AI learns which market signals to trust

FAYETTEVILLE, GA, UNITED STATES, May 25, 2026 /EINPresswire.com/ -- A new study presents a neural-network method for building efficient portfolios while keeping the allocation process interpretable. Instead of treating asset correlations as a black box, the model learns how to clean noisy collective market patterns and respects the symmetries of covariance matrices. Tested on U.S. equities, it delivered higher returns per unit of risk under realistic trading conditions.

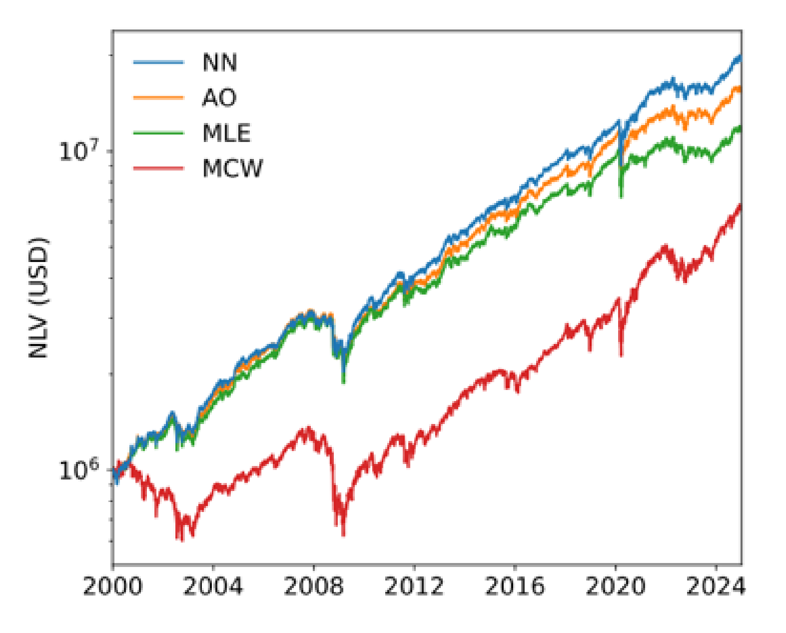

Backtests on U.S. equities show the neural-network portfolio compared with standard covariance-cleaning methods under realistic trading assumptions.

Building an efficient portfolio starts with a deceptively simple question: how do assets move together? In large markets, the answer is never observed cleanly. The data mix genuine collective movements with sampling noise, and small errors in this risk map can lead to unstable allocations.

In a study published in The Journal of Finance and Data Science, researchers from CentraleSupélec, the University of Catania and the University of Palermo, present a neural-network approach that learns how to clean these co-movement patterns before the portfolio is built. The goal is not to let AI replace portfolio theory, but to make AI work inside a transparent allocation framework.

The method focuses on global minimum-variance portfolios designed to control risk. The neural network is trained on the risk realized after allocation, so covariance cleaning is optimized for the portfolio it produces while the main steps of the process remain explicit.

A central idea of the study is that a covariance matrix is more than a table of pairwise correlations. It also contains collective market patterns. Some reflect broad market movements, some capture more specific structures, while others are mostly noise. Cleaning the matrix means deciding how much trust each collective pattern should receive before it affects the allocation.

This view makes the problem broader than portfolio construction. Whenever a noisy covariance matrix is transformed before being used in a decision, the transformation can be understood as a correction of these collective patterns. The study shows how a neural network can learn such corrections while remaining tied to the mathematical structure of the problem.

The main design choice is to build in the symmetries of covariance matrices. The result should not depend on the order in which stocks are listed, or on an arbitrary representation of the same risk structure. By respecting these invariances, the network learns a general cleaning rule rather than memorizing a fixed universe of assets.

In out-of-sample tests on U.S. equities from 2000 to 2024, a model calibrated on a few hundred stocks was applied, without retraining, to about one thousand stocks. The resulting portfolios achieved lower realized volatility, smaller drawdowns and higher Sharpe ratios than competing covariance estimators, including state-of-the-art nonlinear shrinkage methods. These gains persisted in a realistic trading simulation that included transaction costs, slippage, exchange fees and financing costs.

The study suggests that neural networks can become more useful in finance when they are not treated as black boxes, but designed around the symmetries and constraints of the system they are meant to learn.

References
DOI
10.1016/j.jfds.2026.100179

Original Source URL
https://doi.org/10.1016/j.jfds.2026.100179

Funding Information
This work was performed using HPC resources from the “Mésocentre” computing center of CentraleSupélec and École Normale Supérieure Paris-Saclay supported by CNRS and Région Île-de-France (http://mesocentre.centralesupelec.fr/). Efstratios Manolakis acknowledges a fellowship funded by PNRR for the PhD DOT1608375 in Sistemi complessi per le scienze fisiche, socio-economiche e della vita of Catania University.

Lucy Wang
BioDesign Research
email us here