Local Extrema Predictor (LEAP)

A machine learning-based market regime detection and mean reversion strategy

This paper presents the Local Extrema Predictor (LEAP), a novel strategy for forecasting mean reversion (MR) in foreign exchange (FX) markets. LEAP addresses key shortcomings of existing MR methodologies, including dependence on subjective labelling, rigid model structures, limited adaptability to the dynamic nature of FX markets, high-dimensional input spaces and narrow feature selection.

Mehryar Emambakhsh, Ph.D.
Vice President and
Senior Research Scientist
Mesirow Currency

Using a broad set of around 568 FX and non-FX data sources, LEAP formulates MR prediction as a ternary supervised classification problem, assigning each data point as a peak, trough, or normal, based on the topographic profile of the target series. The core classification model is based on Support Vector Machines (SVM), with hyperparameters optimised using Bayesian techniques.

To account for the time-varying influence of economic indicators and reduce feature dimensionality, a genetic algorithm dynamically selects the most relevant signals for each currency pair. LEAP has delivered strong performance across 27 USD-denominated currency pairs, spanning both emerging and developed markets, and has been incorporated into Mesirow Currency's Alpha engine since April 2025.

Keywords

- Mean reversion
- FX
- Feature extraction
- Wavelets
- Genetic Algorithm
- Support Vector Machines

1 | Introduction

Mean reversion (MR) is a financial theory that suggests asset prices and historical returns tend to revert to their long-term average over time. MR is also used for market regime detection and risk sentiment analysis. In machine learning (ML), it can be framed as an anomaly detection problem over sequential data.

MR in financial markets is driven by a mix of economic, behavioural, and institutional factors. In foreign exchange (FX) markets, key influential factors include macroeconomic fundamentals (such as interest rate differentials), central bank actions, investor sentiment, and geopolitical events. These factors often interact in complex, non-linear ways. This makes MR detection highly context-dependent and computationally difficult.

This article introduces the Local Extrema Predictor (LEAP), a technical strategy designed to identify MR in financial time series, with a focus on FX spot rates. LEAP uses a wide range of economic data — both FX and non-FX — as input to a supervised ML model. The model classifies each day as a peak, a trough, or a normal point. Since it is a supervised approach, it requires labelled training data. Labels are generated through an objective process that detects local minima and maxima in the target series using their topographic characteristic.

LEAP includes a set of pre-processing and feature extraction techniques that summarise statistical patterns such as momentum and trendiness in the input time series. It uses Support Vector Machines (SVM) for classification, given their strong generalisation performance in high-dimensional anomaly detection tasks. Hyper-parameters are tuned using Bayesian optimisation within a rolling K-fold cross-validation setup. Feature selection is performed with a genetic algorithm (GA), which identifies the most relevant financial indicators. This helps reduce dimensionality while also improving performance. Since influential economic drivers can shift over time and vary across currency pairs, the GA feature selector also acts as an adaptive mechanism, capable of updating the selected inputs as conditions change. Figure 1 shows various steps in a block diagram.

While the method can be extended to detect MR in other financial instruments, we focus here on daily spot rates of 27 USD-based currency pairs from both emerging (EM) and developed (DM) markets.

The rest of the paper is organised as follows. Section 2 reviews related work. Section 3 explains the labelling method. Section 4 introduces the economic indicators used as inputs. Sections 5, 6 and 7 describe the feature construction and classification process. Section 8 presents the experimental results. We conclude in Section 9.

FIGURE 1: **LEAP**

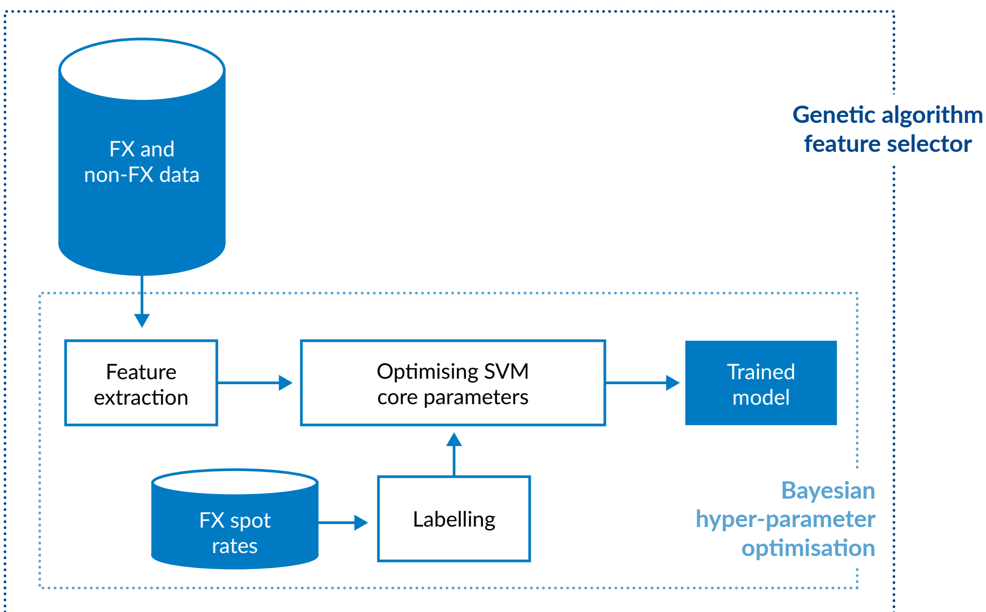

An overview of LEAP — our proposed strategy for detecting MR on currency pairs. The dashed lines indicate iterative processes.

2 | Brief review of relevant work

2.1 | MR DETECTION THROUGH RISK FORMULATION

MR detection and use as a systematic trading strategy for FX markets is a complex problem. Relevant research touches on volatility modelling, market microstructure, tail risk, and sentiment — each offering a unique perspective on MR behaviour.

Andersen and Bondarenko's work[1] with Model-Free and Corridor Implied Volatility (MFIV and CIV) introduces volatility estimates derived from option prices without relying on strict model assumptions. CIV, especially in narrow bands, has outperformed both VIX and MFIV in forecasting volatility. In FX markets, where timely volatility signals can flag potential MR following sharp moves, this is particularly useful.

Chan's risk indicator study on FX markets[2] finds that while crises can break short-term MR patterns, they often pave the way for post-shock corrections once volatility subsides — creating attractive reversion opportunities. Investor sentiment cycles also influence MR. The Goldman Sachs Risk Appetite Index (RAI),[6] for instance, shows that extreme readings (e.g. below -1.5) often precede contrarian reversals, signalling conditions favourable for MR trades. Similarly, BIS and Citi have developed sentiment indicators, which combine cross-asset correlations and volatility ranks.[4,9] These metrics have been linked to currency return predictability, particularly during sentiment-volatility mismatches.

Citi's Early Warning Signal (EWS) tool[5] goes further, integrating macroeconomic and stress indicators to generate hedge signals for EM currencies. Spikes in EWS scores often precede sharp corrections, aligning closely with MR behaviour.

A different point of view is proposed by Lyons' market microstructure framework,[7] suggesting MR as a rational response, not just through statistical patterns: They link short-term price distortions to order flow and dealer inventory imbalances at micro level.

2.2 | TOWARDS AN ML FRAMEWORK FOR MR DETECTION

While effective in certain controlled scenarios, these approaches often fall short in highly dynamic and noisy environments where nonlinear interactions dominate. The limitations of such models — especially in capturing high-dimensional dependencies and adaptive behaviour — have paved the way for data-driven, ML methodologies. As the literature suggests, the manifestation of MR is neither time-invariant nor uniformly distributed across market conditions.[6,9] Therefore, a robust ML-based framework must integrate diverse signals — from both FX and non-FX domains — while retaining interpretability, scalability, and predictive accuracy.

Implementing ML for MR classification and forecasting in FX markets, however, entails three core challenges:

1. Data labelling
In supervised learning, a labelled dataset is required. However, defining what constitutes a "mean-reverting event" is inherently subjective and lacks a universal standard.

2. Feature extraction
Identifying predictive features for MR is notoriously difficult. Unlike tasks with clearly defined temporal or spatial signals, MR events do not arise from a fixed set of indicators. Economic drivers, volatility bursts, sentiment shifts, and structural imbalances can all influence reversions — but at varying timescales and intensities.

