VOL. 3, NO. 6

OMI Research Newsletter – March 2024

The Oxford-Man Institute of Quantitative Finance (OMI) is a world-leading centre for interdisciplinary research in financial markets. The OMI addresses fundamental problems in quantitative finance with a strong focus on machine learning and data driven models.

*The **OMI Research Newsletter** is a curated summary that covers the most recent research, insights, and tools related to financial decision making.*

Highlights

Research focus. Key contributions in this month's edition include works from Harvard University, University of Toronto, and University of Alberta. Curry, Fan and Parkes (2024) frame optimal market making as an optimal transport problem and study optimal mechanisms for multiple goods under adverse selection. Jaimungal and Shi (2024) study the price of information for an agent who maximizes expected utility of terminal wealth. Finally, Degris et al. (2024) study step size in continual learning, i.e., the choice of knowledge to keep and knowledge to forget.

Machine Learning & Statistics

Large (and Deep) Factor Models

B. KELLY, B. KUZNETSOV, S. MALAMUD, T. A. XU
EPFL, YALE SCHOOL OF MANAGEMENT

We open up the black box behind Deep Learning for portfolio optimization and prove that a sufficiently wide and arbitrarily deep neural network (DNN) trained to maximize the Sharpe ratio of the Stochastic Discount Factor (SDF) is equivalent to a large factor model (LFM): A linear factor pricing model that uses many nonlinear characteristics. The nature of these characteristics depends on the architecture of the DNN in an explicit, tractable fashion. This makes it possible to derive end-to-end trained DNN-based SDFs in closed form for the first time. We evaluate LFMs empirically and show how various architectural choices impact SDF performance. We document the virtue of depth complexity: With enough data, the out-of-sample performance of DNN-SDF is increasing in the NN depth, saturating at huge depths of around 100 hidden layers.

Neural option pricing for rough Bergomi model

C. TENG, G. LI
THE UNIVERSITY OF HONG KONG

The rough Bergomi (rBergomi) model can accurately describe the historical and implied volatilities, and has gained much attention in the past few years. However, there are many hidden unknown parameters or even functions in the model. In this work, we investigate the potential of learning the forward variance curve in the rBergomi model using a neural SDE. To construct an efficient solver for the neural SDE, we propose a novel numerical scheme for simulating the volatility process using the modified summation of exponentials. Using the Wasserstein 1-distance to define the loss function, we show that the learned forward variance curve is capable of calibrating the price process of the underlying asset and the price of the European-style options simultaneously. Several numerical tests are provided to demonstrate its performance.

Reinforcement Learning for Optimal Execution when Liquidity is Time-Varying

A. MACRÌ, F. LILLO
SCUOLA NORMALE SUPERIORE, UNIVERSITY OF BOLOGNA

Optimal execution is an important problem faced by any trader. Most solutions are based on the assumption of constant market impact, while liquidity is known to be dynamic. Moreover, models with time-varying liquidity typically assume that it is observable, despite the fact that, in reality, it is latent and hard to measure

in real time. In this paper we show that the use of Double Deep Q-learning, a form of Reinforcement Learning based on neural networks, is able to learn optimal trading policies when liquidity is time-varying. Specifically, we consider an Almgren-Chriss framework with temporary and permanent impact parameters following several deterministic and stochastic dynamics. Using extensive numerical experiments, we show that the trained algorithm learns the optimal policy when the analytical solution is available, and overcomes benchmarks and approximated solutions when the solution is not available.

Step-size Optimization for Continual Learning

T. Degris, K. Javed, A. Sharifnassab, Y. Liu, R. Sutton

DeepMind, University of Alberta

In continual learning, a learner has to keep learning from the data over its whole life time. A key issue is to decide what knowledge to keep and what knowledge to let go. In a neural network, this can be implemented by using a step-size vector to scale how much gradient samples change network weights. Common algorithms, like RMSProp and Adam, use heuristics, specifically normalization, to adapt this step-size vector. In this paper, we show that those heuristics ignore the effect of their adaptation on the overall objective function, for example by moving the step-size vector away from better step-size vectors. On the other hand, stochastic meta-gradient descent algorithms, like IDBD (Sutton, 1992), explicitly optimize the step-size vector with respect to the overall objective function. On simple problems, we show that IDBD is able to consistently improve step-size vectors, where RMSProp and Adam do not. We explain the differences between the two approaches and their respective limitations. We conclude by suggesting that combining both approaches could be a promising future direction to improve the performance of neural networks in continual learning.

