CELLREPROGRAMMER: PHENOTYPE-GUIDED THERAPEUTIC TARGET DISCOVERY FOR CELL REPROGRAMMING

Anonymous authors
Paper under double-blind review

ABSTRACT

Therapeutic target discovery requires interventions that improve disease-associated cellular phenotypes while remaining feasible and limiting disruption to other cellular programs. Existing methods often stop at gene ranking or response prediction, while combination approaches may rely on full-transcriptome matching. To address these limitations, we introduce **CellReprogrammer**, a differentiable framework that directly optimizes combinatorial target sets while jointly considering effectiveness, feasibility, and toxicity. Specifically, CellReprogrammer couples a disease-state classifier that defines phenotype improvement with a perturbation surrogate and a differentiable combination optimizer. The optimizer selects a fixed-size target set, incorporates druggability, penalizes high-risk targets, and caps the magnitude of downstream transcriptional changes, without requiring paired disease–treatment profiles from the target cohort. Across five disorders and three cell types, phenotype-guided optimization substantially reduced downstream transcriptional disruption relative to full-transcriptome matching at matched predicted effectiveness. Across the broader benchmark, CellReprogrammer outperformed 20 baselines in predicted reprogramming effectiveness while selecting compact target sets with 93.7% druggability and lower predicted target and downstream toxicity. Finally, the selected interventions were disease- and cell-type-specific and recovered coherent biological programs and established cross-disease relationships. Together, these results establish phenotype-guided combinatorial optimization as a practical approach to therapeutic target discovery.

1 INTRODUCTION

A central challenge in therapeutic target discovery is to identify perturbations that improve disease-associated phenotypes while preserving cellular functions and identities, which requires correcting disease-associated features rather than reproducing a full control transcriptome. Compounding this challenge, dysregulation in complex diseases spans multiple gene programs and cell types, motivating multi-gene interventions. Their combined effects, however, cannot be inferred from individual perturbations because genetic interactions may be non-additive. Candidate target sets must therefore be evaluated for both phenotype improvement and disruption of other cellular programs. Accordingly, target discovery must balance three criteria: **effectiveness**, improvement in the disease-associated phenotype; **feasibility**, achieving this effect with a small set of druggable genes; and **toxicity**, the risk of perturbing essential genes or inducing excessive downstream transcriptional changes. Together, these requirements define the computational challenge of selecting target combinations that improve disease phenotypes while remaining feasible and minimally disruptive.

Computational target discovery has progressed from association-based prioritization to perturbation modeling and intervention selection. Network-based methods identify disease-relevant targets but not perturbation-induced cellular changes (Guney et al., 2016). To address this limitation, Geneformer evaluates representation shifts after in-silico perturbation, whereas scGPT and GEARS predict expression responses to perturbations, including combinations (Theodoris et al., 2023; Cui et al., 2024; Roohani et al., 2024). Response prediction, however, does not determine which interventions best improve disease phenotypes. To guide selection, scGO and SIDISH evaluate changes in

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.

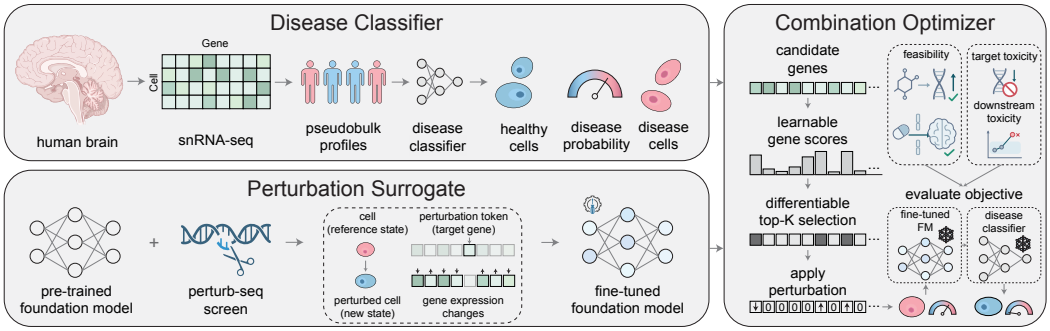

Figure 1: **Overview of CellReprogrammer.** A disease classifier learns disease-associated state variation from donor-level pseudobulk profiles and provides the phenotype objective for state reversal. A pretrained foundation model is fine-tuned on Perturb-seq data as a surrogate for transcriptional responses to candidate interventions. The combination optimizer then performs differentiable top-K selection over candidate genes, using the predicted responses to jointly optimize disease-state improvement while incorporating feasibility and toxicity constraints. Together, these components enable phenotype-guided selection of small combinatorial interventions.

predicted disease states or risk, but focus on individual targets or selected pairs (Wu et al., 2025; Jolasun et al., 2025). Extending selection to combinations, PDGrapher uses paired diseased–treated profiles and an objective including treated-transcriptome reconstruction (Gonzalez et al., 2025). Yet full-transcriptome matching may favor changes unrelated to phenotype improvement and omits explicit druggability or toxicity criteria. Thus, the remaining challenge is phenotype-guided combination selection that jointly considers effectiveness, feasibility, and toxicity without full-transcriptome matching or paired therapeutic-response profiles from the target cohort.

To address these limitations, we introduce **CellReprogrammer**, a differentiable framework for phenotype-guided therapeutic target discovery (**Fig. 1**). It optimizes target sets for disease-phenotype improvement rather than full-transcriptome matching, while jointly considering effectiveness, feasibility, and toxicity. CellReprogrammer has three components. First, a disease classifier distinguishes disease from control profiles, defining phenotype-based effectiveness without specifying a complete target expression profile. Second, a perturbation surrogate predicts joint transcriptional responses to candidate target sets, enabling evaluation of phenotype improvement and downstream expression changes. Third, a combination optimizer selects a fixed-size target set, incorporating druggability for feasibility, penalizing high-risk targets, and capping the magnitude of downstream transcriptional changes to limit predicted toxicity. Together, these components enable phenotype-guided target discovery without paired disease–treatment data from the target cohort. Our main contributions are summarized as follows.

- **Phenotype-guided combinatorial target discovery.** We formulate therapeutic target discovery as direct optimization of target combinations for disease-phenotype improvement rather than full-transcriptome matching, enabling focused reprogramming without prescribing a complete target state or requiring paired disease–treatment profiles.
- **Differentiable optimization of interacting target combinations.** We develop an end-to-end differentiable framework that couples a disease-state classifier and perturbation surrogate with relaxed top-K selection, directly optimizing joint effects that independent gene ranking cannot capture under distributed dysregulation and non-additive interactions.
- **Joint optimization of effectiveness, feasibility, and toxicity.** We integrate all three therapeutic requirements into target selection: phenotype improvement drives effectiveness, compact druggable sets support feasibility, and target-risk penalties and a transcriptomic-change cap limit predicted toxicity.
- **Broad empirical and biological evidence.** Across five disorders, three cell types, and 20 baselines, CellReprogrammer achieves a 38% stronger predicted reprogramming effect while selecting three-target sets with 93.7% druggability and substantially lower predicted toxicity, including 87% lower target risk. The interventions are context-specific and recover coherent biological programs and cross-disease relationships.

