VOL. 4, NO. 3

OMI Research Newsletter – December 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.*

Announcement

The London–Oxford–Warwick Mathematical Finance Workshop will be held on 9–10 January 2025 at the Oxford-Man Institute. Please find the schedule here.

The Oxford–Harvard Conference on Decentralised Finance will be held on 20–21 January 2025 at the Oxford-Man Institute. Please find the schedule and register here.

Highlights

Key contributions in this month's edition include a novel framework for Bayesian online learning in non-stationary environments (Duran-Martin et al., 2024), a time-uniform central limit theorem facilitating inference at arbitrary stopping times (Waudby-Smith et al., 2024), a review of machine learning applications in portfolio decision-making (Guidolin, Panzeri and Pedio, 2024), a new test for identifying abrupt asset price changes using intraday candlestick charts (Bollerslev et al., 2024), an empirical analysis of intraday variations in the factor structure of asset returns (Patton and Simsek, 2024), and a dynamic model for bonding curves in decentralised finance aimed at reducing arbitrage losses for liquidity providers (Cartea et al., 2024).

Machine Learning & Statistics

BONE: A Unifying Framework for Bayesian Online Learning in Non-Stationary Environments

G. DURAN-MARTIN, L. SÁNCHEZ-BETANCOURT, A.Y. SHESTOPALOFF, K. MURPHY

QUEEN MARY UNIVERSITY OF LONDON, OXFORD-MAN INSTITUTE OF QUANTITATIVE FINANCE, UNIVERSITY OF OXFORD, GOOGLE DEEPMIND

We propose a unifying framework for methods that perform Bayesian online learning in non-stationary environments. We call the framework BONE, which stands for (B)ayesian (O)nline learning in (N)on-stationary (E)nvironments. BONE provides a common structure to tackle a variety of problems, including online continual learning, prequential forecasting, and contextual bandits. The framework requires specifying three modelling choices: (i) a model for measurements (e.g., a neural network), (ii) an auxiliary process to model non-stationarity (e.g., the time since the last changepoint), and (iii) a conditional prior over model parameters (e.g., a multivariate Gaussian). The framework also requires two algorithmic choices, which we use to carry out approximate inference under this framework: (i) an algorithm to estimate beliefs (posterior distribution) about the model parameters given the auxiliary variable, and (ii) an algorithm to estimate beliefs about the auxiliary variable. We show how this modularity allows us to write many different existing methods as instances of BONE; we also use this framework to propose a new method. We then experimentally compare existing methods with our proposed new method on several datasets; we provide insights into the situations that make one method more suitable than another for a given task.

Time-Uniform Central Limit Theory and Asymptotic Confidence Sequences

I. WAUDBY-SMITH, D. ARBOUR, R. SINHA, E. H. KENNEDY, A. RAMDAS

CARNEGIE MELLON UNIVERSITY, ADOBE RESEARCH

Confidence intervals based on the central limit theorem

(CLT) are a cornerstone of classical statistics. Despite being only asymptotically valid, they are ubiquitous because they permit statistical inference under weak assumptions and can often be applied to problems even when nonasymptotic inference is impossible. This paper introduces time-uniform analogues of such asymptotic confidence intervals, adding to the literature on confidence sequences (CS) – sequences of confidence intervals that are uniformly valid over time – which provide valid inference at arbitrary stopping times and incur no penalties for "peeking" at the data, unlike classical confidence intervals which require the sample size to be fixed in advance. Existing CSs in the literature are nonasymptotic, enjoying finite-sample guarantees but not the aforementioned broad applicability of asymptotic confidence intervals. This work provides a definition for "asymptotic CSs" and a general recipe for deriving them. Asymptotic CSs forgo nonasymptotic validity for CLT-like versatility and (asymptotic) time-uniform guarantees. While the CLT approximates the distribution of a sample average by that of a Gaussian for a fixed sample size, we use strong invariance principles (stemming from the seminal 1960s work of Strassen) to uniformly approximate the entire sample average process by an implicit Gaussian process. As an illustration, we derive asymptotic CSs for the average treatment effect in observational studies (for which nonasymptotic bounds are essentially impossible to derive even in the fixed-time regime) as well as randomized experiments, enabling causal inference in sequential environments.

