

```

1  # Importing the MNIST dataset from TensorFlow/Keras
2  from tensorflow.keras.datasets import mnist
3
4  # Loading the MNIST dataset into training and test sets
5  # train_images: 60,000 of 28x28 grayscale images of handwritten digits (0-9)
6  # train_labels: Corresponding digit labels (0-9)
7  # test_images: 10,000 of images for testing
8  # test_labels: Corresponding test digit labels
9  (train_images, train_labels), (test_images, test_labels) = mnist.load_data()
10
11 # Examining the shape of training data:
12 # 60,000 images - each of 28x28 pixels
13 train_images.shape
14
15 # Number of training labels (should match number of images)
16 len(train_labels)
17
18 # Display first few training labels (to verify data)
19 train_labels
20
21 # Shape of test data: 10,000 images - each of 28x28 pixels
22 test_images.shape
23
24 # Number of test labels
25 len(test_labels)
26
27 # Display first few test labels
28 test_labels
29
30 # Importing Keras modules for building neural network
31 from tensorflow import keras
32 from tensorflow.keras import layers
33
34 # Creating a sequential neural network model:
35 # - First layer: Dense (fully connected) with 512 neurons and ReLU activation
36 # - Second layer: Dense output layer with 10 neurons (for 10 digits) and softmax
   activation
37 model = keras.Sequential([
38     layers.Dense(512, activation="relu"),
39     layers.Dense(10, activation="softmax")
40 ])
41
42 # Compiling the model with:
43 # - RMSprop optimizer
44 # - Sparse categorical cross-entropy loss (for integer labels)
45 # - Accuracy as the evaluation metric
46 model.compile(optimizer="rmsprop",
47               loss="sparse_categorical_crossentropy",
48               metrics=["accuracy"])
49
50 # Preprocessing the image data:
51 # Reshaping from (60000, 28, 28) to (60000, 784) - flattening 28x28 to 784 pixels
52 train_images = train_images.reshape((60000, 28 * 28))
53 # Normalizing pixel values from 0-255 to 0-1
54 train_images = train_images.astype("float32") / 255
55
56 # Same preprocessing for test images
57 test_images = test_images.reshape((10000, 28 * 28))
58 test_images = test_images.astype("float32") / 255
59
60 # Training the model for 5 epochs with batch size of 128
61 model.fit(train_images, train_labels, epochs=5, batch_size=128)
62
63 # Making predictions on first 10 test images
64 test_digits = test_images[0:10]
65 predictions = model.predict(test_digits)
66

```

```

67 # Displaying the prediction probabilities for first test image
68 predictions[0]
69
70 # Getting the predicted digit (class with highest probability)
71 predictions[0].argmax()
72
73 # Showing the probability score for digit 7 (just as an example)
74 predictions[0][7]
75
76 # Displaying the actual label of first test image
77 test_labels[0]
78
79 # Evaluating model performance on entire test set
80 test_loss, test_acc = model.evaluate(test_images, test_labels)
81 print(f"test_acc: {test_acc}")
82
83 # Demonstrating NumPy arrays and their dimensions
84 import numpy as np
85 x = np.array(12) # scalar (0D tensor)
86 x
87
88 # Checking dimensions of the array
89 x.ndim
90
91 # Creating a 1D array (vector)
92 x = np.array([12, 3, 6, 14, 7])
93 x
94
95 # Checking dimensions
96 x.ndim
97
98 # Creating a 2D array (matrix)
99 x = np.array([[5, 78, 2, 34, 0],
100              [6, 79, 3, 35, 1],
101              [7, 80, 4, 36, 2]])
102 x.ndim
103
104 # Creating a 3D array
105 x = np.array([[[5, 78, 2, 34, 0],
106              [6, 79, 3, 35, 1],
107              [7, 80, 4, 36, 2]],
108              [[5, 78, 2, 34, 0],
109              [6, 79, 3, 35, 1],
110              [7, 80, 4, 36, 2]],
111              [[5, 78, 2, 34, 0],
112              [6, 79, 3, 35, 1],
113              [7, 80, 4, 36, 2]]])
114 x.ndim
115
116 # Reloading MNIST data to demonstrate tensor properties
117 from tensorflow.keras.datasets import mnist
118 (train_images, train_labels), (test_images, test_labels) = mnist.load_data()
119
120 # Checking tensor properties:
121 train_images.ndim # Number of dimensions (3)
122 train_images.shape # Shape (60000, 28, 28)
123 train_images.dtype # Data type (uint8)
124
125 # Visualizing a sample digit
126 import matplotlib.pyplot as plt
127 digit = train_images[4] # Get 5th image
128 plt.imshow(digit, cmap=plt.cm.binary) # Display as grayscale
129 plt.show()
130
131 # Checking the label of this digit
132 train_labels[4]
133