2 Related Work

Therapeutic target discovery. Conventional approaches prioritize disease-relevant genes using differential expression (Love et al., 2014), classifier-based feature selection (Guyon et al., 2002), disease associations (Piñero et al., 2020), and network relationships (Guney et al., 2016). To move beyond association alone, transcriptional connectivity matches disease signatures to measured perturbation signatures, although coverage depends on available libraries (Lamb et al., 2006; Subramanian et al., 2017). More recent approaches extend response-based assessment through foundation-model in-silico perturbation (Zhang et al., 2026), phenotypic reversion in Perturb-seq (Wong et al., 2026), and module-aware simulation (Sun & Zhou, 2026). Building on these advances, phenotype-based methods guide prioritization more directly: scGO evaluates individual perturbations through predicted disease-state changes (Wu et al., 2025), whereas SIDISH ranks knockouts by reductions in predicted high-risk cells and examines selected pairs (Jolasun et al., 2025). These methods link perturbation effects to disease phenotypes but do not optimize fixed-size target sets under joint therapeutic criteria. Moving to combinations, PDGrapher uses paired diseased–treated profiles with an objective including treated-transcriptome reconstruction (Gonzalez et al., 2025). CellReprogrammer instead optimizes fixed-size target sets through a perturbation surrogate, guided by predicted phenotype improvement rather than full-transcriptome matching and without paired profiles from the target cohort. Its optimization jointly considers effectiveness, feasibility, and toxicity through phenotype guidance, druggability and target-risk penalties, and a transcriptomic-change cap.

Single-cell perturbation modeling. Early methods such as scGen (Lotfollahi et al., 2019) predict transcriptional responses through latent-space shifts. CPA (Lotfollahi et al., 2023) extends this idea to perturbation combinations, doses, and cellular contexts through compositional representations, while GEARS (Roohani et al., 2024) incorporates gene–gene relationships to predict unseen single- and multigene perturbations. Complementing these approaches, CellOT learns transport maps between unpaired untreated and perturbed cell populations (Bunne et al., 2023). In parallel, scGPT, Geneformer, scFoundation, UCE, and CellPLM learn transferable cellular representations (Cui et al., 2024; Theodoris et al., 2023; Hao et al., 2024; Rosen et al., 2026; Wen et al., 2024), while State integrates such representations with perturbation-response modeling across contexts (Adduri et al., 2026). Building on these advances, CellReprogrammer uses a replaceable, cell-type-specific surrogate to estimate joint perturbation responses while separating response prediction from target selection, enabling evaluation of both phenotype improvement and downstream transcriptional disruption.

Differentiable subset selection. Gumbel–Softmax and Concrete relaxations enable gradient-based optimization over categorical choices (Jang et al., 2017; Maddison et al., 2016). Building on this principle, reparameterizable sampling without replacement extends differentiable optimization to fixed-size subsets, allowing discrete selections to be optimized jointly (Xie & Ermon, 2019). In biological applications, MetaVelo combines such subset selection with a differentiable surrogate to identify gene sets influencing developmental cell-fate transitions (Liu et al., 2025). Extending this formulation to therapeutic target discovery, CellReprogrammer uses relaxed top-K selection with a phenotype-guided objective that jointly considers effectiveness, feasibility, and toxicity.

3 The CellReprogrammer Framework

CellReprogrammer combines a disease classifier, a perturbation surrogate, and relaxed top-K selection to design cell-type-specific interventions (**Fig. 1**). The classifier and surrogate are complementary components that jointly support combination optimization: the surrogate predicts expression responses, and the classifier scores the resulting disease state. We first define the problem and objective, then describe these components.

3.1 Problem Definition and Method Overview

Formally, for a fixed cell type, let $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^{n}$, $\mathbf{x}_i \in \mathbb{R}^m$, denote the expression profiles of n nuclei over m genes, collected from D donors with diagnostic labels $y_d \in \{0, 1\}$, where $y_d = 1$ denotes disease. We write $\mathcal{C}_d = \{i : d(i) = d\}$ for the nuclei from donor d.

Let $\mathcal{G} \subseteq \{1, \ldots, m\}$ denote the genes available for intervention. Our goal is to identify a size-K subset $S \subseteq \mathcal{G}$ whose predicted perturbation shifts the disease-associated state toward control:

$$S^\star = \underset{S \subseteq \mathcal{G},\, |S|=K}{\arg\min} \; \mathcal{L}(S), \tag{1}$$

where $\mathcal{L}$ measures the disease state after perturbation. Let $h_\psi(\cdot)$ denote the disease classifier and $\tilde{\mathbf{p}}_d(S)$ the predicted transcriptomic state of disease donor d after applying the perturbation set S. The effectiveness (disease-state reversal) objective is

$$\mathcal{L}(S) = \frac{1}{|\mathcal{D}_1|} \sum_{d \in \mathcal{D}_1} \mathrm{BCE}\big(h_\psi(\tilde{\mathbf{p}}_d(S)),\, 0\big), \qquad \mathcal{D}_1 = \{d : y_d = 1\}, \tag{2}$$

The optimizer combines response predictions of g_ϕ with disease-state scores from h_ψ to select S, minimizing this control-directed loss with druggability and target-toxicity penalties (§3.4).

3.2 Disease Classifier

For each donor d, we aggregate the nuclei of the given cell type into a pseudobulk profile $\mathbf{p}_d$ and fit an ℓ_2-regularized logistic classifier h_ψ that separates disease from control donors (binary cross-entropy; full objective in Appendix A.1),

$$\mathbf{p}_d = \frac{1}{|\mathcal{C}_d|} \sum_{i \in \mathcal{C}_d} \mathbf{x}_i \in \mathbb{R}^m, \qquad h_\psi(\mathbf{p}_d) = \sigma(\mathbf{w}^\top \mathbf{p}_d + b). \tag{3}$$

This matches the donor-level diagnosis and avoids treating nuclei from the same donor as independent labeled samples. Each gene gets an intervention direction from the disease–control pseudobulk difference $\mathbf{v} = \bar{\mathbf{p}}^{\text{dis}} - \bar{\mathbf{p}}^{\text{ctrl}}$ (inhibit genes elevated in disease, INH, $v_g > 0$; activate those reduced, ACT, $v_g < 0$). We also augment training with randomly perturbed pseudobulks kept under their donor's original label (Appendix A.1), so the classifier discounts the generic shift shared by all such perturbations and bases its decision on *which* genes are targeted, not on whether any intervention was applied.

3.3 Perturbation Surrogate

Learned perturbation surrogate. We fine-tune a single-cell foundation model into a surrogate $g_\phi : \mathbb{R}^m \times \{0, 1, 2\}^m \to \mathbb{R}^m$, where each entry of $\mathbf{f}$ encodes no intervention (0), ACT (1), or INH (2). We use a *residual* parameterization

$$g_\phi(\mathbf{x}, \mathbf{f}) = \mathbf{x} + \Delta_\phi(\mathbf{x}, \mathbf{f}), \tag{4}$$

so Δ_ϕ captures the perturbation-induced change, trained on Perturb-seq responses to match the *flag-attributable* change $\Delta_\phi(\mathbf{x}, \mathbf{f}) - \Delta_\phi(\mathbf{x}, \mathbf{0})$ to the measured effect (full objective in Appendix A.2).

Surrogate validation. We evaluate the surrogate on Perturb-seq targets held out from fine-tuning. For each assayed gene $g \in \mathcal{G}_{\text{scr}} \subseteq \mathcal{G}$, ρ_Δ correlates predicted flag-induced and measured changes over the 20 genes with the largest observed responses (Eq. (19); Appendix D.1). Table 1 also reports direction accuracy (Acc.) on these genes.

3.4 Combination Optimizer

The combination optimizer optimizes a K-gene intervention using Eq. (2) and the druggability and target-toxicity penalties below, with the classifier and surrogate fixed (Algorithm 1).

