returns a `np.array` with shape C.

Return type float | np.array

MMCLS.MODELS

36.1 models

`mmcls.models.build_backbone(cfg)`
Build backbone.

`mmcls.models.build_head(cfg)`
Build head.

`mmcls.models.build_loss(cfg)`
Build loss.

`mmcls.models.build_neck(cfg)`
Build neck.

36.2 classifiers

class `mmcls.models.classifiers.BaseClassifier`(*init_cfg=None*)
Base class for classifiers.

forward(*img*, *return_loss=True*, ***kwargs*)
Calls either `forward_train` or `forward_test` depending on whether `return_loss=True`.

Note this setting will change the expected inputs. When `return_loss=True`, `img` and `img_meta` are single-nested (i.e. `Tensor` and `List[dict]`), and when `return_loss=False`, `img` and `img_meta` should be double-nested (i.e. `List[Tensor]`, `List[List[dict]]`), with the outer list indicating test time augmentations.

forward_test(*imgs*, ***kwargs*)

Parameters `imgs` (`List[Tensor]`) – the outer list indicates test-time augmentations and inner `Tensor` should have a shape `NxCxHxW`, which contains all images in the batch.

abstract forward_train(*imgs*, ***kwargs*)

Parameters

- `img` (`list[Tensor]`) – List of tensors of shape (1, C, H, W). Typically these should be mean centered and std scaled.
- `kwargs` (*keyword arguments*) – Specific to concrete implementation.

show_result(*img*, *result*, *text_color*='white', *font_scale*=0.5, *row_width*=20, *show*=False, *fig_size*=(15, 10), *win_name*="", *wait_time*=0, *out_file*=None)

Draw *result* over *img*.

Parameters

- **img** (*str* or *ndarray*) – The image to be displayed.
- **result** (*dict*) – The classification results to draw over *img*.
- **text_color** (*str* or *tuple* or *Color*) – Color of texts.
- **font_scale** (*float*) – Font scales of texts.
- **row_width** (*int*) – width between each row of results on the image.
- **show** (*bool*) – Whether to show the image. Default: False.
- **fig_size** (*tuple*) – Image show figure size. Defaults to (15, 10).
- **win_name** (*str*) – The window name.
- **wait_time** (*int*) – How many seconds to display the image. Defaults to 0.
- **out_file** (*str* or *None*) – The filename to write the image. Default: None.

Returns Image with overlaid results.

Return type *img* (*ndarray*)

train_step(*data*, *optimizer*=None, ***kwargs*)

The iteration step during training.

This method defines an iteration step during training, except for the back propagation and optimizer updating, which are done in an optimizer hook. Note that in some complicated cases or models, the whole process including back propagation and optimizer updating are also defined in this method, such as GAN.

Parameters

- **data** (*dict*) – The output of dataloader.
- **optimizer** (*torch.optim.Optimizer* | *dict*, optional) – The optimizer of runner is passed to `train_step()`. This argument is unused and reserved.

Returns

Dict of outputs. The following fields are contained.

- **loss** (*torch.Tensor*): A tensor for back propagation, which can be a weighted sum of multiple losses.
- **log_vars** (*dict*): Dict contains all the variables to be sent to the logger.
- **num_samples** (*int*): Indicates the batch size (when the model is DDP, it means the batch size on each GPU), which is used for averaging the logs.

Return type *dict*

val_step(*data*, *optimizer*=None, ***kwargs*)

The iteration step during validation.

This method shares the same signature as `train_step()`, but used during val epochs. Note that the evaluation after training epochs is not implemented with this method, but an evaluation hook.

Parameters

- **data** (*dict*) – The output of dataloader.

- **optimizer** (`torch.optim.Optimizer` | dict, optional) – The optimizer of runner is passed to `train_step()`. This argument is unused and reserved.

Returns

Dict of outputs. The following fields are contained.

- **loss** (`torch.Tensor`): A tensor for back propagation, which can be a weighted sum of multiple losses.
- **log_vars** (dict): Dict contains all the variables to be sent to the logger.
- **num_samples** (int): Indicates the batch size (when the model is DDP, it means the batch size on each GPU), which is used for averaging the logs.

Return type dict

```
class mmcls.models.classifiers.ImageClassifier(backbone, neck=None, head=None, pretrained=None,
                                              train_cfg=None, init_cfg=None)
```

extract_feat(*img*)

Directly extract features from the backbone + neck.

forward_train(*img*, *gt_label*, ***kwargs*)

Forward computation during training.

Parameters

- **img** (*Tensor*) – of shape (N, C, H, W) encoding input images. Typically these should be mean centered and std scaled.
- **gt_label** (*Tensor*) – It should be of shape (N, 1) encoding the ground-truth label of input images for single label task. It should be of shape (N, C) encoding the ground-truth label of input images for multi-labels task.

Returns a dictionary of loss components

Return type dict[str, Tensor]

simple_test(*img*, *img metas*)

Test without augmentation.

36.3 backbones

```
class mmcls.models.backbones.AlexNet(num_classes=-1)
```

AlexNet backbone.

The input for AlexNet is a 224x224 RGB image.

Parameters **num_classes** (*int*) – number of classes for classification. The default value is -1, which uses the backbone as a feature extractor without the top classifier.

forward(*x*)

Forward computation.

Parameters **x** (*tensor* | *tuple[tensor]*) – x could be a Torch.tensor or a tuple of Torch.tensor, containing input data for forward computation.

```
class mmcls.models.backbones.LeNet5(num_classes=-1)
```

LeNet5 backbone.

The input for LeNet-5 is a 32x32 grayscale image.

Parameters `num_classes` (*int*) – number of classes for classification. The default value is -1, which uses the backbone as a feature extractor without the top classifier.

forward(*x*)

Forward computation.

Parameters `x` (*tensor* | *tuple*[*tensor*]) – `x` could be a `Torch.tensor` or a tuple of `Torch.tensor`, containing input data for forward computation.

```
class mmcls.models.backbones.MlpMixer(arch='b', img_size=224, patch_size=16, out_indices=-1,
                                     drop_rate=0.0, drop_path_rate=0.0, norm_cfg={'type': 'LN'},
                                     act_cfg={'type': 'GELU'}, patch_cfg={}, layer_cfgs={},
                                     init_cfg=None)
```

Mlp-Mixer backbone.

Pytorch implementation of [MLP-Mixer: An all-MLP Architecture for Vision](#)

Parameters

- **arch** (*str* | *dict*) – MLP Mixer architecture Defaults to 'b'.
- **img_size** (*int* | *tuple*) – Input image size.
- **patch_size** (*int* | *tuple*) – The patch size.
- **out_indices** (*Sequence* | *int*) – Output from which layer. Defaults to -1, means the last layer.
- **drop_rate** (*float*) – Probability of an element to be zeroed. Defaults to 0.
- **drop_path_rate** (*float*) – stochastic depth rate. Defaults to 0.
- **norm_cfg** (*dict*) – Config dict for normalization layer. Defaults to `dict(type='LN')`.
- **act_cfg** (*dict*) – The activation config for FFNs. Default GELU.
- **patch_cfg** (*dict*) – Configs of patch embedding. Defaults to an empty dict.
- **layer_cfgs** (*Sequence* | *dict*) – Configs of each mixer block layer. Defaults to an empty dict.
- **init_cfg** (*dict*, *optional*) – Initialization config dict. Defaults to None.

forward(*x*)

Forward computation.

Parameters `x` (*tensor* | *tuple*[*tensor*]) – `x` could be a `Torch.tensor` or a tuple of `Torch.tensor`, containing input data for forward computation.

```
class mmcls.models.backbones.MobileNetV2(widen_factor=1.0, out_indices=(7), frozen_stages=-1,
                                          conv_cfg=None, norm_cfg={'type': 'BN'}, act_cfg={'type':
                                          'ReLU6'}, norm_eval=False, with_cp=False, init_cfg=[{'type':
                                          'Kaiming', 'layer': ['Conv2d']}, {'type': 'Constant', 'val': 1,
                                          'layer': ['_BatchNorm', 'GroupNorm']}]])
```

MobileNetV2 backbone.

Parameters

- **widen_factor** (*float*) – Width multiplier, multiply number of channels in each layer by this amount. Default: 1.0.
- **out_indices** (*None* or *Sequence*[*int*]) – Output from which stages. Default: (7,).
- **frozen_stages** (*int*) – Stages to be frozen (all param fixed). Default: -1, which means not freezing any parameters.

- **conv_cfg** (*dict*, *optional*) – Config dict for convolution layer. Default: None, which means using conv2d.
- **norm_cfg** (*dict*) – Config dict for normalization layer. Default: dict(type='BN').
- **act_cfg** (*dict*) – Config dict for activation layer. Default: dict(type='ReLU6').
- **norm_eval** (*bool*) – Whether to set norm layers to eval mode, namely, freeze running stats (mean and var). Note: Effect on Batch Norm and its variants only. Default: False.
- **with_cp** (*bool*) – Use checkpoint or not. Using checkpoint will save some memory while slowing down the training speed. Default: False.

forward(*x*)

Forward computation.

Parameters *x* (*tensor* | *tuple*[*tensor*]) – *x* could be a Torch.tensor or a tuple of Torch.tensor, containing input data for forward computation.

make_layer(*out_channels*, *num_blocks*, *stride*, *expand_ratio*)

Stack InvertedResidual blocks to build a layer for MobileNetV2.

Parameters

- **out_channels** (*int*) – out_channels of block.
- **num_blocks** (*int*) – number of blocks.
- **stride** (*int*) – stride of the first block. Default: 1
- **expand_ratio** (*int*) – Expand the number of channels of the hidden layer in InvertedResidual by this ratio. Default: 6.

train(*mode=True*)

Set module status before forward computation.

Parameters *mode* (*bool*) – Whether it is train_mode or test_mode

```
class mmcls.models.backbones.MobileNetV3(arch='small', conv_cfg=None, norm_cfg={'eps': 0.001,
                                     'momentum': 0.01, 'type': 'BN'}, out_indices=None,
                                     frozen_stages=-1, norm_eval=False, with_cp=False,
                                     init_cfg=[{'type': 'Kaiming', 'layer': ['Conv2d'], 'nonlinearity':
                                     'leaky_relu'}, {'type': 'Normal', 'layer': ['Linear'], 'std': 0.01},
                                     {'type': 'Constant', 'layer': ['BatchNorm2d'], 'val': 1}])
```

MobileNetV3 backbone.

Parameters

- **arch** (*str*) – Architecture of mobilnetv3, from {small, large}. Default: small.
- **conv_cfg** (*dict*, *optional*) – Config dict for convolution layer. Default: None, which means using conv2d.
- **norm_cfg** (*dict*) – Config dict for normalization layer. Default: dict(type='BN').
- **out_indices** (*None* or *Sequence*[*int*]) – Output from which stages. Default: None, which means output tensors from final stage.
- **frozen_stages** (*int*) – Stages to be frozen (all param fixed). Default: -1, which means not freezing any parameters.
- **norm_eval** (*bool*) – Whether to set norm layers to eval mode, namely, freeze running stats (mean and var). Note: Effect on Batch Norm and its variants only. Default: False.

- **with_cp** (*bool*) – Use checkpoint or not. Using checkpoint will save some memory while slowing down the training speed. Default: False.

forward(*x*)

Forward computation.

Parameters *x* (*tensor* | *tuple*[*tensor*]) – *x* could be a `Torch.tensor` or a tuple of `Torch.tensor`, containing input data for forward computation.

train(*mode=True*)

Set module status before forward computation.

Parameters *mode* (*bool*) – Whether it is `train_mode` or `test_mode`

```
class mmcls.models.backbones.RegNet(arch, in_channels=3, stem_channels=32, base_channels=32,  
                                     strides=(2, 2, 2), dilations=(1, 1, 1, 1), out_indices=(3),  
                                     style='pytorch', deep_stem=False, avg_down=False, frozen_stages=-  
                                     1, conv_cfg=None, norm_cfg={'requires_grad': True, 'type': 'BN'},  
                                     norm_eval=False, with_cp=False, zero_init_residual=True,  
                                     init_cfg=None)
```

RegNet backbone.

More details can be found in [paper](#) .

Parameters

- **arch** (*dict*) – The parameter of RegNets. - *w0* (*int*): initial width - *wa* (*float*): slope of width - *wm* (*float*): quantization parameter to quantize the width - *depth* (*int*): depth of the backbone - *group_w* (*int*): width of group - *bot_mul* (*float*): bottleneck ratio, i.e. expansion of bottleneck.
- **strides** (*Sequence*[*int*]) – Strides of the first block of each stage.
- **base_channels** (*int*) – Base channels after stem layer.
- **in_channels** (*int*) – Number of input image channels. Default: 3.
- **dilations** (*Sequence*[*int*]) – Dilation of each stage.
- **out_indices** (*Sequence*[*int*]) – Output from which stages.
- **style** (*str*) – *pytorch* or *caffe*. If set to “pytorch”, the stride-two layer is the 3x3 conv layer, otherwise the stride-two layer is the first 1x1 conv layer. Default: “pytorch”.
- **frozen_stages** (*int*) – Stages to be frozen (all param fixed). -1 means not freezing any parameters. Default: -1.
- **norm_cfg** (*dict*) – dictionary to construct and config norm layer. Default: `dict(type='BN', requires_grad=True)`.
- **norm_eval** (*bool*) – Whether to set norm layers to eval mode, namely, freeze running stats (mean and var). Note: Effect on Batch Norm and its variants only. Default: False.
- **with_cp** (*bool*) – Use checkpoint or not. Using checkpoint will save some memory while slowing down the training speed. Default: False.
- **zero_init_residual** (*bool*) – whether to use zero init for last norm layer in resblocks to let them behave as identity. Default: True.

Example

```

>>> from mmcls.models import RegNet
>>> import torch
>>> self = RegNet(
    arch=dict(
        w0=88,
        wa=26.31,
        wm=2.25,
        group_w=48,
        depth=25,
        bot_mul=1.0))
>>> self.eval()
>>> inputs = torch.rand(1, 3, 32, 32)
>>> level_outputs = self.forward(inputs)
>>> for level_out in level_outputs:
...     print(tuple(level_out.shape))
(1, 96, 8, 8)
(1, 192, 4, 4)
(1, 432, 2, 2)
(1, 1008, 1, 1)

```

adjust_width_group(*widths*, *bottleneck_ratio*, *groups*)

Adjusts the compatibility of widths and groups.

Parameters

- **widths** (*list[int]*) – Width of each stage.
- **bottleneck_ratio** (*float*) – Bottleneck ratio.
- **groups** (*int*) – number of groups in each stage

Returns The adjusted widths and groups of each stage.

Return type tuple(list)

forward(*x*)

Forward computation.

Parameters *x* (*tensor* | *tuple[torch.tensor]*) – *x* could be a `Torch.tensor` or a tuple of `Torch.tensor`, containing input data for forward computation.

generate_regnet(*initial_width*, *width_slope*, *width_parameter*, *depth*, *divisor=8*)

Generates per block width from RegNet parameters.

Parameters

- **initial_width** (*[int]*) – Initial width of the backbone
- **width_slope** (*[float]*) – Slope of the quantized linear function
- **width_parameter** (*[int]*) – Parameter used to quantize the width.
- **depth** (*[int]*) – Depth of the backbone.
- **divisor** (*int*) – The divisor of channels. Defaults to 8.

Returns

tuple containing:

- list: Widths of each stage.

- `int`: The number of stages.

Return type tuple

get_stages_from_blocks(*widths*)

Gets widths/stage_blocks of network at each stage.

Parameters `widths` (*list[int]*) – Width in each stage.

Returns width and depth of each stage

Return type tuple(list)

static quantize_float(*number, divisor*)

Converts a float to closest non-zero int divisible by divisor.

Parameters

- **number** (*int*) – Original number to be quantized.
- **divisor** (*int*) – Divisor used to quantize the number.

Returns quantized number that is divisible by divisor.

Return type int

```
class mmcls.models.backbones.RepVGG(arch, in_channels=3, base_channels=64, out_indices=(3),
                                   strides=(2, 2, 2, 2), dilations=(1, 1, 1, 1), frozen_stages=-1,
                                   conv_cfg=None, norm_cfg={'type': 'BN'}, act_cfg={'type': 'ReLU'},
                                   with_cp=False, deploy=False, norm_eval=False, init_cfg=[{'type':
                                   'Kaiming', 'layer': ['Conv2d']}, {'type': 'Constant', 'val': 1, 'layer':
                                   ['_BatchNorm', 'GroupNorm']}]])
```

RepVGG backbone.