3. High Dimensionality
Initially, incorporating all available indicators might seem beneficial for enhancing performance — practically, this leads to a prohibitively large dimensional feature space. High dimensionality increases the risk of overfitting, particularly when the number of labelled examples is small. Dimensionality reduction methods such as Principal Component Analysis (PCA) or Linear Discriminant Analysis (LDA) can mitigate this issue, but often at the cost of interpretability.

2.3 | PROPOSED FRAMEWORK

Considering these challenges, we propose a supervised ML framework that addresses each of the core issues. Instead of relying on manually defined regimes, the framework derives labels from the trajectory of target spot rates. Specifically, the model assigns each day a label — peak, trough, or neutral — based on local extrema in the time series.

This objective labelling scheme improves reproducibility and reduces subjective bias. The framework incorporates a wide range of both FX and non-FX data, facilitating broader context for each decision point. By combining automated feature selection with GA and a Bayesian hyper-parameter optimisation we propose capturing the most informative feature subsets for the most optimal classifiers.

3 | Labelling time series for an MR problem

We adopt a supervised classification framework, which requires input samples to be labelled as peaks, troughs, or normal points. Labelling is based on local neighbourhood comparisons, deliberately over-resolving to detect a dense set of candidate extrema.

Each extremum is then assigned a prominence score (see endnote for a definition of topographic prominence[10]), borrowing the definition from topography — a measure of how much a peak stands out relative to its surrounding landscape. In our context, prominence quantifies how significantly a peak or trough deviates from recent market movements.

We apply central differencing for derivatives and use the nearest rightmost peak when calculating prominence, meaning both operations are non-causal as future data is used in the computations. To avoid information leakage, clear temporal gaps are maintained between the training, validation, and test sets.

To further refine the labels, we apply non-maximum suppression (NMS), a technique commonly used in computer vision. NMS filters out less prominent extrema, retaining only the most distinct peaks and troughs. This produces a cleaner and more meaningful set of labelled anomalies. Figure 2 illustrates an example of the output of the labelling process.

FIGURE 2: LABELLING

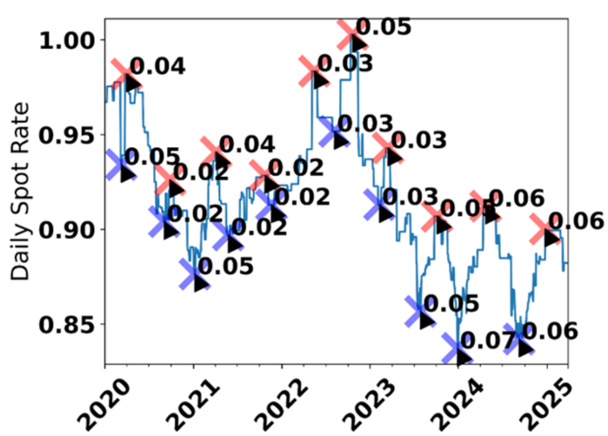

An example labelling USDCHF spot currency rate. The extrema are automatically detected as peaks and troughs over the input target currency pair using topographic formation. The numbers next to each extremum illustrate the prominence of each peak or trough.

4 | Underlying indicators

MR in FX can arise from a wide range of factors, which vary across currency pairs and are not consistently predictive. In this section, we outline commonly observed indicators that influence MR in the FX market.

FIGURE 3: **LEAP INPUT DATA**

An illustration of various categories of FX and non-FX data used during the development of LEAP.

While the next section details how reversion signals are detected, we first summarise the categories of data used in our strategy (also depicted in Figure 3).

1. Currency-related Indicators | These are variables that show clear, direct signals of potential or emerging MR in a given currency pair. Examples include daily or intraday interest rates (USD and non-USD), futures prices, implied volatility, and broader benchmarks like the Dollar Index (DXY).

2. Commodities | Gold and oil prices are often viewed as global risk barometers. Movements in these prices can signal shifts in risk sentiment (risk-on/risk-off), which in turn affect currency values.

3. Global Risk Indicators | This category includes data such as the VIX, economic surprise indices, and even cryptocurrencies like Bitcoin and Ethereum. Though unconventional, digital assets can sometimes reflect shifts in dollar-driven reversion patterns.

4. Stock Market | Equity indices can provide insight into a country's perceived economic strength. Comparing indices across countries helps infer market sentiment around the relative value of their currencies.

5. Credit Indicators | We include metrics such as sovereign credit default swaps (CDS), short-term debt levels, yield curve slopes, and the TED spread. These indicators can offer clues about FX performance.

6. Macro and Microeconomic Data | This group comprises fundamental indicators affecting national economies. Examples include consumer durable spending, housing market data, exports, money supply, and industrial production figures.

Given FX's high efficiency, many of these indicators may act as lagging signals. However, those tied closely to a country's core economic fundamentals can serve as leading indicators. Not all of these indicators are used in the final model, but they form the conceptual foundation of our ML-driven approach to MR detection in FX. In the next section, we explain how these inputs are transformed into abstract features for the model.

5 | Feature generation

FIGURE 4: **FEATURE SPACE CONSTRUCTION**

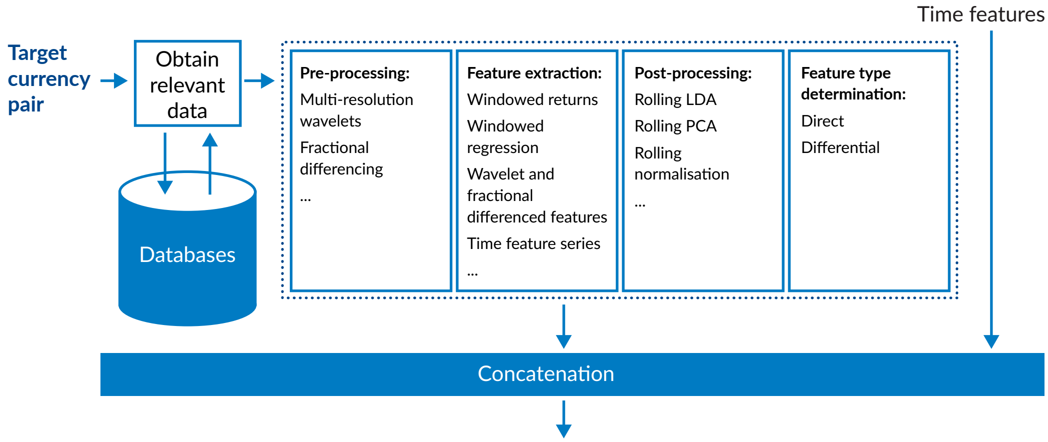

The feature space construction illustrated as a block diagram.

5.1 | PRE-PROCESSING

The data is first pre-processed to prepare it for feature extraction, as shown in Figure 4. Two key techniques used in this phase are fractional differencing and multi-resolution wavelet denoising. Fractional differencing is applied to ensure stationarity while preserving the memory of the time series. The intensity of differencing is adjusted adaptively using the Augmented Dickey-Fuller (ADF) test. Once the *p*-value reaches an acceptable threshold, differencing is halted.

The other pre-processing technique is Wavelet denoising. Applied at multiple resolutions, it reconstructs the signal at each time step using its approximate components. This results in a smoother time series, making the subsequent feature extraction more robust to high-frequency noise and better at detecting trend shifts.

5.2 | FEATURE EXTRACTION TECHNIQUES

Feature extraction transforms raw input data into structured feature vectors. Its goal is to produce an abstract yet simplified representation that captures the underlying structure relevant to the classification or regression task.

We apply several feature extraction algorithms to the pre-processed sequences, aiming to represent each instance in a way that reflects its class. For example, a sequence labelled as a trough should produce a feature vector similar to other troughs, but distinct from peaks or normal points. The two main techniques we use are windowed returns and windowed regression.

In the context of MR, it is essential to quantify trends within the data. Trend detection depends on the chosen look-back window size: longer windows reveal persistent trends, while shorter windows highlight more recent movements. To capture both long- and short-term patterns, we use multiple window sizes.

For windowed returns, we calculate cumulative returns over each window. For windowed regression, we fit a linear regression model within each window and extract the slope. Additionally, we compute the R^2 value for each regression to assess the fit quality. Including R^2 helps differentiate between sequences with similar slopes but varying levels of regression confidence.

5.3 | POST-PROCESSING

The final step involves post-processing the features to produce well-structured vectors. This ensures the feature space is normalised and constrained within a defined range. To achieve this, we apply rolling LDA and PCA. LDA performs supervised dimensionality reduction by projecting the features onto directions that maximise class separability. PCA, on the other hand, offers an unsupervised alternative by projecting the data onto components that capture the highest variance. Both methods help reduce dimensionality while retaining the most informative structure in the data.