Differentiable set selection. We use a continuous top-K relaxation for gradients, while a straight-through estimator applies a hard mask of exactly K genes in the forward pass (Xie & Ermon, 2019). Below, $\mathbf{s}$ denotes this straight-through mask unless stated otherwise (Appendix A.3).

Perturbed-state construction. Each forward-pass $s_g \in \{0, 1\}$ indicates whether gene g is selected in its disease–control-derived direction. The surrogate predicts the change per nucleus; averaging over a donor's nuclei gives the perturbed donor state scored by the classifier:

$$\tilde{\mathbf{p}}_d(\mathbf{s}) = \mathbf{p}_d + \mathrm{cap}_\kappa\big(\boldsymbol{\delta}_d(\mathbf{s})\big), \qquad \boldsymbol{\delta}_d(\mathbf{s}) = \frac{1}{|\mathcal{C}_d|} \sum_{i \in \mathcal{C}_d} \Big[g_\phi\big(\mathbf{x}_i,\, \mathbf{f}_\mathbf{s}\big) - g_\phi\big(\mathbf{x}_i, \mathbf{0}\big) \Big]. \tag{5}$$

Table 1: **Patient data, screen coverage, and model validation.** EXN: excitatory neurons; MG: microglia; AST: astrocytes; IN: inhibitory neurons; OLG: oligodendrocytes; OPC: oligodendrocyte precursors; END: endothelial cells. Bold CTs enter the main intervention comparison.

	Patient nuclei (cohort donors: H/D)					Disease classifier	Perturb-seq / response surrogate					
CT	SCZ (53/47)	BPD (24/24)	ASD (38/40)	MDD (38/32)	PTSD (38/35)	AUROC	Sources	Perturb.	Cells	Controls	ρ_Δ	Acc.
EXN	215,941	234,400	190,304	237,546	238,660	0.69	3	284	78,371	9,720	0.71	89%
MG	4,887	5,392	25,175	17,902	20,248	0.62	3	76	84,093	12,696	0.55	78%
AST	30,501	20,710	55,312	94,676	100,928	0.65	2	62	100,745	10,639	0.47	72%
IN	85,870	68,272	80,193	84,418	84,568	0.70	2	38	7,504	2,791	0.34	17%
OLG	96,591	58,589	88,784	135,831	116,689	0.62	1	34	4,252	2,392	0.04	53%
OPC	13,636	11,469	41,693	38,440	39,919	0.66	0	0	0	0	—	—
END	7,249	1,512	18,055	16,393	16,945	0.63	0	0	0	0	—	—

H/D gives cohort control/disease donor totals; MDD/PTSD share controls. Nuclei counts retain donors with ≥ 8 per CT. AUROC averages five disorders (Table B.2); Cells includes controls. ρ_Δ and direction accuracy retain the original held-out-target estimates, with source-matching and amplitude limitations (Appendix D.1); OPC/END have no matched screen. Data details: Appendix B.

Here $\mathbf{f_s}$ denotes the direction-specific intervention induced by the mask; the relaxation supplies gradients, not fractional discrete flag codes. Differencing against $g_\phi(\mathbf{x}_i, \mathbf{0})$ anchors the response to the measured pseudobulk, so $\mathbf{s}=\mathbf{0}$ recovers $\mathbf{p}_d$ (Appendix A.3). The operator cap_κ bounds $\|\boldsymbol{\delta}_d(\mathbf{s})\|$ by κ times the disease→control shift (Appendix A.3).

Feasibility and toxicity. Feasibility combines a small intervention size K with druggability, quantified by a drug-realizability score $R_g \in [0, 1]$ that measures support for moving gene g in the required direction using LINCS L1000 signatures of brain-penetrant compounds (Appendix A.3). Target toxicity is penalized by $t_g \geq 0$, a proxy for inhibition risk based on DepMap knockout lethality and gnomAD intolerance (Appendix A.3). Downstream toxicity is controlled by the transcriptomic-change cap κ, which bounds the perturbation norm and limits excessive expression changes (Eq. (5)). Thus, t_g constrains *which* genes are selected, whereas κ constrains *how far* the predicted transcriptome may move. We evaluate them separately as target and downstream toxicity (Appendix D.1). The combination optimizer balances reprogramming effectiveness, druggability, and target risk:

$$\mathcal{L}_{\mathrm{c}}(\mathbf{s}) = \mathcal{L}(\mathbf{s}) + \frac{1}{K}\sum_g s_g\left[\lambda_{\mathrm{drug}}(1 - R_g) + \lambda_{\mathrm{tox}}\, t_g\right], \tag{6}$$

where $\mathcal{L}$ is the effectiveness objective of Eq. (2), evaluated on the magnitude-capped transcriptomic state, and $\lambda_{\mathrm{drug}}, \lambda_{\mathrm{tox}} \geq 0$ weight the druggability and target-toxicity penalties, respectively.

We repeat the combination optimizer across seeds and return the K most frequently selected genes, reporting the frequencies themselves as a stability diagnostic (Appendix A.3).

4 Results

4.1 Experimental setup

Experimental overview. We evaluate CellReprogrammer against 20 baselines across three criteria—effectiveness, feasibility, and toxicity—and use ablation analyses to assess the contribution of individual components. The evaluation spans snRNA-seq data from five disorders: schizophrenia (SCZ), bipolar disorder (BPD), autism spectrum disorder (ASD), major depressive disorder (MDD), and post-traumatic stress disorder (PTSD) (Emani et al., 2024; Hwang et al., 2025). For intervention benchmarking, we retain three cell types—excitatory neurons (EXN), microglia (MG), and astrocytes (AST)—based on held-out perturbation-response prediction accuracy in matched Perturb-seq screens (**Table 1**; Appendix B). Their disease relevance is further supported by previous transcriptomic studies (Velmeshev et al., 2019; Hwang et al., 2025).

Baselines. We compare 20 baselines spanning association-based ranking, perturbation-based ranking, and combinatorial selection (**Table 2**; Appendix C). The two ranking approaches score genes independently and select the top K, whereas combinatorial methods evaluate gene sets directly. Association-based methods rank genes using differential expression, classifier weights, network