A PyTorch impl of : [RepVGG: Making VGG-style ConvNets Great Again](#)

Parameters

- **arch** (*str | dict*) – The parameter of RepVGG. If it's a dict, it should contain the following keys:
 - `num_blocks` (*Sequence[int]*): Number of blocks in each stage.
 - `width_factor` (*Sequence[float]*): Width deflator in each stage.
 - `group_layer_map` (*dict | None*): RepVGG Block that declares the need to apply group convolution.
 - `se_cfg` (*dict | None*): Se Layer config
- **in_channels** (*int*) – Number of input image channels. Default: 3.
- **base_channels** (*int*) – Base channels of RepVGG backbone, work with `width_factor` together. Default: 64.
- **out_indices** (*Sequence[int]*) – Output from which stages. Default: (3,).
- **strides** (*Sequence[int]*) – Strides of the first block of each stage. Default: (2, 2, 2, 2).
- **dilations** (*Sequence[int]*) – Dilation of each stage. Default: (1, 1, 1, 1).
- **frozen_stages** (*int*) – Stages to be frozen (all param fixed). -1 means not freezing any parameters. Default: -1.
- **conv_cfg** (*dict | None*) – The config dict for conv layers. Default: None.
- **norm_cfg** (*dict*) – The config dict for norm layers. Default: dict(type='BN').

- **act_cfg** (*dict*) – Config dict for activation layer. Default: `dict(type='ReLU')`.
- **with_cp** (*bool*) – Use checkpoint or not. Using checkpoint will save some memory while slowing down the training speed. Default: `False`.
- **deploy** (*bool*) – Whether to switch the model structure to deployment mode. Default: `False`.
- **norm_eval** (*bool*) – Whether to set norm layers to eval mode, namely, freeze running stats (mean and var). Note: Effect on Batch Norm and its variants only. Default: `False`.
- **init_cfg** (*dict or list[dict], optional*) – Initialization config dict.

forward(*x*)

Forward computation.

Parameters *x* (*tensor | tuple[tensor]*) – *x* could be a `Torch.tensor` or a tuple of `Torch.tensor`, containing input data for forward computation.

train(*mode=True*)

Set module status before forward computation.

Parameters *mode* (*bool*) – Whether it is `train_mode` or `test_mode`

```
class mmls.models.backbones.Res2Net(scales=4, base_width=26, style='pytorch', deep_stem=True,
                                     avg_down=True, init_cfg=None, **kwargs)
```

Res2Net backbone.

A PyTorch implement of : [Res2Net: A New Multi-scale Backbone Architecture](#)

Parameters

- **depth** (*int*) – Depth of Res2Net, choose from {50, 101, 152}.
- **scales** (*int*) – Scales used in Res2Net. Defaults to 4.
- **base_width** (*int*) – Basic width of each scale. Defaults to 26.
- **in_channels** (*int*) – Number of input image channels. Defaults to 3.
- **num_stages** (*int*) – Number of Res2Net stages. Defaults to 4.
- **strides** (*Sequence[int]*) – Strides of the first block of each stage. Defaults to (1, 2, 2, 2).
- **dilations** (*Sequence[int]*) – Dilation of each stage. Defaults to (1, 1, 1, 1).
- **out_indices** (*Sequence[int]*) – Output from which stages. Defaults to (3,).
- **style** (*str*) – “pytorch” or “caffe”. If set to “pytorch”, the stride-two layer is the 3x3 conv layer, otherwise the stride-two layer is the first 1x1 conv layer. Defaults to “pytorch”.
- **deep_stem** (*bool*) – Replace 7x7 conv in input stem with 3 3x3 conv. Defaults to `True`.
- **avg_down** (*bool*) – Use AvgPool instead of stride conv when downsampling in the bottle2neck. Defaults to `True`.
- **frozen_stages** (*int*) – Stages to be frozen (stop grad and set eval mode). -1 means not freezing any parameters. Defaults to -1.
- **norm_cfg** (*dict*) – Dictionary to construct and config norm layer. Defaults to `dict(type='BN', requires_grad=True)`.
- **norm_eval** (*bool*) – Whether to set norm layers to eval mode, namely, freeze running stats (mean and var). Note: Effect on Batch Norm and its variants only. Defaults to `False`.
- **with_cp** (*bool*) – Use checkpoint or not. Using checkpoint will save some memory while slowing down the training speed. Defaults to `False`.

- **zero_init_residual** (*bool*) – Whether to use zero init for last norm layer in resblocks to let them behave as identity. Defaults to True.
- **init_cfg** (*dict or list[dict], optional*) – Initialization config dict. Defaults to None.

Example

```
>>> from mmcls.models import Res2Net
>>> import torch
>>> model = Res2Net(depth=50,
...                 scales=4,
...                 base_width=26,
...                 out_indices=(0, 1, 2, 3))
>>> model.eval()
>>> inputs = torch.rand(1, 3, 32, 32)
>>> level_outputs = model.forward(inputs)
>>> for level_out in level_outputs:
...     print(tuple(level_out.shape))
(1, 256, 8, 8)
(1, 512, 4, 4)
(1, 1024, 2, 2)
(1, 2048, 1, 1)
```

```
class mmcls.models.backbones.ResNeSt(depth, groups=1, width_per_group=4, radix=2, reduction_factor=4,
                                     avg_down_stride=True, **kwargs)
```

ResNeSt backbone.

Please refer to the [paper](#) for details.

Parameters

- **depth** (*int*) – Network depth, from {50, 101, 152, 200}.
- **groups** (*int*) – Groups of conv2 in Bottleneck. Default: 32.
- **width_per_group** (*int*) – Width per group of conv2 in Bottleneck. Default: 4.
- **radix** (*int*) – Radix of SpltAtConv2d. Default: 2
- **reduction_factor** (*int*) – Reduction factor of SplitAttentionConv2d. Default: 4.
- **avg_down_stride** (*bool*) – Whether to use average pool for stride in Bottleneck. Default: True.
- **in_channels** (*int*) – Number of input image channels. Default: 3.
- **stem_channels** (*int*) – Output channels of the stem layer. Default: 64.
- **num_stages** (*int*) – Stages of the network. Default: 4.
- **strides** (*Sequence[int]*) – Strides of the first block of each stage. Default: (1, 2, 2, 2).
- **dilations** (*Sequence[int]*) – Dilation of each stage. Default: (1, 1, 1, 1).
- **out_indices** (*Sequence[int]*) – Output from which stages. If only one stage is specified, a single tensor (feature map) is returned, otherwise multiple stages are specified, a tuple of tensors will be returned. Default: (3,).

- **style** (*str*) – *pytorch* or *caffe*. If set to “pytorch”, the stride-two layer is the 3x3 conv layer, otherwise the stride-two layer is the first 1x1 conv layer.
- **deep_stem** (*bool*) – Replace 7x7 conv in input stem with 3 3x3 conv. Default: False.
- **avg_down** (*bool*) – Use AvgPool instead of stride conv when downsampling in the bottleneck. Default: False.
- **frozen_stages** (*int*) – Stages to be frozen (stop grad and set eval mode). -1 means not freezing any parameters. Default: -1.
- **conv_cfg** (*dict* | *None*) – The config dict for conv layers. Default: None.
- **norm_cfg** (*dict*) – The config dict for norm layers.
- **norm_eval** (*bool*) – Whether to set norm layers to eval mode, namely, freeze running stats (mean and var). Note: Effect on Batch Norm and its variants only. Default: False.
- **with_cp** (*bool*) – Use checkpoint or not. Using checkpoint will save some memory while slowing down the training speed. Default: False.
- **zero_init_residual** (*bool*) – Whether to use zero init for last norm layer in resblocks to let them behave as identity. Default: True.

`class mmcls.models.backbones.ResNeXt(depth, groups=32, width_per_group=4, **kwargs)`
ResNeXt backbone.

Please refer to the [paper](#) for details.

Parameters

- **depth** (*int*) – Network depth, from {50, 101, 152}.
- **groups** (*int*) – Groups of conv2 in Bottleneck. Default: 32.
- **width_per_group** (*int*) – Width per group of conv2 in Bottleneck. Default: 4.
- **in_channels** (*int*) – Number of input image channels. Default: 3.
- **stem_channels** (*int*) – Output channels of the stem layer. Default: 64.
- **num_stages** (*int*) – Stages of the network. Default: 4.
- **strides** (*Sequence[int]*) – Strides of the first block of each stage. Default: (1, 2, 2, 2).
- **dilations** (*Sequence[int]*) – Dilation of each stage. Default: (1, 1, 1, 1).
- **out_indices** (*Sequence[int]*) – Output from which stages. If only one stage is specified, a single tensor (feature map) is returned, otherwise multiple stages are specified, a tuple of tensors will be returned. Default: (3,).
- **style** (*str*) – *pytorch* or *caffe*. If set to “pytorch”, the stride-two layer is the 3x3 conv layer, otherwise the stride-two layer is the first 1x1 conv layer.
- **deep_stem** (*bool*) – Replace 7x7 conv in input stem with 3 3x3 conv. Default: False.
- **avg_down** (*bool*) – Use AvgPool instead of stride conv when downsampling in the bottleneck. Default: False.
- **frozen_stages** (*int*) – Stages to be frozen (stop grad and set eval mode). -1 means not freezing any parameters. Default: -1.
- **conv_cfg** (*dict* | *None*) – The config dict for conv layers. Default: None.
- **norm_cfg** (*dict*) – The config dict for norm layers.

- **norm_eval** (*bool*) – Whether to set norm layers to eval mode, namely, freeze running stats (mean and var). Note: Effect on Batch Norm and its variants only. Default: False.
- **with_cp** (*bool*) – Use checkpoint or not. Using checkpoint will save some memory while slowing down the training speed. Default: False.
- **zero_init_residual** (*bool*) – Whether to use zero init for last norm layer in resblocks to let them behave as identity. Default: True.

```
class mmcls.models.backbones.ResNet(depth, in_channels=3, stem_channels=64, base_channels=64,
                                   expansion=None, num_stages=4, strides=(1, 2, 2, 2), dilations=(1, 1,
                                   1, 1), out_indices=(3), style='pytorch', deep_stem=False,
                                   avg_down=False, frozen_stages=-1, conv_cfg=None,
                                   norm_cfg={'requires_grad': True, 'type': 'BN'}, norm_eval=False,
                                   with_cp=False, zero_init_residual=True, init_cfg=[{'type':
                                   'Kaiming', 'layer': ['Conv2d']}, {'type': 'Constant', 'val': 1, 'layer':
                                   ['_BatchNorm', 'GroupNorm']}]])
```

ResNet backbone.

Please refer to the [paper](#) for details.

Parameters

- **depth** (*int*) – Network depth, from {18, 34, 50, 101, 152}.
- **in_channels** (*int*) – Number of input image channels. Default: 3.
- **stem_channels** (*int*) – Output channels of the stem layer. Default: 64.
- **base_channels** (*int*) – Middle channels of the first stage. Default: 64.
- **num_stages** (*int*) – Stages of the network. Default: 4.
- **strides** (*Sequence[int]*) – Strides of the first block of each stage. Default: (1, 2, 2, 2).
- **dilations** (*Sequence[int]*) – Dilation of each stage. Default: (1, 1, 1, 1).
- **out_indices** (*Sequence[int]*) – Output from which stages. Default: (3,).
- **style** (*str*) – *pytorch* or *caffe*. If set to “pytorch”, the stride-two layer is the 3x3 conv layer, otherwise the stride-two layer is the first 1x1 conv layer.
- **deep_stem** (*bool*) – Replace 7x7 conv in input stem with 3 3x3 conv. Default: False.
- **avg_down** (*bool*) – Use AvgPool instead of stride conv when downsampling in the bottleneck. Default: False.
- **frozen_stages** (*int*) – Stages to be frozen (stop grad and set eval mode). -1 means not freezing any parameters. Default: -1.
- **conv_cfg** (*dict* | *None*) – The config dict for conv layers. Default: None.
- **norm_cfg** (*dict*) – The config dict for norm layers.
- **norm_eval** (*bool*) – Whether to set norm layers to eval mode, namely, freeze running stats (mean and var). Note: Effect on Batch Norm and its variants only. Default: False.
- **with_cp** (*bool*) – Use checkpoint or not. Using checkpoint will save some memory while slowing down the training speed. Default: False.
- **zero_init_residual** (*bool*) – Whether to use zero init for last norm layer in resblocks to let them behave as identity. Default: True.

Example

```

>>> from mmcls.models import ResNet
>>> import torch
>>> self = ResNet(depth=18)
>>> self.eval()
>>> inputs = torch.rand(1, 3, 32, 32)
>>> level_outputs = self.forward(inputs)
>>> for level_out in level_outputs:
...     print(tuple(level_out.shape))
(1, 64, 8, 8)
(1, 128, 4, 4)
(1, 256, 2, 2)
(1, 512, 1, 1)

```

forward(*x*)

Forward computation.

Parameters *x* (*tensor* | *tuple*[*tensor*]) – *x* could be a `Torch.tensor` or a tuple of `Torch.tensor`, containing input data for forward computation.

init_weights()

Initialize the weights.

train(*mode=True*)

Set module status before forward computation.

Parameters *mode* (*bool*) – Whether it is `train_mode` or `test_mode`

class `mmcls.models.backbones.ResNetV1d`(***kwargs*)

`ResNetV1d` backbone.

This variant is described in [Bag of Tricks](#)..

Compared with default `ResNet(ResNetV1b)`, `ResNetV1d` replaces the 7×7 conv in the input stem with three 3×3 convs. And in the downsampling block, a 2×2 avg_pool with stride 2 is added before conv, whose stride is changed to 1.

class `mmcls.models.backbones.ResNet_CIFAR`(*depth*, *deep_stem=False*, ***kwargs*)

`ResNet` backbone for CIFAR.

Compared to standard `ResNet`, it uses `kernel_size=3` and `stride=1` in conv1, and does not apply `MaxPooling` after stem. It has been proven to be more efficient than standard `ResNet` in other public codebase, e.g., <https://github.com/kuangliu/pytorch-cifar/blob/master/models/resnet.py>.

Parameters

- **depth** (*int*) – Network depth, from {18, 34, 50, 101, 152}.
- **in_channels** (*int*) – Number of input image channels. Default: 3.
- **stem_channels** (*int*) – Output channels of the stem layer. Default: 64.
- **base_channels** (*int*) – Middle channels of the first stage. Default: 64.
- **num_stages** (*int*) – Stages of the network. Default: 4.
- **strides** (*Sequence*[*int*]) – Strides of the first block of each stage. Default: (1, 2, 2, 2).
- **dilations** (*Sequence*[*int*]) – Dilation of each stage. Default: (1, 1, 1, 1).

- **out_indices** (*Sequence[int]*) – Output from which stages. If only one stage is specified, a single tensor (feature map) is returned, otherwise multiple stages are specified, a tuple of tensors will be returned. Default: (3,).
- **style** (*str*) – *pytorch* or *caffe*. If set to “pytorch”, the stride-two layer is the 3x3 conv layer, otherwise the stride-two layer is the first 1x1 conv layer.
- **deep_stem** (*bool*) – This network has specific designed stem, thus it is asserted to be False.
- **avg_down** (*bool*) – Use AvgPool instead of stride conv when downsampling in the bottleneck. Default: False.
- **frozen_stages** (*int*) – Stages to be frozen (stop grad and set eval mode). -1 means not freezing any parameters. Default: -1.
- **conv_cfg** (*dict | None*) – The config dict for conv layers. Default: None.
- **norm_cfg** (*dict*) – The config dict for norm layers.
- **norm_eval** (*bool*) – Whether to set norm layers to eval mode, namely, freeze running stats (mean and var). Note: Effect on Batch Norm and its variants only. Default: False.
- **with_cp** (*bool*) – Use checkpoint or not. Using checkpoint will save some memory while slowing down the training speed. Default: False.
- **zero_init_residual** (*bool*) – Whether to use zero init for last norm layer in resblocks to let them behave as identity. Default: True.

forward(*x*)

Forward computation.

Parameters *x* (*tensor | tuple[tensor]*) – *x* could be a `Torch.tensor` or a tuple of `Torch.tensor`, containing input data for forward computation.

class `mmcls.models.backbones.SEResNeXt`(*depth, groups=32, width_per_group=4, **kwargs*)
SEResNeXt backbone.

Please refer to the [paper](#) for details.

Parameters

- **depth** (*int*) – Network depth, from {50, 101, 152}.
- **groups** (*int*) – Groups of conv2 in Bottleneck. Default: 32.
- **width_per_group** (*int*) – Width per group of conv2 in Bottleneck. Default: 4.
- **se_ratio** (*int*) – Squeeze ratio in SELayer. Default: 16.
- **in_channels** (*int*) – Number of input image channels. Default: 3.
- **stem_channels** (*int*) – Output channels of the stem layer. Default: 64.
- **num_stages** (*int*) – Stages of the network. Default: 4.
- **strides** (*Sequence[int]*) – Strides of the first block of each stage. Default: (1, 2, 2, 2).
- **dilations** (*Sequence[int]*) – Dilation of each stage. Default: (1, 1, 1, 1).
- **out_indices** (*Sequence[int]*) – Output from which stages. If only one stage is specified, a single tensor (feature map) is returned, otherwise multiple stages are specified, a tuple of tensors will be returned. Default: (3,).
- **style** (*str*) – *pytorch* or *caffe*. If set to “pytorch”, the stride-two layer is the 3x3 conv layer, otherwise the stride-two layer is the first 1x1 conv layer.

