


# **Experimental Mathematics: Stock trading with hidden Markov models**

Project supervisor: George Kerchev



---

UNIVERSITÉ DU  
LUXEMBOURG

## **Table of contents**

<b>Introduction</b>	<b>3</b>
<b>Part I: Hidden Markov Model</b>	<b>4</b>
Hidden Markov Model	4
Hidden Markov Model for Stock trading	6
Three basic problems of HMMs	6
Forward-Backward algorithm	7
Baum-Welch algorithm	8
Viterbi algorithm	9
<b>Part II: Programming</b>	<b>10</b>
Stock price prediction using Python	10
Conclusion	14

## **Stock trading with hidden Markov models**

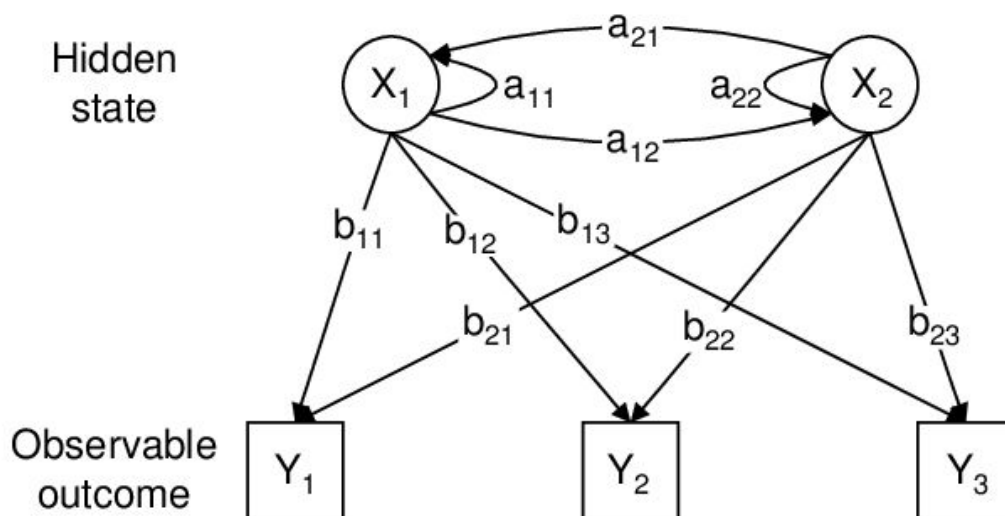
### **Introduction**

The aggregation of buyers and sellers of stocks also called shares is called a stock market. A stock or a share represents ownership claims on businesses. Stock trading consists of buying shares low and selling them high by profiting in doing so. The goal of our project was to be able to predict economic regimes and stock prices using the *Hidden Markov Model* (HMM), a probabilistic model. After gathering economic data from the company *Electronic Arts*, we tried to model it with a HMM. In order to predict the stock prices, we needed to train the model while optimizing the model parameters.

## Part I: Hidden Markov Model

### Hidden Markov Model

Named after the Russian mathematician Andrey Andreyevich, the Hidden Markov Model is a doubly stochastic process where one of the underlying stochastic processes is hidden. The hidden process is a Markov chain going from one state to another but cannot be observed directly. The other process is observable and depends on the hidden states. The goal of HMM is to capture the hidden information from the observables. HMM is widely used in speech recognition and in forecasting of the stock market.



X - hidden states

Y - observables


a - **transition probabilities**: probability of moving from one state to another state

b - **emission probabilities**: probability of an observation being generated from a state

$X_0$  - **initial probability distribution**: probability that the Markov chain will start at a certain state

*Example:*

An application of the HMM is often the determination of the weather which is hidden. The hidden states in our example are sunny and rainy. The activities of a certain person which depend on the weather can be observed. The observables are the following activities: walk, shop or clean. The transition probabilities here represent the change of weather from one day to another. The emission probabilities represent the likelihood of an activity being performed on each day depending on the state of the weather.



