

MINERVA INVESTMENT MANAGEMENT SOCIETY

Quantitative Research Division



Signature Methods in Finance

A Pathwise Approach to Financial Time Series Representation

Ianina Tarca

Sebastian Skjensvold

Filip Borzecki

Head of Quantitative Research Division

Quantitative Research Analyst

Quantitative Research Analyst

Academic Year 2025/2026

Contents

1	Introduction	3
2	Mathematical Framework	3
2.1	Setup: Riemann–Stieltjes integration and bounded variation	4
2.2	Signatures and iterated integrals	5
2.3	Lévy area and a non-cancelling area functional	7
3	Computational Considerations	8
4	Applications	8
5	Conclusion	9

Abstract

Financial time series exhibit strong path-dependent characteristics that are not fully captured by traditional summary statistics. Standard approaches often focus on endpoint values or low-order moments, neglecting the temporal ordering and interaction of increments that drive many market phenomena, such as volatility clustering, leverage effects, and microstructure dynamics.

This paper introduces the signature of a path as a systematic and mathematically grounded feature representation for financial data streams. By encoding a time series through iterated integrals into a purely algebraic object, signatures capture cumulative movements, second-order interactions related to volatility, and higher-order temporal dependencies. These features form a universal and non-parametric basis for approximating path-dependent functionals.

We develop the theoretical framework linking controlled differential equations, Picard iteration, and iterated integrals, and show how signatures naturally arise as the fundamental object summarizing path information. We further discuss computational aspects, including truncation and Chen's theorem, which enables efficient incremental updates in streaming and rolling-window settings.

Finally, we highlight the relevance of signature methods for quantitative finance applications, including volatility forecasting, path-dependent derivative pricing, and market microstructure modeling. The results suggest that signature-based representations provide a flexible bridge between continuous-time stochastic modeling and modern data-driven approaches.

DISCLAIMER

This paper is an academic work prepared exclusively for educational and research purposes. It does not constitute investment advice, financial recommendations, or an offer or solicitation to buy or sell any financial instruments. The opinions expressed herein are solely those of the authors and are subject to change without notice. Any references to specific securities, asset classes, or financial markets are for illustrative purposes only and should not be interpreted as endorsements. The authors and the Minerva Investment Management Society accept no liability for any reliance placed on the information contained in this document.

1 Introduction

Financial time-series data is often not well summarized by endpoint statistics: the final result matters, but how it was achieved is of equal importance. The inherent path-dependence of financial markets appears in many empirical phenomena. Volatility exhibits clustering effects - high (low) volatility days are followed by days with similar magnitudes - and leverage effects - negative shocks tend to have a larger impact on volatility than positive ones. Similar path-dependent characteristics arise in market microstructure, where two markets with similar mid-price can differ substantially in order flow, liquidity, or imbalance. These observations suggest that how increments of a process interact through time can be as important as their final distribution.

Signatures provide a principled way to represent paths. The signature's levels can be interpreted as progressively richer summaries of the underlying stream. The first level captures total increments (for a component representing log-price, this corresponds to cumulative log-returns). Second-order terms capture volatility-related structure and ordering effects between coordinates, while higher-order terms capture increasingly complex interactions that can be crucial in markets, where features are correlated and path interactions matter.

Theoretical results motivate this approach: for bounded-variation paths (and in appropriate extensions), signatures characterize paths up to tree-like equivalence, and linear functionals of truncated signatures provide a universal feature set for approximating a wide class of path functionals. These properties make signatures relevant for modeling path-dependent payoffs, mean-reversion-based trading signals, and volatility dynamics. In addition, Chen's theorem makes signatures computationally attractive in rolling-window settings, where features must be updated as new data arrive. This theorem states that the signature of a concatenation is the tensor product of the signatures. In financial applications this is useful whenever one wants to compute features incrementally - for example, in rolling windows. Rather than recomputing signature features from scratch at each step, one can update the aggregate representation by tensor multiplication (together with the appropriate bookkeeping for truncation). This aligns with real-time workflows in trading and risk, where streaming data must be processed continuously.

2 Mathematical Framework

This section develops the mathematical objects underlying signature methods and explains how iterated integrals naturally arise from controlled differential equations.

