CSE 352 Artificial Intelligence PRESENTATION

Anita Wasilewska

ARTIFICIAL INTELLIGENCE WITHIN FINANCE

Team 1
Li, Roger
Li, Yan
Vinueza, Charles
Wu, Kalvin
Yuan, Reagan

REFERENCES

- Bloomberg Analytic Workshops. "Machine Learning in Finance 2016." April 2016
- Coyle, Damien; Prasad, Girijesh; McGinnity, Thomas. *Faster Self-Organizing Fuzzy Neural Network Training.* IEEE, May 2009.
- Eliaçik, Alpaslan Burak; Erdoğan, Erdoğan. "User-weighted sentiment analysis for financial community on Twitter." IEEE, November 2015.
- Glasserman, Paul. *Monte Carlo Methods in Financial Engineering*. Springer, 2010.
- Halton, John H. "A Retrospective and Prospective Survey of the Monte Carlo Method. January 1970. <http://epubs.siam.org/doi/abs/10.1137/1012001>
- Investopedia. "Monte Carlo Simulation." <http://www.investopedia.com/terms/m/montecarlosimulation.asp>
- Kuremoto, Takashi; Obayashi, Masanao; Kobayashi, Kunikazu. *Forecasting Time Series by SOFNN with Reinforcement Learning*. ResearchGate, September 2013.
- Murphy, Kevin P. *Machine Learning: a Probabilistic Perspective*. MIT Press, 2013.
- Murtha, James A. *Incorporating Historical Data Into Monte Carlo Simulation.* OnePetro, March 1994.
- Pavlyshenko, Bohdan. "Granger Causality Test for Frequent Itemsets in Financial Tweets". October 2015. <https://www.linkedin.com/pulse/granger-causality-test-frequent-itemsets-keywords-bohdan-pavlyshenko/
- Sanders Analytics LLC, Various datasets collected and used. <http://www.sananalytics.com/lab/>
- Stanford, "Deeply Moving: Deep Learning for Sentiment Analysis" <https://nlp.stanford.edu/sentiment/>

OVERVIEW

Applications of AI within Finance:

Sentiment Analysis

- Using natural language processing on news, journals, articles, and social media
- Using sentiment analysis to determine positive/negative sentiment on processed language
- Accurately represent the public opinion of asset valuations
- Correlate sentiment with actual assets/future movements via Granger Causality
- Using SOFNN learning to predict the movements with historic data and sentiment data
- Weighting various sentiment data via self correcting to more accurately aggregate sentiment analysis

Monte Carlo Simulations

- Understanding Monte Carlo simulations via the Unknown Bag approach
- The way influences/signals change Monte Carlo simulations
- Testing various signals against a simulation of a volatile and heavily influenced asset
- Using Bayesian networks to automatically draw correlations with various signals and price movements
- Using Bayesian networks to draw weighted influences based on various signals
- Using weighted signals to adjust Monte Carlo simulations for backtesting

Asset Pricing

How are stocks/assets priced within the market? Why are they worth their value?
- Valuations of the underlying asset (e.g. company) at hand
- Executive decisions and business plans
- Past performance and financial reports
- Technical analysis and pattern detection

All of this leads to an expected **future value** that determines the value of any asset
- Note that people always trade an asset for the value they believe it will be worth in the future, rather than the present value
- Assets are a generalized average of people's expectation of price

What if we could obtain more representative data sets for people's expectations?
More accurately measure and understand distribution of expectations?
More accurately predict future movements?

Purpose

If we could accurately determine the public sentiment, how does this help us?
We can attempt to draw correlations between measurements of sentiment and the actual performance of assets.

Agree with public sentiment:
- If the majority of people believe a stock is going up, we can invest in the stock with the idea that the majority of the public is correct

Disagree with public sentiment:
- If the majority of people believe a stock is going up more than it will, then we assume the price they currently have paid for the stock is higher than it will be

Or we can make use of AI to help make all these conclusions and correlations.

Natural Language Processing

Sentiment Analysis

We can train AI via deep learning to process millions of pieces specific to an asset we are interested in, to get an overall model of how people feel about that asset using natural language processing.

A more basic analysis would determine if an article was positive or negative about an asset.

A sample of NLP performed over several articles:
(using AAPL - Apple Inc.)