5.4 | FEATURE TYPES: DIRECT, DIFFERENTIAL AND TEMPORAL

The extracted feature vectors are represented in the feature space in three ways: direct, differential, and temporal.

- Direct features represent each currency independently, using the extracted values.
- Differential features express the relative impact between the two currencies in a pair by calculating differences across corresponding indicators.
- Temporal features capture the position in time, inspired by Transformer-style positional encoding,[11] allowing the model to detect potential seasonality within the input sequence.

6 | Classification and Bayesian hyper-parameter optimisation

The extracted features and labels are used to train an SVM classifier. SVMs are particularly effective for imbalanced and outlier-prone classification tasks. With strong generalisation ability and significantly fewer parameters than deep neural networks, they are well-suited to high-dimensional problems with limited data. Non-linear class boundaries are handled using kernel functions that project the data into higher-dimensional spaces.

In addition to trainable parameters (support vectors, weights, and bias), SVMs involve several hyperparameters, such as the regularisation term (C), kernel type, and kernel coefficient (γ). These are optimised during the cross-validation phase.

We use a Bayesian optimisation approach to tune these hyperparameters.[12] The tuner begins by randomly sampling a few configurations and training the SVM on each. A Gaussian Process (GP) model is then fitted to the results. At each step, the tuner computes the Expected Improvement (EI) over the best result so far, proposes the next hyperparameter set by maximising the EI, evaluates the SVM, and updates the GP model. After a fixed number of trials, the configuration with the highest average balanced validation accuracy is selected. This is illustrated in Figure 5.

We apply a k-fold cross-validation scheme tailored for time series data, using a rolling approach to prevent information leakage and forward-looking bias.

FIGURE 5: **BAYESIAN HYPER-PARAMETER OPTIMISATION STEPS**

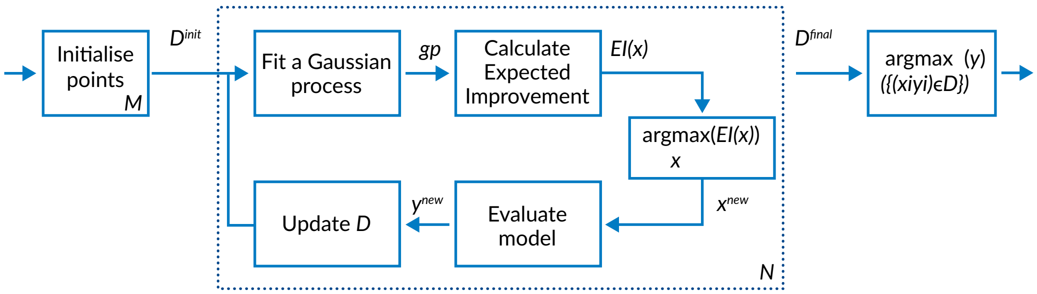

Bayesian hyper-parameter optimisation steps in a block diagram. First the points are randomly initialised over the search space (*D*) and a GP with a Matern kernel is fit on them, which is then used within the EI calculation. Those parameters maximising EI are then used to re-evaluate the model and our current collection of points (*D*) are updated. This process it repeated for a given number of iterations. The parameters corresponding to the highest objective function *y* are used as the selected hyper-parameter. This process is repeated as part of the rolling *k*-fold cross validation.

7 | Feature selection

The feature space described in Section 5 can be extremely high-dimensional relative to the number of samples. This arises from the use of numerous economic indicators, multiple feature extraction techniques, and a variety of feature types. Such high dimensionality increases the risk of overfitting, due to the well-known curse of dimensionality.[13]

To address this, we apply feature selection using a GA. GAs are well-suited for navigating large, complex search spaces, making them ideal for identifying optimal feature subsets in high-dimensional settings, where exhaustive search would be computationally infeasible.

GA is a population-based optimisation method inspired by natural selection. In our context, each individual in the population is a binary vector representing a subset of features — 1 indicates inclusion, 0 indicates exclusion. The length of each vector equals the total number of features, as illustrated in Figure 6.

The process follows standard GA steps:

- **Fitness Evaluation:** For each individual, an SVM model is trained, validated, and tested, using Bayesian hyperparameter tuning (as described in Section 6). The model's average balanced accuracy serves as the fitness score.
- **Selection:** Individuals with higher fitness scores are more likely to be chosen as parents for the next generation.
- **Crossover and Mutation:** Offspring are generated by combining parental vectors (crossover) and introducing occasional random changes (mutation) to encourage diversity and avoid local optima.
- **Replacement:** The offspring replace the existing population, and the process repeats for a fixed number of generations.

Key GA parameters include the population size, which controls solution diversity, and the number of generations, which determines the search depth. This iterative optimisation progressively enhances feature selection by favouring combinations that improve predictive performance.

FIGURE 6: **GA FEATURE SELECTION PROCEDURE**

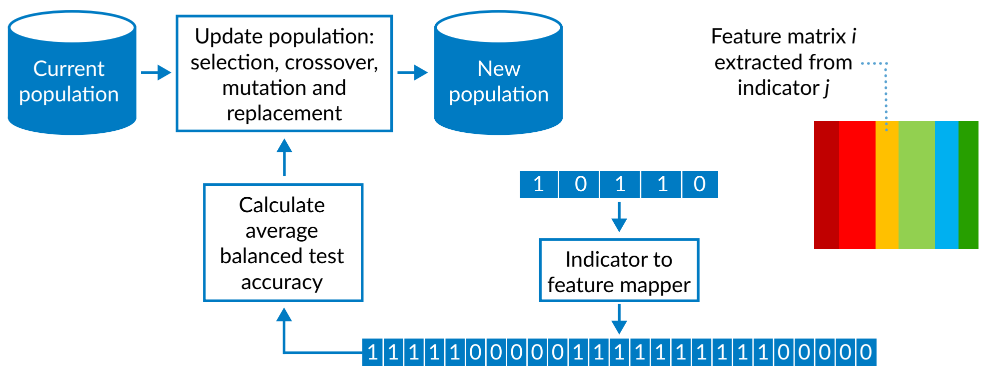

8 | Experimental results and performance

The proposed strategy is evaluated using daily FX spot rates for 27 USD-based currency pairs, covering both EM and DM countries. The rolling training period begins in February 2011, with an initial training window of five years, setting the testing period to start in February 2016. The data from February 2016 to the end of 2022 forms the testing phase, during which feature selection, hyperparameter tuning, and other model adjustments are performed. The period from January 2023 to December 2024 serves as the out-of-sample evaluation, where the model is held fixed and assessed without further tuning. From January 2025 to mid-April 2025, the model enters a paper trading phase, simulating live deployment without affecting actual portfolios. The strategy was deployed live from mid-April 2025. Throughout all periods, the model is explicitly re-tuned every quarter.

8.1 | FEATURE SELECTION RESULTS

Figure 7 illustrates the evolution of the fitness function across generations of the GA for four currency pairs: USDCAD, USDCHF, USDINR, and USDMXN. These optimisations run for 36 generations each with a population size of 28.

For each individual in every generation, Bayesian hyperparameter tuning is performed using rolling cross-validation as described in Section 6. This process is computationally intensive. As a result, increasing the number of generations arbitrarily is not feasible.

To balance this constraint, we prioritise exploration over exploitation by setting relatively high probabilities for mutation (0.2) and crossover (0.7). Additionally, to retain the best-performing individuals throughout the run, we use a hall of fame mechanism with an elite size of 3.

As the optimisation progresses, the GA identifies individuals that yield higher average balanced accuracies. To visualise the impact of feature selection, we use *t*-distributed stochastic neighbour embedding (*t*-SNE), a stochastic technique for projecting high-dimensional data into lower dimensions.

