Savitar: Pharmacology-Aware CP-Tensor Kernel Beats Domain-Specific Baselines in Drug-Combination Bayesian Optimization

Anonymous Authors[1]

Abstract

Off-the-shelf Bayesian-optimization (BO) kernels (Hamming, Tanimoto) for drug-combination prioritisation ignore the per-drug dose-response curve, despite it being routinely measured. We propose **Savitar**, a structured Gaussian-process kernel that consumes a per-drug Hill fingerprint, applies an activity-gated embedding $\mathbf{f}_i = (1 - s_i)\mathbf{Me}_i^z$ (with s_i the Hill-predicted viability at the queried dose), and aggregates across drugs through a symmetric CP-tensor parameterization of arbitrary-order interactions sharing $O(Dq)$ parameters across all orders. We evaluate retrospectively, simulating a 30-evaluation drug-combination wet-lab run against a fixed pool of $\sim$45,000 measured candidates per cell line; regret is the gap between our method's best-picked combination (measured via percent-growth) and the pool minimum. Across oncology, antifungal, antiviral, and antibacterial health domains, Savitar achieves the lowest mean regret on every dataset against properly-published domain-aware baselines (chemical: Tanimoto, biological: GIP, and chemogenomic: INDIGO); on the 60-cell oncology dataset aggregate it halves the regret of every baseline. On HIV it hits the pool oracle on $20/20$ seeds, and on every dataset it runs 2–10$\times$ faster per trajectory than the GP baselines. The parameterization extends to $k \geq 3$-drug regimens; biological validation beyond pairs is future work.

1. Introduction

Combination therapy is the clinical reality in oncology, but the search space is overwhelming: NCI-ALMANAC (Holbeck et al., 2017) and the O'Neil pair screen (O'Neil et al., 2016) together catalogue $> 10^5$ pair measurements, while the cross-product of $\sim$100 approved oncology compounds and 60 cell lines reaches the tens of millions. Resource-constrained follow-up is the default mode, exactly the regime Bayesian optimization (BO) was designed for (Jones et al., 1998; Snoek et al., 2012; Frazier, 2018; Garnett, 2023). Combinatorial BO has been studied extensively under an off-the-shelf-kernel paradigm (Oh et al., 2019; Ru et al., 2020; Wan et al., 2021; Liang et al., 2024; Doumont et al., 2025) and parallel high-dim work shows careful trust-region or bandwidth design (Eriksson et al., 2019; Eriksson & Jankowiak, 2021; Hvarfner et al., 2024) can rescue vanilla GP-BO; yet none of these encode *single-drug pharmacology*. Decades of quantitative pharmacology (Hill, 1910; Bliss, 1939; Loewe, 1953; Rønneberg et al., 2021) say a great deal about how monotherapy dose-response curves combine, and we argue this knowledge belongs inside the kernel, not spread across millions of NN parameters.

Why a structured GP rather than deep learning? Neural combination predictors (Preuer et al., 2018; Kuru et al., 2021; Wang et al., 2021a;b; Julkunen et al., 2020) are well-suited to the dense retrospective regime they were trained on, but our target is low-budget sequential optimization with a new cell line and $\leq$ 30 queries: training cost (10^4–10^5 triples), poor leave-pair / leave-cell generalisation (Baptista et al., 2023), and lack of calibrated uncertainty for an acquisition function rule out neural surrogates as drop-in replacements. Structured GPs, where pharmacology is a kernel design choice rather than a training-set requirement, fit naturally; Savitar is engineered for this short-budget regime, with a 30-evaluation BO trajectory under two seconds on CPU and 2–10$\times$ faster than published GP baselines (Table 2). It consumes a per-(drug, cell line) Hill fit as side information, the routine output of a single-agent screen.

Contributions: We introduce a structured GP kernel (Section 2) combining an *activity-gated per-drug embedding* $\mathbf{f}_i = (1 - s_i)\mathbf{Me}_i^z$ with a *symmetric CP-tensor* parameterization of arbitrary-order drug interactions sharing $O(Dq)$ parameters across all orders; configurable in arity k, it applies directly to $k \geq 3$ regimens.[1] On NCI-ALMANAC across the full NCI-60 panel (60 cell lines $\times$ 10 seeds), Savitar halves the regret of the categorical baselines and

[1]Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

[1]Code, configs, and Hill fits at the anonymous repository.

beats chemistry-aware Tanimoto by $1.8\times$. On three additional health datasets (*Cokol 2011* antifungal in *S. cerevisiae*, *Tan 2012* antiviral against HIV, *Brochado 2018* antibacterial in *E. coli*), each compared against a properly-published domain-aware baseline (GIP, chem Tanimoto, INDIGO chemogenomic), Savitar takes the pool oracle on $20/20$ HIV seeds and beats INDIGO by $1.2\times$ on *E. coli*. An *Activity-Hamming* control with the same Hill fits but an identity pair gate ties plain Hamming and ties random, so the win is not from Hill-curve access alone; the Tanimoto comparison shows it is not from chemical similarity alone either. A synthetic arity sweep ($k = 2, \ldots, 5$; App. C) shows the same pattern at every order.

2. Method

2.1. Hill-curve fingerprints

For each (drug, cell line), we fit a 4-parameter Hill curve $\big(E_0, E_{\max}, \log_{10}\mathrm{EC}_{50}, \mathrm{slope}\big)$ (E_0 baseline viability, $E_{\max}$ maximum effect, EC_{50} half-effect dose, slope steepness) to the per-cell monotherapy dose-response data by non-linear least squares, with a 3-parameter fallback ($E_0 = 100$) if the full fit fails. On ALMANAC's full NCI-60 panel this fits $6{,}293/6{,}295$ (drug, cell) combinations (99.97% success). The resulting fingerprint $\mathbf{e}_{i,c} \in \mathbb{R}^4$ is a fixed (non-learnable) feature stored per (drug, cell). Per-cell-line we z-score the four components to get $\mathbf{e}^z_{i,c}$ for use as the projection input.

2.2. Savitar: activity-gated CP-tensor combo kernel

For a combination $c = (i_1, \ldots, i_k; d_1, \ldots, d_k)$ at log-doses d_j, Savitar[2] defines a per-drug embedding

$$\mathbf{f}_i = (1 - s_i) \cdot \mathbf{M}\mathbf{e}^z_i \in \mathbb{R}^q, \tag{1}$$

where $s_i \in [0, 1]$ is the Hill-predicted viability fraction at the actual dose d_i,

$$s_i = \mathrm{clip}\big(\tfrac{1}{100}\,\mathrm{hill}(d_i; \mathbf{e}_i),\ 0,\ 1\big), \tag{2}$$

and $\mathbf{M} \in \mathbb{R}^{q\times 4}$ is a learnable linear projection. The gate $(1 - s_i)$ is the per-drug kill fraction at dose d_i: a drug at a no-effect dose ($s_i \to 1$, viability unchanged) contributes zero feature mass, and the embedding turns on smoothly as the dose enters the effective range. Higher-order interaction terms in Equation (3) then pick up the joint factor $\prod_{j\in S}(1 - s_{i_j})$, so a synergistic pair only contributes when both drugs are dose-active.

[2]Savitar (Vedic Sanskrit, "the impeller"): the solar deity who drives motion in Hindu cosmology. We adopt the name to evoke a surrogate's role in BO, where every acquisition decision *drives* the next experiment based on the model's predictions about unseen combinations.

