

Unit 5 Lecture 2: Neural Networks

November 16, 2023

In this R demo, we'll be fitting fully-connected neural networks to the MNIST handwritten digit data.

First let's load some libraries:

```
library(keras)      # for deep learning
library(cowplot)   # for side-by-side plots
library(stat471)   # for deep learning helper functions
library(tidyverse) # for everything else
```

Next let's load the MNIST data and do some reshaping and rescaling:

```
# load the data
mnist <- dataset_mnist()

# extract information about the images
num_train_images <- dim(mnist$train$x)[1] # number of training images
num_test_images  <- dim(mnist$test$x)[1]  # number of test images
img_rows         <- dim(mnist$train$x)[2]  # rows per image
img_cols         <- dim(mnist$train$x)[3]  # columns per image
num_pixels       <- img_rows * img_cols    # pixels per image
num_classes      <- length(unique(mnist$train$y)) # number of image classes
max_intensity    <- 255                    # max pixel intensity

# normalize and reshape the images
x_train <- array_reshape(
  mnist$train$x / max_intensity,
  c(num_train_images, img_rows, img_cols, 1)
)
x_test  <- array_reshape(
  mnist$test$x / max_intensity,
  c(num_test_images, img_rows, img_cols, 1)
)

# extract the responses from the training and test data
g_train <- mnist$train$y
g_test  <- mnist$test$y

# recode response labels using "one-hot" representation
y_train <- to_categorical(g_train, num_classes)
y_test  <- to_categorical(g_test, num_classes)
```

Quick note that will be helpful for the homework: `plot_grayscale()` accommodates named classes. For example:


```
# tibble of class names, with columns class and name
class_names <- tribble(
  ~class, ~name,
```

```

0, "zero",
1, "one",
2, "two",
3, "three",
4, "four",
5, "five",
6, "six",
7, "seven",
8, "eight",
9, "nine"
)

p1 <- plot_grayscale(
  image_array = x_train[1, , , ],
  label = g_train[1],
  class_names = class_names
)
p2 <- plot_grayscale(
  image_array = x_train[2, , , ],
  label = g_train[2],
  class_names = class_names
)
p3 <- plot_grayscale(
  image_array = x_train[3, , , ],
  label = g_train[3],
  class_names = class_names
)
plot_grid(p1, p2, p3, nrow = 1)

```



Next, we define a neural network model with one hidden layer with 256 units and dropout rate 0.5.

```

model_nn <- keras_model_sequential() |>
  layer_flatten(input_shape = c(img_rows, img_cols, 1)) |> # flatten the input
  layer_dense(units = 256, activation = "relu") |> # first hidden layer
  layer_dropout(rate = 0.5) |> # specify dropout as "layer"
  layer_dense(units = 10, activation = "softmax") # output layer with 10 units

```

Let's print the summary of this neural network:

```
summary(model_nn)
```

```
## Model: "sequential"
## -----
## Layer (type)                Output Shape          Param #
## =====
## flatten (Flatten)           (None, 784)           0
## dense_1 (Dense)             (None, 256)           200960
## dropout (Dropout)           (None, 256)           0
## dense (Dense)               (None, 10)            2570
## =====
## Total params: 203530 (795.04 KB)
## Trainable params: 203530 (795.04 KB)
## Non-trainable params: 0 (0.00 Byte)
## -----
```

How do we arrive at the total number of parameters in this network?

To train this neural network, we must first define what loss function to use, which optimizer to use, and which metrics to track. We do this by *compiling* the model.

```
model_nn |> compile(
  loss = "categorical_crossentropy", # which loss function to use
  optimizer = optimizer_adagrad(),   # which flavor of stochastic gradient descent to use
  metrics = c("accuracy")           # which metric to track on validation set
)
```

Finally, we can train the model! We use 10 epochs, (mini-)batch size 128, and reserve 20% of our training data for validation.

```
history <- model_nn |>
  fit(
    x_train,           # supply training features
    y_train,           # supply training responses
    epochs = 10,       # number of epochs for stochastic gradient descent
    batch_size = 128,  # mini-batch size for stochastic gradient descent
    validation_split = 0.2 # use 20% of the training data for validation
  )
```

```
Epoch 1/10
375/375 [=====] - 3s 6ms/step - loss: 0.9042 - accuracy: 0.7397 - val_loss: 0.4263 - val_accuracy: 0.8947
Epoch 2/10
375/375 [=====] - 3s 7ms/step - loss: 0.4929 - accuracy: 0.8600 - val_loss: 0.3365 - val_accuracy: 0.9104
Epoch 3/10
375/375 [=====] - 3s 7ms/step - loss: 0.4153 - accuracy: 0.8813 - val_loss: 0.2967 - val_accuracy: 0.9190
Epoch 4/10
375/375 [=====] - 2s 6ms/step - loss: 0.3709 - accuracy: 0.8933 - val_loss: 0.2707 - val_accuracy: 0.9260
Epoch 5/10
375/375 [=====] - 3s 7ms/step - loss: 0.3395 - accuracy: 0.9040 - val_loss: 0.2521 - val_accuracy: 0.9302
Epoch 6/10
375/375 [=====] - 3s 7ms/step - loss: 0.3177 - accuracy: 0.9095 - val_loss: 0.2362 - val_accuracy: 0.9356
Epoch 7/10
375/375 [=====] - 3s 7ms/step - loss: 0.2984 - accuracy: 0.9155 - val_loss: 0.2229 - val_accuracy: 0.9391
Epoch 8/10
375/375 [=====] - 2s 6ms/step - loss: 0.2816 - accuracy: 0.9213 - val_loss: 0.2123 - val_accuracy: 0.9415
Epoch 9/10
375/375 [=====] - 2s 6ms/step - loss: 0.2678 - accuracy: 0.9248 - val_loss: 0.2040 - val_accuracy: 0.9438
Epoch 10/10
375/375 [=====] - 2s 6ms/step - loss: 0.2601 - accuracy: 0.9265 - val_loss: 0.1960 - val_accuracy: 0.9463
```