- **deep_stem** (*bool*) – Replace 7x7 conv in input stem with 3 3x3 conv. Default: False.
- **avg_down** (*bool*) – Use AvgPool instead of stride conv when downsampling in the bottleneck. Default: False.
- **frozen_stages** (*int*) – Stages to be frozen (stop grad and set eval mode). -1 means not freezing any parameters. Default: -1.
- **conv_cfg** (*dict* | *None*) – The config dict for conv layers. Default: None.
- **norm_cfg** (*dict*) – The config dict for norm layers.
- **norm_eval** (*bool*) – Whether to set norm layers to eval mode, namely, freeze running stats (mean and var). Note: Effect on Batch Norm and its variants only. Default: False.
- **with_cp** (*bool*) – Use checkpoint or not. Using checkpoint will save some memory while slowing down the training speed. Default: False.
- **zero_init_residual** (*bool*) – Whether to use zero init for last norm layer in resblocks to let them behave as identity. Default: True.

`class mmcls.models.backbones.SEResNet(depth, se_ratio=16, **kwargs)`
SEResNet backbone.

Please refer to the [paper](#) for details.

Parameters

- **depth** (*int*) – Network depth, from {50, 101, 152}.
- **se_ratio** (*int*) – Squeeze ratio in SELayer. Default: 16.
- **in_channels** (*int*) – Number of input image channels. Default: 3.
- **stem_channels** (*int*) – Output channels of the stem layer. Default: 64.
- **num_stages** (*int*) – Stages of the network. Default: 4.
- **strides** (*Sequence[int]*) – Strides of the first block of each stage. Default: (1, 2, 2, 2).
- **dilations** (*Sequence[int]*) – Dilation of each stage. Default: (1, 1, 1, 1).
- **out_indices** (*Sequence[int]*) – Output from which stages. If only one stage is specified, a single tensor (feature map) is returned, otherwise multiple stages are specified, a tuple of tensors will be returned. Default: (3,).
- **style** (*str*) – *pytorch* or *caffe*. If set to “pytorch”, the stride-two layer is the 3x3 conv layer, otherwise the stride-two layer is the first 1x1 conv layer.
- **deep_stem** (*bool*) – Replace 7x7 conv in input stem with 3 3x3 conv. Default: False.
- **avg_down** (*bool*) – Use AvgPool instead of stride conv when downsampling in the bottleneck. Default: False.
- **frozen_stages** (*int*) – Stages to be frozen (stop grad and set eval mode). -1 means not freezing any parameters. Default: -1.
- **conv_cfg** (*dict* | *None*) – The config dict for conv layers. Default: None.
- **norm_cfg** (*dict*) – The config dict for norm layers.
- **norm_eval** (*bool*) – Whether to set norm layers to eval mode, namely, freeze running stats (mean and var). Note: Effect on Batch Norm and its variants only. Default: False.
- **with_cp** (*bool*) – Use checkpoint or not. Using checkpoint will save some memory while slowing down the training speed. Default: False.

- **zero_init_residual** (*bool*) – Whether to use zero init for last norm layer in resblocks to let them behave as identity. Default: True.

Example

```
>>> from mmcls.models import SEResNet
>>> import torch
>>> self = SEResNet(depth=50)
>>> self.eval()
>>> inputs = torch.rand(1, 3, 224, 224)
>>> level_outputs = self.forward(inputs)
>>> for level_out in level_outputs:
...     print(tuple(level_out.shape))
(1, 64, 56, 56)
(1, 128, 28, 28)
(1, 256, 14, 14)
(1, 512, 7, 7)
```

```
class mmcls.models.backbones.ShuffleNetV1(groups=3, widen_factor=1.0, out_indices=(2),
                                          frozen_stages=-1, conv_cfg=None, norm_cfg={'type': 'BN'},
                                          act_cfg={'type': 'ReLU'}, norm_eval=False, with_cp=False,
                                          init_cfg=None)
```

ShuffleNetV1 backbone.

Parameters

- **groups** (*int*) – The number of groups to be used in grouped 1x1 convolutions in each ShuffleUnit. Default: 3.
- **widen_factor** (*float*) – Width multiplier - adjusts the number of channels in each layer by this amount. Default: 1.0.
- **out_indices** (*Sequence[int]*) – Output from which stages. Default: (2,)
- **frozen_stages** (*int*) – Stages to be frozen (all param fixed). Default: -1, which means not freezing any parameters.
- **conv_cfg** (*dict, optional*) – Config dict for convolution layer. Default: None, which means using conv2d.
- **norm_cfg** (*dict*) – Config dict for normalization layer. Default: dict(type='BN').
- **act_cfg** (*dict*) – Config dict for activation layer. Default: dict(type='ReLU').
- **norm_eval** (*bool*) – Whether to set norm layers to eval mode, namely, freeze running stats (mean and var). Note: Effect on Batch Norm and its variants only. Default: False.
- **with_cp** (*bool*) – Use checkpoint or not. Using checkpoint will save some memory while slowing down the training speed. Default: False.

forward(x)

Forward computation.

Parameters **x** (*tensor | tuple[tensor]*) – x could be a Torch.tensor or a tuple of Torch.tensor, containing input data for forward computation.

init_weights()

Initialize the weights.

make_layer(*out_channels, num_blocks, first_block=False*)

Stack ShuffleUnit blocks to make a layer.

Parameters

- **out_channels** (*int*) – out_channels of the block.
- **num_blocks** (*int*) – Number of blocks.
- **first_block** (*bool*) – Whether is the first ShuffleUnit of a sequential ShuffleUnits. Default: False, which means using the grouped 1x1 convolution.

train(*mode=True*)

Set module status before forward computation.

Parameters mode (*bool*) – Whether it is train_mode or test_mode

```
class mmcls.models.backbones.ShuffleNetV2(widen_factor=1.0, out_indices=(3), frozen_stages=-1,
                                           conv_cfg=None, norm_cfg={'type': 'BN'}, act_cfg={'type':
                                           'ReLU'}, norm_eval=False, with_cp=False, init_cfg=None)
```

ShuffleNetV2 backbone.

Parameters

- **widen_factor** (*float*) – Width multiplier - adjusts the number of channels in each layer by this amount. Default: 1.0.
- **out_indices** (*Sequence[int]*) – Output from which stages. Default: (0, 1, 2, 3).
- **frozen_stages** (*int*) – Stages to be frozen (all param fixed). Default: -1, which means not freezing any parameters.
- **conv_cfg** (*dict, optional*) – Config dict for convolution layer. Default: None, which means using conv2d.
- **norm_cfg** (*dict*) – Config dict for normalization layer. Default: dict(type='BN').
- **act_cfg** (*dict*) – Config dict for activation layer. Default: dict(type='ReLU').
- **norm_eval** (*bool*) – Whether to set norm layers to eval mode, namely, freeze running stats (mean and var). Note: Effect on Batch Norm and its variants only. Default: False.
- **with_cp** (*bool*) – Use checkpoint or not. Using checkpoint will save some memory while slowing down the training speed. Default: False.

forward(*x*)

Forward computation.

Parameters x (*tensor | tuple[tensor]*) – x could be a Torch.tensor or a tuple of Torch.tensor, containing input data for forward computation.

init_weights()

Initialize the weights.

train(*mode=True*)

Set module status before forward computation.

Parameters mode (*bool*) – Whether it is train_mode or test_mode

```
class mmcls.models.backbones.SwinTransformer(arch='T', img_size=224, in_channels=3, drop_rate=0.0,
                                             drop_path_rate=0.1, out_indices=(3),
                                             use_abs_pos_embed=False, auto_pad=False,
                                             with_cp=False, norm_cfg={'type': 'LN'}, stage_cfgs={},
                                             patch_cfg={}, init_cfg=None)
```

Swin Transformer A PyTorch implement of : [Swin Transformer: Hierarchical Vision Transformer using Shifted](#)

Windows

Inspiration from <https://github.com/microsoft/Swin-Transformer>

Parameters

- **arch** (*str* | *dict*) – Swin Transformer architecture Defaults to ‘T’.
- **img_size** (*int* | *tuple*) – The size of input image. Defaults to 224.
- **in_channels** (*int*) – The num of input channels. Defaults to 3.
- **drop_rate** (*float*) – Dropout rate after embedding. Defaults to 0.
- **drop_path_rate** (*float*) – Stochastic depth rate. Defaults to 0.1.
- **use_abs_pos_embed** (*bool*) – If True, add absolute position embedding to the patch embedding. Defaults to False.
- **with_cp** (*bool*, *optional*) – Use checkpoint or not. Using checkpoint will save some memory while slowing down the training speed. Defaults to False.
- **auto_pad** (*bool*) – If True, auto pad feature map to fit window_size. Defaults to False.
- **norm_cfg** (*dict*, *optional*) – Config dict for normalization layer at end of backone. Defaults to dict(type='LN')
- **stage_cfgs** (*Sequence* | *dict*, *optional*) – Extra config dict for each stage. Defaults to empty dict.
- **patch_cfg** (*dict*, *optional*) – Extra config dict for patch embedding. Defaults to empty dict.
- **init_cfg** (*dict*, *optional*) – The Config for initialization. Defaults to None.

Examples

```
>>> from mmcls.models import SwinTransformer
>>> import torch
>>> extra_config = dict(
>>>     arch='tiny',
>>>     stage_cfgs=dict(downsample_cfg={'kernel_size': 3,
>>>                                     'expansion_ratio': 3}),
>>>     auto_pad=True)
>>> self = SwinTransformer(**extra_config)
>>> inputs = torch.rand(1, 3, 224, 224)
>>> output = self.forward(inputs)
>>> print(output.shape)
(1, 2592, 4)
```

forward(x)

Forward computation.

Parameters x (*tensor* | *tuple[tensor]*) – x could be a Torch.tensor or a tuple of Torch.tensor, containing input data for forward computation.

init_weights()

Initialize the weights.

```
class mmcls.models.backbones.T2T_ViT(img_size=224, in_channels=3, embed_dims=384, t2t_cfg={},
                                     drop_rate=0.0, num_layers=14, out_indices=-1, layer_cfgs={},
                                     drop_path_rate=0.0, norm_cfg={'type': 'LN'}, final_norm=True,
                                     output_cls_token=True, init_cfg=None)
```

Tokens-to-Token Vision Transformer (T2T-ViT)

A PyTorch implementation of [`Tokens-to-Token ViT: Training Vision Transformers from Scratch on ImageNet`](https://arxiv.org/abs/2101.11986) <<https://arxiv.org/abs/2101.11986>>`_

Parameters

- **img_size** (*int*) – Input image size.
- **in_channels** (*int*) – Number of input channels.
- **embed_dims** (*int*) – Embedding dimension.
- **t2t_cfg** (*dict*) – Extra config of Tokens-to-Token module. Defaults to an empty dict.
- **drop_rate** (*float*) – Dropout rate after position embedding. Defaults to 0.
- **num_layers** (*int*) – Num of transformer layers in encoder. Defaults to 14.
- **out_indices** (*Sequence | int*) – Output from which stages. Defaults to -1, means the last stage.
- **layer_cfgs** (*Sequence | dict*) – Configs of each transformer layer in encoder. Defaults to an empty dict.
- **drop_path_rate** (*float*) – stochastic depth rate. Defaults to 0.
- **norm_cfg** (*dict*) – Config dict for normalization layer. Defaults to `dict(type='LN')`.
- **final_norm** (*bool*) – Whether to add a additional layer to normalize final feature map. Defaults to True.
- **output_cls_token** (*bool*) – Whether output the cls_token. Defaults to True.
- **init_cfg** (*dict, optional*) – The Config for initialization. Defaults to None.

forward(*x*)

Forward computation.

Parameters *x* (*tensor | tuple[tensor]*) – *x* could be a `Torch.tensor` or a tuple of `Torch.tensor`, containing input data for forward computation.

init_weights()

Initialize the weights.

```
class mmcls.models.backbones.TIMMBackbone(model_name, pretrained=False, checkpoint_path="",
                                           in_channels=3, init_cfg=None, **kwargs)
```

Wrapper to use backbones from timm library. More details can be found in [timm](#).

Parameters

- **model_name** (*str*) – Name of timm model to instantiate.
- **pretrained** (*bool*) – Load pretrained weights if True.
- **checkpoint_path** (*str*) – Path of checkpoint to load after model is initialized.
- **in_channels** (*int*) – Number of input image channels. Default: 3.
- **init_cfg** (*dict, optional*) – Initialization config dict
- ****kwargs** – Other timm & model specific arguments.

forward(*x*)

Forward computation.

Parameters *x* (*tensor* | *tuple*[*tensor*]) – *x* could be a `Torch.tensor` or a tuple of `Torch.tensor`, containing input data for forward computation.

```
class mmcls.models.backbones.TNT(arch='b', img_size=224, patch_size=16, in_channels=3, ffn_ratio=4,
    qkv_bias=False, drop_rate=0.0, attn_drop_rate=0.0,
    drop_path_rate=0.0, act_cfg={'type': 'GELU'}, norm_cfg={'type': 'LN'},
    first_stride=4, num_fcs=2, init_cfg=[{'type': 'TruncNormal', 'layer':
    'Linear', 'std': 0.02}, {'type': 'Constant', 'layer': 'LayerNorm', 'val': 1.0,
    'bias': 0.0}])
```

Transformer in Transformer A PyTorch implement of : [Transformer in Transformer](#)

Inspiration from <https://github.com/rwightman/pytorch-image-models/blob/master/timm/models/tnt.py>

Parameters

- **arch** (*str* | *dict*) – Vision Transformer architecture Default: 'b'
- **img_size** (*int* | *tuple*) – Input image size. Default to 224
- **patch_size** (*int* | *tuple*) – The patch size. Deault to 16
- **in_channels** (*int*) – Number of input channels. Default to 3
- **ffn_ratio** (*int*) – A ratio to calculate the hidden_dims in ffn layer. Default: 4
- **qkv_bias** (*bool*) – Enable bias for qkv if True. Default False
- **drop_rate** (*float*) – Probability of an element to be zeroed after the feed forward layer. Default 0.
- **attn_drop_rate** (*float*) – The drop out rate for attention layer. Default 0.
- **drop_path_rate** (*float*) – stochastic depth rate. Default 0.
- **act_cfg** (*dict*) – The activation config for FFNs. Defaults to GELU.
- **norm_cfg** (*dict*) – Config dict for normalization layer. Default layer normalization
- **first_stride** (*int*) – The stride of the conv2d layer. We use a conv2d layer and a unfold layer to implement image to pixel embedding.
- **num_fcs** (*int*) – The number of fully-connected layers for FFNs. Default 2
- **init_cfg** (*dict*, *optional*) – Initialization config dict

forward(*x*)

Forward computation.

Parameters *x* (*tensor* | *tuple*[*tensor*]) – *x* could be a `Torch.tensor` or a tuple of `Torch.tensor`, containing input data for forward computation.

```
class mmcls.models.backbones.VGG(depth, num_classes=- 1, num_stages=5, dilations=(1, 1, 1, 1, 1),
    out_indices=None, frozen_stages=- 1, conv_cfg=None, norm_cfg=None,
    act_cfg={'type': 'ReLU'}, norm_eval=False, ceil_mode=False,
    with_last_pool=True, init_cfg=[{'type': 'Kaiming', 'layer': ['Conv2d']},
    {'type': 'Constant', 'val': 1.0, 'layer': ['_BatchNorm']}, {'type': 'Normal',
    'std': 0.01, 'layer': ['Linear']}]])
```

VGG backbone.

Parameters

- **depth** (*int*) – Depth of vgg, from {11, 13, 16, 19}.

- **with_norm** (*bool*) – Use BatchNorm or not.
- **num_classes** (*int*) – number of classes for classification.
- **num_stages** (*int*) – VGG stages, normally 5.
- **dilations** (*Sequence[int]*) – Dilation of each stage.
- **out_indices** (*Sequence[int], optional*) – Output from which stages. When it is None, the default behavior depends on whether num_classes is specified. If num_classes <= 0, the default value is (4,), output the last feature map before classifier. If num_classes > 0, the default value is (5,), output the classification score. Default: None.
- **frozen_stages** (*int*) – Stages to be frozen (all param fixed). -1 means not freezing any parameters.
- **norm_eval** (*bool*) – Whether to set norm layers to eval mode, namely, freeze running stats (mean and var). Note: Effect on Batch Norm and its variants only. Default: False.
- **ceil_mode** (*bool*) – Whether to use ceil_mode of MaxPool. Default: False.
- **with_last_pool** (*bool*) – Whether to keep the last pooling before classifier. Default: True.

forward(*x*)

Forward computation.

Parameters *x* (*tensor | tuple[tensor]*) – *x* could be a Torch.tensor or a tuple of Torch.tensor, containing input data for forward computation.

train(*mode=True*)

Set module status before forward computation.

Parameters *mode* (*bool*) – Whether it is train_mode or test_mode

```
class mmcls.models.backbones.VisionTransformer(arch='b', img_size=224, patch_size=16, out_indices=-1, drop_rate=0.0, drop_path_rate=0.0, norm_cfg={'eps': 1e-06, 'type': 'LN'}, final_norm=True, output_cls_token=True, interpolate_mode='bicubic', patch_cfg={}, layer_cfgs={}, init_cfg=None)
```

Vision Transformer.