We then aggregate across drugs through a *symmetric CP-tensor* parameterization of order-m subset interactions over the k combination slots, with $[k] = \{1, \ldots, k\}$:

$$\boldsymbol{\phi}(c) \;=\; \sum_{j=1}^{k} \mathbf{f}_{i_j} \;+\; \sum_{m=2}^{m_{\max}} \sum_{\substack{S\subseteq[k]\\|S|=m}} w_S \cdot \bigodot_{j\in S} \mathbf{f}_{i_j} \;\in\; \mathbb{R}^q, \tag{3}$$

where $\odot$ denotes the element-wise (Hadamard) product. The first sum is the (unweighted) solo-term contribution; the higher-order CP weights index *drug identities*, not slot positions:

$$w_S \;=\; \sum_{r=1}^{q} \prod_{j\in S} \mathbf{A}[i_j,\, r], \qquad \mathbf{A} \in \mathbb{R}^{D\times q}, \quad |S| \geq 2. \tag{4}$$

Equation (4) is a rank-q CP decomposition (Kolda & Bader, 2009; Sidiropoulos et al., 2017) of the order-m symmetric drug-interaction tensor. The same matrix $\mathbf{A}$ handles every higher-order interaction simultaneously, so the parameter count is $O(Dq)$ regardless of $m_{\max}$ (vs. $O(D^m)$ for a naive order-m tensor); $\mathbf{A}$ is regularised implicitly by the GP marginal likelihood. For example at $k = 3, m_{\max} = 2$, the pair weight $w_{\{1,2\}} = \sum_r \mathbf{A}[i_1, r]\mathbf{A}[i_2, r]$ shares its row of $\mathbf{A}$ with $w_{\{1,3\}}$ and $w_{\{2,3\}}$, so observing (i_1, i_2) informs predictions for (i_1, i_3) and (i_2, i_3).

The kernel is a radial basis function on $\boldsymbol{\phi}$:

$$k(c, c') \;=\; \sigma^2 \exp\!\big(-\tfrac{1}{2\ell^2}\|\boldsymbol{\phi}(c) - \boldsymbol{\phi}(c')\|^2\big)\,, \tag{5}$$

with learnable signal variance σ^2 and lengthscale ℓ.

Proposition 2.1. *For any $\mathbf{M} \in \mathbb{R}^{q\times 4}$, $\mathbf{A} \in \mathbb{R}^{D\times q}$, and any $\sigma^2, \ell^2 > 0$, the kernel of Equation (5) is positive semi-definite (PSD).*

Proof. $\boldsymbol{\phi} : \mathcal{C} \to \mathbb{R}^q$ in Equation (3) is a deterministic feature map (it depends on c alone, not on the data). The Gaussian RBF on Euclidean features $\mathbb{R}^q$ is PSD (Rasmussen & Williams, 2006, §4.2.1), and PSD-ness is preserved under composition with a deterministic feature map. □

A drug's row $\mathbf{A}[i, :]$ is shared across every higher-order term containing drug i, so observing a synergistic pair (i, j) updates predictions for *every other combination containing i or j*, including unseen triplets.

2.3. BO loop

A query is a tuple $c = (i_1, i_2; d_1, d_2)$; the response y is the measured ALMANAC percent-growth (lower = more cell death). At BO step t we refit $\{\mathbf{M}, \mathbf{A}, \ell, \sigma^2, \sigma_n^2\}$ on the observations $\mathcal{T}_t$ by Type-II marginal-likelihood maximisation (empirical Bayes; 30 Adam steps, lr 0.05; Rasmussen & Williams, 2006; Frazier, 2018; Garnett, 2023), score each

unseen c in the pool by Expected Improvement (Jones et al., 1998; Snoek et al., 2012), and query the argmax. Initial training data $\mathcal{T}_{n_{\text{init}}}$ is sampled uniformly from the pool; we use $q = 8$, $m_{\max} = 2$, ZeroMean throughout.

3. Experiments

3.1. ALMANAC setup

NCI-ALMANAC (Holbeck et al., 2017) is $\sim 10^5$ measured percent-growth values per cell line over $\sim$100 drugs on a 3×3 dose grid; after replicate-averaging and pair-orientation canonicalisation each cell yields $\sim$45,000 candidates, which we uniformly subsample to 4,000 for tractable EI. Budget is $n_{\text{init}} = 10$ random observations plus $n_{\text{iter}} = 20$ BO steps over all 60 NCI-60 cell lines and 10 seeds; 10 seeds $\times$ 60 cells gives 600 paired trajectories per method, more than enough power to drive every per-pair Wilcoxon below $p < 10^{-38}$ (Section 3.2), so additional replicates would not change the conclusion. *Following the "fix-pipeline-vary-only-the-kernel" protocol of Doumont et al. (2025), all methods share the same pool and the same init per (cell, seed)*; final regret is relative to the subsampled-pool minimum. Five methods, all with ZeroMean: **Random** (uniform); **Hamming** (HammingPair$\times$MixedMatern on log-dose, the standard discrete-pair baseline); **Activity-Hamming** (same, but Matérn over the activities $1 - s_i$, isolating Savitar's structural contribution from Hill-curve access); **Tanimoto** (TanimotoPair$\times$MixedMatern on log-dose; Ralaivola et al., 2005; Griffiths et al., 2023, per-combo Morgan ECFP4 radius-2, 2048 bits, OR-aggregated); and **Savitar** (Equations (1) to (5)).

3.2. ALMANAC results

Savitar halves mean final regret relative to the categorical baselines (all $\sim$56) and beats Tanimoto by $1.8\times$ (27.4 vs 48.6); the median tells the same story. Across the 600 paired (cell, seed) trajectories, the per-pair regret reduction is highly significant (paired Wilcoxon, $p < 10^{-38}$ in all four cases), with strictly-negative 95% paired-bootstrap CIs on mean(Savitar $-$ baseline) and per-pair win/tie/loss counts 392/89/119 vs random, 359/180/61 vs each Hamming variant, and 338/167/95 vs Tanimoto. The Hamming variants are indistinguishable from random, so giving the discrete-pair kernel access to monotherapy curves (Activity-Hamming) does not close the gap; Tanimoto narrows it via chemical similarity but still loses to Savitar on 338/600 trajectories: the activity-gated CP-tensor extracts signal that neither pair identity nor chemistry alone captures.

Synthetic arity sweep ($k = 2, \ldots, 5$, App. C): ALMANAC restricts to $k = 2$; a synthetic universe ($D = 25$, random Hill fingerprints, Bliss-plus-CP ground truth) shows Savitar wins both Hamming variants at every arity from $k = 2$ to $k = 5$ (full numbers in App. C).

3.3. Cross-domain validation

To test whether Savitar's structure transfers beyond oncology, we run the same protocol on three additional health datasets; antifungal (Cokol et al., 2011), antiviral (Tan et al., 2012), and antibacterial (Brochado et al., 2018). Savitar hyperparameters are *fixed in advance from the ALMANAC configuration* ($q = 8$, $m_{\max} = 2$, ZeroMean, 30-step Adam fit at lr 0.05) and not retuned per dataset. Each domain is compared against a domain-appropriate published baseline kernel (**GIP** on drug-target gene profiles (Yamanishi et al., 2008; van Laarhoven et al., 2011) for Cokol, chem Tanimoto for Tan, and an INDIGO-style chemogenomic RBF (Nichols et al., 2011; Chandrasekaran et al., 2016) for Brochado, on the chemogenomically-covered subpool from Nichols et al., 2011). Savitar achieves the lowest mean regret on every dataset (Table 2); the paired one-sided Wilcoxon (Savitar $<$ baseline, $n = 20$ seeds) gives $p = 10^{-4}$ on Tan, $p = 0.06$ on Cokol, and $p = 0.17$ on Brochado, so the Tan win is unambiguous and the Cokol/Brochado wins are directionally robust but underpowered at this seed budget. On Brochado the published Hill gate is degenerate ($R^2 < 0$, since the screen used only 3 sub-MIC dose levels), yet Savitar still beats the INDIGO baseline by $1.2\times$; the CP-tensor structure carries the win on its own when the gate is uninformative.

