
Professional Certificate in Risk Modeling with Machine Learning

Machine Learning For Portfolio Optimization

Alpha – Expected excess return of a portfolio relative to a benchmark. Related: Beta, Sharpe ratio, CAPM. In portfolio optimization, alpha is the objective to maximize; models estimate it from factor exposures. Challenge: Estimating stable alpha in noisy data.

Alpha Decay – The reduction of predictive power of a model over time. Related: Concept drift, model retraining. Machine-learning models may capture transient patterns that fade, requiring periodic recalibration to maintain performance.

Amplitude Modulation – Technique to encode information by varying signal amplitude. Related: Time-series preprocessing, feature engineering. Rarely used directly in finance but can inspire methods for handling heteroscedasticity.

ARIMA (AutoRegressive Integrated Moving Average) – Classical time-series model combining autoregression, differencing, and moving average components. Related: SARIMA, Box-Jenkins methodology. Serves as a baseline for forecasting asset returns before applying more complex ML models.

Backtesting – Simulated evaluation of a strategy on historical data. Related: Walk-forward analysis, out-of-sample testing. Essential to assess portfolio optimization algorithms; must avoid look-ahead bias and overfitting.

Bayesian Optimization – Sequential model-based method for hyperparameter tuning. Related: Gaussian processes, acquisition function. Used to select regularization strength or number of trees in ensemble models for portfolio construction.

Beta – Sensitivity of a portfolio's returns to market movements. Related: Alpha, systematic risk, CAPM. In risk-aware optimization, beta constraints limit exposure to market volatility.

Black-Litterman Model – Framework combining equilibrium market returns with investor views. Related: Reverse optimization, prior distribution. Machine-learning can generate the view vector from predictive models, improving robustness.

Boosting – Ensemble technique that sequentially adds weak learners to correct errors. Related: AdaBoost, Gradient Boosting, XGBoost. Popular for predicting asset returns; must guard against overfitting due to high variance in financial data.

Box-Cox Transformation – Power transformation to stabilize variance. Related: Yeo-Johnson, log transformation. Applied to skewed return series before feeding into linear or tree-based models.

Brownian Motion – Continuous-time stochastic process with independent Gaussian increments. Related: Geometric Brownian motion, Wiener process. Underlies many asset-price models and provides simulated scenarios for Monte-Carlo portfolio evaluation.

Clustering – Unsupervised learning to group similar assets. Related: K-means, hierarchical clustering, DBSCAN. Helps create sector-aware portfolios or identify regimes for regime-switching models.

Concept Drift – Change in underlying data distribution over time. Related: Non-stationarity, adaptive learning. In portfolio optimization, drift may render a once-effective model obsolete, necessitating online learning or periodic retraining.

Cross-Validation – Technique to assess model performance on unseen data by partitioning the dataset. Related: K-fold, time-series split, rolling window. For financial data, use a forward-looking split to preserve temporal order.

CVaR (Conditional Value at Risk) – Expected loss beyond the VaR threshold. Related: VaR, tail risk, coherent risk measures. Optimization often minimizes CVaR to control extreme downside risk.

Data Snooping – Inadvertent use of information from the test set during model development. Related: Look-ahead bias, overfitting. Leads to inflated performance estimates; rigorous separation of training and evaluation periods mitigates this risk.

Deep Learning – Neural networks with multiple hidden layers. Related: LSTM, CNN, transformer. Capable of capturing complex nonlinear relationships in high-dimensional asset data, but require large datasets and careful regularization.

Dimensionality Reduction – Process of reducing the number of variables while preserving information. Related: PCA, t-SNE, autoencoders. Helps alleviate the curse of dimensionality in large asset universes and improves model stability.

Dropout – Regularization technique that randomly omits neurons during training. Related: L1/L2 regularization, early stopping. Prevents overfitting in deep neural networks used for return prediction.

