



Controlling the Chaos: Reproducible Deep Learning Experiments Using Lightning and Hydra

Refik Can MALLI



Orobix

- **We are 45 People**
 - 10 working in software development
 - 17 working in data science
- **Young team with average of 33 years old**
- **We are located at:**
 - Bergamo
 - Brescia
- **We are applying AI to**
 - Manufacturing
 - Healthcare
 - Games (Reinforcement Learning)



The Happy Data Scientist's Checklist

- Python Ecosystem Integrated Tools
- High-Level API for computations
- Hardware Acceleration
- Production Models
- Data Handling
- Distributed Training
- Logging, Tracking, Checkpointing
- Visualization, Performance Metrics
- Configuring Hyperparameters



PyTorch

- **Ease of use:** It offers a user-friendly interface and a simple and intuitive API
- **Dynamic computational graph:** It allows for flexible and efficient model building and automatic gradient computation.
- **Research focus:** It was developed by researchers for researchers and is widely used in the academic community.
- **Flexible deployment options:** TorchScript, ONNX, TorchServe, C++ Frontend



PyTorch

Data Loading

```
1 import torch
2 import torchvision
3 import torchvision.transforms as transforms
4
5 transform = transforms.Compose([transforms.ToTensor(), transforms.Normalize((0.5, 0.5, 0.5), (0.5, 0.5, 0.5))])
6
7 batch_size = 4
8
9 trainset = torchvision.datasets.CIFAR10(root="./data", train=True, download=True, transform=transform)
10 trainloader = torch.utils.data.DataLoader(trainset, batch_size=batch_size, shuffle=True, num_workers=2)
11
12 testset = torchvision.datasets.CIFAR10(root="./data", train=False, download=True, transform=transform)
13 testloader = torch.utils.data.DataLoader(testset, batch_size=batch_size, shuffle=False, num_workers=2)
14
15 classes = ("plane", "car", "bird", "cat", "deer", "dog", "frog", "horse", "ship", "truck")
16
```



PyTorch

Model Building

```
1 import torch
2 import torch.nn as nn
3 import torch.nn.functional as F
4 import torch.optim as optim
5
6 class Net(nn.Module):
7     def __init__(self):
8         super().__init__()
9         self.conv1 = nn.Conv2d(3, 6, 5)
10        self.pool = nn.MaxPool2d(2, 2)
11        self.conv2 = nn.Conv2d(6, 16, 5)
12        self.fc1 = nn.Linear(16 * 5 * 5, 120)
13        self.fc2 = nn.Linear(120, 84)
14        self.fc3 = nn.Linear(84, 10)
15
16    def forward(self, x):
17        x = self.pool(F.relu(self.conv1(x)))
18        x = self.pool(F.relu(self.conv2(x)))
19        x = torch.flatten(x, 1)
20        x = F.relu(self.fc1(x))
21        x = F.relu(self.fc2(x))
22        x = self.fc3(x)
23        return x
24
25 net = Net()
26 criterion = nn.CrossEntropyLoss()
27 optimizer = optim.SGD(net.parameters(), lr=0.001, momentum=0.9)
```

Training

```
1 for epoch in range(2):
2     running_loss = 0.0
3     for i, data in enumerate(trainloader):
4         # data is a list of [inputs, labels]
5         inputs, labels = data
6
7         # zero the parameter gradients
8         optimizer.zero_grad()
9
10        # forward + backward + optimize
11        outputs = net(inputs)
12        loss = criterion(outputs, labels)
13        loss.backward()
14        optimizer.step()
```

Saving

```
1 PATH = './cifar_net.pth'
2 torch.save(net.state_dict(), PATH)
3 torch.onnx.export(net, inputs, "cifar_net.onnx")
```



PyTorch Lightning

- **Extends PyTorch:** Built top of PyTorch.
- **Less Boilerplate Code:** Handles most of the engineering part.
- **Hardware Acceleration:** GPUs/TPUs/HPUs without code changes.
- **Easy Distributed Training:** Mostly single line setup for different distributed backends.
- **Ready-to-use Toolbox:** Callbacks, Loggers, Profilers and more...