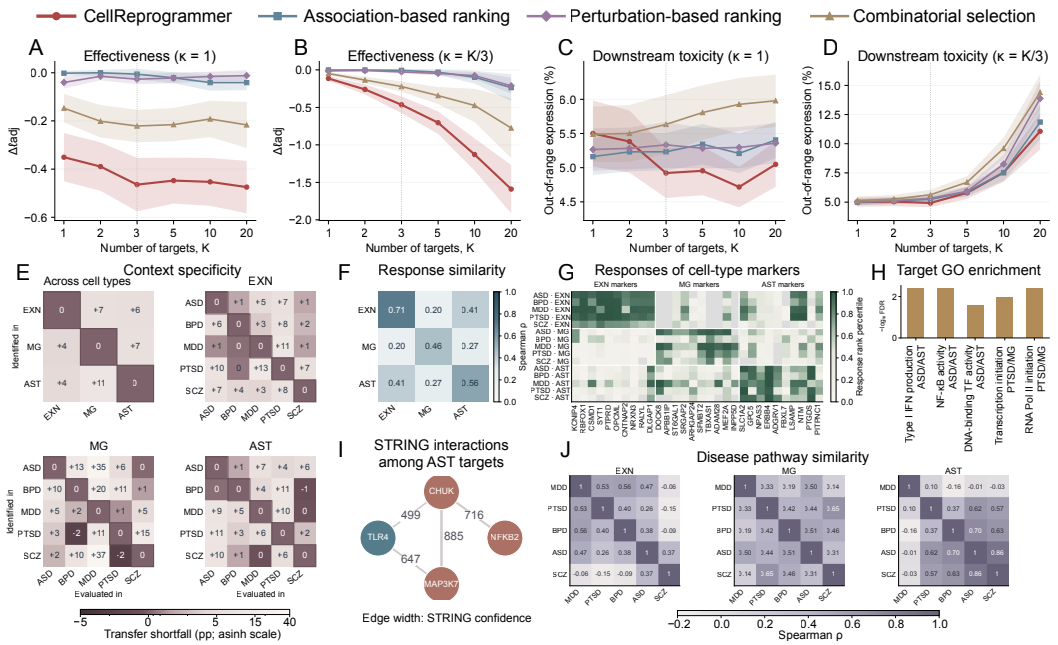

Figure 2: **Predicted reprogramming effectiveness, context specificity, and biological relevance.** *(A–D)* Reprogramming effectiveness and downstream toxicity across intervention sizes under fixed and size-dependent transcriptomic-change caps (mean±SE). *(E–G)* Disease- and cell-type-specific interventions and downstream responses. *(H,I)* Functional enrichment and protein interactions among selected targets. *(J)* Shared disease patterns from downstream responses (Appendix D.4).

proximity, or disease annotations (Love et al., 2014; Guyon et al., 2002; Guney et al., 2016; Piñero et al., 2020). To incorporate intervention effects, perturbation-based methods use scGO disease-classifier reversal (Wu et al., 2025), measured single-gene effects, Connectivity-Map signature reversal (Lamb et al., 2006), or responses predicted by GEARS, scGen, CPA, and single-cell representation models (Roohani et al., 2024; Lotfollahi et al., 2019; 2023; Cui et al., 2024; Hao et al., 2024; Theodoris et al., 2023; Wen et al., 2024; Heimberg et al., 2025). Finally, combinatorial methods include PDGrapher (Gonzalez et al., 2025) and four direct gene-set searches: simulated annealing, genetic search, random search, and greedy selection. Because paired disease–treatment profiles are unavailable, PDGrapher uses disease profiles paired with surrogate-predicted perturbations as proxy treatment data, with the perturbed gene as the intervention label.

4.2 CellReprogrammer improves predicted reprogramming effectiveness

We first benchmark CellReprogrammer against 20 baselines across intervention sizes, and then examine the context specificity and biological relevance of the selected interventions. Across these analyses, CellReprogrammer predicts stronger reprogramming effects, achieves most of its predicted effect with three targets, and identifies disease- and cell-type-specific interventions that recover relevant biological programs.

Effectiveness across intervention sizes. At a fixed downstream transcriptomic-change budget ($\kappa = 1$), CellReprogrammer predicts stronger reprogramming effects than all three baseline categories across intervention sizes ($K = 1, 2, 3, 5, 10, 20$; effect measured by control-adjusted disease-classifier logit change, $\Delta\ell_{\text{adj}}$, with more negative values indicating stronger reprogramming; **Fig. 2A**; Appendix D.1). This advantage emerges at $K = 1$ and grows through $K = 3$, where most of the predicted effect is achieved (**Fig. 2A**). We next asked whether larger intervention sets could yield further gains when allowed to induce larger transcriptomic changes. Relaxing the cap with intervention size ($\kappa = K/3$; **Fig. 2B–D**) further improves predicted effectiveness, but at a clear cost in downstream toxicity: the fraction of post-perturbation expression values outside the reference range increases from 4.92% at $K = 3$ to 11.1% at $K = 20$, compared with 4.7–5.5% under the fixed cap

Table 2: **Predicted effectiveness at $K = 3$.** Overall: mean±SE; Twd.: settings with $\Delta\ell_{\text{adj}} < 0$. Bold/dark/light: ranks 1/2/3 (20 baselines + penalized CellReprogrammer; penalty-free unranked).

	EXN					MG					AST						
Method	ASD	BPD	MDD	PTSD	SCZ	ASD	BPD	MDD	PTSD	SCZ	ASD	BPD	MDD	PTSD	SCZ	Overall	Twd.
Association-based ranking																	
DE	+0.08	−0.24	+0.01	+0.07	+0.05	−0.00	+0.03	+0.02	+0.00	−0.03	+0.01	−0.00	−0.01	+0.02	−0.01	−0.00 ± 0.02	6/15
clf-w	−0.14	+0.18	−0.09	+0.04	+0.11	−0.00	−0.00	−0.77	+0.00	−0.03	−0.08	+0.00	−0.00	+0.02	−0.01	−0.05 ± 0.05	9/15
network prox.	+0.09	−0.32	−0.21	−0.07	−0.05	−0.00	+0.57	+0.02	+0.00	+0.44	+0.29	−0.09	+0.23	+0.03	+0.26	+0.08 ± 0.06	6/15
risk (DisGeNET)	−0.02	−0.27	−0.26	+0.06	+0.28	−0.00	−0.23	+0.02	+0.00	−0.34	+0.01	+0.01	−0.00	+0.02	−0.01	−0.05 ± 0.04	8/15
Perturbation-based ranking																	
scGO	−0.09	−0.33	−0.26	+0.03	+0.09	−0.00	+0.03	+0.02	+0.00	−0.03	−0.00	+0.00	−0.00	+0.02	−0.01	−0.04 ± 0.03	8/15
linear-rank	−0.28	+0.10	−0.19	+0.04	−0.08	−0.21	**−0.98**	−1.51	−0.20	**−0.34**	−0.06	−0.11	+0.01	−0.32	−0.03	−0.28 ± 0.11	12/15
CMap	−0.16	+0.22	−0.29	+0.09	−0.07	−0.13	−0.66	−1.21	+0.35	−0.15	+0.07	−0.12	+0.05	**−0.58**	−0.03	−0.17 ± 0.10	10/15
GEARS	−0.21	+0.12	−0.31	+0.11	+0.02	−0.09	−0.49	+1.79	−0.12	+0.30	−0.23	−0.02	+0.06	−0.23	+0.40	+0.07 ± 0.14	8/15
scGen	−0.14	+0.21	+0.02	−0.01	+0.07	−0.19	+0.44	−0.45	−0.31	+0.43	+0.35	+0.11	−0.03	−0.28	+0.10	+0.02 ± 0.07	7/15
CPA	−0.03	+0.00	+0.04	−0.02	+0.06	+0.08	−0.85	−0.25	+0.31	+0.43	+0.35	−0.15	−0.04	−0.29	+0.07	−0.02 ± 0.08	7/15
scGPT	−0.28	+0.08	+0.13	**−0.34**	−0.08	−0.00	+0.03	+1.76	+0.00	−0.03	+0.01	−0.08	−0.00	+0.02	+0.22	+0.10 ± 0.12	7/15
scFoundation	−0.07	+0.24	+0.11	+0.09	+0.09	−0.00	+0.03	+0.02	+0.00	−0.03	+0.01	+0.01	−0.00	+0.02	−0.01	+0.03 ± 0.02	5/15
Geneformer	+0.08	+0.18	−0.15	+0.07	−0.07	−0.00	+0.03	+0.02	+0.00	−0.03	+0.01	+0.01	−0.01	+0.02	−0.01	+0.01 ± 0.02	6/15
CellPLM	−0.12	−0.31	+0.04	+0.08	+0.09	−0.00	+0.03	+0.02	+0.00	−0.03	+0.01	+0.01	−0.00	+0.02	−0.01	−0.01 ± 0.02	6/15
SCimilarity	+0.09	−0.31	−0.10	−0.07	+0.07	−0.00	+0.03	+0.02	+0.00	+0.08	+0.01	+0.01	−0.00	+0.02	−0.01	−0.01 ± 0.02	6/15
Combinatorial selection																	
PDGrapher	−0.12	−0.34	−0.20	−0.06	+0.05	−0.00	−0.85	+0.01	+0.00	+0.42	+0.00	−0.13	−0.00	+0.02	+0.25	−0.06 ± 0.07	8/15
annealing	−0.16	+0.05	−0.13	+0.03	−0.04	−0.04	−0.21	−1.25	−0.05	−0.15	−0.44	−0.04	+0.00	−0.23	−0.18	−0.19 ± 0.08	12/15
genetic	−0.14	−0.01	−0.22	+0.01	**−0.09**	−0.11	−0.40	−0.88	−0.05	−0.05	−0.37	−0.07	−0.02	−0.21	−0.12	−0.18 ± 0.06	14/15
random	−0.21	−0.12	−0.25	−0.04	−0.09	−0.17	−0.90	−1.42	−0.33	−0.20	−0.50	−0.11	−0.07	−0.41	−0.20	−0.33 ± 0.10	15/15
greedy	−0.26	**−0.47**	−0.03	−0.08	−0.01	**−0.38**	−0.71	−1.52	−0.44	−0.24	+0.01	−0.19	−0.20	−0.29	**−0.24**	−0.34 ± 0.10	14/15
CellReprogrammer	−0.44	−0.46	−0.45	−0.47	−0.28	−0.34	−1.05	−1.96	−0.81	−0.29	−0.70	−0.36	−0.72	−0.64	−0.42	−0.63 ± 0.11	15/15
+ **target penalties**	**−0.31**	−0.42	**−0.35**	−0.19	−0.08	−0.23	−0.80	**−1.85**	**−0.52**	−0.29	**−0.65**	**−0.22**	**−0.42**	−0.49	−0.15	−0.47 ± 0.11	15/15

