



# Buy or Sell?

Chris Lavin

901-658-9401

12/19/03

<b><i>Introduction</i></b>	<b>2</b>
<b>My Data:</b>	<b>2</b>
Price:	2
P/E Ratio:	2
Volume:	3
Williams %R:	3
<b>My stocks</b>	<b>4</b>
General Electric:	4
3M:	5
Wal-Mart:	6
Microsoft:	6
<b>Neural Network Program:</b>	<b>7</b>
<b><i>Analysis and Preparation of Data</i></b>	<b>7</b>
<b>Data Sets Chosen:</b>	<b>9</b>
<b>Laying out the data:</b>	<b>9</b>
<b><i>Experimenting</i></b>	<b>10</b>
<b>Initial Results:</b>	<b>10</b>
<b>Refining:</b>	<b>10</b>
<b>Final Settings:</b>	<b>11</b>
<b><i>Final Results</i></b>	<b>12</b>
<b>Results:</b>	<b>12</b>
<b>Reversed Data;</b>	<b>12</b>
<b>Training with GE:</b>	<b>13</b>
<b>Baseline Test:</b>	<b>13</b>
<b><i>Explanation of Results</i></b>	<b>14</b>
<b>3M Results:</b>	<b>14</b>
<b>Wal-Mart Results:</b>	<b>15</b>
<b>GE Resultes:</b>	<b>15</b>
<b><i>Conclusions:</i></b>	<b>15</b>
<b><i>Buy of Sell?</i></b>	<b>16</b>

## Introduction

For my project I wanted to use a back propagation neural network to decide if you should buy or sell a stock on a one year outlook. I'm something of a stock addict. I've been playing the stock market since I was in eighth grade. During the early years you could probably classify me as a near day trader (school required me to take a few days off now and then). As the market went south I managed to thrive, but I came to realize I'd have gray hair by the time I was 30 if I had any hair left if I continued day trading. Since that point I've been trying to make myself into a long term investor.

When I thought about predicting the stock market for this project I realized that I could use it as an opportunity to try to get myself into the right long term mindset. I choose a year time frame; which is about as long term as I can think about right now, when you are use to trading times measured in days if not hours a year is a very long long time. I also set a goal that I wanted, using results from the last twenty years to be able to predict how the market performed during both the bull market of the late nineties and the collapse around 2000. Based on this I choose to use the years 1996 – 2003 as my testing zone.

I was also interested in finding out how stocks in very different sectors compared to one another in terms of patterns the network could detect. I've always had a theory that regardless of what sector you are in on average stocks performs similarly in similar situations.

My final goal for this project was to prove that you don't need to use a hundred different numbers to get decent results. I've often found that there are two kinds of investors, those who get bogged down by the thousands of different stats you can get about a stock and those who just ignore all the numbers. Neither approach seems right to me. New traders are often highly intimidated by all the data they are presented with. When I first started to trade I use to be a numbers man. I would keep track of dozens of indicators for each stock I followed. Over the years though I've come to realize that it is much more important to know what a few select numbers are telling you then knowing what current stats are for all of them. So for this project I decided I would base all of my calculations off of just four pieces of data.

All though this sounds extreme; I've been very successful in the stock market these last few years by following only a few more indicators. I suspected that a neural network would need even fewer data points to find a pattern. Based on my results I would have to say my hypothesis was justified.

My Data:

Price:

This is probably the most obvious choice for stocks. The price you pay for a stock and the price you sell a stock at determine how well you do in the stock market.

P/E Ratio:

This is my favorite indicator. P/E stands for price to earnings. The Price to Earnings ratio is calculated as the current market price of a company's common stock divided by that company's earnings per share in the previous 12-month period.

The PE ratio allows for the simplest comparison between different shares, as companies within a particular industry generally fall within a certain PE range. Comparisons between companies in different industries, however, are generally not appropriate using the Price to Earnings ratio. I learned the hard way during the dot bust that a PE ratio of a 100 really isn't a deal.

Through the years I've found that PE often is one of the best indications of whether a stock is overbought or oversold for the long run. When you see rapid changes in PE Ratio I've found that it is normally a good time to start looking into either buying or selling depending on the situation.

#### Volume:

Volume plainly put is the number of shares bought and sold for a given period. A large percentage price increase accompanied by a higher than average volume is a strong indicator of future price movements. A large percentage price movement accompanied by lower than average volume is a very weak indicator of higher prices, and is, in fact, an indicator that a correction in prices is possible.