PyTorch Lightning – Data Modules

```
1 import torch
2 from torchvision import datasets, transforms
3 from pytorch_lightning import LightningDataModule
4
5 class CIFAR10DataModule(LightningDataModule):
6     def __init__(self, batch_size=4):
7         super().__init__()
8         self.batch_size = batch_size
9         self.transform = transforms.Compose([
10             transforms.ToTensor(),
11             transforms.Normalize((0.5, 0.5, 0.5), (0.5, 0.5, 0.5))
12         ])
13
14     def prepare_data(self):
15         datasets.CIFAR10(root='./data', train=True, download=True)
16         datasets.CIFAR10(root='./data', train=False, download=True)
17
18     def setup(self, stage=None):
19         self.trainset = datasets.CIFAR10(root='./data', train=True, transform=self.transform)
20         self.testset = datasets.CIFAR10(root='./data', train=False, transform=self.transform)
21
22     def train_dataloader(self):
23         return torch.utils.data.DataLoader(self.trainset, batch_size=self.batch_size, shuffle=True, num_workers=2)
24
25     def val_dataloader(self):
26         return torch.utils.data.DataLoader(self.testset, batch_size=self.batch_size, shuffle=False, num_workers=2)
```

Runs on main process

Runs on every process

Used for training

Used for validation



PyTorch Lightning - Training

```
1 class CNNModel(LightningModule):
2     def __init__(self, learning_rate, momentum):
3         super().__init__()
4         self.save_hyperparameters()
5         self.conv1 = nn.Conv2d(3, 6, 5)
6         self.pool = nn.MaxPool2d(2, 2)
7         self.conv2 = nn.Conv2d(6, 16, 5)
8         self.fc1 = nn.Linear(16 * 5 * 5, 120)
9         self.fc2 = nn.Linear(120, 84)
10        self.fc3 = nn.Linear(84, 10)
11
12    def forward(self, x):
13        x = self.pool(F.relu(self.conv1(x)))
14        x = self.pool(F.relu(self.conv2(x)))
15        x = torch.flatten(x, 1)
16        x = F.relu(self.fc1(x))
17        x = F.relu(self.fc2(x))
18        x = self.fc3(x)
19        return x
20
21    def training_step(self, batch, batch_idx):
22        inputs, targets = batch
23        outputs = self(inputs)
24        loss = F.cross_entropy(outputs, targets)
25        self.log("train_loss", loss)
26        return loss
27
28    def configure_optimizers(self):
29        optimizer = torch.optim.SGD(
30            self.parameters(),
31            lr=self.hparams.learning_rate,
32            momentum=self.hparams.momentum,
33        )
34        scheduler = StepLR(optimizer, step_size=30, gamma=0.1)
35        return [optimizer], [scheduler]
```

Single Step

Optimizer and Scheduler

```
1 if __name__ == "__main__":
2     parser = argparse.ArgumentParser()
3     parser.add_argument("--batch_size", type=int, default=4)
4     parser.add_argument("--learning_rate", type=float, default=0.001)
5     parser.add_argument("--momentum", type=float, default=0.9)
6     parser.add_argument("--max_epochs", type=int, default=2)
7     parser.add_argument("--gpus", type=int, default=None)
8     parser.add_argument("--accelerator", type=str, default=None)
9     parser.add_argument("--patience", type=int, default=3)
10    parser.add_argument("--min_delta", type=float, default=0.01)
11    parser.add_argument("--log_dir", type=str, default="./logs")
12    args = parser.parse_args()
13
14    model = CNNModel(args.learning_rate, args.momentum)
15    data_module = CIFAR10DataModule(args.batch_size)
16    early_stopping = callbacks.EarlyStopping(
17        monitor="train_loss",
18        min_delta=args.min_delta,
19        patience=args.patience,
20        verbose=True,
21    )
22    logger = CSVLogger(args.log_dir, name="cifar10_logs")
23    trainer = Trainer(
24        max_epochs=args.max_epochs,
25        gpus=args.gpus,
26        accelerator=args.accelerator,
27        callbacks=[early_stopping],
28        logger=logger,
29    )
30    trainer.fit(model, datamodule=data_module)
```



Where the Chaos begins

```
1 python main.py \  
2     --batch_size 8 \  
3     --learning_rate 0.01 \  
4     --momentum 0.8 \  
5     --max_epochs 10 \  
6     --gpus 1 \  
7     --accelerator 'gpu' \  
8     --patience 3 \  
9     --min_delta 0.01 \  
10    --log_dir './logs'
```



Where the Chaos begins

```
1 python main.py \  
2     --batch_size 8 \  
3     --learning_rate 0.01 \  
4     --momentum 0.8 \  
5     --max_epochs 10 \  
6     --gpus 1 \  
7     --accelerator 'gpu' \  
8     --patience 3 \  
9     --min_delta 0.01 \  
10    --log_dir './logs'
```

What if we want to:

- Add another model?
- Check if augmentation helps?
- Try different hyperparameters?
- Integrate different optimizer?
- Try with different dataset?
- Run many different configurations?



