Master Engineering ELITE Programme **Sabatino Costanzo & Loren Trigo**
University of Toronto

ARTIFICIAL INTELLIGENCE IN FINANCE: FROM TREES TO DEEP LEARNING (APS 1052 H)

Important Notice:

Be advised that, due to the circumstances surrounding this edition of the course, ALL of the following administrative and academic guidelines are subject to adjustments and/or changes

Lectures:

Fridays / 12:00 M to 3:00 PM / May 29 to Aug 14 / 2020.

Instructors:

Sabatino Costanzo & Loren Trigo

Office Hours:

Every week on the day of the class we'll be available 'live' during 6 hours answering questions in 'Piazza' (from noon time to 6:00 PM). We may open another channel (conversation-oriented technology still to be selected) for less technical and/or less formal questions. Private and/or administrative questions will be handled through a dedicated email that will be provided on the first day of class.

Communication:

- The class slides and the careful transcription of the instructor's detailed comments on the slides (as well as the corresponding optional voice-over) will be uploaded to Quercus on the day of the class, at the time the class starts
- Technical questions will be addressed via 'Piazza'.
- Administrative issues of a more private nature will be addressed via a dedicated email address that will be provided by the instructors during the first class.
- General administrative questions may be addressed via an interactive 'voice' technology, still to be selected.

References:

The main literature will be comprised by the lecture slides and the meticulous transcriptions of the accompanying explanatory comments made by the instructors during previous live editions of this

course. If needed, these transcriptions can be 'listened-to' via the 'text-to-voice' option uploaded with the written course material. In short, there are no required textbooks for this course; the slides and the transcriptions of the slides' comments will suffice as such, and a bibliography will be provided.

Grading Scheme:

- Weekly Homeworks (Cumulative): 40%
- In-class Summary Exercises (Cumulative): 10%
- Final Project Deliverables: 50%

Homework:

- We will have homework assignments on most weeks.
- Homeworks will usually drill the techniques learned during the week in which the homework was assigned.
- You can work by yourself or as part of a team.
- Homeworks are very important because they give you the opportunity to apply the theory you've learned.

Teams:

We'll start the course with a survey that will allow us to build (and suggest to you) possible 'skill-balanced' teams, but you'll always have the option of working as an individual or coming-up with your own.

Attendance & In-Class Summary Exercises:

- We'll consider that you have "attended" the 'nth' class, as far as you can show that you have gone through the session corresponding to week 'n', and this is how you are expected to do it:
- We'll upload the 'nth session' of this course to Quercus on the day and at the time at which the session starts. Therefore, from that moment on you'll have available: (i) An extensive presentation of the topic planned for that week (ii) The transcripts of the comments we used to make in class on each one of the slides (iii) The option of activating a voice-over to read the explanatory text for you while you go through the slides making notes.
- As you go through the nth class (Slides & Text), once every few slides you'll be asked (in the text) to summarize in your own words (and to jot them down in word-pad or a similar text file), the main take away idea(s) presented so far and the main question(s) --if any-- that you may like to ask later in 'Piazza'.
- Once the class is over and you have completed the summary exercises corresponding to that class (each exercise should take about 10 minutes to complete), **you'll have until midnight of the next day** to upload them (all in a single document or txt file) to Quercus.
- Uploading **all** of the summary exercises you did on that day 'on time' (**i.e., before midnight of the next day**), guarantees to us that: (i) You went through the whole class attentively (ii) You stopped, reflected and wrote down periodically your insights and your questions at the crucial moments of the class.

- Typically, there will be 3 to 5 of those summary exercises per class, and their total cumulative value will constitute 10% of your final grade.
- To make sure that everything goes smoothly on the day of the class, we'll be present, answering questions 'live' in 'Piazza', from noon to 6:00 PM.
- At a certain point we were considering teaching this class 'live' online in order to maximize its 'interactive nature', but after testing the technology available we felt that giving a 'recitation' lecture through an imperfect audio system, an often unstable connection and having as an illustration resource a lagging computer screen, could hardly become a positive interactive experience.
- So we thought that it would be much more effective (and efficient) to provide an exhaustive account of [the best content from our class slides] + [the notes of our explanations & comments about those slides, meticulously collected along many editions of this course], and make the package available to you (the slides, the comments and a voice-over of those comments), so that you could control the speed of the information flow--, while, at the same time, on the day of the class, we made ourselves available 'live' for 6 straight hours to give support and answer questions about the content.

