

Approximation of stochastic insurer balance-sheet results using signatures of economic scenarios.

Hervé Andrès¹, Alexandre Boumezoued¹, Arthur Bourdon^{1,2}, Benjamin Jourdain²

¹Milliman R&D, Paris, France

²CERMICS, ENPC, Institut Polytechnique de Paris, CNRS, Marne-la-Vallée, France & MATHRISK team-project, Inria Paris, France.

Contributed Talk at the *"European Actuarial Journal Conference 2026"*, Istanbul, Türkiye.

September 10, 2026



For regulatory purposes, insurers need to value

1. their assets $\rightsquigarrow$ "easy", usually model-independent (existence of liquid markets),
2. their liabilities (BE, VIF) $\rightsquigarrow$ "hard", more challenging task (no more liquid markets).

Following the risk-neutral valuation principle:

Fair value of liabilities = $\mathbb{E}[\text{discounted cashflows}] \rightsquigarrow$ no closed form, requires Monte Carlo approach.

How to sample the discounted cash flows?

discounted cashflows = $\Phi(\text{economic scenarios})$,

where the economic scenarios are sampled from an ESG, and the deterministic mapping

$\Phi : \text{economic scenarios} \mapsto \text{discounted cashflows}$,

is the so-called ALM model.

Main difficulties:

- ALM models involve complex management rules, profit-sharing mechanisms, surrender behaviors and regulatory constraints $\rightsquigarrow$ **strong nonlinearities, barriers, and path-dependence.**
- The inherently sequential nature of the computations $\rightsquigarrow$ **precludes parallelization across time.**
- The presence of conditional mechanisms (e.g. management rules) $\rightsquigarrow$ **limits vectorization and efficient implementation.**

$$\text{total cost} = \text{cost per scenario} \times \underbrace{\text{number of economic scenarios}}_{\text{large and kept fixed}}$$

reduce the total cost $\rightsquigarrow$ reduce the cost per economic scenario $\rightsquigarrow$ **develop a surrogate for the ALM model functional Φ .**

What do we approximate?

Let

- $T > 0$ denote the (finite) horizon of all projections,
- $X := (X_t)_{t \in [0, T]}$ denote the (possibly random) economic scenario,
- $CF = \Phi(X)$ denote the associated discounted cashflows.

We want to approximate the functional $\Phi : (X_t)_{t \in [0, T]} \mapsto CF \rightsquigarrow$ find another functional $\tilde{\Phi}$ such that $\Phi \simeq \tilde{\Phi}$.

Approximation of a functional defined on a path space!

Signature of paths: a powerful family of features

For a *smooth* path $X := (X_t^1, \dots, X_t^d)_{t \in [0, T]}$, and for each word $i_1 \dots i_n$ with $i_k \in \{0, \dots, d\}$ and $n \in \mathbb{N}^*$ a positive integer, we define:

$$\widehat{\mathbb{X}}_T^{i_1 \dots i_n} = \int_0^T \int_0^{t_1} \dots \int_0^{t_{n-1}} d\widehat{X}_{t_n}^{i_1} \dots d\widehat{X}_{t_1}^{i_n},$$

and

$$\widehat{\mathbb{X}}_T^\emptyset = 1,$$

where $\emptyset$ denotes the empty word, $\widehat{X}_t^0 = t$ and $\widehat{X}_t^i = X_t^i$ for $i \in \{1, \dots, d\}$.

- The **time-augmented signature** of X is $\widehat{\mathbb{X}}_T := (\widehat{\mathbb{X}}_T^{i_1 \dots i_n})_{(i_1, \dots, i_n) \in \{0, \dots, d\}^n, n \in \mathbb{N}}$.
- For $N \in \mathbb{N}$, the **N -step truncated time-augmented signature** of X is $\widehat{\mathbb{X}}_T^{\leq N} := (\widehat{\mathbb{X}}_T^{i_1 \dots i_n})_{(i_1, \dots, i_n) \in \mathcal{A}^n, n \leq N}$
 $\rightsquigarrow \widehat{\mathbb{X}}_T^{\leq N} \in \mathbb{R}^{\frac{(d+1)^{N+1}-1}{d}} \rightsquigarrow$ **finite family of features**.
- Libraries for efficiently computing signatures include: *iisignature* [4], *signatory*, *esig*.

When X is a **continuous semimartingale**: we define its time-augmented signature using **Stratonovich** integration theory.

Signature of paths: a powerful family of features

Example. Let $(X_t)_{t \in [0, T]}$ be a *smooth* real-valued path, so that the alphabet is $\{0, 1\}$.

Level	Components
0	$\widehat{\mathbb{X}}_T^\emptyset = 1$
1	$\widehat{\mathbb{X}}_T^0 = \int_0^T dt = T, \quad \widehat{\mathbb{X}}_T^1 = \int_0^T dX_t = X_T - X_0$
2	$\widehat{\mathbb{X}}_T^{00} = \int_0^T \int_0^t ds dt = \frac{T^2}{2}, \quad \widehat{\mathbb{X}}_T^{01} = \int_0^T \int_0^t ds dX_t = \int_0^T t dX_t,$ $\widehat{\mathbb{X}}_T^{10} = \int_0^T \int_0^t dX_s dt = \int_0^T (X_t - X_0) dt, \quad \widehat{\mathbb{X}}_T^{11} = \int_0^T \int_0^t dX_s dX_t = \frac{(X_T - X_0)^2}{2}$

Think of the signature as a polynomial features, but for paths.

