Koopman Model Predictive Control for Portfolio Rebalancing

Amine Kobeissi [1] Aidan Li [1,2] Yue Xin Li [1]

[1]Université de Montréal [2]Mila

Introduction

The Challenge: Dynamic portfolio rebalancing aims to maximize risk-adjusted returns while strictly adhering to transaction costs and turnover limits. However, financial markets exhibit complex, nonlinear dynamics, making them hard to forecast with standard SSM MPC.

Solution: We propose a **Koopman-MPC** framework bridging Deep Learning and Control Theory:

- **Global Linearization:** A Koopman Autoencoder (KAE) lifts nonlinear market observations into a high-dimensional latent space where dynamics become approximately linear.
- **Convex Control:** We solve a convex Model Predictive Control (MPC) problem using the forecast market returns to determine optimal portfolio weights.

Key Results: Our Koopman MPC controller achieves a +173% total return, outperforming all benchmarks on both **absolute return** and **risk-adjusted return** over the 2021-2024 test period.

Problem Formulation

Let t denote the current rebalancing time. We solve a finite-horizon control problem over a prediction horizon H. The dynamics of the log-returns x are learned by the Koopman Autoencoder.

- **Decision Variables (Controls):** $w_k \in \mathbb{R}^{N+1}$ for $k = 0, \ldots, H-1$. Represents the target portfolio weights (Allocations to N risky assets + 1 Cash asset).
- **System Dynamics (Koopman):** We approximate the unknown market physics by lifting the state $\mathcal{X}_t$ into a latent space z where evolution is linear:
$$z_{k+1} = \mathcal{K} z_k, \quad \text{with } z_0 = \phi(\mathcal{X}_t)$$
- **Observation Model:** The log-returns $\hat{x}$ are recovered from the latent state via the learned decoder ψ:
$$\hat{x}_{t+k+1} = \psi(z_{k+1})$$
- **Forecasts:** We define the **full return vector** $\hat{R}_k \in \mathbb{R}^{N+1}$ used in the objective as:
$$\hat{R}_k = \begin{bmatrix} \exp(\hat{x}_{t+k+1}) \\ \exp(r_f) \end{bmatrix}$$
where r_f is the risk-free rate.
- **Initial State:** w_{t-1} is the current portfolio allocation at time t.

Objective Function (Log-Utility MPC):

$$\underset{w_0,\ldots,w_{H-1}}{\text{maximize}} \quad \sum_{k=0}^{H-1} \left(\underbrace{\log(w_k^T \hat{R}_k)}_{\text{Expected Growth}} - \underbrace{\lambda \|w_k - w_{k-1}\|_1}_{\text{Transaction Cost}} \right)$$
$$\text{subject to} \quad \mathbf{1}^T w_k = 1, \quad \forall k$$
$$w_k \geq 0, \quad \forall k \quad \text{(No Short Selling)}$$
$$\|w_k - w_{k-1}\|_1 \leq \delta_{max}, \quad \forall k \quad \text{(Turnover Limit)}$$

Data

- **Source:** Yahoo Finance API (Daily/Weekly Adjusted Close).
- **Assets:** Diverse basket of 20 US stocks (e.g. AAPL, BAC, GOOGL, GS).
- **Period:** Training (2012-2018), Validation (2019-2020), Test (2021-2024).