Ensemble Methods – Combining multiple models to improve predictive performance. Related: Bagging, stacking, voting. In portfolio optimization, ensembles often yield more robust forecasts than single models.

Feature Engineering – Creation of informative variables from raw data. Related: Lagged returns, volatility measures, macro indicators. Critical for ML models; poorly engineered features can degrade performance dramatically.

Feature Selection – Process of choosing a subset of relevant features. Related: LASSO, recursive feature elimination, mutual information. Reduces overfitting risk and computational load, especially important when the asset universe is large.

Fisher Information – Measure of the amount of information a random variable carries about an unknown parameter. Related: Cramér-Rao bound, information matrix. Appears in portfolio theory when assessing estimator efficiency.

GBM (Geometric Brownian Motion) – Model for asset price dynamics assuming log-normal returns. Related: Black-Scholes, Monte-Carlo simulation. Provides baseline scenario generation for stress testing optimized

portfolios.

Gaussian Process – Non-parametric Bayesian model defining a distribution over functions. Related: Kernel methods, Kriging. Useful for quantifying predictive uncertainty in return forecasts, enabling risk-aware allocation.

Gradient Boosting – Boosting variant that fits residuals using gradient descent. Related: XGBoost, LightGBM, CatBoost. Dominant in return prediction competitions; hyperparameter tuning crucial to avoid overfitting to noisy financial signals.

Hedging – Constructing positions that offset potential losses. Related: Delta hedging, risk parity. ML-driven optimization can determine optimal hedge ratios given forecasted covariances.

Horizon – Investment period over which returns are evaluated. Related: Holding period, rebalancing frequency. Choice of horizon influences model design; short-horizon models may focus on high-frequency features, long-horizon on macro fundamentals.

Imputation – Filling missing values in datasets. Related: Forward fill, K-NN imputation. In finance, missing price data must be handled carefully to avoid biasing model estimates.

Information Ratio – Ratio of portfolio alpha to tracking error. Related: Sharpe ratio, Sortino ratio. Optimization may target maximization of the information ratio, balancing excess return against active risk.

Kernel Methods – Algorithms that operate in transformed feature spaces. Related: Support vector machines, Gaussian kernels. Enable nonlinear separation of asset characteristics without explicit feature mapping.

KNN (K-Nearest Neighbors) – Instance-based learning that predicts based on closest training points. Related: Distance metrics, weighting schemes. Simple baseline for return forecasting; suffers from curse of dimensionality in large asset sets.

Lagged Features – Past values of a variable used as predictors. Related: Autoregressive terms, rolling windows. Common in time-series models to capture momentum or mean-reversion effects.

Leverage – Use of borrowed capital to amplify returns. Related: Margin, risk-adjusted return. Optimization models may include leverage constraints to respect regulatory or risk limits.

LSTM (Long Short-Term Memory) – Recurrent neural network architecture designed for sequence data. Related: GRU, RNN. Effective at modeling temporal dependencies in price series, yet requires careful regularization to prevent overfitting.

Markowitz Mean-Variance Optimization – Classical framework balancing expected return against variance. Related: Efficient frontier, quadratic programming. ML enhancements replace the estimation of means and covariances with data-driven forecasts.

Maximum Drawdown – Largest peak-to-trough decline over a period. Related: Drawdown risk, recovery time. Often incorporated as a constraint in portfolio construction to limit tail risk.

Monte-Carlo Simulation – Repeated random sampling to assess portfolio outcomes. Related: Scenario analysis, stochastic modeling. Generates distributions of returns based on ML-predicted parameters, aiding risk assessment.

Momentum – Tendency of assets that performed well to continue doing so in the short term. Related: Trend following, mean reversion. Feature engineering frequently includes momentum indicators such as 12-month cumulative returns.

Multicollinearity – High correlation among predictor variables. Related: Variance inflation factor, ridge regression. Can inflate coefficient variance in linear models; regularization or dimensionality reduction mitigates the issue.