FIGURE 7: **FITNESS FUNCTION EVOLUTION**

(a) USDCAD

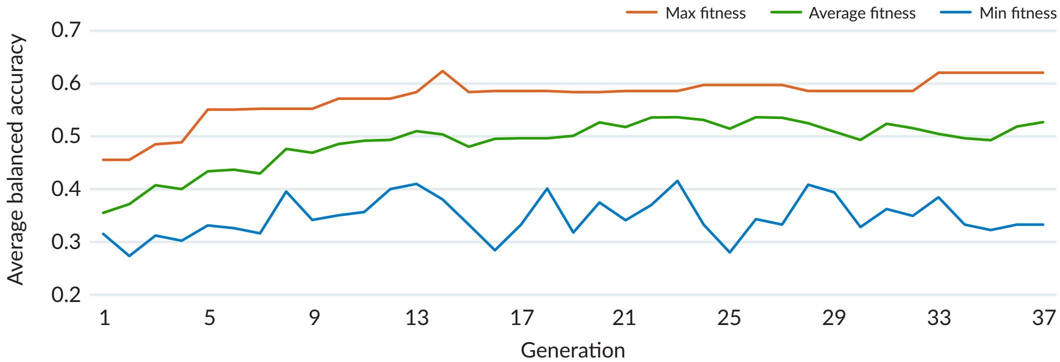

(b) USDCHF

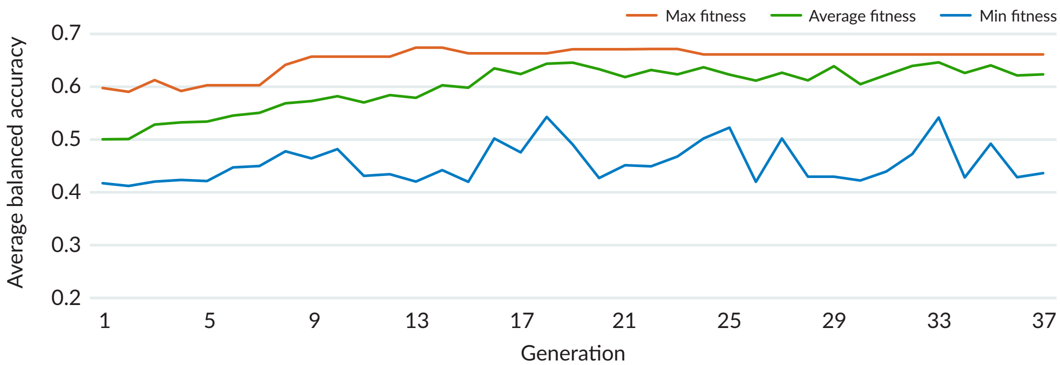

(c) USDINR

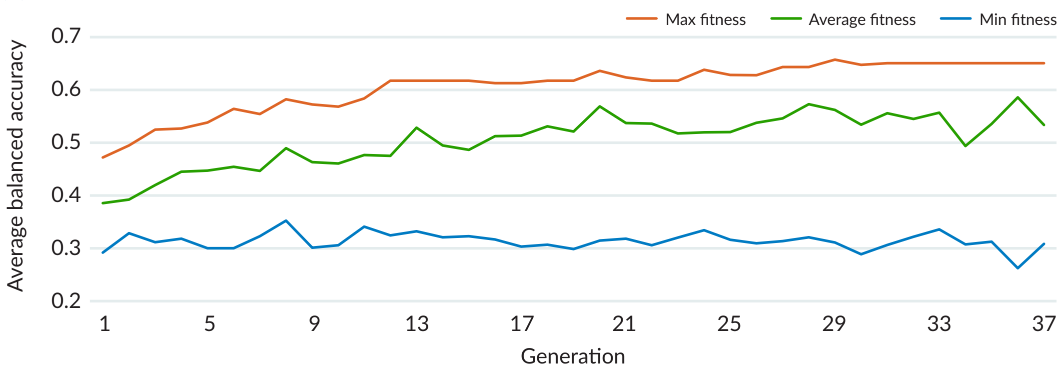

(d) USDMXN

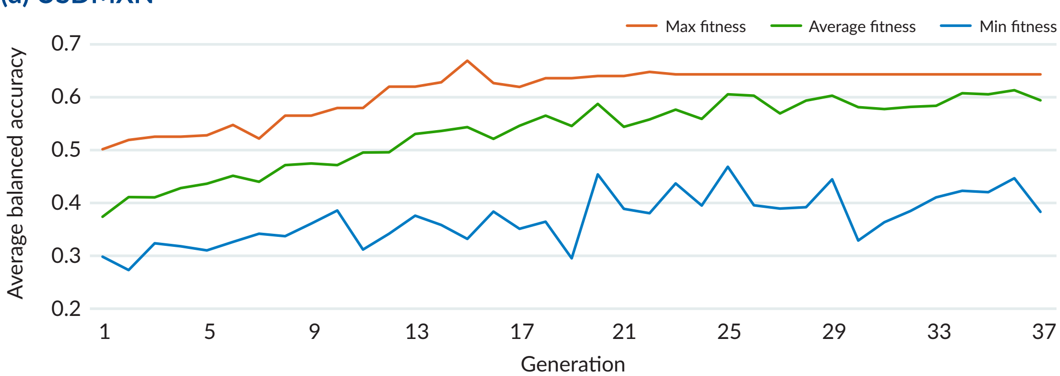

Figure 7 GA evolution for four example currency pairs: (a) USDCAD, (b) USDCHF, (c) USDINR and (d) USDMXN.

Figure 8 -a and -b present *t*-SNE visualisations of the feature space before and after GA-based feature selection for USDMXN and USDCAD, respectively. Although *t*-SNE is an unsupervised method, we display the samples in different colours based on their labels for visual clarity.

These figures illustrate that feature selection leads to better class separation. The classes appear more concentrated and distinct, indicating reduced within-class variance and increased between-class separation. This suggests that the selected features contribute more strongly to the structure relevant to classification.

FIGURE 8A: **FEATURE SPACE VISUALISATION BEFORE AND AFTER – USDMXN**

(a) Before

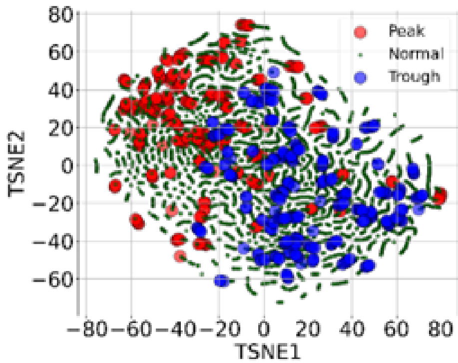

(a) After

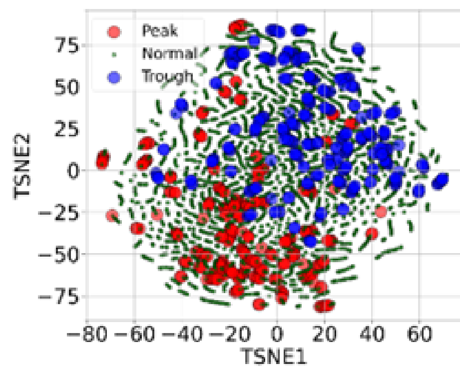

FIGURE 8B: **FEATURE SPACE VISUALISATION BEFORE AND AFTER – USDCAD**

(a) Before

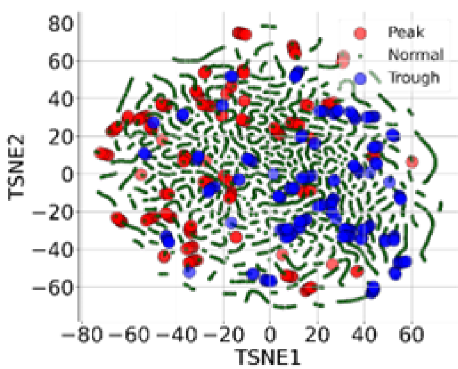

(a) After

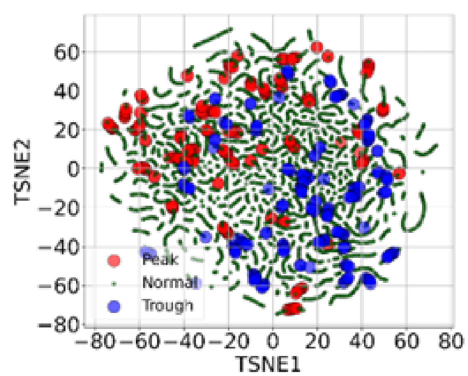

Feature space visualisation before (left column) and after (right column) applying GA using *t*-SNE for two example currency pairs:
(a) USDMXN and
(b) USDCAD

Figure 9 shows the average balanced accuracy for each currency pair after GA-based feature selection. As this is a ternary classification task, the baseline balanced accuracy (random guessing) is 33%. It is encouraging to observe that all currency pairs perform well above this baseline, with an overall average of approximately 58%. While the lowest-performing pairs are mainly from EM, the top-performing group includes a mix of EM and DM currencies, indicating that the model generalises well across different market types.

