
DenseTorch

Jul 07, 2020

Contents

1 DenseTorch: PyTorch Wrapper for Smooth Workflow with Dense Per-Pixel Tasks	1
1.1 Installation	1
1.2 Examples	1
1.3 Motivation behind the library	2
1.4 Future Work	2
1.5 Documentation	2
1.6 Citation	2
2 Single-Task Training Example	3
2.1 Prepare Data	3
2.2 Training	3
3 Multi-Task Training Example	5
3.1 Prepare Data	5
3.2 Training	5
4 Code Documentation	7
4.1 densetorch.nn	7
4.2 densetorch.engine	9
4.3 densetorch.data	11
4.4 densetorch.misc	12
5 Indices and tables	15
Python Module Index	17
Index	19

DenseTorch: PyTorch Wrapper for Smooth Workflow with Dense Per-Pixel Tasks

[Build Status](#) [Docs Status](#)

This library aims to ease typical workflows involving dense per-pixel tasks in PyTorch. The progress in such tasks as semantic image segmentation and depth estimation have been significant over the last years, and in this library we provide an easy-to-setup environment for experimenting with given (or your own) models that reliably solve these tasks.

1.1 Installation

Python \geq 3.6.7 is supported.

```
git clone https://github.com/drsleep/densetorch.git
cd densetorch
pip install -e .
```

1.2 Examples

Currently, we provide several models for single-task and multi-task setups:

- `resnet` ResNet-18/34/50/101/152.
- `mobilenet-v2` MobileNet-v2.
- `xception-65` Xception-65.
- `deeplab-v3+` DeepLab-v3+.
- `lwrfr` Light-Weight RefineNet.
- `mtlwrfr` Multi-Task Light-Weight RefineNet.

Examples are given in the `examples/` directory. Note that the provided examples do not necessarily reproduce the results achieved in corresponding papers, rather their goal is to demonstrate what can be done using this library.

1.3 Motivation behind the library

As my everyday research is concerned with dense per-pixel tasks, I found myself oftentimes re-writing and updating (occasionally improving upon) my own code for each project. With the number of projects being on the rise recently, such an approach was no longer easy to manage. Hence, I decided to create a simple to use and simple to extend upon backbone (pun is not intended) structure, which I would be able to share with the community and, hopefully, ease the experience for others in the field.

1.4 Future Work

This library is still work-in-progress. More examples and more models will be added. Contributions are welcome.

1.5 Documentation

Is available [here](#).

1.6 Citation

If you found this library useful in your research, please consider citing

```
@misc{Nekrasov19,  
  author = {Nekrasov, Vladimir},  
  title = {DenseTorch},  
  year = {2019},  
  publisher = {GitHub},  
  journal = {GitHub repository},  
  howpublished = {\url{https://github.com/drsleep/densetorch}}  
}
```

Single-Task Training Example

In this example, we are going to train DeepLab-v3+ with the Xception-65 backbone for the task of semantic segmentation on NYUDv2.

2.1 Prepare Data

Considering that you successfully installed the `DenseTorch` package, the next step is to download the NYUDv2 dataset with segmentation and depth masks. The dataset can be downloaded by following the [link](#).

After downloading and unpacking the archive, create the `datasets` folder and link the `nyudv2` directory in the archive to the just created folder:

```
mkdir datasets
ln -s /path/to/nyudv2 ./datasets/
```

2.2 Training

Now you are ready to run the example script. To do so, simply execute `python train.py`. After it is finished, the best model will be stored in the corresponding `pth.tar` file.

Multi-Task Training Example

In this example, we are going to train Multi-Task Light-Weight RefineNet for joint semantic segmentation and depth estimation. Note that inference examples together with pre-trained weights can be found in the official [repository](#).