Neural Architecture Search – Automated process of discovering optimal network structures. Related: Hyperparameter optimization, AutoML. Can produce bespoke models for specific asset classes, though computationally intensive.

Non-Stationarity – Property of a time series whose statistical characteristics change over time. Related: Unit root, differencing. Requires techniques like rolling windows or adaptive learning to maintain model relevance.

Normalization – Scaling features to a common range. Related: Standardization, min-max scaling. Essential for algorithms sensitive to magnitude, such as neural networks and distance-based methods.

Overfitting – Model captures noise rather than underlying pattern. Related: Regularization, cross-validation. In finance, overfitted models often produce spectacular backtest results but fail in live trading.

Parsimony – Preference for simpler models that achieve comparable performance. Related: Occam's razor, model complexity. Encouraged to improve interpretability and reduce overfitting risk in portfolio optimization.

Passive Investing – Strategy that tracks a benchmark rather than actively selecting securities. Related: Index funds, tracking error. ML can be used to construct smart-beta indices that blend passive exposure with factor tilts.

Performance Attribution – Decomposition of portfolio returns into source components. Related: Factor attribution, sector attribution. Helps evaluate whether ML-driven signals are adding value beyond market movements.

Portfolio Turnover – Frequency of trading required to maintain target weights. Related: Transaction costs, liquidity. Optimization often penalizes turnover to balance expected gains against trading expenses.

Predictive Modeling – Building statistical or ML models to forecast future outcomes. Related: Regression, classification, time-series forecasting. Core of ML-based portfolio construction; quality of predictions drives allocation decisions.

Probabilistic Forecasting – Producing a full distribution rather than a point estimate. Related: Quantile

regression, Bayesian inference. Enables risk-aware optimization by incorporating uncertainty directly into the objective.

Quadratic Programming – Optimization method for problems with a quadratic objective and linear constraints. Related: Convex optimization, interior-point methods. Underpins many mean-variance formulations; scalable solvers are critical for large asset universes.

Random Forest – Ensemble of decision trees trained on bootstrapped samples. Related: Bagging, feature importance. Robust to overfitting and provides interpretable variable importance metrics for return predictors.

Rebalancing Frequency – Interval at which portfolio weights are adjusted. Related: Daily, monthly, quarterly. Trade-off between capturing new signals and incurring transaction costs; ML can recommend optimal frequency based on market regime.

Regularization – Adding penalty terms to loss functions to discourage complexity. Related: L1 (LASSO), L2 (Ridge), Elastic Net. Crucial for stabilizing coefficient estimates when training data is limited relative to predictors.

Risk Parity – Allocation method that equalizes risk contribution across assets. Related: Volatility weighting, diversification. ML can estimate forward-looking risk measures to improve risk-parity allocations.

Rolling Window – Moving time window used for model training and evaluation. Related: Expanding window, walk-forward analysis. Provides up-to-date parameter estimates while preserving temporal order.

Scenario Analysis – Examination of portfolio performance under predefined market conditions. Related: Stress testing, macro shocks. ML-generated scenarios can reflect realistic joint movements of assets.

Sharpe Ratio – Excess return per unit of total risk (standard deviation). Related: Information ratio, Sortino ratio. Often used as an objective function; however, it assumes normally distributed returns.

Shrinkage Estimator – Technique that pulls sample covariance matrix toward a structured target. Related: Ledoit-Wolf, Bayesian shrinkage. Improves stability of covariance estimates used in mean-variance optimization.

Signal-to-Noise Ratio – Measure of predictive strength relative to random variation. Related: Information coefficient, t-statistic. Low ratios in finance demand robust modeling and careful validation.

Simplified Portfolio Theory (SPT) – Approximation that treats assets as independent for tractable optimization. Related: Factor models, diagonal covariance. ML can estimate factor sensitivities to relax independence assumptions.

Smoothing – Technique to reduce volatility in time-series signals. Related: Moving average, exponential smoothing. Helps stabilize input features for ML models but may lag behind rapid market changes.