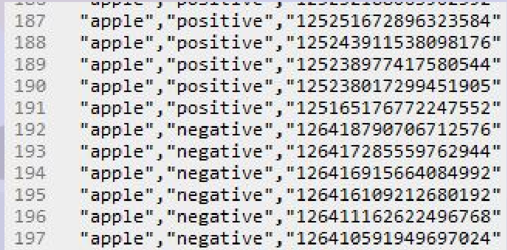

```
187   "apple","positive","125251672896323584"
188   "apple","positive","125243911538098176"
189   "apple","positive","125238977417580544"
190   "apple","positive","125238017299451905"
191   "apple","positive","125165176772247552"
192   "apple","negative","126418790706712576"
193   "apple","negative","126417285559762944"
194   "apple","negative","126416915664084992"
195   "apple","negative","126416109212680192"
196   "apple","negative","126411162622496768"
197   "apple","negative","126410591949697024"
```

Sample results from http://www.sananalytics.com/lab/

For a more fine tuned system, we use a scale to measure the given sentiment of a piece. We also use a tree which is built using natural language processing, and then evaluate the tree by using sentiment analysis on each branch

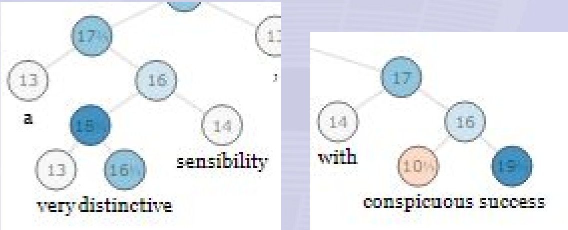

Samples from https://nlp.stanford.edu/sentiment/treebank.html

Sentiment Trees

We first extensively train an AI program via deep learning. We feed it various pieces and linguistics that allow it to determine how to analyze the sentiment of a given word, phrase, sentence, or article.

If we run the following sentence: "Apple has solid fundamentals and will rise in price."

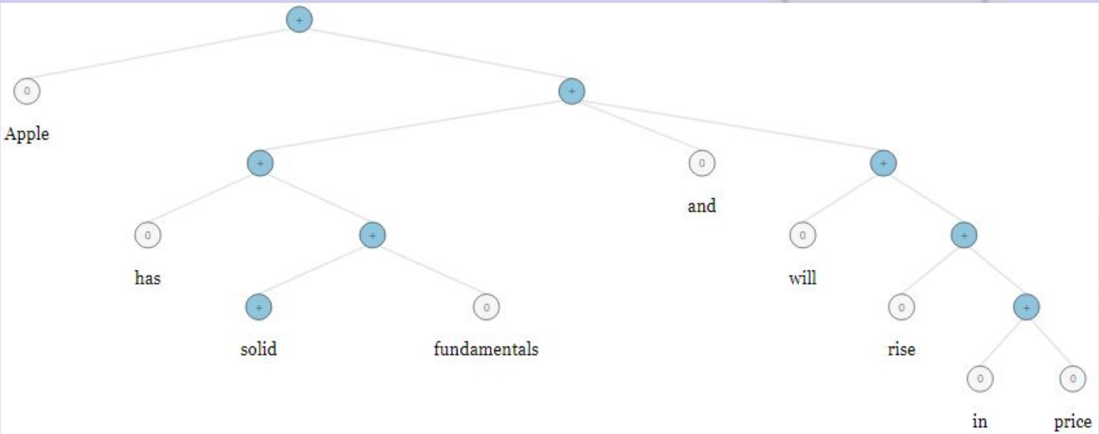

Stanford Natural Language Processing software

Notice the way the sentiment trees determines which nodes and branches are positive or neutral, and uses that to climb up the tree and return an overall result: positive

Sentiment Trees

"Apple has excellent fundamentals and will exceedingly outperform."

Notice the extremely positive result instead of the just positive result

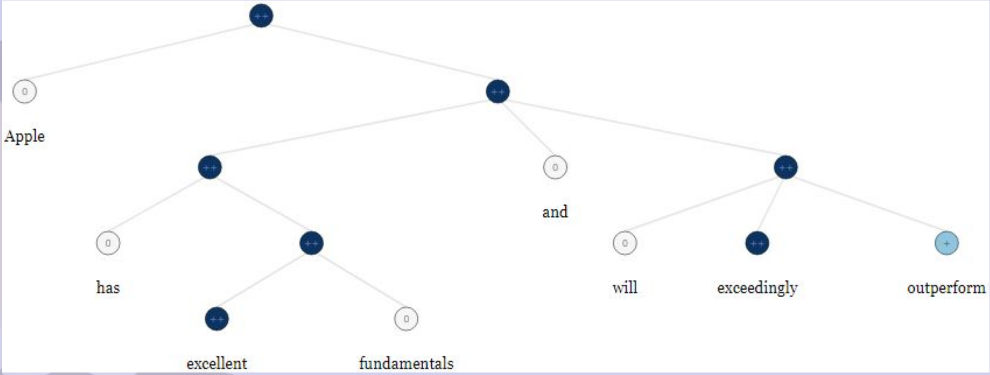

Stanford Natural Language Processing software

"Apple has horrible execution and will disastrously lose you money."