Similarly, a large downward price movement accompanied by higher than average volume is a strong indicator that the stock will continue to move downward.

Most intriguing of all, however, is higher than average volume accompanied by no price movement. This generally indicates something happening behind the scenes, such as a news event or rumor, but the buying is not accompanied by market orders. Determining what is happening when accumulation of this kind occurs can be difficult, but sometimes rewarding.

#### Williams %R:

The Williams Percent probably requires a bit of explanation. Developed by Larry Williams, Williams %R is a momentum indicator. It is especially popular for measuring overbought and oversold levels. The scale ranges from 0 to -100 with readings from 0 to -20 considered overbought, and readings from -80 to -100 considered oversold.

Williams %R, sometimes referred to as %R, shows the relationship of the close relative to the high-low range over a set period of time. The nearer the close is to the top of the range, the nearer to zero (higher) the indicator will be. The nearer the close is to the bottom of the range, the nearer to -100 (lower) the indicator will be. If the close equals the high of the high-low range, then the indicator will show 0 (the highest reading). If the close equals the low of the high-low range, then the result will be -100 (the lowest reading).

It is important to remember that overbought does not necessarily imply time to sell and oversold does not necessarily imply time to buy. A security can be in a downtrend, become oversold and remain oversold as the price continues to trend lower. Once a security becomes overbought or oversold, traders should wait for a signal that a price reversal has occurred. One method might be to wait for Williams %R to cross above or

below -50 for confirmation. Price reversal confirmation can also be accomplished by using other indicators or aspects of technical analysis in conjunction with Williams %R.

One method of using Williams %R might be to identify the underlying trend and then look for trading opportunities in the direction of the trend. In an uptrend, traders may look to oversold readings to establish long positions. In a downtrend, traders may look to overbought readings to establish short positions.


### **My stocks**

For this project I decided to use only a handful of stocks. There are several reasons for this. First collecting twenty years of data on a stock is a very time consuming task. Although I pay a lot for access to that data most services don't offer the raw data, they provide the information in graphs, which requires a great deal of time to compute into data points. Second it is rather hard to find companies that have been around for twenty plus years are publicly traded entities and where all the necessary data was maintained. Volume in particular was hard to find in the eighties. The final reason I choose these stocks is that they are several of my larger holding at the moment in my long term portfolio. I should also note that I also decided to take the averages of all of my indicators for each quarter. I wanted to have somewhat manageable data files, 260 sets of data for each year would have been a headache.

#### General Electric:

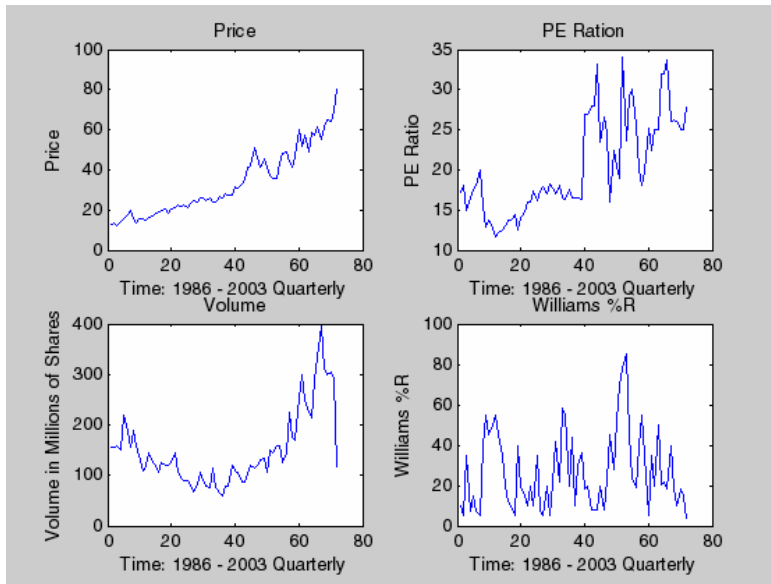
GE is a diversified technology and services company dedicated to creating products that make life better from aircraft engines and power generation to financial services, medical imaging, television programming and plastics. GE operates in more than 100 countries and employs more than 315,000 people worldwide.

