

Unit 5 Lecture 1: Deep Learning Preliminaries

November 14, 2023

In this R demo, we'll get warmed up for deep learning by fitting a multi-class logistic regression to the MNIST handwritten digit data. The inputs are 28 pixel by 28 pixel images of handwritten digits (a total of 784 pixels), and the output is one of the ten categories 0, 1, ..., 9.

First let's load some libraries:

```
library(keras)      # for deep learning
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)
```

Let's plot a few of the digits:


```
# plot a few of the digits
p1 <- plot_grayscale(x_train[1, , , ], g_train[1])
plot(p1)
```

5




```
p2 <- plot_grayscale(x_train[2, , , ], g_train[2])  
plot(p2)
```

0



```
p3 <- plot_grayscale(x_train[3, , , ], g_train[3])  
plot(p3)
```



Now, we can define the class of model we want to train (multi-class logistic model):

```
model_lr <- keras_model_sequential() |>
  layer_flatten(
    input_shape =          # flatten the images
      c(img_rows, img_cols, 1)
  ) |>
  layer_dense(
    units = num_classes,  # number of outputs
    activation = "softmax" # type of activation function
  )
```

We can get a summary of this model as follows:

```
summary(model_lr)
```

```
## Model: "sequential"
## -----
## Layer (type)                Output Shape          Param #
## =====
## flatten (Flatten)           (None, 784)           0
## dense (Dense)               (None, 10)            7850
## =====
## Total params: 7850 (30.66 KB)
## Trainable params: 7850 (30.66 KB)
## Non-trainable params: 0 (0.00 Byte)
## -----
```

Now we need to *compile* the model, which adds information to the object about which loss we want to use, which way we want to optimize the loss, and how we will evaluate validation error:

```
model_lr |>                                # note: modifying model_lr "in place"
  compile(
    loss = "categorical_crossentropy", # which loss to use
    optimizer = optimizer_rmsprop(),  # how to optimize the loss
```

```
metrics = c("accuracy")           # how to evaluate the fit
)
```

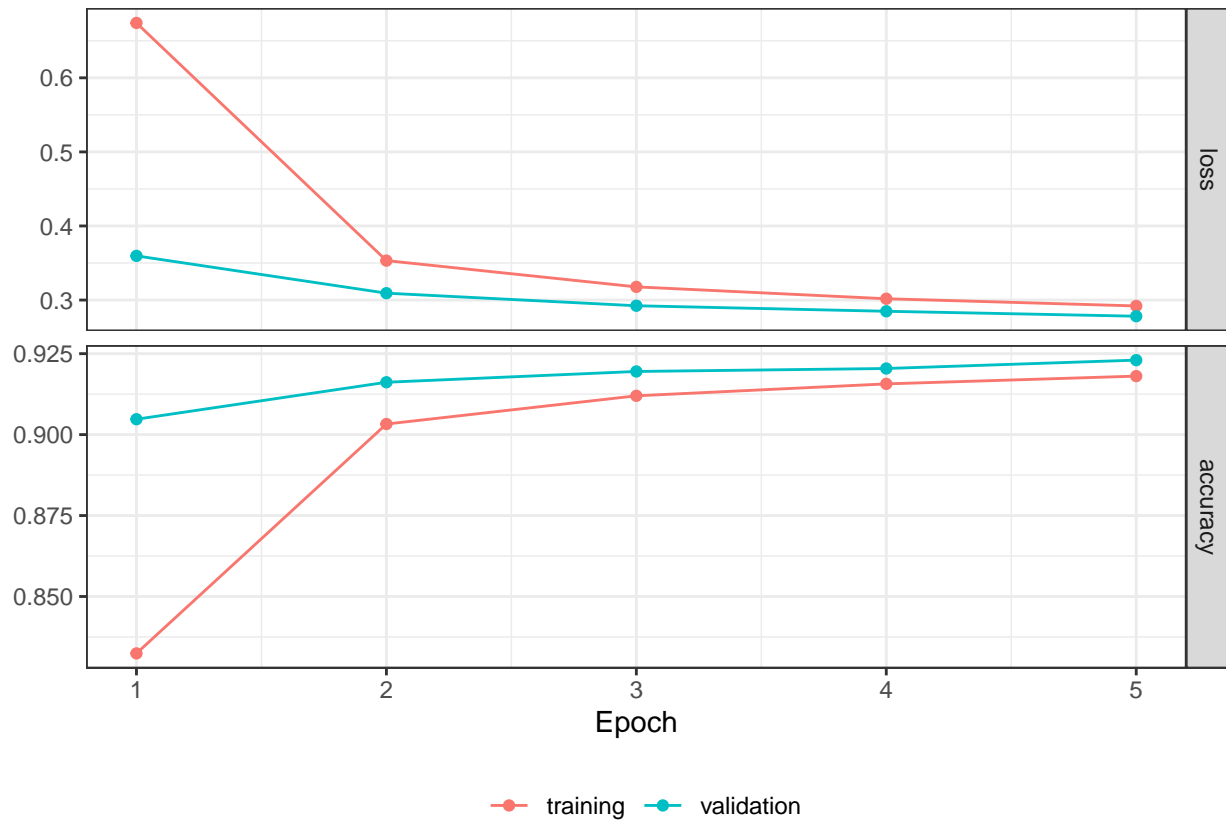
Finally, we can train the model! Let's use a small number of epochs (gradient steps) so the run time is manageable.

```
model_lr |>
  fit(
    x_train,           # supply training features
    y_train,           # supply training responses
    epochs = 5,        # an epoch is a gradient step
    batch_size = 128,  # we will learn about batches in Lecture 2
    validation_split = 0.2 # use 20% of the training data for validation
  )
```

```
## Epoch 1/5
## 375/375 - 1s - loss: 0.6741 - accuracy: 0.8324 - val_loss: 0.3596 - val_accuracy: 0.9047 - 1s/epoch
## Epoch 2/5
## 375/375 - 1s - loss: 0.3532 - accuracy: 0.9033 - val_loss: 0.3093 - val_accuracy: 0.9162 - 1s/epoch
## Epoch 3/5
## 375/375 - 1s - loss: 0.3178 - accuracy: 0.9120 - val_loss: 0.2923 - val_accuracy: 0.9195 - 1s/epoch
## Epoch 4/5
## 375/375 - 1s - loss: 0.3018 - accuracy: 0.9157 - val_loss: 0.2848 - val_accuracy: 0.9204 - 1s/epoch
## Epoch 5/5
## 375/375 - 1s - loss: 0.2920 - accuracy: 0.9181 - val_loss: 0.2782 - val_accuracy: 0.9230 - 1s/epoch
```

