Applications of Machine Learning Algorithms to Finance
L'application des algorithmes de l'apprentissage automatique en finance
(Org: **Michael Chen** (York University) and/et **George Lai** (Wilfrid Laurier University))

FREDERIC GODIN, Concordia University
Equal Risk Pricing of Derivatives with Reinforcement Learning

The equal risk pricing methodology for derivatives pricing is introduced. The implementation of the approach based on reinforcement learning with neural networks is discussed. Results from numerical experiments assessing the impact of the risk measure serving as the objective function, the underlying asset model choice and the selection of hedging instruments are presented. The approach is also benchmarked against traditional pricing methods.

The talk is based on the following papers:

Carbonneau, A., & Godin, F. (2021). Equal risk pricing of derivatives with deep hedging. Quantitative Finance, 21(4), 593-608.

Carbonneau, A., & Godin, F. (2021). Deep Equal Risk Pricing of Financial Derivatives with Multiple Hedging Instruments. arXiv preprint arXiv:2102.12694.

Carbonneau, A., & Godin, F. (2021). Deep equal risk pricing of financial derivatives with non-translation invariant risk measures. Working paper.

ANASTASIS KRATSIOS, ETH Zurich
Universal Probability Measure-Valued Deep Neural Networks

We introduce deep neural architecture types with inputs from a separable and locally-compact metric space X and outputs in the Wasserstein- 1 space over a separable metric space Y. We establish the density of our architecture type in C(X; P1(Y)), quantitatively. NB that our results are new even in the case where X and Y are Euclidean, in which case, we find that many commonly used types such as MDNs and MGANs are universal special cases of our model type. We show that our models approximate functions in C(X; P1(Y)) by implementing e-metric projections in the Wasserstein- metric onto the hull of certain finite families of measures therein. If the target function can be represented as a mixture of finitely many functions, each taking values in a finite-dimensional topological submanifold of the Wasserstein space, we find that the approximating networks can be assumed to have bounded width. As applications of our results, we address the following problems. We show that, under mild conditions, our architecture can approximate any regular conditional distribution of an X-valued random element X depending on a Y -valued random element Y with arbitrarily high probability. Consequentially, we show that once our approximation of this regular conditional distribution is learned, any conditional expectation of the form E[f(X; Y)|Y = y] for Caratheodory f with uniformly-Lipschitz first component and a uniformly-bounded second component, is approximable by standard Monte-Carlo sampling against the learned measure. We illustrate our theory in the context of stochastic processes.

GEORGE LAI, Wilfrid Laurier University
Portfolio Selection with Deep Learning

In this talk, we propose a two-stage framework to construct portfolios based on deep learning algorithms. Both NASDAQ100 and CSI300 are selected as representatives of developed markets and emerging markets, respectively. At stage 1, once an stock index is selected, for each component stock in the index, we use principal component analysis (PCA), auto-encoder(AE) and restricted Boltzmann machine (RBM) as data representation methods to reconstruct the stock prices, and select outstanding stocks to enter the portfolio according to the characteristics of data reconstruction. At stage 2, taking the selected stock index as the target, we train the artificial neural networks to construct portfolios and to test investment strategies by validation. Our initial results show that (1) there is no significant difference in the performance of different data representation methods; (2) the contribution of communal information to the optimal portfolio decreases with the number of selected stocks; (3) the

characteristics of different types of markets obtained by deep learning are different; (4) this approach achieves good results for different trading frequency data. This is a joint work with C. Zhang, R. Chen and G. Wang.

RICHARD LE, York University
Reducing Systemic Risk in a Multi-Layer Network Using Reinforcement Learning

Inter-bank lending relationships are an essential part of a working financial system. However, in the case of an extreme economic shock, lending relationships act as a mechanism for contagion. Here, we use reinforcement learning as a method to reduce the systemic risk of a financial network by modifying the lending relationships. We model a financial network as a multi-layered complex network where the nodes represent banks and the directed edges represent their lending relationships. We use deep deterministic policy gradient (DDPG), a model-free off-policy learning algorithm to reconfigure the simulated network. For the purpose of measuring the systemic risk of the network we consider its DebtRank, the potential economic loss of the network due to the economic shock experienced by individual banks. In the multi-layered network case we use an extension of the DebtRank by incorporating the level of distress from other layers of the network.

RUI LIANG, York University
Credit default risk prediction with machine learning

The credit default risk prediction can be considered as a binary classification problem. Hence machine learning models used for binary classification problems can be applied to credit default risk prediction as well. However, challenges remain. Dataset for credit default risk prediction is often highly imbalanced and incomplete. In this talk, we investigate the use of Random Forest and Neural Network for predicting credit default risk on a high dimensional dataset. The Random Forest model based on synthetic minority over-sampling technique (SMOTE) and a random search cross validation is employed on hyperparameters optimization. The neural network is ensembled with Adaptive boosting (Adaboost). The performances of these two approaches are interpreted in the talk. This is a joint work with Michael Chen and Hongmei Zhu

---

TERESA CRISTINA DE SA LIMA, York University
Categorizing Canadian Mutual Funds: An application of Machine Learning Classifiers

Mutual funds and exchange-traded funds (ETFs) are some of the most popular options for investors who are looking to diversify their financial portfolios. The popularity of these instruments has increased the competition intensely among fund providers, who manage a variety of funds and strive to outperform their peers. Given the growth of this market, classifying similar funds has become more important than ever and also a complex problem. The vendor-provider categorization developed by Morningstar is considered the most reliable in the Canadian and US markets. This methodology relies not only on data but also on a qualitative screen performed by a team of human experts, which is likely to add emotional biases to the process. The existing literature on fund categorization has been focused on unsupervised clustering. In this talk, however, we aim to learn a well-regarded categorization system using supervised machine learning algorithms, creating, therefore a fully data-driven method to classify Canadian Mutual funds.