```

```

134 # Demonstrating tensor slicing operations
135 my_slice = train_images[10:100] # Get images 10-99
136 my_slice.shape
137
138 # Equivalent slice with explicit dimensions
139 my_slice = train_images[10:100, :, :]
140 my_slice.shape
141
142 # More specific slicing
143 my_slice = train_images[10:100, 0:28, 0:28]
144 my_slice.shape
145
146 # Center-cropping digits (keeping center 14x14 pixels)
147 my_slice = train_images[:, 14:, 14:] # All images, bottom-right 14x14
148
149 # Alternative center crop (keeping center 14x14 pixels)
150 my_slice = train_images[:, 7:-7, 7:-7]
151
152 # Batch processing examples
153 batch = train_images[:128] # First batch of 128 images
154 batch = train_images[128:256] # Second batch of 128 images
155
156 # General batch selection
157 n = 3
158 batch = train_images[128 * n:128 * (n + 1)] # nth batch
159
160 # Implementing ReLU activation function manually
161 def naive_relu(x):
162     assert len(x.shape) == 2 # Only works for 2D arrays
163     x = x.copy() # Do not modify original
164     for i in range(x.shape[0]): # Rows
165         for j in range(x.shape[1]): # Columns
166             x[i, j] = max(x[i, j], 0) # ReLU: max(0, x)
167     return x
168
169 # Implementing matrix addition manually
170 def naive_add(x, y):
171     assert len(x.shape) == 2 # Only works for 2D arrays
172     assert x.shape == y.shape # Must be same shape
173     x = x.copy() # Do not modify original
174     for i in range(x.shape[0]): # Rows
175         for j in range(x.shape[1]): # Columns
176             x[i, j] += y[i, j] # Element-wise addition
177     return x
178
179 # Performance comparison between NumPy and manual operations
180 import time
181 x = np.random.random((20, 100)) # Random 20x100 matrices
182 y = np.random.random((20, 100))
183
184 # Time NumPy operations (1000 iterations)
185 t0 = time.time()
186 for _ in range(1000):
187     z = x + y # Matrix addition
188     z = np.maximum(z, 0.) # ReLU
189 print("Took: {0:.2f} s".format(time.time() - t0))
190
191 # Time manual operations (1000 iterations)
192 t0 = time.time()
193 for _ in range(1000):
194     z = naive_add(x, y) # Manual addition
195     z = naive_relu(z) # Manual ReLU
196 print("Took: {0:.2f} s".format(time.time() - t0))
197
198 # Demonstrating broadcasting with NumPy
199 import numpy as np
200 X = np.random.random((32, 10)) # 32 samples, 10 features

```

```

201 y = np.random.random((10,)) # Single vector of 10 values
202
203 # Demonstrate broadcasting by expanding dimensions
204 y = np.expand_dims(y, axis=0) # Now shape (1, 10)
205 Y = np.concatenate([y] * 32, axis=0) # Now shape (32, 10) via tiling
206
207 # Manual implementation of matrix + vector addition
208 def naive_add_matrix_and_vector(x, y):
209     assert len(x.shape) == 2 # Must be matrix
210     assert len(y.shape) == 1 # Must be vector
211     assert x.shape[1] == y.shape[0] # Compatible dimensions
212     x = x.copy() # Do not modify original
213     for i in range(x.shape[0]): # Rows
214         for j in range(x.shape[1]): # Columns
215             x[i, j] += y[j] # Add vector element to each row
216     return x
217
218 # More broadcasting examples
219 import numpy as np
220 x = np.random.random((64, 3, 32, 10)) # 4D tensor
221 y = np.random.random((32, 10)) # 2D tensor
222 z = np.maximum(x, y) # Broadcasting across first two dimensions
223
224 # Vector dot product example
225 x = np.random.random((32,))
226 y = np.random.random((32,))
227 z = np.dot(x, y) # Dot product
228
229 # Manual implementation of vector dot product
230 def naive_vector_dot(x, y):
231     assert len(x.shape) == 1 # Must be vectors
232     assert len(y.shape) == 1
233     assert x.shape[0] == y.shape[0] # Same length
234     z = 0. # Initialize result
235     for i in range(x.shape[0]):
236         z += x[i] * y[i] # Sum of element-wise products
237     return z
238
239 # Manual matrix-vector multiplication
240 def naive_matrix_vector_dot(x, y):
241     assert len(x.shape) == 2 # Must be matrix and vector
242     assert len(y.shape) == 1
243     assert x.shape[1] == y.shape[0] # Compatible dimensions
244     z = np.zeros(x.shape[0]) # Initialize output vector
245     for i in range(x.shape[0]): # For each row
246         for j in range(x.shape[1]): # For each column
247             z[i] += x[i, j] * y[j] # Dot product with vector
248     return z
249
250 # Alternative implementation using previous function
251 def naive_matrix_vector_dot_new(x, y):
252     z = np.zeros(x.shape[0])
253     for i in range(x.shape[0]):
254         z[i] = naive_vector_dot(x[i, :], y) # Use vector dot product
255     return z
256
257 # Manual matrix multiplication
258 def naive_matrix_dot(x, y):
259     assert len(x.shape) == 2 # Must be matrices
260     assert len(y.shape) == 2
261     assert x.shape[1] == y.shape[0] # Compatible dimensions
262     z = np.zeros((x.shape[0], y.shape[1])) # Result matrix
263     for i in range(x.shape[0]): # Rows of x
264         for j in range(y.shape[1]): # Columns of y
265             row_x = x[i, :] # Get row from x
266             column_y = y[:, j] # Get column from y
267             z[i, j] = naive_vector_dot(row_x, column_y) # Dot product

