

Signature-Informed Transformer for Asset Allocation

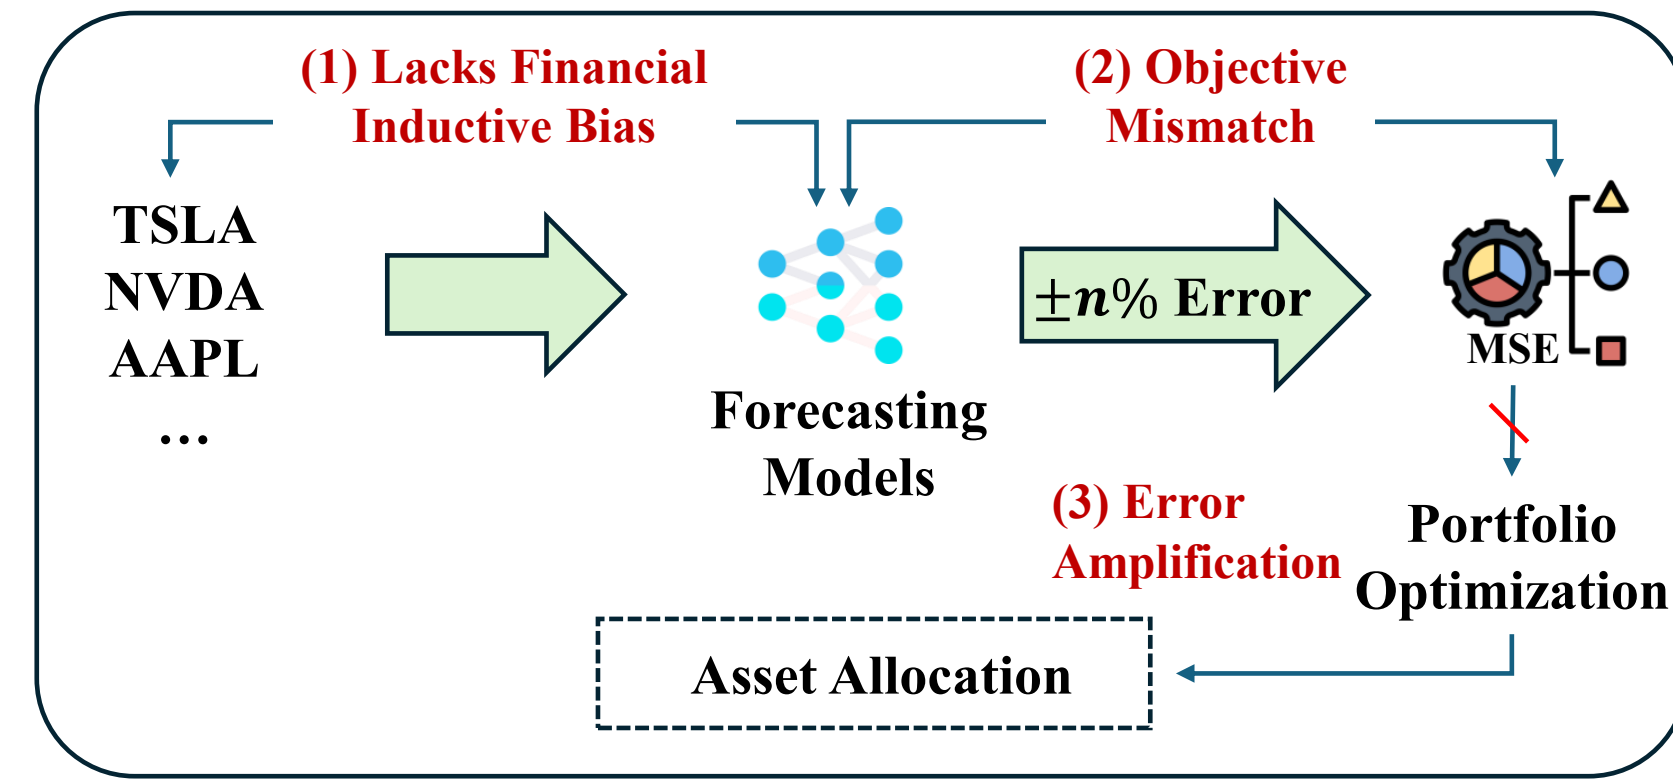
Yoontae Hwang¹ • Stefan Zohren²

¹Pusan National University, Busan, Korea • ²University of Oxford, Oxford, UK

The Problem: A Fragile Two-Stage Pipeline

The dominant deep-learning recipe for portfolio construction is **decoupled**: a general-purpose model first *forecasts* returns, and an optimizer then *allocates*. This pipeline breaks down in three compounding ways:

- ▶ **No financial inductive bias** — generic forecasters cannot model market structure such as *lead-lag* relations between assets.
- ▶ **Objective mismatch** — minimizing MSE is agnostic to the downstream decision; it also over-rewards easy-to-predict assets and distorts the allocation.
- ▶ **Error amplification** — the optimizer magnifies tiny ranking errors into volatile, impractical portfolio weights.



