SPIN-BO: SOURCE-GUIDED PROTEIN–PROTEIN INTERACTION DISCOVERY FOR NEW TARGETS

Anonymous authors
Paper under double-blind review

ABSTRACT

Protein engineering often requires optimizing a common mutational library against multiple interaction partners, but experimentally mapping a complete interaction landscape for every new partner is costly. Existing landscapes from related partners offer valuable prior knowledge, yet transferring this knowledge is challenging because shared mutational trends coexist with partner-specific effects. We introduce SPIN-BO, a source-guided Bayesian optimization framework for discovering high-response variants against a new interaction partner under a limited target-assay budget. SPIN-BO reuses previous interaction landscapes through two complementary transfer mechanisms. First, an adaptively calibrated source-response mean uses sparse target assays to learn how the source profiles should be weighted and combined for the new partner. Second, a source-supervised protein-pair representation, learned from pretrained protein language model embeddings, defines the geometry of a target-specific Gaussian-process residual, enabling sparse target measurements to generalize across the mutant library while capturing target-specific structure not explained by transferred source patterns. As target assays accumulate, SPIN-BO continually recalibrates how source landscapes contribute to prediction, updates predictive uncertainty, and selects the next variants to test. Controlled semi-synthetic experiments show that response-level and representation-level transfer provide complementary benefits across different source–target relationships. Across 54 retrospective discovery tasks from the human bZIP interaction network, SPIN-BO achieves a mean optimum-success rate of 95.0% at a budget of 200 target assays, compared with 73.3% for target-only Bayesian optimization. These results show that previously measured interaction landscapes can be systematically reused to accelerate discovery for new interaction partners.

1 INTRODUCTION

Protein–protein interactions (PPIs) connect molecular recognition to cellular function, and many protein-engineering campaigns seek to reprogram one side of an interaction while holding the other partner fixed. In a common experimental setting, a partner protein is challenged with a library of variants of the interacting protein, with the goal of identifying mutations that improve affinity, specificity, or another interaction-related phenotype. Such campaigns can produce substantial functional gains: for example, engineering the natural protease inhibitor N-TIMP2 against MMP-14 yielded variants with nearly 900-fold stronger affinity and up to 16,000-fold greater specificity over related metalloproteinases (Arkadash et al., 2017). Yet mutational effects are inherently partner dependent. Related partners may share broad mutational preferences, while individual substitutions can have markedly different effects across interaction contexts. Quantitative N-TIMP2 interaction landscapes against homologous MMPs illustrate this combination of shared mutational patterns and partner-specific effects (Aharon et al., 2020). Such interaction maps are highly informative, but comprehensive screening across mutant–partner pairs is experimentally intensive. This raises a natural transfer question: when complete interaction profiles are already available for a common mutant library against one or more related partners, can they be reused to accelerate discovery for a new partner without assuming that the new partner follows the same mutational landscape?

We study this as a *cross-partner discovery* problem, in which a common mutant library has been comprehensively assayed against source interaction partners, while only a small number of sequen-

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

tial assays can be collected for a new target partner. Addressing this problem requires a search strategy that can exploit reusable source knowledge while continually adapting its relevance to the new partner as target measurements accumulate. Bayesian optimization (BO) provides an efficient framework for such expensive experimental search by using a probabilistic surrogate to decide which variants to measure next (Frazier, 2018). In protein engineering, pretrained protein language models (PLMs) provide informative representations of sequence variation (Alley et al., 2019), while recent active learning systems have demonstrated the value of iterative model-guided experimentation (Yang et al., 2025). These advances, however, do not directly address how information from related interaction landscapes should be incorporated when optimizing a new target. Standard target-specific optimization largely treats each new partner as a separate search problem, leaving available source interaction landscapes unused. At the other extreme, directly transferring source-response patterns can be unreliable because broad mutational trends may be shared while individual variant effects remain partner specific. This creates two complementary challenges: determining how source-response patterns should inform prediction for the new partner, and learning the remaining target-specific structure from only a limited number of target measurements.

To address these two challenges, we introduce **SPIN-BO** (**S**ource-Guided **P**rotein–Protein **I**nteraction Discovery for **N**ew Targets), a source-guided Bayesian optimization framework that transfers information from existing interaction landscapes at both the response and representation levels. At the response level, an adaptively calibrated source-response mean uses sparse target assays to learn how the source profiles should be weighted and combined for the new partner. At the representation level, a source-supervised protein-pair representation, learned from pretrained PLM embeddings, defines the geometry of a target-specific Gaussian-process residual, enabling sparse target measurements to generalize across the mutant library while capturing target-specific structure not explained by transferred source patterns. As new target assays are collected, SPIN-BO continually recalibrates its use of source information, updates predictive uncertainty, and uses the resulting surrogate to guide the next batch of experiments.

Our main contributions are:

- We formulate assay-efficient optimization for a new interaction partner as a *cross-partner discovery* problem, where complete interaction landscapes are available for several source partners over a common mutant library, while only a limited number of target assays can be acquired sequentially for the new partner.
- We introduce SPIN-BO, a source-guided Bayesian optimization framework that transfers cross-partner information at both the response and representation levels. As target assays accumulate, SPIN-BO adaptively re-estimates the source-response weights to learn which source landscapes remain relevant and how they should be combined, while source supervision provides an interaction-aware representation for target-specific residual modeling.
- We demonstrate that previously measured interaction landscapes can accelerate discovery for new partners and characterize when different forms of transfer are beneficial. Controlled semi-synthetic experiments reveal complementary roles of response-level and representation-level transfer, while retrospective evaluation on the human bZIP interaction network shows consistent improvements over target-only optimization.

2 Related Work

Protein engineering and cross-partner discovery. Protein engineering increasingly uses iterative, model-guided experimental design under limited assay budgets. Machine-learning-guided directed evolution learns sequence–function relationships from measured variants (Yang et al., 2019), while systems such as ALDE (Yang et al., 2025) and EVOLVEpro (Jiang et al., 2025) combine iterative selection with uncertainty-aware models or pretrained PLM representations. BioBO further incorporates multimodal biological representations and enrichment-derived priors into perturbation design (Li et al., 2026). These approaches improve sample-efficient optimization but largely treat each target landscape as a separate discovery problem. A complementary line of work transfers information across previously measured protein landscapes to predict variant effects or reconstruct fitness landscapes for new proteins (Gray et al., 2018; Jagota et al., 2023; Chen et al., 2023). We instead study a cross-partner discovery problem, where a common mutant library has been comprehensively measured against source interaction partners and the goal is to identify high-response

variants for a new target using only a small number of sequential assays. This objective differs fundamentally from landscape reconstruction: accurate recovery of an entire landscape is neither required nor necessarily sufficient for rapid discovery of its best variants.

Transfer Bayesian optimization across related tasks. Transfer BO reuses historical tasks in several ways. Multi-task BO models source and target functions through learned cross-task covariance (Swersky et al., 2013), while representation-based approaches learn features shared across optimization problems (Perrone et al., 2018); both transfer through learned task or feature structure rather than using aligned source-response profiles as explicit target predictors. RGPE adaptively weights fixed source surrogates according to agreement with target observations (Feurer et al., 2022), whereas BOLT transfers previous optimization experience primarily through initialization (Zeng et al., 2026). Our cross-partner setting provides distinct candidate-level structure: every mutant has a complete, aligned profile of experimentally measured source responses across partners, allowing SPIN-BO to reuse the measured source landscapes directly and continually recalibrate their contribution as target assays accumulate, rather than relying on transfer mediated through task covariance, learned representations, source surrogates, or initialization.

Adaptive source transfer and target-specific discrepancy. The response-level component of SPIN-BO is related to a broader statistical lineage of calibration and discrepancy modeling, which combines an auxiliary response with a residual process (Kennedy & O'Hagan, 2001). More directly, DeltaBO models a target as a source function plus a source–target difference process (Lin et al., 2026); the source contribution is fixed during target optimization, while new target observations update the discrepancy process. SPIN-BO instead makes the transferred response itself target-adaptive: it learns a target-specific combination of the measured source landscapes and repeatedly re-estimates their contributions as target assays accumulate, while a target-specific GP residual captures remaining structure. This adaptivity also differs from DynMeanBO, which places external expert knowledge in the GP mean and progressively attenuates its influence according to a prescribed decay schedule (Qu et al., 2026); SPIN-BO instead uses target observations to learn which source patterns remain relevant and how they should be combined.

3 Proposed Method

3.1 Problem Definition

Let $\mathcal{B} = \{B_m\}_{m=1}^{M}$ denote a common library of M protein mutants, $A_1^{\mathcal{S}}, \ldots, A_S^{\mathcal{S}}$ a set of source partners, and $A^{\mathcal{T}}$ a new target partner. For each source partner $A_s^{\mathcal{S}}$, a complete response profile $\{Y_{sm}^{\mathcal{S}}\}_{m=1}^{M}$ is available over the common mutant library, with $S \geq 1$. We assume that no target responses are observed before the campaign begins. The campaign first collects an initial set of target measurements, after which additional measurements are acquired adaptively in batches.

For mutant B_m, let $f_m = f(A^{\mathcal{T}}, B_m)$ denote its latent interaction response with the target partner. We orient the response so that larger values correspond to more desirable interactions. Querying B_m reveals

$$Y_m^{\mathcal{T}} = f_m + \varepsilon_m, \qquad \varepsilon_m \sim \mathcal{N}(0, \sigma_\varepsilon^2),$$

with independent observation noise. At any stage of the campaign, let $O \subseteq \{1, \ldots, M\}$ denote the indices of mutants already measured against the target and $U = \{1, \ldots, M\} \setminus O$ the remaining candidate pool. Mutants are selected from U in batches, and their responses are observed only after all mutants in the batch have been assayed.

The global optimum is the highest-response mutant in the library,

$$m^\star = \arg \max_{m \in \{1,\ldots,M\}} f(A^{\mathcal{T}}, B_m) = \arg \max_{m \in \{1,\ldots,M\}} f_m.$$

Our objective is *budgeted discovery*: given a total target-assay budget B_{tot}, identify $m^\star$ and other high-response variants using only a limited number of target measurements. Throughout, the source response profiles are treated as available source knowledge and do not count toward the target-assay budget.