The company traces its beginnings to Thomas A. Edison, who established Edison Electric Light Company in 1878. In 1892, a merger of Edison General Electric Company and Thomson-Houston Electric Company created General Electric Company. GE is the only company listed in the Dow Jones Industrial Index today that was also included in the original index in 1896.



3M:

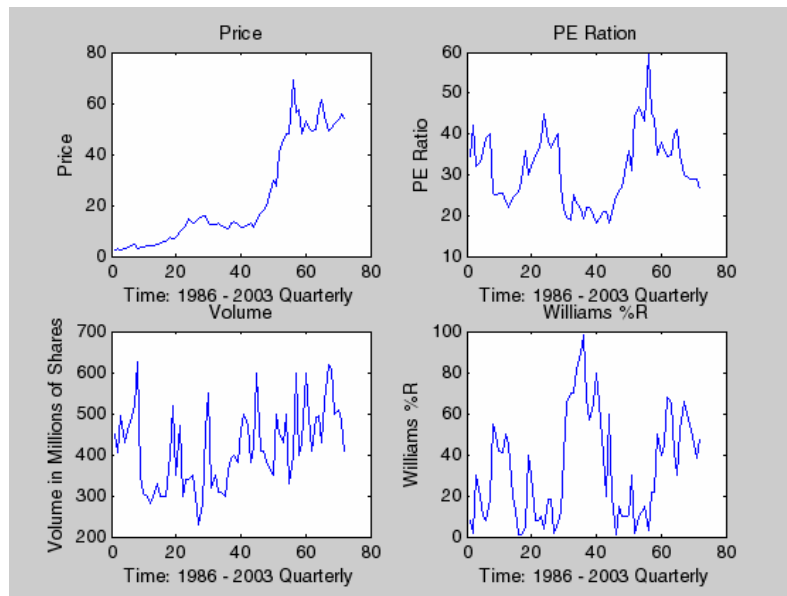
3M is a \$16 billion diversified technology company with leading positions in consumer and office; display and graphics; electronics and telecommunications; health care; industrial; safety, security and protection services; transportation and other businesses. Headquartered in St. Paul, Minnesota, the company has operations in more than 60 countries and serves customers in nearly 200 countries. 3M is one of the 30 stocks that make up the Dow Jones Industrial Average and also is a component of the Standard & Poor's 500 Index.



### Wal-Mart:

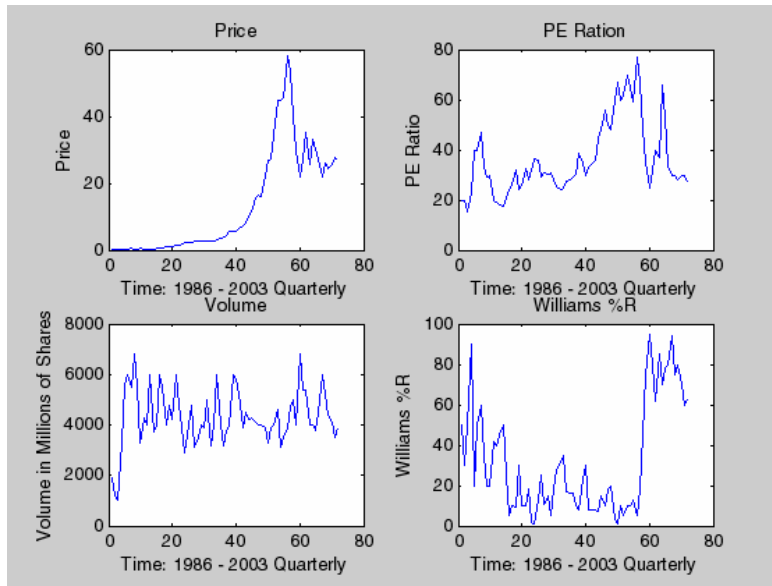
Wal-Mart Stores, Inc. is the world's largest retailer, with \$244.5 billion in sales in the fiscal year ending Jan. 31, 2003. The company employs more than 1.3 million associates worldwide through more than 3,200 facilities in the United States and more than 1,100 units in Mexico, Puerto Rico, Canada, Argentina, Brazil, China, Korea, Germany and the United Kingdom. More than 100 million customers per week visit Wal-Mart stores worldwide.