Let V, W be Banach spaces, $\mathcal{L}(V, W)$ the space of bounded linear maps $L : V \rightarrow W$, and let $J \subset \mathbb{R}$ be a compact interval (for convenience we typically take $J = [0, T]$ unless specified otherwise). For maps $X : J \rightarrow V$, $Y : J \rightarrow W$, and a sufficiently regular function $f : W \rightarrow \mathcal{L}(V, W)$ so that

$$\int_0^t f(Y_s) dX_s \tag{2.1}$$

is well-defined, any solution of the (controlled) differential equation

$$dY_t = f(Y_t) dX_t, \quad Y_0 = y_0 \in W \tag{2.2}$$

must satisfy the integral equation

$$Y_t = y_0 + \int_0^t f(Y_s) dX_s. \quad (2.3)$$

Define the Picard operator Φ by

$$(\Phi Y)_t = y_0 + \int_0^t f(Y_s) dX_s. \quad (2.4)$$

Solutions to (2.2) are precisely fixed points of Φ . Under suitable regularity assumptions (for instance, Lipschitz f and bounded-variation X) the operator becomes a contraction on an appropriate space (locally in time, and globally by patching), implying existence and uniqueness by Banach's fixed point theorem. In that setting, for any initial guess $\tilde{Y}$,

$$\lim_{n \rightarrow \infty} \Phi^{(n)} \tilde{Y} = Y,$$

where Y is the unique solution and $\Phi^{(n)}$ denotes n -fold iteration.

A key observation is that iterating (2.4) produces higher-order integral expressions in the increments of X . These *iterated integrals* provide progressively richer information about the geometry and temporal structure of the driving signal X . This becomes especially important in finance, where Brownian-like roughness makes classical Riemann–Stieltjes theory insufficient. Rough path theory extends integration to such irregular signals and identifies the *signature* of a path as the central object encoding this iterated-integral information.

2.1 Setup: Riemann–Stieltjes integration and bounded variation

We begin by making precise the integral in (2.1) in a classical setting.

Definition 2.1. Let $J = [0, T]$ be a compact interval. A *partition* P of J is a set of points

$$P = (0 = t_0, t_1, \dots, t_n = T)$$

with $t_i < t_{i+1}$ for all i .

A path $X : J \rightarrow V$ is said to have *bounded variation* if

$$\|X\|_{\text{var}; J} := \sup_{P \subset J} \sum_{k=0}^{n-1} \|X_{t_{k+1}} - X_{t_k}\| < \infty,$$

where $\|\cdot\|$ is the norm on V . In that case, the Riemann–Stieltjes integral is defined by the limit

$$\int_0^t f(Y_s) dX_s = \lim_{|P| \rightarrow 0} \sum_{k=0}^{n-1} f(Y_{t_k}) (X_{t_{k+1}} - X_{t_k}),$$

whenever the right-hand side converges, with $|P| = \max_k |t_{k+1} - t_k|$.

Theorem 2.1 (Picard–Lindelöf for BV drivers). *Let $f : W \rightarrow \mathcal{L}(V, W)$ be L -Lipschitz and let X have bounded variation on J . Then for every $y_0 \in W$ there exists a unique solution $Y : J \rightarrow W$ to (2.2).*

In contrast, for a Brownian motion B we have increments

$$B_{t+\delta} - B_t \sim \mathcal{N}(0, \delta),$$

so heuristically,

$$\sup_{P \subset J} \sum_{k=0}^{n-1} |B_{t_{k+1}} - B_{t_k}| \approx \lim_{n \rightarrow \infty} n \cdot \sqrt{\frac{1}{n}} = \infty,$$

and indeed Brownian motion has infinite variation almost surely. Classical Riemann–Stieltjes integration is therefore not equipped to handle Brownian roughness.

A more refined notion is p -variation.

Definition 2.2. The p -variation of a path $X : J \rightarrow V$ is

$$\|X\|_{p\text{-var}; J} = \left(\sup_P \sum_{k=0}^{n-1} \|X_{t_{k+1}} - X_{t_k}\|^p \right)^{1/p}.$$

Write $\mathcal{V}^p(J, V)$ for the space of paths with finite p -variation on J .

It can be shown that Brownian motion belongs to $\mathcal{V}^p(J, \mathbb{R}^d)$ for all $p > 2$ (and for no $p \leq 2$), motivating rough path integration and the need to enrich the signal X with iterated-integral data.

2.2 Signatures and iterated integrals

Return to the differential equation (2.2). In sufficiently regular settings (and in the rough-path extension), solutions admit expansions obtained from Picard iteration. To see the structure, let $V = \mathbb{R}^d$ with basis $\{e_1, \dots, e_d\}$ and write $X = \sum_i X^i e_i$. Define

$$f^i(y) := f(y)(e_i) \in W, \quad i = 1, \dots, d.$$

(Thus each $f^i : W \rightarrow W$ is a vector field in the Banach-space sense.)