The hyperparameters set here are not the same as used in the corresponding paper, hence the results will differ. Please refer to the paper below for more information on the model and the training regime.

```
Real-Time Joint Semantic Segmentation and Depth Estimation Using Asymmetric_  
↪Annotations  
Vladimir Nekrasov, Thanuja Dharmasiri, Andrew Spek, Tom Drummond, Chunhua Shen, Ian_  
↪Reid  
In ICRA 2019
```

3.1 Prepare Data

Considering that you successfully installed the `DenseTorch` package, the next step is to download the NYUDv2 dataset with segmentation and depth masks. The dataset can be downloaded by following the [link](#).

After downloading and unpacking the archive, create the `datasets` folder and link the `nyudv2` directory in the archive to the just created folder:

```
mkdir datasets  
ln -s /path/to/nyudv2 ./datasets/
```

3.2 Training

Now you are ready to run the example script. To do so, simply execute `python train.py`. After it is finished, the best model will be stored in the corresponding `pth.tar` file. Note that it would be the model that improves upon the previous checkpoint both in terms of mean IoU (for segmentation) and linear RMSE (for depth estimation).

4.1 densenetorch.nn

The *nn* module implements a range of well-established encoders and decoders.

```
class densenetorch.nn.decoders.DLv3plus (input_sizes, num_classes, skip_size=48,  
                                         agg_size=256, rates=(6, 12, 18), **kwargs)
```

DeepLab-v3+ for Semantic Image Segmentation.

ASPP with decoder. Allows to have multiple skip-connections. More information about the model: <https://arxiv.org/abs/1802.02611>

Parameters

- **input_sizes** (*int*, or *list*) – number of channels for each input. Last value represents the input to ASPP, other values are for skip-connections.
- **num_classes** (*int*) – number of output channels.
- **skip_size** (*int*) – common filter size for skip-connections.
- **agg_size** (*int*) – common filter size.
- **rates** (*list of ints*) – dilation rates in the ASPP module.

forward (*xs*)

