

# Economics 706/420

November 17, 2025

R. Davidson

## Assignment 2

This is exercise 10 from Chapter 15 of Géron's book, reproduced here for ease of reference.

Download the Bach chorales dataset from <https://homl.info/bach> and unzip it. It is composed of 382 chorales composed by Johann Sebastian Bach. Each chorale is 100 to 640 time steps long, and each time step contains 4 integers, where each integer corresponds to a note's index on a piano (except for the value 0, which means that no note is played). Train a model – recurrent, convolutional, or both – that can predict the next time step (four notes), given a sequence of time steps from a chorale. Then use this model to generate Bach-like music, one note at a time: you can do this by giving the model the start of a chorale and asking it to predict the next time step, then appending these time steps to the input sequence and asking the model for the next note, and so on. Also make sure to check out Google's Coconet model from <https://homl.info/coconet>. It was used for a nice Google doodle about Bach. It will give you lots of ideas about how to undertake this assignment, and, perhaps, extend it.