Our Approach: SIT

We unify feature extraction and decision making into a **single end-to-end policy** built on three synergistic pillars:

1. **Path-wise representation.** Rough-path *signatures* summarise each asset's price trajectory — trends, oscillations, and signed areas.
2. **Signature-augmented attention.** A novel cross-asset attention bias from *pairwise cross-signatures* injects geometric lead-lag structure.
3. **Decision alignment.** Training directly minimizes portfolio **CVaR** — the objective is the financial goal, closing the predict-then-optimize gap.

A long-only, fully-invested allocation learned solely from a tail-risk objective — no auxiliary prediction loss.

Path Signatures: A Geometric Prior

For a path $\mathbf{X} : [s, t] \rightarrow \mathbb{R}^d$, the order- M truncated signature is

$$\text{Sig}^M(\mathbf{X}_{[s,t]}) = \left(1, \int_s^t d\mathbf{X}_u, \int_s^t \int_s^u d\mathbf{X}_r \otimes d\mathbf{X}_u, \dots \right).$$

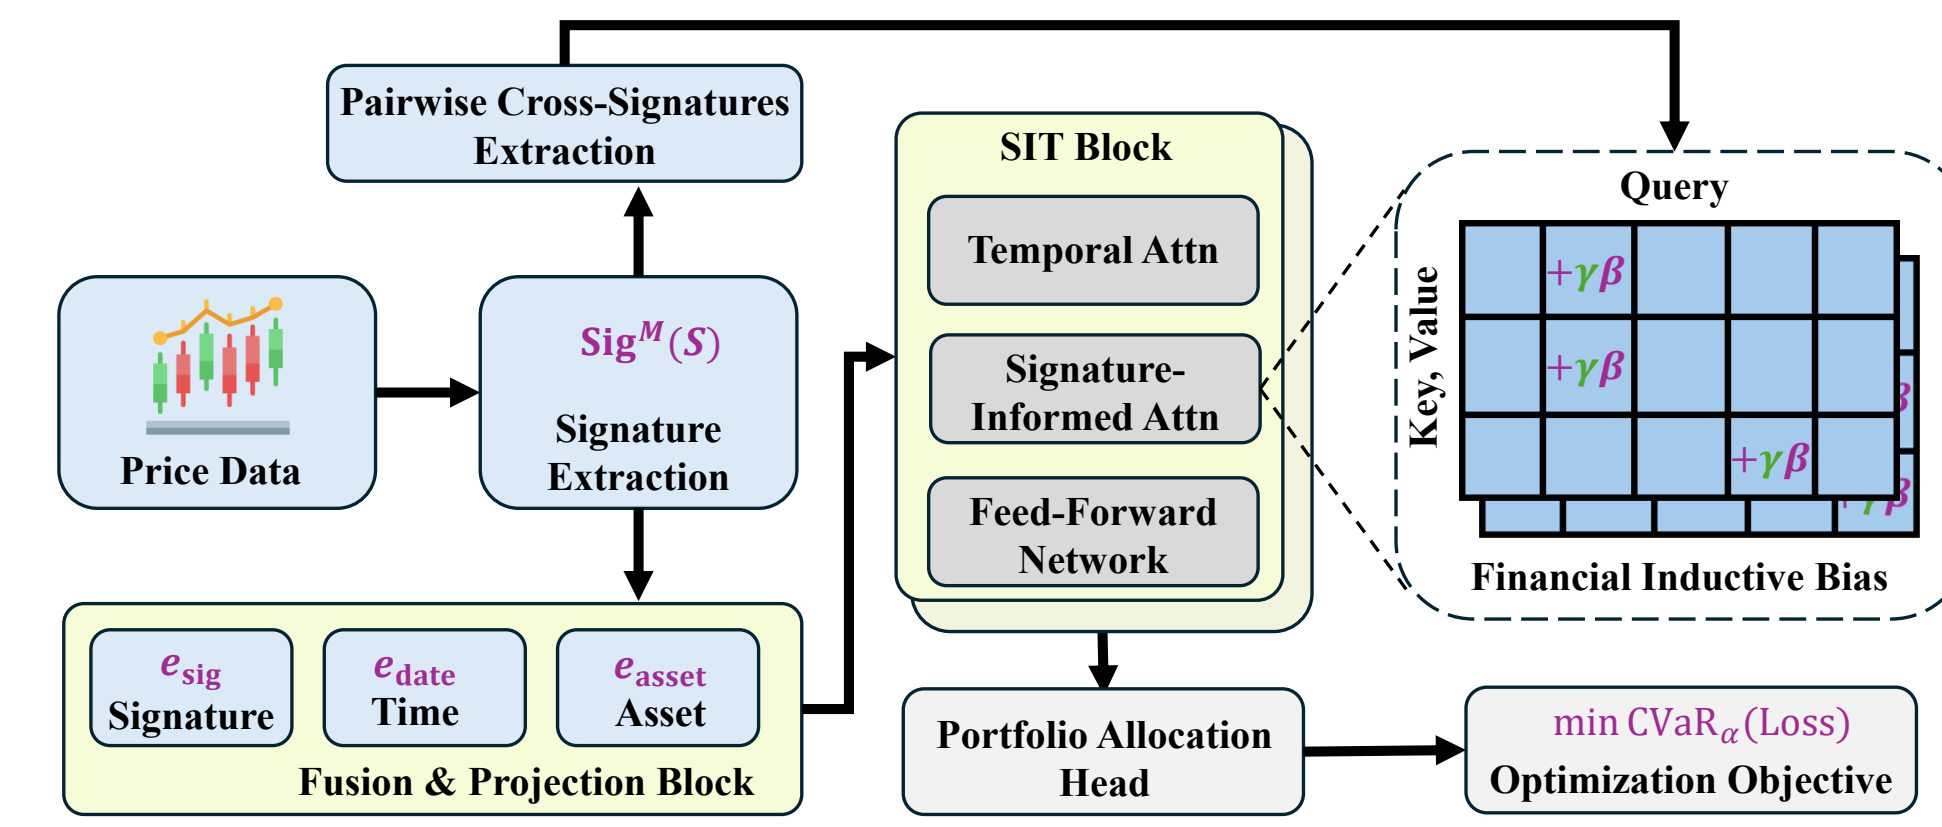
First-order terms capture net increments; **second-order** terms encode *signed areas* that reveal non-linear correlations and directional lead-lag. We use $M = 2$ — the lowest order carrying signed-area information — and update incrementally via **Chen's relation**, so signatures cost $\mathcal{O}(1)$ per step with no recomputation over history.