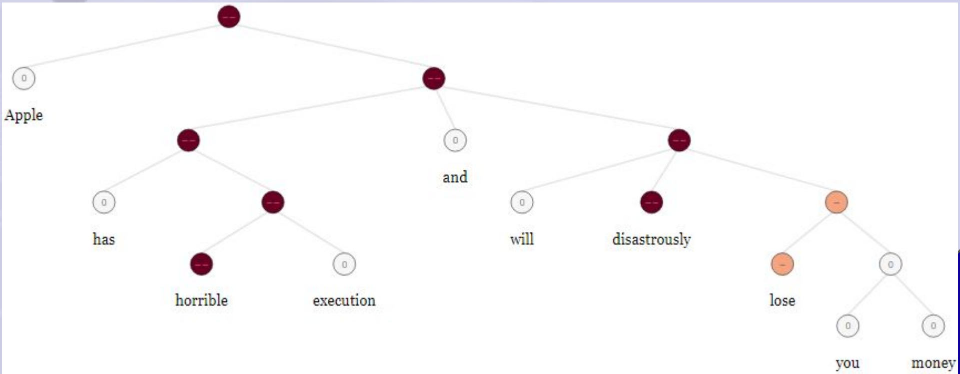

Stanford Natural Language Processing software

Granger Causality

The sentiment data we've collected: is it actually useful?
We can use a granger causality test: a method of determining how correlatable two data series are

A time series X_1 is considered a Granger-cause for a time series X_2 if we can show that movements of X_2 have some correlation or influence from movements in X_1

The following set of data shows Granger analysis on a sentiment time series of discussions on the Apple stock in comparison with the actual historical price of AAPL (Apple).

The resulting p-value was 0.002103, indicating the two charts represented matching data sets with on average 2-14 days of lag.

Notice that the sentiment series actually lags behind the stock chart, showing that the sentiment analysis could in theory have predicted a good majority of these price movements

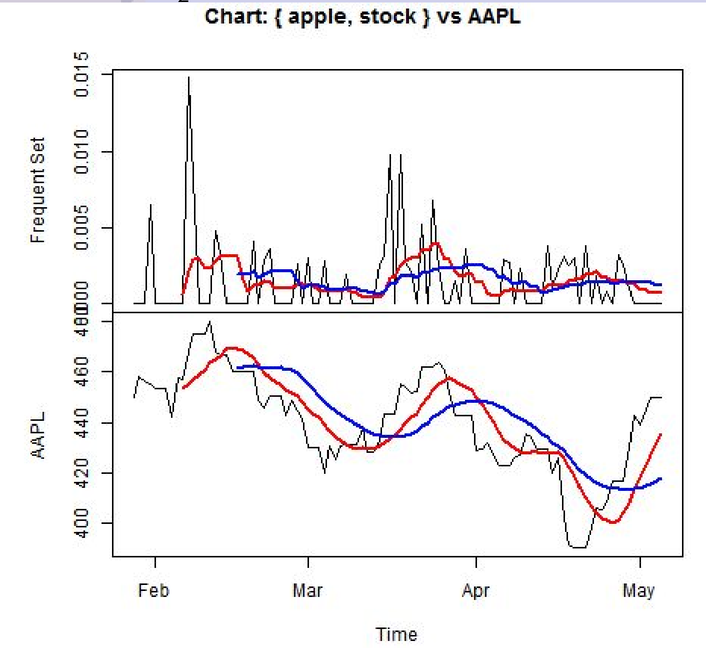

Granger Causality tests on Item Sets - Pavlyshenko

SOFNN

Self Organized Fuzzy Neural Networks use Stochastic Gradient Ascent, a learning algorithm that uses positive and negative reinforcement, to try to predict a future time series. The algorithm sets up a series of possible movements, and weighs each of them by using past experience and pattern detection.

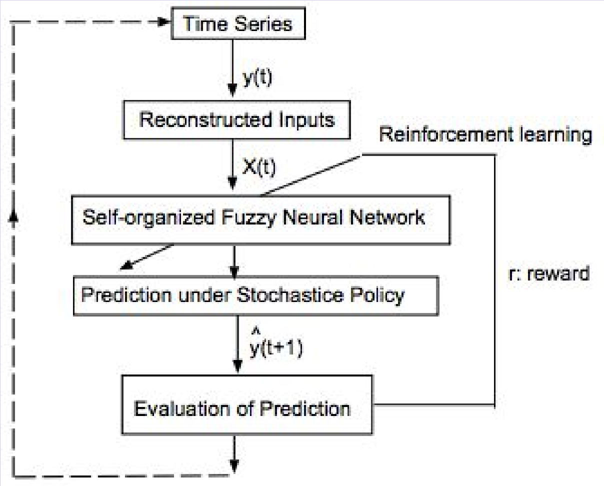

Forecasting Time Series by SOFNN - Kuremoto, Obayashi

Using the past two months of data for AAPL prices, we will use SOFNN to backtest and attempt to predict the chart

Data courtesy of NASDAQ

SOFNN Backtesting

Sentiment Analysis

By running a neural network on the historic price data for AAPL from 08/14 to 09/13, we can begin predicting price data on the day of 09/14. From there, we not only try to predict data, but continue training the neural network as time passes.