FIGURE 9: **AVERAGE BALANCED ACCURACY**

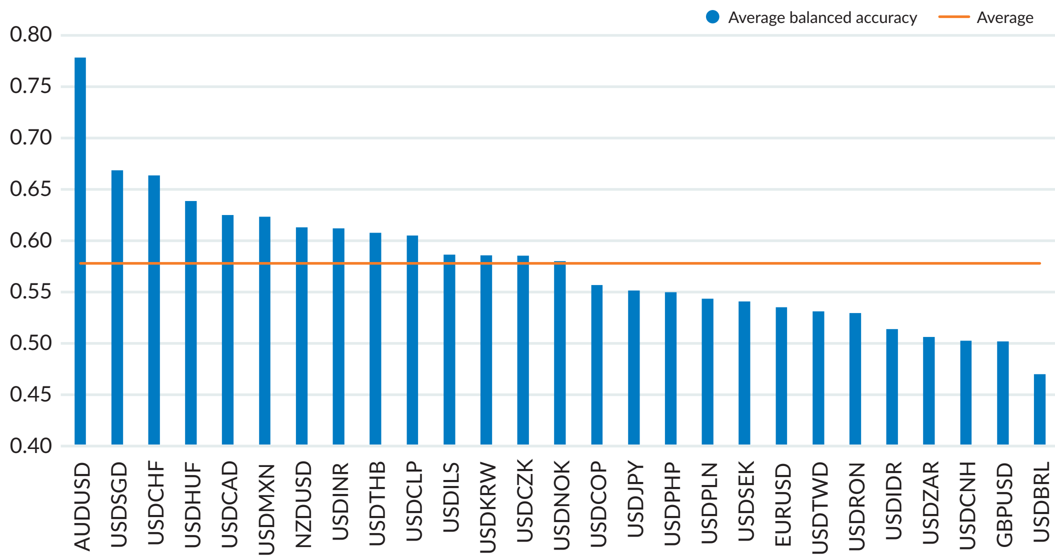

Averaged balanced accuracy over all currencies in our basket – an illustration of LEAP's performance.

Figure 10 presents a typical signal generated by LEAP, where $s_t^{(i)}$ denotes the MR signal produced at t for the i-th currency pair. Alongside the signal itself, it is valuable to have an indication of the model's confidence in its prediction. As SVM are not inherently probabilistic and do not directly provide confidence measures, we employ Platt scaling[14] to transform the decision function outputs into probability estimates. These probabilities are stored for each signal as $p_t^{(i)}$, corresponding to currency pair i at time t.

FIGURE 10: **SIGNAL GENERATED BY LEAP**

Signal_date	CCY_pair	Signal	Probability
YYYY-MM-DD	XXXYYY Curncy	$s_t^{(i)}$	$p_t^{(i)}$

A typical LEAP signal at time t for currency pair I. $s_t^{(i)}$ and $p_t^{(i)}$ are the signal and its corresponding confidence, respectively

As mentioned earlier, the period from January 2025 to mid-April 2025 serves as our paper trading phase. One way to assess whether the returns during this period are consistent with those observed in the in-sample and out-of-sample phases is through a quantile–quantile (Q–Q) plot. Figure 11 displays the Q–Q plot, which shows that the returns exhibit only minor skewness from the identity line ($y = x$), suggesting strong consistency with the historical distribution.

FIGURE 11: **Q-Q PLOT OF THE PAPER TRADING VS THE TESTING (IN-SAMPLE) PERIOD**

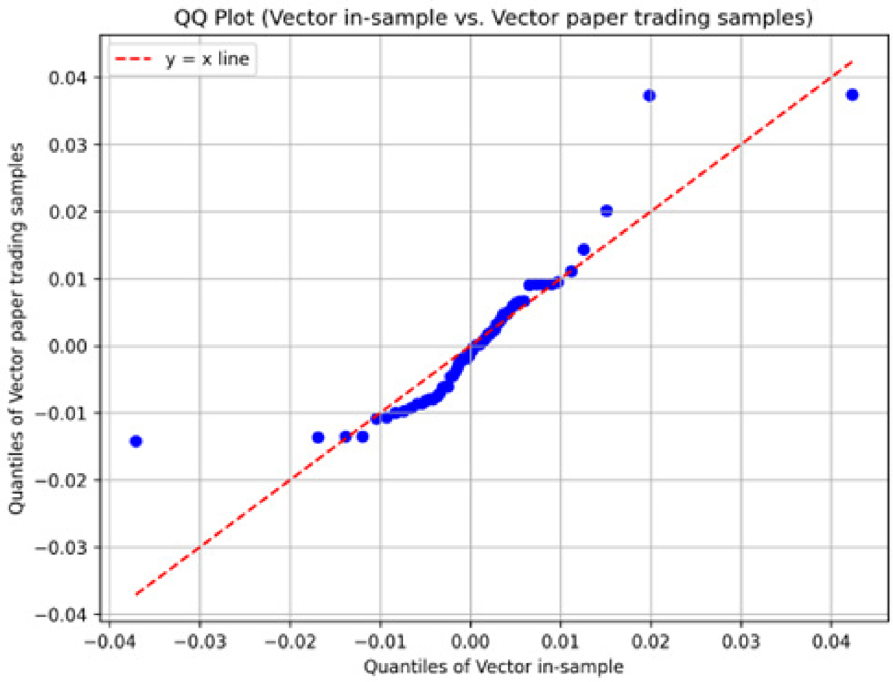

LEAP makes a notable contribution to performance when integrated into Mesirow Currency's Alpha framework, as summarised in Table 1 and Table 2. Four currency baskets are evaluated: EM (18 emerging market currencies), Asia (11 Asian currencies), DM (9 developed market currencies), and Extended (all 27 currencies), all quoted against USD.

Table 1 reports performance with LEAP as a standalone strategy. The out-of-sample IR are 1.38 (Extended), 1.12 (DM), 1.04 (Asia), and 1.36 (EM).

During this period, the US dollar exhibited stronger mean-reverting behaviour (see Figure 12, which overlays the DXY index on daily cumulative returns for each basket). LEAP effectively captured these reversals, contributing to its strong annualised returns.

In contrast, the in-sample period featured a more trending dollar, with fewer reversals. This structural difference explains the performance gap between the in-sample and out-of-sample phases.

Table 2 shows the performance of LEAP when integrated into Mesirow Currency's Alpha strategies: Extended, Asia, and Emerging. In all three cases, LEAP contributes positively to the overall IR. As illustrated in Figure 13, LEAP also demonstrates relatively low correlation with other strategies. Consistent with expectations for an MR strategy, it exhibits negative correlation with trend-following strategies and positive correlation with other MR strategies, such as Mesirow Currency's Short-Term Reversion and Momentum models.

TABLE 1: **LEAP ONLY BACKTEST, DURING TESTING AND OUT-OF-SAMPLE PERIODS**

LEAP only back tests	Annualised returns (%)	IR	Annualised returns (%) – OoS period	IR – OoS period
Extended	6.56	0.55	16.18	1.38
Only DM	6.77	0.54	14.20	1.12
Only Asia	4.74	0.37	13.15	1.04
Only EM	4.52	0.38	15.75	1.36