End-to-End Policy Learning of a Statistical Arbitrage Autoencoder Architecture

Fabian Krause, Jan-Peter Calliess

University of Oxford

In Statistical Arbitrage (StatArb), classical mean reversion trading strategies typically hinge on asset-pricing or PCA based models to identify the mean of a synthetic asset. Once such a (linear) model is identified, a separate mean reversion strategy is then devised to generate a trading signal. With a view of generalising such an approach and turning it truly data-driven, we study the utility of Autoencoder architectures in StatArb. As a first approach, we employ a standard Autoencoder trained on US stock returns to derive trading strategies based on the Ornstein-Uhlenbeck (OU) process. To further enhance this model, we take a policy-learning approach and embed the Autoencoder network into a neural network representation of a space of portfolio trading policies. This integration outputs portfolio allocations directly and is end-to-end trainable by backpropagation of the risk-adjusted returns of the neural policy. Our findings demonstrate that this innovative end-to-end policy learning approach not only simplifies the strategy development process, but also yields superior gross returns over its competitors illustrating the potential of end-to-end training over classical two-stage approaches.

Repeat After Me: Transformers are Better than State Space Models at Copying

S. Jelassi, D. Brandfonbrener, S. M. Kakade, E. Malach

Harvard University

Transformers are the dominant architecture for sequence modeling, but there is growing interest in models that use a fixed-size latent state that does not depend on the sequence length, which we refer to as "generalized state space models" (GSSMs). In this paper we show that while GSSMs are promising in terms of inference-time efficiency, they are limited compared to transformer models on tasks that require copying from the input context. We start with a theoretical analysis of the simple task of string copying and prove that a two layer transformer can copy strings of exponential length while GSSMs are fundamentally limited by their fixed-size latent state. Empirically, we find that transformers outperform GSSMs in terms of efficiency and generalization on synthetic tasks that require copying the context. Finally, we evaluate pretrained large language models and find that transformer models dramatically outperform state space models at copying and retrieving information from context. Taken together, these results suggest a fundamental gap between transformers and GSSMs on tasks of practical interest.

AutoFT: Robust Fine-Tuning by Optimizing Hyperparameters on OOD Data

C. Choi, Y. Lee, A. Chen, A. Zhou, A. Raghunathan, C. Finn

Stanford University, Carnegie Mellon University

Foundation models encode rich representations that can be adapted to a desired task by fine-tuning on task-specific data. However, fine-tuning a model on one particular data distribution often compromises the model's original performance on other distributions. Current methods for robust fine-tuning utilize hand-crafted regularization techniques to constrain the fine-tuning process towards the base foundation model. Yet,

it is hard to precisely specify what characteristics of the foundation model to retain during fine-tuning, as this depends on how the pre-training, fine-tuning, and evaluation data distributions relate to each other. We propose AutoFT, a data-driven approach for guiding foundation model fine-tuning. AutoFT optimizes fine-tuning hyperparameters to maximize performance on a small out-of-distribution (OOD) validation set. To guide fine-tuning in a granular way, AutoFT searches a highly expressive hyperparameter space that includes weight coefficients for many different losses, in addition to learning rate and weight decay values. We evaluate AutoFT on nine natural distribution shifts which include domain shifts and subpopulation shifts. Our experiments show that AutoFT significantly improves generalization to new OOD data, outperforming existing robust fine-tuning methods. Notably, AutoFT achieves new state-of-the-art performance on the WILDS-iWildCam and WILDS-FMoW benchmarks, outperforming the previous best methods by 6.0% and 1.5%, respectively.