By the turn of the century Wal-Mart had been named "Retailer of the Century" by Discount Store News; made FORTUNE magazine's lists of the "Most Admired Companies in America" and the "100 Best Companies To Work For;" and was ranked on Financial Times' "Most Respected in the World" list. In 2002, Wal-Mart became No. 1 on the FORTUNE 500 list and was presented with the Ron Brown Award for Corporate Leadership, a presidential award that recognizes companies for outstanding achievement in employee and community relations. In 2003, Wal-Mart was named FORTUNE magazine's Most Admired Company in America.



### Microsoft:

Microsoft is of course the maker of Windows and office. Through the years Microsoft came to represent the technology age. As such it is probably my hardest to classify stock as it sky rocketed during the late nineties and took a nose dive after that.

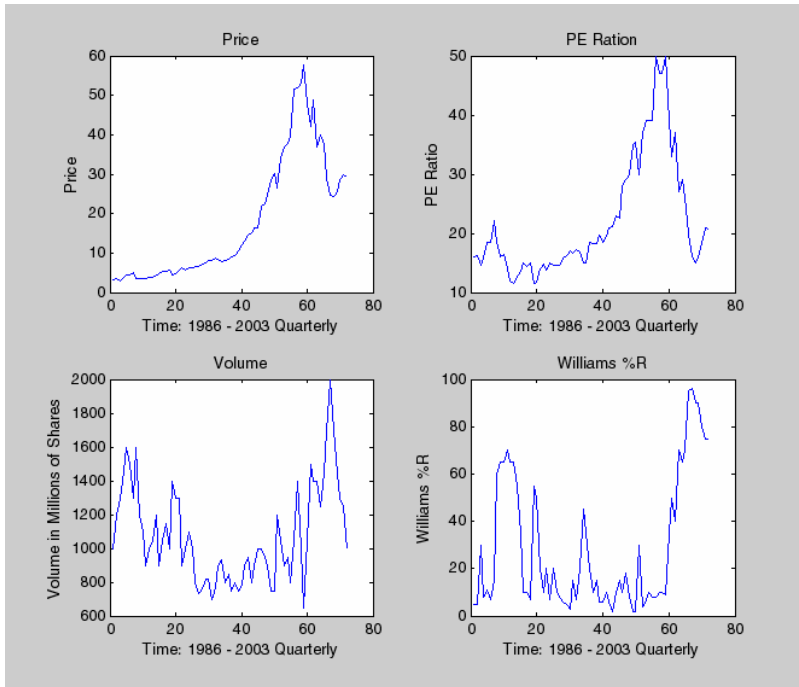


### Neural Network Program:

I choose to implement my MLP with a back propagation algorithm designed by Professor Hu. For the most part I kept the same program the professor designed; I did do some modifications so I could output predictions instead of classification. I picked BP because I felt it could detect the patterns better than the other algorithms I've learned about in class.