Final Project:

- The Final Project is definitely the highest point of the course.
- As soon as the theoretical bases have been covered in the first 6 or 7 class sessions and the basic practical experience has been acquired by doing the first 6 or 7 homeworks, the student will be ready to start deciding the subject of their Final Project.
- In due time we'll provide a vast, detailed list of possible projects to choose from, but we'll be always be open to listen to and to evaluate the viability of a student's own proposed project.
- Details about the expected deliverables regarding those projects will also be provided in due time.

--COURSE DESCRIPTION—

COURSE MAIN GOALS:

- **First:** *To provide a demystifying review of the conceptual basis of Machine Learning through the discussion of the intuition behind the powerful and innovative ideas underlying the main algorithms: from simple regression to the perceptron, and from there to the latest neural networks with memory*

- **Second:** *To provide a sound but accessible description of these algorithm's underlying mathematical structure, so that their potential innovation could be conceived and/or authoritatively discussed at the deepest level, at any forum or interview*

- **Third:** *To provide fluency at the handling of the coding tools used to implement, to test rigorously --and to adapt to the participant's needs-- the different working machine learning programs that will be provided by the instructors.*

- **Fourth:** *To teach through a hands-on approach how to operate these working machine learning programs to trade different asset categories (Stocks, Currencies, Crypto Currencies, ETFs, etc.) under real market constraints*

SUGGESTED TEXTBOOKS THAT ARE HELPFUL BUT NOT REQUIRED:

Sebastian Raschka and Vahid Mirjalili Python Machine Learning Packt edition.
We McKinney Python for Data Analysis O'Reilly edition.

COURSE OUTLINE

Artificial Intelligence as Machine Learning, evolving from single-layer Neural Networks to Deep Learning, is fast becoming relevant to stock market forecasting. Major advances in hardware and algorithmic improvements in the last 10 years has made it possible to run deep learning models --on a laptop-- that would have been intractable before. In this course we'll give an overview of several applications of machine learning to stock market forecasting (including high frequency trading), beginning with regressions, two "shallow" machine learning models (Support Vector Machines and basic Neural Networks) and ending with a deep learning model (Long Short Term Memory Networks). Each model is discussed in detail as to what input variables and what architecture is used (rationale), how the model's *learning progress* is evaluated and how machine learning scientists and stock market traders evaluate the model's *final performance*, so that by the end of the course, the students should be able to identify the main features of a machine learning model for stock market forecasting and to evaluate if it is likely to be useful and if it is structured efficiently in terms of inputs and outputs.

COURSE PREREQUISITES

Even though our course intends to be self-contained, it would be useful if the participant has previously seen ECE421H1S in which a linear version of some of the tools taught in our course are presented. However, if needed, we will provide a quick review of the content needed. Since the examples of our course will come from finance, it would be desirable for the student to be familiar with the basic structure of Capital Markets. It will be also helpful to know Python, but the objective of the course is not to learn to program (shallow & deep) machine learning models from scratch, but rather, to understand how they work and to learn to adapt them to the particular needs of the user and to optimize their application to stock market forecasting. The math. foundations of the basic

machine learning models (regression, neural networks & support vector machines) will be discussed and followed by a panoramic view of the inputs that are most likely to provide valuable information for stock market forecasting. Standard benchmarking methods used in the industry will be also covered. Subsequently, a number of basic --already programmed-- models will be discussed in detail and their performance evaluated.