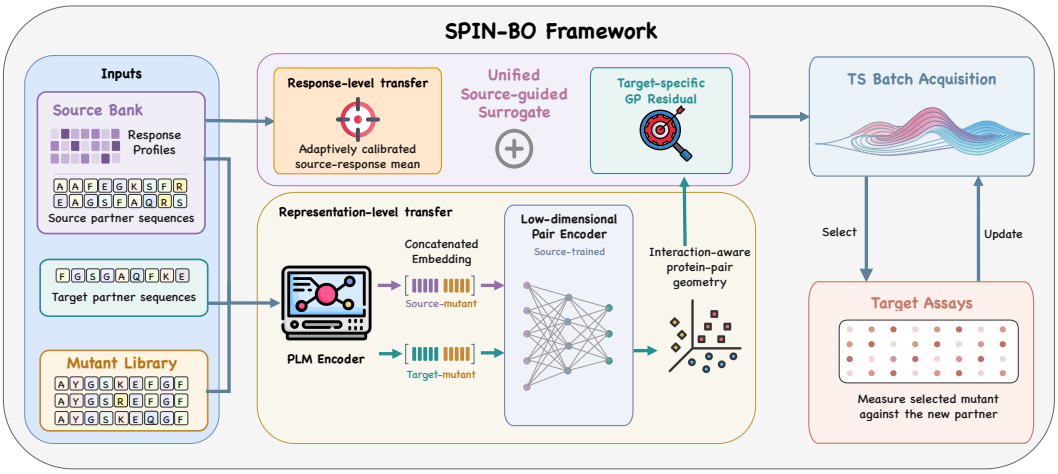

Figure 1: **Overview of SPIN-BO for source-guided PPI discovery.** Source interaction landscapes are reused through two complementary transfer mechanisms: response-level transfer via an adaptively calibrated source-response mean, and representation-level transfer via a source-supervised protein-pair geometry. These components are combined in a unified source-guided surrogate with a target-specific GP residual. The surrogate drives batch Thompson-sampling acquisition, and newly observed target assays are used to update the model for the next round.

3.2 Method Overview

SPIN-BO leverages existing interaction landscapes through two complementary transfer mechanisms, corresponding to the two forms of information available from the source partners. At the *response level*, complete source response profiles provide candidate-specific information about how each mutant behaves across previously measured partners. Sparse target assays are used to adaptively calibrate how these profiles should be weighted and combined for the new partner. At the *representation level*, source partner–mutant interactions supervise a protein-pair representation built from PLM embeddings. This learned representation defines the geometry of a target-specific GP residual, allowing sparse target measurements to generalize across the mutant library.

The two transfer mechanisms are combined in a unified source-guided target surrogate: the calibrated source-response mean captures transferable response patterns, while the target-specific GP residual captures structure not explained by those transferred patterns. As new target assays are collected, the surrogate is recalibrated and its predictive uncertainty is updated. The resulting joint predictive distribution is then used to select the next batch of mutants to assay. Figure 1 summarizes the overall procedure.

3.3 Response-Level Transfer

The first transfer mechanism directly reuses the measured source interaction landscapes. For each mutant B_m, let $s_m = (Y^{\mathcal{S}}_{1m}, \ldots, Y^{\mathcal{S}}_{Sm})^\top$ denote its responses across the S source partners, and define $h_m = (1, s_m^\top)^\top$. We represent the transferred response-level information through an adaptively calibrated mean

$$a_{\mathcal{S}}(B_m; \beta) = h_m^\top \beta = \beta_0 + \sum_{s=1}^{S} \beta_s Y^{\mathcal{S}}_{sm}.$$

Instead of directly carrying a source landscape over to the new partner, the coefficients β are learned from the target assays collected so far and are recalibrated as new target measurements accumulate. This adaptivity allows the model to continuously adjust the relevance of each source profile to the new target. With multiple source partners, the model learns how their response profiles should be weighted and combined for the new interaction partner.

This mechanism transfers information at the response level: the influence of each source profile on target prediction adapts to its agreement with the target assays observed so far, allowing source

contributions to evolve as the campaign progresses rather than remain fixed. The calibration is performed jointly with the target-specific residual model introduced below. Additional details on the source-response design and initialization are provided in Appendix A.1.

3.4 Representation-Level Transfer

The second transfer mechanism uses source interaction data to learn how partner–mutant pairs should be represented. Each protein is first embedded using the pretrained ESM2-650M model (Lin et al., 2023), and the partner and mutant embeddings are concatenated to form the pair representation $r(A, B) = [e(A); e(B)]$. A low-dimensional encoder g_θ is then trained on source partner–mutant pairs using their measured interaction responses as supervision, refining the pretrained PLM space toward interaction-relevant variation observed in the source data. Before target optimization begins, the encoder is frozen and applied to each target–mutant pair to obtain $z_m = g_{\hat{\theta}}(r(A^{\mathcal{T}}, B_m))$.

These representations define the geometry used by the target-specific residual model. Target measurements for one mutant can therefore inform predictions for other mutants that are nearby in the learned protein-pair space. The residual GP itself remains target specific: source interaction data determine the geometry through the frozen encoder, whereas the residual response structure is learned from target assays collected during the campaign. Architecture, training, and preprocessing details are provided in Appendix A.2.

Importantly, the transferred components play different roles during target optimization. The source-supervised protein-pair encoder is learned once from the source interactions and then frozen, providing a stable interaction-aware geometry, whereas the source-response calibration and target-specific residual process are repeatedly updated as target assays accumulate. This separation preserves reusable source structure while allowing source–target relevance and target-specific variation to adapt to the new partner.

3.5 Unified Source-Guided Target Surrogate

The two transfer mechanisms are combined in a single target surrogate. For mutant B_m, we model the latent target response as

$$f_m = \underbrace{a_{\mathcal{S}}(B_m; \beta)}_{\text{response-level transfer}} + \underbrace{\delta(z_m)}_{\text{target-specific residual}}, \qquad \delta \sim \mathcal{GP}(0, K_\eta). \tag{1}$$

The first term is the adaptively calibrated source-response mean. It directly transfers candidate-level response information from the complete source landscapes and learns from the observed target assays how those source profiles should be combined for the new partner. The second term is a target-specific GP residual that captures response structure not explained by the transferred source patterns. Its covariance is defined in the source-supervised protein-pair representation introduced above, so that nearby target–mutant pairs can share information through the residual GP. We use an isotropic Matérn-$5/2$ kernel in the learned representation space.

Suppose that $n = |O|$ target mutants have been observed. Let H collect their source-response features, $y \in \mathbb{R}^n$ their observed target responses, and C_η denote the covariance matrix of the observed target responses induced by the residual GP and observation noise. For fixed covariance parameters η, the source-response mean is calibrated by generalized least squares,

$$\widehat{\beta}(\eta) = \left(H^\top C_\eta^{-1} H\right)^{-1} H^\top C_\eta^{-1} y. \tag{2}$$

The covariance parameters are estimated by restricted maximum likelihood (REML) (Patterson & Thompson, 1971), with $\widehat{\beta}$ recomputed as the covariance parameters change.

Conditioning on the observed target assays yields a joint posterior $f_U \mid y \sim \mathcal{N}(\mu_U, \Sigma_U)$ over the unmeasured candidate pool. The posterior mean μ_U combines the adaptively calibrated source-response mean with a correction learned from target residuals, while the posterior covariance Σ_U captures uncertainty in both the target-specific residual process and the calibration of the source-response mean from limited target data. Thus, uncertainty in response-level transfer is propagated into prediction rather than treating the calibrated source contribution as fixed. The explicit form of C_η, the kernel specification, the posterior mean and covariance, and the REML derivation are provided in Appendix A.3.

Batch acquisition. The resulting joint posterior is used to select the next batch of mutants through fantasy-conditioned Thompson sampling (Russo & Van Roy, 2016; Desautels et al., 2014). Within each batch, a Thompson sample is drawn over the remaining candidates and its maximizer is selected. The selected mutant is then assigned its current posterior mean as a fantasy observation before drawing a fresh Thompson sample for the next batch position. Once the batch has been assayed, the observed target responses are added to the data and the source-response calibration and target-specific residual model are refitted. Full acquisition and conditioning details are provided in Appendix A.4.

4 Controlled Simulations

We use semi-synthetic interaction landscapes to isolate the two transfer mechanisms in SPIN-BO: response-level transfer through an adaptively calibrated source-response mean, and representation-level transfer through a PLM-informed, source-supervised protein-pair representation for target-specific residual modeling. By independently varying source–target response association and smooth response structure aligned with shared ESM-derived coordinates, we examine when each form of transfer is useful and when combining them improves discovery.

Controlled landscapes. Motivated by N-TIMP2–MMP specificity studies (Arkadash et al., 2017; Aharon et al., 2020), we construct a candidate library containing wild-type human N-TIMP2 and all single and double substitutions at seven interface sites, giving $M = 7{,}715$ variants. MMP-10 serves as the target partner, and eight other human MMPs form the source panel. Protein sequences and ESM2 embeddings (Lin et al., 2023) are real, while the interaction responses are synthetic dimensionless scores designed to independently control source–target response association and smooth response structure.

Specifically, let $g_0, \ldots, g_8$ denote smooth functions of three shared ESM principal-component coordinates, and let $v_0, \ldots, v_8$ denote unstructured candidate-level variation. After centering and orthogonalization, we define

$$u_j = \sqrt{\gamma}\, g_j + \sqrt{1-\gamma}\, v_j, \qquad f_T = u_0, \qquad f_s = \rho_s u_0 + \sqrt{1-\rho_s^2}\, u_s. \tag{3}$$

The generator separates the two sources of transfer benefit targeted by SPIN-BO. The parameter $\rho_s = \mathrm{Corr}_M(f_s, f_T)$ controls how directly source s predicts the target response and therefore the information available for response-level transfer through the source-response mean. In contrast, γ controls the fraction of response variation that is smooth in the shared ESM-derived coordinate system, determining how much structure source supervision can reveal for representation-level transfer. We cross $\rho \in \{0, 0.3, 0.6, 0.9\}$ with $\gamma \in \{0, 0.3, 0.7, 0.9, 1\}$, assigning two sources $\rho_s = \rho$ and the remaining six $\rho_s = 0$. This design creates regimes in which useful transfer arises primarily at the response level, at the representation level, through both mechanisms, or through neither. Additional construction details are provided in Appendix B.2.

Protocol, ablations, and evaluation. Each noiseless campaign begins with 24 target assays and acquires 24 additional batches of 16, for a total target budget of 408 assays. We evaluate five independently generated landscapes, with two optimization runs per landscape using different initialization and acquisition seeds; initial target observations are shared across methods.

Four GP variants isolate response-level and representation-level transfer. **SPIN-BO (Both)** uses both the adaptive source-response mean and the source-supervised encoder. **SPIN-BO (Response Transfer)** retains response-level transfer but removes the learned encoder, using raw pretrained ESM geometry for the residual GP. **SPIN-BO (Representation Transfer)** retains representation-level transfer through SPIN-BO (Both)'s frozen encoder but replaces the source-response mean with an intercept-only mean. **Target-only BO** removes both transfer mechanisms, using an intercept-only mean and raw ESM geometry without source responses. All four retain the same target-specific residual GP structure and share the fitting and acquisition procedures in Section 3. As external target-only comparators, we additionally include ALDE-style (Yang et al., 2025), EVOLVEpro-style (Jiang et al., 2025), and random search, which samples unmeasured candidates uniformly without replacement. Additional comparison-method details are provided in Appendix B.3.