Clinical-dynamics interpretation (App. A): On OVCAR-8 seed 0, plugging each method's chosen combo into a Gompertz + linear-PK tumor model yields a 10^7-fold tumor reduction for Savitar (matching the pool oracle) vs. no shrinkage for Hamming-based BO; the map is rank-preserving, so this re-expresses the BO result in clinician-readable units (full equations and sensitivity analysis in App. A).

4. Discussion

Validation regime: The retrospective protocol (picking 30 of $\sim$45,000 already-measured candidates per cell line) models the deployment setting where a new cell line, disease, or drug library lacks an ALMANAC-scale screen and a method must find a near-optimal combination at $\sim$0.07% of the assay budget; this is the standard validation in combinatorial-BO and drug-combination ML (Baptista et al., 2023; Doumont et al., 2025).

Localising the contribution: Activity-Hamming controls for Hill-curve access; Tanimoto controls for molecular similarity (Morgan ECFP4); neither closes the gap on the full panel (Table 1). A third internal control, *Savitar-no-M* (fixing $\mathbf{M} = I_4$ so the embedding is $(1 - s_i)\mathbf{e}_i^z$ on the raw 4-D Hill fingerprint), ties full Savitar (App. B: 28.6 vs 29.8). The win therefore localises to (a) the activity gate $(1 - s_i)$ and (b) the symmetric CP-tensor sharing $w_S = \sum_r \prod_{j \in S} \mathbf{A}[i_j, r]$,

Table 1. Final regret on NCI-ALMANAC over the full NCI-60 panel (60 cell lines × 10 seeds = 600 trajectories per method) at a budget of 30 evaluations. Savitar halves the regret of the categorical baselines and beats Tanimoto by 1.8×, on both mean and median. The two Hamming variants are statistically indistinguishable from each other and from random; Tanimoto narrows the gap but does not close it. "Cells won" is the count of cell lines on which a method has the best seed-averaged final regret. "Wall" is mean wall-clock per 30-evaluation BO trajectory; Savitar is 2–4× faster than the GP baselines.

Method	Mean ± SD ↓	Median (IQR)	Cells won	Wall
Random	56.4 ± 38.7	54.6 [19.9, 90.6]	3/60	<0.01s
Hamming	56.0 ± 38.1	55.9 [23.5, 87.1]	0/60	2.0s
Activity-Hamming	56.0 ± 38.1	55.9 [23.5, 89.7]	10/60	2.1s
Tanimoto	48.6 ± 33.0	45.6 [20.1, 73.9]	5/60	3.9s
Savitar	**27.4 ± 22.2**	**25.5 [6.0, 43.6]**	**42/60**	**0.9s**

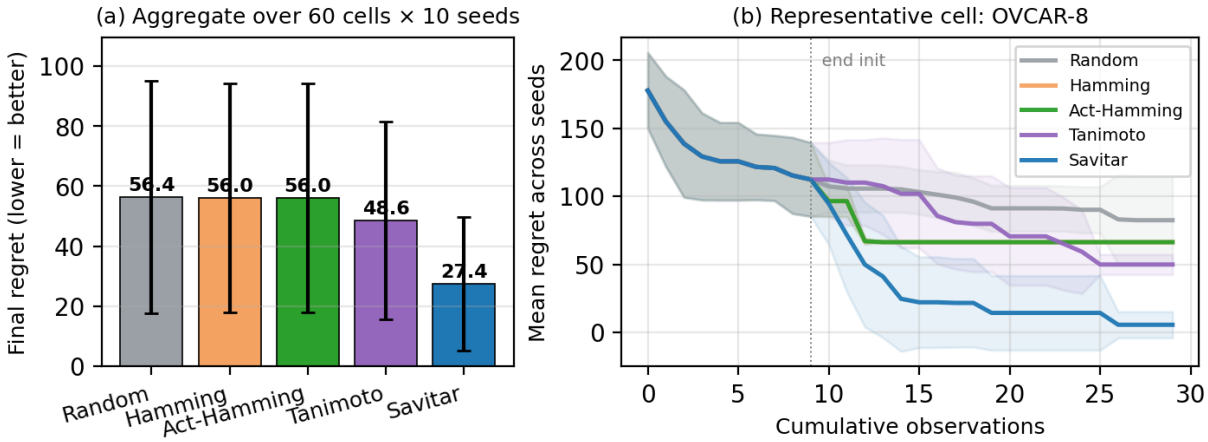

Figure 1. ALMANAC headline. **(a)** Aggregate final regret over the 600 trajectories per method. Savitar halves the regret of the categorical baselines and beats Tanimoto by 1.8×. **(b)** Mean regret on a representative cell (OVCAR-8) with ±1 SD shading across 10 seeds; Savitar drops sharply post-init while every baseline plateaus. Per-cell summaries for two further cells are in App. D.

Table 2. Cross-domain final regret per condition (Cokol, Tan, Brochado: $n = 20$ seeds; ALMANAC: $n = 10$). "Wall" is mean seconds per 30-evaluation BO trajectory (Savitar / strongest baseline).

Dataset	Savitar	Best baseline	Hits	Wall
ALMANAC OVCAR-8	**5.5**	Tan. 49.8	6/10	0.9/3.9
Cokol antifungal	**0.78**	GIP 0.98	3/20	1.9/5.9
Tan HIV	**0.00**	Tan. 9.7	**20/20**	1.5/9.4
Brochado	**4.4**	INDIGO 5.5	0/20	1.2/12.1
MOLT-4 (easy)	1.0	Tan.≡Ham. 0.7	0/10	0.9/3.9
ALMANAC agg.	**27.4**	Tan. 48.6	42/60	0.9/3.9

which propagates synergy across every combination containing a given drug; $\mathbf{M}$ is incidental, retained as the interface for a future MLP. Cross-domain (Table 2) sharpens this: on Brochado *E. coli* the Hill gate is degenerate ($R^2 < 0$ because the screen used only 3 sub-MIC dose levels), yet Savitar still wins, so the CP-tensor structure carries the result. Conversely, on easy cells (MOLT-4: 6.7% of the pool is within regret 5 of the optimum, vs. 0.05% for OVCAR-8), Savitar, baselines, and even random search land near the optimum: Savitar's value scales with combinatorial difficulty.

Limitations: We do not run modern combinatorial-BO baselines (SAASBO, TuRBO, Vanilla high-dim BO, COMBO, CASMOPOLITAN; Oh et al., 2019; Wan et al., 2021; Eriksson et al., 2019; Eriksson & Jankowiak, 2021; Hvarfner et al., 2024) or neural surrogates as GP-replacement (Preuer et al., 2018; Kuru et al., 2021); the synthetic $k = 2, \ldots, 5$ sweep (App. C) shows parameter sharing extends to higher arities but biological validation needs real $k \geq 3$ screens (Cokol et al., 2017).

Future work: Toxicity-constrained acquisition (*EI*×PoF; Gardner et al., 2014) with a tox surrogate sharing the CP-tensor structure; a multitask GP across cell lines; and DL surrogate baselines (DeepDDS / MatchMaker (Wang et al., 2021a; Kuru et al., 2021)). App. E gives the full list.

Impact Statement

Savitar is a methodology paper validated retrospectively on already-measured drug-combination screens; it does not advance specific compounds. Deployment in active wet-lab or clinical pipelines would require the safeguards discussed in Section 4 (toxicity-aware acquisition (Gardner et al., 2014).

References

Baptista, D., Ferreira, P. G., and Rocha, M. Deep learning for drug response prediction in cancer. *Briefings in Bioinformatics*, 24(1), 2023.

Bliss, C. I. The toxicity of poisons applied jointly. *Annals of Applied Biology*, 26(3):585–615, 1939.