The field `model_lr$history$history` contains the progress during training, and can be plotted via

```
# plot the history
plot_model_history(model_lr$history$history)
```



We can get the fitted probabilities using the `predict()` function, and extract the classes with highest predicted probability using `k_argmax()`

```
# get fitted probabilities
model_lr |>
  predict(x_test) |>
  head()
```

```
## 313/313 - 0s - 374ms/epoch - 1ms/step
```

```
##           [,1]      [,2]      [,3]      [,4]      [,5]
## [1,] 1.462808e-05 1.160702e-09 3.399561e-05 2.135880e-03 2.030271e-06
## [2,] 8.087418e-04 8.483121e-06 9.892879e-01 1.020995e-03 8.766349e-11
## [3,] 1.040302e-05 9.747548e-01 9.142877e-03 3.605245e-03 3.094041e-04
## [4,] 9.994370e-01 1.073498e-11 9.841297e-05 9.129473e-06 5.880265e-08
## [5,] 7.693336e-04 7.383549e-07 4.425820e-03 1.747560e-04 9.242624e-01
## [6,] 6.184019e-07 9.884081e-01 2.390383e-03 2.538760e-03 3.634184e-05
##           [,6]      [,7]      [,8]      [,9]      [,10]
## [1,] 0.0000190821 1.372994e-08 9.970676e-01 3.131502e-05 6.953650e-04
## [2,] 0.0017641168 6.745132e-03 2.064690e-12 3.646166e-04 1.098275e-09
## [3,] 0.0007950441 1.893050e-03 1.848017e-03 7.292730e-03 3.484255e-04
## [4,] 0.0002433288 1.001752e-04 4.719529e-05 5.071840e-05 1.412544e-05
## [5,] 0.0002254635 3.227136e-03 6.438272e-03 1.257549e-02 4.790055e-02
## [6,] 0.0000697338 4.292931e-05 2.153714e-03 3.924290e-03 4.350991e-04
```

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

```
## 313/313 - 0s - 227ms/epoch - 727us/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] 0.0786
```

```
# accuracy
```

```
mean(predicted_classes == g_test)
```

```
## [1] 0.9214
```

Or we can use a shortcut and call evaluate:

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

```
##      loss  accuracy
```

```
## 0.2817411 0.9214000
```

Finally, let's take a look at one of the misclassified digits:

```
misclassifications <- which(predicted_classes != g_test)
```

```
idx <- misclassifications[1]
```

```
plot_grayscale(x_test[idx, , ] +  
  ggtitle(sprintf(  
    "Predicted class %d; True class %d",  
    predicted_classes[idx],  
    g_test[idx]  
  ))
```

Predicted class 6; True class 5