Our chart tends to lag. The algorithm seems to be accurately mapping the price movements, but later than we'd like.
Furthermore, our chart shows incredibly amounts of volatility. While there may seem to be a general trend following, we actually have a rather large number of incorrect day to day predictions with this chart.

Drawing a correlation chart, we notice there are quite a number of negatives - when our predictions go the exact opposite direction of the stock, which would net an obvious loss. 5 failures out of 20 is an incredibly unmanageable level of risk.

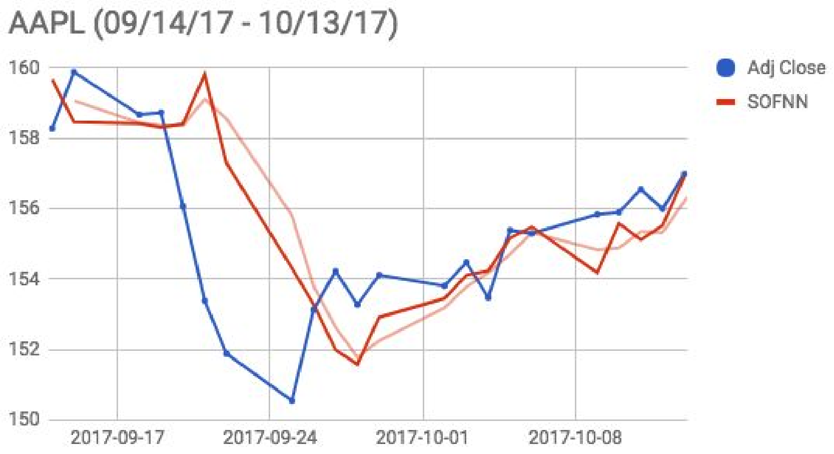

Faster SOFNN training - Coyle, Prasad - provided R program

SOFNN Backtesting

We now instead, train the SOFNN network with a sentiment analysis time series.

Notice the significant improvement of the algorithm, which manages to predict price movements on a day to day basis and shows no lag. This can be attributed to the fact that the price data is taken from closing price, the end of day price of AAPL. The sentiment analysis is performed throughout the day, allowing a "foreshadowing" of the movements.

By giving the sentiment analysis heavy weighting within the SOFNN learning process, we can also smooth out the prediction more as it relies less on historic volatile movements.

In the corresponding correlation chart, we also notice only a single failure to predict, a ~5% risk level from an albeit small sample size.

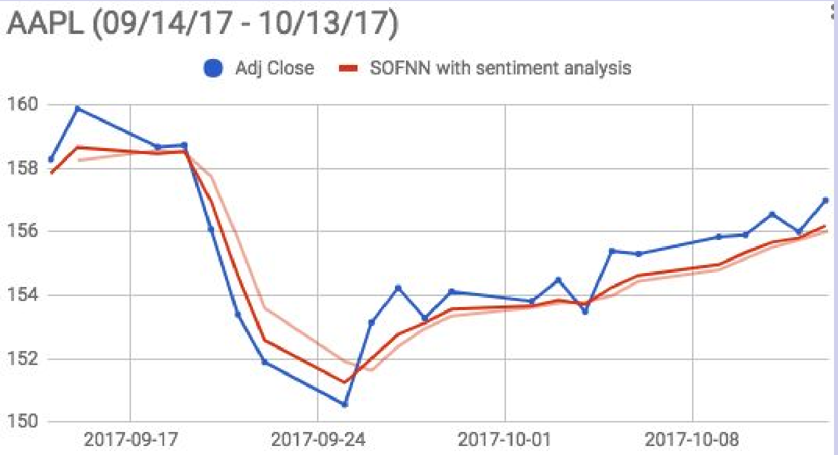

Faster SOFNN training - Coyle, Prasad - provided R program

Individual Sentiment Weighting

We have only been considering the overall sentiment that we receive, that is, the group public opinion of an asset.

Now, consider splitting researchers, publishers, and market analysts based on their sentiment and weighting them based on their results.