Table 1: **Complementary transfer mechanisms at $(\rho, \gamma) = (0.6, 0.9)$.** Mean (standard error) across five landscapes. OptSuccess (408) denotes optimum success evaluated after 408 target assays, and $R_{20}(n)$ denotes Recall@20 at target budget n. Bold marks the best mean.

Method	INR ↓	OptSuccess (408) ↑	$R_{20}(200)$ ↑	$R_{20}(408)$ ↑
SPIN-BO (Both)	**0.0096** (0.0018)	**1.000** (0.000)	**0.810** (0.066)	**0.935** (0.030)
SPIN-BO (Response Transfer)	0.0138 (0.0011)	0.700 (0.200)	0.655 (0.106)	0.880 (0.084)
SPIN-BO (Representation Transfer)	0.0347 (0.0052)	0.600 (0.245)	0.350 (0.093)	0.550 (0.084)
Target-only BO	0.0831 (0.0134)	0.000 (0.000)	0.035 (0.010)	0.105 (0.012)
ALDE-style	0.0870 (0.0171)	0.200 (0.200)	0.055 (0.022)	0.155 (0.060)
EVOLVEpro-style	0.0630 (0.0126)	0.200 (0.122)	0.120 (0.029)	0.240 (0.057)
Random	0.1018 (0.0167)	0.000 (0.000)	0.025 (0.016)	0.045 (0.020)

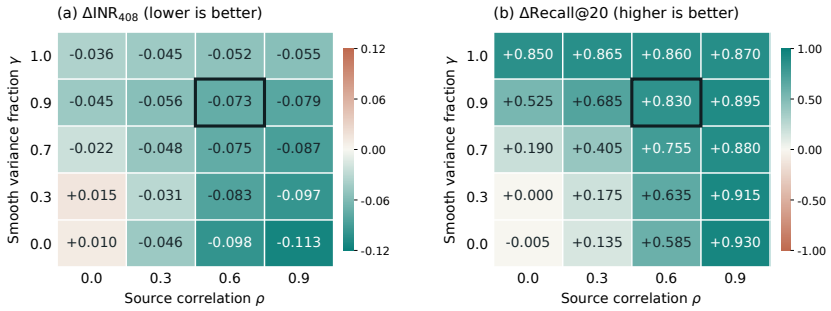

Figure 2: **Transfer gains depend on source–target association and smooth response structure.** Mean differences (SPIN-BO (Both) minus Target-only BO) across five landscapes over the (ρ, γ) grid. Negative ΔINR and positive ΔRecall@20 at 408 assays favor SPIN-BO (Both). The outlined cell marks the focal setting $(0.6, 0.9)$. Absolute performance is reported in Appendix Figure 4.

We evaluate both discovery speed and recovery of high-response variants. Integrated normalized regret (INR) summarizes the normalized regret trajectory over budgets 24–408, with lower values indicating faster discovery. Optimum success is the proportion of runs that have assayed a global optimum by a given budget, and Recall@20 is the fraction of the true top 20 target variants that have been assayed by that budget. Results are summarized across the five landscapes using the mean and standard error. Formal definitions of the evaluation metrics are provided in Appendix B.4.

Combining both transfer mechanisms improves discovery. We use $(\rho, \gamma) = (0.6, 0.9)$ as a focal setting where both source–target response association and smooth ESM-aligned structure are informative. SPIN-BO (Both) achieves an INR of 0.0096, compared with 0.0138 for SPIN-BO (Response Transfer), 0.0347 for SPIN-BO (Representation Transfer), and 0.0831 for Target-only BO (Table 1). It also reaches an optimum-success rate of 1.00 by 408 assays, versus 0.70, 0.60, and 0.00, respectively. The advantage is already clear at 200 assays, where SPIN-BO (Both) recovers 81.0% of the true top-20 variants compared with 65.5% for SPIN-BO (Response Transfer); at 408 assays, the corresponding recalls are 93.5% and 88.0%. Response-level transfer guides candidate prioritization through the calibrated source-response mean, while representation-level transfer provides geometry for target-specific residual modeling, yielding complementary gains when both signals are informative.

The two transfer mechanisms are useful in different regimes. The controlled grid separates direct predictive information in the source profiles from smooth structure that supports representation-level transfer. When source–target association is absent but smooth structure is strong, $(\rho, \gamma) = (0, 0.9)$, SPIN-BO (Representation Transfer) reduces INR from 0.0831 for Target-only BO to 0.0329, showing that source-supervised geometry can improve target-specific residual modeling even without direct source–target response association. Conversely, when source–target association is strong but smooth structure is absent, $(0.9, 0)$, SPIN-BO (Response Transfer) and SPIN-BO

(Both) recover all top-20 variants, whereas SPIN-BO (Representation Transfer) recovers only 9.0%, indicating that transfer in this regime is driven primarily by the adaptive source-response mean. When neither signal is present, $(0, 0)$, SPIN-BO (Both) provides no improvement over Target-only BO. Across the full grid, SPIN-BO lowers mean INR in 18 of 20 conditions, with the two exceptions occurring at $\rho = 0$ and $\gamma \leq 0.3$. Absolute performance is strongest when both source–target association and smooth structure are high, although stronger smooth structure also improves Target-only BO and therefore narrows the relative gap. Complete transfer-mechanism ablations across all 20 conditions are reported in Appendix Table 5.

5 Retrospective Discovery

Dataset and discovery task. We evaluate SPIN-BO on the human bZIP deep mutational scan of Bendel et al. (2025), which measures mutant–partner interactions using bindingPCA. Each retrospective campaign fixes a mutant background and a target partner, then searches over single-substitution variants to maximize the recorded pair-level log fold change. The complete target profile is withheld during optimization and used only for retrospective evaluation; querying a variant reveals its target interaction score. Complete interaction profiles for the remaining source partners are available throughout the campaign as source knowledge.

Biologically grouped source panels. We evaluate ten partner panels across four mutant backgrounds: BATF2, CEBPB, DDIT3, and JUNB. Partners are grouped by bZIP family membership, sequence similarity, and known interaction biology (Vinson et al., 2002; Amoutzias et al., 2007; Wingender et al., 2018). This reflects a realistic transfer setting in which source partners would typically be chosen because they are biologically related to the target and therefore more likely to provide informative prior knowledge. Within each panel, we retain variants measured across all partners, yielding candidate libraries of 501–614 variants. Each partner serves in turn as the target, with all remaining panel members used jointly as sources, giving 54 discovery problems with 3–9 source partners. Panel composition and data-processing details are provided in Appendix C.

Evaluation protocol. Each discovery problem uses ten seeds, 20 shared initial target assays, and nine adaptive batches of 20, for a total target budget of 200 assays. We compare SPIN-BO with the same methods used in Section 4. We report normalized best score, optimum success, Recall@5, and Recall@20. Results are aggregated equally across the ten partner panels; full aggregation details are provided in Appendix C.2.

Discovery efficiency and hit recovery. SPIN-BO improves discovery over Target-only BO throughout the assay trajectory (Figure 3). The advantage is already apparent at 100 target assays: SPIN-BO achieves a mean optimum-success rate of 82.2% versus 53.6% for Target-only BO, with Recall@20 of 72.3% versus 48.8% (Appendix Table 7). At 200 assays, the corresponding optimum-success rates are 95.0% and 73.3%, while SPIN-BO recovers 89.1% of the true top 5 variants and 88.7% of the top 20 (Table 2).

SPIN-BO (Response Transfer) achieves a 93.9% optimum-success rate at 200 assays, indicating that response-level transfer accounts for a substantial portion of the improvement in these bZIP panels. Representation-level transfer is also useful on its own: SPIN-BO (Representation Transfer) outperforms Target-only BO across all four final-budget metrics. Combining the two mechanisms gives SPIN-BO (Both) the highest normalized best score, optimum-success rate, and Recall@5, while SPIN-BO (Response Transfer) attains a slightly higher Recall@20 (89.2% versus 88.7%). Overall, response-level transfer provides most of the gain in this dataset, while representation-level transfer remains independently beneficial and further improves several discovery endpoints when combined with response-level transfer. A complementary single-source analysis shows that SPIN-BO (Both) also improves discovery over Target-only BO when only one source interaction profile is available, indicating that useful transfer can arise even from a single source partner. Full single-source discovery trajectories are reported in Appendix Figure 5.

Robustness to biology-agnostic source panels. To examine whether SPIN-BO can benefit from source transfer without biologically curated partner groups, we construct ten panels without using bZIP family membership, sequence similarity, or known interaction biology. For each panel, we

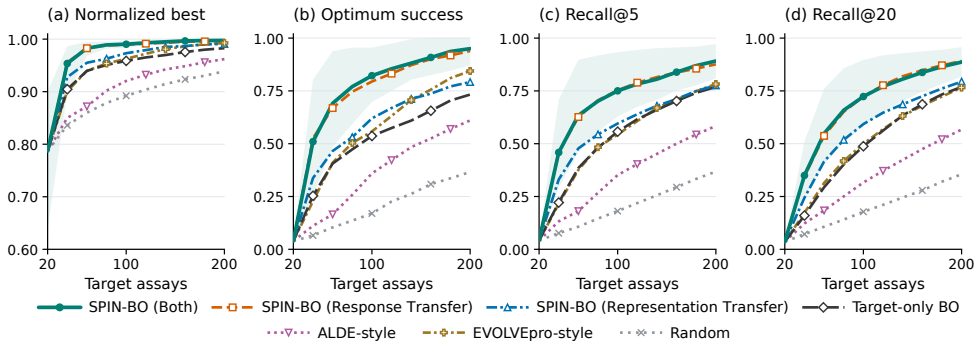

Figure 3: **Multi-source bZIP discovery.** Curves show equal-panel means for normalized best score, optimum success, Recall@5, and Recall@20 as the target-assay budget increases. Shading shows descriptive 95% bootstrap intervals for SPIN-BO (Both).

Table 2: **Multi-source discovery at a target budget of 200 assays.** Equal-panel means across retrospective discovery tasks. OptSuccess denotes optimum success at 200 assays. Higher is better; bold marks the best mean.