FIGURE 12: **CUMULATIVE RETURNS OF LEAP OVER USD-BASED CURRENCY PAIRS**

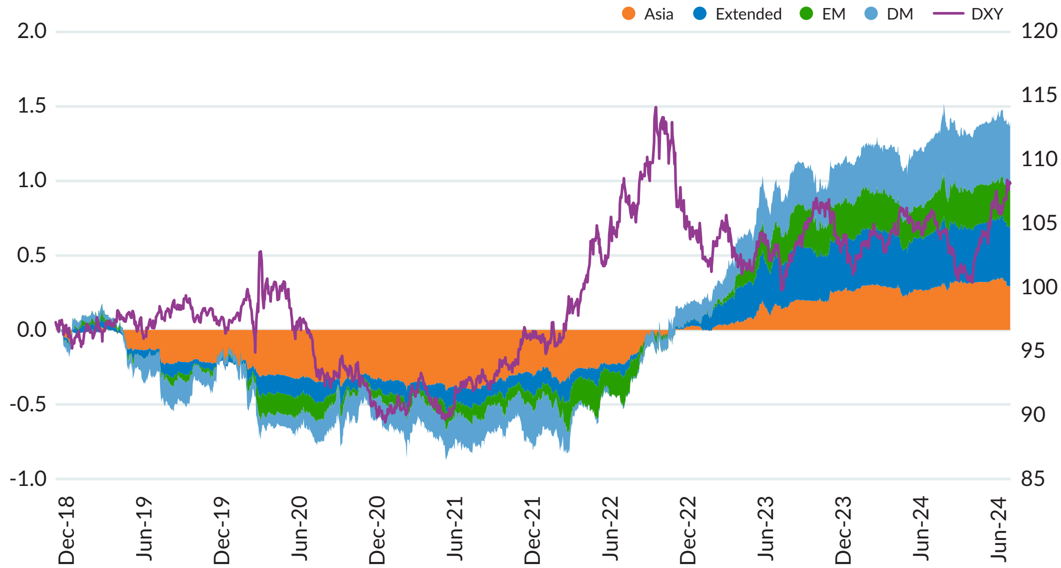

FIGURE 13: **CORRELATION OF LEAP V. OTHER ALPHA MODEL**

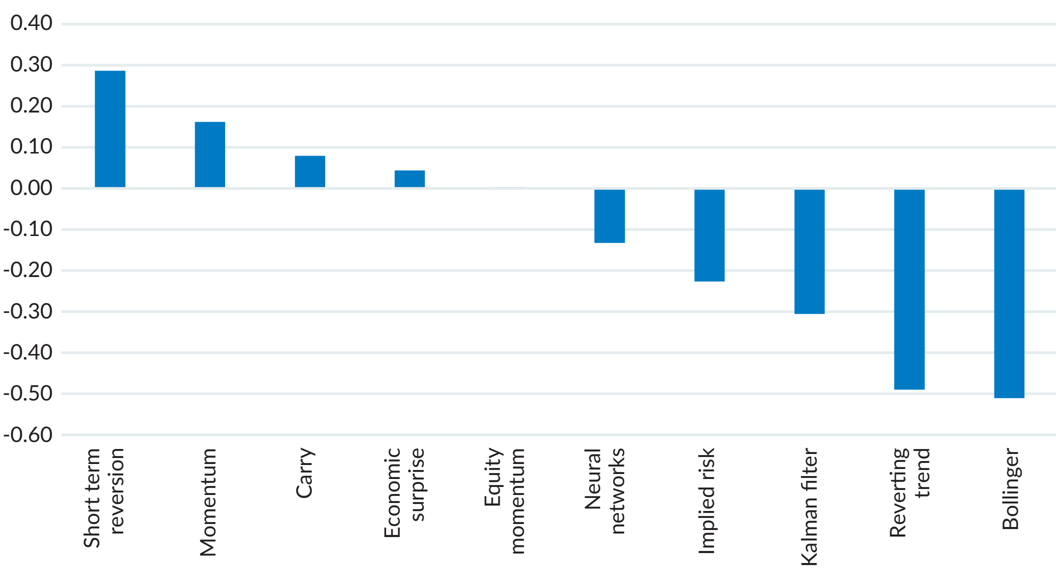

Figure 13 Correlation of LEAP against some of the other Alpha models, assuming the Extended currency basket.

TABLE 2: **STRATEGY PERFORMANCE WHEN LEAP IS ADDED**

IR increase

Strategy specific back test	2019 to 2024	OoS period
Extended strategy	+0.14	+0.32
Asia strategy	+0.12	+0.35
Emerging strategy	+0.05	+0.20

8.2 | DEEPER ANALYSIS OF GA RESULTS

In this section, we delve deeper into the feature selection results to understand which types of features were most frequently selected for each currency pair or basket. One of the key advantages of using GA for feature selection — compared to dimensionality reduction methods like PCA or LDA — is the explainability of the selected inputs.

Excluding indicators based on cross-currency and USD-quoted pairs, the remaining input indicators can be grouped into the following categories:

- **Vol:** Implied volatility indicators for the target currency pair
- **Global indicators:** Measures of global financial risk
- **Equity markets:** Stock market performance in the target country
- **Interest rates:** Overnight index swaps and 2-/10-year interest rates for the target currency
- **Differentials:** Economic indicator differences between the base and quoted currencies
- **Time features:** Encoded timestamps, independent of indicator values, used to capture potential seasonality
- **CDS:** Sovereign credit default swap spreads for the target currency

Figure 14 shows the frequency with which each feature category was selected by the GA for the EM, DM, and Asian currency baskets, respectively. Across all three baskets, Differentials and Global Indicators emerge as the most frequently selected features. This highlights the importance of geopolitical context and cross-country comparisons in identifying MR in FX. Interestingly, Equity Markets and Interest Rates are the least selected categories. This may suggest that these indicators lag in providing relevant signals for currency-level MR.

Another notable finding is the consistent selection of Time Features across all baskets, indicating potential seasonal patterns in MR behaviour.

FIGURE 14: **FEATURE TYPE FREQUENCY**

(a) EM currency basket

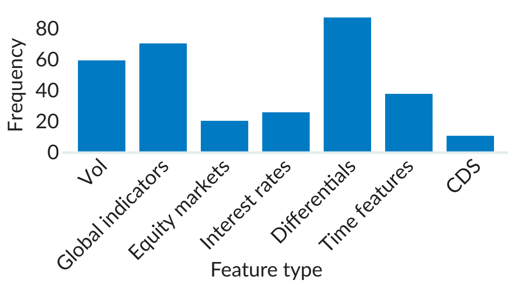

(b) DM currency basket

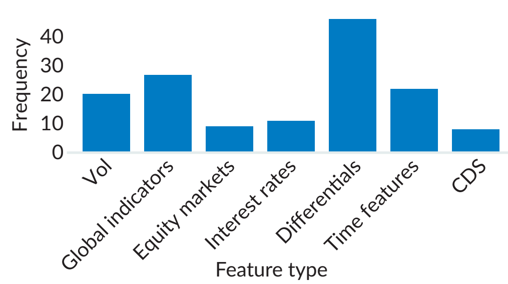

(c) Asia currency basket

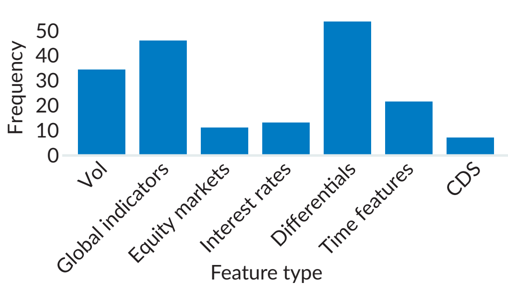

Figure 14 The frequency of each feature type selected by GA over the: (a) EM, (b) DM and (c) Asia currency baskets.

Another important input to LEAP comes from cross-currency pairs calculated for each target currency. For example, when predicting reversions in USDCAD, LEAP also computes all CAD-based crosses against the other 26 currencies in the basket — e.g. CADGBP, CADBRL, CADCHF, CADTWD, and so on. These cross rates are treated as separate feature subsets, and the GA selects the most informative ones.

Once this process is completed across all 27 USD-quoted currency pairs, we can analyse which currencies are most frequently referenced when constructing cross-currency features. These relationships can be visualised as directed graphs, where each node represents a currency, and edges indicate that a currency (child node) was used to generate a cross rate for the target currency (its parent node).

This analysis was performed across all four currency baskets: DM, Asia, EM, and Extended, as shown in Figure 15-a, -b, -c, and -d, respectively.

In the DM basket, EUR is the most frequently referenced currency, while CHF is the least. This may reflect EUR's greater transparency, higher liquidity, and lower likelihood of central bank intervention, making it more informative for MR detection. And therefore, other DM currencies tend to use it more often to construct cross pairs.

Within the Asian basket, AUD is most often used in cross pairs, suggesting its prominence in capturing regional FX dynamics.

For EM currencies, CNH stands out as the most frequently referenced — an intriguing result, given that CNH is a heavily managed currency. This may suggest that the patterns in its managed behaviour still convey useful signals, particularly around global sentiment shifts and reversion dynamics.

9 | Conclusions

This paper has introduced LEAP, a novel strategy for predicting MR in FX markets. LEAP addresses several limitations inherent in existing MR approaches, including reliance on subjective labelling, fixed modelling assumptions, lack of adaptability to the highly dynamic FX environment, high data dimensionality, and the use of narrow feature sets. By incorporating a wide range of FX and non-FX data sources, LEAP frames MR prediction as a ternary supervised classification task, categorising each training instance as a peak, trough, or normal based on the topographic characteristics of the target series.

To construct the feature space, a comprehensive software package has been developed, supporting not only feature extraction but also pre- and post-processing of the data. Model training is conducted through rolling *k*-fold cross-validation, with SVM hyperparameters optimised via Bayesian techniques. To ensure adaptive learning and reduce dimensionality, a GA-based feature selection mechanism is employed as a wrapper around the entire pipeline.