Where the Chaos begins

```
1 parser = argparse.ArgumentParser(description='PyTorch ImageNet Training')
2 parser.add_argument('data', metavar='DIR', nargs='?', default='imagenet',
3                     help='path to dataset (default: imagenet)')
4 parser.add_argument('-a', '--arch', metavar='ARCH', default='resnet18',
5                     choices=model_names,
6                     help='model architecture: ' +
7                         ' | '.join(model_names) +
8                         ' (default: resnet18)')
9 parser.add_argument('-j', '--workers', default=4, type=int, metavar='N',
10                    help='number of data loading workers (default: 4)')
11 parser.add_argument('--epochs', default=90, type=int, metavar='N',
12                    help='number of total epochs to run')
13 parser.add_argument('--start-epoch', default=0, type=int, metavar='N',
14                    help='manual epoch number (useful on restarts)')
15 parser.add_argument('-b', '--batch-size', default=256, type=int,
16                    metavar='N',
17                    help='mini-batch size (default: 256), this is the total '
18                        'batch size of all GPUs on the current node when '
19                        'using Data Parallel or Distributed Data Parallel')
20 parser.add_argument('--lr', '--learning-rate', default=0.1, type=float,
21                    metavar='LR', help='initial learning rate', dest='lr')
22 parser.add_argument('--momentum', default=0.9, type=float, metavar='M',
23                    help='momentum')
24 parser.add_argument('--wd', '--weight-decay', default=1e-4, type=float,
25                    metavar='W', help='weight decay (default: 1e-4)',
26                    dest='weight_decay')
27 parser.add_argument('-p', '--print-freq', default=10, type=int,
28                    metavar='N', help='print frequency (default: 10)')
29 parser.add_argument('--resume', default='', type=str, metavar='PATH',
30                    help='path to latest checkpoint (default: none)')
31 parser.add_argument('-e', '--evaluate', dest='evaluate', action='store_true',
32                    help='evaluate model on validation set')
33 parser.add_argument('--pretrained', dest='pretrained', action='store_true',
34                    help='use pre-trained model')
35 parser.add_argument('--world-size', default=-1, type=int,
36                    help='number of nodes for distributed training')
37 parser.add_argument('--rank', default=-1, type=int,
38                    help='node rank for distributed training')
39 parser.add_argument('--dist-url', default='tcp://224.66.41.62:23456', type=str,
40                    help='url used to set up distributed training')
41 parser.add_argument('--dist-backend', default='nccl', type=str,
42                    help='distributed backend')
43 parser.add_argument('--seed', default=None, type=int,
44                    help='seed for initializing training. ')
45 parser.add_argument('--gpu', default=None, type=int,
46                    help='GPU id to use.')
47 parser.add_argument('--multiprocessing-distributed', action='store_true',
48                    help='Use multi-processing distributed training to launch '
49                        'N processes per node, which has N GPUs. This is the '
50                        'fastest way to use PyTorch for either single node or '
51                        'multi node data parallel training')
52 parser.add_argument('--dummy', action='store_true', help="use fake data to benchmark")
```

- When the project becomes complex, argument list gets out of control.
- Classes should be instantiated based on arguments. (string -> class)
- Some arguments might be set up in a group.
- **Dataclasses** can help but do not solve all problems.