Consider the following shallow-learning model:
- We record a set of sentiment data from a various number of people.
- We can assign each person a starting weight
- We adjust that weight depending on the correctness of their sentiment in comparison with actual market movements.
- We then adjust our sentiment analysis to account for the weights

Sentiment Weighting Testing

We attempt to test weighting with AAPL (Apple Inc) for the fiscal year of 2014.
To simplify, we focus on the fiscal quarters which correspondingly returned:
- Q1-14: negative movement
- Q2-14: positive movement
- Q3-14: positive movement
- Q4-14: positive movement

Assuming we were sitting at the start of the fourth quarter, we could average their sentiments and make a decision on whether to buy AAPL or not within this timeframe based on that average.
The resulting average sentiment is slightly negative at -0.3, which would cause us to not buy the stock, missing out on a profit.

We instead compare the sentiment across each quarter (noted **S1-S4**) with the actual price movements.

This gives us a chart that shows a person's correctness (noted **C1-C3**) based on their sentiment across the first three fiscal quarters.

User	S1	S2	S3	S4	C1	C2	C3
1382248	1	0	1	0	N	N	Y
1575012	-1	1	0	0	Y	Y	N
1675375	-2	0	1	-1	Y	N	Y
1032794	0	-1	-3	-3	N	N	N
1502480	0	2	3	2	N	Y	Y
1731104	3	3	2	0	N	Y	Y
1709259	-2	-1	-2	-1	Y	N	N

Sentiment Weighting Testing

With our weighting system, we assign everyone an initial weight **W(i)** of 5.
At each quarter:
- if their sentiment is correct, we give them +2 weight
- if their sentiment is incorrect, we take -3 weight

We are purposefully more punishing to clear any inaccurate forecasters.

User	S1	S2	S3	S4	C1	C2	C3	W(i)	W1	W2	W3
1382248	1	0	1	0	N	N	Y	5	2	-1	1
1575012	-1	1	0	0	Y	Y	N	5	7	9	6
1675375	-2	0	1	-1	Y	N	Y	5	7	4	6
1032794	0	-1	-3	-3	N	N	N	5	2	-1	-4
1502480	0	2	3	2	N	Y	Y	5	2	4	6
1731104	3	3	2	0	N	Y	Y	5	2	4	6
1709259	-2	-1	-2	-1	Y	N	N	5	7	4	1
1361124	0	1	3	1	N	Y	Y	5	2	4	6
1773782	3	2	2	3	N	Y	Y	5	2	4	6
1141669	0	-2	-2	-3	N	N	N	5	2	-1	-4
1166438	2	0	-2	-3	N	N	N	5	2	-1	-4

Sentiment Weighting Testing

We now have a final weight for all our users. We ignore anybody with a negative weight.

User	S1	S2	S3	S4	C1	C2	C3	W(i)	W1	W2	W3	Final Weight	Weighted Sentiment
1382248	1	0	1	0	N	N	Y	5	2	-1	1	1	0
1575012	-1	1	0	0	Y	Y	N	5	7	9	6	6	0
1675375	-2	0	1	-1	Y	N	Y	5	7	4	6	6	-6
1032794	0	-1	-3	-3	N	N	N	5	2	-1	-4	0	0
1502480	0	2	3	2	N	Y	Y	5	2	4	6	6	12
1731104	3	3	2	0	N	Y	Y	5	2	4	6	6	0
1709259	-2	-1	-2	-1	Y	N	N	5	7	4	1	1	-1
1361124	0	1	3	1	N	Y	Y	5	2	4	6	6	6
1773782	3	2	2	3	N	Y	Y	5	2	4	6	6	18

By averaging the weighted sentiment for Q4, we are left with:

-0.3448275862	**Non Weighted Sentiment**
3.379310345	**Weighted Sentiment**

Notice that our weighted sentiment is overwhelmingly positive, as opposed to our original faulty non-weighted sentiment.
This tells us to buy the stock, a positive result given AAPL's positive movement in Q4.

Sentiment Weighting

Other applications of AI currently used with sentiment weighting:

Using certain sentiment keywords as indicators of better or worse prediction. Algorithms adapt to recognize certain sentiment keywords as not only holding more opinion, but also as being more strongly correlated with price movements.

Pattern recognition among sentiment data. An algorithm can notice cycles in sentiment, for example, as is common in the consumer industry. Sentiment can be accordingly considered more or less important depending on its position within any recognized sentiment pattern.

Future of Sentiment Analysis

The biggest consideration of sentiment analysis is the fact that while sometimes it may show very accurate results - it often is nothing more than a helper tool. It draws minor trends in price movements, but is an addition to the traditional fundamental analysis that is generally applied to asset pricing.

The greatest consideration is mapping all sentiment data to the corresponding fundamental analysis of any asset that helps determine exactly why and how sentiment analysis can help any trader make the right decisions.

While sentiment analysis is a well developed field, it is still a difficult goal to use AI to draw a solid correlation.