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 densenetorch.nn.decoders.LWRefineNet (input_sizes, collapse_ind, num_classes,  
                                           agg_size=256, n_crp=4)
```

Light-Weight RefineNet for Semantic Image Segmentation.

More information about the model: <https://arxiv.org/abs/1810.03272>

Parameters

- **input_sizes** (*int*, or *list*) – number of channels for each input.
- **collapse_ind** (*list*) – which input layers should be united together (via element-wise summation) before CRP.
- **num_classes** (*int*) – number of output channels.
- **agg_size** (*int*) – common filter size.
- **n_crp** (*int*) – number of CRP layers in a single CRP block.

forward (*xs*)

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 `densetorch.nn.decoders.MTLWRefineNet` (*input_sizes*, *collapse_ind*, *num_classes*,
agg_size=256, *n_crp=4*)

Multi-Task Light-Weight RefineNet for Dense per-pixel tasks.

More information about the model: <https://arxiv.org/abs/1809.04766>

Parameters

- **input_sizes** (*int*, or *list*) – number of channels for each input.
- **collapse_ind** (*list*) – which input layers should be united together (via element-wise summation) before CRP.
- **num_classes** (*int* or *list*) – number of output channels per each head.
- **agg_size** (*int*) – common filter size.
- **n_crp** (*int*) – number of CRP layers in a single CRP block.

forward (*xs*)

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.

`densetorch.nn.mobilenetv2.mobilenetv2` (*pretrained=True*, ***kwargs*)

Constructs the mobilenet-v2 network.

Parameters **pretrained** (*bool*) – whether to load pre-trained weights.

Returns *nn.Module* instance.

`densetorch.nn.resnet.resnet18` (*pretrained=False*, ***kwargs*)

Constructs the ResNet-18 model.

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

Returns *nn.Module* instance.

`densetorch.nn.resnet.resnet34` (*pretrained=False, **kwargs*)
Constructs the ResNet-34 model.

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

Returns *nn.Module* instance.

`densetorch.nn.resnet.resnet50` (*pretrained=False, **kwargs*)
Constructs the ResNet-50 model.

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

Returns *nn.Module* instance.

`densetorch.nn.resnet.resnet101` (*pretrained=False, **kwargs*)
Constructs the ResNet-101 model.

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

Returns *nn.Module* instance.

`densetorch.nn.resnet.resnet152` (*pretrained=False, **kwargs*)
Constructs the ResNet-152 model.

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

Returns *nn.Module* instance.

`densetorch.nn.xception.xception65` (*pretrained=False, **kwargs*)
Constructs the Xception-65 network.

Parameters `pretrained` (*bool*) – whether to load pre-trained weights.

Returns *nn.Module* instance.

4.2 densetorch.engine

The *engine* module contains metrics and losses typically used for the tasks of semantic segmentation and depth estimation. Also contains training and validation functions.

class `densetorch.engine.losses.InvHuberLoss` (*ignore_index=0*)
Inverse Huber Loss for depth estimation.

The setup is taken from <https://arxiv.org/abs/1606.00373>

Parameters `ignore_index` (*float*) – value to ignore in the target when computing the loss.

forward (*x, target*)

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 `densetorch.engine.metrics.MeanIoU` (*num_classes*)
Mean-IoU computational block for semantic segmentation.

Parameters `num_classes` (*int*) – number of classes to evaluate.

name
descriptor of the estimator.

Type str

class `densetorch.engine.metrics.RMSE` (*ignore_val=0*)
Root Mean Squared Error computational block for depth estimation.

Parameters `ignore_val` (*float*) – value to ignore in the target when computing the metric.

name
descriptor of the estimator.

Type str

`densetorch.engine.trainval.maybe_cast_target_to_long` (*target*)
Torch losses usually work on Long types

`densetorch.engine.trainval.train` (*model, opts, crits, dataloader, loss_coeffs=(1.0,), freeze_bn=False, grad_norm=0.0*)

Full Training Pipeline.

Supports multiple optimisers, multiple criteria, multiple losses, multiple outputs. Assumes that the `model.eval()` property has been set up properly before the function call, that the dataloader outputs have the correct type, that the model outputs do not require any post-processing bar the upsampling to the target size. Criteria, `loss_coeff`, and model's outputs all must have the same length, and correspond to the same keys as in the ordered dict of dataloader's sample.