Unveiling the Hidden Structure of Self-Attention via Kernel Principal Component Analysis

R. S. Y. Teo, T. M. Nguyen

National University of Singapore

The remarkable success of transformers in sequence modeling tasks, spanning various applications in natural language processing and computer vision, is attributed to the critical role of self-attention. Similar to the development of most deep learning models, the construction of these attention mechanisms relies on heuristics and experience. In our work, we derive self-attention from kernel principal component analysis (kernel PCA) and show that self-attention projects its query vectors onto the principal component axes of its key matrix in a feature space. We then formulate the exact formula for the value matrix in self-attention, theoretically and empirically demonstrating that this value matrix captures the eigenvectors of the Gram matrix of the key vectors in self-attention. Leveraging our kernel PCA framework, we propose Attention with Robust Principal Components (RPC-Attention), a novel class of robust attention that is resilient to data contamination. We empirically demonstrate the advantages of RPC-Attention over softmax attention on the ImageNet-1K object classification, WikiText-103 language modeling, and ADE20K image segmentation task.

What Type of Inference is Planning?

M. Lázaro-Gredilla, L. Y. Ku, K. P. Murphy, D. George

Google Deepmind

Multiple types of inference are available for probabilistic graphical models, e.g., marginal, maximum-a-posteriori, and even marginal maximum-a-posteriori. Which one do researchers mean when they talk about “planning as inference”? There is no consistency in the literature, different types are used, and their ability to do planning is further entangled with specific approximations or additional constraints. In this work we use the variational framework to show that, just like all commonly used types of inference correspond to different weightings of the entropy terms in the variational problem, planning corresponds exactly to a different set of weights. This means that all the tricks of variational inference are readily applicable to planning. We develop an analogue of loopy belief propagation that allows us to perform approximate planning in factored-state Markov decisions processes without incurring intractability due to the exponentially large state space. The variational perspective shows that the previous types of inference for planning are only adequate in environments with low stochasticity, and allows us to characterize each type by its own merits, disentangling the type of inference from the additional approximations that its practical use requires. We validate these results empirically on synthetic MDPs and tasks posed in the International Planning Competition.

SigDiffusions: Score-Based Diffusion Models for Long Time Series via Log-Signature Embeddings

B. Barancikova, Z. Huang, C. Salvi

Imperial College London, University of Cambridge

Score-based diffusion models have recently emerged as state-of-the-art generative models for a variety of data modalities. Nonetheless, it remains unclear how to adapt these models to generate long multivariate time series. Viewing a time series as the discretization of an underlying continuous process, we introduce SigDiffusion, a novel diffusion model operating on log-signature embeddings of the data. The forward and backward processes gradually perturb and denoise log-signatures preserving their algebraic structure. To recover a signal from its log-signature, we provide new closed-form inversion formulae expressing the coefficients obtained by expanding the signal in a given basis (e.g. Fourier or orthogonal polynomials) as explicit polynomial functions of the log-signature. Finally, we show that combining

SigDiffusion with these inversion formulae results in highly realistic time series generation, competitive with the current state-of-the-art on various datasets of synthetic and real-world examples.

Hymba: A Hybrid-head Architecture for Small Language Models

X. Dong, Y. Fu, S. Diao, W. Byeon, Z. Chen, A. S. Mahabaleshwarkar, S.-Y. Liu, M. Van Keirsbilck, M.-H. Chen, Y. Suhara, Y. Lin, J. Kautz, P. Molchanov

NVIDIA Research, Georgia Institute of Technolog, HKUST

We propose Hymba, a family of small language models featuring a hybrid-head parallel architecture that integrates transformer attention mechanisms with state space models (SSMs) for enhanced efficiency. Attention heads provide high-resolution recall, while SSM heads enable efficient context summarization. Additionally, we introduce learnable meta tokens that are prepended to prompts, storing critical information and alleviating the "forced-to-attend" burden associated with attention mechanisms. This model is further optimized by incorporating cross-layer key-value (KV) sharing and partial sliding window attention, resulting in a compact cache size. During development, we conducted a controlled study comparing various architectures under identical settings and observed significant advantages of our proposed architecture. Notably, Hymba achieves state-of-the-art results for small LMs: Our Hymba-1.5B-Base model surpasses all sub-2B public models in performance and even outperforms Llama-3.2-3B with 1.32% higher average accuracy, an 11.67x cache size reduction, and 3.49x throughput.

