

PyTorch Lightning

The key pieces of PyTorch Lightning, which strips away training boilerplate and organizes your model, loops, and logging around a Trainer.

The LightningModule

One class that bundles model, data flow, and training logic.

PURPOSE	COMMAND	RESULT
Subclass the base module.	<pre>class LitModel(L.LightningModule):</pre>	
Build layers in the constructor.	<pre>def __init__(self): super().__init__() self.net = nn.Sequential(...)</pre>	
Save constructor args as hparams.	<pre>self.save_hyperparameters()</pre>	
Define the forward pass.	<pre>def forward(self, x): return self.net(x)</pre>	
Instantiate the model.	<pre>model = LitModel(lr=1e-3)</pre>	

a LightningModule is a torch.nn.Module with fixed slots that say what to do, not how to loop

Step Hooks

Lightning calls these per-batch; you just return the loss.

PURPOSE	COMMAND	RESULT
Train on one batch.	<pre>def training_step(self, batch, batch_idx): loss = F.cross_entropy(self(x), y);return loss</pre>	
Validate on one batch.	<pre>def validation_step(self, batch, batch_idx): self.log("val_loss", loss)</pre>	
Test on one batch.	<pre>def test_step(self, batch, batch_idx):</pre>	
Predict on one batch.	<pre>def predict_step(...): return self(batch)</pre>	
Hook into the epoch end.	<pre>def on_train_epoch_end(self):</pre>	
The signature contract.	<pre>(self, batch, batch_idx)</pre>	

training_step returns the loss (Lightning calls backward + step); other steps log metrics or return predictions

Optimizers & Schedulers

Return optimizers from one method; Lightning steps them.

PURPOSE	COMMAND	RESULT
Configure one optimizer.	<pre>def configure_optimizers(self): return torch.optim.Adam(self.parameters())</pre>	
Add an LR scheduler.	<pre>return {"optimizer": opt, "lr_scheduler": sched}</pre>	
Scheduler step interval.	<pre>{"scheduler": s, "interval": "epoch", "monitor": ..}</pre>	
Lightning runs the loop.	<pre>loss.backward(); opt.step() # automatic</pre>	
Multiple optimizers (manual).	<pre>self.automatic_optimization = False</pre>	
Clip gradients.	<pre>L.Trainer(gradient_clip_val=1.0)</pre>	

by default Lightning runs zero_grad, backward, step for you; clipping and accumulation are set on the Trainer

The Trainer







The engine that owns the loop and the hardware.

PURPOSE	COMMAND	RESULT
Build the trainer.	<pre>trainer = L.Trainer(max_epochs=10)</pre>	
Pick hardware automatically.	<pre>L.Trainer(accelerator="auto", devices="auto")</pre>	
Run the full fit loop.	<pre>trainer.fit(model, train_dl, val_dl)</pre>	
Evaluate on a held-out set.	<pre>trainer.test(model, test_dl)</pre>	
Generate predictions.	<pre>preds = trainer.predict(model, dl)</pre>	
Smoke-test the whole pipeline.	<pre>L.Trainer(fast_dev_run=True)</pre>	

device placement, precision, and distribution live in the Trainer, so the same module runs anywhere

Logging Metrics







Call `self.log`; Lightning aggregates and routes it.

PURPOSE	COMMAND	RESULT
Log a scalar.	<code>self.log("train_loss", loss)</code>	
Show it in the progress bar.	<code>self.log("val_acc", acc, prog_bar=True)</code>	
Control step vs epoch reduce.	<code>self.log(.., on_step=True, on_epoch=True)</code>	
Log several at once.	<code>self.log_dict({"acc": acc, "f1": f1})</code>	
Pick a backend logger.	<code>L.Trainer(logger=CSVLogger("logs"))</code>	
Track a real metric object.	<code>self.acc = torchmetrics.Accuracy(...)</code>	

`self.log` aggregates across batches and devices; a `torchmetrics` object accumulates the correct value per epoch

Callbacks

Plug behavior into the loop without touching the model.

PURPOSE	COMMAND	RESULT
Save the best checkpoint.	<code>ModelCheckpoint(monitor=.., mode="min", save_top_k=1)</code>	
Stop when it plateaus.	<code>EarlyStopping(monitor="val_loss", patience=3)</code>	
Log the learning rate.	<code>LearningRateMonitor(logging_interval="step")</code>	
Attach them to the trainer.	<code>L.Trainer(callbacks=[ckpt, es, lr])</code>	
Resume from a checkpoint.	<code>trainer.fit(model, dl, ckpt_path="last.ckpt")</code>	
Write your own.	<code>class MyCb(L.Callback): on_train_epoch_end(...)</code>	

attach callbacks as a list to `Trainer(callbacks=[...])`; `ckpt_path` resumes optimizer state too

DataModule

Package download, split, and dataloaders in one class.

PURPOSE	COMMAND	RESULT
Subclass the data base.	<code>class DataMod(L.LightningDataModule):</code>	
Download once, on rank 0.	<code>def prepare_data(self):</code>	
Build splits per process.	<code>def setup(self, stage): self.train, self.val = random_split(...)</code>	
Expose the train loader.	<code>def train_dataloader(self): DataLoader(bs=32)</code>	
Add val / test / predict loaders.	<code>def val_dataloader(self): ... (test_, predict_)</code>	
Hand it to the trainer.	<code>trainer.fit(model, datamodule=dm)</code>	

prepare_data downloads once; setup(stage) splits per process; pass datamodule=dm instead of loose loaders

Save, Load & Export

Checkpoints capture everything; export for serving.

PURPOSE	COMMAND	RESULT
Save a checkpoint by hand.	<code>trainer.save_checkpoint("model.ckpt")</code>	
Reload weights + hparams.	<code>LitModel.load_from_checkpoint("model.ckpt")</code>	
Override hparams on load.	<code>load_from_checkpoint("model.ckpt", lr=1e-4)</code>	
Get just the weights dict.	<code>torch.load("model.ckpt", weights_only=False)["state_dict"]</code>	
Export to ONNX.	<code>model.to_onnx("model.onnx", sample)</code>	
Export the graph (TorchScript-free).	<code>torch.export.export(model, (sample,))</code>	

a checkpoint stores weights + hparams + optimizer + epoch; to_torchscript is deprecated, prefer torch.export