Parameters

- **model** – PyTorch model object.
- **opts** – list of optimisers.
- **crits** – list of criterions.
- **dataloader** – iterable over samples. Each sample must contain *image* key and ≥ 1 optional keys.
- **loss_coeffs** – list of coefficients for each loss term.
- **freeze_bn** – whether to freeze batch norm parameters in the module.
- **grad_norm** – if > 0 , clip gradients' norm to this value.

`densetorch.engine.trainval.trainbal` (*model, dataloader*)

Full Training Pipeline with balanced model.

Assumes that the `model.eval()` property has been set up properly before the function call, that the dataloader outputs have the correct type, that the model outputs do not require any post-processing bar the upsampling to the target size.

Parameters

- **model** – PyTorch model object.
- **dataloader** – iterable over samples. Each sample must contain *image* key and ≥ 1 optional keys.

`densetorch.engine.trainval.validate` (*model, metrics, dataloader*)

Full Validation Pipeline.

Support multiple metrics (but 1 per modality), multiple outputs. Assumes that the dataloader outputs have the correct type, that the model outputs do not require any post-processing bar the upsampling to the target size. Metrics and model's outputs must have the same length, and correspond to the same keys as in the ordered dict of dataloader's sample.

Parameters

- **model** – PyTorch model object.
- **metrics** – list of metric classes. Each metric class must have update and val functions, and must have ‘name’ attribute.
- **dataloader** – iterable over samples. Each sample must contain *image* key and ≥ 1 optional keys.

4.3 densetorch.data

The *data* module implements datasets and relevant utilities used for data pre-processing. It supports multi-modal data.

class `densetorch.data.datasets.MMDataset` (*data_file*, *data_dir*, *line_to_paths_fn*, *masks_names*, *transform=None*)

Multi-Modality dataset.

Works with any datasets that contain image and any number of 2D-annotations.

Parameters

- **data_file** (*string*) – Path to the data file with annotations.
- **data_dir** (*string*) – Directory with all the images.
- **line_to_paths_fn** (*callable*) – function to convert a line of *data_file* into paths (*img_relpath*, *msk1_relpath*, *msk2_relpath*).
- **masks_names** (*list of strings*) – keys for each annotation mask (e.g., ‘segm’, ‘depth’). Must be in the same order as outputs of *line_to_paths_fn*
- **transform** (*callable, optional*) – Optional transform to be applied on a sample.

static read_image (*x*)

Simple image reader

Parameters *x* (*str*) – path to image.

Returns image as *np.array*.

class `densetorch.data.utils.Normalise` (*scale*, *mean*, *std*, *depth_scale=1.0*)

Normalise a tensor image with mean and standard deviation. Given mean: (R, G, B) and std: (R, G, B), will normalise each channel of the torch.*Tensor, i.e. $\text{channel} = (\text{scale} * \text{channel} - \text{mean}) / \text{std}$

Parameters

- **scale** (*float*) – Scaling constant.
- **mean** (*sequence*) – Sequence of means for R,G,B channels respectively.
- **std** (*sequence*) – Sequence of standard deviations for R,G,B channels respectively.
- **depth_scale** (*float*) – Depth divisor for depth annotations.

class `densetorch.data.utils.Pad` (*size*, *img_val*, *msk_vals*)

Pad image and mask to the desired size.

Parameters

- **size** (*int*) – minimum length/width.
- **img_val** (*array*) – image padding value.
- **msk_vals** (*list of ints*) – masks padding value.

class `densetorch.data.utils.RandomCrop` (*crop_size*)
Crop randomly the image in a sample.

Parameters `crop_size` (*int*) – Desired output size.

class `densetorch.data.utils.RandomMirror`
Randomly flip the image and the mask

class `densetorch.data.utils.ResizeAndScale` (*side*, *low_scale*, *high_scale*, *shorter=True*)
Resize shorter/longer side to a given value and randomly scale.

Parameters

- **side** (*int*) – shorter / longer side value.
- **low_scale** (*float*) – lower scaling bound.
- **high_scale** (*float*) – upper scaling bound.
- **shorter** (*bool*) – whether to resize shorter / longer side.

class `densetorch.data.utils.ToTensor`
Convert ndarrays in sample to Tensors.

`densetorch.data.utils.albumentations2densetorch` (*augmentation*)