Finance and Economics

Machine Learning in Portfolio Decisions

M. Guidolin, G. Panzeri, M. Pedio

Bocconi University, University of Bristol

Machine learning is significantly shaping the advancement of various fields, and among them, notably, finance, where its range of applications and efficiency impacts are seemingly boundless. Contemporary techniques, particularly in reinforcement learning, have prompted both practitioners and academics to contemplate the potential of an artificial intelligence revolution in portfolio management. In this paper, we provide an overview of the primary methods in machine learning currently utilized in portfolio decision-making. We delve into discussions surrounding the existing limitations of machine learning algorithms and explore prevailing hypotheses regarding their future expansions. Specifically, we categorize and analyze the applications of machine learning in systematic trading strategies, portfolio weight optimization, smart beta and passive investment strategies, textual analysis, and trade execution, each separately surveyed for a comprehensive understanding.

Illuminating Important Economic News by Candlesticks: Optimal Testing Meets Technical Analysis

T. Bollerslev, J. Li, Y. Li, Q. Zhang

Duke University, Singapore Management University, University of Manchester, University of International Business and Economics

Building on time-honored ideas from technical analysis, we propose a simple new test, termed the Marubozu test, for detecting abrupt asset price changes over short time intervals based on the information embedded in readily available high-frequency candlesticks. Using an infill asymptotic framework, combined with a novel coupling-based method, we show that the test is correctly sized and optimal, or nearly so, against a variety of empirically plausible alternatives. Underscoring the practical relevance of the test, we demonstrate its ability to illuminate significant economic shocks triggered by specific statements made by the Fed chair in post-FOMC announcement press conferences.

Re(Visiting) Large Language Models in Finance

E. Rahimikia, F. Drinkall

University of Manchester, Oxford-Man Institute of Quantitative Finance, University of Oxford

This study introduces a novel suite of historical large language models (LLMs) pre-trained specifically for accounting and finance, utilising a diverse set of major textual resources. The models are unique in that they are year-specific, spanning from 2007 to 2023, effectively eliminating look-ahead bias, a limitation present in other LLMs. Empirical analysis reveals that, in trading, these specialised models outperform much larger models, including the state-of-the-art LLaMA 1, 2, and 3, which are approximately 50 times their size. The findings are further validated through a range of robustness checks, confirming the superior performance of these LLMs.

Intraday Variation in Systematic Risks and Information Flows

A. J. Patton, Y. Simsek

Duke University

This paper analyzes variation in the factor structure of asset returns within a trade day by combining non-parametric kernel methods with principal component

analysis. We estimate the model on a collection of over 400 high frequency U.S. equity returns over the period 1996-2020 and show that the proposed model has superior explanatory power relative to a collection of well-known observable factor models and standard PCA. We present a stylized model of asset prices and information flows and show that the factor structure of asset returns varies with the arrival of news. Using data on individual firm earnings announcements, FOMC announcements, and other macroeconomic announcements, we provide evidence consistent with our stylized model, that the superior performance of the proposed model is due to time variation in the factor structure of asset returns around times of information flows.

Do Limits to Arbitrage Explain Portfolio Gains from Asset Mispricing?

N. Lassance, A. Martin-Utrera
UCLouvain, Iowa State University

We study the mean-variance efficiency gains from exploiting asset mispricing encapsulated in a large set of firm characteristics through the lenses of an optimal arbitrage portfolio. We document significant efficiency gains from the arbitrage portfolio and show that limits to arbitrage, such as estimation risk, transaction costs, shorting constraints, and post-publication effects, can only account for these gains during low-sentiment regimes. We provide evidence that the low-beta anomaly can partly explain the portfolio gains from exploitable asset mispricing and attribute these gains to compensation for business-cycle risk.