Then (2.3) can be written as

$$Y_t = y_0 + \sum_{i=1}^d \int_0^t f^i(Y_s) dX_s^i.$$

Picard iteration then produces iterated integrals multiplied by *iterated directional derivatives* of the vector fields. To state this cleanly, define first-order differential operators $\mathcal{D}_i$ acting on C^1 maps $\phi : W \rightarrow W$ by

$$(\mathcal{D}_i \phi)(y) := D\phi(y)[f^i(y)],$$

where $D\phi(y) \in \mathcal{L}(W, W)$ denotes the Fréchet derivative. In particular, for the identity map $\text{Id}_W(y) = y$ we have

$$(\mathcal{D}_i \text{Id}_W)(y) = f^i(y).$$

With this notation, the formal iterated-integral expansion reads

$$Y_t = y_0 + \sum_{k=1}^{\infty} \sum_{i_1, \dots, i_k=1}^d (\mathcal{D}_{i_k} \cdots \mathcal{D}_{i_1} \text{Id}_W)(y_0) \int_{0 < t_1 < \dots < t_k < t} dX_{t_1}^{i_1} \cdots dX_{t_k}^{i_k}, \quad (2.5)$$

where each coefficient $(\mathcal{D}_{i_k} \cdots \mathcal{D}_{i_1} \text{Id}_W)(y_0) \in W$ is well-defined (under sufficient smoothness of f). For instance, the second-order coefficient expands as

$$(\mathcal{D}_i \mathcal{D}_j \text{Id}_W)(y_0) = Df^j(y_0)[f^i(y_0)] \in W,$$

so the level-two term is

$$\sum_{i,j=1}^d Df^j(y_0)[f^i(y_0)] \int_{0 < t_1 < t_2 < t} dX_{t_1}^i dX_{t_2}^j.$$

The iterated integrals

$$S(X)_{[0,t]}^{(i_1, \dots, i_k)} := \int_{0 < t_1 < \dots < t_k < t} dX_{t_1}^{i_1} \cdots dX_{t_k}^{i_k}$$

are the coordinates of the *signature* of the path X .

These multilinearity and noncommutativity of these objects make it natural to consider them as tensors.

Definition 2.3. The k th tensor power of a vector space V is

$$T^k(V) = \text{span}\{v_1 \otimes \cdots \otimes v_k \mid v_1, \dots, v_k \in V\},$$

where $\otimes$ is multilinear and distributive. The *tensor algebra* of V is

$$T(V) = \bigoplus_{k=0}^{\infty} T^k(V) = \mathbb{R} \oplus V \oplus (V \otimes V) \oplus (V^{\otimes 3}) \oplus \cdots.$$

Definition 2.4. The *signature* of a (sufficiently regular) path $X : J \rightarrow \mathbb{R}^d$ over $J = [0, T]$ is the element $S(X)_J \in T(\mathbb{R}^d)$ given by

$$S(X)_J = (1, S(X)_J^1, \dots, S(X)_J^d, S(X)_J^{1,1}, S(X)_J^{1,2}, \dots),$$

where for multi-indices $(i_1, \dots, i_k)$,

$$S(X)_J^{(i_1, \dots, i_k)} := \int_{0 < t_1 < \dots < t_k < T} dX_{t_1}^{i_1} \cdots dX_{t_k}^{i_k}.$$

Intuitively, the signature records *all* iterated-integral interactions among the coordinates of X , so it captures more than endpoints: it captures ordering and multi-scale geometry of the path. In applications one truncates $S(X)$ at some level m .

A crucial computational and algebraic property comes from concatenation.

Definition 2.5 (Concatenation (with time reparameterization)). Given two paths X, Y on $[0, T]$, define their *concatenation* $X * Y$ on $[0, T]$ by

$$(X * Y)(s) = \begin{cases} X(2s), & s \in [0, T/2], \\ Y(2s - T) + X(T) - Y(0), & s \in [T/2, T]. \end{cases}$$

This definition keeps the common time domain $[0, T]$ by reparameterizing each half; for paths of bounded variation the signature is invariant under increasing reparameterizations, so this is equivalent to the more standard concatenation on $[0, T_1 + T_2]$.

Theorem 2.2 (Chen). *For suitable paths X, Y (e.g. bounded variation, or in the rough path setting with enhancement), the signature is a homomorphism from concatenation to tensor multiplication:*

$$S(X * Y) = S(X) \otimes S(Y).$$

Computationally, if we partition J into sub-intervals and write $X = X^{(1)} * \dots * X^{(m)}$ as a concatenation of path segments, then Chen's theorem allows incremental computation of $S(X)$ by computing signatures on segments and tensoring them together (useful, for instance, for rolling windows).

2.3 Lévy area and a non-cancelling area functional

We now discuss geometric meaning at low levels and introduce a cancellation-robust area functional.

Let $X \in \mathcal{V}^p(J, \mathbb{R}^k)$. The first level records net displacement:

$$S(X)_J^i = \int_0^T dX_t^i = X_T^i - X_0^i.$$

Thus level one cannot distinguish paths that share endpoints but differ in how they move in between.