The Unknown Bag

Suppose you are given a bag of infinite marbles.
You are allowed to pull out a marble. If it's blue, you win a dollar. If it's red, you lose a dollar. Do you grab a marble? How many do you grab?

What if you were allowed to grab out as many marbles as you wanted to without taking that bet? In essence, you could do trial runs.

It would make sense to pull out a ton of marbles, count the colors, and use that data to determine whether or not the odds are in your favor.
While you will never know the definitive distribution of the blue or red marbles in this bag, you can run as many trial runs as possible to get a better and better idea.

Monte Carlo Simulations are a process of running a large amount of trials in an attempt to understand the performance of a random variable or time series.

Monte Carlo Simulations

The application of Monte Carlo Simulations to asset pricing prediction is much like running trial runs for the unknown bag.

Our aim is to look at all the various internal and external influences that can change the price of an asset. Internal influences can include executive decisions within a company, while external influences can be political power shifts such as the presidential election.

A monte carlo simulation can be run on these various movement influences to get an overall distribution of where the stock price will be in the future.

To the right is a sample of 30 possible projections for TWX (Time Warner Inc) over the course of 25 days

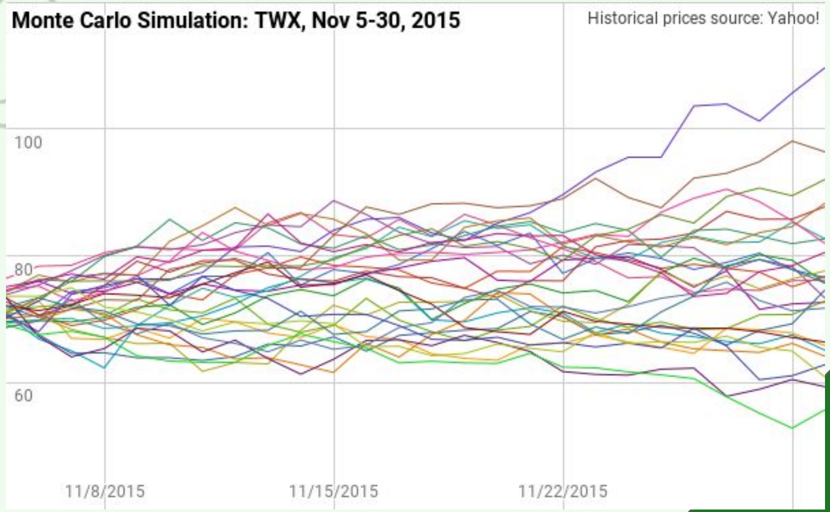

Projections taken from investopedia's sample simulation

Influences

Consider the various factors that are responsible for a stock's price moving:

- change in intrinsic value of the underlying company
- change in public sentiment regarding a company
- public information released, such as earnings report, merger announcements, business strategies, etc.
- shift in executive positions
- change in interest rates
- change in the global market
- change in political affairs
- risk of accidents, natural disasters, etc.

We draw our own Monte Carlo simulation across 30 days of USO (United States Oil Fund) by using some of these factors.
We simulate day to day changes, and use various influences to simulate how those changes are applied.

USO Influences/Signals

Monte Carlo

Influence #**1**: Sentiment fluctuation. We expect people to always have a slight difference in opinions of USO on various days, so we account for this by adding a daily percentage change with a standard deviation of 3%

Influence #**2**: On any given day, OPEC (Organization of the Petroleum Exporting Countries) could cut or increase oil imports from other countries. We say there is a 2% chance of a policy change on a day, and if there is, then the price would either go up or down by 10-15%.

Influence #**3**: Because of hurricane season, on any given day, there is a 0.2% chance that there is a hurricane that destroys a good portion of US oil mining facilities, causing the price to go down 30-50%.

Influence #**4**: A public oil report is expected on Day 20, which details the performance of oil production, earnings, revenue, and other finances. We allow USO a price move with a deviation of 20% on this day.

These influences are considered "signals." In accurate simulations, they would not be fixed percentages, but complex random variables that are dependent on other real world data. For example, the hurricane signal would be reliant on actual weather data and forecasted models.

USO Simulations

Monte Carlo

In the following chart you can see 25 separate simulations ran from day 0 to day 30.

Notice that around day 20, we see fluctuations due to our consideration of the earnings report.
We also see a massive fall in the pink line which fell to a possible hurricane around day 7.