Method	Normalized best	OptSuccess	Recall@5	Recall@20
SPIN-BO (Both)	**0.998**	**0.950**	**0.891**	0.887
SPIN-BO (Response Transfer)	0.996	0.939	0.876	**0.892**
SPIN-BO (Representation Transfer)	0.991	0.792	0.776	0.794
Target-only BO	0.983	0.733	0.769	0.769
ALDE-style	0.962	0.610	0.584	0.567
EVOLVEpro-style	0.993	0.844	0.782	0.766
Random	0.939	0.366	0.368	0.355

randomly select a mutant background and retain ten partners with shared mutant coverage. We evaluate a 9-source setting, where each partner serves in turn as the target and the other nine as sources, and a 5-source setting, where five of the remaining nine partners are randomly selected as sources. Using the same assay protocol, SPIN-BO continues to outperform Target-only BO, although the gains are smaller than with biologically grouped sources. This suggests that biological relatedness strengthens cross-partner transfer, while useful gains remain possible without such curation. Full results are reported in Appendix C.4 (Figures 6 and 7; Tables 8 and 9).

6 Conclusion

We introduced SPIN-BO, a source-guided Bayesian optimization framework for assay-efficient PPI discovery with existing interaction landscapes from related partners. The key idea is to use source knowledge in two complementary ways: to guide prediction through an adaptively calibrated source-response mean and to inform target-specific residual modeling through a source-supervised protein-pair representation. By combining these components in a unified probabilistic surrogate, SPIN-BO can borrow useful cross-partner information while continuing to adapt to the new target as additional assays are collected. Controlled semi-synthetic experiments demonstrate that response-level and representation-level transfer are useful under different source–target relationships and can provide complementary gains when combined. Retrospective experiments on the human bZIP interaction network further show improved discovery over target-only optimization under limited assay budgets.

Overall, our results support the use of existing interaction landscapes as reusable knowledge for discovering high-response variants against new partners, rather than treating each target as an independent optimization problem. Extending this framework from the current single-target setting to joint discovery across multiple target partners is a potential direction for future work, where measurements acquired for one target could also inform optimization of others.

AI USE STATEMENT

In this work, we used generative AI tools to provide feedback on experimental design and to assist with manuscript writing and revision. We also used generative AI for limited support with code development. All AI-assisted suggestions, code, and text were reviewed and verified by the authors, who made all final scientific and editorial decisions. We take responsibility for the final content of this work, including any material produced with the aid of generative AI.

REPRODUCIBILITY STATEMENT

To support reproducibility, Section 3 describes SPIN-BO's response-level and representation-level transfer mechanisms, unified source-guided target surrogate, and fantasy-conditioned batch acquisition. Sections 4 and 5 present the controlled semi-synthetic simulations and retrospective bZIP discovery experiments, including their evaluation protocols. Appendix A provides implementation details and derivations for response-level calibration, representation learning, unified surrogate fitting and posterior prediction, and batch acquisition. Appendix B details library construction, landscape generation, transfer ablations, comparison methods, and evaluation metrics. Appendix C documents bZIP data processing, source-panel construction, and evaluation and aggregation procedures, including the single-source and biology-agnostic analyses. An anonymized code repository containing the implementation and scripts used to reproduce the experiments is available at https://anonymous.4open.science/r/SPIN-BO-7861/.

REFERENCES

Lidan Aharon, Shay-Lee Aharoni, Evette S. Radisky, and Niv Papo. Quantitative mapping of binding specificity landscapes for homologous targets by using a high-throughput method. *Biochemical Journal*, 477(9):1701–1719, 2020. doi: 10.1042/BCJ20200188.

Ethan C. Alley, Grigory Khimulya, Surojit Biswas, Mohammed AlQuraishi, and George M. Church. Unified rational protein engineering with sequence-based deep representation learning. *Nature Methods*, 16(12):1315–1322, 2019. doi: 10.1038/s41592-019-0598-1.

G. D. Amoutzias, A. S. Veron, J. Weiner, III, M. Robinson-Rechavi, E. Bornberg-Bauer, S. G. Oliver, and D. L. Robertson. One billion years of bZIP transcription factor evolution: conservation and change in dimerization and DNA-binding site specificity. *Molecular Biology and Evolution*, 24 (3):827–835, 2007. doi: 10.1093/molbev/msl211.

Valeria Arkadash, Gal Yosef, Jason Shirian, Itay Cohen, Yuval Horev, Moran Grossman, Irit Sagi, Evette S. Radisky, Julia M. Shifman, and Niv Papo. Development of high affinity and high specificity inhibitors of matrix metalloproteinase 14 through computational design and directed evolution. *Journal of Biological Chemistry*, 292(8):3481–3495, 2017. doi: 10.1074/jbc.M116.756718.

A. M. Bendel, B. Eichenberger, G. Kempf, D. Klein, K. Shimada, M. Franco, S. Jaaks-Kraatz, S. Durdu, V. Herrmann, G. Roth, C. Soneson, S. Smallwood, N. H. Thomä, D. Schübeler, M. B. Stadler, F. Zenke, and G. Diss. The genetic architecture of the human bZIP interaction network. bioRxiv, 2025. URL https://doi.org/10.1101/2025.08.21.671354. Preprint, version posted August 22, 2025.

Lin Chen, Zehong Zhang, Zhenghao Li, Rui Li, Ruifeng Huo, Lifan Chen, Dingyan Wang, Xiaomin Luo, Kaixian Chen, Cangsong Liao, and Mingyue Zheng. Learning protein fitness landscapes with deep mutational scanning data from multiple sources. *Cell Systems*, 14(8):706–721.e5, 2023. doi: 10.1016/j.cels.2023.07.003.

Thomas Desautels, Andreas Krause, and Joel W. Burdick. Parallelizing exploration-exploitation tradeoffs in Gaussian process bandit optimization. *Journal of Machine Learning Research*, 15(119):4053–4103, 2014. URL https://www.jmlr.org/papers/v15/desautels14a.html.

Matthias Feurer, Benjamin Letham, Frank Hutter, and Eytan Bakshy. Practical transfer learning for Bayesian optimization, 2022. URL https://arxiv.org/abs/1802.02219v4. arXiv preprint arXiv:1802.02219, version 4, revised October 24, 2022.

Peter I. Frazier. A tutorial on Bayesian optimization, 2018. URL https://arxiv.org/abs/1807.02811. arXiv preprint arXiv:1807.02811.

Vanessa E. Gray, Ronald J. Hause, Jens Luebeck, Jay Shendure, and Douglas M. Fowler. Quantitative missense variant effect prediction using large-scale mutagenesis data. *Cell Systems*, 6(1): 116–124.e3, 2018. doi: 10.1016/j.cels.2017.11.003.

Milind Jagota, Chengzhong Ye, Carlos Albors, Ruchir Rastogi, Antoine Koehl, Nilah Ioannidis, and Yun S. Song. Cross-protein transfer learning substantially improves disease variant prediction. *Genome Biology*, 24(1):182, 2023. doi: 10.1186/s13059-023-03024-6.

Kaiyi Jiang, Zhaoqing Yan, Matteo Di Bernardo, Samantha R. Sgrizzi, Lukas Villiger, Alisan Kayabolen, B. J. Kim, Josephine K. Carscadden, Masahiro Hiraizumi, Hiroshi Nishimasu, Jonathan S. Gootenberg, and Omar O. Abudayyeh. Rapid in silico directed evolution by a protein language model with EVOLVEpro. *Science*, 387(6732):eadr6006, 2025. doi: 10.1126/science.adr6006.

Marc C. Kennedy and Anthony O'Hagan. Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(3):425–464, 2001. doi: 10.1111/1467-9868.00294.

Yanke Li, Tianyu Cui, Tommaso Mansi, Mangal Prakash, and Rui Liao. BioBO: Biology-informed Bayesian optimization for perturbation design. In *International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=CF3kJrAwmV.

Haitao Lin, Boxin Zhao, Mladen Kolar, and Chong Liu. Provable accelerated bayesian optimization with knowledge transfer. In *Proceedings of the 29th International Conference on Artificial Intelligence and Statistics*, volume 300 of *Proceedings of Machine Learning Research*, pp. 2359–2367. PMLR, 2026. URL https://proceedings.mlr.press/v300/lin26c.html.

Zeming Lin, Halil Akin, Roshan Rao, Brian Hie, Zhongkai Zhu, Wenting Lu, Nikita Smetanin, Robert Verkuil, Ori Kabeli, Yaniv Shmueli, Allan dos Santos Costa, Maryam Fazel-Zarandi, Tom Sercu, Salvatore Candido, and Alexander Rives. Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science*, 379(6637):1123–1130, 2023. doi: 10.1126/science.ade2574.

H. D. Patterson and R. Thompson. Recovery of inter-block information when block sizes are unequal. *Biometrika*, 58(3):545–554, 1971. doi: 10.1093/biomet/58.3.545.

Valerio Perrone, Rodolphe Jenatton, Matthias W. Seeger, and Cédric Archambeau. Scalable hyperparameter transfer learning. In *Advances in Neural Information Processing Systems*, volume 31, 2018. URL https://papers.nips.cc/paper_files/paper/2018/hash/14c879f3f5d8ed93a09f6090d77c2cc3-Abstract.html.

Chongqi Qu, Meiqin Liu, Jian Lan, Shanling Dong, and Zhunga Liu. Incorporating expert priors into Bayesian optimization via dynamic mean decay. In *International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=ylX5PWfK0v.

Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems*, volume 20, 2007. URL https://papers.nips.cc/paper/2007/hash/013a006f03dbc5392effeb8f18fda755-Abstract.html.

Daniel Russo and Benjamin Van Roy. An information-theoretic analysis of Thompson sampling. *Journal of Machine Learning Research*, 17(68):1–30, 2016. URL https://jmlr.org/papers/v17/14-087.html.

Kevin Swersky, Jasper Snoek, and Ryan P. Adams. Multi-task Bayesian optimization. In *Advances in Neural Information Processing Systems*, volume 26, 2013. URL https://papers.nips.cc/paper_files/paper/2013/hash/f33ba15effa5c10e873bf3842afb46a6-Abstract.html.

Charles Vinson, Max Myakishev, Asha Acharya, Alain A. Mir, Jonathan R. Moll, and Maria Bonovich. Classification of human B-ZIP proteins based on dimerization properties. *Molecular and Cellular Biology*, 22(18):6321–6335, 2002. doi: 10.1128/MCB.22.18.6321-6335.2002.

Edgar Wingender, Torsten Schoeps, Martin Haubrock, Mathias Krull, and Jürgen Dönitz. TFClass: expanding the classification of human transcription factors to their mammalian orthologs. *Nucleic Acids Research*, 46(D1):D343–D347, 2018. doi: 10.1093/nar/gkx987.

Jason Yang, Ravi G. Lal, James C. Bowden, Raul Astudillo, Mikhail A. Hameedi, Sukhvinder Kaur, Matthew Hill, Yisong Yue, and Frances H. Arnold. Active learning-assisted directed evolution. *Nature Communications*, 16:714, 2025. doi: 10.1038/s41467-025-55987-8.