Although LEAP has been evaluated on USD-denominated FX rates, the framework is readily generalisable to other currency crosses and can operate at arbitrary time resolutions. Furthermore, the underlying philosophy of LEAP lends itself well to applications beyond FX, potentially extending to other classes of financial instruments.

LEAP has been incorporated into Mesirow Currency's Alpha suite, deployed in live trading since mid-April 2025.

FIGURE 15: **TOTAL NUMBER OF TIMES EACH CURRENCY IS REFERRED TO**

(a) DM

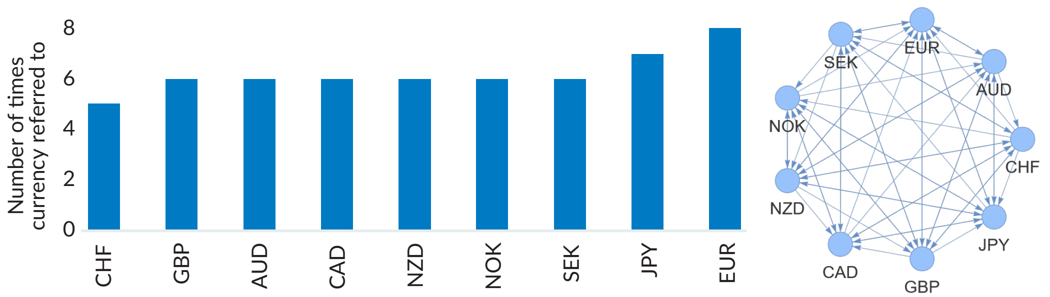

(b) Asia

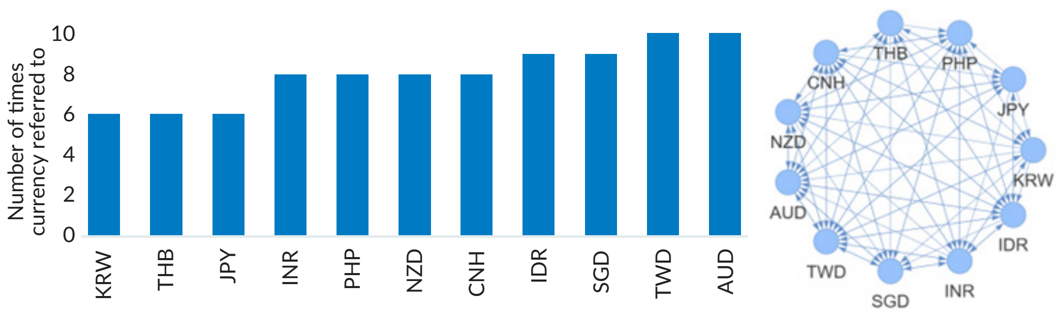

(c) EM

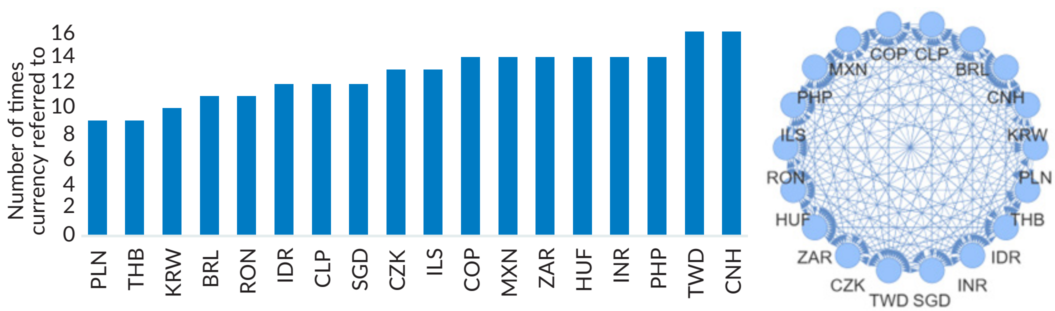

(d) Extended

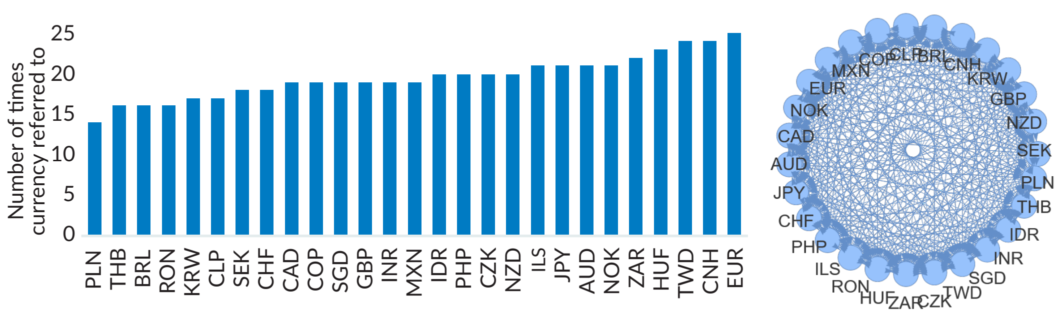

Figure 15 Total number of times each currency is referred to, to construct a cross-rate in GA's outputs among: (a) DM, (b) Asia, (c) EM and (d) Extended currency baskets.

About Mesirow

Mesirow is an independent, employee-owned financial services firm founded in 1937. Headquartered in Chicago, with offices around the world, we serve clients through a personal, custom approach to reaching financial goals and acting as a force for social good. With capabilities spanning Private Capital & Currency, Capital Markets & Investment Banking, and Advisory Services, we invest in what matters: our clients, our communities and our culture. To learn more, visit mesirow.com, follow us on LinkedIn and subscribe to Spark, our quarterly newsletter.

1. T. G. Andersen and O. Bondarenko, "Construction and interpretation of model-free implied volatility," 2007. | 2. E. Chan, Discovering risk indicators in the FX markets, White paper, 2012. | 3. K. V. Chalupa, "Foreign currency futures: reducing foreign exchange risk," Economic Perspectives, vol. 6, no. Win, pp. 3–11, 1982. | 4. Citigroup Global Markets Limited, Citi Risk Aversion Indicator: Index methodology, Bloomberg CIISRAI <Index>, 2015. | 5. Citigroup Global Markets Limited, Early Warning Signal: CitiFX hedging tool for emerging markets, Market Commentary, 30 June 2011. | 6. Goldman Sachs Global Investment Research, Disentangling risk appetite: Drivers of mini-cycles and asset allocation implications, Global Strategy Paper No. 33, Apr. 4, 2019. | 7. R. K. Lyons, The microstructure approach to exchange rates, vol. 333, Cambridge, MA: MIT Press, 2001. | 8. S. Y. Novak, Extreme value methods with applications to finance, Monographs on Statistics and Applied Probability, vol. 122, p. 22, 2011. | 9. B. Uhlenbrock, "Financial markets' appetite for risk – and the challenge of assessing its evolution by risk appetite indicators," in IFC Bulletin No. 31, Bank for International Settlements, 2009, pp. 221–226. | 10. Topographic prominence, Wikipedia. [Online]. Available: https://en.wikipedia.org/wiki/Topographic_prominence. [Accessed: Jun. 2, 2025, 16:17 BST]. | 11. A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in Advances in Neural Information Processing Systems, vol. 30, 2017. [Online]. Available: https://arxiv.org/abs/1706.03762 | 12. J. Snoek, H. Larochelle, and R. P. Adams, "Practical Bayesian Optimization of Machine Learning Algorithms," in Advances in Neural Information Processing Systems, vol. 25, 2012. [Online]. Available: https://arxiv.org/abs/1206.2944 | 13. R. O. Duda, P. E. Hart, and D. G. Stork, Pattern Classification, 2nd ed. Hoboken, NJ, USA: Wiley-Interscience, 2001. | 14. J. Platt, "Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods," Advances in Large Margin Classifiers, pp. 61–74, 1999.

Past performance is not an indication of future results. Actual results may be materially different from the results achieved historically.

Nothing herein is meant to be taken as a recommendation to buy or sell a particular asset or invest in a particular strategy.