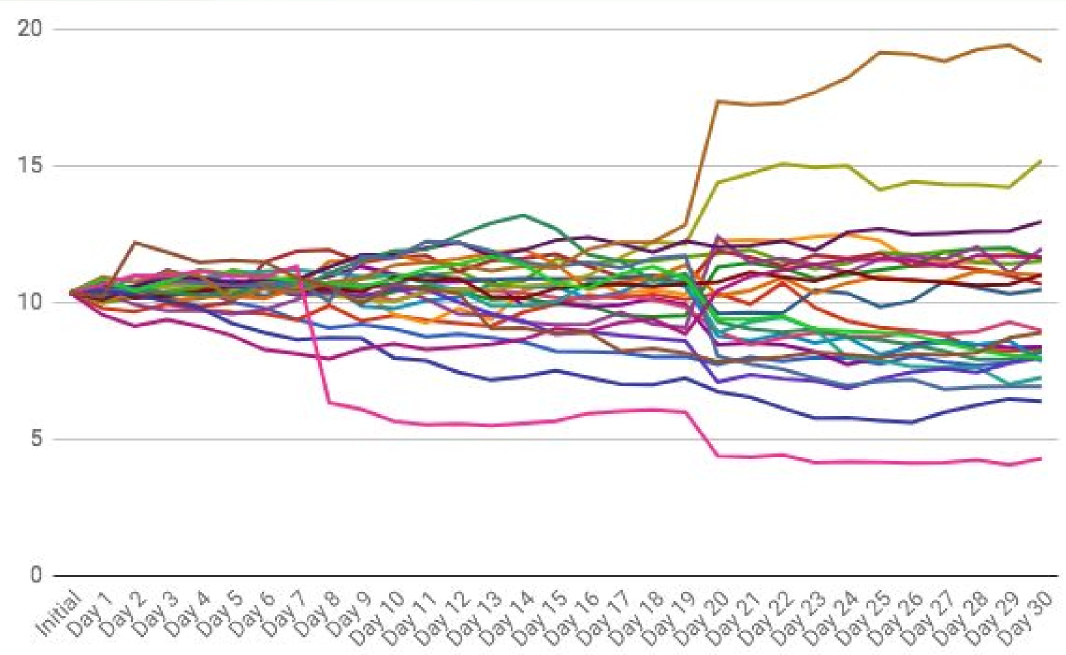

Lack of Signals

Theoretically, if we had properly accounted for all signals in our simulation, and then used the correct random distributions for all our signals, we would have a "perfect" Monte Carlo simulation, one that was able to perfectly model the future possibilities of an asset.

Unfortunately, each asset has an unimaginable number of signals that trigger price movements. Furthermore, even if we could list all the triggers, we would need to figure out the exact influence; how would this trigger cause an asset to move?

The application of AI can enable us to use machine learning to determine what real world triggers act as signals for our asset, and also how these triggers influence our assets' price movements.

Bayesian Signals

Bayesian networks are probabilistic models that draw dependencies between various random variables.
If we processed data related to various people's incomes, jobs, and their spending habits, we might end up with:

While some of these correlations are obvious, note that we did not have to inform the machine of these correlations, or even know them ourselves. The machine can automatically draw these correlations from data it has been trained on.

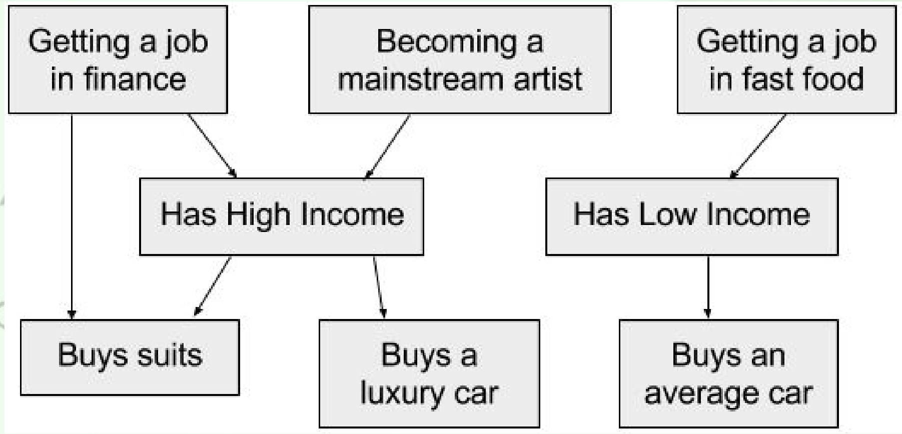

Sample Bayesian model from "Monte Carlo Methods in Financial Engineering"

Similarly, we can train a bayesian network with all raw data regarding changes in the world related to weather, politics, sports, or anything else, along with price charts. Ideally, the AI model can figure out if there exist any correlations.

Bayesian Modeling for USO

We run a Bayesian model specifically against historic price movements of USO (US Oil Fund). From the model, we note all the signal dependencies that it draws.