Let's suppose that for two consecutive days this person went for a walk. We will furthermore assume that it rained the first day and the second day it was sunny.

$$\begin{aligned} P((\text{walk}, \text{walk}), (\text{rainy}, \text{sunny})) &= P((\text{walk}, \text{walk}) | (\text{rainy}, \text{sunny})) P(\text{rainy}, \text{sunny}) \\ &= P(\text{walk} | \text{rainy}) P(\text{walk} | \text{sunny}) P(\text{sunny} | \text{rainy}) P(\text{rainy}) \\ &= 0.1 * 0.6 * 0.3 * 0.6 = 0.0108 \end{aligned}$$

## Hidden Markov Model for Stock trading


HMM are capable of predicting and analyzing time-based phenomena, hence, they are very useful for financial market prediction. The price of the stock, in this case our observable, is impacted by hidden volatility regimes. The hidden states are namely low volatility regimes, high volatility regimes and neutral volatility regimes. The goal here is to estimate the most likely regime, including the associated time varying means and volatilities, in order to create a reliable predictive model. So by analyzing the stock behavioral pattern in the past, we can forecast its future outcomes.

### Three basic problems of HMMs

- Evaluation: computation of the probability of the observation sequence  
→ Forward-Backward algorithm
- Learning: determination of the parameters of the model  
→ Baum-Welch algorithm
- Decoding: determination of the most probable state sequence  
→ Viterbi algorithm

## Forward-Backward algorithm

The Forward-Backward algorithm computes the posterior marginals of all hidden state variables. This means that it gives the probability distribution with relevant evidence and each background being taken into account. To achieve this it makes use of the principle of dynamic programming, which means simplifying a complicated problem by breaking it down into simpler sub-problems, as we can see on this image.



The algorithm uses two passes, the first pass goes forward in time and the second pass goes backward. The first pass computes a set of forward probabilities which provide the probability of ending up in any particular state. In the second pass, the algorithm computes a set of backward probabilities which provide the probability of observing the remaining observations given any starting point. So we have two sets of probabilities that can then be combined to obtain the distribution over states at any specific point in time given.


### *Conclusion:*

The Forward-Backward algorithm is used to find the most likely state for any point in time.

## Baum-Welch algorithm


The Baum-Welch algorithm is used to find the unknown parameters of a hidden Markov model. It's a special case of the EM algorithm (expectation-maximization algorithm) which is a method to find maximum a posteriori estimates of parameters in a statistical model.

To achieve this, the algorithm makes use of the Forward-Backward algorithm.



So we have two steps, the E-step, 'E' stands for expectation, re-estimating the responsibility profile  $\pi$  given the current HMM parameters.

(emitted string, ?, Parameters)  $\rightarrow \pi$



The other step is the M-step, 'M' stands for maximization, re-estimating the HMM parameters given the current responsibility profile  $\pi$ .

(emitted string,  $\pi$ , ?)  $\rightarrow$  Parameters

The Baum-Welch algorithm is needed because the state paths are hidden, and the equations cannot be solved analytically.

### *Conclusion:*


The Baum-Welch algorithm attempts to find the model that assigns the training data the highest likelihood.



## Viterbi algorithm

The Viterbi algorithm is most useful when one wants to calculate the most likely path through the state transitions of these models over time. Let's say we have  $N$  states and  $T$  moments in time, calculating the probabilities of all transitions over time would be  $N^T$  probability calculations. This algorithm is used to speed all these calculations up.

The observation made by the Viterbi algorithm is that for any state at time  $T$ , there is only one most likely path to that state. Therefore, if several paths converge at a particular state at time  $T$ , instead of redoing them all when calculating the transitions from this state to states at time  $T+1$ , one can discard the less likely paths, and only use the most likely one.



### *Conclusion:*

Using the Viterbi algorithm we can identify the most likely sequence of hidden states given the sequence of observations.