COURSE OBJECTIVES

This course intends:

- To present the intuitive ideas behind the most important machine learning algorithms, their underlying mathematical structure and their code implementation
- To expose the student to "linear" and "kernelized support vector machines" and to provide an introduction to their use for more advanced classification tasks in finance
- To introduce the student to basic neural nets: the intuition behind them and their underlying mathematical structure; to their evolution from classic and logistic regression to the perceptron; to the use these tools for time-series forecasting and classification tasks in finance.
- To compare the results of time-series forecasting through examples of regressions, support vector machines and neural networks.
- To show how to select data inputs for time-series forecasting by introducing machine learning data pre-processing (Data Smoothing and Data Scaling & Normalizing Techniques) and feature engineering techniques (Feature Selection, Feature Combination, Principal Component Analysis)
- To introduce the main concepts related to regime indicators: "non-stationarity" of stock market returns & multiple regimes, unsupervised data clustering techniques like Gaussian mixes (continuous case), and k-means clustering
- To expose the student to model evaluation and improvement: cross validation, grid search, statistical evaluation metrics and scoring, "trading-based" evaluation metrics in model selection, and "horse race-type" tests: White's reality check for data snooping, Timmerman & Pessaran and Anatolyev & Gerko tests.
- To show how Deep Learning (Multiple Layer NNS, Convolutional NNS, Simple Recurrent NNS and Long Short Term Memory NNS) is used for time-series forecasting.
- How to do data pre-processing, feature extraction and model evaluation and improvement without breaking the continuity of the time series.
- If time allows, to introduce the basics of market microstructure theory: how a high-frequency limit order book works: iceberg orders and predictive features that can be extracted from the order book (e.g. bid-ask spread, bid-ask asymmetry etc.) to be used as inputs for Deep Neural Network algorithms.

COURSE STRUCTURE AND CONTENT:

This course consists of four parts, each one discussed in several weekly sessions:

- **First Part:** Reviewing the theoretical basis and the intuition behind Machine Learning
  (Overview of the intuitive ideas behind the basic machine learning algorithms, their underlying mathematical structure and their implementation in Python)
- **Second Part:** General applications of machine learning to financial markets
  (Simple regressions compared to support vector machines, trees & neural networks when forecasting time-series; data processing and 'feature extraction'; market model evaluation techniques)
- **Third Part:** Deep Learning in time-series forecasting, multi-layer neural networks

- **Fourth Part (3 Sessions):** Neural Nets with Memory and their applications to financial markets *(Co-Integration & Variance Ratio as statistical metrics of autocorrelation; machine learning algorithms with memory: RNNs, LSTMs and their uses in time-series forecasting, evaluation with walk-forward methods)*

LEARNING OUTCOMES:

By the end of the course students should be able to demonstrate that they:

- Are familiar with the **intuitive ideas behind basic neural networks and their underlying mathematical structure**, their evolution from the simple regression to the perceptron and from there to linear & non-linear (kernelized) support vector machines; that they understand the corresponding algorithms and how are these used for basic forecasting and classification tasks.
- Understand the differences between **simple regressions (benchmark) and support vector machines & neural networks when forecasting stock prices**; are familiar with the main techniques of **data processing and 'feature extraction'** through fundamental, technical and market regime indicators, the main **market model evaluation techniques** and their applications to market model improvement.
- Understand the basics of **statistical metrics of autocorrelation**. Why autocorrelation of stock returns creates the need for **machine learning algorithms with memory**. Understand the main idea behind the solution to the "exploding/vanishing gradient problem" which made "long short term memory neural networks" (LSTM) possible, and the intuition & structure behind **LSTMs and their new uses in finance.**
- Understand the basics of **market micro-structure theory**: how the limit order book works and how the **extraction of predictive features from the order-flow data** (which can be used as inputs for Deep Neural Networks models of price action), is done. **Understand model behavior during training through history plotting and how to reduce over-fitting with dropout regularization and learning rate schedules.**