Dependency	Signal	Correlation
PRICE_MVMT_DOWN	NATURAL_HURRICANE	0.8615
PRICE_VOL_UP	NATURAL_HURRICANE	0.9780
PRICE_MVMT_ANY	NATURAL_WTHRCOND_ANY	0.1883
PRICE_MVMT_ANY	POLICY_CHANGE_OPEC_OIL_IMPORTS	0.4971
PRICE_MVMT_ANY	POLICY_CHANGE_OPEC_OIL_EXPORTS	0.4389
PRICE_MVMT_ANY	POLICY_CHANGE_USA_OIL	0.2711
PRICE_MVMT_ANY	POLICY_CHANGE_IRAN_OIL	0.1460
PRICE_MVMT_FOLLOW	ASSET_MVMT_NYSE:RIG	0.2457

"Monte Carlo Methods in Financial Engineering" - Glasserman

The AI model was capable of picking up several signals that we had previously not seen. Note that the price movements are correlated with changes in weather condition, not just natural disasters. Furthermore, there was correlation among other policy changes related to oil outside of just OPEC and the US. We also note that the index NYSE:RIG, which is Transocean, an oil rig company, was correlatable with some price movements of the oil index itself.

Bayesian Weighted Simulations

Monte Carlo

By allowing the Bayesian model to not only use AI to draw correlations, but to also use AI to draw the strength of various correlations, we can attempt to draw more accurate Monte Carlo simulations based on these new triggers and resulting influenced movements.

We adjust the simulation model to include new random variables that dictate day to day changes, including weather forecasts, foreign oil policy, and stock indices of other oil-related assets such as rig production companies, energy companies, etc.

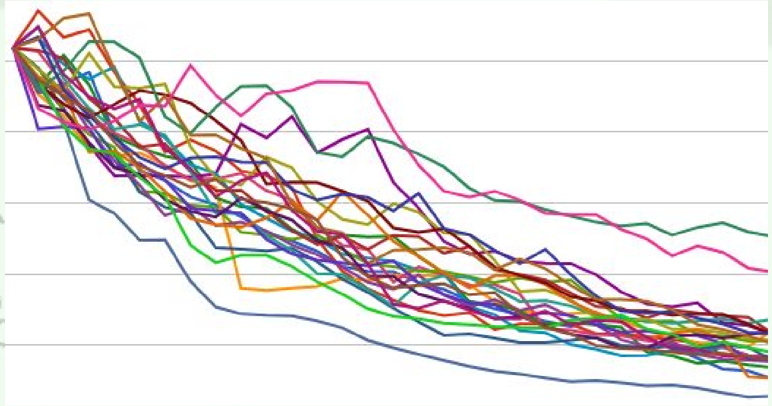

Backtesting the simulations onto historic price data for USO (specifically June 2014), we notice that our simulation tends towards the actual price movements of the stock.

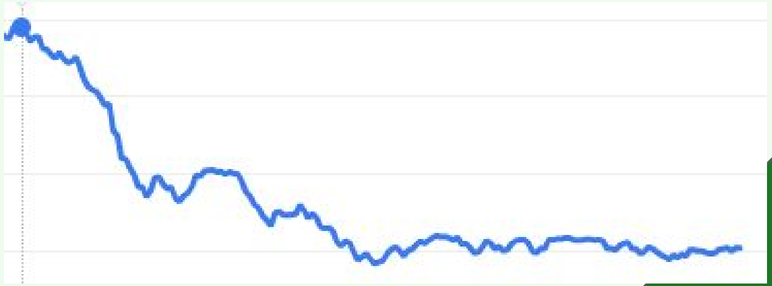

Future of Monte Carlo Simulations

The field of Monte Carlo simulations is one of the most widely used approaches in asset price modeling. It also holds a certain appeal due to the fact that if someone could accurately include all the signals and influences when modeling their simulation, they could theoretically perfectly predict the stock market.

While some obvious signals and resulting influences are obvious to us, they are never the only signals, and so the real problem is figuring out all the different signals, some of which may seem completely uncorrelated, go into the actual price movements of any asset.

With the recent trend in AI, machine learning, and model training, opportunities have opened up to making breakthroughs in furthering the accuracy of Monte Carlo simulations, by drawing correlations no other algorithm or human could make.

CONCLUSION

Finance has always been a field regarding strategic decision making and varying forms of analysis. The increase in technology has helped advance the field, but there still remains issues with market uncertainty and risk management.

The growing trend in AI has shown to have massive amounts of potential in regards to helping people with their analysis and drawing solutions within the field previously not thought possible with the use of large data processing, various machine learning algorithms, and neural networks.

As AI becomes more fine tuned and accurate to various data models, it could completely overtake any other form of financial analysis. Sentiment analysis and monte carlo simulations are a fraction of AI's applications within finance.