```

```

268     return z
269
270 # Reshaping demonstration
271 train_images = train_images.reshape((60000, 28 * 28)) # Flatten images
272
273 # More reshaping examples
274 x = np.array([[0., 1.],
275              [2., 3.],
276              [4., 5.]])
277 x.shape
278
279 x = x.reshape((6, 1)) # Reshape to 6x1
280 x
281
282 # Transposition example
283 x = np.zeros((300, 20)) # Create 300x20 matrix
284 x = np.transpose(x) # Transpose to 20x300
285 x.shape
286
287 # TensorFlow automatic differentiation example
288 import tensorflow as tf
289 x = tf.Variable(0.) # Create variable with value 0
290 with tf.GradientTape() as tape:
291     y = 2 * x + 3 # Define computation
292     grad_of_y_wrt_x = tape.gradient(y, x) # Compute gradient dy/dx
293
294 # Matrix case
295 x = tf.Variable(tf.random.uniform((2, 2))) # 2x2 variable
296 with tf.GradientTape() as tape:
297     y = 2 * x + 3 # Simple computation
298     grad_of_y_wrt_x = tape.gradient(y, x) # Gradient is also 2x2
299
300 # More complex differentiation example
301 W = tf.Variable(tf.random.uniform((2, 2))) # Weights
302 b = tf.Variable(tf.zeros((2,))) # Bias
303 x = tf.random.uniform((2, 2)) # Input
304 with tf.GradientTape() as tape:
305     y = tf.matmul(x, W) + b # Linear transformation
306     grad_of_y_wrt_W_and_b = tape.gradient(y, [W, b]) # Gradients for W and b
307
308 # Reloading and preparing MNIST data again
309 (train_images, train_labels), (test_images, test_labels) = mnist.load_data()
310 train_images = train_images.reshape((60000, 28 * 28)) # Flatten
311 train_images = train_images.astype("float32") / 255 # Normalize
312 test_images = test_images.reshape((10000, 28 * 28)) # Same for test
313 test_images = test_images.astype("float32") / 255
314
315 # Rebuilding the model
316 model = keras.Sequential([
317     layers.Dense(512, activation="relu"), # Hidden layer
318     layers.Dense(10, activation="softmax") # Output layer
319 ])
320
321 # Compiling the model
322 model.compile(optimizer="rmsprop",
323              loss="sparse_categorical_crossentropy",
324              metrics=["accuracy"])
325
326 # Training the model
327 model.fit(train_images, train_labels, epochs=5, batch_size=128)
328
329 # Custom layer implementation
330 import tensorflow as tf
331 class NaiveDense:
332     def __init__(self, input_size, output_size, activation):
333         self.activation = activation
334         # Initialize weights with small random values