(**Fig. 2B–D**). Thus, larger perturbations can produce stronger predicted effects, but with substantially greater transcriptional disruption. This trade-off motivates $K = 3$ and $\kappa = 1$ as the default setting.

Benchmarking at $K = 3$. Following the intervention-size analysis, we compared methods at $K = 3$ and $\kappa = 1$ to assess predicted effectiveness under a common intervention budget. With target constraints included, CellReprogrammer achieved the strongest mean reprogramming effect ($\Delta\ell_{\text{adj}} = -0.47$) and improved over matched random perturbations in all 15 disease–cell-type settings (**Table 2**). This advantage extended across all three baseline categories: the strongest association-based, perturbation-based, and combinatorial methods achieved mean effects of -0.05, -0.28, and -0.34, respectively (**Table 2**). Removing target constraints further strengthened the mean effect to -0.63 (**Table 2**). Thus, although these constraints reduced the effect magnitude, CellReprogrammer retained an effectiveness advantage while accounting for feasibility and toxicity.

Context-specific targets and responses. We next tested whether CellReprogrammer identifies interventions specific to disease and cell type. At the target level, context-matched interventions produced stronger predicted effects than transferred target sets in 84 of 90 comparisons (**Fig. 2E**). At the downstream level, predicted responses retained cell-type structure across diseases, with correlations of 0.71, 0.46, and 0.56 within EXN, MG, and AST, respectively (**Fig. 2F**). Cell-type marker genes showed corresponding patterns (**Fig. 2G**). Together, these results support context-specific target selection and characteristic downstream responses within each cell type (Appendix D.3).

Selected interventions recover coherent biological programs. We further asked whether the selected interventions were biologically coherent at both the target and downstream-response levels. At the target level, selected genes were enriched for NF-κB and type I interferon regulation in ASD–AST and transcription initiation in PTSD–MG (global FDR < 0.05; **Fig. 2H**; Appendix D.4). Protein-interaction analysis provided complementary support, revealing connections among proteins encoded by *TLR4*, *CHUK*, *NFKB2*, and *MAP3K7* in astrocytes (**Fig. 2I**). Beyond the selected targets, predicted downstream responses also recovered established cross-disease relationships. Pathway-level response profiles were correlated between MDD and PTSD in EXN ($\rho = 0.53$) and between ASD and SCZ in AST ($\rho = 0.86$; **Fig. 2J**), consistent with shared molecular patterns reported in post-mortem transcriptomic studies (Daskalakis et al., 2024; Gandal et al., 2018). Together, these findings support the biological relevance of the selected interventions at both the target and downstream-response levels.

4.3 CellReprogrammer remains effective with feasibility and low toxicity

Having established CellReprogrammer's reprogramming effectiveness, we next asked whether this advantage can be retained as intervention requirements become more stringent, from compact and druggable target sets to lower target and downstream toxicity.

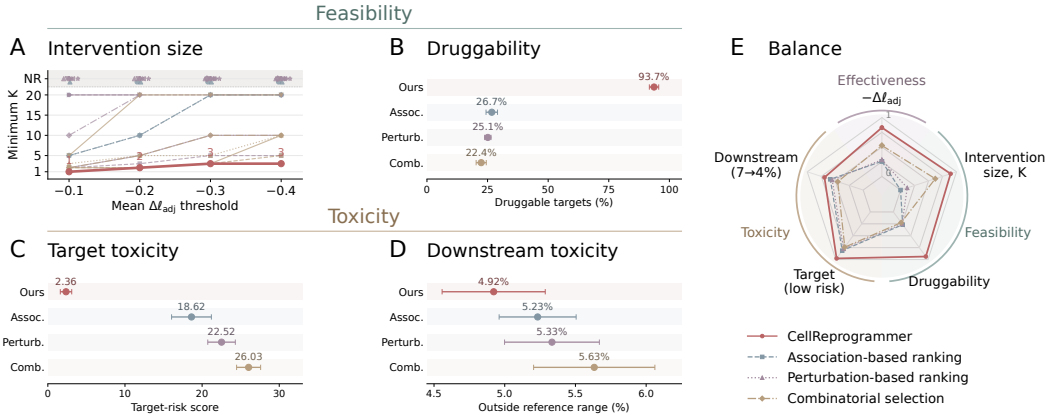

Figure 3: **Effectiveness, feasibility, and toxicity trade-offs.** *(A)* Minimum intervention size required to reach a common reprogramming-effect threshold ($\kappa = K/3$; NR, not reached). *(B–D)* Druggability, target-risk score, and downstream toxicity at the default setting ($K = 3$, $\kappa = 1$; mean $\pm$ SE). *(E)* Joint five-metric comparison. Radar scales: Appendix D.1.

Feasibility. We first asked how many targets are required to reach a common effectiveness threshold. CellReprogrammer reached $\Delta\ell_{\text{adj}} \leq -0.30$ with three targets, whereas measured single-gene ranking, CMap, and classifier-weight ranking required five, ten, and twenty targets, respectively; half of the baselines ($10/20$) did not reach the threshold within $K \leq 20$ (**Fig. 3A**). Compactness alone, however, does not ensure feasibility if the selected targets are difficult to drug. We therefore compared target druggability at a fixed size of $K = 3$. CellReprogrammer achieved a druggable target fraction of 93.7%, compared with 9–38% across the 20 baselines (**Fig. 3B**; Appendices A.3 and D.1). Thus, CellReprogrammer requires fewer targets to reach a common predicted-effect threshold and at $K = 3$, it also achieves a higher fraction of targets with direction-matched LINCS support.