Conclusion

SIT decisively outperforms traditional strategies and advanced forecasting baselines, and remains superior under realistic transaction costs — performance is driven by the *calibrated, signature-based* architecture, not raw predictive accuracy.

Aligning the training objective with the financial objective is necessary to turn predictive signals into reliable portfolios.

The SIT Architecture



Unified embeddings. At each decision time, every token fuses three sources — the truncated path signature \mathbf{e}_{sig} , calendar features \mathbf{e}_{date} , and a learnable asset embedding $\mathbf{e}_{\text{asset}}$ — concatenated and projected to $\mathbb{R}^{d_{\text{model}}}$. A factorised Transformer then runs causal *temporal* attention followed by *cross-asset* attention within each layer.

Signature-Informed Self-Attention

Cross-asset queries/keys/values $\mathbf{Q}, \mathbf{K}, \mathbf{V}$ are augmented with an *additive, data-dependent bias*. The pairwise cross-signature $\mathbf{c}_{i,j,l}$ is embedded by MLP_β into $\beta_{i,j,l}$; a dynamic query $\mathbf{q}_{k,j}^{\text{dyn}} = \text{MLP}_q(\mathbf{x}'_{k,j})$ encodes the query asset's *informational need*. Their inner product forms the bias $b_{k,h,j,l} = \langle \mathbf{q}_{k,j}^{\text{dyn}}, \beta_{i,j,l} \rangle$, scaled by a learnable gate $\gamma = \text{softplus}(\hat{\gamma}) > 0$:

$$\text{Logits}_{k,h} = \frac{\mathbf{Q}_{k,h} \mathbf{K}_{k,h}^\top}{\sqrt{d_k}} + \gamma \mathbf{B}_{k,h}.$$

Factorisation keeps the cost at $\mathcal{O}(dH^2d_{\text{model}} + Hd^2d_{\text{model}} + d^2C_{\text{cross-sig}}(M))$.

Decision-Focused CVaR Objective

Weights are a tempered softmax of predicted returns, $\mathbf{w}_t^{(k)} = \text{softmax}(\hat{\mathbf{r}}_t^{(k)} / \tau)$, with per-step loss $L^{(k)} = -(\mathbf{w}_t^{(k)})^\top \mathbf{r}_{t+k}$. The *sole* training signal is the expected CVaR of the K -step loss, optimised via its dual and empirical form:

$$\begin{aligned} \mathcal{L}(\theta) &= \mathbb{E}_{\omega \sim \mathcal{D}} [\text{CVaR}_\alpha(\{L^{(k)}(\theta_\omega)\}_{k=1}^K)] \\ &= \mathbb{E}_\omega \left[\min_{v_\omega} v_\omega + \frac{1}{(1-\alpha)K} \sum_k (L^{(k)} - v_\omega)^+ \right] \\ &\approx \frac{1}{N} \sum_i \min_{v_i} v_i + \frac{1}{(1-\alpha)K} \sum_k (L_{\omega_i}^{(k)} - v_i)^+. \end{aligned}$$

No prediction loss is used; gradients flow only from the portfolio objective, removing objective mismatch by construction.