Mesirow Currency ("MC") is a division of Mesirow Financial Investment Management, Inc. ("MFIM") a SEC registered investment advisor, a CFTC registered commodity trading advisor and a member of the NFA and Mesirow Financial International UK, Ltd. ("MFIUK"), authorized and regulated by the FCA. The information contained herein is intended for institutional clients, Qualified Eligible Persons and Eligible Contract Participants or the equivalent classification in the recipient's jurisdiction and is for informational purposes only. This information has been obtained from sources believed to be reliable, but is not necessarily complete and its accuracy cannot be guaranteed. Any opinions expressed are subject to change without notice. It should not be assumed that any recommendations incorporated herein will be profitable or will equal past performance. Mesirow does not render tax or legal advice. Nothing contained herein constitutes an offer to sell or a solicitation of an offer to buy an interest in any Mesirow investment vehicle(s). Any offer can only be made through the appropriate Offering Memorandum. The Memorandum contains important information concerning risk factors and other material aspects of theinvestment and should be read carefully before an investment decision is made.

Australian Investors: The information contained herein is intended for Wholesale Clients only and is for informational purposes only. This document is not a prospectus or product disclosure statement under the Corporations Act 2001 (Cth) (Corporations Act) and does not constitute a recommendation to acquire, an invitation to apply for, an offer to apply for or buy, an offer to arrange the issue or sale of, or an offer for issue or sale of, any securities or investment service in Australia, except as set out below. The strategy has not authorised nor taken any action to prepare or lodge with the Australian Securities & Investments Commission an Australian law compliant prospectus or product disclosure statement. Accordingly, this strategy and document may not be issued or distributed in Australia other than by way of or pursuant to an offer or invitation that does not need disclosure to investors under Part 6D.2 or Part 7.9 of the Corporations Act, whether by reason of the investor being a 'wholesale client' (as defined in section 761G of the Corporations Act and applicable regulations) or otherwise. This document does not constitute or involve a recommendation to acquire, an offer or invitation for issue or sale, an offer or invitation to arrange the issue or sale, or an issue or sale, of any strategy or investment service to a 'retail client' (as defined in section 761G of the Corporations Act and applicable regulations) in Australia.

Canadian Investors: The information contained herein is intended for Permitted Clients only and is for informational purposes only. This confidential material pertains to the offering of the currency strategies described herein only in those jurisdictions and to those persons where and to whom they may be lawfully offered for sale, and only by persons permitted to sell such strategies. This material is not, and under no circumstances is to be construed as, an advertisement or a public offering of the strategies described herein in Canada. No securities commission or similar authority in Canada has reviewed or in any way passed upon this document or the merits of the strategies described herein, and any representation to the contrary is an offence.

Netherlands and Luxembourg Investors: The information contained herein is intended for Professional Clients as the term is defined by MiFID II and is for informational purposes only. Recipients that are classified under MiFID II as retail clients must opt up to Professional Clients before receiving any services from Mesirow Currency.

Japanese Investors: Mesirow Currency provides discretionary investment management services to managed accounts held on behalf of qualified investors only. MC will not act as agent or intermediary in respect of the execution of a discretionary investment management agreement. Please note that this presentation is intended for educational purposes and solely for the addressee and may not be distributed.

Hong Kong Investors: The contents of this document have not been reviewed by any regulatory authority in Hong Kong. You are advised to exercise caution in relation to the contents of this document. You should obtain independent professional advice prior to considering or making any investment. The investment is not authorized under Section 104 of the Securities and Futures Ordinance of Hong Kong by the Securities and Futures Commission of Hong Kong. Accordingly, the distribution of this Presentation Material and discretionary management services in Hong Kong are restricted. This Presentation Material is only for the use of the addressee and may not be distributed, circulated or issued to any other person or entity.

South Korean Investors: Upon attaining a client, Mesirow Financial Investment Management, Inc. ("MFIM") will apply for the appropriate licenses and retain the services of a local licensed intermediary (a Korean financial investment company). In the interim, MFIM will rely on and sub-delegate to Mesirow Advanced Strategies, Inc. ("MAS").

This material is not intended for investment nor distribution purposes, but is being provided to the addressee for educational purposes only. This book may contain privileged and/or confidential information and is intended solely for the use of the addressee. If this information was received in error, you are strictly prohibited from disclosing, copying, distributing or using any of this information and are requested to contact the sender immediately and destroy the material in its entirety, whether electronic or hardcopy. Nothing contained herein constitutes an offer to sell or a solicitation of an offer to buy an interest in any Mesirow investment vehicle or to invest in any strategy. Any offer can only be made to Qualified Professional Investors through the appropriate Investment Management Agreement or Offering Memorandum, which contains important information concerning risk factors and other material aspects of the investment and should be read carefully before an investment decision is made.

Mesirow Financial Investment Management, Inc. ("MFIM") is not making any representation with respect to the eligibility of any potential investors or recipients of this material to invest in or acquire any interests therein under the laws of Korea, including but without limitation the Foreign Exchange Transaction Act and Regulations thereunder. The investment or interests may only be offered to Qualified Professional Investors, as such term is defined under the Financial investment Services and Capital Markets Act, and no investment or any of the interests may be offered, sold or delivered, or offered or sold to any person for re-offering or resale, directly or indirectly, in Korea or to any resident of Korea except pursuant to applicable laws and regulations of Korea.

Singapore Investors: Mesirow Currency provides discretionary investment management services to managed accounts held on behalf of qualified investors only. MC will not act as agent or intermediary in respect of the execution of a discretionary investment management agreement. Please note that this presentation is intended for educational purposes and solely for the addressee and may not be distributed.

Swiss Investors: The information contained herein is intended for Professional Clients as the term is defined by MiFID II and is for informational purposes only. Recipients that are classified under MiFID II as retail clients must opt up to Professional Clients before receiving any services from Mesirow. Services are only offered to Regulated Qualified Investors, as defined in Article 10 of the Swiss Collective Investment Scheme Act. There can be no guarantee investment advice will be profitable or meet its investment objectives.

United Kingdom Investors: In the United Kingdom, this communication is directed only at persons who fall within the definition of a "professional client" as defined in COBS 3.5 of the FCA handbook.

Additional Information: Currency strategies are only suitable and appropriate for sophisticated investors that are able to lose all of their capital investment. This communication may contain privileged and/or confidential information. It is intended solely for the use of the addressee. If this information was received in error, you are strictly prohibited from disclosing, copying, distributing or using any of this information and are requested to contact the sender immediately and destroy the material in its entirety, whether electronic or hardcopy.

Certain strategies discussed throughout the document are based on proprietary models of MC's or its affiliates. No representation is being made that any account will or is likely to achieve profits or losses similar to those referenced.

Performance pertaining to the Currency Alpha and Macro strategies may be stated gross of fees or net of fees. Performance information that is provided net of fees reflects the deduction of implied management and performance fees. Performance information that is provided gross of fees does not reflect the deduction of advisory fees. Client returns will be reduced by such fees and other expenses that may be incurred in the management of the account. Advisory fees are described in Part II of Form ADV of Mesirow Financial Investment Management, Inc. For any performance information that is provided gross of fees, MFIM will provide performance results net of such fees upon request. Hypothetical model performance information and results do not reflect actual trading or asset or fund advisory management and the results may not reflect the impact that material economic and market factors may have had. Hypothetical performance results have many inherent limitations, some of which are described herein. No representation is being made that any account will or is likely to achieve profits or losses similar to those shown. In fact, there are frequently sharp differences between hypothetical performance results and the actual results subsequently achieved by any particular trading program. One of the limitations of hypothetical performance results is that they are generally prepared with the benefit of hindsight. In addition, hypothetical trading does not involve financial risk, and no hypothetical trading record can completely account for the impact of financial risk in actual trading. For example, the ability to withstand losses or adhere to a particular trading program in spite of trading losses are material points which can also adversely affect actual trading results. There are numerous other factors related to the markets in general or to the implementation of any specific trading program which cannot be fully accounted for in the preparation of hypothetical performance results and all of which can adversely affect actual trading results.

Comparisons to any indices referenced herein are for illustrative purposes only and are not meant to imply that a strategy's returns or volatility will be similar to the indices. The strategy is compared to the indices because they are widely used performance benchmarks.

Mesirow refers to Mesirow Financial Holdings, Inc. and its divisions, subsidiaries and affiliates. The Mesirow name and logo are registered service marks of Mesirow Financial Holdings, Inc., © 2025, Mesirow Financial Holdings, Inc. All rights reserved. Investment management services provided through Mesirow Financial Investment Management, Inc., a SEC registered investment advisor, a CFTC registered commodity trading advisor and member of the NFA, or Mesirow Financial International UK, Ltd. ("MFIUK"), authorized and regulated by the FCA, depending on the jurisdiction.