Koopman Model Predictive Control

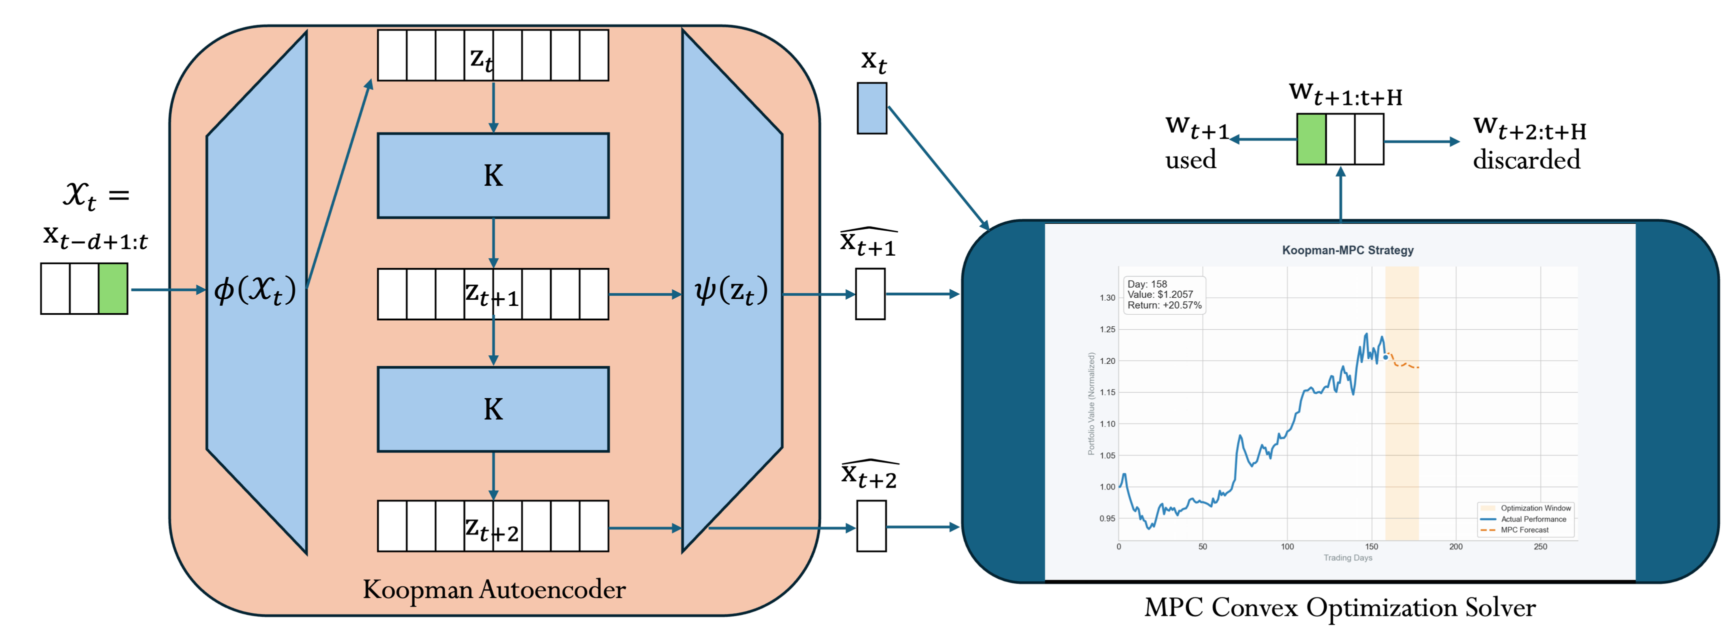

Figure 1. Simplified Koopman-MPC schematic for one asset. The last 6 weeks of price changes are used by the KAE to forecast the next 3 weeks, then the MPC provides controls to rebalance weights for the forcast period. The first set of controls are applied and the rest are discarded.

1. Learning Linear Koopman Dynamics

Leaning on Koopman operator theory, we lift nonlinear market states into a high-dimensional latent space where dynamics evolve approximately linearly.

- **Constructing the State:** We construct the state $\mathcal{X}_t$ using a sliding window of past log-returns ($d = 6$ and $N = 20$ assets): $\mathcal{X}_t = [x_{t-d+1}, \ldots, x_{t-1}, x_t] \in \mathbb{R}^{d \times N}$
- **Koopman Autoencoder (KAE):** Learn an encoder ϕ s.t. $z_t = \phi(\mathcal{X}_t)$, decoder ψ, and a global matrix $\mathcal{K}$ that advances the high-dimensional latent state: $z_{t+1} \approx \mathcal{K} z_t$.

2. Koopman Training Objectives

We optimize a composite loss to ensure the model captures stable, long-term market dynamics.

- **Reconstruction Loss:** $\|\mathcal{X} - \psi(\phi(\mathcal{X}))\|_2^2$ ensures latent state preserves information.
- **Dynamics Loss:** $\|\phi(\mathcal{X}_{t+1}) - \mathcal{K}\phi(\mathcal{X}_t)\|_2^2$ enforces the Koopman dynamics in latent space.
- **Prediction Loss:** $\sum_{k=1}^{H} \|\mathcal{X}_{t+k} - \psi(\mathcal{K}^k \phi(\mathcal{X}_t))\|_2^2$ minimizes error over a trajectory of length $H = 3$ to prevent error accumulation during rollout.
- **Sparsity:** Apply L_1 regularization to the encoder to learn interpretable market factors.

3. Koopman-MPC Algorithm

At each step t, we forecast $H = 3$ weeks of returns using the learned operator. These forecasts feed into the **Koopman MPC** solver to find optimal weights. We execute only the first action (w_0), shift the window, and repeat (Receding Horizon Control).