KAI LIU, University of Prince Edward Island
Real-time Portfolio Optimization Under Multivariate Affine Generalized Hyperbolic Distributions

This paper focuses on the real-time portfolio optimization which captures the impacts of leptokurtic phenomenon and heterogeneous preferences in higher moments on asset allocation. To achieve this, we propose a utility maximization asset allocation framework under the multivariate affine generalized hyperbolic (MAGH) asset prices dynamics. With the investor's preference given by the exponential utility, we derive the closed-form optimal asset allocations for mixed multivariate affine Normal-Inverse-Gaussian-Normal model and mixed multivariate affine Variance-Gamma-Normal model, which cover Markowitz's mean-variance model as our special case. In addition, the optimization could be combined with dimensionality reduction methodology. The new algorithm reduces the estimation error and has real-time application. Extensive empirical studies are conducted to assess the

effectiveness of the proposed asset allocation models relative to other portfolio strategies based on Markowitz's mean-variance theory and the equally weighted 1/N rule.

ROGEMAR MAMON, University of Western Ontario
A multivariate hidden Markov model in the analysis of financial stability

We develop multivariate hidden Markov model (HMM) filters in recursive forms. These filters then provide optimal estimates for the state of the Markov chain as well as estimates of other related and auxiliary processes. The motivation behind our modelling is to capture regimes of financial stability through the filtering of relevant indices. In particular, these indices contain information regarding the systemic stress levels in the financial and business cycles. Through the recursive filters, parameters are updated instantaneously when new index information is available. Actual data is considered in our model validation. An early-warning signal system is also developed to produce alerts on the potential occurrence of some financial-crisis events.

BIHAI SU, Shanghai University of Finance and Economics
Solving a Seal's type partial intergro-differential equation with general jump by employing deep neural networks

In this paper, we study the problem of the Seal's type equation under the classical compound Poisson risk model. We propose a data-driven deep neural networks (DNNs) method to calculate the finite time survival probability. An alternative scheme under the exponential claim distribution is also discussed. This method is extended to the numerical solution of generalized partial integro-differential equations. Finally, the numerical approximation results under different claim distributions are given. Numerical results show that the proposed scheme can obtain accurate results under different claim distributions and has better computational efficiency.

YAODE SUI, WIlfrid Laurier University
Prediction of bitcoin trading signals

This study explores the prediction of bitcoin trading signals based on the historical bitcoin prices from 2017 to 2021. With some technical indicators, we use ML and DL methods to predict the trading signal of bitcoin. Compared with the traditional machine learning approaches such as the decision tree and support vector machine, the neural network could be a more suitable methodology than traditional statistics for predicting the cryptocurrency trading signal. In this research, we also use the NLP technique to find out how twitter influences the prices of bitcoin. The results show that the sentiment analysis of twitter data improves the accuracy of prediction of trading signal. This is a joint work with Y. Lai.

MINGFU WANG, York University
Risk-sensitive Policies for Portfolio Management Problems

Portfolio management is the decision-making process of allocating investment into different financial investment products. However, seeking the optimal trading strategy in a complex and dynamic stock market is challenging due to high uncertainty and massive noise in the financial market. Nowadays, Artificial Intelligence (AI) is well-developed in the financial markets. Trained an AI model as an automated agent can beat experienced human traders. Deep Deterministic Policy Gradient (DDPG), proposed by Google DeepMind, is an actor-critic Deep Reinforcement Learning (RL) algorithm. Although the DDPG achieves a remarkable performance in financial market trading, most of the literature surprisingly ignores the possible risk of rare occurrences of catastrophic events and the effects of the worst-case scenarios on trading decisions. It consists of two-level policies: the lower-level policy and the higher-level policy. The lower-level policy aims to maximize the expected future discounted rewards when the portfolio risk under a certain level of risk, and when the portfolio risk excesses a certain level of risk, the higher-level policy is to adjust the action from the lower-level policy to reduce the portfolio risk and provides a very conservative trading strategy. In addition, we are the first to apply the distributional DDPG to the portfolio management problems for maximizing the protection of investors. It refers to seek a risk-sensitive policy for portfolio management problems by modeling the distribution of future returns and maximizing the worst-case performance under this distribution. The learned risk-sensitive policy can map the same state to different actions depending on the propensity for risk.

SHENGNAN WANG, Zhejiang University of Finance and Economics

DAVID XU, Ryerson University

Delta-Gamma like Hedging with transaction cost under reinforcement learning technique

Option hedging is critical in financial risk management. The traditional methods to determine the hedging position require assumptions of a frictionless market and continuous hedging. In this paper, we remove these two assumptions and propose a hedging strategy based on the reinforcement learning technique. Our new strategy maximizes the present value of accounting and realized profits of hedging portfolio, and simultaneously minimizes the sensitivity of hedging position to the changes of the underlying asset. Finally, we test the performance of our method on option trading data of S&P 500, S&P 100, and DIJA from 2004 to 2020.