## Part II: Programming

### Stock price prediction using Python

In this section, we'll talk about the program of which the main goal was to predict stock prices using a Hidden Markov Model to fit the data, and we'll give some detailed explanations on the functioning of the code.

```
import datetime
import numpy as np
from pandas_datareader import data
from hmmlearn.hmm import GaussianHMM
import matplotlib.pyplot as plt
import matplotlib.cm as cm
from matplotlib.dates import YearLocator, MonthLocator
import warnings
warnings.filterwarnings('ignore')

ticker = 'EA'
start_date = datetime.date(2019, 1, 1)
end_date = datetime.date.today()
Stocks = data.DataReader(ticker, 'yahoo', start_date, end_date)

Stocks.reset_index(inplace=True, drop=False)
Stocks.drop(['Open', 'High', 'Low', 'Adj Close', 'Volume'], axis=1, inplace=True)
Stocks['Date'] = Stocks['Date'].apply(datetime.datetime.toordinal)
Stocks = list(Stocks.itertuples(index=False, name=None))

dates = np.array([q[0] for q in Stocks], dtype=int)
close_val = np.array([q[1] for q in Stocks])
close = [close_val[j]/close_val[j-1] for j in range(1, len(close_val))]

fig, ax = plt.subplots(figsize=(12,8))
plt.title('Stock Price over Time (EA)', fontsize = 14)
ax.plot_date(dates, close_val, '-', label='Close')
plt.gca().xaxis.set_major_locator(MonthLocator())
plt.xlabel('Date', fontsize = 14)
plt.ylabel('Close (price)', fontsize = 14)
plt.legend()
ax.autoscale_view()
ax.grid(True)
fig.autofmt_xdate()
plt.show()
```

First of all we have to import several packages in order to be able to run the code properly. The most important ones are the “pandas\_datareader” package, which is used to get our stock data from the ‘Yahoo Finance’ website, as well as the ‘hmmlearn.hmm’ package, which is used to fit the data with the Hidden Markov Model.

As mentioned above, the stock data that we use comes from the company *Electronic Arts*. We collect the data from the 1st January 2019 until the current date and convert it into a table using the ‘data.datareader()’ function.

This will look like that:

```

RESTART: C:\Users\benim\Desktop\final.py
High      Low      Open      Close     Volume    Adj Close
Date
2019-01-02  80.639999  77.199997  77.220001  80.370003  3176600  80.370003
2019-01-03  80.349998  77.620003  79.449997  78.709999  3747100  78.709999
2019-01-04  85.250000  80.040001  80.820000  84.419998  7254400  84.419998
2019-01-07  88.099998  83.440002  84.889999  87.589996  6611300  87.589996
2019-01-08  89.739998  87.690002  88.459999  89.019997  6468300  89.019997
...
...
...
2020-05-18  117.919998  114.730003  116.709999  117.120003  4336300  117.120003
2020-05-19  119.370003  116.720001  116.769997  118.540001  2517900  118.540001
2020-05-20  121.760002  119.059998  119.099998  119.550003  3530400  119.550003
2020-05-21  119.879997  117.199997  119.879997  117.360001  2257100  117.360001
2020-05-22  119.190002  116.614998  117.220001  119.070000  731565  119.070000

[351 rows x 6 columns]
Returns


```

Now we rearrange the table. We only want to have the data that we'll need later on. That means we remove every column except the one with the closing prices of each day, which is the 'Close' column and the 'Date' column containing the dates.

Then we transform the dates into Gregorian ordinals which are representing the number of days passed since 1/1/0001. So for example the first day of our data, the 2nd of January 2019 (on the 1st of January the stock market is closed) would be represented by the number 737061.

After that we transform the table into a list containing the Gregorian ordinals and the closing prices for each day. We also define our lists 'dates' which are the Georgian ordinals, 'close\_val' which are the closing prices and 'close'. For the 'close' list we take the closing price of a given day and divide it by the closing price of the previous day in order to have an indicator of the current stock situation.

At last, we plot the closing prices and get the following plot:



We will now proceed to the actual 'prediction' part of the program.

```
N = 5
train_x = list(map(lambda el:[el], close[:N]))

model = GaussianHMM(n_components=2, covariance_type='diag', n_iter=1000)

predicted_prices = []
predicted_dates = []

for idx in range(100):
    model.fit(train_x)
    state = model.predict(train_x)[-1]
    means = model.means_[state]
    current_price = Stocks[N+idx][1]
    current_date = datetime.date.fromordinal(dates[N+idx])
    predicted_date = current_date + datetime.timedelta(days=1)
    predicted_dates.append(predicted_date)
    train_x.append([means[0]])
    predicted_prices.append(current_price * means[0])
    train_x.pop(0)

print('Actual:',close_val[N+1:-246], '\n','Prediction:', predicted_prices )

plt.figure(figsize=(12, 8))
plt.title('Close (price)', fontsize = 14)
plt.plot(predicted_dates,close_val[N+1:-246])
plt.plot(predicted_dates,predicted_prices)
plt.legend(['Actual','Predicted'])
plt.grid(True)
plt.show()
```

We start by defining the set of our data which will be used to train our model. For reasons of simplification we only use the first 5 days of our data as the training set. Then we convert every element of our training set 'train\_x' into a list, so that we get a list of lists. This step is necessary in order to be able to train our model on our data.