REFERENCES

- Muller (2017) Introduction to Machine Learning with Python. Ch 3, 4, 5
- Aronson(2007) Evidence based Technical Analysis Chapter 6
- Han (2018) Aggregate Implied Volatility Spread and Stock Market Returns Rotman School of Management Working Paper No. 3047528 https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3047528##
- Zhang (2004) Prediction of financial time series with hidden markov models https://www.cs.sfu.ca/~anoop/students/rzhang/rzhang_msc_thesis.pdf
- Antonachi (2011) Optimal Momentum: A Global Cross Asset Approach https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1833722
- Piotroski (2000) Value Investing: The Use of Historical Financial Statement Information to Separate Winners from Losers https://www.ivey.uwo.ca/cmsmedia/3775523/value_investing_the_use_of_historical_financial_statement_information.pdf
- Beneish (1999) The Detection of Earnings Manipulation http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.195.3676&rep=rep1&type=pdf

- Altman (2000) Predicting Financial Distress of Companies https://pdfs.semanticscholar.org/3a40/ad1e6e88fc05ae19564fbd90bccae48accd1.pdf
- White, Timmerman (2002) Data-Snooping, Technical Trading Rule Performance, and the Bootstrap https://www.researchgate.net/publication/2550469_Data-Snooping_Technical_Trading_Rule_Performance_and_the_Bootstrap
- Anatolyev, Gerko (2005) A Trading Approach to Testing for Predictability https://amstat.tandfonline.com/doi/abs/10.1198/073500104000000640#.WuRq54jOWM8

- Naes et. al. (2006) Is the market microstructure of stock markets important? file:///C:/Users/Rosario/Documents/NeuralNetworksMachineLearning/microstructure-stock-market.pdf
- Borcheding (2017) The Value of True Liquidity https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3047528)
- Kercheval et.al.(2013) Modeling high-frequency limit order book dynamics with support vector machines https://www.math.fsu.edu/~aluffi/archive/paper462.pdf
- Han et. Al. (2015) Machine learning techniques for price change forecast using the limit order book data http://jcyhong.github.io/assets/machine-learning-price-movements.pdf
- Tsantekidis (2017) Forecasting Stock Prices from the Limit Order Book using Convolutional Neural Network https://www.researchgate.net/publication/319220815_Forecasting_Stock_Prices_from_the_Limit_Order_Book_Using_Convolutional_Neural_Networks
- Yoshida (2015) Short term stock price analysis based on order book information https://www.researchgate.net/profile/Kenichi_Yoshida4/publication/282989314_Short-term_Stock_Price_Analysis_Based_on_Order_Book_Information/links/599e3cb445851574f4b35f29/Short-term-Stock-Price-Analysis-Based-on-Order-Book-Information.pdf

--COURSE LAYOUT—

FIRST PART: REVIEWING THE THEORETICAL BASIS OF MACHINE LEARNING (3 SESSIONS):

Demystifying neural networks. Understanding the intuitive ideas behind them and their underlying math structure. Learning about their evolution from the simple regression to the perceptron and from there to the linear & non-linear support vector machines. Understanding how their algorithms work and learning how to use them for forecasting and classification tasks.

TOPICS OF DISCUSSION:

- Regression for classification and time series modelling.
- Neural network models for regression and classification, linear support vector machines.

SESSIONS:

- **Session 1:** introduction to basic neural nets: the intuition behind them and their underlying mathematical structure. Evolution from classic and logistic regression to the perceptron. How to use these tools for basic forecasting and basic classification tasks in finance.
- **Session 2:** the intuition and the underlying mathematical structure behind "linear support vector machines"; an introduction to their use for classification tasks in finance.
- **Session 3:** the intuition and the underlying mathematical structure behind "kernelized" support vector machines and an introduction to their use for advanced classification tasks in finance.

REFERENCES:

- Taylor (2017) Neural Networks, a visual introduction Stages1-5
- Muller (2017) Introduction to Machine Learning with Python. Ch. 2
- Scikit-learn Documentation on Support Vector Machines http://scikit-learn.org/stable/documentation.html

SECOND PART: GENERAL APPLICATIONS OF MACHINE LEARNING TO FINANCIAL MARKETS (4 SESSIONS):

Comparing the effectiveness of simple regressions (benchmark) with support vector machines & neural networks when forecasting stock prices. Learning the main techniques of data processing and 'feature extraction' through fundamental, technical and market regime indicators. Learning how to apply the main market model evaluation techniques and their applications in market model improvement.