Source: <https://github.com/pytorch/examples/blob/main/imagenet/main.py>



Hydra

- **Configuration composition:** Structure your configurations into hierarchical YAML files. Mix and match configuration files.
- **Defaults and overrides:** easy management of default values and allows for overriding them from the command line.
- **Dynamic configurations:** Setting variables from environment variables, other configuration files or sub-folders.
- **Plugin system:** Launchers (Joblib, Submitit, Sequential), Sweepers (Optuna, Nevergrad)



Hydra

```
1 # conf/config.yaml
2 defaults:
3   - model: default
4   - data: default
5   - trainer: default
6
7 # trainer/default.yaml
8 max_epochs: 10
9 gpus: 1
10 accelerator: "gpu"
11 patience: 3
12 min_delta: 0.01
13 log_dir: "./logs"
14
15 # trainer/cpu.yaml
16 max_epochs: 2
17 accelerator: "cpu"
18 patience: 3
19 min_delta: 0.01
20 log_dir: "./logs"
21
22 # model/default.yaml
23 learning_rate: 0.001
24 momentum: 0.9
25
26 # data/default.yaml
27 batch_size: 32
```

```
1 @hydra.main(config_path="conf", config_name="config")
2 def cli_main(cfg: DictConfig) -> None:
3     model = CNNModel(cfg.model.learning_rate, cfg.model.momentum)
4     data_module = CIFAR10DataModule(cfg.data.batch_size)
5     early_stopping = callbacks.EarlyStopping(
6         monitor='train_loss',
7         min_delta=cfg.trainer.min_delta,
8         patience=cfg.trainer.patience,
9         verbose=True
10    )
11    logger = CSVLogger(cfg.trainer.log_dir, name="cifar10_logs")
12    trainer = Trainer(max_epochs=cfg.trainer.max_epochs,
13                     gpus=cfg.trainer.gpus,
14                     accelerator=cfg.trainer.accelerator,
15                     callbacks=[early_stopping],
16                     logger=logger)
17    trainer.fit(model, datamodule=data_module)
18
19 if __name__ == '__main__':
20     cli_main()
```



Hydra

```
1 # conf/config.yaml
2 defaults:
3   - model: default
4   - data: default
5   - trainer: default
6
7 # trainer/default.yaml
8 max_epochs: 10
9 gpus: 1
10 accelerator: "gpu"
11 patience: 3
12 min_delta: 0.01
13 log_dir: "./logs"
14
15 # trainer/cpu.yaml
16 max_epochs: 2
17 accelerator: "cpu"
18 patience: 3
19 min_delta: 0.01
20 log_dir: "./logs"
21
22 # model/default.yaml
23 learning_rate: 0.001
24 momentum: 0.9
25
26 # data/default.yaml
27 batch_size: 32
```

```
1 @hydra.main(config_path="conf", config_name="config")
2 def cli_main(cfg: DictConfig) -> None:
3     model = CNNModel(cfg.model.learning_rate, cfg.model.momentum)
4     data_module = CIFAR10DataModule(cfg.data.batch_size)
5     early_stopping = callbacks.EarlyStopping(
6         monitor='train_loss',
7         min_delta=cfg.trainer.min_delta,
8         patience=cfg.trainer.patience,
9         verbose=True
10    )
11    logger = CSVLogger(cfg.trainer.log_dir, name="cifar10_logs")
12    trainer = Trainer(max_epochs=cfg.trainer.max_epochs,
13                    gpus=cfg.trainer.gpus,
14                    accelerator=cfg.trainer.accelerator,
15                    callbacks=[early_stopping],
16                    logger=logger)
17    trainer.fit(model, datamodule=data_module)
18
19 if __name__ == '__main__':
20     cli_main()
```

```
1 python main.py trainer.max_epochs=5
2 python main.py trainer=default,cpu model.learning_rate=0.01,0.1,0.001 --multirun
```



Hydra

Instantiate the class

```
1 # optimizer/adam.yaml
2 _target_: torch.optim.Adam
3 lr: 0.001
4 betas: [0.9, 0.999]
5 eps: 1e-08
6 weight_decay: 0
```

```
1 param_list = model.parameters()
2 optimizer = hydra.utils.instantiate(config.optimizer, param_list)
```

Resolve configuration

```
1 def as_tuple(*args: Any) -> Tuple[Any, ...]:
2     """Resolves a list of arguments to a tuple."""
3     return tuple(args)
4 OmegaConf.register_new_resolver("as_tuple", as_tuple)
```

Access Environment Variables

```
1 tracking_uri: ${oc.env:MLFLOW_TRACKING_URI}
2 data_path: ${oc.env:HOME}/.quadra/datasets/MNIST
```



Checklist Revisited ...