A PyTorch implement of : **[`An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale<https://arxiv.org/abs/2010.11929>`](https://arxiv.org/abs/2010.11929)**

Parameters

- **arch** (*str | dict*) – Vision Transformer architecture Default: 'b'
- **img_size** (*int | tuple*) – Input image size
- **patch_size** (*int | tuple*) – The patch size
- **out_indices** (*Sequence | int*) – Output from which stages. Defaults to -1, means the last stage.
- **drop_rate** (*float*) – Probability of an element to be zeroed. Defaults to 0.
- **drop_path_rate** (*float*) – stochastic depth rate. Defaults to 0.
- **norm_cfg** (*dict*) – Config dict for normalization layer. Defaults to dict(type='LN').
- **final_norm** (*bool*) – Whether to add a additional layer to normalize final feature map. Defaults to True.

- **output_cls_token** (*bool*) – Whether output the cls_token. If set True, *with_cls_token* must be True. Defaults to True.
- **interpolate_mode** (*str*) – Select the interpolate mode for position embedding vector resize. Defaults to “bicubic”.
- **patch_cfg** (*dict*) – Configs of patch embedding. Defaults to an empty dict.
- **layer_cfgs** (*Sequence | dict*) – Configs of each transformer layer in encoder. Defaults to an empty dict.
- **init_cfg** (*dict, optional*) – Initialization config dict. Defaults to None.

forward(*x*)

Forward computation.

Parameters *x* (*tensor | tuple[tensor]*) – *x* could be a `Torch.tensor` or a tuple of `Torch.tensor`, containing input data for forward computation.

init_weights()

Initialize the weights.

static **resize_pos_embed**(*pos_embed, src_shape, dst_shape, mode='bicubic'*)

Resize `pos_embed` weights.

Parameters

- **pos_embed** (*torch.Tensor*) – Position embedding weights with shape [1, L, C].
- **src_shape** (*tuple*) – The resolution of downsampled origin training image.
- **dst_shape** (*tuple*) – The resolution of downsampled new training image.
- **mode** (*str*) – Algorithm used for upsampling: 'nearest' | 'linear' | 'bilinear' | 'bicubic' | 'trilinear'. Default: 'bicubic'

Returns The resized `pos_embed` of shape [1, L_new, C]

Return type `torch.Tensor`

36.4 necks

class `mmcls.models.necks.GlobalAveragePooling`(*dim=2*)

Global Average Pooling neck.

Note that we use *view* to remove extra channel after pooling. We do not use *squeeze* as it will also remove the batch dimension when the tensor has a batch dimension of size 1, which can lead to unexpected errors.

Parameters **dim** (*int*) – Dimensions of each sample channel, can be one of {1, 2, 3}. Default: 2

forward(*inputs*)

Defines the computation performed at every call.

Should be overridden by all subclasses.

Note: Although the recipe for forward pass needs to be defined within this function, one should call the `Module` instance afterwards instead of this since the former takes care of running the registered hooks while the latter silently ignores them.

36.5 heads

```
class mmcls.models.heads.ClsHead(loss={'loss_weight': 1.0, 'type': 'CrossEntropyLoss'}, topk=(1),
                                cal_acc=False, init_cfg=None)
```

classification head.

Parameters

- **loss** (*dict*) – Config of classification loss.
- **topk** (*int* | *tuple*) – Top-k accuracy.
- **cal_acc** (*bool*) – Whether to calculate accuracy during training. If you use Mixup/CutMix or something like that during training, it is not reasonable to calculate accuracy. Defaults to False.

```
simple_test(cls_score)
```

Test without augmentation.

```
class mmcls.models.heads.LinearClsHead(num_classes, in_channels, init_cfg={'layer': 'Linear', 'std': 0.01,
                                'type': 'Normal'}, *args, **kwargs)
```

Linear classifier head.

Parameters

- **num_classes** (*int*) – Number of categories excluding the background category.
- **in_channels** (*int*) – Number of channels in the input feature map.
- **init_cfg** (*dict* | *optional*) – The extra init config of layers. Defaults to use dict(type='Normal', layer='Linear', std=0.01).

```
simple_test(x)
```

Test without augmentation.

```
class mmcls.models.heads.MultiLabelClsHead(loss={'loss_weight': 1.0, 'reduction': 'mean', 'type':
                                'CrossEntropyLoss', 'use_sigmoid': True}, init_cfg=None)
```

Classification head for multilabel task.

Parameters **loss** (*dict*) – Config of classification loss.

```
class mmcls.models.heads.MultiLabelLinearClsHead(num_classes, in_channels, loss={'loss_weight': 1.0,
                                'reduction': 'mean', 'type': 'CrossEntropyLoss',
                                'use_sigmoid': True}, init_cfg={'layer': 'Linear',
                                'std': 0.01, 'type': 'Normal'})
```

Linear classification head for multilabel task.

Parameters

- **num_classes** (*int*) – Number of categories.
- **in_channels** (*int*) – Number of channels in the input feature map.
- **loss** (*dict*) – Config of classification loss.
- **init_cfg** (*dict* | *optional*) – The extra init config of layers. Defaults to use dict(type='Normal', layer='Linear', std=0.01).

```
simple_test(x)
```

Test without augmentation.

```
class mmcls.models.heads.StackedLinearClsHead(num_classes: int, in_channels: int, mid_channels:
                                             Sequence, dropout_rate: float = 0.0, norm_cfg:
                                             Optional[Dict] = None, act_cfg: Dict = {'type': 'ReLU'},
                                             **kwargs)
```

Classifier head with several hidden fc layer and a output fc layer.

Parameters

- **num_classes** (*int*) – Number of categories excluding the background category.
- **in_channels** (*int*) – Number of channels in the input feature map.
- **mid_channels** (*Sequence*) – Number of channels in the hidden fc layers.
- **dropout_rate** (*float*) – Dropout rate after each hidden fc layer, except the last layer. Defaults to 0.
- **norm_cfg** (*dict, optional*) – Config dict of normalization layer after each hidden fc layer, except the last layer. Defaults to None.
- **act_cfg** (*dict, optional*) – Config dict of activation function after each hidden layer, except the last layer. Defaults to use “ReLU”.

init_weights()

Initialize the weights.

simple_test(x)

Test without augmentation.

```
class mmcls.models.heads.VisionTransformerClsHead(num_classes, in_channels, hidden_dim=None,
                                                  act_cfg={'type': 'Tanh'}, init_cfg={'layer': 'Linear',
                                                  'type': 'Constant', 'val': 0}, *args, **kwargs)
```

Vision Transformer classifier head.

Parameters

- **num_classes** (*int*) – Number of categories excluding the background category.
- **in_channels** (*int*) – Number of channels in the input feature map.
- **hidden_dim** (*int*) – Number of the dimensions for hidden layer. Only available during pre-training. Default None.
- **act_cfg** (*dict*) – The activation config. Only available during pre-training. Defaults to Tanh.

init_weights()

Initialize the weights.

simple_test(x)

Test without augmentation.

36.6 losses

class `mmcls.models.losses.Accuracy`(*topk=(1)*)

forward(*pred, target*)

Forward function to calculate accuracy.

Parameters

- **pred** (*torch.Tensor*) – Prediction of models.
- **target** (*torch.Tensor*) – Target for each prediction.

Returns The accuracies under different topk criterions.

Return type `list[float]`

class `mmcls.models.losses.AsymmetricLoss`(*gamma_pos=0.0, gamma_neg=4.0, clip=0.05, reduction='mean', loss_weight=1.0*)

asymmetric loss.

Parameters

- **gamma_pos** (*float*) – positive focusing parameter. Defaults to 0.0.
- **gamma_neg** (*float*) – Negative focusing parameter. We usually set `gamma_neg > gamma_pos`. Defaults to 4.0.
- **clip** (*float, optional*) – Probability margin. Defaults to 0.05.
- **reduction** (*str*) – The method used to reduce the loss into a scalar.
- **loss_weight** (*float*) – Weight of loss. Defaults to 1.0.

forward(*pred, target, weight=None, avg_factor=None, reduction_override=None*)

asymmetric loss.

class `mmcls.models.losses.CrossEntropyLoss`(*use_sigmoid=False, use_soft=False, reduction='mean', loss_weight=1.0, class_weight=None, pos_weight=None*)

Cross entropy loss.

Parameters

- **use_sigmoid** (*bool*) – Whether the prediction uses sigmoid of softmax. Defaults to False.
- **use_soft** (*bool*) – Whether to use the soft version of CrossEntropyLoss. Defaults to False.
- **reduction** (*str*) – The method used to reduce the loss. Options are “none”, “mean” and “sum”. Defaults to ‘mean’.
- **loss_weight** (*float*) – Weight of the loss. Defaults to 1.0.
- **class_weight** (*List[float], optional*) – The weight for each class with shape (C), C is the number of classes. Default None.
- **pos_weight** (*List[float], optional*) – The positive weight for each class with shape (C), C is the number of classes. Only enabled in BCE loss when `use_sigmoid` is True. Default None.

forward(*cls_score, label, weight=None, avg_factor=None, reduction_override=None, **kwargs*)

Defines the computation performed at every call.

Should be overridden by all subclasses.

Note: Although the recipe for forward pass needs to be defined within this function, one should call the `Module` instance afterwards instead of this since the former takes care of running the registered hooks while the latter silently ignores them.

`class mmcls.models.losses.FocalLoss(gamma=2.0, alpha=0.25, reduction='mean', loss_weight=1.0)`
Focal loss.

Parameters

- **gamma** (*float*) – Focusing parameter in focal loss. Defaults to 2.0.
- **alpha** (*float*) – The parameter in balanced form of focal loss. Defaults to 0.25.
- **reduction** (*str*) – The method used to reduce the loss into a scalar. Options are “none” and “mean”. Defaults to ‘mean’.
- **loss_weight** (*float*) – Weight of loss. Defaults to 1.0.

`forward(pred, target, weight=None, avg_factor=None, reduction_override=None)`
Sigmoid focal loss.

Parameters

- **pred** (*torch.Tensor*) – The prediction with shape (N, *).
- **target** (*torch.Tensor*) – The ground truth label of the prediction with shape (N, *), N or (N,1).
- **weight** (*torch.Tensor, optional*) – Sample-wise loss weight with shape (N, *). Defaults to None.
- **avg_factor** (*int, optional*) – Average factor that is used to average the loss. Defaults to None.
- **reduction_override** (*str, optional*) – The method used to reduce the loss into a scalar. Options are “none”, “mean” and “sum”. Defaults to None.

Returns Loss.

Return type torch.Tensor

`class mmcls.models.losses.LabelSmoothLoss(label_smooth_val, num_classes=None, mode=None, reduction='mean', loss_weight=1.0)`

Initializer for the label smoothed cross entropy loss.

Refers to [Rethinking the Inception Architecture for Computer Vision](#)

This decreases gap between output scores and encourages generalization. Labels provided to forward can be one-hot like vectors (NxC) or class indices (Nx1). And this accepts linear combination of one-hot like labels from mixup or cutmix except multi-label task.

Parameters

- **label_smooth_val** (*float*) – The degree of label smoothing.
- **num_classes** (*int, optional*) – Number of classes. Defaults to None.
- **mode** (*str*) – Refers to notes, Options are ‘original’, ‘classy_vision’, ‘multi_label’. Defaults to ‘classy_vision’
- **reduction** (*str*) – The method used to reduce the loss. Options are “none”, “mean” and “sum”. Defaults to ‘mean’.
- **loss_weight** (*float*) – Weight of the loss. Defaults to 1.0.

Notes

if the mode is “original”, this will use the same label smooth method as the original paper as:

$$(1 - \epsilon)\delta_{k,y} + \frac{\epsilon}{K}$$

where epsilon is the `label_smooth_val`, `K` is the `num_classes` and `delta(k,y)` is Dirac delta, which equals 1 for `k=y` and 0 otherwise.

if the mode is “classy_vision”, this will use the same label smooth method as the facebookresearch/ClassyVision repo as:

$$\frac{\delta_{k,y} + \epsilon/K}{1 + \epsilon}$$

if the mode is “multi_label”, this will accept labels from multi-label task and smoothing them as:

$$(1 - 2\epsilon)\delta_{k,y} + \epsilon$$

forward(*cls_score*, *label*, *weight=None*, *avg_factor=None*, *reduction_override=None*, ***kwargs*)
Label smooth loss.

Parameters

- **pred** (*torch.Tensor*) – The prediction with shape (N, *).
- **label** (*torch.Tensor*) – The ground truth label of the prediction with shape (N, *).
- **weight** (*torch.Tensor*, *optional*) – Sample-wise loss weight with shape (N, *). Defaults to None.
- **avg_factor** (*int*, *optional*) – Average factor that is used to average the loss. Defaults to None.
- **reduction_override** (*str*, *optional*) – The method used to reduce the loss into a scalar. Options are “none”, “mean” and “sum”. Defaults to None.

Returns Loss.

Return type torch.Tensor

generate_one_hot_like_label(*label*)

This function takes one-hot or index label vectors and computes one-hot like label vectors (float)

class `mmcls.models.losses.SeesawLoss`(*use_sigmoid=False*, *p=0.8*, *q=2.0*, *num_classes=1000*, *eps=0.01*, *reduction='mean'*, *loss_weight=1.0*)

Implementation of seesaw loss.

Refers to [Seesaw Loss for Long-Tailed Instance Segmentation \(CVPR 2021\)](#)

Parameters

- **use_sigmoid** (*bool*) – Whether the prediction uses sigmoid of softmax. Only False is supported. Defaults to False.
- **p** (*float*) – The p in the mitigation factor. Defaults to 0.8.
- **q** (*float*) – The q in the compensation factor. Defaults to 2.0.
- **num_classes** (*int*) – The number of classes. Default to 1000 for the ImageNet dataset.
- **eps** (*float*) – The minimal value of divisor to smooth the computation of compensation factor, default to 1e-2.

- **reduction** (*str*) – The method that reduces the loss to a scalar. Options are “none”, “mean” and “sum”. Default to “mean”.
- **loss_weight** (*float*) – The weight of the loss. Defaults to 1.0

forward(*cls_score, labels, weight=None, avg_factor=None, reduction_override=None*)
Forward function.

Parameters

- **cls_score** (*torch.Tensor*) – The prediction with shape (N, C).
- **labels** (*torch.Tensor*) – The learning label of the prediction.
- **weight** (*torch.Tensor, optional*) – Sample-wise loss weight.
- **avg_factor** (*int, optional*) – Average factor that is used to average the loss. Defaults to None.
- **reduction** (*str, optional*) – The method used to reduce the loss. Options are “none”, “mean” and “sum”.

Returns The calculated loss

Return type torch.Tensor

`mmcls.models.losses.accuracy(pred, target, topk=1, thrs=0.0)`
Calculate accuracy according to the prediction and target.

Parameters

- **pred** (*torch.Tensor | np.array*) – The model prediction.
- **target** (*torch.Tensor | np.array*) – The target of each prediction
- **topk** (*int | tuple[int]*) – If the predictions in topk matches the target, the predictions will be regarded as correct ones. Defaults to 1.
- **thrs** (*Number | tuple[Number], optional*) – Predictions with scores under the thresholds are considered negative. Default to 0.

Returns

Accuracy

- float: If both **topk** and **thrs** is a single value.
- list[float]: If one of **topk** or **thrs** is a tuple.
- list[list[float]]: If both **topk** and **thrs** is a tuple. And the first dim is topk, the second dim is thrs.

Return type float | list[float] | list[list[float]]

`mmcls.models.losses.asymmetric_loss(pred, target, weight=None, gamma_pos=1.0, gamma_neg=4.0, clip=0.05, reduction='mean', avg_factor=None)`

asymmetric loss.

Please refer to the [paper](#) for details.

Parameters

- **pred** (*torch.Tensor*) – The prediction with shape (N, *).
- **target** (*torch.Tensor*) – The ground truth label of the prediction with shape (N, *).
- **weight** (*torch.Tensor, optional*) – Sample-wise loss weight with shape (N,). Defaults to None.

- **gamma_pos** (*float*) – positive focusing parameter. Defaults to 0.0.
- **gamma_neg** (*float*) – Negative focusing parameter. We usually set `gamma_neg > gamma_pos`. Defaults to 4.0.
- **clip** (*float, optional*) – Probability margin. Defaults to 0.05.
- **reduction** (*str*) – The method used to reduce the loss. Options are “none”, “mean” and “sum”. If reduction is ‘none’, loss is same shape as pred and label. Defaults to ‘mean’.
- **avg_factor** (*int, optional*) – Average factor that is used to average the loss. Defaults to None.

Returns Loss.

Return type torch.Tensor

```
mmcls.models.losses.binary_cross_entropy(pred, label, weight=None, reduction='mean',
                                          avg_factor=None, class_weight=None, pos_weight=None)
```

Calculate the binary CrossEntropy loss with logits.

Parameters

- **pred** (*torch.Tensor*) – The prediction with shape (N, *).
- **label** (*torch.Tensor*) – The gt label with shape (N, *).
- **weight** (*torch.Tensor, optional*) – Element-wise weight of loss with shape (N,). Defaults to None.
- **reduction** (*str*) – The method used to reduce the loss. Options are “none”, “mean” and “sum”. If reduction is ‘none’, loss is same shape as pred and label. Defaults to ‘mean’.
- **avg_factor** (*int, optional*) – Average factor that is used to average the loss. Defaults to None.
- **class_weight** (*torch.Tensor, optional*) – The weight for each class with shape (C), C is the number of classes. Default None.
- **pos_weight** (*torch.Tensor, optional*) – The positive weight for each class with shape (C), C is the number of classes. Default None.