TOPICS OF DISCUSSION:

- Regressive methods compared to "Machine learning based" stock price forecasting
    - Forecasting Stock Prices with regressions

    - Forecasting Stock Prices with Support Vector Machines
    - Forecasting Stock Prices with Neural Networks
- Data pre-processing techniques, feature extraction
    - Technical indicators: MACD, AMX-DMI, RSI, SAR, HURST exponent, regime indicators etc.
    - Fundamental indicators: Put Call implied volatility spread, Piotroski, Altman and Beneish scores, variations in cash flow from operations.
- Unsupervised data clustering techniques
    - Principal component analysis
    - Regime indicators: Hidden Markov Models (HMMs), Gaussian mixes, K-Means Clustering
- Model Evaluation and Improvement
    - Cross Validation
    - Grid Search
    - Evaluation Metrics and Scoring
        - Metrics for Binary Classification and Regression
        - Using Evaluation Metrics in Model Selection
        - "Horse race-type" tests: White's Reality Check for Data Snooping, Timmerman & Pessaran and Anatolyev & Gerko tests.

SESSIONS:

- **Session 4:** comparing regression with "machine learning based" stock price forecasting: stock price forecasting with regressions, with support vector machines and with neural networks.
- **Session 5:** data selection for stock price forecasting, machine learning data pre-processing techniques, feature extraction through technical indicators (e.g. MACD, AMX-DMI, RSI, SAR, Hurst exponent as a regime indicator etc.), fundamental indicators (put-call implied volatility spread, Piotroski, Altman and Beneish scores, etc.), principal component analysis for feature and portfolio selection.
- **Session 6:** regime indicators: the non-stationarity of stock market returns & multiple regimes, unsupervised data clustering techniques, hidden Markov models (HMMs), Gaussian mixes (continuous case), and k-means clustering
- **Session 7:** model evaluation and improvement: cross validation, grid search, evaluation metrics and scoring, which includes: metrics for binary classification & regression using trading-based evaluation metrics in model selection, and "horse race-type" tests: White's reality check for data snooping.

REFERENCES:

- Muller (2017) Introduction to Machine Learning with Python. Ch 3, 4, 5
- Aronson(2007) Evidence based Technical Analysis Chapter 6
- Han (2018) Aggregate Implied Volatility Spread and Stock Market Returns Rotman School of Management Working Paper No. 3047528 https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3047528##
- Zhang (2004) Prediction of financial time series with hidden markov models https://www.cs.sfu.ca/~anoop/students/rzhang/rzhang_msc_thesis.pdf
- Antonachi (2011) Optimal Momentum: A Global Cross Asset Approach https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1833722
- Piotroski (2000) Value Investing: The Use of Historical Financial Statement Information to Separate Winners from Losers

https://www.ivey.uwo.ca/cmsmedia/3775523/value_investing_the_use_of_historical_financial_statement_information.pdf

- Beneish (1999) The Detection of Earnings Manipulation
  http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.195.3676&rep=rep1&type=pdf
- Altman (2000) Predicting Financial Distress of Companies
  https://pdfs.semanticscholar.org/3a40/ad1e6e88fc05ae19564fbd90bccae48accd1.pdf
- White, Timmerman (2002) Data-Snooping, Technical Trading Rule Performance, and the Bootstrap
  https://www.researchgate.net/publication/2550469_Data-Snooping_Technical_Trading_Rule_Performance_and_the_Bootstrap
- Anatolyev, Gerko (2005) A Trading Approach to Testing for Predictability
  https://amstat.tandfonline.com/doi/abs/10.1198/073500104000000640#.WuRq54jOWM8

THIRD PART: DEEP NEURAL NETWORKS IN HIGH FREQUENCY TRADING (2 SESSIONS):

Understanding market micro-structure theory: how the limit order book works and how to extract predictive features from the order-flow data that can be used as inputs for machine learning models of price action. Understanding model behavior during training through history plotting; using dropout regularization and learning rate schedules to improve performance.