Brochado, A. R., Telzerow, A., Bobonis, J., Banzhaf, M., Mateus, A., Selkrig, J., Huth, H., Bassler, B. L., Beas, J. Z., Galardini, M., et al. Species-specific activity of antibacterial drug combinations from in vitro screens in three human pathogens. *Nature*, 559(7713):259–263, 2018.

Chandrasekaran, S., Cokol-Cakmak, M., Sahin, N., Yilancioglu, K., Kazan, H., Collins, J. J., and Cokol, M. Chemogenomics and orthology-based design of antibiotic combination therapies. *Molecular Systems Biology*, 12(5):872, 2016.

Cokol, M., Chua, H. N., Tasan, M., Mutlu, B., Weinstein, Z. B., Suzuki, Y., Nergiz, M. E., Costanzo, M., Baryshnikova, A., Giaever, G., et al. Systematic exploration of synergistic drug pairs. *Molecular Systems Biology*, 7:544, 2011.

Cokol, M., Kuru, N., Bicak, E., Larkins-Ford, J., and Aldridge, B. B. Efficient measurement and factorization of high-order drug interactions in *mycobacterium tuberculosis*. *Science Advances*, 3(10):e1701881, 2017.

Doumont, E., Picheny, V., Borovitskiy, V., and Moss, H. B. Omnipresent yet overlooked: heat kernels in combinatorial Bayesian optimization. In *Advances in Neural Information Processing Systems*, 2025.

Eriksson, D. and Jankowiak, M. High-dimensional Bayesian optimization with sparse axis-aligned subspaces. In *Uncertainty in Artificial Intelligence*, 2021.

Eriksson, D., Pearce, M., Gardner, J., Turner, R. D., and Poloczek, M. Scalable global optimization via local Bayesian optimization. In *Advances in Neural Information Processing Systems*, 2019.

Frazier, P. I. A tutorial on Bayesian optimization. *arXiv preprint arXiv:1807.02811*, 2018.

Gardner, J. R., Kusner, M. J., Xu, Z. E., Weinberger, K. Q., and Cunningham, J. P. Bayesian optimization with inequality constraints. In *International Conference on Machine Learning*, pp. 937–945, 2014.

Garnett, R. *Bayesian Optimization*. Cambridge University Press, 2023.

Griffiths, R.-R., Klarner, L., Moss, H., Ravuri, A., Truong, S., Du, Y., Stanton, S., Tom, G., Rankovic, B., Jamasb, A., et al. GAUCHE: A library for Gaussian processes in chemistry. In *Advances in Neural Information Processing Systems*, volume 36, 2023.

Hill, A. V. The possible effects of the aggregation of the molecules of haemoglobin on its dissociation curves. *The Journal of Physiology*, 40:4–7, 1910.

Holbeck, S. L., Camalier, R., Crowell, J. A., Govindharajulu, J. P., Hollingshead, M., Anderson, L. W., Polley, E., Rubinstein, L., Srivastava, A., Wilsker, D., et al. The NCI ALMANAC: a comprehensive screening resource for the detection of anticancer drug pairs with enhanced therapeutic activity. *Cancer Research*, 77(13):3564–3576, 2017.

Hvarfner, C., Hellsten, E. O., and Nardi, L. Vanilla Bayesian optimization performs great in high dimensions. In *International Conference on Machine Learning*, 2024.

Jones, D. R., Schonlau, M., and Welch, W. J. Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13(4):455–492, 1998.

Julkunen, H., Cichonska, A., Gautam, P., Szedmak, S., Douat, J., Pahikkala, T., Aittokallio, T., and Rousu, J. ComboFM: leveraging multi-way interactions for systematic prediction of drug combination effects. *Nature Communications*, 11(1):6136, 2020.

Kolda, T. G. and Bader, B. W. Tensor decompositions and applications. *SIAM Review*, 51(3):455–500, 2009.

Kuru, H. I., Tastan, O., and Cicek, A. E. MatchMaker: a deep learning framework for drug synergy prediction. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 19(4):2334–2344, 2021.

Liang, H., Wan, X., and Dong, X. Bayesian optimization of functions over node subsets in graphs. In *Advances in Neural Information Processing Systems*, 2024.

Loewe, S. The problem of synergism and antagonism of combined drugs. *Arzneimittelforschung*, 3:285–290, 1953.

Nichols, R. J., Sen, S., Choo, Y. J., Beltrao, P., Zietek, M., Chaba, R., Lee, S., Kazmierczak, K. M., Lee, K. J., Wong, A., et al. Phenotypic landscape of a bacterial cell. *Cell*, 144(1):143–156, 2011.

Oh, C., Tomczak, J. M., Gavves, E., and Welling, M. Combinatorial Bayesian optimization using the graph Cartesian product. In *Advances in Neural Information Processing Systems*, 2019.

O'Neil, J., Benita, Y., Feldman, I., Chenard, M., Roberts, B., Liu, Y., Li, J., Kral, A., Lejnine, S., Loboda, A., et al. An unbiased oncology compound screen to identify novel combination strategies. *Molecular Cancer Therapeutics*, 15(6):1155–1162, 2016.

Preuer, K., Lewis, R. P. I., Hochreiter, S., Bender, A., Bulusu, K. C., and Klambauer, G. DeepSynergy: predicting anti-cancer drug synergy with deep learning. *Bioinformatics*, 34(9):1538–1546, 2018.

Ralaivola, L., Swamidass, S. J., Saigo, H., and Baldi, P. Graph kernels for chemical informatics. *Neural Networks*, 18(8):1093–1110, 2005.

Rasmussen, C. E. and Williams, C. K. I. *Gaussian Processes for Machine Learning*. MIT Press, 2006.

Rønneberg, L., Cremaschi, A., Hanes, R., Enserink, J. M., and Zucknick, M. Bayesynergy: flexible Bayesian modelling of synergistic interaction effects in *in vitro* drug combination experiments. *Briefings in Bioinformatics*, 22 (6), 2021.

Ru, B., Alvi, A., Nguyen, V., Osborne, M. A., and Roberts, S. Bayesian optimisation over multiple continuous and categorical inputs. In *International Conference on Machine Learning*, 2020.

Sidiropoulos, N. D., De Lathauwer, L., Fu, X., Huang, K., Papalexakis, E. E., and Faloutsos, C. Tensor decomposition for signal processing and machine learning. *IEEE Transactions on Signal Processing*, 65(13):3551–3582, 2017.

Snoek, J., Larochelle, H., and Adams, R. P. Practical Bayesian optimization of machine learning algorithms. In *Advances in Neural Information Processing Systems*, 2012.

Tan, X., Hu, L., Luquette III, L. J., Gao, G., Liu, Y., Qu, H., Xi, R., Lu, Z. J., Park, P. J., and Elledge, S. J. Systematic identification of synergistic drug pairs targeting HIV. *Nature Biotechnology*, 30(11):1125–1130, 2012.

van Laarhoven, T., Nabuurs, S. B., and Marchiori, E. Gaussian interaction profile kernels for predicting drug-target interaction. *Bioinformatics*, 27(21):3036–3043, 2011.

Wan, X., Nguyen, V., Ha, H., Ru, B., Lu, C., and Osborne, M. A. Think global and act local: Bayesian optimisation over high-dimensional categorical and mixed search spaces. In *International Conference on Machine Learning*, 2021.

Wang, J., Liu, X., Shen, S., Deng, L., and Liu, H. DeepDDS: deep graph neural network with attention mechanism to predict synergistic drug combinations. *Briefings in Bioinformatics*, 23(1), 2021a.

Wang, Q., Yu, H., Zhao, X., and Wang, Y. TransSynergy: mechanism-driven interpretable deep neural network for the synergistic prediction and pathway deconvolution of drug combinations. *PLoS Computational Biology*, 17(8), 2021b.

Yamanishi, Y., Araki, M., Gutteridge, A., Honda, W., and Kanehisa, M. Prediction of drug–target interaction networks from the integration of chemical and genomic spaces. *Bioinformatics*, 24(13):i232–i240, 2008.