Kevin K. Yang, Zachary Wu, and Frances H. Arnold. Machine-learning-guided directed evolution for protein engineering. *Nature Methods*, 16(8):687–694, 2019. doi: 10.1038/s41592-019-0496-6.

Yimeng Zeng, Natalie Maus, Haydn Jones, Jeffrey Tao, Fangping Wan, Marcelo Der Torossian Torres, Cesar de la Fuente, Ryan Marcus, Osbert Bastani, and Jacob Gardner. Scaling multi-task Bayesian optimization with large language models. In *International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=ekmUkRYnkN.

A METHOD DETAILS AND DERIVATIONS

This appendix provides additional details for response-level calibration, representation-level transfer, unified surrogate fitting and posterior prediction, and batch acquisition described in Section 3.

A.1 RESPONSE-LEVEL TRANSFER AND INITIALIZATION

For each mutant B_m, recall that $h_m = (1, Y^{\mathcal{S}}_{1m}, \ldots, Y^{\mathcal{S}}_{Sm})^\top$ collects an intercept and its responses across the S source partners. Let $q = S + 1$ denote the number of mean coefficients. For an observed target set O, let H_O denote the design matrix whose rows are $h_m^\top$ for $m \in O$. The response-level transfer component is therefore $H_O\beta$, with the coefficients β recalibrated as target assays accumulate.

To ensure that the source-response mean is identifiable at the start of the campaign, we require the initial design matrix to have full column rank. The initial target set is sampled uniformly without replacement with size $n_0 > q$, subject to this rank condition. The same accepted initial target assays are used across matched comparison methods, and all initial assays count toward the target-assay budget.

A.2 REPRESENTATION-LEVEL TRANSFER

As in Section 3.4, a partner–mutant pair is represented from pretrained ESM2-650M embeddings as $r(A, B) = [e(A); e(B)]$, where $e(\cdot)$ is obtained from the 33rd layer of `esm2_t33_650M_UR50D`. The pretrained PLM remains fixed.

Let $\mathcal{D}_{\text{enc}}$ denote the source partner–mutant training pairs. The encoder g_θ and a temporary scalar prediction head q_ψ are trained by

$$(\widehat{\theta}, \widehat{\psi}) = \arg\min_{\theta,\psi} \frac{1}{|\mathcal{D}_{\text{enc}}|} \sum_{(s,m)\in\mathcal{D}_{\text{enc}}} \left[Y^{\mathcal{S}}_{sm} - q_\psi\big(g_\theta(r(A^{\mathcal{S}}_s, B_m))\big)\right]^2.$$

The concatenated PLM input has dimension 2560, and the encoder uses the architecture $2560 \to 128 \to 32$ with a GELU activation, followed during training by the scalar linear prediction head. Training uses AdamW with learning rate 10^{-3}, weight decay 10^{-4}, and batch size 256.

After training, the encoder is frozen and the prediction head is discarded. The target–mutant representation is then $z_m = g_{\widehat{\theta}}(r(A^{\mathcal{T}}, B_m))$. Thus, source interaction responses supervise the learned protein-pair geometry, while target responses do not update the representation during optimization.

A.3 UNIFIED SURROGATE FITTING AND POSTERIOR DERIVATION

Section 3.5 combines the adaptively calibrated source-response mean with a target-specific GP residual defined in the frozen protein-pair representation. We first give the covariance specification and REML fitting procedure, followed by the posterior expressions used for prediction and acquisition.

The residual GP uses an isotropic Matérn-$5/2$ covariance,

$$K_\eta(m, m') = \tau^2\left(1 + \sqrt{5}d_{mm'} + \frac{5d^2_{mm'}}{3}\right)\exp\left(-\sqrt{5}d_{mm'}\right), \qquad d_{mm'} = \frac{\|z_m - z_{m'}\|_2}{\ell},$$

where τ^2 is the residual variance and ℓ is the length scale.

Suppose that n target mutants indexed by O have been observed. Let H denote their source-response design matrix, $y \in \mathbb{R}^n$ their observed target responses, and

$$C_\eta = K_{OO} + \sigma^2_\varepsilon I_n, \qquad A_\eta = H^\top C_\eta^{-1} H.$$

Here, C_η is the covariance matrix of the observed target responses induced by the target-specific residual GP and observation noise. For fixed covariance parameters η, the generalized least-squares estimate of the response-level calibration coefficients is

$$\widehat{\beta}(\eta) = A_\eta^{-1} H^\top C_\eta^{-1} y,$$

which is Equation (2).

Define

$$P_\eta = C_\eta^{-1} - C_\eta^{-1} H A_\eta^{-1} H^\top C_\eta^{-1}.$$

The covariance parameters are estimated by maximizing the restricted log likelihood

$$\ell_R(\eta) = -\frac{1}{2}\left\{\log|C_\eta| + \log|A_\eta| + y^\top P_\eta y + (n-q)\log(2\pi)\right\},$$

where q is the number of columns of H. The GLS coefficients are recomputed at each covariance-parameter evaluation.

We next derive the joint posterior over the unmeasured candidate set U. All quantities below are evaluated at the fitted covariance parameters, with this dependence suppressed for readability. Define

$$D_U = H_U - K_{UO} C^{-1} H, \qquad A = H^\top C^{-1} H.$$

Conditional on β, Gaussian conditioning gives

$$\mathbb{E}[f_U \mid y, \beta] = K_{UO} C^{-1} y + D_U \beta,$$
$$\mathrm{Cov}(f_U \mid y, \beta) = K_{UU} - K_{UO} C^{-1} K_{OU}.$$

Using the flat prior $p(\beta) \propto 1$ for the mean coefficients,

$$\beta \mid y \sim \mathcal{N}\left(\widehat{\beta}, A^{-1}\right).$$

Applying total expectation and total covariance therefore yields

$$f_U \mid y \sim \mathcal{N}(\mu_U, \Sigma_U),$$

where

$$\mu_U = H_U \widehat{\beta} + K_{UO} C^{-1}\left(y - H\widehat{\beta}\right),$$

and

$$\Sigma_U = K_{UU} - K_{UO} C^{-1} K_{OU} + D_U A^{-1} D_U^\top.$$

The posterior mean combines the adaptively calibrated source-response mean with a correction learned from target residuals. In the posterior covariance, the first two terms give the usual conditional GP covariance, while the final term propagates uncertainty in calibrating the response-level transfer component from the available target measurements.

The predictive distribution is empirical-Bayes: it conditions on the frozen encoder and fitted covariance parameters, while retaining uncertainty in the source-response calibration coefficients. Uncertainty from encoder training and covariance-parameter estimation is not propagated.

For the retrospective bZIP experiments, the covariance parameters are constrained to $\tau \in [10^{-3}, 10]$, $\ell \in [0.1 d_{\mathrm{med}}, 10 d_{\mathrm{med}}]$, and $\sigma_\varepsilon \in [0.03, 2]$, where d_{med} is the median nonzero distance in the frozen representation. The controlled simulations are noiseless and fix $\sigma_\varepsilon = 0$.

A.4 Fantasy-Conditioned Batch Acquisition

The joint posterior from Section 3.5 is used to construct each batch sequentially. We use fantasy-conditioned Thompson sampling, with the full within-batch procedure detailed here.

Let V denote the set of candidates currently eligible for selection. At each within-batch step, we draw a joint Thompson sample

$$\widetilde{f}_V \sim \mathcal{N}(\mu_V, \Sigma_V),$$

and select

$$m = \arg\max_{j \in V} \widetilde{f}_j$$

as the next mutant in the batch (Russo & Van Roy, 2016). Because the true target response of m is not observed until the entire batch has been assayed, we assign the selected mutant its current posterior mean as a fantasy observation (Desautels et al., 2014) and condition the predictive distribution before selecting the next batch member.

For a selected mutant m and the remaining candidate set R, write $c = \Sigma_{Rm}$ and $v = \Sigma_{mm}$. Under the current predictive distribution, the joint distribution of the remaining latent responses and the planned noisy assay at m is

$$\begin{pmatrix} f_R \\ Y_m \end{pmatrix} \sim \mathcal{N}\left[\begin{pmatrix} \mu_R \\ \mu_m \end{pmatrix}, \begin{pmatrix} \Sigma_{RR} & c \\ c^\top & v + \widehat{\sigma}_\varepsilon^2 \end{pmatrix} \right].$$

Conditioning on the posterior-mean fantasy $Y_m = \mu_m$ gives

$$\mu^+ = \mu_R, \qquad \Sigma^+ = \Sigma_{RR} - \frac{\Sigma_{Rm}\Sigma_{mR}}{\Sigma_{mm} + \widehat{\sigma}_\varepsilon^2}.$$

Because the fantasy value equals the current posterior mean, the predictive mean remains unchanged, while the covariance contracts according to the posterior dependence between the selected mutant and the remaining candidates. In particular, the marginal variance of a remaining candidate j decreases by

$$\frac{\Sigma_{jm}^2}{\Sigma_{mm} + \widehat{\sigma}_\varepsilon^2}.$$

After conditioning, a fresh joint Thompson sample is drawn from $\mathcal{N}(\mu^+, \Sigma^+)$ over the remaining candidates, and its maximizer is selected for the next batch position. Fantasy conditioning and resampling are repeated until the batch is complete. Previously queried mutants and candidates already selected into the current batch are excluded from subsequent selections.

No target outcomes are revealed during within-batch construction. The encoder and covariance parameters remain fixed while the batch is assembled. Once the complete batch has been assayed, the observed target responses are added to the data, the source-response calibration and target-specific residual GP are refitted, and the next BO round begins.

We use this fantasy-conditioned acquisition rule as a practical batch policy and do not rely on standard Thompson-sampling regret guarantees for the deterministic posterior-mean fantasy construction.

B Simulation Details

This appendix provides additional details on the protein library, semi-synthetic landscape construction, comparison methods, evaluation metrics, and full-grid results for the controlled simulations in Section 4.

B.1 Protein Library and Sequence Embeddings

The N-TIMP2 scaffold uses mature residues 1–127 (UniProt P16035, precursor positions 27–153). The seven mutagenized sites are S4, I35, N38, S68, V71, H97, and T99 (Arkadash et al., 2017; Aharon et al., 2020). The library contains one wild type, 133 single mutants, and 7,581 double mutants, giving $M = 7{,}715$ candidates.

MMP-10 serves as the target partner, and the remaining eight MMP constructs in Table 3 form the source panel. We use 1,280-dimensional embeddings from the 33rd layer of `esm2_t33_650M_UR50D` (Lin et al., 2023). Mutant embeddings are obtained by averaging the seven designated interface positions, while partner embeddings are averaged over all construct residues.