```

```

335     w_shape = (input_size, output_size)
336     w_initial_value = tf.random.uniform(w_shape, minval=0, maxval=1e-1)
337     self.W = tf.Variable(w_initial_value)
338     # Initialize biases with zeros
339     b_shape = (output_size,)
340     b_initial_value = tf.zeros(b_shape)
341     self.b = tf.Variable(b_initial_value)
342
343     # Forward pass
344     def __call__(self, inputs):
345         return self.activation(tf.matmul(inputs, self.W) + self.b)
346
347     # Property to access weights
348     @property
349     def weights(self):
350         return [self.W, self.b]
351
352     # Custom sequential model implementation
353     class NaiveSequential:
354         def __init__(self, layers):
355             self.layers = layers # List of layers
356
357         # Forward pass through all layers
358         def __call__(self, inputs):
359             x = inputs
360             for layer in self.layers:
361                 x = layer(x)
362             return x
363
364         # Property to access all weights
365         @property
366         def weights(self):
367             weights = []
368             for layer in self.layers:
369                 weights += layer.weights
370             return weights
371
372     # Creating an instance of our custom model
373     model = NaiveSequential([
374         NaiveDense(input_size=28 * 28, output_size=512, activation=tf.nn.relu),
375         NaiveDense(input_size=512, output_size=10, activation=tf.nn.softmax)
376     ])
377     assert len(model.weights) == 4 # 2 weight matrices + 2 bias vectors
378
379     # Batch generator for custom training loop
380     import math
381     class BatchGenerator:
382         def __init__(self, images, labels, batch_size=128):
383             assert len(images) == len(labels) # Sanity check
384             self.index = 0 # Current position
385             self.images = images
386             self.labels = labels
387             self.batch_size = batch_size
388             self.num_batches = math.ceil(len(images) / batch_size) # Total batches
389
390         # Get next batch
391         def next(self):
392             images = self.images[self.index : self.index + self.batch_size]
393             labels = self.labels[self.index : self.index + self.batch_size]
394             self.index += self.batch_size # Advance position
395             return images, labels
396
397     # Single training step implementation
398     def one_training_step(model, images_batch, labels_batch):
399         # Forward pass
400         with tf.GradientTape() as tape:
401             predictions = model(images_batch) # Get predictions

```

```

402     # Compute loss
403     per_sample_losses = tf.keras.losses.sparse_categorical_crossentropy(
404         labels_batch, predictions)
405     average_loss = tf.reduce_mean(per_sample_losses) # Average loss
406
407     # Backward pass
408     gradients = tape.gradient(average_loss, model.weights) # Compute gradients
409     update_weights(gradients, model.weights) # Update weights
410     return average_loss # Return loss for monitoring
411
412 # Learning rate for weight updates
413 learning_rate = 1e-3
414
415 # Simple weight update function
416 def update_weights(gradients, weights):
417     for g, w in zip(gradients, weights):
418         # w = w - learning_rate * gradient
419         w.assign_sub(g * learning_rate)
420
421 # Alternative using Keras optimizer
422 from tensorflow.keras import optimizers
423 optimizer = optimizers.SGD(learning_rate=1e-3)
424
425 # Optimizer-based weight update
426 def update_weights_new(gradients, weights):
427     optimizer.apply_gradients(zip(gradients, weights))
428
429 # Full training loop
430 def fit(model, images, labels, epochs, batch_size=128):
431     for epoch_counter in range(epochs):
432         print(f"Epoch {epoch_counter}")
433         batch_generator = BatchGenerator(images, labels, batch_size)
434         for batch_counter in range(batch_generator.num_batches):
435             images_batch, labels_batch = batch_generator.next()
436             loss = one_training_step(model, images_batch, labels_batch)
437             # Print progress periodically
438             if batch_counter % 100 == 0:
439                 print(f"loss at batch {batch_counter}: {loss:.2f}")
440
441 # Reload and prepare MNIST data
442 from tensorflow.keras.datasets import mnist
443 (train_images, train_labels), (test_images, test_labels) = mnist.load_data()
444 train_images = train_images.reshape((60000, 28 * 28))
445 train_images = train_images.astype("float32") / 255
446 test_images = test_images.reshape((10000, 28 * 28))
447 test_images = test_images.astype("float32") / 255
448
449 # Train our custom model
450 fit(model, train_images, train_labels, epochs=10, batch_size=128)
451
452 # Evaluating custom model
453 predictions = model(test_images) # Get predictions
454 predictions = predictions.numpy() # Convert to NumPy
455 predicted_labels = np.argmax(predictions, axis=1) # Get class predictions
456 matches = predicted_labels == test_labels # Compare with true labels
457 print(f"accuracy: {matches.mean():.2f}") # Calculate accuracy

```