The number 375 represents the number of mini-batches in one epoch. Why are there 375 of these? The output printed while training gives us information about the metrics on the training and validation data,

as well as the time (in seconds) for each epoch and the average time (in milliseconds) for each stochastic gradient step for each epoch.

Now that we've had the patience to wait for this model to train, let's go ahead and save it, along with its history, so we don't need to train it again:


```
# save model  
save_model_hdf5(model_nn, "model_nn.h5")  
  
# save history  
saveRDS(model_nn$history$history, "model_nn_hist.RDS")
```

We can then load the model and its history into memory again:

```
# load model  
model_nn <- load_model_hdf5("model_nn.h5")  
  
# load history  
model_nn_hist <- readRDS("model_nn_hist.RDS")
```

We can plot the training history using `plot_model_history()` from `stat471`:

```
plot_model_history(model_nn_hist)
```



Did we observe any overfitting?

As before, we can get the fitted probabilities and predicted classes for the test set using `predict()` and `k_argmax()`:

```
# get fitted probabilities  
model_nn |>  
  predict(x_test) |>
```

```
head()
```

```
## 313/313 - 1s - 1s/epoch - 3ms/step
##           [,1]      [,2]      [,3]      [,4]      [,5]
## [1,] 4.085056e-05 3.616896e-07 0.0003544441 1.454709e-03 3.249738e-06
## [2,] 9.674003e-04 1.736223e-04 0.9818693995 8.310557e-03 6.199219e-07
## [3,] 1.040745e-04 9.779668e-01 0.0054318639 2.363973e-03 5.282222e-04
## [4,] 9.988591e-01 4.228233e-07 0.0001090119 4.886339e-05 6.358399e-07
## [5,] 9.557388e-04 9.433114e-05 0.0033908419 5.678684e-04 9.439621e-01
## [6,] 1.145565e-05 9.915801e-01 0.0013057842 7.227233e-04 8.456492e-05
##           [,6]      [,7]      [,8]      [,9]      [,10]
## [1,] 0.0001290360 4.686711e-07 9.971980e-01 4.798083e-05 7.706612e-04
## [2,] 0.0018138436 5.058184e-03 7.479182e-07 1.805018e-03 6.083085e-07
## [3,] 0.0012418190 2.088551e-03 5.391436e-03 4.228448e-03 6.547123e-04
## [4,] 0.0006133148 1.673248e-04 1.317946e-04 4.104923e-05 2.858582e-05
## [5,] 0.0012209849 4.288028e-03 4.148317e-03 2.349667e-03 3.902213e-02
## [6,] 0.0001423117 1.021783e-04 3.922658e-03 1.859010e-03 2.692516e-04
```

```
# get predicted classes
predicted_classes <- model_nn |>
  predict(x_test) |>
  k_argmax() |>
  as.integer()
```

```
## 313/313 - 0s - 451ms/epoch - 1ms/step
head(predicted_classes)
```

```
## [1] 7 2 1 0 4 1
```

We can extract the misclassification error / accuracy manually:

```
# misclassification error
mean(predicted_classes != g_test)
```

```
## [1] 0.0547
```

```
# accuracy
mean(predicted_classes == g_test)
```

```
## [1] 0.9453
```

Or we can use a shortcut and call `evaluate`:

```
evaluate(model_nn, x_test, y_test, verbose = FALSE)
```

```
##      loss accuracy
## 0.1982096 0.9453000
```

In addition to the accuracy / misclassification error, we can take a look at the *confusion matrix* of this classifier using the `plot_confusion_matrix()` function from `stat471`:

```
plot_confusion_matrix(
  predicted_responses = predicted_classes,
  actual_response = g_test
)
```

		Predicted Response									
		0	1	2	3	4	5	6	7	8	9
Actual Response	0	965	0	1	1	0	3	7	1	2	0
	1	0	1118	3	2	1	0	3	1	7	0
	2	9	2	968	8	10	1	7	7	17	3
	3	3	1	15	950	1	11	2	12	11	4
	4	1	1	3	0	936	0	10	2	3	26
	5	7	1	2	26	6	815	12	3	14	6
	6	10	3	4	1	11	6	920	0	3	0
	7	2	9	21	6	6	0	1	959	2	22
	8	5	4	5	19	6	8	13	11	895	8
	9	11	7	1	11	34	4	1	10	3	927

Let's take a look at an 8 that was misclassified as a 2:

```

misclassifications <- which(predicted_classes == 2 & g_test == 8)
idx <- misclassifications[1]
plot_grayscale(image_array = x_test[idx, , ], label = "An 8 that was classified as a 2 :(")

```

An 8 that was classified as a 2 :(