Toxicity. Feasible interventions may still be undesirable if they inhibit high-risk genes or induce excessive downstream changes. We therefore evaluated toxicity at both levels using the mean inhibition-associated risk score and the fraction of expression values outside the reference range (Appendices A.3 and D.1). CellReprogrammer yielded a mean target-risk score of 2.36, compared with 18.6–26.0 across baseline categories (**Fig. 3C**). Its downstream expression-outlier fraction was also lower, at 4.92% versus 5.23–5.63% (**Fig. 3D**). Thus, after meeting the effectiveness and feasibility requirements, CellReprogrammer further reduces both target-associated risk and downstream transcriptional disruption. Together, these results show that its effectiveness is achieved with interventions that are not only compact and druggable, but also lower in both forms of toxicity.

Overall balance. Having established the feasibility and toxicity advantages, we finally quantified the effectiveness cost of imposing these constraints. Adding feasibility and target-risk penalties reduced the mean predicted effect from $\Delta\ell_{\text{adj}} = -0.63$ to -0.47, yet the full framework remained more effective than the strongest baseline (-0.34). Thus, at the default $K = 3$ setting, CellReprogrammer retains an effectiveness advantage while selecting compact, druggable target sets with lower toxicity-related scores (**Fig. 3E**).

4.4 Ablation of key design choices

We examined four design choices that determine what CellReprogrammer optimizes, how targets are selected, how perturbation effects are modeled, and how therapeutic constraints enter the objective.

Optimization objective. We first tested whether phenotype guidance reduces unnecessary transcriptional changes relative to control-profile matching. At matched reprogramming effectiveness, phenotype guidance produced substantially fewer large downstream changes (approximately 2 versus 10 genes exceeding twofold change at $\Delta\ell_{\text{adj}} = -0.1$; **Fig. 4A**), with the same pattern across additional matched-effect levels (Appendix D.2). Thus, optimizing disease-state reversal rather than transcriptome-wide similarity preserves reprogramming effect with less downstream disruption.

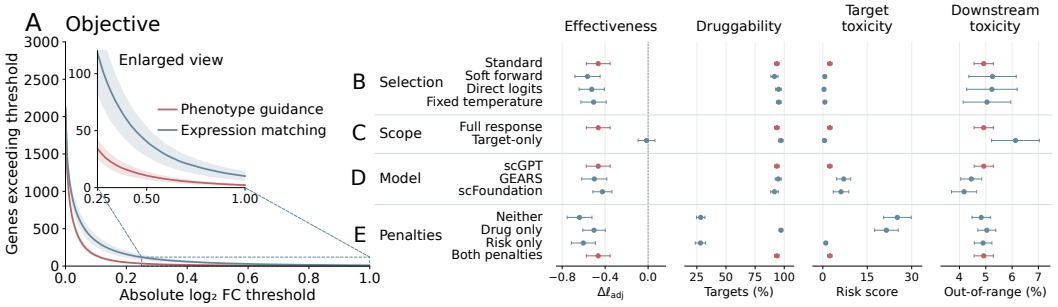

Figure 4: **Ablation of key design choices.** *(A)* Downstream transcriptional changes under phenotype guidance and control-profile matching at matched predicted effectiveness ($\Delta\ell_{\text{adj}} = -0.1$). *(B)* Target-selection variants with different forward masks, gene-score parameterizations, and temperature schedules. *(C)* Full-response versus target-only selection. *(D)* Response-model substitutions using scGPT, GEARS, and scFoundation. *(E)* Individual and combined effects of druggability and target-risk penalties. Values show mean±SE across disease–cell-type contexts.

Differentiable target selection. We varied the forward mask, gene-score parameterization, and temperature schedule of differentiable top-K selection. Soft forward yielded stronger effects but lower druggability ($\Delta\ell_{\text{adj}} = -0.564$, 91.6%; standard reference: -0.47, 93.7%; **Fig. 4B**).

Perturbation-response modeling. Target selection also depends on accurately modeling the consequences of each candidate intervention. We therefore tested both the scope and the model used for response prediction. Target-only selection yielded a much weaker effect ($\Delta\ell_{\text{adj}} = -0.015$; **Fig. 4C**), showing the importance of changes beyond the perturbed genes. Replacing scGPT with GEARS or scFoundation yielded mean $\Delta\ell_{\text{adj}} = -0.50$ and -0.43, respectively, across 15 contexts (scGPT reference: -0.47; **Fig. 4D**; Table 2). Druggability remained similar across models, while scGPT had the lowest mean target-risk score (2.36 versus 7.05 and 6.14). We therefore retain scGPT as the default surrogate while the framework remains effective across alternative response models.

Target penalties. Finally, we tested whether the feasibility and target-risk penalties terms meaningfully alter the selected interventions. Relative to target selection without either penalty, the main-benchmark setting increased druggability (28.3% to 93.7%) and reduced the target-risk score (25.15 to 2.36), while retaining 74% of the penalty-free effect (**Fig. 4E**). Thus, these penalties shift target selection toward more feasible and lower-risk interventions without eliminating effectiveness.

Together, these ablations support the main design choices of CellReprogrammer: phenotype guidance limits unnecessary transcriptional change, differentiable selection provides an effective target-set optimizer, downstream response modeling is necessary for evaluating interventions, and therapeutic penalties improve feasibility and toxicity while preserving substantial effectiveness.

5 Conclusion

This work introduced CellReprogrammer, a phenotype-guided framework for combinatorial therapeutic target discovery. Rather than matching a complete control transcriptome, CellReprogrammer optimizes target sets for disease-phenotype improvement while jointly considering effectiveness, feasibility, and toxicity. Across five disorders and three cell types, it achieved stronger predicted reprogramming effects than the evaluated baselines, identified compact and highly druggable target sets with lower predicted toxicity, and recovered disease- and cell-type-specific biological patterns. These results support phenotype-guided combination selection as a practical formulation for computational therapeutic target discovery. CellReprogrammer currently generates computational hypotheses whose accuracy depends on the disease classifier, perturbation surrogate, and available perturbation data. Experimental validation is therefore required to establish the therapeutic relevance of these predictions. As perturbation resources and disease atlases expand, the framework can incorporate broader cell types, disease contexts, and therapeutic constraints. CellReprogrammer provides a path from gene prioritization toward systematic, context-specific intervention design.

Reproducibility Statement

Code and evaluation commands are available at `https://anonymous.4open.science/r/CellReprogrammer-B0F5`. Data access, implementation details and figure protocols are provided in the appendices and repository documentation.

The Use of Large Language Models

LLM-based assistants were used to support data visualization, code debugging, and manuscript polishing. The authors verified all analyses and claims and are responsible for the final content of the manuscript.

References

Abhinav K Adduri, Dhruv Gautam, Beatrice Bevilacqua, Mohsen Naghipourfar, Alishba Imran, Rohan Shah, Noam Teyssier, Rishi Verma, Christopher Carpenter, Basak Eraslan, et al. Predicting cellular responses to perturbation across diverse contexts with state. *Cell*, 2026.

Broad DepMap. DepMap 24Q4 Public. Figshare+ dataset, 2024. URL `https://doi.org/10.25452/figshare.plus.27993248.v1`. DepMap portal: `https://depmap.org/portal/`.

Charlotte Bunne, Stefan G Stark, Gabriele Gut, Jacobo Sarabia Del Castillo, Mitch Levesque, Kjong-Van Lehmann, Lucas Pelkmans, Andreas Krause, and Gunnar Rätsch. Learning single-cell perturbation responses using neural optimal transport. *Nature methods*, 20(11):1759–1768, 2023.