Returns The calculated loss

Return type torch.Tensor

```
mmcls.models.losses.convert_to_one_hot(targets: torch.Tensor, classes) → torch.Tensor
```

This function converts target class indices to one-hot vectors, given the number of classes.

Parameters

- **targets** (*Tensor*) – The ground truth label of the prediction with shape (N, 1)
- **classes** (*int*) – the number of classes.

Returns Processed loss values.

Return type Tensor

```
mmcls.models.losses.cross_entropy(pred, label, weight=None, reduction='mean', avg_factor=None,
                                   class_weight=None)
```

Calculate the CrossEntropy loss.

Parameters

- **pred** (*torch.Tensor*) – The prediction with shape (N, C), C is the number of classes.

- **label** (*torch.Tensor*) – The gt label of the prediction.
- **weight** (*torch.Tensor, optional*) – Sample-wise loss weight.
- **reduction** (*str*) – The method used to reduce the loss.
- **avg_factor** (*int, optional*) – Average factor that is used to average the loss. Defaults to None.
- **class_weight** (*torch.Tensor, optional*) – The weight for each class with shape (C), C is the number of classes. Default None.

Returns The calculated loss

Return type torch.Tensor

`mmcls.models.losses.reduce_loss(loss, reduction)`

Reduce loss as specified.

Parameters

- **loss** (*Tensor*) – Elementwise loss tensor.
- **reduction** (*str*) – Options are “none”, “mean” and “sum”.

Returns Reduced loss tensor.

Return type Tensor

`mmcls.models.losses.sigmoid_focal_loss(pred, target, weight=None, gamma=2.0, alpha=0.25, reduction='mean', avg_factor=None)`

Sigmoid focal loss.

Parameters

- **pred** (*torch.Tensor*) – The prediction with shape (N, *).
- **target** (*torch.Tensor*) – The ground truth label of the prediction with shape (N, *).
- **weight** (*torch.Tensor, optional*) – Sample-wise loss weight with shape (N,). Defaults to None.
- **gamma** (*float*) – The gamma for calculating the modulating factor. Defaults to 2.0.
- **alpha** (*float*) – A balanced form for Focal Loss. Defaults to 0.25.
- **reduction** (*str*) – The method used to reduce the loss. Options are “none”, “mean” and “sum”. If reduction is ‘none’, loss is same shape as pred and label. Defaults to ‘mean’.
- **avg_factor** (*int, optional*) – Average factor that is used to average the loss. Defaults to None.

Returns Loss.

Return type torch.Tensor

`mmcls.models.losses.weight_reduce_loss(loss, weight=None, reduction='mean', avg_factor=None)`

Apply element-wise weight and reduce loss.

Parameters

- **loss** (*Tensor*) – Element-wise loss.
- **weight** (*Tensor*) – Element-wise weights.
- **reduction** (*str*) – Same as built-in losses of PyTorch.
- **avg_factor** (*float*) – Average factor when computing the mean of losses.

Returns Processed loss values.

Return type Tensor

`mmcls.models.losses.weighted_loss(loss_func)`
Create a weighted version of a given loss function.

To use this decorator, the loss function must have the signature like `loss_func(pred, target, **kwargs)`. The function only needs to compute element-wise loss without any reduction. This decorator will add weight and reduction arguments to the function. The decorated function will have the signature like `loss_func(pred, target, weight=None, reduction='mean', avg_factor=None, **kwargs)`.

Example

```
>>> import torch
>>> @weighted_loss
>>> def l1_loss(pred, target):
>>>     return (pred - target).abs()
```

```
>>> pred = torch.Tensor([0, 2, 3])
>>> target = torch.Tensor([1, 1, 1])
>>> weight = torch.Tensor([1, 0, 1])
```

```
>>> l1_loss(pred, target)
tensor(1.3333)
>>> l1_loss(pred, target, weight)
tensor(1.)
>>> l1_loss(pred, target, reduction='none')
tensor([1., 1., 2.])
>>> l1_loss(pred, target, weight, avg_factor=2)
tensor(1.5000)
```

36.7 utils

`class mmcls.models.utils.Augments(augments_cfg)`

Data augments.

We implement some data augmentation methods, such as mixup, cutmix.

Parameters (`list[mmcv.ConfigDict]` | `obj` (`augments_cfg`) – `mmcv.ConfigDict`): Config dict of augments

Example

```
>>> augments_cfg = [
    dict(type='BatchCutMix', alpha=1., num_classes=10, prob=0.5),
    dict(type='BatchMixup', alpha=1., num_classes=10, prob=0.3)
]
>>> augments = Augments(augments_cfg)
>>> imgs = torch.randn(16, 3, 32, 32)
>>> label = torch.randint(0, 10, (16, ))
>>> imgs, label = augments(imgs, label)
```

To decide which augmentation within Augments block is used the following rule is applied. We pick augmentation based on the probabilities. In the example above, we decide if we should use BatchCutMix with probability 0.5, BatchMixup 0.3. As Identity is not in augments_cfg, we use Identity with probability $1 - 0.5 - 0.3 = 0.2$.

```
class mmcls.models.utils.HybridEmbed(backbone, img_size=224, feature_size=None, in_channels=3,
                                     embed_dims=768, conv_cfg=None, init_cfg=None)
```

CNN Feature Map Embedding.

Extract feature map from CNN, flatten, project to embedding dim.

Parameters

- **backbone** (*nn.Module*) – CNN backbone
- **img_size** (*int* | *tuple*) – The size of input image. Default: 224
- **feature_size** (*int* | *tuple*, *optional*) – Size of feature map extracted by CNN backbone. Default: None
- **in_channels** (*int*) – The num of input channels. Default: 3
- **embed_dims** (*int*) – The dimensions of embedding. Default: 768
- **conv_cfg** (*dict*, *optional*) – The config dict for conv layers. Default: None.
- **init_cfg** (*mmcv.ConfigDict*, *optional*) – The Config for initialization. Default: None.

forward(*x*)

Defines the computation performed at every call.

Should be overridden by all subclasses.

Note: Although the recipe for forward pass needs to be defined within this function, one should call the `Module` instance afterwards instead of this since the former takes care of running the registered hooks while the latter silently ignores them.

```
class mmcls.models.utils.InvertedResidual(in_channels, out_channels, mid_channels, kernel_size=3,
                                           stride=1, se_cfg=None, conv_cfg=None, norm_cfg={'type':
                                           'BN'}, act_cfg={'type': 'ReLU'}, with_cp=False,
                                           init_cfg=None)
```

Inverted Residual Block.

Parameters

- **in_channels** (*int*) – The input channels of this Module.
- **out_channels** (*int*) – The output channels of this Module.
- **mid_channels** (*int*) – The input channels of the depthwise convolution.
- **kernel_size** (*int*) – The kernel size of the depthwise convolution. Default: 3.
- **stride** (*int*) – The stride of the depthwise convolution. Default: 1.
- **se_cfg** (*dict*) – Config dict for se layer. Default: None, which means no se layer.
- **conv_cfg** (*dict*) – Config dict for convolution layer. Default: None, which means using conv2d.
- **norm_cfg** (*dict*) – Config dict for normalization layer. Default: dict(type='BN').
- **act_cfg** (*dict*) – Config dict for activation layer. Default: dict(type='ReLU').

- **with_cp** (*bool*) – Use checkpoint or not. Using checkpoint will save some memory while slowing down the training speed. Default: False.

Returns The output tensor.

Return type Tensor

forward(*x*)

Defines the computation performed at every call.

Should be overridden by all subclasses.

Note: Although the recipe for forward pass needs to be defined within this function, one should call the Module instance afterwards instead of this since the former takes care of running the registered hooks while the latter silently ignores them.

```
class mmcls.models.utils.MultiheadAttention(embed_dims, num_heads, input_dims=None,
                                           attn_drop=0.0, proj_drop=0.0,
                                           dropout_layer={'drop_prob': 0.0, 'type': 'Dropout'},
                                           qkv_bias=True, qk_scale=None, proj_bias=True,
                                           v_shortcut=False, init_cfg=None)
```

Multi-head Attention Module.

This module implements multi-head attention that supports different input dims and embed dims. And it also supports a shortcut from value, which is useful if input dims is not the same with embed dims.

Parameters

- **embed_dims** (*int*) – The embedding dimension.
- **num_heads** (*int*) – Parallel attention heads.
- **input_dims** (*int, optional*) – The input dimension, and if None, use `embed_dims`. Defaults to None.
- **attn_drop** (*float*) – Dropout rate of the dropout layer after the attention calculation of query and key. Defaults to 0.
- **proj_drop** (*float*) – Dropout rate of the dropout layer after the output projection. Defaults to 0.
- **dropout_layer** (*dict*) – The dropout config before adding the shortcut. Defaults to `dict(type='Dropout', drop_prob=0.)`.
- **qkv_bias** (*bool*) – If True, add a learnable bias to q, k, v. Defaults to True.
- **qk_scale** (*float, optional*) – Override default qk scale of `head_dim ** -0.5` if set. Defaults to None.
- **proj_bias** (*bool*) – Defaults to True.
- **v_shortcut** (*bool*) – Add a shortcut from value to output. It's usually used if `input_dims` is different from `embed_dims`. Defaults to False.
- **init_cfg** (*dict, optional*) – The Config for initialization. Defaults to None.

forward(*x*)

Defines the computation performed at every call.

Should be overridden by all subclasses.

Note: Although the recipe for forward pass needs to be defined within this function, one should call the `Module` instance afterwards instead of this since the former takes care of running the registered hooks while the latter silently ignores them.

```
class mmcls.models.utils.PatchEmbed(img_size=224, in_channels=3, embed_dims=768, norm_cfg=None,  
conv_cfg=None, init_cfg=None)
```

Image to Patch Embedding.

We use a conv layer to implement PatchEmbed.

Parameters

- **img_size** (*int* | *tuple*) – The size of input image. Default: 224
- **in_channels** (*int*) – The num of input channels. Default: 3
- **embed_dims** (*int*) – The dimensions of embedding. Default: 768
- **norm_cfg** (*dict*, *optional*) – Config dict for normalization layer. Default: None
- **conv_cfg** (*dict*, *optional*) – The config dict for conv layers. Default: None
- **init_cfg** (*mmcv.ConfigDict*, *optional*) – The Config for initialization. Default: None

forward(*x*)

Defines the computation performed at every call.

Should be overridden by all subclasses.

Note: Although the recipe for forward pass needs to be defined within this function, one should call the `Module` instance afterwards instead of this since the former takes care of running the registered hooks while the latter silently ignores them.

```
class mmcls.models.utils.PatchMerging(input_resolution, in_channels, expansion_ratio, kernel_size=2,  
stride=None, padding=0, dilation=1, bias=False,  
norm_cfg={'type': 'LN'}, init_cfg=None)
```

Merge patch feature map.

This layer use nn.Unfold to group feature map by kernel_size, and use norm and linear layer to embed grouped feature map.

Parameters

- **input_resolution** (*tuple*) – The size of input patch resolution.
- **in_channels** (*int*) – The num of input channels.
- **expansion_ratio** (*Number*) – Expansion ratio of output channels. The num of output channels is equal to $\text{int}(\text{expansion_ratio} * \text{in_channels})$.
- **kernel_size** (*int* | *tuple*, *optional*) – the kernel size in the unfold layer. Defaults to 2.
- **stride** (*int* | *tuple*, *optional*) – the stride of the sliding blocks in the unfold layer. Defaults to be equal with kernel_size.
- **padding** (*int* | *tuple*, *optional*) – zero padding width in the unfold layer. Defaults to 0.
- **dilation** (*int* | *tuple*, *optional*) – dilation parameter in the unfold layer. Defaults to 1.

- **bias** (*bool, optional*) – Whether to add bias in linear layer or not. Defaults to False.
- **norm_cfg** (*dict, optional*) – Config dict for normalization layer. Defaults to dict(type='LN').
- **init_cfg** (*dict, optional*) – The extra config for initialization. Defaults to None.

forward(*x*)

x: B, H*W, C

```
class mmcls.models.utils.SELayer(channels, squeeze_channels=None, ratio=16, divisor=8, bias='auto',
                                conv_cfg=None, act_cfg=({'type': 'ReLU'}, {'type': 'Sigmoid'}),
                                init_cfg=None)
```

Squeeze-and-Excitation Module.

Parameters

- **channels** (*int*) – The input (and output) channels of the SE layer.
- **squeeze_channels** (*None or int*) – The intermediate channel number of SELayer. Default: None, means the value of `squeeze_channels` is `make_divisible(channels // ratio, divisor)`.
- **ratio** (*int*) – Squeeze ratio in SELayer, the intermediate channel will be `make_divisible(channels // ratio, divisor)`. Only used when `squeeze_channels` is None. Default: 16.
- **divisor** (*int*) – The divisor to true divide the channel number. Only used when `squeeze_channels` is None. Default: 8.
- **conv_cfg** (*None or dict*) – Config dict for convolution layer. Default: None, which means using `conv2d`.
- **act_cfg** (*dict or Sequence[dict]*) – Config dict for activation layer. If `act_cfg` is a dict, two activation layers will be configured by this dict. If `act_cfg` is a sequence of dicts, the first activation layer will be configured by the first dict and the second activation layer will be configured by the second dict. Default: (dict(type='ReLU'), dict(type='Sigmoid'))

forward(*x*)

Defines the computation performed at every call.

Should be overridden by all subclasses.

Note: Although the recipe for forward pass needs to be defined within this function, one should call the `Module` instance afterwards instead of this since the former takes care of running the registered hooks while the latter silently ignores them.

```
class mmcls.models.utils.ShiftWindowMSA(embed_dims, input_resolution, num_heads, window_size,
                                         shift_size=0, qkv_bias=True, qk_scale=None, attn_drop=0,
                                         proj_drop=0, dropout_layer=({'drop_prob': 0.0, 'type':
                                         'DropPath'}, auto_pad=False, init_cfg=None)
```

Shift Window Multihead Self-Attention Module.

Parameters

- **embed_dims** (*int*) – Number of input channels.
- **input_resolution** (*Tuple[int, int]*) – The resolution of the input feature map.
- **num_heads** (*int*) – Number of attention heads.
- **window_size** (*int*) – The height and width of the window.

- **shift_size** (*int*, *optional*) – The shift step of each window towards right-bottom. If zero, act as regular window-msa. Defaults to 0.
- **qkv_bias** (*bool*, *optional*) – If True, add a learnable bias to q, k, v. Default: True
- **qk_scale** (*float* | *None*, *optional*) – Override default qk scale of head_dim ** -0.5 if set. Defaults to None.
- **attn_drop** (*float*, *optional*) – Dropout ratio of attention weight. Defaults to 0.0.
- **proj_drop** (*float*, *optional*) – Dropout ratio of output. Defaults to 0.
- **dropout_layer** (*dict*, *optional*) – The dropout_layer used before output. Defaults to dict(type='DropPath', drop_prob=0.).
- **auto_pad** (*bool*, *optional*) – Auto pad the feature map to be divisible by window_size, Defaults to False.
- **init_cfg** (*dict*, *optional*) – The extra config for initialization. Default: None.

forward(*query*)

Defines the computation performed at every call.

Should be overridden by all subclasses.

Note: Although the recipe for forward pass needs to be defined within this function, one should call the Module instance afterwards instead of this since the former takes care of running the registered hooks while the latter silently ignores them.

`mmcls.models.utils.channel_shuffle(x, groups)`

Channel Shuffle operation.

This function enables cross-group information flow for multiple groups convolution layers.

Parameters

- **x** (*Tensor*) – The input tensor.
- **groups** (*int*) – The number of groups to divide the input tensor in the channel dimension.

Returns The output tensor after channel shuffle operation.

Return type Tensor

`mmcls.models.utils.make_divisible(value, divisor, min_value=None, min_ratio=0.9)`

Make divisible function.

This function rounds the channel number down to the nearest value that can be divisible by the divisor.

Parameters

- **value** (*int*) – The original channel number.
- **divisor** (*int*) – The divisor to fully divide the channel number.
- **min_value** (*int*, *optional*) – The minimum value of the output channel. Default: None, means that the minimum value equal to the divisor.
- **min_ratio** (*float*) – The minimum ratio of the rounded channel number to the original channel number. Default: 0.9.

Returns The modified output channel number

Return type int

MMCLS.DATASETS

37.1 datasets

class `mmcls.datasets.BaseDataset`(*data_prefix*, *pipeline*, *classes=None*, *ann_file=None*, *test_mode=False*)
Base dataset.

Parameters

- **data_prefix** (*str*) – the prefix of data path
- **pipeline** (*list*) – a list of dict, where each element represents a operation defined in `mmcls.datasets.pipelines`
- **ann_file** (*str* | *None*) – the annotation file. When `ann_file` is `str`, the subclass is expected to read from the `ann_file`. When `ann_file` is `None`, the subclass is expected to read according to `data_prefix`
- **test_mode** (*bool*) – in train mode or test mode

property `class_to_idx`

Map mapping class name to class index.