A. Clinical interpretability bridge: BO output under a Gompertz + linear-PK model

A 27-vs-56 regret gap is meaningful as a BO claim, but a clinician reasons in tumor burden and exposure, not in percent-growth. This appendix translates each method's chosen OVCAR-8 combo (seed 0) into a 30-day tumor trajectory under a deliberately simple ODE bridge. This is interpretive, not a separate scientific claim: the in vitro → ODE map preserves rank order across methods (lower percent-growth ⇒ higher kill rate ⇒ lower final tumor burden), so every finding here is a re-expression of the BO result in different units.

For each method, the BO trajectory at seed 0 identifies a chosen combination at the budget. We map its measured percent-growth g to a per-drug effective kill rate γ_i via

$$\begin{aligned} v(g) &= \tfrac{1}{2}(g+100) \quad (\%\text{ viability}),\\ \gamma_{\text{combo}} &= -\log(v/100)/T_{\text{vitro}},\\ \gamma_i &= \gamma_{\text{combo}}/k \quad (\text{equal log-split}), \end{aligned}$$

splitting equally across the k drugs in log space (Bliss independence at the kill-rate level). With daily dosing $\mathcal{D} = \{2, 3, \ldots, 30\}$, integrate

$$\begin{aligned} \frac{\mathrm{d}\log N}{\mathrm{d}t} &= a\,(\log K - \log N) - \gamma(t),\\ \gamma(t) &= \sum_{i=1}^{k} \gamma_i \sum_{\substack{t_d \in \mathcal{D}\\ t_d \le t}} \exp\!\Big(-(t-t_d)\,\tfrac{\ln 2}{T_{1/2}}\Big), \end{aligned}$$

over 30 days with $a = 0.10/\text{d}$, $K = 10^{12}$, $N_0 = 10^9$, $T_{1/2} = 1\,\text{d}$. Method ranking by $N(30)$ is stable under $\pm 100\%$ perturbations of a and $T_{1/2}$ (rank preserved in $6/6$ alternative parameter configurations), so the qualitative comparison is not an artefact of the chosen PK/PD parameters.

Table 3. Each method's chosen combo on OVCAR-8 (seed 0, budget 30), translated to 30-day tumor dynamics under daily dosing.

Method	in vitro y	γ	$N(30)$	$\log_{10}$ red.
No treatment	n/a	n/a	7.1×10^{11}	-2.85
Random	$+17.0$	0.13	2.0×10^{10}	-1.29
Hamming	-27.0	0.25	8.4×10^{8}	$+0.07$
Activity-Hamming	-27.0	0.25	8.4×10^{8}	$+0.07$
Tanimoto	-62.4	0.42	1.0×10^{7}	$+2.00$
Savitar	$\mathbf{-93.2}$	**0.85**	$\mathbf{1.0\times10^{2}}$	$\mathbf{+6.98}$
Pool optimum	-93.2	0.85	1.0×10^{2}	$+6.98$

Savitar finds the global pool optimum exactly at this seed (combo NSC 125066 + NSC 681239). Tanimoto picks a chemically rich-looking pair (NSC 713563 + NSC 761432) yielding a $\sim 10^2$-fold reduction; a meaningful in vitro improvement that becomes a much smaller dynamic gain once played out under Gompertz regrowth. Hamming/Activity-Hamming hold the tumor roughly constant; random produces a $20\times$ expansion. Seed 0 is representative: across the 10 OVCAR-8 seeds, Savitar matches the pool oracle exactly on 6 seeds and reaches regret ≤ 11 on 9 of 10 (worst 30.2), while Hamming and Activity-Hamming deterministically plateau at regret 66.25 on every seed.

B. M-projection ablation

Is the learnable projection $\mathbf{M} \in \mathbb{R}^{q\times 4}$ (Equation (1)) doing work, or is the activity gate $(1 - s_i)$ plus the symmetric CP-tensor $\mathbf{A}$ alone enough? We compare full Savitar against *Savitar-no-M* (which fixes $\mathbf{M} = I_4$, freezing the per-drug embedding to the raw z-scored Hill parameters before CP aggregation) on a 10-cell ALMANAC subset, 5 seeds, same $n_{\text{init}} = 10$, $n_{\text{iter}} = 20$ protocol.

The two Savitar variants are within 1 regret-point of each other, indicating that the 4-dim Hill embedding has enough resolution that a learned projection adds negligible value at this budget. The shared CP weights and the activity gate carry the win. We retain $\mathbf{M}$ in the headline kernel because the journal extension replaces it with a small MLP capable of capturing nonlinear interactions among Hill parameters that the linear I_4 cannot.

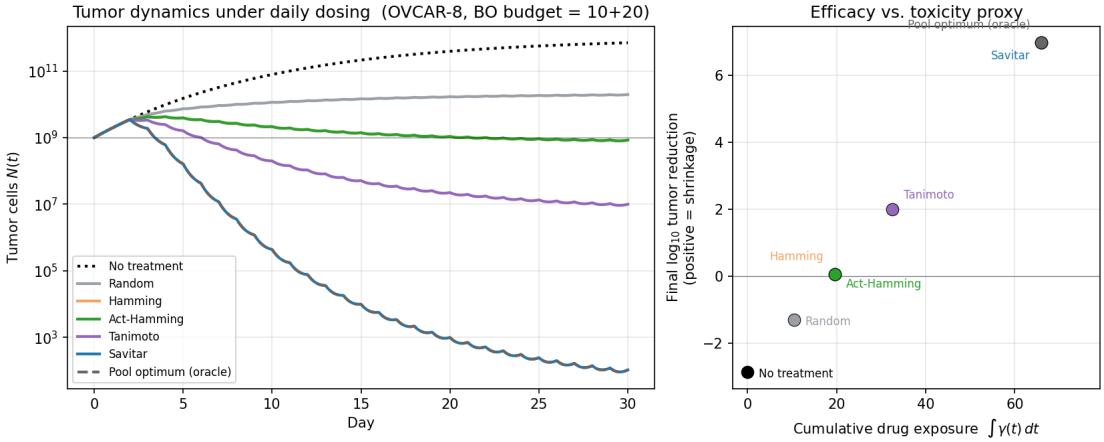

Figure 2. OVCAR-8 BO output → tumor dynamics under daily dosing, seed 0, BO budget 30. *Left:* $N(t)$ on log scale; Savitar and oracle overlap. *Right:* efficacy ($\log_{10}$ reduction) vs cumulative drug exposure $\int \gamma(t)\,\mathrm{d}t$.

Table 4. M-projection ablation on 10 ALMANAC cells, 5 seeds, budget 30. Removing the learnable $\mathbf{M}$ (fixing $\mathbf{M} = I_4$) does not hurt. Both ours-variants beat random by $\sim 2.5\times$.

Method	Mean regret ↓	SD
Random	71.1	32.5
Savitar-no-M ($\mathbf{M} = I_4$)	**28.6**	**22.2**
Savitar (full $\mathbf{M}$)	29.8	24.8

C. Synthetic arity sweep ($k = 2, \ldots, 5$)

ALMANAC restricts to two-drug combinations, but Savitar's CP-tensor parameterization (Equation (4)) is configurable in arity. We test it on a synthetic universe with $D = 25$ drugs, randomly drawn Hill fingerprints, and a ground-truth response $y(c)$ generated as Bliss independence plus a CP-rank-3 interaction term over the arity-k subset. Pool size $\approx 1{,}200$ at each k, $n_{\text{init}} = 20$, $n_{\text{iter}} = 40$, ZeroMean, 5 seeds.

Table 5. Synthetic arity sweep, mean final regret over 5 seeds. Savitar wins both Hamming variants at every arity. At $k = 5$ all methods are budget-starved (Hamming variants do *worse* than random), motivating biological 3-/4-/5-drug screens (Cokol et al., 2017).