TOPICS OF DISCUSSION:

- Market microstructure theory: physiology of a high-frequency limit order book
  - Iceberg orders
  - Predictive features that can be extracted from the order book (e.g. bid-ask spread, bid-ask asymmetry etc.)
- Deep Neural Networks for high frequency price forecasting
- Understanding model behavior during training through history plotting
- Reducing over-fitting with Dropout Regularization
- Lifting Performance through Learning Rate Schedules

SESSIONS:

- **Session 8:** market microstructure theory: how a high-frequency limit order book works: iceberg orders and predictive features that can be extracted from the order book (e.g. bid-ask spread, bid-ask asymmetry etc.) to be used as inputs in machine learning algorithms.
- **Session 9:** Deep Neural Networks for high frequency price forecasting: data pre-processing of limit order book data, feature extraction, model evaluation and subsequent performance improvement through Dropout Regularization and Learning Rate Schedules.

REFERENCES:

- Naes et. al. (2006) Is the market microstructure of stock markets important? file:///C:/Users/Rosario/Documents/NeuralNetworksMachineLearning/microstructure-stock-market.pdf
- Borcheding (2017) The Value of True Liquidity https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3047528)
- Kercheval et.al.(2013) Modeling high-frequency limit order book dynamics with support vector machines https://www.math.fsu.edu/~aluffi/archive/paper462.pdf
- Han et. Al. (2015) Machine learning techniques for price change forecast using the limit order book data http://jcyhong.github.io/assets/machine-learning-price-movements.pdf
- Tsantekidis (2017) Forecasting Stock Prices from the Limit Order Book using Convolutional Neural Network https://www.researchgate.net/publication/319220815_Forecasting_Stock_Prices_from_the_Limit_Order_Book_Using_Convolutional_Neural_Networks
- Yoshida (2015) Short term stock price analysis based on order book information https://www.researchgate.net/profile/Kenichi_Yoshida4/publication/282989314_Short-term_Stock_Price_Analysis_Based_on_Order_Book_Information/links/599e3cb445851574f4b35f29/Short-term-Stock-Price-Analysis-Based-on-Order-Book-Information.pdf

FOURTH PART: OTHER DEEP LEARNING APPLICATIONS TO FINANCIAL MARKETS (3 SESSIONS):

Understanding statistical metrics of autocorrelation. Understanding why autocorrelation of stock returns creates the need for machine learning algorithms with memory, the solution to the exploding/vanishing gradient problem that made (functional) "long short term memory neural networks" (LTSM) possible, and the intuition & structure behind LSTMs and their uses in finance.

TOPICS OF DISCUSSION:

- Understanding statistical metrics of autocorrelation.
- Introduction to LSTMs and time series forecasting (single & multi-variable examples) using KERAS.
- LSTM neural network for stock price forecasting.
- LSTM, Convolutional NN, Gated Recurrent Unit for crypto currency price forecasting.
- Data pre-processing, feature extraction, over-fitting and benchmarking for deep learning models.

SESSIONS:

- **Session 10:** understanding statistical metrics of autocorrelation. Introduction to LSTMs and time series forecasting (single & multi-variable examples) with KERAS.
- **Session 11:** LSTM neural networks for stock price forecasting.
- **Session 12:** LSTMs for crypto currency price forecasting, data pre-processing, feature extraction, model evaluation and improvement without breaking the continuity of the time series.

REFERENCES:

- Chollet (2018) Deep Learning With Python, Chapter 6: Deep learning for text and time series