Stochastic Gradient Descent (SGD) – Iterative optimization algorithm using random subsets of data. Related: Mini-batch, learning rate. Preferred for training large neural networks on high-frequency financial data.

Stress Testing – Evaluation of portfolio resilience under extreme but plausible conditions. Related: Scenario analysis, tail risk. ML can generate stress scenarios by perturbing model inputs in realistic ways.

Structural Break – Point at which the statistical properties of a series change abruptly. Related: Regime shift, Chow test. Detecting breaks is vital for updating ML models to maintain predictive accuracy.

Supervised Learning – Learning paradigm where models are trained on labeled data. Related: Regression, classification. Most return-prediction tasks fall under supervised learning, requiring historical return labels.

Support Vector Machine (SVM) – Classification/regression algorithm that maximizes margin between classes. Related: Kernel trick, soft margin. Can be applied to predict binary outcomes such as up/down movements; careful tuning needed for financial noise.

Swaption – Option granting the right to enter an interest-rate swap. Related: Derivatives, volatility surface. Pricing models may incorporate ML-estimated volatility surfaces to improve hedging decisions.

Technical Indicator – Quantitative measure derived from price and volume data. Related: RSI, MACD, Bollinger Bands. Frequently used as features; must be validated for predictive relevance.

Temporal Fusion Transformer (TFT) – Attention-based architecture for multi-horizon time-series forecasting. Related: Transformers, variable selection networks. Demonstrated strong performance on financial series where multiple covariates evolve over time.

Transaction Cost Model – Quantitative representation of fees, slippage, and market impact. Related: Linear cost, non-linear impact. Incorporating realistic cost models into optimization prevents over-trading driven by ML signals.

Training-Test Split – Division of data into subsets for model fitting and evaluation. Related: Hold-out, validation set. In finance, splits must respect chronological order to avoid contaminating future information.

Transfer Learning – Reusing knowledge from a source domain to improve learning in a target domain. Related: Fine-tuning, domain adaptation. Pre-trained models on large market data can be adapted to niche asset classes.

Tree-Based Models – Algorithms that partition feature space using decision trees. Related: CART, random forest, gradient boosting. Offer interpretability via feature importance and handle mixed data types well.

Trend Following – Strategy that invests in assets exhibiting persistent price movements. Related: Momentum, breakout. ML can enhance trend detection by combining multiple time-scale signals.

Unsupervised Learning – Learning from data without explicit labels. Related: Clustering, dimensionality reduction. Useful for discovering latent structures such as hidden market regimes.

Variance Inflation Factor (VIF) – Metric quantifying multicollinearity among predictors. Related: Condition number, ridge regression. High VIF values suggest the need for feature elimination or regularization.

Volatility Clustering – Phenomenon where high-volatility periods tend to cluster together. Related: GARCH,

stochastic volatility. Modeling this effect improves risk forecasts used in portfolio construction.

Weighted Least Squares (WLS) – Regression technique assigning different weights to observations. Related: Heteroscedasticity, robust regression. Can prioritize recent data points when estimating factor returns.

Yield Curve – Graph of interest rates across different maturities. Related: Term structure, Nelson-Siegel model. ML can capture dynamics of the curve for fixed-income portfolio optimization.

Z-Score Normalization – Centering data to zero mean and unit variance. Related: Standardization, scaling. Common preprocessing step for models sensitive to feature magnitude, such as SVMs and neural networks.

Zero-Mean Portfolio – Portfolio constructed to have zero expected return, focusing on risk reduction. Related: Market neutral, hedged. ML can identify assets that offset each other's exposures, achieving neutrality.

Alpha-Beta Decomposition – Separation of portfolio return into market-related (beta) and stock-specific (alpha) components. Related: Factor model, residual return. Useful for evaluating the contribution of ML-generated signals beyond systematic exposure.