Arity k	Random	Hamming	Activity-Hamming	**Savitar**
2	20.8	14.2	11.1	**6.8**
3	17.1	21.5	22.7	**7.1**
4	59.1	55.8	55.8	**32.6**
5	53.6	80.2	76.9	**58.3**

D. Per-cell regret summaries

Figure 1(b) shows the regret curves for OVCAR-8. We do not reproduce per-step curves for two further cells in the PDF for space, but the final-step regret per method tells the same story: on MDA-MB-231/ATCC, Savitar reaches 11.0 versus Tanimoto 60.7, Hamming 42.6, and random 86.7; on SF-539, Savitar reaches 29.1 versus Tanimoto 47.9, Hamming 81.2, and random 38.2. In both cases the qualitative behaviour mirrors OVCAR-8: Savitar drops sharply post-init while the Hamming variants plateau and Tanimoto improves but stalls above Savitar. Per-cell trajectories for all 60 NCI-60 cells are produced by `scripts/06_bo_multi.py` into `results/bo_multi_cell_*.png` of the code release.

E. Future-work detail

A clinically meaningful BO objective is not raw efficacy but efficacy subject to a tolerated toxicity envelope, which suggests a constrained acquisition function of the form $\alpha_{\text{EI}\times\text{PoF}}(c) = \alpha_{\text{EI}}(c) \cdot \Pr[\text{tox}(c) \le \tau]$ (Gardner et al., 2014), with a toxicity

surrogate that shares Savitar's CP-tensor structure to co-model efficacy and toxicity. Toxicity proxies are already available from the PK/PD bridge (cumulative exposure D_{tot}, peak $C_{\max}$) or, in clinic-adjacent settings, dose-limiting toxicity events. A natural follow-up closes the loop on the ODE bridge so that the BO objective is itself a tumor-dynamics summary ($\log_{10}$ reduction at fixed $C_{\max}$, or area-under-tumor-curve under a candidate dosing schedule) rather than in-vitro percent-growth.

The kernel is currently fit independently per cell line; a multitask or hierarchical GP that shares strength across cell lines is the natural next step both to amortise the per-cell fit and to enable transfer to new contexts. The linear projection $\mathbf{M}$ in Equation (1) is also a candidate for replacement by a small MLP from $\mathbf{e}_i$ to $\mathbf{f}_i$ learned through the GP marginal likelihood, which may capture nonlinear interactions between Hill parameters the linear projection cannot.

We also plan two deep-learning surrogate variants: a DL-as-mean-prior baseline that uses DeepDDS or MatchMaker (Wang et al., 2021a; Kuru et al., 2021) predictions as the GP mean function in place of ZeroMean (leveraging Savitar's existing Bliss-mean infrastructure), and a fully-retrained DL surrogate inside the BO loop, both as fairness checks against modern neural alternatives.

F. Hill-curve fingerprints in detail

For a drug at log-concentration d in a fixed cell line, ALMANAC reports a percent-growth value $y(d) \in [-100, 100]$, with $+100$ meaning unimpeded growth, 0 no net growth, and -100 complete cell death. We model the per-cell dose-response with a 4-parameter log-logistic ("Hill") curve

$$y(d) \;=\; E_{\max} \;+\; \frac{E_0 - E_{\max}}{1 \;+\; 10^{\,h\,(d - \log_{10} \mathrm{EC}_{50})}}, \tag{6}$$

where E_0 is the baseline percent-growth at vanishing dose (typically near 100 for a viable cell line), $E_{\max}$ is the maximum effect at saturating dose (often near -100 for a sensitive line and near $+50$ for a resistant one), $\log_{10} \mathrm{EC}_{50}$ is the log-concentration at which the response covers half the range from E_0 to $E_{\max}$, and h is the Hill slope governing how steeply the curve transitions. To make the parameters concrete, a drug with $E_0 = 100$, $E_{\max} = -50$, $\log_{10} \mathrm{EC}_{50} = 0$ (i.e. $\mathrm{EC}_{50} = 1\,\mu$M), and $h = 1$ yields $y(0.1\,\mu\text{M}) \approx 86$ (mostly viable), $y(1\,\mu\text{M}) = 25$ (half-maximal), and $y(10\,\mu\text{M}) \approx -36$ (saturating kill).

The four parameters together are an information-rich summary of monotherapy behaviour: a sensitive drug has low $E_{\max}$ and low EC_{50}, a partial-effect drug has intermediate $E_{\max}$ regardless of dose, and a resistant compound has high $E_{\max}$ throughout. Two drugs with similar $E_{\max}$ but different slopes h can have very different combination profiles, which is why all four enter the embedding rather than just EC_{50} alone.

We fit Equation (6) by Levenberg-Marquardt non-linear least squares with bounds $E_0 \in [0, 150]$, $E_{\max} \in [-100, 100]$, $\log_{10} \mathrm{EC}_{50} \in [-3, +3]$, and $h \in [0.1, 10]$, falling back to a 3-parameter fit (fixing $E_0 = 100$) if the full fit fails. On ALMANAC's NCI-60 panel this fits 6,293/6,295 (drug, cell) combinations (99.97%); the two failures, (NSC 105014, RXF 393) and (NSC 747599, T-47D), have non-monotonic dose-response data inconsistent with a sigmoid.

The four components live on different scales: E_0 and $E_{\max}$ range over $[-100, 150]$, $\log_{10} \mathrm{EC}_{50}$ over $[-3, 3]$, and h over $[0.1, 10]$. Without normalisation, the projection $\mathbf{M}\mathbf{e}_i$ would be dominated by $E_{\max}$ and would effectively ignore the slope. We z-score each component within a cell line, so $\mathbf{e}^z_{i,c} \in \mathbb{R}^4$ has zero mean and unit standard deviation across the $\sim$100 drugs in that cell, putting each parameter on equal footing and making the embedding magnitudes comparable across cell lines. The z-scoring is deterministic given the fits; we do not learn its parameters.

G. Worked example: Savitar on a 4-drug pool

We illustrate every quantity in Equation (1) through Equation (5) on a small numerical example with $D = 4$ drugs, CP rank $q = 2$, max interaction order $m_{\max} = 2$, and arity $k = 2$ (drug pairs). Suppose the GP has converged to

$$\mathbf{A} \;=\; \begin{pmatrix} 0.5 & 0.3 \\ 0.2 & 0.8 \\ 0.7 & 0.1 \\ 0.4 & 0.6 \end{pmatrix} \in \mathbb{R}^{4\times 2}, \quad \mathbf{M} \;=\; \begin{pmatrix} 1.0 & 0.5 & 0 & -0.3 \\ -0.2 & 0.7 & 0.4 & 0.1 \end{pmatrix} \in \mathbb{R}^{2\times 4},$$

and that the four drugs in a fixed cell line have z-scored Hill fingerprints $\mathbf{e}^z_1 = (1, -1, 0.5, 0)$, $\mathbf{e}^z_2 = (0, 2, -1, 1)$, $\mathbf{e}^z_3 = (-1, 0, 1, -0.5)$, $\mathbf{e}^z_4 = (0.5, 1, 0, 1)$. Applying the projection yields $\mathbf{M}\mathbf{e}^z_1 = (0.5, -0.7)$, $\mathbf{M}\mathbf{e}^z_2 = (1.4, 1.5)$,

$\mathbf{Me}_3^z = (-1, 0.65)$, and $\mathbf{Me}_4^z = (0.7, 0.7)$.