- **Python Ecosystem Integrated Tools**
- **High-Level API for computations**
- **Hardware Acceleration**
- **Production Models**
- **Data Handling**
- **Distributed Training**
- **Losses, Performance Metrics**
- **Logging, Tracking, Checkpointing**
- **Configuring Hyperparameters**



Checklist Revisited ...

- Python Ecosystem Integrated Tools
- High-Level API for computations
- Hardware Acceleration
- Production Models
- Data Handling
- Distributed Training
- Losses, Performance Metrics
- Logging, Tracking, Checkpointing
- Configuring Hyperparameters
- Using other model repositories
- Implementing different Tasks
- Reusing submodules
- Reproducibility
- Visualizing Model Predictions
- Sharing Experiment Configurations



Checklist Revisited ...

- Python Ecosystem Integrated Tools
- High-Level API for computations
- Hardware Acceleration
- Production Models
- Data Handling
- Distributed Training
- Losses, Performance Metrics
- Logging, Tracking, Checkpointing
- Configuring Hyperparameters
- Using other model repositories
- Implementing different Tasks
- Reusing submodules
- Reproducibility
- Visualizing Model Predictions
- Sharing Experiment Configurations



Quadra



Configuration Folders

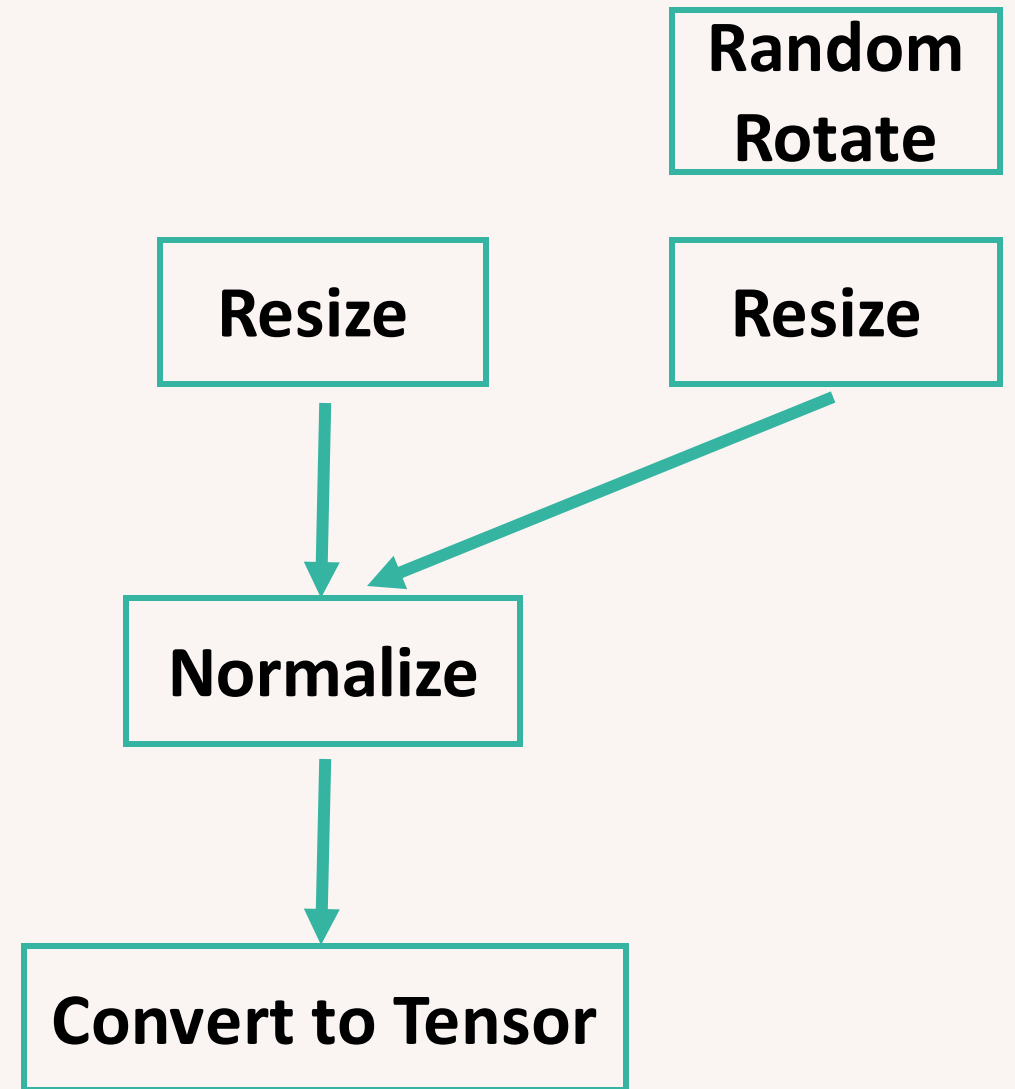
```
1 configs/  
2 |— backbone  
3 |— callbacks  
4 |— core  
5 |— datamodule  
6 |— experiment  
7 |— hydra  
8 |— logger  
9 |— loss  
10 |— model  
11 |— optimizer  
12 |— scheduler  
13 |— task  
14 |— trainer  
15 |— transforms
```