PDE-refiner: Achieving Accurate Long Rollouts with Neural PDE Solvers

P. Lippe, B. Veeling, P. Perdikaris, R. Turner, J. Brandstetter

Microsoft Research AI4Science

Time-dependent partial differential equations (PDEs) are ubiquitous in science and engineering. Recently, mostly due to the high computational cost of traditional solution techniques, deep neural network based surrogates have gained increased interest. The practical utility of such neural PDE solvers relies on their ability to provide accurate, stable predictions over long time horizons, which is a notoriously hard problem. In this work, we present a large-scale analysis of common temporal rollout strategies, identifying the neglect of non-dominant spatial frequency information, often associated with high frequencies in PDE solutions, as the primary pitfall limiting stable, accurate rollout performance. Based on these insights, we draw inspiration from recent advances in diffusion models to introduce PDE-Refiner; a novel model class that enables more accurate modeling of all frequency components via a multistep refinement process. We validate PDE-Refiner on challenging benchmarks of complex fluid dynamics, demonstrating stable and accurate rollouts that consistently outperform state-of-the-art models, including neural, numerical, and hybrid neural-numerical architectures. We further demonstrate that PDE-Refiner greatly enhances data efficiency, since the denoising objective implicitly induces a novel form of spectral data augmentation. Finally, PDE-Refiner's connection to diffusion models enables an accurate and efficient assessment of the model's predictive uncertainty, allowing us to estimate when the surrogate becomes inaccurate.

Finance and Economics

Optimal Automated Market Makers: Differentiable Economics and Strong Duality

M. J. Curry, Z. Fan, D. C. Parkes

Harvard University

The role of a market maker is to simultaneously offer to buy and sell quantities of goods, often a financial asset such as a share, at specified prices. An automated market maker (AMM) is a mechanism that offers to trade according to some predetermined schedule; the best choice of this schedule depends on the market maker's goals. The literature on the design of AMMs has mainly focused on prediction markets with the goal of information elicitation. More recent work motivated by DeFi has focused instead on the goal of profit maximization, but considering only a single type of good (traded with a numeraire), including under adverse selection (Milionis et al. 2022). Optimal market making in the presence of multiple goods, including the possibility of complex bundling behavior, is not well understood. In this paper, we show that finding an optimal market maker is dual to an optimal transport problem, with specific geometric constraints on the transport plan in the dual. We show that optimal mechanisms for multiple goods and under adverse selection can take advantage of bundling, both improved prices for bundled purchases and sales as well as sometimes accepting payment "in kind." We present conjectures of optimal mechanisms in additional settings which show further complex behavior. From a methodological perspective, we make essential use of the tools of differentiable economics to generate conjectures of optimal mechanisms, and give a proof-of-concept for the use of such tools in guiding theoretical investigations.

Efficiently Breaking the Folk Theorem by Reliably Communicating Long Term Commitments

D. K. Levine

Royal Holloway

The introduction of artificially intelligent algorithms in pricing decisions by firms has triggered a literature in industrial organization asking if the use of these algorithms will lead to collusive outcomes. In a simple repeated game environment it is shown that if algorithms can be reliably communicated or inferred the folk theorem breaks and the long-run outcome must be collusive.

Perfect Bayesian Persuasion

E. Lipnowski, D. Ravid, D. Shishkin
Columbia University, University of Chicago, UC San Diego

A sender commits to an experiment to persuade a receiver. Accounting for the sender's experiment-choice incentives, and not presupposing a receiver tie-breaking rule when indifferent, we characterize when the sender's equilibrium payoff is unique and so coincides with her "Bayesian persuasion" value. A sufficient condition in finite models is that every action which is receiver-optimal at some belief is uniquely optimal at some other belief – a generic property. We similarly show the equilibrium sender payoff is typically unique in ordered models. In an extension, we show uniqueness generates robustness to imperfect sender commitment.

Persuading a Learning Agent

T. Lin, Y. Chen
Harvard University

We study a repeated Bayesian persuasion problem (and more generally, any generalized principal-agent problem with complete information) where the principal does not have commitment power and the agent uses algorithms to learn to respond to the principal's signals. We reduce this problem to a one-shot generalized principal-agent problem with an approximately-best-responding agent. This reduction allows us to show that: if the agent uses contextual no-regret learning algorithms, then the principal can guarantee a utility that is arbitrarily close to the principal's optimal utility in the classic non-learning model with commitment; if the agent uses contextual no-swap-regret learning algorithms, then the principal cannot obtain any utility significantly more than the optimal utility in the non-learning model with commitment. The difference between the principal's obtainable utility in the learning model and the non-learning model is bounded by the agent's regret (swap-regret). If the agent uses mean-based learning algorithms (which can be no-regret but not no-swap-regret), then the principal can do significantly better than the non-learning model. These conclusions hold not only for Bayesian persuasion, but also for any generalized principal-agent problem with complete information, including Stackelberg games and contract design.