Returns mapping from class name to class index.

Return type dict

evaluate(*results*, *metric='accuracy'*, *metric_options=None*, *logger=None*)

Evaluate the dataset.

Parameters

- **results** (*list*) – Testing results of the dataset.
- **metric** (*str* | *list[str]*) – Metrics to be evaluated. Default value is `accuracy`.
- **metric_options** (*dict*, *optional*) – Options for calculating metrics. Allowed keys are `'topk'`, `'thrs'` and `'average_mode'`. Defaults to `None`.
- **logger** (*logging.Logger* | *str*, *optional*) – Logger used for printing related information during evaluation. Defaults to `None`.

Returns evaluation results

Return type dict

get_cat_ids(*idx: int*) → List[int]

Get category id by index.

Parameters `idx` (*int*) – Index of data.

Returns Image category of specified index.

Return type `cat_ids` (`List[int]`)

classmethod `get_classes` (`classes=None`)

Get class names of current dataset.

Parameters `classes` (`Sequence[str] | str | None`) – If `classes` is `None`, use default `CLASSES` defined by builtin dataset. If `classes` is a string, take it as a file name. The file contains the name of classes where each line contains one class name. If `classes` is a tuple or list, override the `CLASSES` defined by the dataset.

Returns Names of categories of the dataset.

Return type `tuple[str]` or `list[str]`

get_gt_labels()

Get all ground-truth labels (categories).

Returns categories for all images.

Return type `list[int]`

class `mmcls.datasets.CIFAR10` (`data_prefix, pipeline, classes=None, ann_file=None, test_mode=False`)

`CIFAR10` Dataset.

This implementation is modified from <https://github.com/pytorch/vision/blob/master/torchvision/datasets/cifar.py>

class `mmcls.datasets.CIFAR100` (`data_prefix, pipeline, classes=None, ann_file=None, test_mode=False`)

`CIFAR100` Dataset.

class `mmcls.datasets.ClassBalancedDataset` (`dataset, oversample_thr`)

A wrapper of repeated dataset with repeat factor.

Suitable for training on class imbalanced datasets like LVIS. Following the sampling strategy in², in each epoch, an image may appear multiple times based on its “repeat factor”.

The repeat factor for an image is a function of the frequency the rarest category labeled in that image. The “frequency of category *c*” in $[0, 1]$ is defined by the fraction of images in the training set (without repeats) in which category *c* appears.

The dataset needs to implement `self.get_cat_ids()` to support `ClassBalancedDataset`.

The repeat factor is computed as followed.

1. For each category *c*, compute the fraction $f(c)$ of images that contain it.
2. For each category *c*, compute the category-level repeat factor

$$r(c) = \max\left(1, \sqrt{\frac{t}{f(c)}}\right)$$

3. For each image *I* and its labels $L(I)$, compute the image-level repeat factor

$$r(I) = \max_{c \in L(I)} r(c)$$

² <https://arxiv.org/pdf/1908.03195.pdf>

References

Parameters

- **dataset** (`CustomDataset`) – The dataset to be repeated.
- **oversample_thr** (`float`) – frequency threshold below which data is repeated. For categories with $f_c \geq \text{oversample_thr}$, there is no oversampling. For categories with $f_c < \text{oversample_thr}$, the degree of oversampling following the square-root inverse frequency heuristic above.

class `mmcls.datasets.ConcatDataset(datasets)`

A wrapper of concatenated dataset.

Same as `torch.utils.data.dataset.ConcatDataset`, but add `get_cat_ids` function.

Parameters `datasets` (`list[Dataset]`) – A list of datasets.

class `mmcls.datasets.DistributedSampler(dataset, num_replicas=None, rank=None, shuffle=True, round_up=True)`

class `mmcls.datasets.FashionMNIST(data_prefix, pipeline, classes=None, ann_file=None, test_mode=False)`
Fashion-MNIST Dataset.

class `mmcls.datasets.ImageNet(data_prefix, pipeline, classes=None, ann_file=None, test_mode=False)`
ImageNet Dataset.

This implementation is modified from <https://github.com/pytorch/vision/blob/master/torchvision/datasets/imagenet.py>

class `mmcls.datasets.ImageNet21k(data_prefix, pipeline, classes=None, ann_file=None, multi_label=False, recursion_subdir=False, test_mode=False)`

ImageNet21k Dataset.

Since the dataset ImageNet21k is extremely big, contains 21k+ classes and 1.4B files. This class has improved the following points on the basis of the class `ImageNet`, in order to save memory usage and time

required :

- Delete the samples attribute
- using ‘slots’ create a `Data_item` tp replace dict
- Modify setting `info` dict from function `load_annotations` to function `prepare_data`
- using `int` instead of `np.array(..., np.int64)`

Args: `data_prefix` (`str`): the prefix of data path `pipeline` (`list`): a list of dict, where each element represents a operation defined in `mmcls.datasets.pipelines`

ann_file (`str` | `None`): the annotation file. When **ann_file** is `str`, the subclass is expected to read from the `ann_file`. When `ann_file` is `None`, the subclass is expected to read according to `data_prefix`

`test_mode` (`bool`): in train mode or test mode `multi_label` (`bool`): use multi label or not. `recursion_subdir` (`bool`): whether to use sub-directory pictures, which

are meet the conditions in the folder under category directory.

get_cat_ids(`idx: int`) → `List[int]`

Get category id by index.

Parameters `idx` (`int`) – Index of data.

Returns Image category of specified index.

Return type `cat_ids` (`List[int]`)

load_annotations()

load dataset annotations.

class `mmcls.datasets.MNIST`(`data_prefix`, `pipeline`, `classes=None`, `ann_file=None`, `test_mode=False`)

MNIST Dataset.

This implementation is modified from <https://github.com/pytorch/vision/blob/master/torchvision/datasets/mnist.py>

class `mmcls.datasets.MultiLabelDataset`(`data_prefix`, `pipeline`, `classes=None`, `ann_file=None`, `test_mode=False`)

Multi-label Dataset.

evaluate(`results`, `metric='mAP'`, `metric_options=None`, `logger=None`, `**deprecated_kwargs`)

Evaluate the dataset.

Parameters

- **results** (`list`) – Testing results of the dataset.
- **metric** (`str` | `list[str]`) – Metrics to be evaluated. Default value is ‘mAP’. Options are ‘mAP’, ‘CP’, ‘CR’, ‘CF1’, ‘OP’, ‘OR’ and ‘OF1’.
- **metric_options** (`dict`, `optional`) – Options for calculating metrics. Allowed keys are ‘k’ and ‘thr’. Defaults to None
- **logger** (`logging.Logger` | `str`, `optional`) – Logger used for printing related information during evaluation. Defaults to None.
- **deprecated_kwargs** (`dict`) – Used for containing deprecated arguments.

Returns evaluation results

Return type `dict`

get_cat_ids(`idx: int`) → `List[int]`

Get category ids by index.

Parameters `idx` (`int`) – Index of data.

Returns Image categories of specified index.

Return type `cat_ids` (`List[int]`)

class `mmcls.datasets.RepeatDataset`(`dataset`, `times`)

A wrapper of repeated dataset.

The length of repeated dataset will be *times* larger than the original dataset. This is useful when the data loading time is long but the dataset is small. Using RepeatDataset can reduce the data loading time between epochs.

Parameters

- **dataset** (`Dataset`) – The dataset to be repeated.
- **times** (`int`) – Repeat times.

class `mmcls.datasets.VOC`(`**kwargs`)

Pascal VOC Dataset.

load_annotations()

Load annotations.

Returns Annotation info from XML file.

Return type list[dict]

`mmcls.datasets.build_data_loader(dataset, samples_per_gpu, workers_per_gpu, num_gpus=1, dist=True, shuffle=True, round_up=True, seed=None, pin_memory=True, persistent_workers=True, **kwargs)`

Build PyTorch DataLoader.

In distributed training, each GPU/process has a dataloader. In non-distributed training, there is only one dataloader for all GPUs.

Parameters

- **dataset** (*Dataset*) – A PyTorch dataset.
- **samples_per_gpu** (*int*) – Number of training samples on each GPU, i.e., batch size of each GPU.
- **workers_per_gpu** (*int*) – How many subprocesses to use for data loading for each GPU.
- **num_gpus** (*int*) – Number of GPUs. Only used in non-distributed training.
- **dist** (*bool*) – Distributed training/test or not. Default: True.
- **shuffle** (*bool*) – Whether to shuffle the data at every epoch. Default: True.
- **round_up** (*bool*) – Whether to round up the length of dataset by adding extra samples to make it evenly divisible. Default: True.
- **pin_memory** (*bool*) – Whether to use pin_memory in DataLoader. Default: True
- **persistent_workers** (*bool*) – If True, the data loader will not shutdown the worker processes after a dataset has been consumed once. This allows to maintain the workers Dataset instances alive. The argument also has effect in PyTorch>=1.7.0. Default: True
- **kwargs** – any keyword argument to be used to initialize DataLoader

Returns A PyTorch dataloader.

Return type DataLoader

37.2 pipelines

class `mmcls.datasets.pipelines.AutoAugment(policies, hparams={'pad_val': 128})`

Auto augmentation.

This data augmentation is proposed in [AutoAugment: Learning Augmentation Policies from Data](#).

Parameters

- **policies** (*list[list[dict]]*) – The policies of auto augmentation. Each policy in policies is a specific augmentation policy, and is composed by several augmentations (dict). When AutoAugment is called, a random policy in policies will be selected to augment images.
- **hparams** (*dict*) – Configs of hyperparameters. Hyperparameters will be used in policies that require these arguments if these arguments are not set in policy dicts. Defaults to use `_HPARAMS_DEFAULT`.

class `mmcls.datasets.pipelines.AutoContrast(prob=0.5)`

Auto adjust image contrast.

Parameters **prob** (*float*) – The probability for performing invert therefore should be in range [0, 1]. Defaults to 0.5.

class `mmcls.datasets.pipelines.Brightness`(*magnitude*, *prob=0.5*, *random_negative_prob=0.5*)
Adjust images brightness.

Parameters

- **magnitude** (*int* / *float*) – The magnitude used for adjusting brightness. A positive magnitude would enhance the brightness and a negative magnitude would make the image darker. A magnitude=0 gives the origin img.
- **prob** (*float*) – The probability for performing contrast adjusting therefore should be in range [0, 1]. Defaults to 0.5.
- **random_negative_prob** (*float*) – The probability that turns the magnitude negative, which should be in range [0,1]. Defaults to 0.5.

class `mmcls.datasets.pipelines.CenterCrop`(*crop_size*, *efficientnet_style=False*, *crop_padding=32*, *interpolation='bilinear'*, *backend='cv2'*)

Center crop the image.

Parameters

- **crop_size** (*int* / *tuple*) – Expected size after cropping with the format of (h, w).
- **efficientnet_style** (*bool*) – Whether to use efficientnet style center crop. Defaults to False.
- **crop_padding** (*int*) – The crop padding parameter in efficientnet style center crop. Only valid if `efficientnet_style` is True. Defaults to 32.
- **interpolation** (*str*) – Interpolation method, accepted values are ‘nearest’, ‘bilinear’, ‘bicubic’, ‘area’, ‘lanczos’. Only valid if `efficientnet_style` is True. Defaults to ‘bilinear’.
- **backend** (*str*) – The image resize backend type, accepted values are `cv2` and `pillow`. Only valid if `efficientnet_style` is True. Defaults to `cv2`.

Notes

- If the image is smaller than the crop size, return the original image.
- If `efficientnet_style` is set to False, the pipeline would be a simple center crop using the `crop_size`.
- If `efficientnet_style` is set to True, the pipeline will be to first to perform the center crop with the `crop_size_` as:

$$\text{crop_size_} = \frac{\text{crop_size}}{\text{crop_size} + \text{crop_padding}} \times \text{short_edge}$$

And then the pipeline resizes the img to the input crop size.

class `mmcls.datasets.pipelines.Collect`(*keys*, *meta_keys=('filename', 'ori_filename', 'ori_shape', 'img_shape', 'flip', 'flip_direction', 'img_norm_cfg')*)

Collect data from the loader relevant to the specific task.

This is usually the last stage of the data loader pipeline. Typically `keys` is set to some subset of “img” and “gt_label”.

Parameters

- **keys** (*Sequence[str]*) – Keys of results to be collected in data.

- **meta_keys** (*Sequence[str], optional*) – Meta keys to be converted to `mmcv.DataContainer` and collected in `data[img metas]`. Default: ('filename', 'ori_shape', 'img_shape', 'flip', 'flip_direction', 'img_norm_cfg')

Returns

The result dict contains the following keys

- keys in `self.keys`
- `img_metas` if available

Return type dict

class `mmcls.datasets.pipelines.ColorJitter`(*brightness, contrast, saturation*)

Randomly change the brightness, contrast and saturation of an image.

Parameters

- **brightness** (*float*) – How much to jitter brightness. `brightness_factor` is chosen uniformly from `[max(0, 1 - brightness), 1 + brightness]`.
- **contrast** (*float*) – How much to jitter contrast. `contrast_factor` is chosen uniformly from `[max(0, 1 - contrast), 1 + contrast]`.
- **saturation** (*float*) – How much to jitter saturation. `saturation_factor` is chosen uniformly from `[max(0, 1 - saturation), 1 + saturation]`.

class `mmcls.datasets.pipelines.ColorTransform`(*magnitude, prob=0.5, random_negative_prob=0.5*)

Adjust images color balance.

Parameters

- **magnitude** (*int | float*) – The magnitude used for color transform. A positive magnitude would enhance the color and a negative magnitude would make the image grayer. A magnitude=0 gives the origin img.
- **prob** (*float*) – The probability for performing `ColorTransform` therefore should be in range `[0, 1]`. Defaults to 0.5.
- **random_negative_prob** (*float*) – The probability that turns the magnitude negative, which should be in range `[0,1]`. Defaults to 0.5.

class `mmcls.datasets.pipelines.Compose`(*transforms*)

Compose a data pipeline with a sequence of transforms.

Parameters `transforms` (*list[dict | callable]*) – Either config dicts of transforms or transform objects.

class `mmcls.datasets.pipelines.Contrast`(*magnitude, prob=0.5, random_negative_prob=0.5*)

Adjust images contrast.

Parameters

- **magnitude** (*int | float*) – The magnitude used for adjusting contrast. A positive magnitude would enhance the contrast and a negative magnitude would make the image grayer. A magnitude=0 gives the origin img.
- **prob** (*float*) – The probability for performing contrast adjusting therefore should be in range `[0, 1]`. Defaults to 0.5.
- **random_negative_prob** (*float*) – The probability that turns the magnitude negative, which should be in range `[0,1]`. Defaults to 0.5.

class `mmcls.datasets.pipelines.Cutout`(*shape*, *pad_val=128*, *prob=0.5*)

Cutout images.

Parameters

- **shape** (*int* | *float* | *tuple(int | float)*) – Expected cutout shape (h, w). If given as a single value, the value will be used for both h and w.
- **pad_val** (*int*, *Sequence[int]*) – Pixel *pad_val* value for constant fill. If it is a sequence, it must have the same length with the image channels. Defaults to 128.
- **prob** (*float*) – The probability for performing cutout therefore should be in range [0, 1]. Defaults to 0.5.

class `mmcls.datasets.pipelines.Equalize`(*prob=0.5*)

Equalize the image histogram.

Parameters **prob** (*float*) – The probability for performing invert therefore should be in range [0, 1]. Defaults to 0.5.

class `mmcls.datasets.pipelines.Invert`(*prob=0.5*)

Invert images.

Parameters **prob** (*float*) – The probability for performing invert therefore should be in range [0, 1]. Defaults to 0.5.

class `mmcls.datasets.pipelines.Lighting`(*eigval*, *eigvec*, *alphastd=0.1*, *to_rgb=True*)

Adjust images lighting using AlexNet-style PCA jitter.

Parameters

- **eigval** (*list*) – the eigenvalue of the covariance matrix of pixel values, respectively.
- **eigvec** (*list[list]*) – the eigenvector of the covariance matrix of pixel values, respectively.
- **alphastd** (*float*) – The standard deviation for distribution of alpha. Defaults to 0.1
- **to_rgb** (*bool*) – Whether to convert img to rgb.

class `mmcls.datasets.pipelines.LoadImageFromFile`(*to_float32=False*, *color_type='color'*,
file_client_args={'backend': 'disk'})

Load an image from file.

Required keys are “img_prefix” and “img_info” (a dict that must contain the key “filename”). Added or updated keys are “filename”, “img”, “img_shape”, “ori_shape” (same as *img_shape*) and “img_norm_cfg” (means=0 and stds=1).

Parameters

- **to_float32** (*bool*) – Whether to convert the loaded image to a float32 numpy array. If set to False, the loaded image is an uint8 array. Defaults to False.
- **color_type** (*str*) – The flag argument for `mmcv.imfrombytes()`. Defaults to ‘color’.
- **file_client_args** (*dict*) – Arguments to instantiate a `FileClient`. See `mmcv.fileio.FileClient` for details. Defaults to `dict(backend='disk')`.

class `mmcls.datasets.pipelines.Normalize`(*mean*, *std*, *to_rgb=True*)

Normalize the image.