The model we use is the 'GaussianHMM' model which uses the algorithms explained earlier in this report (Forward-Backward, Baum-Welch and Viterbi). As already mentioned, we train the model on the set 'train\_x'.

Now we define the 2 lists 'predicted\_prices' and 'predicted\_dates' that are evidently empty in the beginning. In order to fill these lists, we're using a for loop with range 100.

We start the loop by fitting our model and then looking at the hidden states. As we only want to predict the price for 1 day in the future, we just look at what the hidden state corresponds to for the last day of the training set. As an example, let's say that for the last day the hidden state corresponds to 1. Then we look at the mean parameter of the hidden state corresponding to the one we found (our expected return), so in that case we would be getting the mean parameter for the hidden state 1.

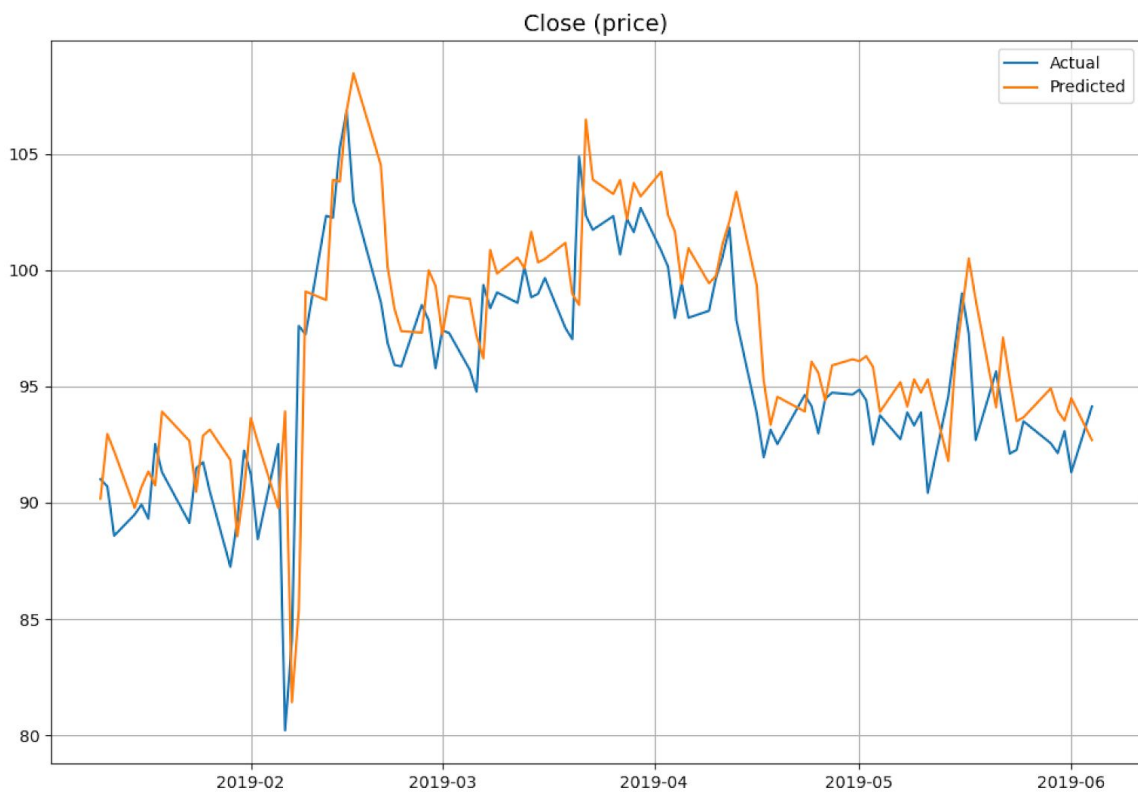
As we want to compare the actual prices with our predicted prices, we first set our current date to be the last day of our training set and look at the closing price for that day. Our predicted date will then be the following day.

Now we want to update our training set after every run of the loop. We add to the list the element '[means[0]]' and we remove the first element.