Table 3: **Human protein constructs used in the controlled simulations.** Intervals are one-based UniProt precursor positions.

Role	Protein	UniProt	Interval	Length
Scaffold	N-TIMP2	P16035	27–153	127
Target	MMP-10	P09238	99–263	165
Source	MMP-1	P03956	100–269	170
Source	MMP-2	P08253	110–449	340
Source	MMP-3	P08254	100–272	173
Source	MMP-7	P09237	95–267	173
Source	MMP-8	P22894	101–262	162
Source	MMP-9	P14780	107–443	337
Source	MMP-13	P45452	104–268	165
Source	MMP-14	P50281	112–287	176

B.2 Controlled Landscape Generation

The semi-synthetic generator is designed to vary the information available to response-level and representation-level transfer independently. We center the mutant embedding matrix, retain 64 principal components, and use standardized PC4, PC8, and PC12 as three generating coordinates $\xi_m \in \mathbb{R}^3$. These coordinates are fixed before landscape generation and are hidden from all optimization methods.

For each landscape, smooth response components are generated from 32 random cosine features (Rahimi & Recht, 2007),

$$\phi_r(\xi_m) = \cos(\omega_r^\top \xi_m + b_r), \qquad \omega_r \sim \mathcal{N}(0, 1.5^2 I_3), \qquad b_r \sim \mathrm{Unif}(0, 2\pi). \tag{4}$$

Nine smooth profiles are formed as independent random linear combinations of these features and then centered and orthogonalized to obtain $g_0, \ldots, g_8$ with zero mean and unit variance. Independent Gaussian candidate-level profiles are similarly orthogonalized to obtain $v_0, \ldots, v_8$, with the smooth and unstructured components mutually orthogonal.

The final target and source landscapes follow Equation (3). Because the smooth and unstructured components are standardized and orthogonal, γ determines the fraction of response variation attributable to the smooth ESM-aligned component, while $\rho_s = \mathrm{Corr}_M(f_s, f_T)$ determines the source–target response association. For each (ρ, γ) condition, two sources are assigned $\rho_s = \rho$ and the remaining six $\rho_s = 0$, as in Section 4. Thus, ρ primarily controls the information available for response-level transfer, whereas γ controls the smooth structure that can support representation-level transfer. Source correlations are assigned independently of partner-sequence similarity.

B.3 Transfer Ablations and Comparison Methods

Representation learning follows Section 3.4 and Appendix A.2. Simulation-specific preprocessing centers and scales mutant and partner embeddings using statistics computed from the source data. A fixed set of 4,096 candidates is used for encoder training, and the encoder is trained for 100 AdamW epochs before being frozen for target optimization.

Table 4 summarizes the four matched GP variants used to isolate response-level and representation-level transfer. The variants differ only in whether the adaptive source-response mean is retained and whether the frozen source-supervised encoder defines the residual geometry.

All four GP variants use the same target-specific residual GP, fitting procedure, and batch acquisition rule described in Section 3. Simulation responses are noiseless, so $\sigma_\varepsilon = 0$.

As external target-only comparators, ALDE-style (Yang et al., 2025) uses one-hot mutation features with an ensemble neural predictor and uncertainty-aware batch selection, whereas EVOLVEpro-style (Jiang et al., 2025) uses ESM2 mutant embeddings with a random-forest predictor and selects candidates with the highest predicted responses. Random search samples unmeasured candidates uniformly without replacement. These three comparators use only target measurements collected during the campaign.

Table 4: **Transfer mechanisms retained by the matched GP variants.**

Method	Target mean	Residual geometry
Target-only BO	Intercept only	Raw ESM2
SPIN-BO (Response Transfer)	Adaptive source-response mean	Raw ESM2
SPIN-BO (Representation Transfer)	Intercept only	Frozen source encoder
SPIN-BO (Both)	Adaptive source-response mean	Frozen source encoder

B.4 Evaluation Metrics

We evaluate both discovery speed and recovery of high-response variants. Let O_n denote the set of target variants assayed by target budget n, and let $f_T(m)$ denote the target response of candidate m.

Normalized regret. At target budget n, normalized regret is

$$\widetilde{r}_n = \frac{\max_m f_T(m) - \max_{m \in O_n} f_T(m)}{\max_m f_T(m) - \min_m f_T(m)}.$$

It equals zero once a global optimum has been assayed, while larger values indicate that the best variant discovered so far remains farther from the optimum.

Integrated normalized regret. Integrated normalized regret (INR) summarizes the full discovery trajectory from the initial budget $n_0 = 24$ to the final budget $B_{\text{tot}} = 408$:

$$\text{INR} = \frac{1}{B_{\text{tot}} - n_0} \int_{n_0}^{B_{\text{tot}}} \widetilde{r}(n)\, dn.$$

Lower INR indicates earlier discovery of high-response variants over the course of the campaign.

Optimum success. Let

$$\mathcal{M}^\star = \arg\max_m f_T(m)$$

denote the set of global optima. For a collection of evaluation runs $\mathcal{R}$, optimum success at budget n is

$$\text{OptSuccess}(n) = \frac{1}{|\mathcal{R}|} \sum_{r \in \mathcal{R}} \mathbb{I}\left\{ O_n^{(r)} \cap \mathcal{M}^\star \neq \varnothing \right\}.$$

Thus, optimum success is the proportion of campaigns that have discovered at least one global optimum by budget n.

Recall@K. Let $\mathcal{T}_K$ denote the set of the K highest-response target variants. For a run evaluated at budget n,

$$\text{Recall@}K(n) = \frac{|O_n \cap \mathcal{T}_K|}{K}.$$

We report Recall@20 in the controlled simulations at intermediate and final assay budgets.

B.5 Full-Grid Evaluation

The campaign protocol follows Section 4. Each (ρ, γ) condition uses five independently generated landscapes and two optimization runs per landscape. The two runs are first averaged within each landscape, after which we report the mean and standard error across the five landscapes.

Figure 4 shows the absolute performance of SPIN-BO (Both) over the full (ρ, γ) grid. Absolute discovery generally improves as either source–target response association or smooth ESM-aligned structure becomes more informative, with the strongest performance occurring when both are high.

Table 5 reports INR, optimum success, and Recall@20 for the four matched GP variants across all 20 simulation conditions. The full grid supports the main-text interpretation: response-level transfer

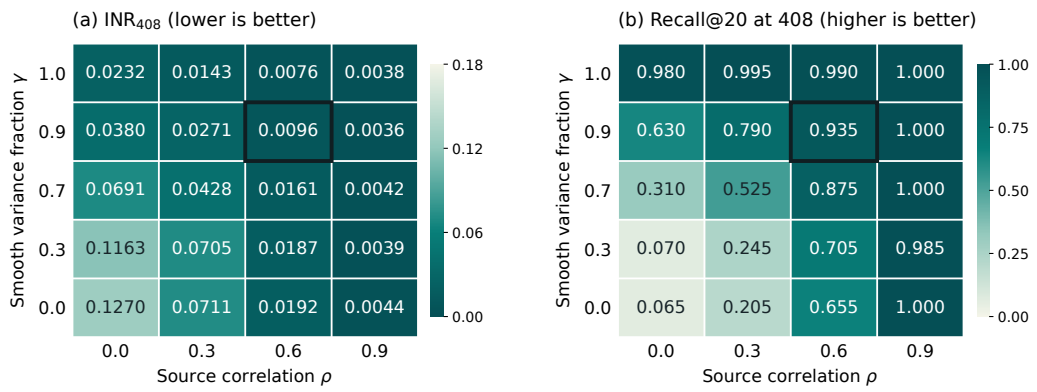

Figure 4: **Absolute SPIN-BO (Both) performance.** Mean INR and Recall@20 at 408 assays across five landscapes over the (ρ, γ) grid. The focal condition $(0.6, 0.9)$ is outlined.

becomes more informative as source–target response association increases, whereas representation-level transfer is most useful when substantial smooth ESM-aligned structure is available. Their combination provides complementary gains in regimes where both forms of information are informative.

Table 5: **Transfer-mechanism ablations across all 20 (ρ, γ) combinations.** Entries are mean (SE) across five landscapes, after averaging two runs per landscape. INR integrates budgets 24–408; optimum success is the fraction of campaigns finding the global optimum by 408 assays; Recall@20 is the fraction of the top 20 candidates observed by 200 or 408 assays. Bold marks the best mean among the four GP variants, including ties at the displayed precision.

(a) INR ↓

ρ	γ	SPIN-BO (Both)	SPIN-BO (Response Transfer)	SPIN-BO (Representation Transfer)	Target-only BO
0.0	0.0	0.1270 (0.0207)	0.1227 (0.0096)	**0.1151** (0.0095)	0.1169 (0.0171)
	0.3	0.1163 (0.0275)	0.1064 (0.0117)	0.1383 (0.0092)	**0.1014** (0.0204)
	0.7	0.0691 (0.0094)	0.0931 (0.0150)	**0.0608** (0.0152)	0.0910 (0.0125)
	0.9	0.0380 (0.0046)	0.0795 (0.0154)	**0.0329** (0.0054)	0.0831 (0.0134)
	1.0	**0.0232** (0.0045)	0.0872 (0.0230)	0.0238 (0.0060)	0.0592 (0.0146)
0.3	0.0	**0.0711** (0.0081)	0.0835 (0.0070)	0.1027 (0.0178)	0.1169 (0.0171)
	0.3	**0.0705** (0.0142)	0.0736 (0.0111)	0.1123 (0.0146)	0.1014 (0.0204)
	0.7	**0.0428** (0.0085)	0.0649 (0.0039)	0.0606 (0.0098)	0.0910 (0.0125)
	0.9	**0.0271** (0.0042)	0.0543 (0.0066)	0.0459 (0.0107)	0.0831 (0.0134)
	1.0	**0.0143** (0.0024)	0.0444 (0.0039)	0.0227 (0.0059)	0.0592 (0.0146)
0.6	0.0	**0.0192** (0.0054)	0.0233 (0.0089)	0.1358 (0.0156)	0.1169 (0.0171)
	0.3	0.0187 (0.0058)	**0.0171** (0.0046)	0.1091 (0.0110)	0.1014 (0.0204)
	0.7	**0.0161** (0.0069)	0.0180 (0.0049)	0.0670 (0.0079)	0.0910 (0.0125)
	0.9	**0.0096** (0.0018)	0.0138 (0.0011)	0.0347 (0.0052)	0.0831 (0.0134)
	1.0	**0.0076** (0.0015)	0.0133 (0.0026)	0.0217 (0.0066)	0.0592 (0.0146)
0.9	0.0	0.0044 (0.0002)	**0.0041** (0.0003)	0.1165 (0.0150)	0.1169 (0.0171)
	0.3	**0.0039** (0.0005)	0.0040 (0.0005)	0.1092 (0.0142)	0.1014 (0.0204)
	0.7	**0.0042** (0.0004)	**0.0042** (0.0007)	0.0575 (0.0084)	0.0910 (0.0125)
	0.9	**0.0036** (0.0005)	0.0038 (0.0005)	0.0414 (0.0092)	0.0831 (0.0134)
	1.0	**0.0038** (0.0004)	0.0040 (0.0003)	0.0226 (0.0066)	0.0592 (0.0146)