Parameters

- **mean** (*sequence*) – Mean values of 3 channels.
- **std** (*sequence*) – Std values of 3 channels.

- **to_rgb** (*bool*) – Whether to convert the image from BGR to RGB, default is true.

```
class mmcls.datasets.pipelines.Pad(size=None, pad_to_square=False, pad_val=0,
                                   padding_mode='constant')
```

Pad images.

Parameters

- **size** (*tuple[int] | None*) – Expected padding size (h, w). Conflicts with `pad_to_square`. Defaults to None.
- **pad_to_square** (*bool*) – Pad any image to square shape. Defaults to False.
- **pad_val** (*Number | Sequence[Number]*) – Values to be filled in padding areas when `padding_mode` is 'constant'. Default to 0.
- **padding_mode** (*str*) – Type of padding. Should be: constant, edge, reflect or symmetric. Default to "constant".

```
class mmcls.datasets.pipelines.Posterize(bits, prob=0.5)
```

Posterize images (reduce the number of bits for each color channel).

Parameters

- **bits** (*int | float*) – Number of bits for each pixel in the output img, which should be less or equal to 8.
- **prob** (*float*) – The probability for posterizing therefore should be in range [0, 1]. Defaults to 0.5.

```
class mmcls.datasets.pipelines.RandAugment(policies, num_policies, magnitude_level,
                                           magnitude_std=0.0, total_level=30, hparams={'pad_val':
                                                                                       128})
```

Random augmentation.

This data augmentation is proposed in [RandAugment: Practical automated data augmentation with a reduced search space](#).

Parameters

- **policies** (*list[dict]*) – The policies of random augmentation. Each policy in `policies` is one specific augmentation policy (dict). The policy shall at least have key `type`, indicating the type of augmentation. For those which have magnitude, (given to the fact they are named differently in different augmentation,) `magnitude_key` and `magnitude_range` shall be the magnitude argument (str) and the range of magnitude (tuple in the format of (val1, val2)), respectively. Note that val1 is not necessarily less than val2.
- **num_policies** (*int*) – Number of policies to select from policies each time.
- **magnitude_level** (*int | float*) – Magnitude level for all the augmentation selected.
- **total_level** (*int | float*) – Total level for the magnitude. Defaults to 30.
- **magnitude_std** (*Number | str*) – Deviation of magnitude noise applied.
 - If positive number, magnitude is sampled from normal distribution (mean=magnitude, std=magnitude_std).
 - If 0 or negative number, magnitude remains unchanged.
 - If str "inf", magnitude is sampled from uniform distribution (range=[min, magnitude]).
- **hparams** (*dict*) – Configs of hyperparameters. Hyperparameters will be used in policies that require these arguments if these arguments are not set in policy dicts. Defaults to use `_HPARAMS_DEFAULT`.

Note: `magnitude_std` will introduce some randomness to policy, modified by <https://github.com/rwightman/pytorch-image-models>.

When `magnitude_std=0`, we calculate the magnitude as follows:

$$\text{magnitude} = \frac{\text{magnitude_level}}{\text{total_level}} \times (\text{val2} - \text{val1}) + \text{val1}$$

class `mmcls.datasets.pipelines.RandomCrop`(*size*, *padding=None*, *pad_if_needed=False*, *pad_val=0*, *padding_mode='constant'*)

Crop the given Image at a random location.

Parameters

- **size** (*sequence or int*) – Desired output size of the crop. If size is an int instead of sequence like (h, w), a square crop (size, size) is made.
- **padding** (*int or sequence, optional*) – Optional padding on each border of the image. If a sequence of length 4 is provided, it is used to pad left, top, right, bottom borders respectively. If a sequence of length 2 is provided, it is used to pad left/right, top/bottom borders, respectively. Default: None, which means no padding.
- **pad_if_needed** (*boolean*) – It will pad the image if smaller than the desired size to avoid raising an exception. Since cropping is done after padding, the padding seems to be done at a random offset. Default: False.
- **pad_val** (*Number | Sequence[Number]*) – Pixel pad_val value for constant fill. If a tuple of length 3, it is used to pad_val R, G, B channels respectively. Default: 0.
- **padding_mode** (*str*) – Type of padding. Defaults to “constant”. Should be one of the following:
 - constant: Pads with a constant value, this value is specified with pad_val.
 - edge: pads with the last value at the edge of the image.
 - reflect: Pads with reflection of image without repeating the last value on the edge. For example, padding [1, 2, 3, 4] with 2 elements on both sides in reflect mode will result in [3, 2, 1, 2, 3, 4, 3, 2].
 - symmetric: Pads with reflection of image repeating the last value on the edge. For example, padding [1, 2, 3, 4] with 2 elements on both sides in symmetric mode will result in [2, 1, 1, 2, 3, 4, 4, 3].

static `get_params`(*img*, *output_size*)

Get parameters for `crop` for a random crop.

Parameters

- **img** (*ndarray*) – Image to be cropped.
- **output_size** (*tuple*) – Expected output size of the crop.

Returns

Params (*xmin*, *ymin*, *target_height*, *target_width*) to be passed to `crop` for random crop.

Return type `tuple`

```
class mmcls.datasets.pipelines.RandomErasing(erase_prob=0.5, min_area_ratio=0.02,
max_area_ratio=0.4, aspect_range=(0.3,
3.3333333333333335), mode='const', fill_color=(128,
128, 128), fill_std=None)
```

Randomly selects a rectangle region in an image and erase pixels.

Parameters

- **erase_prob** (*float*) – Probability that image will be randomly erased. Default: 0.5
- **min_area_ratio** (*float*) – Minimum erased area / input image area Default: 0.02
- **max_area_ratio** (*float*) – Maximum erased area / input image area Default: 0.4
- **aspect_range** (*sequence* | *float*) – Aspect ratio range of erased area. if float, it will be converted to (aspect_ratio, 1/aspect_ratio) Default: (3/10, 10/3)
- **mode** (*str*) – Fill method in erased area, can be:
 - const (default): All pixels are assign with the same value.
 - rand: each pixel is assigned with a random value in [0, 255]
- **fill_color** (*sequence* | *Number*) – Base color filled in erased area. Defaults to (128, 128, 128).
- **fill_std** (*sequence* | *Number*, *optional*) – If set and mode is 'rand', fill erased area with random color from normal distribution (mean=fill_color, std=fill_std); If not set, fill erased area with random color from uniform distribution (0~255). Defaults to None.

Note: See [Random Erasing Data Augmentation](#)

This paper provided 4 modes: RE-R, RE-M, RE-0, RE-255, and use RE-M as default. The config of these 4 modes are:

- RE-R: RandomErasing(mode='rand')
 - RE-M: RandomErasing(mode='const', fill_color=(123.67, 116.3, 103.5))
 - RE-0: RandomErasing(mode='const', fill_color=0)
 - RE-255: RandomErasing(mode='const', fill_color=255)
-

```
class mmcls.datasets.pipelines.RandomFlip(flip_prob=0.5, direction='horizontal')
```

Flip the image randomly.

Flip the image randomly based on flip probability and flip direction.

Parameters

- **flip_prob** (*float*) – probability of the image being flipped. Default: 0.5
- **direction** (*str*) – The flipping direction. Options are 'horizontal' and 'vertical'. Default: 'horizontal'.

```
class mmcls.datasets.pipelines.RandomGrayscale(gray_prob=0.1)
```

Randomly convert image to grayscale with a probability of gray_prob.

Parameters **gray_prob** (*float*) – Probability that image should be converted to grayscale. Default: 0.1.

Returns Image after randomly grayscale transform.

Return type ndarray

Notes

- If input image is 1 channel: grayscale version is 1 channel.
- If input image is 3 channel: grayscale version is 3 channel with r == g == b.

```
class mmcls.datasets.pipelines.RandomResizedCrop(size, scale=(0.08, 1.0), ratio=(0.75,  
1.3333333333333333), max_attempts=10,  
efficientnet_style=False, min_covered=0.1,  
crop_padding=32, interpolation='bilinear',  
backend='cv2')
```

Crop the given image to random size and aspect ratio.

A crop of random size (default: of 0.08 to 1.0) of the original size and a random aspect ratio (default: of 3/4 to 4/3) of the original aspect ratio is made. This crop is finally resized to given size.

Parameters

- **size** (*sequence* / *int*) – Desired output size of the crop. If size is an int instead of sequence like (h, w), a square crop (size, size) is made.
- **scale** (*tuple*) – Range of the random size of the cropped image compared to the original image. Defaults to (0.08, 1.0).
- **ratio** (*tuple*) – Range of the random aspect ratio of the cropped image compared to the original image. Defaults to (3. / 4., 4. / 3.).
- **max_attempts** (*int*) – Maximum number of attempts before falling back to Central Crop. Defaults to 10.
- **efficientnet_style** (*bool*) – Whether to use efficientnet style Random ResizedCrop. Defaults to False.
- **min_covered** (*Number*) – Minimum ratio of the cropped area to the original area. Only valid if efficientnet_style is true. Defaults to 0.1.
- **crop_padding** (*int*) – The crop padding parameter in efficientnet style center crop. Only valid if efficientnet_style is true. Defaults to 32.
- **interpolation** (*str*) – Interpolation method, accepted values are ‘nearest’, ‘bilinear’, ‘bicubic’, ‘area’, ‘lanczos’. Defaults to ‘bilinear’.
- **backend** (*str*) – The image resize backend type, accepted values are *cv2* and *pillow*. Defaults to *cv2*.

```
static get_params(img, scale, ratio, max_attempts=10)
```

Get parameters for `crop` for a random sized crop.

Parameters

- **img** (*ndarray*) – Image to be cropped.
- **scale** (*tuple*) – Range of the random size of the cropped image compared to the original image size.
- **ratio** (*tuple*) – Range of the random aspect ratio of the cropped image compared to the original image area.
- **max_attempts** (*int*) – Maximum number of attempts before falling back to central crop. Defaults to 10.

Returns

Params (ymin, xmin, ymax, xmax) to be passed to `crop` for a random sized crop.

Return type tuple

```
static get_params_efficientnet_style(img, size, scale, ratio, max_attempts=10, min_covered=0.1,  
crop_padding=32)
```

Get parameters for `crop` for a random sized crop in efficientnet style.

Parameters

- **img** (*ndarray*) – Image to be cropped.
- **size** (*sequence*) – Desired output size of the crop.
- **scale** (*tuple*) – Range of the random size of the cropped image compared to the original image size.
- **ratio** (*tuple*) – Range of the random aspect ratio of the cropped image compared to the original image area.
- **max_attempts** (*int*) – Maximum number of attempts before falling back to central crop. Defaults to 10.
- **min_covered** (*Number*) – Minimum ratio of the cropped area to the original area. Only valid if `efficientnet_style` is true. Defaults to 0.1.
- **crop_padding** (*int*) – The crop padding parameter in efficientnet style center crop. Defaults to 32.

Returns

Params (ymin, xmin, ymax, xmax) to be passed to `crop` for a random sized crop.

Return type tuple

```
class mmcls.datasets.pipelines.Resize(size, interpolation='bilinear', adaptive_side='short',  
backend='cv2')
```

Resize images.

Parameters

- **size** (*int | tuple*) – Images scales for resizing (h, w). When size is int, the default behavior is to resize an image to (size, size). When size is tuple and the second value is -1, the image will be resized according to `adaptive_side`. For example, when size is 224, the image is resized to 224x224. When size is (224, -1) and `adaptive_side` is “short”, the short side is resized to 224 and the other side is computed based on the short side, maintaining the aspect ratio.
- **interpolation** (*str*) – Interpolation method. For “cv2” backend, accepted values are “nearest”, “bilinear”, “bicubic”, “area”, “lanczos”. For “pillow” backend, accepted values are “nearest”, “bilinear”, “bicubic”, “box”, “lanczos”, “hamming”. More details can be found in `mmcv.image.geometric`.
- **adaptive_side** (*str*) – Adaptive resize policy, accepted values are “short”, “long”, “height”, “width”. Default to “short”.
- **backend** (*str*) – The image resize backend type, accepted values are `cv2` and `pillow`. Default: `cv2`.

```
class mmcls.datasets.pipelines.Rotate(angle, center=None, scale=1.0, pad_val=128, prob=0.5,  
random_negative_prob=0.5, interpolation='nearest')
```

Rotate images.

Parameters

- **angle** (*float*) – The angle used for rotate. Positive values stand for clockwise rotation.

- **center** (*tuple[float], optional*) – Center point (w, h) of the rotation in the source image. If None, the center of the image will be used. Defaults to None.
- **scale** (*float*) – Isotropic scale factor. Defaults to 1.0.
- **pad_val** (*int, Sequence[int]*) – Pixel pad_val value for constant fill. If a sequence of length 3, it is used to pad_val R, G, B channels respectively. Defaults to 128.
- **prob** (*float*) – The probability for performing Rotate therefore should be in range [0, 1]. Defaults to 0.5.
- **random_negative_prob** (*float*) – The probability that turns the angle negative, which should be in range [0,1]. Defaults to 0.5.
- **interpolation** (*str*) – Interpolation method. Options are ‘nearest’, ‘bilinear’, ‘bicubic’, ‘area’, ‘lanczos’. Defaults to ‘nearest’.

class `mmcls.datasets.pipelines.Sharpness`(*magnitude, prob=0.5, random_negative_prob=0.5*)
Adjust images sharpness.

Parameters

- **magnitude** (*int | float*) – The magnitude used for adjusting sharpness. A positive magnitude would enhance the sharpness and a negative magnitude would make the image blur. A magnitude=0 gives the origin img.
- **prob** (*float*) – The probability for performing contrast adjusting therefore should be in range [0, 1]. Defaults to 0.5.
- **random_negative_prob** (*float*) – The probability that turns the magnitude negative, which should be in range [0,1]. Defaults to 0.5.

class `mmcls.datasets.pipelines.Shear`(*magnitude, pad_val=128, prob=0.5, direction='horizontal', random_negative_prob=0.5, interpolation='bicubic'*)

Shear images.

Parameters

- **magnitude** (*int | float*) – The magnitude used for shear.
- **pad_val** (*int, Sequence[int]*) – Pixel pad_val value for constant fill. If a sequence of length 3, it is used to pad_val R, G, B channels respectively. Defaults to 128.
- **prob** (*float*) – The probability for performing Shear therefore should be in range [0, 1]. Defaults to 0.5.
- **direction** (*str*) – The shearing direction. Options are ‘horizontal’ and ‘vertical’. Defaults to ‘horizontal’.
- **random_negative_prob** (*float*) – The probability that turns the magnitude negative, which should be in range [0,1]. Defaults to 0.5.
- **interpolation** (*str*) – Interpolation method. Options are ‘nearest’, ‘bilinear’, ‘bicubic’, ‘area’, ‘lanczos’. Defaults to ‘bicubic’.

class `mmcls.datasets.pipelines.Solarize`(*thr, prob=0.5*)
Solarize images (invert all pixel values above a threshold).

Parameters

- **thr** (*int | float*) – The threshold above which the pixels value will be inverted.
- **prob** (*float*) – The probability for solarizing therefore should be in range [0, 1]. Defaults to 0.5.

class `mmcls.datasets.pipelines.SolarizeAdd`(*magnitude*, *thr=128*, *prob=0.5*)

SolarizeAdd images (add a certain value to pixels below a threshold).

Parameters

- **magnitude** (*int* | *float*) – The value to be added to pixels below the thr.
- **thr** (*int* | *float*) – The threshold below which the pixels value will be adjusted.
- **prob** (*float*) – The probability for solarizing therefore should be in range [0, 1]. Defaults to 0.5.

class `mmcls.datasets.pipelines.Translate`(*magnitude*, *pad_val=128*, *prob=0.5*, *direction='horizontal'*, *random_negative_prob=0.5*, *interpolation='nearest'*)

Translate images.

Parameters

- **magnitude** (*int* | *float*) – The magnitude used for translate. Note that the offset is calculated by `magnitude * size` in the corresponding direction. With a magnitude of 1, the whole image will be moved out of the range.
- **pad_val** (*int*, *Sequence[int]*) – Pixel `pad_val` value for constant fill. If a sequence of length 3, it is used to `pad_val` R, G, B channels respectively. Defaults to 128.
- **prob** (*float*) – The probability for performing translate therefore should be in range [0, 1]. Defaults to 0.5.
- **direction** (*str*) – The translating direction. Options are 'horizontal' and 'vertical'. Defaults to 'horizontal'.
- **random_negative_prob** (*float*) – The probability that turns the magnitude negative, which should be in range [0,1]. Defaults to 0.5.
- **interpolation** (*str*) – Interpolation method. Options are 'nearest', 'bilinear', 'bicubic', 'area', 'lanczos'. Defaults to 'nearest'.

`mmcls.datasets.pipelines.to_tensor`(*data*)

Convert objects of various python types to `torch.Tensor`.

Supported types are: `numpy.ndarray`, `torch.Tensor`, `Sequence`, `int` and `float`.

MMCLS.UTILS

`mmcls.utils.collect_env()`

Collect the information of the running environments.

`mmcls.utils.load_json_logs(json_logs)`

load and convert json_logs to log_dicts.