Backpropagation – Algorithm for computing gradients in neural networks. Related: Chain rule, learning rate. Core of training deep learning models for financial forecasting.

Bayesian Inference – Statistical method that updates prior beliefs with observed data. Related: Posterior distribution, conjugate priors. Provides a natural framework for incorporating uncertainty in model parameters used for portfolio allocation.

Bootstrap Resampling – Generating multiple datasets by sampling with replacement. Related: Confidence intervals, out-of-bag error. Enables estimation of variability in performance metrics for ML models in finance.

Capital Asset Pricing Model (CAPM) – Linear model linking expected return to systematic risk (beta). Related: Market portfolio, risk-free rate. Serves as a benchmark for evaluating the added value of ML-driven alpha.

CatBoost – Gradient-boosting library that handles categorical features natively. Related: LightGBM, XGBoost. Often yields superior performance on tabular financial data with mixed variable types.

Change Point Detection – Identifying times where statistical properties shift. Related: Bayesian online change point detection, CUSUM. Alerts practitioners to moments when model retraining may be required.

Conditional Expectation – Expected value of a variable given certain information. Related: Law of iterated expectations, regression. Central to mean-variance optimization where future returns are conditioned on predictive signals.

Correlation Matrix – Square matrix showing pairwise correlations among assets. Related: Covariance matrix, eigenvalue decomposition. Accurate estimation is vital for risk budgeting; shrinkage techniques improve stability.

Cross-Asset Modeling – Simultaneous modeling of multiple asset classes. Related: Multi-task learning, hierarchical models. Allows ML to capture inter-market dynamics that improve diversification benefits.

Decile Portfolio – Grouping assets into ten equal-size buckets based on a ranking metric. Related: Quantile sorting, factor portfolios. Facilitates evaluation of predictive signals by comparing performance across deciles.

Deep Reinforcement Learning – Combination of deep learning with reinforcement learning for sequential decision making. Related: Q-learning, policy gradient. Emerging approach for dynamic portfolio rebalancing under transaction cost constraints.

Diffusion Model – Stochastic process describing how a variable evolves continuously. Related: Ornstein-Uhlenbeck, mean-reversion. Provides analytical forms for asset dynamics used in scenario generation.

Elastic Net – Regularization method blending L1 and L2 penalties. Related: LASSO, Ridge. Useful for feature selection when predictors are highly correlated, as common in financial factor libraries.

Ensemble Kalman Filter – Recursive algorithm for estimating hidden states in nonlinear systems. Related: Data assimilation, particle filter. Can be employed to update latent factor estimates as new market data arrive.

Factor Model – Representation of asset returns as a linear combination of common risk factors. Related: Fama-French, Barra. ML can estimate factor loadings and residual variances more flexibly than traditional OLS.

Gaussian Mixture Model (GMM) – Probabilistic model assuming data are generated from a mixture of Gaussian distributions. Related: EM algorithm, clustering. Captures multimodal return distributions for scenario analysis.

Hierarchical Risk Parity (HRP) – Portfolio construction method that uses hierarchical clustering to allocate risk. Related: Tree-based allocation, diversification. Avoids matrix inversion, making it robust to estimation error; ML can refine clustering inputs.

Information Coefficient (IC) – Correlation between predicted and realized returns. Related: Rank IC, predictive power. Serves as a performance metric for ML models; low IC values indicate limited usefulness.

Kernel Density Estimation (KDE) – Non-parametric method to estimate probability density functions. Related: Bandwidth selection, smoothing. Useful for visualizing return distributions and for constructing empirical risk measures.

Knock-out Options – Derivatives that become worthless if the underlying asset breaches a barrier. Related: Barrier options, exotic derivatives. Pricing models may incorporate ML-estimated volatility surfaces for more accurate valuation.

Laguerre Polynomials – Orthogonal polynomials used in function approximation. Related: Basis expansion,

spectral methods. Occasionally employed to model term structure dynamics in fixed-income portfolio optimization.