Then we just need to calculate our predicted price by multiplying the current price with our expected returns.

This process gets repeated 100 times in order to predict the prices for 100 days.

Finally we print and plot our actual and our predicted prices.



```

Actual: [ 91.01000214 90.69999695 88.58000183 89.48999786 89.91999817
89.30999756 92.51999664 91.30999756 89.12999725 91.48000336
91.73999786 90.47000122 87.23999786 89.23000336 92.23999786
91.22000122 88.43000031 92.51999664 80.20999908 84.09999847
97.59999847 97.23999786 102.33000183 102.25 105.25
106.83999634 102.95999908 98.62999725 96.86000061 95.91999817
95.86000061 98.5 97.83999634 95.77999878 97.41000366
97.29000092 95.72000122 94.76999664 99.36000061 98.36000061
99.04000092 98.58999634 100.12999725 98.83000183 98.98000336
99.66000366 97.51000214 97.02999878 104.87999725 102.33999634
101.73000336 102.31999969 100.66999817 102.19999695 101.62999725
102.66999817 100.83999634 100.15000153 97.94000244 99.43000031
97.94999695 98.25 99.62999725 100.58000183 101.83000183
97.84999847 93.83000183 91.94999695 93.13999939 92.51999664
94.62999725 94.15000153 92.98000336 94.47000122 94.73000336
94.65000153 94.86000061 94.40000153 92.5 93.76000214
92.73000336 93.87999725 93.31999969 93.87999725 90.41999817
94.56999969 96.76999664 99. 97.26000214 92.69000244
95.65000153 93.83999634 92.11000061 92.26999664 93.5
92.55000305 92.12999725 93.08000183 91.30999756 94.12999725]
Prediction: [90.17383737347944, 92.95382197730841, 92.21195405469606, 89.775193
65401092, 90.6796872292341, 91.34106884078668, 90.7480423072053, 93.914581442640
38, 92.65714737432498, 90.46748867239536, 92.8840074730246, 93.13961823624541, 9
1.83495897516062, 88.556469231731, 90.58251156327832, 93.64092462422666, 92.6025
2080257331, 89.768457525475, 93.92125428596802, 81.42570218731032, 85.3746427621
3859, 99.07863685684364, 98.71306753079476, 103.88044547524349, 103.799353206210
48, 106.84472881133104, 108.45872083757789, 104.51995737819523, 100.124396704860
59, 98.32758499093408, 97.37332092597757, 97.31240641150261, 99.99241316814954,
99.32241524856767, 97.231203605172, 98.88590219515439, 98.76408125800317, 97.170
29546252141, 96.20589735060854, 100.86544575341208, 99.85029396417056, 100.54059
744496301, 100.08377490355775, 101.64710898378203, 100.32741665802078, 100.47969
090358183, 101.16999423354764, 98.98741715520482, 98.50014104587777, 106.4690781
2631052, 103.8905925956868, 103.27135736451572, 103.87029298500124, 102.19529158
844878, 103.74847202042221, 103.1698360226679, 104.22559442998349, 102.367865472
09422, 101.66741625821292, 99.42393254956316, 100.93650599705872, 99.43407848777
43, 99.73862701161404, 101.13953318289772, 102.1039316789112, 103.37287095391751
, 99.33256489464635, 95.25165959415533, 93.34316996782451, 94.55120263599176, 93
.9218059674204, 96.06377608314261, 95.57650773875082, 94.38878243619017, 95.9013
5588334324, 96.1652974211316, 96.08408344873567, 96.29726431752935, 95.830295593
74445, 93.901506346812, 95.18059930456585, 94.13499458119973, 95.30241251816754,
94.73393020138394, 95.30241251816763, 91.78998953448574, 96.00286947634082, 98.
23619950226873, 100.49999057658778, 98.73362927444515, 94.09438759500229, 97.099
23486870139, 95.2618055319701, 93.50559791262451, 93.66801811242384, 94.91665776
677735, 93.95226701580637, 93.52589753322792, 94.4902960291896, 92.6934736786558
]
>>> |

```

## Conclusion

The purpose of our project was to use the Hidden Markov Model as a tool in financial forecasting. We successfully predicted the closing prices of a stock using HMM. Looking at the last chart about the predicted and actual prices, we see that the model is not that bad. It was able to catch the massive rise of the stock.