The second level begins to capture missing geometry. Its components are

$$S(X)_J^{i,j} = \int_{0 < t_1 < t_2 < T} dX_{t_1}^i dX_{t_2}^j,$$

and generally $S(X)_J^{i,j} \neq S(X)_J^{j,i}$. In the planar case $X = (X^1, X^2)$, the antisymmetric part has a direct geometric interpretation: the (signed) Lévy area

$$A_{[0,T]}(X) := \frac{1}{2} \left(S(X)_{[0,T]}^{1,2} - S(X)_{[0,T]}^{2,1} \right),$$

which equals the signed area enclosed by the curve together with the straight chord closing X_T back to X_0 (in the bounded-variation/smooth setting; in rough paths it is provided by the second-level enhancement).

The sign depends on orientation (clockwise vs. anticlockwise), so cancellations can occur: a path that loops in opposite directions over different periods can have small signed Lévy area even though it exhibits substantial rotational activity.

To quantify *how much* rotation occurs while removing directional cancellation, it is natural to consider the *area process*

$$A_t(X) := \frac{1}{2} \left(S(X)_{[0,t]}^{1,2} - S(X)_{[0,t]}^{2,1} \right), \quad t \in [0, T],$$

and then take its total variation.

Definition 2.6 (Non-cancelling (non-signed) Lévy area). Define the non-cancelling area functional by

$$A_{[0,T]}^{\text{ns}}(X) := \|A(\cdot; X)\|_{\text{var};[0,T]} = \sup_{P \subset [0,T]} \sum_{k=0}^{n-1} |A_{t_{k+1}}(X) - A_{t_k}(X)|.$$

This definition is intrinsic (partition-independent): it measures the total accumulated signed area *in absolute increments*, so opposite orientations contribute additively rather than cancelling.

In computation one may approximate A^{ns} by choosing a partition $0 = t_0 < \dots < t_m = T$ and summing

$$\sum_{\ell=1}^m |A_{t_\ell}(X) - A_{t_{\ell-1}}(X)|,$$

refining the partition until the estimate stabilizes. In a trading context, A^{ns} can be read as a robust “total rotational interaction” between two price processes (or two coordinates derived from them), which is precisely the type of path-dependent information signatures are designed to encode.

3 Computational Considerations

Since signatures are infinite objects, practical work focuses on truncated signatures (cutting off at level m). The number of coordinates grows exponentially with m , creating a trade-off between expressiveness and computational cost. Moreover, market data are discrete, so one must interpolate between observations to define iterated integrals. A standard approach is to linearly interpolate observations to obtain a piecewise-linear path. Iterated integrals can then be computed exactly for that interpolation. This is also the framework implemented in common software libraries (e.g. `iisignature` or `signatory`).

4 Applications

Signature methods provide a flexible, path-wise feature representation for financial time series, making them particularly suitable for settings where temporal ordering and higher-order interactions are essential.

In volatility modeling and forecasting, signatures can be used to encode rolling windows of returns into feature vectors that capture not only cumulative movements but also second-order and higher-order dependencies. This allows learning-based models to incorporate nonlinear and asymmetric effects that are difficult to capture with traditional parametric approaches such as GARCH-type models.

In the context of path-dependent derivatives, signatures offer a non-parametric framework for approximating payoffs that depend on the entire trajectory of an underlying asset. By expressing such payoffs as functions of truncated signature terms, one can construct flexible pricing and hedging schemes that do not rely on specific model assumptions about the underlying dynamics.

Finally, in market microstructure applications, signature features can encode the sequential structure of order flow, trades, and order book updates. Since these settings

are inherently sensitive to the ordering and interaction of events, signature-based representations provide a natural way to design signals for tasks such as optimal execution, liquidity modeling, and short-horizon prediction.

5 Conclusion

This paper presented truncated path signatures as a principled representation of financial time series that retains information about how prices evolve, not only where they end. By encoding cumulative returns, volatility structure, and higher-order interactions in a systematic hierarchy, signatures offer a flexible feature language for phenomena that are inherently path-dependent in markets, from volatility clustering and asymmetry to microstructure effects. Theoretical results motivate their use by showing that signatures characterize paths up to a natural equivalence and provide a universal basis for approximating path functionals using linear combinations of truncated terms. We also highlighted the practical relevance of Chen’s theorem for incremental and rolling-window computation, and discussed the computational trade-offs created by truncation and dimensional growth. Overall, signature representations provide a mathematically grounded bridge between continuous-time path structure and modern data-driven methods in quantitative finance, with clear avenues for empirical validation in volatility forecasting, path-dependent derivatives, and trading signal design.

References

- [1] T. Lyons, M. Caruana, and T. Lévy, *Differential Equations Driven by Rough Paths*, Springer, 2007.
- [2] B. Hambly and T. Lyons, *Uniqueness for the signature of a path of bounded variation and the reduced path group*, Annals of Mathematics, 2010.
- [3] P. Friz and M. Hairer, *A Course on Rough Paths*, Springer, 2014.