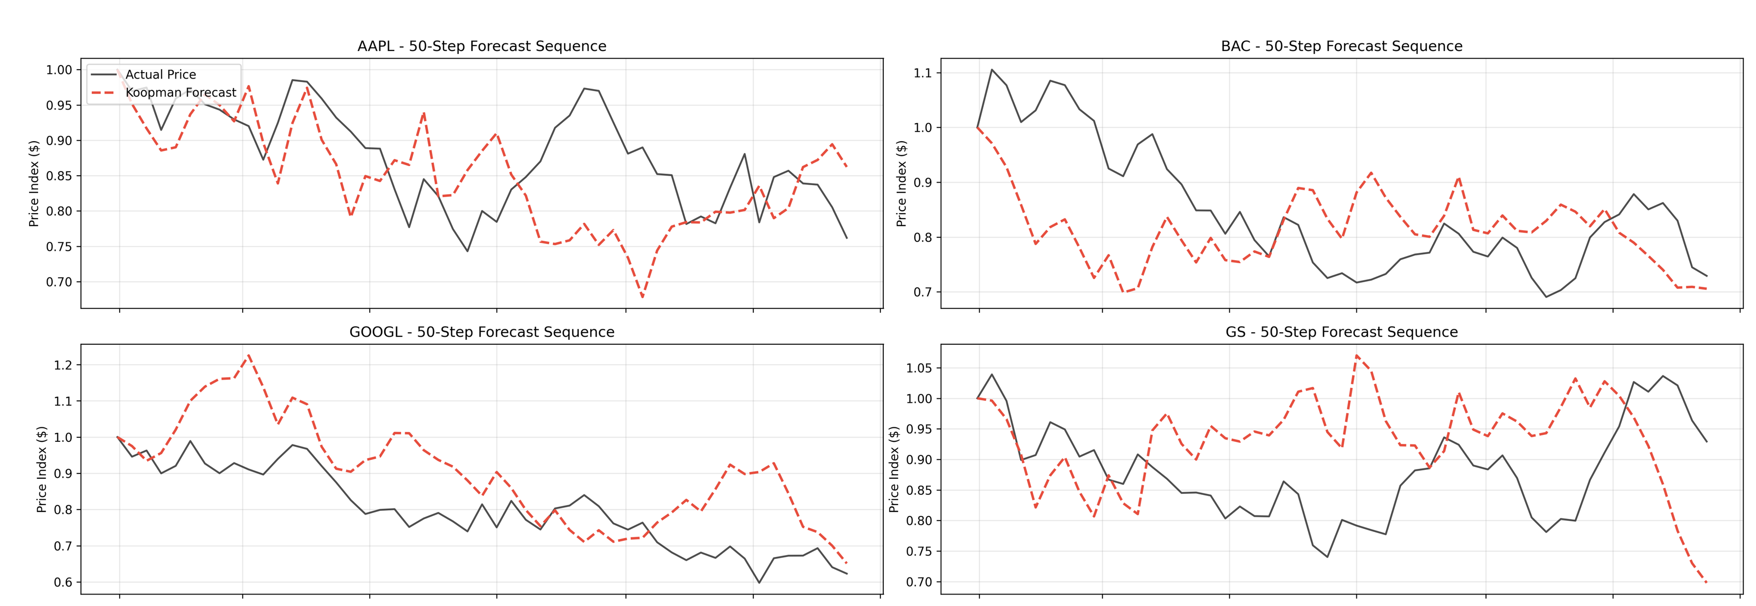

Figure 2. Koopman Predictions: This plot visualizes the predictions made by the KAE on 4 stocks in the test period.

Results

We evaluated the Koopman-MPC strategy against standard baselines from the literature over an unseen test period (2021-2024). We test our methods on the following evaluation metrics.

- **Sharpe Ratio (Risk-Adjusted Return):** Measures return per unit of volatility.
- **Max Drawdown (Risk):** The largest percentage drop from a peak to a trough.
- **Turnover (Cost):** The percentage of the portfolio traded per period.

Table 1. Performance comparison of portfolio rebalancing strategies.

Strategy	Sharpe	Total Return (%)	Max Drawdown (%)	Avg Turnover
Buy & Hold	1.129	+71.75	**-15.99**	0.000
Markowitz	0.905	+131.73	-19.86	0.092
DMD-MPC	0.810	+74.98	-48.03	0.200
VAR-MPC	1.243	+139.01	-28.46	0.200
Koopman-MPC (Ours)	**1.378**	**+173.08**	-29.56	0.200