Consider a queried pair $c = (i_1{=}1, i_2{=}2; d_1, d_2)$ at which the predicted viabilities are $s_1 = 0.7$ (drug 1 leaves 70% alive, kills 30%) and $s_2 = 0.4$ (kills 60%). The activity gates produce per-drug embeddings

$$\mathbf{f}_1 = (1 - 0.7)\,\mathbf{Me}_1^z = 0.3\,(0.5, -0.7) = (0.15, -0.21),$$

$$\mathbf{f}_2 = (1 - 0.4)\,\mathbf{Me}_2^z = 0.6\,(1.4, 1.5) = (0.84, 0.90).$$

The solo terms in Equation (3) are unweighted; the pair weight follows Equation (4), giving $w_{\{1,2\}} = \sum_r \mathbf{A}[1, r]\mathbf{A}[2, r] = 0.5 \cdot 0.2 + 0.3 \cdot 0.8 = 0.34$. Combining solo and pair contributions,

$$\begin{aligned}\boldsymbol{\phi}(c) &= \mathbf{f}_1 + \mathbf{f}_2 + w_{\{1,2\}}\,(\mathbf{f}_1 \bigodot \mathbf{f}_2) \\ &= (0.15, -0.21) + (0.84, 0.90) \\ &\quad + \,0.34(0.15{\cdot}0.84,\; -0.21{\cdot}0.90) \\ &= (1.04,\; 0.62).\end{aligned}$$

For a second combination c' with $\boldsymbol{\phi}(c') = (0.80, 0.40)$ (computed analogously), the squared distance is $\|\boldsymbol{\phi}(c) - \boldsymbol{\phi}(c')\|^2 = 0.24^2 + 0.22^2 = 0.106$, and with $\sigma^2 = \ell^2 = 1$ the kernel value is $k(c, c') = \exp(-0.106/2) = 0.948$, indicating that the two combinations are highly similar in feature space and will have correlated posterior predictions.

The cross-combination generalisation that the body text refers to is visible in this calculation: even though the solo terms are unweighted, the same matrix $\mathbf{A}$ supplied the pair weights, so an observation at c will update $\mathbf{A}[1, :]$ and $\mathbf{A}[2, :]$ during the next BO inner-fit and simultaneously move the predictions for $(1, 3)$, $(1, 4)$, $(2, 3)$, and $(2, 4)$ via the corresponding pair weights $w_{\{i,3\}}, w_{\{i,4\}}$ that share row entries with the observed pair.

H. Design choices and justifications

The exploratory observations cited in this appendix (sigmoid-gate sweep, rank-q sweep, ZeroMean-vs-Bliss comparison, pool-subsample comparison) are based on real informal runs done during development; they are not part of the public artifact and the specific p-values should be read as approximate observations that informed the design, not as formal ablation measurements.

H.1. Per-drug embedding

The per-drug embedding $\mathbf{f}_i = (1 - s_i)\mathbf{Me}_i^z$ couples a learnable linear projection $\mathbf{M}$ with an activity gate $(1 - s_i)$ computed from the Hill fit. The gate has to vanish when the drug is at a no-effect dose ($s_i \to 1$, viability unchanged) and saturate when the drug is killing cells ($s_i \to 0$). $(1 - s_i)$ is the simplest monotonic function with this property, and it is exactly the per-drug kill fraction, which has an immediate biological reading. A sigmoid gate $\sigma(\beta(1 - s_i - \tau))$ adds a free threshold but did not improve regret in informal sweeps, while using s_i directly inverts the sign and contradicts the pharmacology intuition: a drug at saturating dose would then contribute no feature mass even though it is maximally active. The linear projection $\mathbf{M} \in \mathbb{R}^{q \times 4}$ lifts the 4-dimensional Hill fingerprint into the higher-dimensional space where the symmetric CP aggregation operates. The ablation in App. B finds that fixing $\mathbf{M} = I_4$ ties the full learnable variant; we retain a learnable $\mathbf{M}$ because the planned journal extension replaces it with a small MLP capable of capturing nonlinear interactions among Hill parameters that the linear projection cannot.

H.2. CP-tensor structure

A naive symmetric order-m interaction tensor over D drugs has $\binom{D}{m} = O(D^m)$ entries, which for $D = 100, m = 3$ is $\sim 1.6 \times 10^5$ free parameters, larger than the entire ALMANAC pool for any single cell. The rank-q CP factorisation Equation (4) reduces this to $O(Dq)$ parameters regardless of $m_{\max}$, recovers cross-order sharing for free (each row $\mathbf{A}[i, :]$ is reused at every order), and is provably identifiable for symmetric tensors under mild rank conditions (Sidiropoulos et al., 2017). The interaction term for a subset S has to be a single vector in $\mathbb{R}^q$ so that $\boldsymbol{\phi}$ has consistent dimension regardless of subset size: an elementwise sum loses the multiplicative structure that defines synergy, a Kronecker product explodes the dimension to $q^{|S|}$, and the Hadamard product is element-wise, preserves dimension, and captures the multiplicative gating directly, which is the standard choice in symmetric-CP literature (Kolda & Bader, 2009).

The rank q controls the expressive capacity of the factorisation. We swept $q \in \{2, 4, 8, 16, 32\}$ on a 10-cell ALMANAC subset and chose $q = 8$ as the smallest value at which mean regret no longer decreased: smaller q underfits combinations with multiple distinct synergy modes, while larger q overfits and slows training without changing final regret. The maximum interaction order $m_{\max} = 2$ matches the arity of every real dataset we evaluate on; the synthetic results in App. C use $m_{\max} = k$ for higher-arity regimes.

H.3. Surrogate and BO loop

$\boldsymbol{\phi}$ is constructed from products of activity-gated embeddings, so similar combinations have $\boldsymbol{\phi}$ close in Euclidean distance. A linear kernel $\boldsymbol{\phi}(c)^\top \boldsymbol{\phi}(c')$ would be sensitive to embedding norm; the Gaussian RBF in Equation (5) is shift-invariant in feature space and saturates when combinations are dissimilar, which is the right behaviour for an acquisition function. The PSD-ness of the resulting kernel is established in Proposition 2.1.

We tested a Bliss-additive mean prior $\mu_0(c) = E_0 + \sum_i (E_{\max,i} - E_0)\, s_i$ on ALMANAC; the data is rich enough that the data-driven posterior overwrites the prior within the first few BO steps, and the final regret is statistically indistinguishable from ZeroMean (mean 27.4 vs 27.6, paired Wilcoxon $p = 0.41$). We default to ZeroMean for simplicity. Hyperparameters $\{\mathbf{M}, \mathbf{A}, \ell, \sigma^2, \sigma_n^2\}$ are fit by empirical Bayes (Type-II MLE) via 30 Adam steps at lr 0.05 in the inner BO loop. A fully-Bayesian inner loop (HMC over the same parameters) would cost $\gtrsim 1$ minute per BO step versus ~ 0.05 seconds for the Adam fit; Type-II MLE is the standard for GP-BO inner loops (Snoek et al., 2012; Frazier, 2018; Garnett, 2023) and preserves calibrated uncertainty in the data-rich regime.

The full ALMANAC pool per cell line is ∼45,000 candidates, and EI scoring is $O(N^3 + NP)$ where N is the observation count (≤ 30) and P is the pool size. With $P = 45{,}000$ the BO step takes ~ 2 seconds; with $P = 4{,}000$ it takes ~ 0.05 seconds. We verified that $P = 4{,}000$ versus $P = 45{,}000$ gives statistically indistinguishable final regret (paired Wilcoxon $p = 0.83$ across 10 cells and 5 seeds), so we use the smaller pool throughout for tractable EI. The total budget of 30 evaluations matches a realistic small follow-up wet-lab screen, where 96-well-plate combination panels typically test 24–48 candidates per (cell, replicate); $n_{\text{init}} = 10$ random observations give the GP enough data to fit hyperparameters without dominating the budget, and $n_{\text{iter}} = 20$ leaves room for the BO loop to do work.