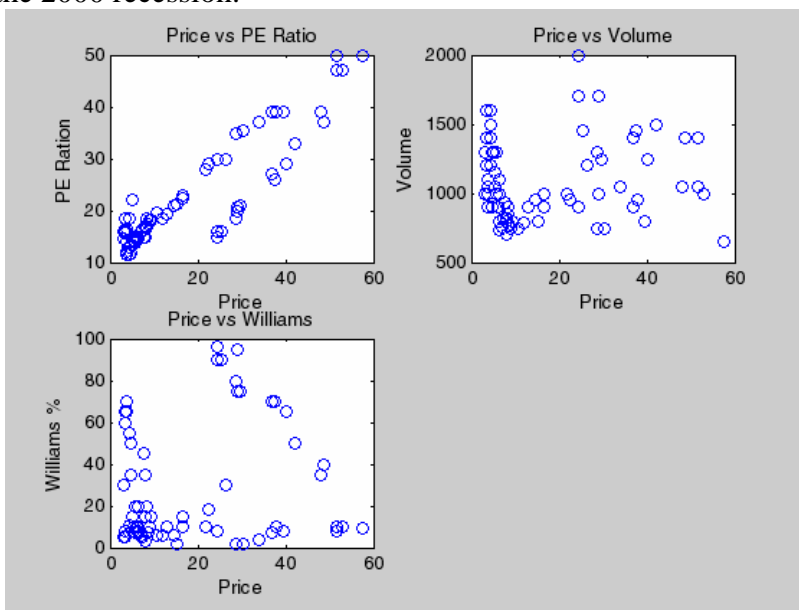
### Analysis and Preparation of Data

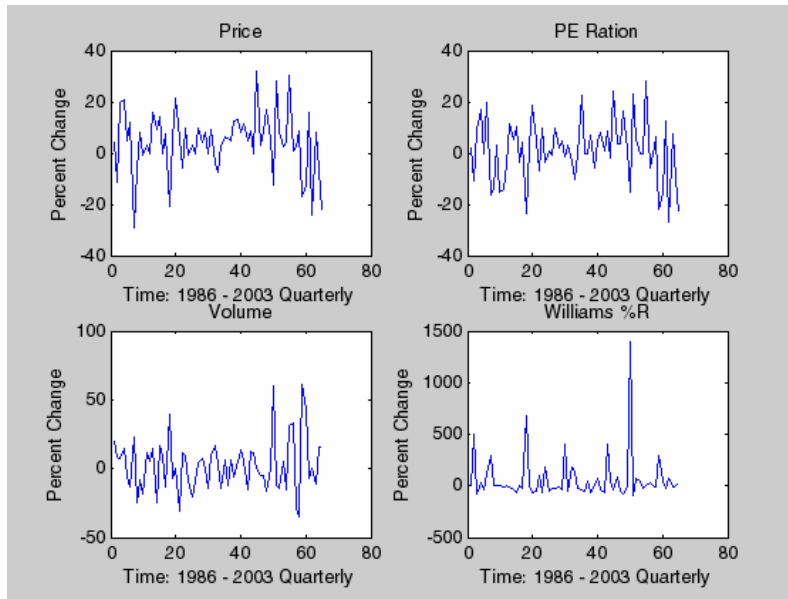
For the majority of my analysis I decided to use GE, although I did much of the analysis below for the other stocks as well, I tried the most things with GE. Following is the actual data collected quarterly over the period of 1986 to 2003. For the majority of my analysis I was looking for trends and patterns that might be useful to isolate. I also attempted to use clustering, but the data just didn't seem to lean one way or the other.



GE	Variance	Mean	Minimum	Maximum
Price	242.49	17.723	3	57.7
PE Ratio	100.69	22.066	11.5	50
Volume	81723	1080.1	650	2000
Williams	792.77	29.042	2	96

Right away I noticed that price and PE seem to trend together, which doesn't surprise me, you can also see the dip in the center and at the end that represents the early nineties and the 2000 recession.





When I started analyzing the data I began to realize that it might be difficult for a neural network to predict future performance when prices have been in a steady upswing over the last two decades. So I started to study the percent changes.

% Change	Variance	Mean	Minimum	Maximum
Price	143.18	4.1395	-29.412	32.121
PE Ratio	140.39	1.0473	-27.027	28.205
Volume	374.23	2.5332	-35	61.538
Williams	48908	62.17	-86.667	1400

#### Data Sets Chosen:

Looking at the data I decided that I wanted to try out four different data sets. The first set I decided I would just supply it the raw data for a full year and a buy or sell recommendation. To do this I used a simple matlab loop to compile the data into rows for each quarter. When I was looking at the data I noticed that percent change seemed like a reasonable idea. When you are comparing prices from 1986 to 2003 using the price by itself just didn't seem to make sense. The percent change for the other three items also looked interesting. So for my second data set I used matlab to create a percent change for each item over the period of a year. Finally when I was looking at the data I started to wonder if the Williams %R should be viewed as a percent change. It seemed to have incredibly large swings (max change 1400%), I started to wonder if I should take the percent change of the first three items but not Williams. So I set up a third data set to look at the percent change price, PE, and volume and the raw Williams number. For the final set I wanted to see if just the change in Williams number might give me better results.

#### Laying out the data:

I collected all of my data through stock charts and saved it in excel tables. Below is a small part of my GE data set:

	Price	PE	Volume	Williams
1986	3.25	16	1000	5
	3.38	16.3	1200	5
	3	14.5	1300	30
	3.6	15.8	1400	8
1987	4.35	18.5	1600	11
	4.55	18.5	1500	7
	5.1	22.1	1300	15
	3.6	18.5	1600	60
1988	3.37	16	1200	65

To prepare my data I first ran the raw data through matlab, comparing the price one year into the future to the “current price” if the price increased I marked it as a buy, if it went down I marked it as a sell.

For the first data set I just placed four consecutive sets of quarter data into one line and found the buy sell lines. The final quarter in each line would be compared to 4 quarters later to figure out whether to buy or sell. For the second data set I computed the percent change between consecutive quarter. The other data sets were designed in similar fashion.