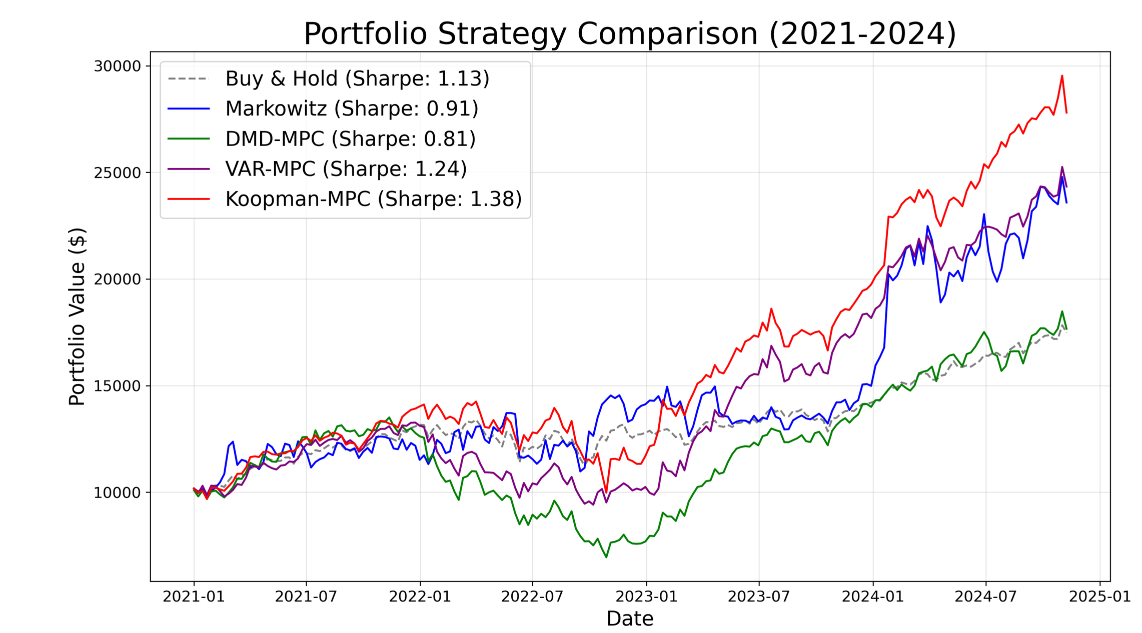

Figure 3. Equity Curve: Koopman-MPC strategy (Red) grows steadily with lower volatility than benchmarks.

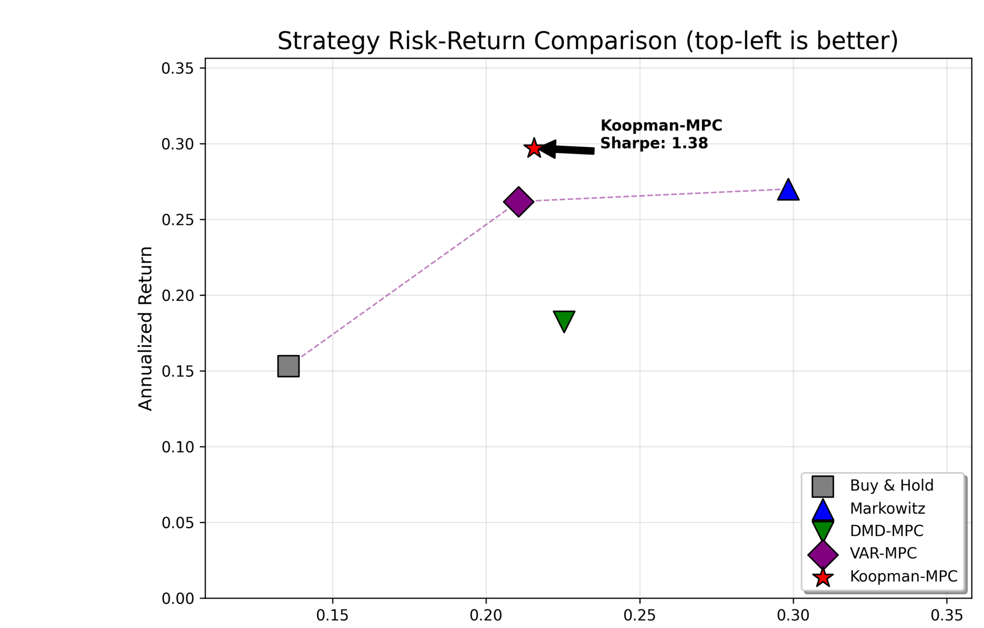

Figure 4. Trade-off between Risk and Return. Koopman-MPC (star) achieves the best trade-off.

Conclusion & Future Work

Our approach integrates deep representation learning with the safety guarantees of convex control, offering an alternative to black-box RL methods for portfolio selection. By lifting non-linear financial data into a high-dimensional latent space, we achieved superior forecasts for MPC decisions compared to linear baselines.

In future work, we plan to incorporate methods to better handle varying volatility regimes, such as gating different Koopman operators. Our strategy should also be tested on different asset classes and a larger stock universe over a longer time horizon to test generalization.