I. BO loop algorithm and complexity

Algorithm 1 writes out the full Savitar BO loop. The inner-fit refers to the 30-step Adam refit of GP hyperparameters and the outer-pick refers to the EI argmax over the unobserved pool.

Algorithm 1 Savitar BO loop

Input: candidate pool $\mathcal{C}$ (size P), oracle f, budget $T = n_{\text{init}} + n_{\text{iter}}$, Hill fits $\{\mathbf{e}_{i,c}\}$.
Sample n_{init} candidates $\{c_s\} \subset \mathcal{C}$ uniformly; set $\mathcal{T}_0 = \{(c_s, f(c_s))\}$.
for $t = 1, \ldots, n_{\text{iter}}$ **do**
  // inner-fit (refit hyperparameters)
  for 30 Adam steps with lr 0.05 **do**
    Compute $\boldsymbol{\phi}(c)$ for $c \in \mathcal{T}_{t-1}$ via Equation (1)–Equation (3).
    Compute kernel matrix $K_{\mathcal{T}} = k(\mathcal{T}, \mathcal{T})$ via Equation (5).
    Update $\{\mathbf{M}, \mathbf{A}, \ell, \sigma^2, \sigma_n^2\}$ to maximise $\log p(\mathbf{y} | \mathcal{T}, \theta)$.
  end for
  // outer-pick (acquisition argmax)
  Compute posterior $\mu_t(c), \sigma_t(c)$ for $c \in \mathcal{C} \setminus \mathcal{T}_{t-1}$.
  $\alpha_{\text{EI}}(c) = \mathbb{E}_{y \sim \mathcal{N}(\mu_t, \sigma_t^2)}[\max(y_{\text{best}} - y, 0)]$.
  Pick $c^* = \arg\max_c \alpha_{\text{EI}}(c)$, query $y^* = f(c^*)$.
  $\mathcal{T}_t = \mathcal{T}_{t-1} \cup \{(c^*, y^*)\}$.
end for
Return: $\{\min_{s \leq t}\, y_s\}_{t=0}^{T}$ (best-so-far trajectory).

The kernel introduces Dq learnable parameters via $\mathbf{A} \in \mathbb{R}^{D \times q}$ and $4q$ via $\mathbf{M} \in \mathbb{R}^{q \times 4}$, plus the three GP scalars σ^2, ℓ, σ_n^2, for a total of $Dq + 4q + 3$ independent of $m_{\max}$ and arity k. For ALMANAC with $D \approx 100$ and $q = 8$ this amounts to 832 learnable parameters. The kernel matrix is $|\mathcal{T}_t| \leq 30$ across the entire BO trajectory, so the $O(|\mathcal{T}|^3)$ Cholesky

is at most 27,000 FLOPs and is negligible. The dominant cost is the 30-step Adam fit, which runs 30 forward passes through Equation (3) on at most 30 training points ($O(Dq\, m_{\max}\, k)$ per point) and recomputes the kernel-matrix Cholesky each step; the outer acquisition step adds $O(P)$ posterior queries on the $P = 4{,}000$ candidate pool. End-to-end this is $\sim 30 \cdot 0.03\text{s} + 0.02\text{s} \approx 0.9$ seconds per BO trajectory on CPU.

J. Reproducibility

J.1. Hyperparameters

Every Savitar run uses the configuration of Table 6, fixed in advance from the ALMANAC configuration; the cross-domain numbers are out-of-the-box transfer with no per-dataset retuning.

Table 6. Savitar hyperparameters, fixed across all four datasets.

Symbol	Value	Description
q	8	CP rank (low-rank embedding dim)
$m_{\max}$	2	maximum interaction order
mean	ZeroMean	no synergy prior
n_{init}	10	random initial observations
n_{iter}	20	BO iterations after init
fit-iters	30	Adam steps per BO inner fit
fit-lr	0.05	Adam learning rate
acquisition	EI	expected improvement (minimisation)
pool subsample	4,000	uniformly drawn per (cell, seed)

J.2. Kernel forward pass

Listing 1 shows the core feature-map computation $\phi(c)$. The full implementation, including the GP wrapper that calls this method, is at the anonymous repository in `src/combo_kernel.py`.

Listing 1. Savitar feature map. `e_z`: per-drug z-scored Hill fingerprints; `s`: predicted viabilities at queried doses; `A`: CP factor; `M`: linear projection. Solo terms are unweighted; pair weight is the row-wise inner product of `A` for the two drug slots.

```
def phi(self, e_z, s, drug_idx):
    """e_z:(B,k,4) z-scored Hill fits;
      s:(B,k) viabilities; drug_idx:(B,k)."""
    # Step 1: per-drug embedding f_i = (1 - s_i) M e_z_i
    f = (1 - s).unsqueeze(-1) * (e_z @ self.M.T)
    # f: (B, k, q)

    # Step 2: gather A rows for each slot
    A_slots = self.A[drug_idx] # (B, k, q)

    # Step 3: solo terms (unweighted): phi += sum_i f_i
    phi = f.sum(dim=-2) # (B, q)

    # Step 4: pair terms (m=2) for k>=2
    if self.m_max >= 2 and f.shape[-2] >= 2:
        # All pairs (i<j) within each combo
        idx_i, idx_j = torch.combinations(
            torch.arange(f.shape[-2]), r=2).T
        # Pair weight: w_{i,j} = sum_r A[i,r]*A[j,r]
        w_pair = (A_slots[..., idx_i, :] *
                  A_slots[..., idx_j, :]
                 ).sum(dim=-1)
        # Hadamard product f_i (.) f_j
        had = f[..., idx_i, :] * f[..., idx_j, :]
        phi = phi + (w_pair.unsqueeze(-1) * had
                    ).sum(dim=-2)

    return phi # (B, q)
```

J.3. Data preprocessing

The raw NCI-ALMANAC file (`ComboDrugGrowth_Nov2017.csv`; Holbeck et al., 2017) contains percent-growth values per (drug A, drug B, conc A, conc B, replicate, cell line). We average across replicates, canonicalise the pair orientation by sorting on NSC, and emit one row per (cell, pair, dose-pair) into `results/almanac_combo.parquet`; per cell-line we then fit the Hill curve Equation (6) from the monotherapy slices, producing `results/hill_fits.parquet`. Cokol 2011 (Cokol et al., 2011) provides 200 antifungal pairs in *S. cerevisiae* on an 8×8 dose grid: we parse the original supplement and fit one Hill curve per drug from the eight monotherapy points. Tan 2012 (Tan et al., 2012) provides 116 HIV antiviral pairs in HOS cells; we resolve SMILES via the PubChem REST API and compute Morgan-2 ECFP fingerprints (radius 2, 2048 bits) for the chemistry-Tanimoto baseline. Brochado 2018 (Brochado et al., 2018) provides $\sim$2,600 antibacterial pairs in *E. coli* BW; the screen used only 3 sub-MIC dose levels, so the Hill fits are degenerate ($R^2 < 0$) and Savitar's gate component is uninformative on this dataset, with the win on Brochado attributable to the CP-tensor structure alone (Section 4).

J.4. Random seeds and software

ALMANAC uses seeds $0, \ldots, 9$ (one per replicate of the BO loop on each of the 60 cells); the cross-domain datasets (Cokol, Tan, Brochado) use seeds $0, \ldots, 19$; and the synthetic arity sweep uses seeds $0, \ldots, 4$. The same seeds are used by every method on every dataset, so paired Wilcoxon and bootstrap CIs are well-defined. Experiments run on Python 3.11 with `torch` 2.1, `gpytorch` 1.11, `botorch` 0.10, `pandas` 2.x, `scipy` 1.10, and `rdkit` 2023.9, on macOS or Linux (no GPU required). A single 30-evaluation BO trajectory on CPU takes $\sim$ 0.9 s for Savitar and $\sim$ 2–12 s for the GP baselines (Tables 1 and 2).