Signature of paths: why should this work?

Sufficiently rich linear combinations of signature terms approximate continuous path functionals [2]:

$$\Phi(X) \simeq \underbrace{\langle \beta, \widehat{\mathbb{X}}_T^{\leq N} \rangle}_{\tilde{\Phi}(X)}, \quad \text{for large } N \text{ and some } \beta \in \mathbb{R}^{\frac{(d+1)N+1-1}{d}}$$

nonlinear path-dependent ALM functional $\rightsquigarrow$ linear model in nonlinear path features

$$\beta_N^* := \arg \min_{\beta} \mathbb{E}[|\Phi(X) - \langle \beta, \widehat{\mathbb{X}}_T^{\leq N} \rangle|^2]$$

The number of features grows like $(d + 1)^{N+1} \rightsquigarrow$ **problematic when d is large!**

Consider an economic scenario (ESG output) containing:

- the risk-free curve across 30 maturities $\rightsquigarrow$ **each maturity is encoded by a letter,**
- a basket of 5 equities $\rightsquigarrow$ **each equity is encoded by a letter,**

$d = 35$ $\rightsquigarrow$ 1333 features ($N = 2$), 47989 features ($N = 3$)

Realistic ESG outputs can contain many more coordinates/risk factors (≈ 300)!



A preliminary dimension-reduction step is therefore required

In what follows, $X = (X_t^1, \dots, X_t^d)_{t \in [0, T]}$ denotes the **compressed/reduced** economic scenario (**d is moderate**).




Figure: Signature-based model. Parameters: μ, D, β_0, β . Hyper-parameters: φ, N . Blocks circled in red are optimized during training; blocks circled in gray are specified prior to training.

ALM setup:

1. $CF = \sum_{t=1}^T D(0, t) P\&L_t$, $VIF = \mathbb{E}[CF]$,
2. ALM model: Python implementation of an open source model [1].

ESG setup:




1. 2000 economic scenarios $\rightsquigarrow$ split 70%/30% into Train/Test,
2. Risk factors:
 - risk-free rate curve $\rightsquigarrow$ 40 maturities,
 - 2 equity indexes,
 - 7 risky rate curves $\rightsquigarrow$ 40 maturities for each rating,
 - rating transition probabilities.

Computing the 2000 cashflow outputs with the full ALM model $\rightsquigarrow$ 3 hours

Order	R^2 (%)	Time (s)
2	80.44	0.7
3	92.72	5.8
4	94.63	7.1
5	95.25	56.8




Table: Out-of-sample results for each truncation order.

Surrogate training time (excluding ESG generation)




Empirical study of cashflow replication. Results are obtained out-of-sample with the truncation order equal to 4.

VIF relative error $\rightsquigarrow -0.2\%$



Empirical study of residuals. Results are obtained out-of-sample with the truncation order equal to 4.

Comparison with other ML models





Benchmark. Comparison with methods from the *aeon* library [3].

- $\beta^* = \arg \min_{\beta} \mathbb{E}[|\Phi(X) - \langle \beta, \hat{\mathbb{X}}_T^{\leq N} \rangle|^2] \rightsquigarrow$ sensitive to the distribution of X .

What happens if β^* is estimated under a given ESG distribution (e.g. calibrated to current market data), and the ESG distribution is subsequently perturbed without retraining the surrogate? (changes in market conditions or in the ESG specification) ?

Robustness with respect to the ESG distribution



Robustness under ESG parameter perturbations. The boxplots report the distribution of out-of-sample R^2 values and VIF relative errors across 50 perturbed ESG parameterizations. For each parameterization, 1000 new economic scenarios are generated and evaluated with the full ALM model. The signature-based surrogate is kept fixed after calibration on the central ESG distribution. The dotted horizontal line corresponds to the value obtained under the central ESG distribution on the reference dataset of 600 scenarios.

1. Signatures provide a natural representation of path-dependent ALM outputs.

A complex nonlinear functional of an economic scenario can be approximated by a **linear model on signature features**.

2. Low-order signatures already provide high predictive accuracy at very low computational cost.

In our experiment:

$$N = 4 \quad \implies \quad R^2 = 94.63\%, \quad \text{training time} = 7.1 \text{ s}, \quad \text{VIF error} = -0.2\%.$$

3. The surrogate can be reused under moderate shifts away from the ESG distribution on which it was trained.

It remains accurate under moderate ESG perturbations **without retraining**.

- [1] Aurélien Alfonsi, Adel Cherchali, and Jose Arturo Infante Acevedo.
A synthetic model for asset-liability management in life insurance, and analysis of the SCR with the standard formula.
European Actuarial Journal, 10(2):457–498, 2020.
- [2] C. Cuchiero, F. Primavera, and S. Svaluto-Ferro.
Universal approximation theorems for continuous functions of càdlàg paths and Lévy-type signature models.
Finance and Stochastics, 29(2):289–342, 2025.
- [3] Matthew Middlehurst, Ali Ismail-Fawaz, Antoine Guillaume, Christopher Holder, David Guijo-Rubio, Guzal Bulatova, Leonidas Tsaprounis, Lukasz Mentel, Martin Walter, Patrick Schäfer, et al.
aeon: a Python toolkit for learning from time series.
Journal of Machine Learning Research, 25(289):1–10, 2024.
- [4] J. Reizenstein and B. Graham.
Efficient calculation of iterated-integral signatures and log signatures.
ACM Transactions on Mathematical Software, 2020.

Thanks for your attention!

Questions?