## Experimenting

### Initial Results:

For my initial experiments I simply used the most generic design for BP I could think of, which meant going with as many defaults as possible. For each test I used the data from 1986 to 1995 to train and 1996 to 2003 for testing. To get an initial idea of which data set is the best I ran each data set 10 times with the default settings listed below.

### Default Settings:

Scaling, Input Layer = 2, Number of Neurons per layer = # of inputs, Default Activation function, Pattern Classification, Use Entire Training Set, Output Nodes Default, alpha = 0.1, mom = 0.8, nepochs = 2000, epoch size = default, Convergence Check = 5, n0 = 100

Results (All Results Averaged 10 runs):

Data Set 1: 62.069%

Data Set 2: 83.3333%

Data Set 3: 34.4828%

Data Set 4: 42.4325%

### Refining:

It was rather apparent that my third and fourth ideas were not going to work. Considering the high success rate of the second data set I decided to look further into data set 2. Although I tried several more designs I never found one that was nearly as good as data set 2. For the rest of the experiment I used data set 2’s design for all of the stocks. Below are a few of the results I got from running the GE stock through BP with data set 2.

layers		Alpha when layers = 3	
2	83.33%	0.001	80.56%
3	83.33%	0.01	91.67%
4	80.56%	0.015	86.11%
5	83.33%	0.02	88.89%
		0.5	83.33%
		0.2	75.00%
mom			
0.1	86.11%		
0.2	88.89%		
0.5	83.33%		

All of the above results were averaged over 10 runs. I tried dozens of different arrangements. I found though that with 3 layers (2 hidden, 1 output) and with an alpha of 0.01 as apposed to 0.1 and with all other defaults I consistently got the best results. After I got these results I decided to use the same settings for the other 3 stocks.

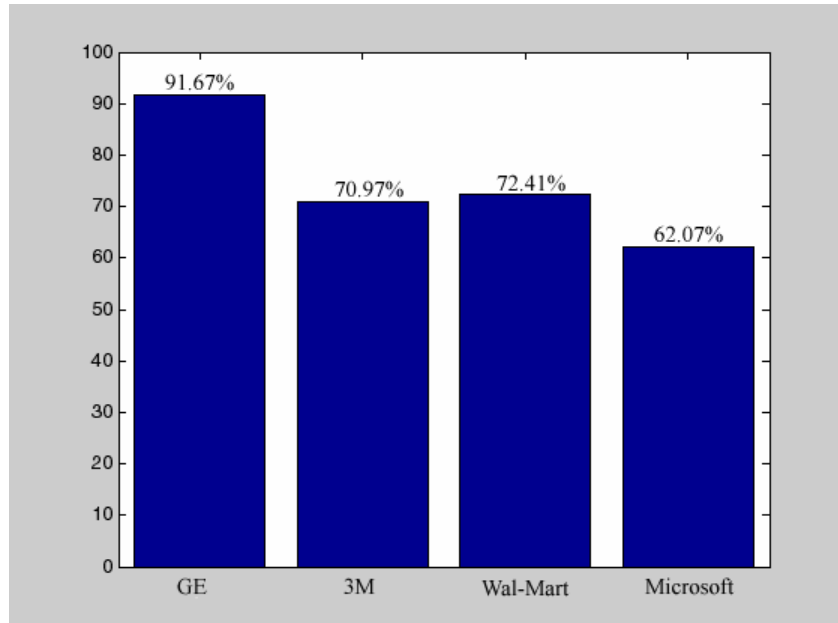
### Final Settings:

<i>Default Settings</i>	<i>Changes</i>	<i>Settings</i>
Scaling	Number of Layers	3
Tanh activation	Number of hidde	12
Pattern classification	Alpha	0.01
Use Entire Training Set	Number Epochs	2000
Scale Output	Convergence Check	5
Mom	No improvement	100
Epoch Size		

For the remaining stocks I initially ran them with the same settings, and found that the results I got were near or the best possible for all four stocks. This suggested to me that I had found a near optimal solution. As an interesting side note I also tried to train on GE stock and predict the results of another company's stock.