(b) Optimum success at 408 assays ↑

ρ	γ	SPIN-BO (Both)	SPIN-BO (Response Transfer)	SPIN-BO (Representation Transfer)	Target-only BO
0.0	0.0	**0.100** (0.100)	0.000 (0.000)	**0.100** (0.100)	**0.100** (0.100)
	0.3	**0.300** (0.200)	0.000 (0.000)	0.000 (0.000)	0.100 (0.100)
	0.7	0.200 (0.200)	0.100 (0.100)	**0.500** (0.224)	0.100 (0.100)
	0.9	0.700 (0.122)	0.000 (0.000)	**0.900** (0.100)	0.000 (0.000)
	1.0	**1.000** (0.000)	0.200 (0.122)	**1.000** (0.000)	0.100 (0.100)
0.3	0.0	**0.300** (0.200)	0.100 (0.100)	0.100 (0.100)	0.100 (0.100)
	0.3	**0.400** (0.187)	**0.400** (0.187)	0.100 (0.100)	0.100 (0.100)
	0.7	**0.700** (0.200)	0.100 (0.100)	0.400 (0.187)	0.100 (0.100)
	0.9	**1.000** (0.000)	0.200 (0.200)	0.800 (0.122)	0.000 (0.000)
	1.0	**1.000** (0.000)	0.600 (0.187)	0.800 (0.200)	0.100 (0.100)
0.6	0.0	**0.800** (0.122)	0.600 (0.245)	0.000 (0.000)	0.100 (0.100)
	0.3	**1.000** (0.000)	**1.000** (0.000)	0.100 (0.100)	0.100 (0.100)
	0.7	**1.000** (0.000)	**1.000** (0.000)	0.600 (0.245)	0.100 (0.100)
	0.9	**1.000** (0.000)	0.700 (0.200)	0.600 (0.245)	0.000 (0.000)
	1.0	**1.000** (0.000)	**1.000** (0.000)	**1.000** (0.000)	0.100 (0.100)
0.9	0.0	**1.000** (0.000)	**1.000** (0.000)	0.000 (0.000)	0.100 (0.100)
	0.3	**1.000** (0.000)	**1.000** (0.000)	0.100 (0.100)	0.100 (0.100)
	0.7	**1.000** (0.000)	**1.000** (0.000)	0.600 (0.187)	0.100 (0.100)
	0.9	**1.000** (0.000)	**1.000** (0.000)	0.500 (0.224)	0.000 (0.000)
	1.0	**1.000** (0.000)	**1.000** (0.000)	0.900 (0.100)	0.100 (0.100)

Table 5 (continued). **Transfer-mechanism ablations across all 20 (ρ, γ) combinations.** Entries are mean (SE); bold marks the best mean among the four GP variants, including ties at the displayed precision.

(c) Recall@20 at 200 assays ↑

ρ	γ	SPIN-BO (Both)	SPIN-BO (Response Transfer)	SPIN-BO (Representation Transfer)	Target-only BO
0.0	0.0	0.045 (0.027)	**0.050** (0.014)	0.030 (0.005)	0.045 (0.015)
	0.3	0.040 (0.019)	0.040 (0.017)	0.010 (0.006)	**0.045** (0.017)
	0.7	0.145 (0.072)	0.050 (0.008)	**0.170** (0.061)	0.045 (0.012)
	0.9	**0.335** (0.088)	0.045 (0.022)	0.255 (0.035)	0.035 (0.010)
	1.0	**0.645** (0.096)	0.055 (0.032)	0.490 (0.081)	0.055 (0.012)
0.3	0.0	**0.095** (0.009)	0.085 (0.019)	0.040 (0.017)	0.045 (0.015)
	0.3	**0.115** (0.026)	0.090 (0.017)	0.040 (0.006)	0.045 (0.017)
	0.7	**0.295** (0.060)	0.120 (0.020)	0.130 (0.056)	0.045 (0.012)
	0.9	**0.485** (0.075)	0.170 (0.051)	0.250 (0.085)	0.035 (0.010)
	1.0	**0.820** (0.025)	0.160 (0.059)	0.550 (0.114)	0.055 (0.012)
0.6	0.0	0.485 (0.046)	**0.530** (0.051)	0.025 (0.008)	0.045 (0.015)
	0.3	0.490 (0.026)	**0.500** (0.044)	0.055 (0.009)	0.045 (0.017)
	0.7	**0.670** (0.079)	0.575 (0.085)	0.125 (0.025)	0.045 (0.012)
	0.9	**0.810** (0.066)	0.655 (0.106)	0.350 (0.093)	0.035 (0.010)
	1.0	**0.975** (0.025)	0.735 (0.115)	0.580 (0.120)	0.055 (0.012)
0.9	0.0	0.965 (0.019)	**0.970** (0.020)	0.060 (0.017)	0.045 (0.015)
	0.3	0.975 (0.019)	**0.980** (0.020)	0.040 (0.017)	0.045 (0.017)
	0.7	**0.985** (0.006)	0.980 (0.012)	0.185 (0.030)	0.045 (0.012)
	0.9	**1.000** (0.000)	**1.000** (0.000)	0.290 (0.105)	0.035 (0.010)
	1.0	**1.000** (0.000)	**1.000** (0.000)	0.575 (0.103)	0.055 (0.012)

(d) Recall@20 at 408 assays ↑

ρ	γ	SPIN-BO (Both)	SPIN-BO (Response Transfer)	SPIN-BO (Representation Transfer)	Target-only BO
0.0	0.0	0.065 (0.022)	**0.070** (0.017)	**0.070** (0.009)	**0.070** (0.017)
	0.3	0.070 (0.018)	**0.095** (0.015)	0.035 (0.023)	0.070 (0.017)
	0.7	0.310 (0.090)	0.115 (0.015)	**0.330** (0.076)	0.120 (0.017)
	0.9	0.630 (0.096)	0.205 (0.050)	**0.650** (0.063)	0.105 (0.012)
	1.0	**0.980** (0.015)	0.105 (0.041)	0.885 (0.053)	0.130 (0.028)
0.3	0.0	**0.205** (0.030)	**0.205** (0.042)	0.070 (0.017)	0.070 (0.017)
	0.3	**0.245** (0.033)	0.170 (0.029)	0.075 (0.008)	0.070 (0.017)
	0.7	**0.525** (0.085)	0.280 (0.041)	0.360 (0.095)	0.120 (0.017)
	0.9	**0.790** (0.061)	0.345 (0.073)	0.570 (0.083)	0.105 (0.012)
	1.0	**0.995** (0.005)	0.510 (0.097)	0.870 (0.083)	0.130 (0.028)
0.6	0.0	0.655 (0.053)	**0.705** (0.044)	0.055 (0.005)	0.070 (0.017)
	0.3	0.705 (0.052)	**0.720** (0.035)	0.090 (0.017)	0.070 (0.017)
	0.7	**0.875** (0.037)	0.825 (0.091)	0.365 (0.090)	0.120 (0.017)
	0.9	**0.935** (0.030)	0.880 (0.084)	0.550 (0.084)	0.105 (0.012)
	1.0	**0.990** (0.010)	0.930 (0.047)	0.925 (0.047)	0.130 (0.028)
0.9	0.0	**1.000** (0.000)	**1.000** (0.000)	0.090 (0.015)	0.070 (0.017)
	0.3	0.985 (0.015)	**0.995** (0.005)	0.090 (0.019)	0.070 (0.017)
	0.7	**1.000** (0.000)	**1.000** (0.000)	0.400 (0.075)	0.120 (0.017)
	0.9	**1.000** (0.000)	**1.000** (0.000)	0.620 (0.077)	0.105 (0.012)
	1.0	**1.000** (0.000)	**1.000** (0.000)	0.870 (0.081)	0.130 (0.028)

C RETROSPECTIVE BZIP EVALUATION

This appendix provides additional details on data processing, biologically grouped source panels, evaluation and aggregation procedures, and the complementary single-source and biology-agnostic analyses for the retrospective bZIP experiments in Section 5.

C.1 ASSAY DATA AND CANDIDATE LIBRARIES

We use the human bZIP bindingPCA dataset of Bendel et al. (2025), based on the authors' supplementary Table S1. We retain expected, non-chimeric bait constructs with a single standard-amino-acid substitution, valid bait and partner sequences, and finite pair-level `logFC`. Following the original coverage rule, all three input read counts must exceed 10 and all output counts must be positive. The published `logFC` is used directly as the interaction response, without additional significance filtering, imputation, or partner-wise centering.

For each mutant background and partner panel, the candidate library contains variants with retained measurements across all partners in that panel. When one partner is treated as the target, its complete interaction profile is withheld during optimization and used only for retrospective evaluation; individual target responses are revealed only when queried. The remaining partner profiles are available throughout the campaign as source knowledge.

C.2 BIOLOGICALLY GROUPED SOURCE PANELS AND EVALUATION SETUP

The primary retrospective analysis uses biologically grouped source panels defined using bZIP family membership, dimerization classes, sequence relationships, and known interaction biology (Wingender et al., 2018; Vinson et al., 2002; Amoutzias et al., 2007). The C/EBP groups contain combinations of CEBPA, CEBPB, CEBPD, CEBPE, CEBPG, and DDIT3; the Fos-related group contains FOS, FOSB, FOSL1, FOSL2, ATF3, and JDP2; and the AP-1/Jun/ATF2 group contains JUN, JUNB, JUND, FOS-family members, ATF2, ATF7, and CREB5.

The evaluation uses ten panels across BATF2, CEBPB, DDIT3, and JUNB mutant backgrounds (Table 6). Within each panel, each partner serves once as the target and all remaining panel members are used jointly as sources, giving 54 multi-source discovery problems with 3–9 source partners.

Campaign settings follow Section 5. Each discovery problem uses ten seeds, 20 shared initial target assays, and nine adaptive batches of 20, for a total target budget of 200 assays. We use the same comparison methods as in Section 4, with model fitting following Section 3. Results are first averaged over seeds within each target problem, then over target problems within each panel, and finally equally across the ten panels. The descriptive 95% intervals in Figures 3 and 5 are obtained by bootstrap resampling the four mutant-background clusters.