Wrapper to use Albumentations within DenseTorch dataset.

Parameters `augmentation` – either a list of augmentations or a single augmentation

Returns A composition of augmentations

`densetorch.data.utils.denormalise` (*tensor_bchw*, *scale*, *mean_c*, *std_c*)
Reversed normalisation

Parameters

- **tensor_bchw** (*torch.tensor*) – 4D tensor of shape BxCxHxW
- **scale** (*float*) – scale value
- **mean_c** (*np.ndarray*) – mean array of shape (C,)
- **std_c** (*np.ndarray*) – standard deviation array of shape (C,)

Returns Un-normalised torch tensor.

`densetorch.data.utils.densetorch2torchvision` (*augmentation*)
Wrapper to use DenseTorch augmentations within torchvision dataset.

Parameters `augmentation` – either a list of augmentations or a single augmentation

Returns A composition of augmentations.

`densetorch.data.utils.get_loaders` (*train_batch_size*, *val_batch_size*, *train_set*, *val_set*,
num_stages=1, *num_workers=8*, *train_shuffle=True*,
val_shuffle=False, *train_pin_memory=False*,
val_pin_memory=False, *train_drop_last=False*,
val_drop_last=False)

Create train and val loaders

4.4 densetorch.misc

The *misc* module has various useful utilities.

class `densetorch.misc.utils.AverageMeter` (*momentum=0.99*)

Simple running average estimator.

Parameters `momentum` (*float*) – running average decay.

update (*val*)

Update running average given a new value.

The new running average estimate is given as a weighted combination of the previous estimate and the current value.

Parameters `val` (*float*) – new value

class `densetorch.misc.utils.Balancer` (*model, opts, crits, loss_coefs*)

Wrapper for balanced multi-GPU training.

When forward and backward passes are fused into a single `nn.Module` object, the multi-GPU consumption is distributed more equally across the GPUs.

Parameters

- **model** (*nn.Module*) – PyTorch module.
- **opts** (*list or single instance of torch.optim*) – optimisers.
- **crits** (*list or single instance of torch.nn or nn.Module*) – criterions.
- **loss_coefs** (*list of single instance of float*) – loss coefficients.

forward (*inp, targets=None*)

Forward and (optionally) backward pass.

When targets are provided, the backward pass is performed. Otherwise only the forward pass is done.

Parameters

- **inp** (*torch.tensor*) – input batch.
- **targets** (*None or torch.tensor*) – targets batch.

Returns Forward output if *targets=None*, else returns the loss value.

class `densetorch.misc.utils.Saver` (*args, ckpt_dir, best_val=0, condition=<function Saver.<lambda>>, save_interval=100, save_several_mode=<built-in function any>*)

Saver class for checkpointing the training progress.

maybe_load (*ckpt_path, keys_to_load*)

Loads existing checkpoint if exists.

Parameters

- **ckpt_path** (*str*) – path to the checkpoint.
- **keys_to_load** (*list of str*) – keys to load from the checkpoint.

Returns the epoch at which the checkpoint was saved.

maybe_save (*new_val, dict_to_save*)

Maybe save new checkpoint

`densetorch.misc.utils.broadcast` (*x, num_times*)

Given an element, broadcast it number of times and return a list. If it is already a list, only the first element will be copied.

Parameters

- **x** – input.
- **num_times** (*int*) – how many times to copy the element.

Returns list of length `num_times`.

`densetorch.misc.utils.compute_params(model)`

Compute the total number of parameters.

Parameters `model` (*nn.Module*) – PyTorch model.

Returns Total number of parameters - both trainable and non-trainable (*int*).

`densetorch.misc.utils.create_optim(optim_type, parameters, **kwargs)`

Initialise optimisers.

Parameters

- **optim_type** (*string*) – type of optimiser - either ‘SGD’ or ‘Adam’.
- **parameters** (*iterable*) – parameters to be optimised.

Returns An instance of `torch.optim`.

Raises `ValueError` if `optim_type` is not either of ‘SGD’ or ‘Adam’.

`densetorch.misc.utils.create_scheduler(scheduler_type, optim, **kwargs)`

Initialise schedulers.

Parameters

- **scheduler_type** (*string*) – type of scheduler – either ‘poly’ or ‘multistep’.
- **optim** (*torch.optim*) – optimiser to which the scheduler will be applied to.

Returns An instance of `torch.optim.lr_scheduler`

Raises `ValueError` if `scheduler_type` is not either of ‘poly’ or ‘multistep’.

`densetorch.misc.utils.ctime()`

Returns current timestamp in the format of hours-minutes-seconds.

`densetorch.misc.utils.get_args(func)`