Christopher L Cardona, Lai Wei, Joonwon Kim, Ellen Angeles, Gunjandeep Singh, et al. High throughput identification of genetic regulators of microglial inflammatory processes in alzheimer's disease. *Journal of Neuroinflammation*, 22:294, 2025.

Haotian Cui, Chloe Wang, Hassaan Maan, Kuan Pang, Fengning Luo, Nan Duan, and Bo Wang. scgpt: toward building a foundation model for single-cell multi-omics using generative ai. *Nature methods*, 21(8):1470–1480, 2024.

Nikolaos P Daskalakis, Artemis Iatrou, Chris Chatzinakos, Aarti Jajoo, Clara Snijders, Dennis Wylie, Christopher P DiPietro, Ioulia Tsatsani, Chia-Yen Chen, Cameron D Pernia, et al. Systems biology dissection of PTSD and MDD across brain regions, cell types, and blood. *Science*, 384 (6698):eadh3707, 2024. doi: 10.1126/science.adh3707.

Nina M Dräger, Sydney M Sattler, Cindy Tzu-Ling Huang, Olivia M Teter, Kun Leng, Sayed Hadi Hashemi, Jason Hong, Giovanni Aviles, Claire D Clelland, Lihong Zhan, et al. A crispri/a platform in human ipsc-derived microglia uncovers regulators of disease states. *Nature neuroscience*, 25(9):1149–1162, 2022.

Prashant S Emani, Jason J Liu, Declan Clarke, Matthew Jensen, Jonathan Warrell, Chirag Gupta, Ran Meng, Che Yu Lee, Siwei Xu, Cagatay Dursun, et al. Single-cell genomics and regulatory networks for 388 human brains. *Science*, 384(6698):eadi5199, 2024.

Michael J Gandal, Pan Zhang, Evi Hadjimichael, et al. Transcriptome-wide isoform-level dysregulation in ASD, schizophrenia, and bipolar disorder. *Science*, 362(6420):eaat8127, 2018. doi: 10.1126/science.aat8127.

Guadalupe Gonzalez, Xiang Lin, Isuru Herath, Kirill Veselkov, Michael Bronstein, and Marinka Zitnik. Combinatorial prediction of therapeutic perturbations using causally inspired neural networks. *Nature Biomedical Engineering*, pp. 1–18, 2025.

Emre Guney, Jörg Menche, Marc Vidal, and Albert-László Barábasi. Network-based in silico drug efficacy screening. *Nature communications*, 7(1):10331, 2016.

Isabelle Guyon, Jason Weston, Stephen Barnhill, and Vladimir Vapnik. Gene selection for cancer classification using support vector machines. *Machine learning*, 46(1):389–422, 2002.

Minsheng Hao, Jing Gong, Xin Zeng, Chiming Liu, Yucheng Guo, Xingyi Cheng, Taifeng Wang, Jianzhu Ma, Xuegong Zhang, and Le Song. Large-scale foundation model on single-cell transcriptomics. *Nature methods*, 21(8):1481–1491, 2024.

Graham Heimberg, Tony Kuo, Daryle J DePianto, Omar Salem, Tobias Heigl, Nathaniel Diamant, Gabriele Scalia, Tommaso Biancalani, Shannon J Turley, Jason R Rock, et al. A cell atlas foundation model for scalable search of similar human cells. *Nature*, 638(8052):1085–1094, 2025.

Ahyeon Hwang, Mario Skarica, Siwei Xu, Jensine Coudriet, Che Yu Lee, Lin Lin, Rosemarie Terwilliger, Alexa-Nicole Sliby, Jiawei Wang, Tuan Nguyen, et al. Single-cell transcriptomic and chromatin dynamics of the human brain in ptsd. *Nature*, 643(8072):744–754, 2025.

Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. In *International Conference on Learning Representations*, 2017. URL `https://openreview.net/forum?id=rkE3y85ee`.

Xin Jin, Sean K Simmons, Amy Guo, Ashwin S Shetty, Michelle Ko, Lan Nguyen, Vahbiz Jokhi, Elise Robinson, Paul Oyler, Nathan Curry, et al. In vivo perturb-seq reveals neuronal and glial abnormalities associated with autism risk genes. *Science*, 370(6520):eaaz6063, 2020.

Yasmin Jolasun, Kailu Song, Yumin Zheng, Jingtao Wang, Gregory J Fonseca, David H Eidelman, and Jun Ding. Sidish integrates single-cell and bulk transcriptomics to identify high-risk cells and guide precision therapeutics through in silico perturbation. *Nature Communications*, 16(1): 11271, 2025.

Konrad J Karczewski, Laurent C Francioli, Grace Tiao, Beryl B Cummings, Jessica Alföldi, Qingbo Wang, Ryan L Collins, Kristen M Laricchia, Andrea Ganna, Daniel P Birnbaum, et al. The mutational constraint spectrum quantified from variation in 141,456 humans. *Nature*, 581(7809): 434–443, 2020.

Justin Lamb, Emily D Crawford, David Peck, Joshua W Modell, Irene C Blat, Matthew J Wrobel, Jim Lerner, Jean-Philippe Brunet, Aravind Subramanian, Kenneth N Ross, et al. The connectivity map: using gene-expression signatures to connect small molecules, genes, and disease. *science*, 313(5795):1929–1935, 2006.

Kun Leng, Indigo VL Rose, Hyosung Kim, Wenlong Xia, Wilber Romero-Fernandez, Brendan Rooney, Mark Koontz, Emmy Li, Yan Ao, Shinong Wang, et al. Crispri screens in human ipsc-derived astrocytes elucidate regulators of distinct inflammatory reactive states. *Nature neuroscience*, 25(11):1528–1542, 2022.

Junhao Liu, Pengpeng Zhang, Martin Renqiang Min, and Jing Zhang. Identifying combinatorial regulatory genes for cell fate decision via reparameterizable subset explanations. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, pp. 1823–1832, 2025.

Mohammad Lotfollahi, F Alexander Wolf, and Fabian J Theis. scgen predicts single-cell perturbation responses. *Nature methods*, 16(8):715–721, 2019.

Mohammad Lotfollahi, Anna Klimovskaia Susmelj, Carlo De Donno, Leon Hetzel, Yuge Ji, Ignacio L Ibarra, Sanjay R Srivatsan, Mohsen Naghipourfar, Riza M Daza, Beth Martin, et al. Predicting cellular responses to complex perturbations in high-throughput screens. *Molecular systems biology*, 19(6):MSB202211517, 2023.

Michael I Love, Wolfgang Huber, and Simon Anders. Moderated estimation of fold change and dispersion for rna-seq data with deseq2. *Genome biology*, 15(12):550, 2014.

Liang Lu, Liu-kun Yang, Jiao Yue, Xin-shang Wang, Jing-yu Qi, Fan Yang, Ban Feng, and Shuibing Liu. Scutellarin alleviates depression-like behaviors induced by lps in mice partially through inhibition of astrocyte-mediated neuroinflammation. *Neuroscience Letters*, 765:136284, 2021.

Chris J Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous relaxation of discrete random variables. *arXiv preprint arXiv:1611.00712*, 2016.

Janet Piñero, Juan Manuel Ramírez-Anguita, Josep Saüch-Pitarch, Francesco Ronzano, Emilio Centeno, Ferran Sanz, and Laura I Furlong. The disgenet knowledge platform for disease genomics: 2019 update. *Nucleic acids research*, 48(D1):D845–D855, 2020.