- **Divided into categories**
- **Contains part of the configuration**
- **Each folder can have default configuration**
- **Default configs can be extended by adding new YAML files**



Transformations

```
1 configs/  
2 |— backbone  
3 |— callbacks  
4 |— core  
5 |— datamodule  
6 |— experiment  
7 |— hydra  
8 |— logger  
9 |— loss  
10 |— model  
11 |— optimizer  
12 |— scheduler  
13 |— task  
14 |— trainer  
15 |— transforms
```



Transformations

```
1 configs/  
2 |— backbone  
3 |— callbacks  
4 |— core  
5 |— datamodule  
6 |— experiment  
7 |— hydra  
8 |— logger  
9 |— loss  
10 |— model  
11 |— optimizer  
12 |— scheduler  
13 |— task  
14 |— trainer  
15 |— transforms
```

```
1 input_height: 256  
2 input_width: 256  
3 mean: [0.485, 0.456, 0.406]  
4 std: [0.229, 0.224, 0.225]  
5  
6 normalize:  
7   _target_: albumentations.Compose  
8   transforms:  
9     - _target_: albumentations.Normalize  
10     mean: ${transforms.mean}  
11     std: ${transforms.std}  
12     always_apply: True  
13     - _target_: albumentations.pytorch.ToTensorV2  
14     always_apply: True  
15  
16 resize:  
17   _target_: albumentations.Resize  
18   height: ${transforms.input_height}  
19   width: ${transforms.input_width}  
20   interpolation: 2  
21   always_apply: True  
22 standard_transform:  
23   _target_: albumentations.Compose  
24   transforms:  
25     - ${transforms.resize}  
26     - ${transforms.normalize}  
27  
28 train_transform: ${transforms.standard_transform}  
29 val_transform: ${transforms.standard_transform}  
30 test_transform: ${transforms.standard_transform}
```



Datamodules

```
1 configs/  
2 |— backbone  
3 |— callbacks  
4 |— core  
5 |— datamodule  
6 |— experiment  
7 |— hydra  
8 |— logger  
9 |— loss  
10 |— model  
11 |— optimizer  
12 |— scheduler  
13 |— task  
14 |— trainer  
15 |— transforms
```

```
1 _target_: quadra.datamodules.classification.ClassificationDataModule  
2 data_path: ???  
3 exclude_filter: ["unwanted_class"]  
4 seed: ${core.seed}  
5 num_workers: 8  
6 batch_size: 16  
7 val_size: 0.2  
8 train_transform: ${transforms.train_transform}  
9 test_transform: ${transforms.test_transform}  
10 val_transform: ${transforms.val_transform}  
11 dataset:  
12   _target_: hydra.utils.get_method  
13   path: quadra.datasets.classification.ClassificationDataset
```

Backbone Modules

```
1  configs/  
2  |— backbone  
3  |— callbacks  
4  |— core  
5  |— datamodule  
6  |— experiment  
7  |— hydra  
8  |— logger  
9  |— loss  
10 |— model  
11 |— optimizer  
12 |— scheduler  
13 |— task  
14 |— trainer  
15 |— transforms
```

```
1  model:  
2    _target_: quadra.models.classification.TorchHubNetworkBuilder  
3    repo_or_dir: facebookresearch/dino:main  
4    model_name: dino_vitb8  
5    pretrained: true  
6    freeze: false  
7  metadata:  
8    input_size: 224  
9    output_dim: 768  
10   patch_size: 8  
11   nb_heads: 12
```

```
1  model:  
2    _target_: quadra.models.classification.TimmNetworkBuilder  
3    model_name: resnet50  
4    pretrained: true  
5    freeze: false  
6  metadata:  
7    input_size: 224  
8    output_dim: 2048
```



