

Create Tensors

Spin up tensors from data, ranges, and random noise.

PURPOSE	COMMAND	RESULT
From a Python list.	<code>torch.tensor([[1, 2], [3, 4]])</code>	
Filled with zeros / ones.	<code>torch.zeros(2,3)</code> <code>torch.ones(2,3)</code>	
Evenly spaced range.	<code>torch.arange(0, 10, 2)</code>	
Random normal noise.	<code>torch.randn(2, 2)</code>	
Bridge from NumPy.	<code>torch.from_numpy(arr)</code>	
Match shape of another.	<code>torch.zeros_like(x)</code>	

a tensor is PyTorch's n-dimensional array; dtype and shape are inferred at creation

Shapes, Indexing & Ops







Reshape, slice, and combine tensors without copying when you can.

PURPOSE	COMMAND	RESULT
Inspect shape / dtype.	<code>x.shape</code> <code>x.dtype</code> <code>x.ndim</code>	
Reshape / flatten.	<code>x.reshape(2, 6)</code> <code>x.view(-1)</code>	
Add / drop a dim.	<code>x.unsqueeze(0)</code> <code>x.squeeze()</code>	
Transpose / permute axes.	<code>x.permute(1, 0)</code>	
Index, slice, mask.	<code>x[:, 1]</code> <code>x[x > 5]</code>	
Concatenate / matrix multiply.	<code>torch.cat([a, b])</code> <code>a @ b.T</code>	

reshape/view reinterpret the same data; @ does the matrix multiply at the heart of every layer

Autograd (gradients)



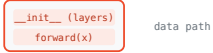



Tensors that track history; `.backward()` fills in the gradients.

PURPOSE	COMMAND	RESULT
Track gradients.	<code>torch.tensor([2.], requires_grad=True)</code>	
Build a computation.	<code>y = w * x + b</code>	
Compute gradients.	<code>y.backward()</code>	
Read a gradient.	<code>w.grad</code>	
Stop tracking (eval / infer).	<code>with torch.no_grad(): ...</code>	
Detach a branch.	<code>z.detach()</code>	

`requires_grad=True` records every op; `.backward()` walks the graph in reverse, filling each leaf's `.grad`

Build a Model (nn.Module)


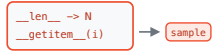




Layers are parameters with a forward; stack them into a module.

PURPOSE	COMMAND	RESULT
A linear (dense) layer.	<code>nn.Linear(10, 20)</code>	
Stack layers in order.	<code>nn.Sequential(nn.Linear(10,20), nn.ReLU(), nn.Linear(20,2))</code>	
Subclass for custom forward.	<code>class Net(nn.Module): def forward(self, x): ...</code>	
Apply an activation.	<code>F.relu(x)</code>	
Pick a loss.	<code>nn.CrossEntropyLoss()</code>	
Count / list parameters.	<code>sum(p.numel() for p in m.parameters())</code>	

layers go in `__init__`, the data path in `forward`; the module auto-registers weights as parameters

Data Loading




A Dataset indexes samples; a DataLoader batches and shuffles them.

PURPOSE	COMMAND	RESULT
Wrap tensors as a dataset.	<code>TensorDataset(X, y)</code>	
Custom dataset (3 methods).	<code>class DS(Dataset): __len__ __getitem__</code>	
Batch + shuffle.	<code>DataLoader(ds, batch_size=16, shuffle=True)</code>	
Iterate batches.	<code>for xb, yb in loader:</code>	
Grab one batch (debug).	<code>xb, yb = next(iter(loader))</code>	
Split train / val.	<code>random_split(ds, [80, 20])</code>	

a Dataset answers `__len__` and `__getitem__`; `next(iter(loader))` sanity-checks one batch's shapes

The Training Loop

Zero grads, forward, loss, backward, step. Repeat each batch.

PURPOSE	COMMAND	RESULT
Pick an optimizer.	<code>optim.AdamW(model.parameters(), lr=1e-3)</code>	
Clear old gradients.	<code>optimizer.zero_grad()</code>	
Forward + loss.	<code>loss = loss_fn(model(xb), yb)</code>	
Backpropagate.	<code>loss.backward()</code>	
Update the weights.	<code>optimizer.step()</code>	
Anneal the rate.	<code>scheduler.step()</code>	

five beats: zero_grad, forward, loss, backward, step; gradients accumulate, so zeroing first is mandatory

Save and Load

Save the state_dict, not the object; reload with weights_only=True.

PURPOSE	COMMAND	RESULT
Save model weights.	<code>torch.save(model.state_dict(), "model.pth")</code>	
Load weights back.	<code>model.load_state_dict(torch.load("model.pth", weights_only=True))</code>	
Save a full checkpoint.	<code>torch.save({"model":..., "opt":..., "epoch":e})</code>	
Set the right mode.	<code>model.eval()</code> <code>model.train()</code>	
Save a plain tensor.	<code>torch.save(t, "t.pt")</code>	
Export a portable graph.	<code>torch.export.export(model, (ex,))</code>	

save the state_dict (a plain dict of tensors); always pass weights_only=True to avoid running pickled code

Devices and Inference

Move model and data to the accelerator; run with no gradients.

PURPOSE	COMMAND	RESULT
Pick the accelerator.	<code>device =torch.accelerator.current_ accelerator() ... else "cpu"</code>	
Move model + data over.	<code>model.to(device)</code> <code>xb.to(device)</code>	
Run inference safely.	<code>withtorch.inference_mode():</code> <code>out = model(x)</code>	
Logits to a prediction.	<code>out.argmax(dim=1)</code>	
Logits to probabilities.	<code>torch.softmax(out, dim=1)</code>	
Make a run reproducible.	<code>torch.manual_seed(0)</code>	

torch.accelerator picks CUDA / MPS / CPU; run predictions under inference_mode, then argmax or softmax