Executive Summary

Machine learning has gained widespread attention in finance, in particular in asset pricing, where Gu et al. (2020a) showed the superior predictive power of algorithms such as neural networks and tree ensembles. We conduct a comparative experiment of six different algorithms across three regions: United States, Europe and Japan, that we also combine in a unique universal balanced data set. We introduce state-of-the-art developments in machine learning in terms of methods (XGBoost and VASA), model explainability (SHAP values) and hyperparameters optimization (Bayesian search), but also in covariance matrix estimation (Analytical Shrinkage) and portfolio construction (Efficient Sorting).

We confirm the supremacy of complex non-linear algorithms over linear ones both from a predictive and financial performance standpoints. Neural networks are outperformed by tree ensembles in region-specific universes, while they are more competitive in the universal one. We show how restricting only to stocks with full history of returns, in the balanced data set, can decrease overall performances, because of the absence of stocks with smaller capitalization, and stock-specific metrics can highlight extreme errors made by our models. However, we can achieve similar financial results with Long-Only strategies, while Efficient Sorting can enhance the performances of Long-Short portfolios. We perform statistical tests to assess the significance of differences in predictive power and strategies' risk-adjusted returns, showing how even large gaps can often be seen as random fluctuations. Finally, we make an extensive use of SHAP values for non-linear algorithms, to show how they can provide invaluable insights on their behaviour, with a solid theoretical framework derived from game theory, that unifies several other approaches used in the field of explainability.