Get function’s arguments.

Parameters `func` (*callable*) – input function.

Returns List of positional and keyword arguments.

`densetorch.misc.utils.make_list(x)`

Returns the given input as a list.

`densetorch.misc.utils.polyschedule(max_epochs, gamma=0.9)`

Poly-learning rate policy popularised by DeepLab-v2: <https://arxiv.org/abs/1606.00915>

Parameters

- **max_epochs** (*int*) – maximum number of epochs, at which the multiplier becomes zero.
- **gamma** (*float*) – decay factor.

Returns Callable that takes the current epoch as an argument and returns the learning rate multiplier.

`densetorch.misc.utils.set_seed(seed)`

Setting the random seed across `torch`, `numpy` and `random` libraries.

Parameters `seed` (*int*) – random seed value.

CHAPTER 5

Indices and tables

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

d

- `densetorch.data.datasets`, 11
- `densetorch.data.utils`, 11
- `densetorch.engine.losses`, 9
- `densetorch.engine.metrics`, 9
- `densetorch.engine.trainval`, 10
- `densetorch.misc.utils`, 12
- `densetorch.nn.decoders`, 7
- `densetorch.nn.mobilenetv2`, 8
- `densetorch.nn.resnet`, 8
- `densetorch.nn.xception`, 9

A

albumations2densetorch() (in module *densetorch.data.utils*), 12

AverageMeter (class in *densetorch.misc.utils*), 12

B

Balancer (class in *densetorch.misc.utils*), 13

broadcast() (in module *densetorch.misc.utils*), 13

C

compute_params() (in module *densetorch.misc.utils*), 14

create_optim() (in module *densetorch.misc.utils*), 14

create_scheduler() (in module *densetorch.misc.utils*), 14

ctime() (in module *densetorch.misc.utils*), 14

D

denormalise() (in module *densetorch.data.utils*), 12

densetorch.data.datasets (module), 11

densetorch.data.utils (module), 11

densetorch.engine.losses (module), 9

densetorch.engine.metrics (module), 9

densetorch.engine.trainval (module), 10

densetorch.misc.utils (module), 12

densetorch.nn.decoders (module), 7

densetorch.nn.mobilenetv2 (module), 8

densetorch.nn.resnet (module), 8

densetorch.nn.xception (module), 9

densetorch2torchvision() (in module *densetorch.data.utils*), 12

DLv3plus (class in *densetorch.nn.decoders*), 7

F

forward() (*densetorch.engine.losses.InvHuberLoss* method), 9

forward() (*densetorch.misc.utils.Balancer* method), 13

forward() (*densetorch.nn.decoders.DLv3plus* method), 7

forward() (*densetorch.nn.decoders.LWRefineNet* method), 8

forward() (*densetorch.nn.decoders.MTLWRefineNet* method), 8

G

get_args() (in module *densetorch.misc.utils*), 14

get_loaders() (in module *densetorch.data.utils*), 12

I

InvHuberLoss (class in *densetorch.engine.losses*), 9

L

LWRefineNet (class in *densetorch.nn.decoders*), 7

M

make_list() (in module *densetorch.misc.utils*), 14

maybe_cast_target_to_long() (in module *densetorch.engine.trainval*), 10

maybe_load() (*densetorch.misc.utils.Saver* method), 13

maybe_save() (*densetorch.misc.utils.Saver* method), 13

MeanIoU (class in *densetorch.engine.metrics*), 9

MMDataset (class in *densetorch.data.datasets*), 11

mobilenetv2() (in module *densetorch.nn.mobilenetv2*), 8

MTLWRefineNet (class in *densetorch.nn.decoders*), 8

N

name (*densetorch.engine.metrics.MeanIoU* attribute), 9

name (*densetorch.engine.metrics.RMSE* attribute), 10

Normalise (class in *densetorch.data.utils*), 11

P

Pad (class in *densetorch.data.utils*), 11

`polyschedule()` (in module `densetorch.misc.utils`),
14

R

`RandomCrop` (class in `densetorch.data.utils`), 11
`RandomMirror` (class in `densetorch.data.utils`), 12
`read_image()` (`densetorch.data.datasets.MMDataset`
static method), 11
`ResizeAndScale` (class in `densetorch.data.utils`), 12
`resnet101()` (in module `densetorch.nn.resnet`), 9
`resnet152()` (in module `densetorch.nn.resnet`), 9
`resnet18()` (in module `densetorch.nn.resnet`), 8
`resnet34()` (in module `densetorch.nn.resnet`), 9
`resnet50()` (in module `densetorch.nn.resnet`), 9
`RMSE` (class in `densetorch.engine.metrics`), 10

S

`Saver` (class in `densetorch.misc.utils`), 13
`set_seed()` (in module `densetorch.misc.utils`), 14

T

`ToTensor` (class in `densetorch.data.utils`), 12
`train()` (in module `densetorch.engine.trainval`), 10
`trainbal()` (in module `densetorch.engine.trainval`),
10

U

`update()` (`densetorch.misc.utils.AverageMeter`
method), 13

V

`validate()` (in module `densetorch.engine.trainval`),
10

X

`xception65()` (in module `densetorch.nn.xception`), 9