## Final Results


### Results:



The graph above shows the results for predicting the four stocks over the period of 1996 to 2003 based on data from 1986 to 1995 using the following BP settings: Scaling, Input Layer = 3, Number of Neurons per layer = 12, Default Activation function, Pattern Classification, Use Entire Training Set, Output Nodes Default,  $\alpha = 0.01$ ,  $\text{mom} = 0.8$ ,  $\text{nepochs} = 2000$ ,  $\text{epoch size} = \text{default}$ ,  $\text{Convergence Check} = 5$ ,  $n_0 = 100$ .

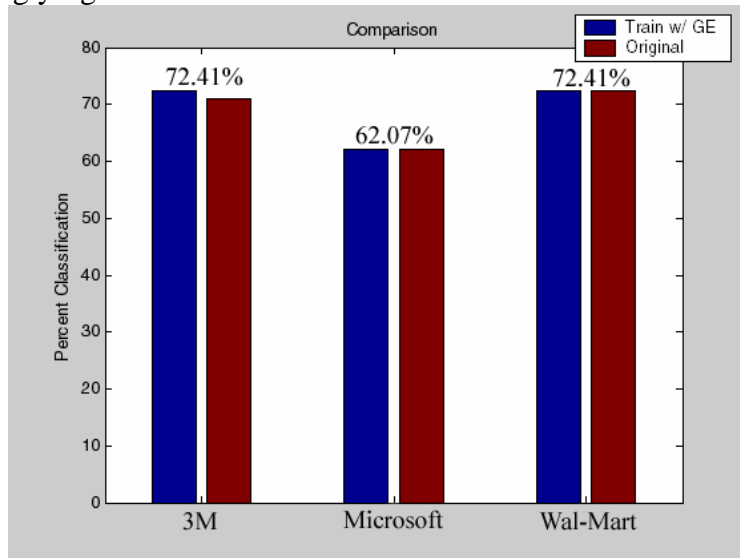
### Reversed Data;

I also tried to reverse the test and training files to see what kind of results I would get.



### Training with GE:


I decided to use GE to train the network and then testing with the other three stocks, surprisingly I got either the same or better results.



### Baseline Test:

I wanted to make sure the design I came up with was better than other possible designs. I decided to go with the professor's k-Nearest Neighbor Classifier. I ran the program for  $k = 1$  to 15 and then took the best classification rate.

	BP	KNN
GE	91.67	77.963
MMM	70.97	65.5172
WMT	72.41	65.0575
MS	62.07	61.6092



My technique won out on all four tests, although the Microsoft one was quite close. This is because Microsoft after the dot com bust is just about impossible to predict considering the data we have.

## Explanation of Results

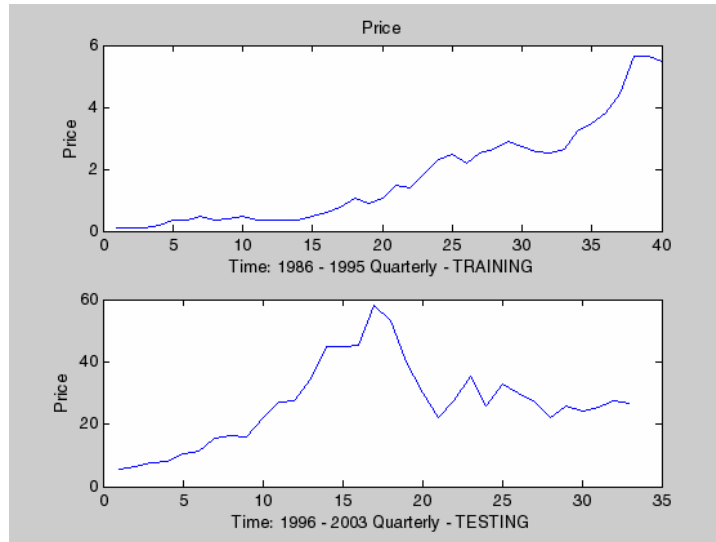
Although I was pleased with most of my results, I still would have liked higher numbers for some of the stocks. I was pleased with the results of 3M and Wal-Mart. Those two show that they probably just need a bit more training data. I doubt I have found the optimal configuration, but even after I tried these stocks on different designs my original results were almost always better than or about the same.