Learning Rate – Hyperparameter controlling step size in gradient-based optimization. Related: Decay schedule, Adam optimizer. Too high a learning rate can cause divergence; too low slows convergence, especially in noisy financial data.

Liquidity Risk – Risk arising from the inability to trade assets without affecting price. Related: Market impact, bid-ask spread. Optimization models often include liquidity constraints to avoid excessive exposure to illiquid securities.

Markov Chain Monte Carlo (MCMC) – Class of algorithms for sampling from complex probability distributions. Related: Metropolis-Hastings, Gibbs sampling. Enables Bayesian estimation of portfolio weights when analytical solutions are unavailable.

Mean Reversion – Tendency of asset prices to return to a long-term average. Related: Ornstein-Uhlenbeck process, pair trading. ML models may capture mean-reverting signals through lagged features and regime detection.

Monte-Carlo Dropout – Technique that interprets dropout at inference time as a Bayesian approximation. Related: Uncertainty quantification, predictive intervals. Provides a cheap way to assess confidence in neural-network forecasts for portfolio risk.

Multitask Learning – Training a model on several related tasks simultaneously. Related: Shared representation, transfer learning. In finance, can jointly predict returns for multiple assets, improving data efficiency.

Neural Tangent Kernel (NTK) – Analytical tool describing the behavior of infinitely wide neural networks. Related: Kernel methods, convergence. Offers insight into why certain deep architectures generalize well on financial data.

Noise-Robust Loss – Loss functions less sensitive to outliers. Related: Huber loss, quantile loss. Useful when financial returns contain extreme events that would otherwise dominate training.

Optimal Stopping – Decision problem of choosing the time to take a particular action to maximize expected reward. Related: American option pricing, dynamic programming. Reinforcement-learning agents use this framework for trade execution.

Out-of-Sample Performance – Measure of model accuracy on data not used during training. Related: Backtesting, validation set. Critical for assessing whether ML-driven portfolio strategies will succeed in live markets.

Partial Least Squares (PLS) – Dimensionality reduction technique that projects predictors and response onto latent variables. Related: Canonical correlation analysis, regression. Handles collinear features common in macro-economic datasets used for asset allocation.

Portfolio Optimization – Process of selecting asset weights to achieve a desired risk-return trade-off. Related: Mean-variance, robust optimization. Machine learning enhances the estimation of expected returns and covariances, the core inputs to optimization.

Predictor Variable – Input feature used to forecast a target. Related: Independent variable, explanatory variable. In finance, predictors range from technical indicators to macroeconomic releases.

Quantile Regression – Regression technique estimating conditional quantiles of the response variable. Related: Expectile regression, VaR estimation. Provides asymmetric loss functions useful for tail-risk-focused portfolio construction.

Recursive Feature Elimination (RFE) – Wrapper method that iteratively removes least important features. Related: Feature importance, model selection. Helps identify a parsimonious set of predictors for return forecasting.

Regularized Portfolio Optimization – Incorporating penalty terms (e.g., L1 for sparsity) into the objective function. Related: Elastic net, constrained optimization. Produces more stable and interpretable allocations, especially when the number of assets exceeds observations.

Risk Budgeting – Allocating risk contributions rather than capital. Related: Risk parity, volatility targeting. ML forecasts of future volatilities feed directly into risk-budget calculations.

Robust Optimization – Optimization approach that accounts for uncertainty in parameters. Related: Ambiguity set, worst-case scenario. Addresses estimation error in ML-predicted means and covariances, producing portfolios less sensitive to mis-specification.

Rolling Forecast Origin – Strategy where the start point of the training window moves forward in time for each forecast. Related: Expanding window, walk-forward validation. Mimics real-world deployment of models that update continuously.

Sample Covariance Matrix – Empirical estimator of asset return covariances. Related: Shrinkage estimator, factor model. Often ill-conditioned in high-dimensional settings; regularization improves its use in optimization.