Ablation — Every Pillar Matters (Asset-40)

Variant	Sharpe \uparrow	Sortino \uparrow	MDD \downarrow	Wealth \uparrow
SIT (full)	0.672	0.823	0.361	1.790
w/o CVaR	0.569	0.706	0.370	1.641
w/o Asset Attn	0.528	0.658	0.334	1.538
w/o Financial Bias	0.605	0.759	0.343	1.680
w/o Gate γ	0.525	0.649	0.347	1.547

Main Results — S&P 100


Red = best, blue = 2nd within each panel.

Model	Sharpe \uparrow	Sortino \uparrow	MDD \downarrow	Wealth \uparrow
Panel A — Asset-40 Universe (S&P 100)				
CVaR	0.1531	0.2001	0.3516	1.0569
EW	0.5759	0.7153	0.3688	1.6439
GMV	0.4148	0.5337	0.2743	1.3258
HRP	0.4958	0.6171	0.3185	1.4561
Autoformer	0.2499 \pm 0.1405	0.3423 \pm 0.1980	0.3812 \pm 0.0480	1.1809 \pm 0.2403
DLinear	0.3167 \pm 0.1326	0.4513 \pm 0.2005	0.3621 \pm 0.0407	1.2915 \pm 0.2133
FEDformer	0.4006 \pm 0.2317	0.5540 \pm 0.3192	0.3647 \pm 0.0167	1.5198 \pm 0.5703
iTransformer	0.3157 \pm 0.0749	0.4233 \pm 0.0943	0.4136 \pm 0.0326	1.2860 \pm 0.0147
NSformer	0.4074 \pm 0.1151	0.5820 \pm 0.1655	0.4475 \pm 0.0672	1.5129 \pm 0.3010
PatchTST	0.3286 \pm 0.2021	0.4540 \pm 0.2818	0.4523 \pm 0.0838	1.3409 \pm 0.3886
TimesNet	0.3568 \pm 0.0782	0.4959 \pm 0.1019	0.4704 \pm 0.0701	1.3765 \pm 0.1729
RFormer	0.4901 \pm 0.1437	0.6308 \pm 0.1828	0.3415 \pm 0.0482	1.5387 \pm 0.2353
SIT (Ours)	0.6717 \pm 0.0628	0.8232 \pm 0.0792	0.3611 \pm 0.0037	1.7903 \pm 0.1023

Panel B — Asset-50 Universe (S&P 100)				
CVaR	0.2165	0.2858	0.3086	1.1170
EW	0.6008	0.7399	0.3604	1.6683
GMV	0.3947	0.4992	0.2678	1.2845
HRP	0.4637	0.5620	0.3258	1.4021
Autoformer	0.3899 \pm 0.1985	0.5321 \pm 0.2870	0.4356 \pm 0.1256	1.4697 \pm 0.4573
DLinear	0.2540 \pm 0.1215	0.3557 \pm 0.1828	0.3716 \pm 0.0193	1.1883 \pm 0.1979
FEDformer	0.4318 \pm 0.0692	0.6039 \pm 0.1097	0.4039 \pm 0.1143	1.5286 \pm 0.1508
iTransformer	0.5162 \pm 0.1367	0.6761 \pm 0.1770	0.4542 \pm 0.0239	1.7910 \pm 0.3722
NSformer	0.5238 \pm 0.0694	0.7105 \pm 0.1033	0.4992 \pm 0.0975	1.8138 \pm 0.1922
PatchTST	0.3821 \pm 0.1871	0.5134 \pm 0.2635	0.4255 \pm 0.1533	1.4411 \pm 0.3814
TimesNet	0.3050 \pm 0.3439	0.4296 \pm 0.4864	0.5181 \pm 0.1404	1.3737 \pm 0.8857
RFormer	0.5315 \pm 0.2519	0.6671 \pm 0.3255	0.5202 \pm 0.0555	1.8014 \pm 0.6303
SIT (Ours)	0.7715 \pm 0.0627	0.9743 \pm 0.0998	0.3271 \pm 0.0094	1.9215 \pm 0.1792

SIT consistently delivers the strongest risk-adjusted return and final wealth across both S&P 100 universes, while GMV remains the most conservative on maximum drawdown.

Decision-Focused Learning Improves Sharpe



Across both S&P 100 universes, decision-focused learning (DFL) raises Sharpe for most forecasting backbones relative to MSE training, showing that objective alignment matters even with a fixed architecture.