

tb1x0.py

```
1 import torch
2 import torch.nn as nn
3 import torch.optim as optim
4 import torchvision
5 import torchvision.transforms as transforms
6
7
8 # Define the neural network model
9 class NeuralNetwork(nn.Module):
10     def __init__(self, input_size, hidden_size, num_classes):
11         super(NeuralNetwork, self).__init__()
12         self.fc1 = nn.Linear(input_size, hidden_size)
13         self.relu = nn.ReLU()
14         self.fc2 = nn.Linear(hidden_size, num_classes)
15
16     def forward(self, x):
17         out = self.fc1(x)
18         out = self.relu(out)
19         out = self.fc2(out)
20         return out
21
22
23 # Hyperparameters
24 input_size = 784 # 28x28 pixels
25 hidden_size = 128
26 num_classes = 10
27 learning_rate = 0.001
28 batch_size = 100
29 num_epochs = 5
30
31 # Load the MNIST dataset
32 transform = transforms.Compose(
33     [transforms.ToTensor(), transforms.Normalize((0.5,), (0.5,))]
34 )
35 train_dataset = torchvision.datasets.MNIST(
36     root="./data", train=True, transform=transform, download=True
37 )
38 test_dataset = torchvision.datasets.MNIST(
39     root="./data", train=False, transform=transform
40 )
41 train_loader = torch.utils.data.DataLoader(
42     dataset=train_dataset, batch_size=batch_size, shuffle=True
43 )
44 test_loader = torch.utils.data.DataLoader(
45     dataset=test_dataset, batch_size=batch_size, shuffle=False
46 )
47
48 # Initialize the model
49 model = NeuralNetwork(input_size, hidden_size, num_classes)
50
51 # Loss and optimizer
52 criterion = nn.CrossEntropyLoss()
53 optimizer = optim.Adam(model.parameters(), lr=learning_rate)
```

```

54
55 # Training loop
56 total_steps = len(train_loader)
57 for epoch in range(num_epochs):
58     for i, (images, labels) in enumerate(train_loader):
59         # Reshape images to (batch_size, input_size)
60         images = images.reshape(-1, 28 * 28)
61
62         # Forward pass
63         outputs = model(images)
64         loss = criterion(outputs, labels)
65
66         # Backward pass and optimization
67         optimizer.zero_grad()
68         loss.backward()
69         optimizer.step()
70
71         if (i + 1) % 100 == 0:
72             print(
73                 f"Epoch [{epoch+1}/{num_epochs}], Step [{i+1}/{total_steps}], Loss:
74                 {loss.item():.4f}"
75                 )
76 # Test the model
77 with torch.no_grad():
78     correct = 0
79     total = 0
80     for images, labels in test_loader:
81         images = images.reshape(-1, 28 * 28)
82         outputs = model(images)
83         _, predicted = torch.max(outputs.data, 1)
84         total += labels.size(0)
85         correct += (predicted == labels).sum().item()
86
87     print(f"Accuracy of the network on the 10000 test images: {100 * correct / total}%")
88

```