--COURSE INSTRUCTORS—

Sabatino Costanzo-Alvarez

Sabatino Costanzo-Alvarez holds a Masters in Economics and Finance from Brandeis University as well as a Magister Scientiarum, a Magister Philosopharum and a Ph.D. in Mathematics from Yale University, where in 1990 achieved a significant breakthrough by solving a mathematical conjecture which had remained unsolved for more than 3 decades. Taught Mathematics of Finance at Boston University as an Associated Professor for 5 years and later co-founded the Boston Trading Group LLC, designed the trading systems used in the firm's daily Futures Trading Operations and acted as head trader of the team. Holds the licenses "Registered Representative NYSE/NASDAQ" (Series 7), "Registered Financial Advisor", "Registered Uniform State Law Securities Agent", "Registered Managed Futures Fund Representative" in the U.S. and "Canadian Securities Course" & "Conduct and Practices" in Canada, as well as products training at Morgan Stanley in Boston, and later at Merrill Lynch in New York. Chaired the Advanced Management Program for Senior Executives (PAG), an Executive MBA at the US Accredited IESA Institute in Caracas, where he taught Financial Engineering and Investment Management as an Associate Professor, and tutored over 70 MBA dissertations. Acted as Head of Research at Econo Invest C.A., the largest Investment Firm in Venezuela, leading the Investment Strategy Team in charge of generating and executing the U.S. & E.U. investment strategies for Commodities, Fixed Income Instruments and Equities for the firm (published weekly in Bloomberg), as well as generating and maintaining the Sovereign Fixed Income Indexes of Brazil, Colombia, Mexico, Peru, Chile, Uruguay and Venezuela to be used in the design of international financial products. Acted as an Investment Advisor for the International Wealth Management Groups at Morgan Stanley (Boston), Merrill Lynch (NY) and the Royal Bank of Canada(Toronto), and is now a Senior Partner at the Toronto boutique Investment Firm Inter Alea, where he provides state-of-the-art mathematical modeling solutions to portfolio and risk management problems for a select group of corporate and high net worth private clients, designing and managing their investment portfolios based on their specific risk & return requirements. He teaches Portfolio Management, Statistics & Mathematical Modelling and Business Mathematics Courses at the Pilon School of Business, and is the founder and advisor of the Sheridan Students Trading and Investment Association. He is a Lecturer at the U of T Graduate School, where he is teaching Portfolio Management, Blockchain Technology, Cryptocurrencies and Artificial Intelligence applied to Finance.

Rosario Lorenza Trigo-Ferre

Holder of a B. A. in Philosophy (Magna Cum Laude) from Yale University -where she also received training in Math & Physics-, a Ph.D. in Generative Linguistics from Massachusetts Institute of Technology (MIT) and a M. Sc. in Management of Information Systems from Boston University ("Beta Gamma Sigma Honors" award), she was a Professor at Boston University for 8 years. While a Programmer Analyst at Boston University, she designed and developed an application for the management of accounts trading stock and currency futures and co-designed financial applications under the direction of Professor Zvie Bodie at B.U. Co-founder and Trader at the Boston Trading Group and Certified Programmer Analyst in e-commerce by the University Computer Careers Program, she generated the trading signals for currencies and metals futures used in the BTG's market operations; developed an application maximizing the efficiency of trading system for currency and metal futures, and designed a client-server application for the management and operation of trading accounts. Has designed and developed many multi- tiered e-commerce applications dynamically generated from databases. Project leader and senior programmer analyst at IngeDigit, designed and developed internet applications for banking accounts management & operation, and for international transactions between banking accounts and credit cards. She was a Professor at the Department of Production and

Technical Innovation of the IESA Institute, the top -only US accredited- Venezuelan Business School, where taught courses in Information Systems, Simulation in Finance, Operations and Database Marketing. She is the author of many scientific papers in refereed journals and a Permanent Consultant for an international development bank (C.A.F, The Andean Region Development Bank), where she has designed the financial models used to evaluate the profitability, coverage and socio-economic impact of projects like the inclusion of fiber-optic cable in highways in Colombia and Peru. These models led to the enactment of new laws making such inclusion mandatory in the Andean region. Also designed the financial models used to evaluate the profitability of projects in satellite technology in Argentina (specifically the ARSAT program) by estimating the future regional demand for transponders and the impact of the project in the input-output matrix of the country, and is now a Partner at the boutique Investment Firm InterAlea, where she designs, develops, tests and implements trading and risk management strategies based on the entropy analysis of price signals, executed on stock quote-data processed through SQL-Server. She is a Lecturer at the U of T Graduate School, where she is teaching Portfolio Management, Blockchain Technology, Cryptocurrencies and Artificial Intelligence applied to Finance.