Mathematical Finance

Signature Volatility Models: Pricing and Hedging With Fourier

E. Abi Jaber, L-A. Gérard
École Polytechnique, Université Paris 1 Pantheon-Sorbonne

We consider a stochastic volatility model where the dynamics of the volatility are given by a possibly infinite linear combination of the elements of the time extended signature of a Brownian motion. First, we show that the model is remarkably universal, as it includes, but is not limited to, the celebrated Stein-Stein, Bergomi, and Heston models, together with some path-dependent variants. Second, we derive the joint characteristic functional of the log-price and integrated variance provided that some infinite-dimensional extended tensor algebra valued Riccati equation admits a solution. This allows us to price and (quadratically) hedge certain European and path-dependent options using Fourier inversion techniques. We highlight the efficiency and accuracy of these Fourier techniques in a comprehensive numerical study.

The Price of Information

S. Jaimungal, X. Shi
University of Toronto

When an investor is faced with the option to purchase additional information regarding an asset price, how much should she pay? To address this question, we solve for the indifference price of information in a setting where a trader maximizes her expected utility of terminal wealth over a finite time horizon. If she does not purchase the information, then she solves a partial information stochastic control problem, while, if she does purchase the information, then she pays a cost and receives partial information about the asset's trajectory. We further demonstrate that when the investor can purchase the information at any stopping time prior to the end of the trading horizon, she chooses to do so at a deterministic time.

Acknowledgement

The OMI Research Newsletter is curated by Álvaro Cartea, Director of the Oxford-Man Institute, and Fayçal Drissi, postdoctoral researcher at the institute. Paper suggestions are from fellow members at the institute: Álvaro Arroyo, Patrick Chang, Gerardo Durán-Martín, Marcello Monga, and from Harrison Waldon.

References

Abi Jaber, Eduardo, and Louis-Amand Gérard. 2024. "Signature volatility models: pricing and hedging with Fourier." *Available at SSRN*.

Choi, Caroline, Yoonho Lee, Annie Chen, Allan Zhou, Aditi Raghunathan, and Chelsea Finn. 2024. "Aut-

oFT: Robust Fine-Tuning by Optimizing Hyperparameters on OOD Data." *arXiv preprint arXiv:2401.10220*.

Curry, Michael J, Zhou Fan, and David C Parkes. 2024. "Optimal Automated Market Makers: Differentiable Economics and Strong Duality." *arXiv preprint arXiv:2402.09129*.

Degris, Thomas, Khurram Javed, Arsalan Sharifnassab, Yuxin Liu, and Richard Sutton. 2024. "Step-size Optimization for Continual Learning." *arXiv preprint arXiv:2401.17401*.

Jaimungal, Sebastian, and Xiaofei Shi. 2024. "The Price of Information." *arXiv preprint arXiv:2402.11864*.

Jelassi, Samy, David Brandfonbrener, Sham M Kakade, and Eran Malach. 2024. "Repeat after me: Transformers are better than state space models at copying." *arXiv preprint arXiv:2402.01032*.

Kelly, Bryan, Boris Kuznetsov, Semyon Malamud, and Teng Andrea Xu. 2024. "Large (and Deep) Factor Models." *arXiv preprint arXiv:2402.06635*.

Krause, Fabian, and Jan-Peter Calliess. 2024. "End-to-End Policy Learning of a Statistical Arbitrage Autoencoder Architecture." *arXiv preprint arXiv:2402.08233*.

Levine, David K. 2023. "Efficiently Breaking the Folk Theorem by Reliably Communicating Long Term Commitments." David K. Levine.

Lin, Tao, and Yiling Chen. 2024. "Persuading a Learning Agent." *arXiv preprint arXiv:2402.09721*.

Lipnowski, Elliot, Doron Ravid, and Denis Shishkin. 2024. "Perfect Bayesian Persuasion." *arXiv preprint arXiv:2402.06765*.

Lippe, Phillip, Bas Veeling, Paris Perdikaris, Richard Turner, and Johannes Brandstetter. 2024. "Pderefiner: Achieving accurate long rollouts with neural pde solvers." *Advances in Neural Information Processing Systems*, 36.

Macrì, Andrea, and Fabrizio Lillo. 2024. "Reinforcement Learning for Optimal Execution when Liquidity is Time-Varying." *arXiv preprint arXiv:2402.12049*.

Teng, Changqing, and Guanglian Li. 2024. "Neural option pricing for rough Bergomi model." *arXiv preprint arXiv:2402.02714*.