Parameters `json_logs` (*str*) – paths of json_logs.

Returns

dict(): **key is epoch, value is a sub dict keys of** sub dict is different metrics, e.g. memory, bbox_mAP, value of sub dict is a list of corresponding values of all iterations.

Return type list[dict(int

INDICES AND TABLES

- [genindex](#)
- [search](#)

PYTHON MODULE INDEX

m

- `mmcls.apis`, 111
- `mmcls.core.evaluation`, 113
- `mmcls.datasets`, 153
- `mmcls.datasets.pipelines`, 157
- `mmcls.models`, 117
 - `mmcls.models.backbones`, 119
 - `mmcls.models.classifiers`, 117
 - `mmcls.models.heads`, 139
 - `mmcls.models.losses`, 141
 - `mmcls.models.necks`, 138
 - `mmcls.models.utils`, 147
- `mmcls.utils`, 169

A

Accuracy (class in *mmcls.models.losses*), 141
 accuracy() (in module *mmcls.models.losses*), 144
 adjust_width_group() (in module *mmcls.models.backbones.RegNet* method), 123
 AlexNet (class in *mmcls.models.backbones*), 119
 asymmetric_loss() (in module *mmcls.models.losses*), 144
 AsymmetricLoss (class in *mmcls.models.losses*), 141
 Augments (class in *mmcls.models.utils*), 147
 AutoAugment (class in *mmcls.datasets.pipelines*), 157
 AutoContrast (class in *mmcls.datasets.pipelines*), 157
 average_performance() (in module *mmcls.core.evaluation*), 113
 average_precision() (in module *mmcls.core.evaluation*), 113

B

BaseClassifier (class in *mmcls.models.classifiers*), 117
 BaseDataset (class in *mmcls.datasets*), 153
 binary_cross_entropy() (in module *mmcls.models.losses*), 145
 Brightness (class in *mmcls.datasets.pipelines*), 158
 build_backbone() (in module *mmcls.models*), 117
 build_data_loader() (in module *mmcls.datasets*), 157
 build_head() (in module *mmcls.models*), 117
 build_loss() (in module *mmcls.models*), 117
 build_neck() (in module *mmcls.models*), 117

C

calculate_confusion_matrix() (in module *mmcls.core.evaluation*), 114
 CenterCrop (class in *mmcls.datasets.pipelines*), 158
 channel_shuffle() (in module *mmcls.models.utils*), 152
 CIFAR10 (class in *mmcls.datasets*), 154
 CIFAR100 (class in *mmcls.datasets*), 154
 class_to_idx (in module *mmcls.datasets.BaseDataset* property), 153
 ClassBalancedDataset (class in *mmcls.datasets*), 154
 ClsHead (class in *mmcls.models.heads*), 139

Collect (class in *mmcls.datasets.pipelines*), 158
 collect_env() (in module *mmcls.utils*), 169
 ColorJitter (class in *mmcls.datasets.pipelines*), 159
 ColorTransform (class in *mmcls.datasets.pipelines*), 159
 Compose (class in *mmcls.datasets.pipelines*), 159
 ConcatDataset (class in *mmcls.datasets*), 155
 Contrast (class in *mmcls.datasets.pipelines*), 159
 convert_to_one_hot() (in module *mmcls.models.losses*), 145
 cross_entropy() (in module *mmcls.models.losses*), 145
 CrossEntropyLoss (class in *mmcls.models.losses*), 141
 Cutout (class in *mmcls.datasets.pipelines*), 159

D

DistEvalHook (class in *mmcls.core.evaluation*), 113
 DistributedSampler (class in *mmcls.datasets*), 155

E

Equalize (class in *mmcls.datasets.pipelines*), 160
 EvalHook (class in *mmcls.core.evaluation*), 113
 evaluate() (*mmcls.datasets.BaseDataset* method), 153
 evaluate() (*mmcls.datasets.MultiLabelDataset* method), 156
 extract_feat() (in module *mmcls.models.classifiers.ImageClassifier* method), 119

F

f1_score() (in module *mmcls.core.evaluation*), 114
 FashionMNIST (class in *mmcls.datasets*), 155
 FocalLoss (class in *mmcls.models.losses*), 142
 forward() (*mmcls.models.backbones.AlexNet* method), 119
 forward() (*mmcls.models.backbones.LeNet5* method), 120
 forward() (*mmcls.models.backbones.MlpMixer* method), 120
 forward() (*mmcls.models.backbones.MobileNetV2* method), 121
 forward() (*mmcls.models.backbones.MobileNetV3* method), 122

- [forward\(\)](#) (*mmcls.models.backbones.RegNet method*), 123
[forward\(\)](#) (*mmcls.models.backbones.RepVGG method*), 125
[forward\(\)](#) (*mmcls.models.backbones.ResNet method*), 129
[forward\(\)](#) (*mmcls.models.backbones.ResNet_CIFAR method*), 130
[forward\(\)](#) (*mmcls.models.backbones.ShuffleNetV1 method*), 132
[forward\(\)](#) (*mmcls.models.backbones.ShuffleNetV2 method*), 133
[forward\(\)](#) (*mmcls.models.backbones.SwinTransformer method*), 134
[forward\(\)](#) (*mmcls.models.backbones.T2T_ViT method*), 135
[forward\(\)](#) (*mmcls.models.backbones.TIMMBackbone method*), 135
[forward\(\)](#) (*mmcls.models.backbones.TNT method*), 136
[forward\(\)](#) (*mmcls.models.backbones.VGG method*), 137
[forward\(\)](#) (*mmcls.models.backbones.VisionTransformer method*), 138
[forward\(\)](#) (*mmcls.models.classifiers.BaseClassifier method*), 117
[forward\(\)](#) (*mmcls.models.losses.Accuracy method*), 141
[forward\(\)](#) (*mmcls.models.losses.AsymmetricLoss method*), 141
[forward\(\)](#) (*mmcls.models.losses.CrossEntropyLoss method*), 141
[forward\(\)](#) (*mmcls.models.losses.FocalLoss method*), 142
[forward\(\)](#) (*mmcls.models.losses.LabelSmoothLoss method*), 143
[forward\(\)](#) (*mmcls.models.losses.SeesawLoss method*), 144
[forward\(\)](#) (*mmcls.models.necks.GlobalAveragePooling method*), 138
[forward\(\)](#) (*mmcls.models.utils.HybridEmbed method*), 148
[forward\(\)](#) (*mmcls.models.utils.InvertedResidual method*), 149
[forward\(\)](#) (*mmcls.models.utils.MultiheadAttention method*), 149
[forward\(\)](#) (*mmcls.models.utils.PatchEmbed method*), 150
[forward\(\)](#) (*mmcls.models.utils.PatchMerging method*), 151
[forward\(\)](#) (*mmcls.models.utils.SELayer method*), 151
[forward\(\)](#) (*mmcls.models.utils.ShiftWindowMSA method*), 152
[forward_test\(\)](#) (*mmcls.models.classifiers.BaseClassifier method*), 117
[forward_train\(\)](#) (*mmcls.models.classifiers.BaseClassifier method*), 117
[forward_train\(\)](#) (*mmcls.models.classifiers.ImageClassifier method*), 119
G
[generate_one_hot_like_label\(\)](#) (*mmcls.models.losses.LabelSmoothLoss method*), 143
[generate_regnet\(\)](#) (*mmcls.models.backbones.RegNet method*), 123
[get_cat_ids\(\)](#) (*mmcls.datasets.BaseDataset method*), 153
[get_cat_ids\(\)](#) (*mmcls.datasets.ImageNet21k method*), 155
[get_cat_ids\(\)](#) (*mmcls.datasets.MultiLabelDataset method*), 156
[get_classes\(\)](#) (*mmcls.datasets.BaseDataset class method*), 154
[get_gt_labels\(\)](#) (*mmcls.datasets.BaseDataset method*), 154
[get_params\(\)](#) (*mmcls.datasets.pipelines.RandomCrop static method*), 162
[get_params\(\)](#) (*mmcls.datasets.pipelines.RandomResizedCrop static method*), 164
[get_params_efficientnet_style\(\)](#) (*mmcls.datasets.pipelines.RandomResizedCrop static method*), 165
[get_stages_from_blocks\(\)](#) (*mmcls.models.backbones.RegNet method*), 124
[GlobalAveragePooling](#) (*class in mmcls.models.necks*), 138
H
[HybridEmbed](#) (*class in mmcls.models.utils*), 148
I
[ImageClassifier](#) (*class in mmcls.models.classifiers*), 119
[ImageNet](#) (*class in mmcls.datasets*), 155
[ImageNet21k](#) (*class in mmcls.datasets*), 155
[inference_model\(\)](#) (*in module mmcls.apis*), 111
[init_model\(\)](#) (*in module mmcls.apis*), 111
[init_weights\(\)](#) (*mmcls.models.backbones.ResNet method*), 129
[init_weights\(\)](#) (*mmcls.models.backbones.ShuffleNetV1 method*), 132
[init_weights\(\)](#) (*mmcls.models.backbones.ShuffleNetV2 method*), 133

- `init_weights()` (*mmcls.models.backbones.SwinTransformer method*), 134
- `init_weights()` (*mmcls.models.backbones.T2T_ViT method*), 135
- `init_weights()` (*mmcls.models.backbones.VisionTransformer method*), 138
- `init_weights()` (*mmcls.models.heads.StackedLinearClsHead method*), 140
- `init_weights()` (*mmcls.models.heads.VisionTransformerClsHead method*), 140
- `Invert` (*class in mmcls.datasets.pipelines*), 160
- `InvertedResidual` (*class in mmcls.models.utils*), 148
- ## L
- `LabelSmoothLoss` (*class in mmcls.models.losses*), 142
- `LeNet5` (*class in mmcls.models.backbones*), 119
- `Lighting` (*class in mmcls.datasets.pipelines*), 160
- `LinearClsHead` (*class in mmcls.models.heads*), 139
- `load_annotations()` (*mmcls.datasets.ImageNet21k method*), 156
- `load_annotations()` (*mmcls.datasets.VOC method*), 156
- `load_json_logs()` (*in module mmcls.utils*), 169
- `LoadImageFromFile` (*class in mmcls.datasets.pipelines*), 160
- ## M
- `make_divisible()` (*in module mmcls.models.utils*), 152
- `make_layer()` (*mmcls.models.backbones.MobileNetV2 method*), 121
- `make_layer()` (*mmcls.models.backbones.ShuffleNetV1 method*), 132
- `mAP()` (*in module mmcls.core.evaluation*), 114
- `MlpMixer` (*class in mmcls.models.backbones*), 120
- `mmcls.apis`
module, 111
- `mmcls.core.evaluation`
module, 113
- `mmcls.datasets`
module, 153
- `mmcls.datasets.pipelines`
module, 157
- `mmcls.models`
module, 117
- `mmcls.models.backbones`
module, 119
- `mmcls.models.classifiers`
module, 117
- `mmcls.models.heads`
module, 139
- `mmcls.models.losses`
module, 141
- `mmcls.models.necks`
module, 138
- `mmcls.models.utils`
module, 147
- `mmcls.utils`
module, 169
- `MNIST` (*class in mmcls.datasets*), 156
- `MobileNetV2` (*class in mmcls.models.backbones*), 120
- `MobileNetV3` (*class in mmcls.models.backbones*), 121
- module
- `mmcls.apis`, 111
 - `mmcls.core.evaluation`, 113
 - `mmcls.datasets`, 153
 - `mmcls.datasets.pipelines`, 157
 - `mmcls.models`, 117
 - `mmcls.models.backbones`, 119
 - `mmcls.models.classifiers`, 117
 - `mmcls.models.heads`, 139
 - `mmcls.models.losses`, 141
 - `mmcls.models.necks`, 138
 - `mmcls.models.utils`, 147
 - `mmcls.utils`, 169
- `multi_gpu_test()` (*in module mmcls.apis*), 111
- `MultiheadAttention` (*class in mmcls.models.utils*), 149
- `MultiLabelClsHead` (*class in mmcls.models.heads*), 139
- `MultiLabelDataset` (*class in mmcls.datasets*), 156
- `MultiLabelLinearClsHead` (*class in mmcls.models.heads*), 139
- ## N
- `Normalize` (*class in mmcls.datasets.pipelines*), 160
- ## P
- `Pad` (*class in mmcls.datasets.pipelines*), 161
- `PatchEmbed` (*class in mmcls.models.utils*), 150
- `PatchMerging` (*class in mmcls.models.utils*), 150
- `Posterize` (*class in mmcls.datasets.pipelines*), 161
- `precision()` (*in module mmcls.core.evaluation*), 115
- `precision_recall_f1()` (*in module mmcls.core.evaluation*), 115
- ## Q
- `quantize_float()` (*mmcls.models.backbones.RegNet static method*), 124
- ## R
- `RandAugment` (*class in mmcls.datasets.pipelines*), 161
- `RandomCrop` (*class in mmcls.datasets.pipelines*), 162
- `RandomErasing` (*class in mmcls.datasets.pipelines*), 162
- `RandomFlip` (*class in mmcls.datasets.pipelines*), 163

- RandomGrayscale (class in *mmcls.datasets.pipelines*), 163
- RandomResizedCrop (class in *mmcls.datasets.pipelines*), 164
- recall() (in module *mmcls.core.evaluation*), 116
- reduce_loss() (in module *mmcls.models.losses*), 146
- RegNet (class in *mmcls.models.backbones*), 122
- RepeatDataset (class in *mmcls.datasets*), 156
- RepVGG (class in *mmcls.models.backbones*), 124
- Res2Net (class in *mmcls.models.backbones*), 125
- Resize (class in *mmcls.datasets.pipelines*), 165
- resize_pos_embed() (*mmcls.models.backbones.VisionTransformer* static method), 138
- ResNeSt (class in *mmcls.models.backbones*), 126
- ResNet (class in *mmcls.models.backbones*), 128
- ResNet_CIFAR (class in *mmcls.models.backbones*), 129
- ResNetV1d (class in *mmcls.models.backbones*), 129
- ResNeXt (class in *mmcls.models.backbones*), 127
- Rotate (class in *mmcls.datasets.pipelines*), 165
- ## S
- SeesawLoss (class in *mmcls.models.losses*), 143
- SELayer (class in *mmcls.models.utils*), 151
- SEResNet (class in *mmcls.models.backbones*), 131
- SEResNeXt (class in *mmcls.models.backbones*), 130
- set_random_seed() (in module *mmcls.apis*), 111
- Sharpness (class in *mmcls.datasets.pipelines*), 166
- Shear (class in *mmcls.datasets.pipelines*), 166
- ShiftWindowMSA (class in *mmcls.models.utils*), 151
- show_result() (*mmcls.models.classifiers.BaseClassifier* method), 117
- show_result_pyplot() (in module *mmcls.apis*), 112
- ShuffleNetV1 (class in *mmcls.models.backbones*), 132
- ShuffleNetV2 (class in *mmcls.models.backbones*), 133
- sigmoid_focal_loss() (in module *mmcls.models.losses*), 146
- simple_test() (*mmcls.models.classifiers.ImageClassifier* method), 119
- simple_test() (*mmcls.models.heads.ClsHead* method), 139
- simple_test() (*mmcls.models.heads.LinearClsHead* method), 139
- simple_test() (*mmcls.models.heads.MultiLabelLinearClsHead* method), 139
- simple_test() (*mmcls.models.heads.StackedLinearClsHead* method), 140
- simple_test() (*mmcls.models.heads.VisionTransformerClsHead* method), 140
- Solarize (class in *mmcls.datasets.pipelines*), 166
- SolarizeAdd (class in *mmcls.datasets.pipelines*), 166
- StackedLinearClsHead (class in *mmcls.models.heads*), 139
- support() (in module *mmcls.core.evaluation*), 116
- SwinTransformer (class in *mmcls.models.backbones*), 133
- ## T
- T2T_ViT (class in *mmcls.models.backbones*), 134
- TIMMBackbone (class in *mmcls.models.backbones*), 135
- TNT (class in *mmcls.models.backbones*), 136
- to_tensor() (in module *mmcls.datasets.pipelines*), 167
- train() (*mmcls.models.backbones.MobileNetV2* method), 121
- train() (*mmcls.models.backbones.MobileNetV3* method), 122
- train() (*mmcls.models.backbones.RepVGG* method), 125
- train() (*mmcls.models.backbones.ResNet* method), 129
- train() (*mmcls.models.backbones.ShuffleNetV1* method), 133
- train() (*mmcls.models.backbones.ShuffleNetV2* method), 133
- train() (*mmcls.models.backbones.VGG* method), 137
- train_step() (*mmcls.models.classifiers.BaseClassifier* method), 118
- Translate (class in *mmcls.datasets.pipelines*), 167
- ## V
- val_step() (*mmcls.models.classifiers.BaseClassifier* method), 118
- VGG (class in *mmcls.models.backbones*), 136
- VisionTransformer (class in *mmcls.models.backbones*), 137
- VisionTransformerClsHead (class in *mmcls.models.heads*), 140
- VOC (class in *mmcls.datasets*), 156
- ## W
- weight_reduce_loss() (in module *mmcls.models.losses*), 146
- weighted_loss() (in module *mmcls.models.losses*), 147