A Machine Learning Algorithm for Finite-Horizon Stochastic Control Problems in Economics

X. Peng, S. Kou, L. Zhang
Peking University, Boston Universit

We propose a machine learning algorithm for solving finite-horizon stochastic control problems based on a deep neural network representation of the optimal policy functions. The algorithm has three features: (1) It can solve high-dimensional (e.g., over 100 dimensions) and finite-horizon time-inhomogeneous stochastic control problems. (2) It has a monotonicity of performance improvement in each iteration, leading to good convergence properties. (3) It does not rely on the Bellman equation. To demonstrate the efficiency of the algorithm, it is applied to solve various finite-horizon time-inhomogeneous problems including recursive utility optimization under a stochastic volatility model, a multi-sector stochastic growth, and optimal control under a dynamic stochastic integration of climate and economy model with eight-dimensional state vectors and 600 time periods.

Mathematical Finance

Strategic Bonding Curves in Automated Market Makers

Á. Cartea, F. Drissi, L. Sánchez-Betancourt, D. Siska, L. Szpruch
Oxford-Man Institute of Quantitative Finance, University of Oxford, University of Edinburgh

The bonding curves of decentralised exchanges (DEXs) define, with mathematical formulae, the relationship between liquidity supply, price impact, and execution prices. Most DEXs use the bonding curves of constant function markets (CFMs) to clear the supply and demand of liquidity. At present, liquidity providers (LPs) operate at a loss in CFMs, on average. We generalise CFMs and introduce decentralised liquidity pools (DLPs) which provide LPs with the tools to design dynamic bonding curves according to strategic preferences. In DLPs, impact functions encode how orders affect prices and quote functions encode how the price of liquidity is determined. To illustrate the strategic flexibility of LPs in DLPs, we develop models for dynamic bonding curves when prices form in a competing trading venue, within the DLP, or across multiple venues. In fragmented markets, the DLP estimates the fundamental price from the trading flow to adjust the bonding curve and reduce arbitrage losses. Our models may be used as hooks in Uniswap v4 when the DLP's impact functions are those implied by the constant product function.

Mirror Descent Algorithms for Risk Budgeting Portfolios

M. A. Iglesias, A. R. Cetingoz, N. Frikha
Université Paris 1 Panthéon-Sorbonne

This paper introduces and examines numerical approximation schemes for computing risk budgeting portfolios associated to positive homogeneous and sub-additive risk measures. We employ Mirror Descent algorithms to determine the optimal risk budgeting weights in both deterministic and stochastic settings, establishing convergence along with an explicit non-asymptotic quantitative rate for the averaged algorithm. A comprehensive numerical analysis follows, illustrating our theoretical findings across various risk measures – including standard deviation, Expected Shortfall, deviation measures, and Variantiles – and comparing the performance with that of the standard stochastic gradient descent method recently proposed in the literature.

Efficient Trading with Price Impact

X. Brokmann, L. Gonon, G. He, D. Itkin, J. Muhle-Karbe

QUBE RESEARCH AND TECHNOLOGIES, IMPERIAL COLLEGE LONDON, UNIVERSITY OF ST. GALLEN, LONDON SCHOOL OF ECONOMICS

This paper studies how to compute trading strategies that trade off expected returns, risk, and the trading costs generated by price impact in an efficient manner. For a general model where impact is nonlinear and decays only gradually, we show that the performance of linear feedback policies can be computed in closed form. These are optimal for linear impact and linear decay dynamics only. However, when the linear impact parameters are optimized for the nonlinear impact model at hand, then the performance of the resulting policy is competitive with state-of-the-art neural network methods unless the market is so illiquid that even the performance of optimized strategies is no longer attractive net of trading costs.