### **3M Results:**

I think there are reasons why I was unable to reach the GE level of accuracy for 3M. During the dot com boom 3M owned Palm, when 3M sold Palm several years ago investors no longer saw its PE ratio as worthy of such high levels. For the first few years that I was predicting 3M it was nearly a tech company and for the last few it was old time industrial giant. The training data that I used was built around a hundred year industrial company, but the testing area was around a tech company and then back to industrial. When I checked just the last few years with the same settings as before my prediction rates were much higher, nearly 80%. This leads me to conclude that although you can predict patterns with neural networks, you can predict one time wonders like the tech boom.

### **Microsoft Results:**

Microsoft was probably the hardest stock to predict. The company went public just before I started tracking the data and for the training period it enjoyed a steady increase. Microsoft came to epitomize the tech boom. As can be seen by the graphs below the Training data's price is the first graph and the Testing data's price is the second graph. Over the training period the price went up just about every quarter. While during the testing period we have a huge upsurge and then a drastic drop in price. Although I had hoped to be able to predict this, I suspect I needed more up and down data to even start to get a feel for detecting the upcoming downturn. This also shows that you can't use one set of years for every company. Many companies haven't even been around for 20 years, so you need to customize the design for each organization.



### **Wal-Mart Results:**

Wal-Mart presents a similar problem to Microsoft, it had an incredible growth pattern over the eighties and early nineties, in little over a decade it went from a small mom and pop store to the largest retailer in the world accounting for 10% of all retail sales in the US.

### **GE Results:**

I believe the reason GE was so successful for me is that investors have treated it as the same company for well over 20 years. GE is the only company that has remained on the Dow Jones for its hundred plus year existence. Microsoft and Wal-Mart only this last decade become top ten companies, before that they were hardly even known. I suspect if I went with older companies like AT&T and Coca-Cola I would get similar results to GE.

When I used GE for my final experiment GE's consistent results over the 80s and early 90s gave at least the same results for my three other companies. This further leads me to suspect that to get the most accurate results you need to compare companies in today's world against stocks that have performed in a similar fashion in the past. GE was one of the largest and most prosperous companies in the world in the 80s and it still is today. In other words I believe the approach I have designed is best suited for mature companies. To predict the future performance of new or emerging companies I suspect that using an index of similar stocks would give better results.

Although I was unable to recreate GE's success with the other three stocks. Seventy percent is still a great number in the stock market.

### **Conclusions:**

Overall my neural network setup was able to accurately predict if you should buy or sell a stock on a year's outlook on average 75% of the time. Compared to the baseline of the knn classifier which had an average of 67% I am quite pleased with the results. Most mutual funds are barely over 50% accurate with their predictions over the last two years.

The results really proved that neural networks could track patterns in the stock market. Obviously they aren't perfect but I don't think anything can predict so many

random events, from new innovations to terrorism. I believe a lot more work needs to be done in this area to get consistent results over a broad range of stocks, but I suspect that an average success rate of 80% or higher would not be impossible with enough research and data.

I'm also proud that I managed to show that you don't need a hundred different facts to predict the market. I managed to make a pretty accurate assessment with four pieces of data. I look at that alone as an accomplishment. I know that major models that the professional use often taken in dozens of different stats to get results like mine.

One of the most interesting things I took out of this was that when I tried to test the stocks based on the GE training I was able to get as good or better results out of all three stocks than I had when I trained on their own data. This just goes to reiterate the fact that there really is a pattern in the market the stocks don't even have to be similar to follow many of the same patterns.

That said I still believe that even more accurate results could be reached by testing a stock based on its entire industry's performance. Although stocks share many similarities from my experience I've seen that investors view stocks in different categories. I will continue to research this area, personally I've been quite lucky to have an average success rate of 70% for picking stocks over the last three years, but it would certainly be nice to find away to get up to 80% or so.

The stock market has always been more of a mental game then anything else. Prices are determined more on physiological value and what people are currently feeling than on any real concrete facts. That's why I think it is so appropriate that a program, designed to imitate the human mind, can not only detect the patterns of the market but actually see them better then the people buying and selling on all of the trading floors in the world.

## Buy or Sell?

So I'm sure if you've managed to read this entire report you've got to be wondering, should you buy or sell these stocks? I ran the most recent data I had into the network and was quite surprised to see that it mirrored the decisions I had already made.

Sell GE

Buy the rest!

I was quite surprised to see the MLP agreed with me on my own choices for the holidays especially on my choice of buying 3M. The stock has been on a tear lately, I've been going on more of a gut feel then any hard data that it will continue growing for the next year. Once exam season is over I think I'm going to enter in more data and see if I can't predict a few other stocks!