Scenario Generation – Creation of plausible future paths for risk factors. Related: Monte-Carlo simulation, copula models. ML techniques such as variational autoencoders can produce realistic joint distributions.

Sharpe Ratio Maximization – Objective of maximizing excess return per unit of total risk. Related: Mean-variance, utility function. Directly optimizing this ratio is non-convex; approximations or surrogate objectives are employed.

Shrinkage Target – Structured matrix toward which the sample covariance is pulled. Related: Identity matrix, constant correlation. Choosing an appropriate target reduces estimation error in high-dimensional portfolios.

Smoothing Parameter – Controls the degree of regularization in non-parametric regression. Related:

Bandwidth, penalty term. Important when fitting splines to noisy financial data.

Spurious Correlation – Apparent relationship caused by chance rather than causality. Related: Multiple testing, data mining bias. ML pipelines must include statistical controls to avoid exploiting such false patterns.

Stochastic Dominance – Preference ordering of distributions based on expected utility. Related: First-order, second-order dominance. Can be incorporated into portfolio selection to ensure dominance over a benchmark.

Structure-Preserving Transformation – Mapping that maintains relationships (e.G., Monotonicity) between variables. Related: Rank transformation, copula. Useful when normalizing returns while retaining dependence structure for risk modeling.

Super-hedging – Constructing a portfolio that dominates the payoff of a derivative under all scenarios. Related: Convex risk measures, dual representation. ML can approximate super-hedging costs by learning worst-case scenarios.

Synthetic Asset – Constructed security whose payoff mimics a desired exposure. Related: Replicating portfolio, factor mimicry. ML can design synthetic assets that capture complex factor tilts not directly investable.

Target Volatility Strategy – Portfolio that adjusts exposure to maintain a constant volatility level. Related: Risk budgeting, volatility scaling. Forecasted volatilities from ML models drive the scaling factor.

Temporal Difference Learning – Reinforcement-learning method that updates value estimates based on successive predictions. Related: TD(λ), Q-learning. Applied to learn optimal trading policies that adapt to market dynamics.

Time-Series Cross-Validation – Validation technique that respects temporal ordering, such as expanding window. Related: Rolling origin, forward chaining. Prevents leakage of future information into model training.

Transaction Cost Analysis (TCA) – Assessment of the cost incurred when executing trades. Related: Implementation shortfall, slippage. Incorporating TCA into the optimization objective prevents eroding ML-generated alpha.

Transformers – Attention-based neural architectures that process sequences in parallel. Related: Self-attention, positional encoding. Emerging models for financial time series, offering improved long-range dependency capture over RNNs.

Unbiased Estimator – Statistic whose expected value equals the true parameter. Related: Bias-variance trade-off, consistency. In finance, unbiasedness may be sacrificed for lower variance via regularization.

Value at Risk (VaR) – Quantile of the loss distribution over a specified horizon. Related: CVaR, tail risk. Often used as a constraint in portfolio optimization to limit potential losses.

Variance-Covariance Matrix – Matrix containing variances along the diagonal and covariances off-diagonal. Related: Correlation matrix, factor decomposition. Accurate estimation is crucial for any risk-based allocation.

Weighted Portfolio – Portfolio where each asset carries a specific weight reflecting its allocation. Related: Capital allocation, fractional shares. Optimization determines the set of weights that best achieve the objective.

Yield Curve Modeling – Statistical representation of interest rates across maturities. Related: Nelson-Siegel, spline models. Machine-learning can capture nonlinear dynamics and improve fixed-income portfolio decisions.

Z-Score Portfolio – Portfolio constructed by ranking assets based on their standardized scores. Related: Mean-variance, ranking strategy. Allows comparison across assets with differing scales, facilitating cross-asset allocation.

Zero-Sum Game – Situation where one participant's gain equals another's loss. Related: Competitive markets, game theory. In portfolio optimization, market impact can be modeled as a zero-sum interaction among traders.