Table 7 reports the multi-source results at the intermediate budget of 100 target assays, complementing the 200-assay results in Table 2. The intermediate-budget comparison shows that the discovery advantage of SPIN-BO (Both) is already apparent well before the final assay budget.

C.3 SINGLE-SOURCE DISCOVERY

The primary multi-source analysis asks how several related interaction profiles can guide discovery for a new target. We additionally consider the more limited setting in which only one source interaction profile is available. Every ordered source–target pair within the same ten biologically grouped panels defines a discovery problem, giving 270 single-source problems. Candidate libraries, target-assay budgets, comparison methods, and aggregation procedures otherwise follow the multi-source evaluation.

With one source partner, SPIN-BO (Both) improves over Target-only BO across the discovery trajectory (Figure 5). Its trajectory-average normalized best score is 0.966 versus 0.948 for Target-only BO, and its optimum-success rate at 200 assays is 85.8% versus 73.3%. Recall@5 and Recall@20 are 83.6% and 84.5%, respectively, compared with 76.9% for both metrics under Target-only BO. The single-source setting shows that one available source profile can already provide useful guidance, while the multi-source setting further increases SPIN-BO's optimum-success rate to 95.0%.

Table 6: **Evaluated biologically grouped bZIP panels.** Each candidate library contains variants measured across all partners in the corresponding panel.

Background	Partner panel	Partners	Mutants
BATF2	C/EBP full	6	504
BATF2	C/EBP extended	4	585
CEBPB	C/EBP core	4	571
CEBPB	C/EBP full	6	534
CEBPB	C/EBP extended	4	614
DDIT3	C/EBP core	4	533
DDIT3	C/EBP full	6	501
DDIT3	C/EBP extended	4	569
JUNB	AP-1 / Jun / ATF2	10	545
JUNB	Fos-related	6	552

Table 7: **Multi-source discovery at a target budget of 100 assays.** Equal-panel means across retrospective discovery tasks. OptSuccess denotes optimum success at 100 assays. Higher is better; bold marks the best mean, including ties at the displayed precision.

Method	Normalized best	OptSuccess	Recall@5	Recall@20
SPIN-BO (Both)	**0.990**	**0.822**	0.750	**0.723**
SPIN-BO (Response Transfer)	**0.990**	0.794	**0.753**	**0.723**
SPIN-BO (Representation Transfer)	0.973	0.620	0.595	0.593
Target-only BO	0.959	0.536	0.557	0.488
ALDE-style	0.920	0.360	0.352	0.316
EVOLVEpro-style	0.963	0.558	0.543	0.500
Random	0.892	0.170	0.181	0.178

C.4 Biology-Agnostic Source Panels

To evaluate transfer without biologically curated source groups, we construct ten additional panels without using bZIP family membership, sequence similarity, or known interaction biology. For each panel, we first randomly select a mutant background and then retain ten partners chosen to maximize the number of variants jointly observed across the panel. The resulting mutant backgrounds are ATF4, BATF2, CEBPB, CREB3, CREM, DBP, FOSB, JUNB, MAFG, and XBP1, with candidate libraries containing 572–657 variants.

We evaluate two source-count settings using the same ten panels and 100 target problems. In the 9-source setting, each partner serves in turn as the target and the remaining nine partners are used as sources. In the 5-source setting, each partner again serves as the target, while five of the remaining nine partners are sampled without replacement once for each background–target pair and held fixed across campaign seeds.

Each campaign begins with 20 shared initial target assays and adds nine batches of 20, for a total budget of 200 assays. These experiments use five campaign seeds per target, compared with ten in the biologically grouped analysis. The seven methods and four discovery metrics are the same as in Section 5. Results are averaged over seeds within each target, targets within each panel, and then equally across the ten panels. Shaded intervals are descriptive 95% percentile intervals from 2,500 bootstrap resamples of the ten mutant-background panels.

Nine-source discovery. Figure 6 shows the discovery trajectories, and Table 8 reports final-budget results. SPIN-BO (Both) improves all four metrics over Target-only BO at 200 assays. SPIN-BO (Response Transfer) attains the highest mean across the four final-budget metrics, indicating that response-level transfer provides substantial benefit in this setting. SPIN-BO (Representation Transfer) also outperforms Target-only BO across all four metrics, showing that representation-level transfer remains useful even without biologically curated source panels.

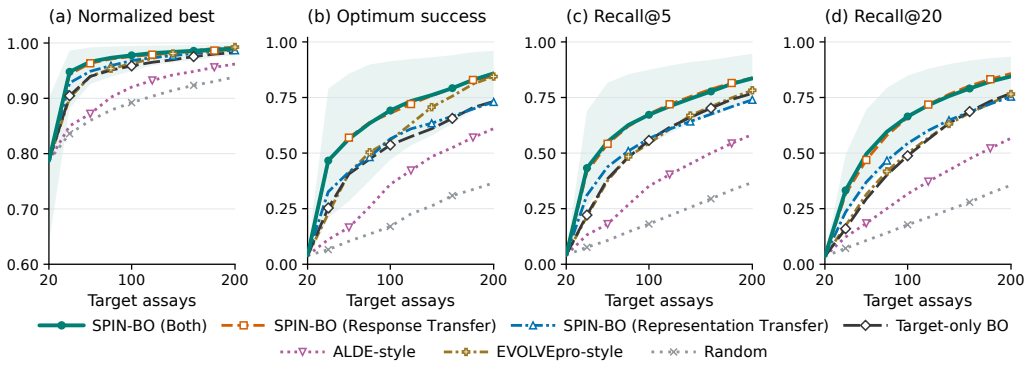

Figure 5: **Single-source bZIP discovery.** Curves show equal-panel means for normalized best score, optimum success, Recall@5, and Recall@20 as the target-assay budget increases. Shading shows descriptive 95% bootstrap intervals for SPIN-BO (Both).

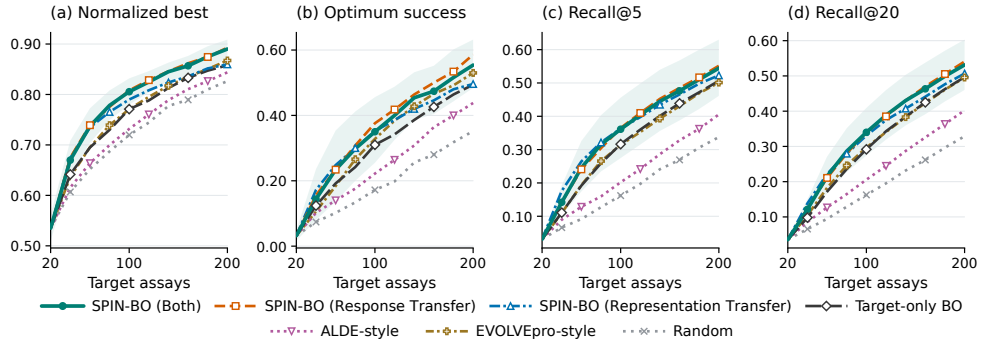

Figure 6: **Nine-source discovery with biology-agnostic panels.** Curves show equal-panel means for normalized best score, optimum success, Recall@5, and Recall@20 as the target-assay budget increases. Shading shows descriptive 95% bootstrap intervals for SPIN-BO (Both).

Five-source discovery. Figure 7 and Table 9 show the corresponding results with five sources. SPIN-BO (Both) improves over Target-only BO across the final-budget metrics, achieving the highest normalized best score and optimum-success rate, while SPIN-BO (Response Transfer) attains the highest Recall@5 and Recall@20. SPIN-BO (Representation Transfer) also improves normalized best and Recall@5 relative to Target-only BO, indicating that representation-level transfer continues to provide useful information in this setting. Overall, useful cross-partner transfer remains possible even without biologically curated source panels.

Table 8: **Nine-source discovery with biology-agnostic panels at a target budget of 200 assays.** Equal-panel means across the ten panels. OptSuccess denotes optimum success at 200 assays. Higher is better; bold marks the best mean, including ties at the displayed precision.

Method	Normalized best	OptSuccess	Recall@5	Recall@20
SPIN-BO (Both)	0.890	0.554	0.544	0.531
SPIN-BO (Response Transfer)	**0.893**	**0.586**	**0.552**	**0.541**
SPIN-BO (Representation Transfer)	0.860	0.496	0.523	0.507
Target-only BO	0.858	0.494	0.506	0.497
ALDE-style	0.844	0.438	0.406	0.402
EVOLVEpro-style	0.867	0.530	0.501	0.495
Random	0.826	0.352	0.338	0.329

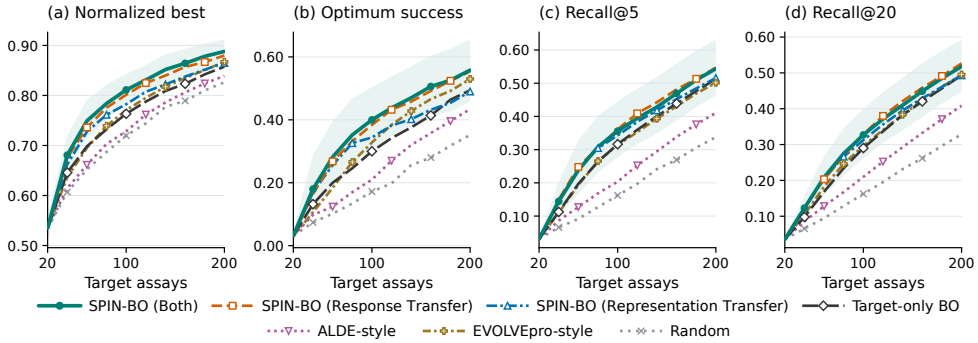

Figure 7: **Five-source discovery with biology-agnostic panels.** Curves show equal-panel means for normalized best score, optimum success, Recall@5, and Recall@20 as the target-assay budget increases. Shading shows descriptive 95% bootstrap intervals for SPIN-BO (Both).

Table 9: **Five-source discovery with biology-agnostic panels at a target budget of 200 assays.** Equal-panel means across the ten panels. OptSuccess denotes optimum success at 200 assays. Higher is better; bold marks the best mean, including ties at the displayed precision.

Method	Normalized best	OptSuccess	Recall@5	Recall@20
SPIN-BO (Both)	**0.888**	**0.558**	0.544	0.518
SPIN-BO (Response Transfer)	0.880	0.554	**0.547**	**0.526**
SPIN-BO (Representation Transfer)	0.865	0.490	0.516	0.493
Target-only BO	0.858	0.494	0.513	0.493
ALDE-style	0.839	0.434	0.411	0.408
EVOLVEpro-style	0.867	0.530	0.501	0.495
Random	0.826	0.352	0.338	0.329