Neural Operators Can Play Dynamic Stackelberg Games

G. ALVAREZ, I. EKREN, A. KRATSIOS, X. YANG

UNIVERSITY OF MICHIGAN, MCMASTER UNIVERSITY, VECTOR INSTITUTE

Dynamic Stackelberg games are a broad class of two-player games in which the leader acts first, and the follower chooses a response strategy to the leader's strategy. Unfortunately, only stylized Stackelberg games are explicitly solvable since the follower's best-response operator (as a function of the control of the leader) is typically analytically intractable. This paper addresses this issue by showing that the *follower's best-response operator* can be approximately implemented by an *attention-based neural operator*, uniformly on compact subsets of adapted open-loop controls for the leader. We further show that the value of the Stackelberg game where the follower uses the approximate best-response operator approximates the value of the original Stackelberg game. Our main result is obtained using our universal approximation theorem for attention-based neural operators between spaces of square-integrable adapted stochastic processes, as well as stability results for a general class of Stackelberg games.

Acknowledgement

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

References

Alvarez, Guillermo, Ibrahim Ekren, Anastasis Kratsios, and Xuwei Yang. 2024. "Neural Operators Can Play Dynamic Stackelberg Games." *arXiv preprint arXiv:2411.09644*.

Barancikova, Barbora, Zhuoyue Huang, and Cristopher Salvi. 2024. "SigDiffusions: Score-Based Diffusion Models for Long Time Series via Log-Signature Embeddings." *arXiv preprint arXiv:2406.10354*.

Bollerslev, Tim, Jia Li, Yifan Li, and Qiushi Zhang. 2024. "Illuminating Important Economic News by Candlesticks: Optimal Testing Meets Technical Analysis." *Working Paper*.

Brokmann, Xavier, Lukas Gonon, Guangyi He, David Itkin, and Johannes Muhle-Karbe. 2024. "Efficient Trading with Price Impact." *Working Paper*.

Cartea, Álvaro, Fayçal Drissi, Leandro Sánchez-Betancourt, David Siska, and Lukasz Szpruch. 2024. "Strategic Bonding Curves in Automated Market Makers." *Working Paper*.

Dong, Xin, Yonggan Fu, Shizhe Diao, Wonmin Byeon, Zijia Chen, Ameya Sunil Mahabaleshwarkar, Shih-Yang Liu, Matthijs Van Keirsbilck, Min-Hung Chen, Yoshi Suhara, et al. 2024. "Hymba: A Hybrid-head Architecture for Small Language Models." *arXiv preprint arXiv:2411.13676*.

Duran-Martin, Gerardo, Leandro Sánchez-Betancourt, Alexander Y Shestopaloff, and Kevin Murphy. 2024. "BONE: A Unifying Framework for Bayesian Online Learning in Non-Stationary Environments." *arXiv preprint arXiv:2411.10153*.

Guidolin, Massimo, Giulia Panzeri, and Manuela Pedio. 2024. "Machine Learning in Portfolio Decisions." *Working Paper*.

Iglesias, Martin Arnaiz, Adil Rengim Cetingoz, and Noufel Frikha. 2024. "Mirror Descent Algorithms for Risk Budgeting Portfolios." *arXiv preprint arXiv:2411.12323*.

Lassance, Nathan, and Alberto Martin-Utrera. 2024. "Do Limits to Arbitrage Explain Portfolio Gains from Asset Mispricing?" *Working Paper*.

Lázaro-Gredilla, Miguel, Li Yang Ku, Kevin P Murphy, and Dileep George. 2024. "What type of inference is planning?" *arXiv preprint arXiv:2406.17863*.

Patton, Adrew J., and Yasin Simsek. 2024. "Intraday Variation in Systematic Risks and Information Flows." *Working Paper*.

Peng, Xianhua, Steven Kou, and Lekang Zhang. 2024. "A Machine Learning Algorithm for Finite-Horizon Stochastic Control Problems in Economics." *arXiv preprint arXiv:2411.08668*.

Rahimikia, Eghbal, and Felix Drinkall. 2024. "Re(Visiting) Large Language Models in Finance." *Working Paper*.

Teo, Rachel SY, and Tan M Nguyen. 2024. "Unveiling the Hidden Structure of Self-Attention via Kernel Principal Component Analysis." *arXiv preprint arXiv:2406.13762*.

Waudby-Smith, Ian, David Arbour, Ritwik Sinha, Edward H Kennedy, and Aaditya Ramdas. 2024. "Time-Uniform Central Limit Theory and Asymptotic Confidence Sequences." *arXiv preprint arXiv:2103.06476*.