Models

```
1 configs/  
2 |— backbone  
3 |— callbacks  
4 |— core  
5 |— datamodule  
6 |— experiment  
7 |— hydra  
8 |— logger  
9 |— loss  
10 |— model  
11 |— optimizer  
12 |— scheduler  
13 |— task  
14 |— trainer  
15 |— transforms
```

```
1 model: ${backbone.model}  
2 num_classes: ???  
3 pre_classifier: null  
4 classifier:  
5   _target_: torch.nn.Linear  
6   in_features: ${backbone.metadata.output_dim}  
7   out_features: ${model.num_classes}  
8 module:  
9   _target_: quadra.modules.classification.ClassificationModule  
10  lr_scheduler_interval: "epoch"  
11  criterion: ${loss}  
12  gradcam: true
```



Callbacks

```
1 configs/  
2 |— backbone  
3 |— callbacks  
4 |— core  
5 |— datamodule  
6 |— experiment  
7 |— hydra  
8 |— logger  
9 |— loss  
10 |— model  
11 |— optimizer  
12 |— scheduler  
13 |— task  
14 |— trainer  
15 |— transforms
```

```
1 model_checkpoint:  
2   _target_: pytorch_lightning.callbacks.ModelCheckpoint  
3   monitor: "val_loss" # name of the logged metric which determines when model is improving  
4   mode: "min" # can be "max" or "min"  
5   save_top_k: 1 # save k best models (determined by above metric)  
6   save_last: True # additionally always save model from last epoch  
7   verbose: False  
8   dirpath: "checkpoints/"  
9   filename: "epoch_{epoch:03d}"  
10  auto_insert_metric_name: False  
11  
12 log_gradients:  
13   _target_: quadra.callbacks.mlflow.LogGradients  
14   norm: 2  
15 lr_monitor:  
16   _target_: pytorch_lightning.callbacks.LearningRateMonitor  
17   logging_interval: "epoch"  
18 progress_bar:  
19   _target_: pytorch_lightning.callbacks.progress.TQDMProgressBar  
20 lightning_trainer_setup:  
21   _target_: quadra.callbacks.lightning.LightningTrainerBaseSetup  
22   log_every_n_steps: 1
```



Hydra Settings

```
1 configs/  
2 |— backbone  
3 |— callbacks  
4 |— core  
5 |— datamodule  
6 |— experiment  
7 |— hydra  
8 |— logger  
9 |— loss  
10 |— model  
11 |— optimizer  
12 |— scheduler  
13 |— task  
14 |— trainer  
15 |— transforms
```

```
1 hydra:  
2   run:  
3     dir: logs/runs/${core.name}/${now:%Y-%m-%d_%H-%M-%S}  
4   sweep:  
5     dir: logs/multiruns/${core.name}/${now:%Y-%m-%d %H-%M-%S}  
6     subdir: ${multirun_subdir_beautify:${hydra.job.override_dirname}}  
7   job:  
8     chdir: true
```

```
1 def multirun_subdir_beautify(subdir: str) -> str:  
2     hydra_cfg = HydraConfig.get()  
3     if hydra_cfg.mode is None or hydra_cfg.mode.name == "RUN":  
4         return subdir  
5     subdir_list = subdir.replace("/", "|").split(",")  
6     subdir = ",".join([x.split("=")[1].replace(" ", "") for x in subdir_list])  
7     return subdir  
8 OmegaConf.register_new_resolver("multirun_subdir_beautify", multirun_subdir_beautify)
```

```
1 quadra experiment=anomaly/padim trainer.batch_size=32,64 --multirun  
2  
3 logs/multiruns/myexperiment/2023-05-28_12-40-00/anomaly|padim,32  
4 logs/multiruns/myexperiment/2023-05-28_12-40-00/anomaly|padim,64
```



Tasks

```
1 configs/  
2 |— backbone  
3 |— callbacks  
4 |— core  
5 |— datamodule  
6 |— experiment  
7 |— hydra  
8 |— logger  
9 |— loss  
10 |— model  
11 |— optimizer  
12 |— scheduler  
13 |— task  
14 |— trainer  
15 |— transforms
```

Available Tasks:

- Classification
- Segmentation
- Self-supervised Learning
- Anomaly Detection

```
1 _target_: quadra.tasks.Classification  
2 export_type: [torchscript]  
3 lr_multiplier: null  
4 output:  
5   example: false  
6 report: false  
7 run_test: true
```



Tasks

```
1 configs/  
2 |— backbone  
3 |— callbacks  
4 |— core  
5 |— datamodule  
6 |— experiment  
7 |— hydra  
8 |— logger  
9 |— loss  
10 |— model  
11 |— optimizer  
12 |— scheduler  
13 |— task  
14 |— trainer  
15 |— transforms
```

Task :

- Prepares modules
- Trains
- Tests
- Exports the model
- Evaluates the exported model
- Create performance report
- Visualize results



Experiments

```
1 configs/  
2 |— backbone  
3 |— callbacks  
4 |— core  
5 |— datamodule  
6 |— experiment  
7 |— hydra  
8 |— logger  
9 |— loss  
10 |— model  
11 |— optimizer  
12 |— scheduler  
13 |— task  
14 |— trainer  
15 |— transforms
```



```
1 # @package _global_  
2 defaults:  
3   - override /backbone: resnet18  
4   - override /datamodule: generic/imagenette/classification/base  
5   - override /loss: cross_entropy  
6   - override /model: classification  
7   - override /optimizer: adam  
8   - override /task: classification  
9   - override /scheduler: rop  
10  - override /transforms: default_resize  
11  
12 core:  
13   tag: "run"  
14   name: classification_imagenette_${trainer.max_epochs}  
15  
16 trainer:  
17   max_epochs: 20  
18  
19 model:  
20   num_classes: 10  
21  
22 logger:  
23   mlflow:  
24     experiment_name: imagenette_classification  
25     run_name: ${core.name}
```



Running Experiments

Classification

```
1 quadra experiment=generic/imagenette/classification/default
```



Running Experiments

Segmentation

```
1 quadra experiment=generic/oxford_pet/segmentation/smp
```



Running Experiments

Anomaly Detection


```
1 quadra experiment=generic/mnist/anomaly/padim
```

F1 threshold: 0.5, Mask_max: 0.616, Anomaly_score: 0.616

Image



Ground Truth



Predicted Heat Map



Predicted Mask



Segmentation Result



Running Experiments

Classification

```
1 quadra experiment=generic/imagenette/classification/default
```

Segmentation

```
1 quadra experiment=generic/oxford_pet/segmentation/smp
```

SSL Self-supervised Learning

```
1 quadra experiment=generic/imagenette/ssl/byol
```

Anomaly Detection

```
1 quadra experiment=generic/mnist/anomaly/padim
```

**Data scientists while
their model is training**



@reddit/ProgrammerHumor

Reproducibility

```
1 logs/.../experiment_name/  
2 |— checkpoints  
3 |   |— ...  
4 |— config_resolved.yaml  
5 |— config_tree.txt  
6 |— data  
7 |   |— datamodule.pkl  
8 |— deployment_model  
9 |   |— ...  
10 |— main.log  
11 |— .hydra  
12 |   |— config.yaml  
13 |   |— overrides.yaml  
14 |— task_specific_outputs...
```

Experiment Folder :

- All checkpoints saved during training
- Configuration file to reload experiment
- Datamodule metadata such as splits
- Git Commit Hash
- Exported model for deployment
- Hydra settings saved
- Images, results, metrics, all data

Quadra



- **Fast Experimenting**
- **Configuration sharing among the developers**
- **Self-contained tracking, saving system with open-source tools**
- **Code reusing for different tasks**
- **Standardized results and export scheme for easy comparison**
- **Future:**
 - Hyperparameter Optimization
 - More tasks, models, configurations
 - Cloud integration




<https://github.com/orobix/quadra>




Shout-Out

- Lorenzo Mammana
- Alessandro Polidori
- Federico Belotti
- Silvia Bianchetti
- Lisa Lozza
- Luca Antiga

TALK 

**ChatGPT and
Minecraft's diamonds
playing agents:
Reinforcement
Learning in a nutshell
with sheeprl**

Intermediate, English

Federico Belotti 

SheepRL - Distributed Reinforcement Learning accelerated by Lightning Fabric



<https://github.com/Eclectic-Sheep/sheeprl>



Thank you!