Yusuf Roohani, Kexin Huang, and Jure Leskovec. Predicting transcriptional outcomes of novel multigene perturbations with gears. *Nature biotechnology*, 42(6):927–935, 2024.

Yanay Rosen, Yusuf Roohani, Ayush Agrawal, Leon Samotorčan, Tabula Sapiens Consortium, Stephen R Quake, and Jure Leskovec. Universal cell embedding provides a foundation model for cell biology. *Nature*, pp. 1–9, 2026.

Aravind Subramanian, Rajiv Narayan, Steven M Corsello, David D Peck, Ted E Natoli, Xiaodong Lu, Joshua Gould, John F Davis, Andrew A Tubelli, Jacob K Asiedu, et al. A next generation connectivity map: L1000 platform and the first 1,000,000 profiles. *Cell*, 171(6):1437–1452, 2017.

George Sun and Yi-Hui Zhou. Computational framework for therapeutic target discovery via perturbation simulation: application to cystic fibrosis airway disease. *Briefings in Bioinformatics*, 27 (2):bbag129, 2026.

Christina V Theodoris, Ling Xiao, Anant Chopra, Mark D Chaffin, Zeina R Al Sayed, Matthew C Hill, Helene Mantineo, Elizabeth M Brydon, Zexian Zeng, X Shirley Liu, et al. Transfer learning enables predictions in network biology. *Nature*, 618(7965):616–624, 2023.

Ruilin Tian, Anthony Abarientos, Jason Hong, Sayed Hadi Hashemi, Rui Yan, Nina Dräger, Kun Leng, Mike A Nalls, Andrew B Singleton, Ke Xu, et al. Genome-wide crispri/a screens in human neurons link lysosomal failure to ferroptosis. *Nature neuroscience*, 24(7):1020–1034, 2021.

Dmitry Velmeshev, Lucas Schirmer, Diane Jung, Maximilian Haeussler, Yonatan Perez, Simone Mayer, Aparna Bhaduri, Nitasha Goyal, David H Rowitch, and Arnold R Kriegstein. Single-cell genomics identifies cell type–specific molecular changes in autism. *Science*, 364(6441):685–689, 2019.

Hongzhi Wen, Wenzhuo Tang, Xinnan Dai, Jiayuan Ding, Wei Jin, Yuying Xie, and Jiliang Tang. Cellplm: Pre-training of cell language model beyond single cells. In *International Conference on Learning Representations*, volume 2024, pp. 5649–5673, 2024.

Daniel R Wong, Mary Piper, Jiao Qiao, Max Russo, Pierre Jean, Djork-Arné Clevert, Jason D Arroyo, and Evanthia Pashos. Phenotypic reversion and target prioritization for cellular inflammation via representation learning with foundation models. *bioRxiv*, 2026.

You Wu, Pengfei Xu, Liyuan Wang, Shuai Liu, Yingnan Hou, Hui Lu, Peng Hu, Xiaofei Li, and Xiang Yu. scgo: interpretable deep neural network for cell status annotation and disease diagnosis. *Briefings in Bioinformatics*, 26(1):bbaf018, 2025.

Sang Michael Xie and Stefano Ermon. Reparameterizable subset sampling via continuous relaxations. *arXiv preprint arXiv:1901.10517*, 2019.

Yujie Zhang, Madhavan S Venkatesh, and Christina V Theodoris. Discovery of candidate therapeutic targets with geneformer. *Nature Protocols*, pp. 1–24, 2026.

Xinhe Zheng, Boli Wu, Yuejia Liu, Sean K Simmons, Kwanho Kim, Grace S Clarke, Abdullah Ashiq, Joshua Park, Jiwen Li, Zhilin Wang, et al. Massively parallel in vivo perturb-seq reveals cell-type-specific transcriptional networks in cortical development. *Cell*, 187(13):3236–3248, 2024.

Appendix contents

A Method Details 13

B Data Details 15

C Baseline Details 15

D Supplementary Evaluation 16

A METHOD DETAILS

A.1 DISEASE CLASSIFIER

Random-perturbation augmentation. A classifier may respond to perturbation-induced distribution shifts that are unrelated to disease reversal. To reduce this sensitivity, we augment classifier training with randomly perturbed pseudobulks. For each donor, we sample R random K-gene perturbation flags $\{\boldsymbol{\xi}^{(r)}\}_{r=1}^{R}$, where $\boldsymbol{\xi}^{(r)} \in \{0,1,2\}^m$, construct the corresponding pseudobulks $\tilde{\mathbf{p}}_d(\boldsymbol{\xi}^{(r)})$, and assign them the donor's original label y_d. The classifier minimizes the augmented objective

$$\mathcal{L}_{\text{cls}}(\mathbf{w}, b) = \sum_d \Big[\mathrm{BCE}\big(\sigma(\mathbf{w}^\top \mathbf{p}_d + b), y_d\big) + \sum_{r=1}^{R} \mathrm{BCE}\big(\sigma(\mathbf{w}^\top \tilde{\mathbf{p}}_d(\boldsymbol{\xi}^{(r)}) + b), y_d\big) \Big] + \frac{\lambda}{2}\|\mathbf{w}\|_2^2. \tag{7}$$

This augmentation reduces sensitivity to generic perturbation-induced shifts.

A.2 PERTURBATION SURROGATE

For perturbation i, let $\boldsymbol{\delta}_i \in \mathbb{R}^m$ denote the observed gene-expression change relative to the corresponding control, and let $\widehat{\boldsymbol{\delta}}_i$ and $\widehat{\boldsymbol{\delta}}_i^{\,0}$ denote the surrogate-predicted changes with and without the perturbation flag, respectively. We train the perturbation surrogate with

$$\mathcal{L}_{\text{surr}}(\phi) = \frac{1}{|\mathcal{B}|} \sum_{i \in \mathcal{B}} \Big[\underbrace{\overline{\boldsymbol{\omega}_i \odot (\widehat{\boldsymbol{\delta}}_i - \boldsymbol{\delta}_i)^2}}_{\text{DE-weighted reconstruction}} + \lambda_0 \underbrace{\overline{(\widehat{\boldsymbol{\delta}}_i^{\,0})^2}}_{\text{no-op}} + \lambda_1 \underbrace{\overline{\boldsymbol{\omega}_i \odot (\widehat{\boldsymbol{\delta}}_i - \widehat{\boldsymbol{\delta}}_i^{\,0} - \boldsymbol{\delta}_i)^2}}_{\text{perturbation effect}} + \lambda_2 \underbrace{\ell_{\text{dir}}(\phi; i)}_{\text{direction}} \Big], \tag{8}$$

Here $\overline{(\cdot)}$ averages over genes, and $\boldsymbol{\omega}_i = \mathbf{1} + \gamma|\boldsymbol{\delta}_i|$ weights larger observed responses. The four terms fit the observed change, regularize the flag-off output toward zero, match the flag-induced response to the observed change, and penalize incorrect directions for genes with $|\delta_{ij}| > 0.5$. We use $\gamma = 5$, $\lambda_0 = 0.1$, $\lambda_1 = 1$, and $\lambda_2 = 0.075$. Predictions for candidates outside the assayed target pool $\mathcal{G}_{\text{scr}}$ extrapolate beyond measured perturbations; validation on held-out assayed targets does not establish their accuracy or that of their combinations.

A.3 COMBINATION OPTIMIZER

Differentiable top-K sampler. We parameterize the selection of K genes from the candidate set $\mathcal{G}$ using a learnable embedding matrix $\mathbf{E}_\theta \in \mathbb{R}^{C \times h}$, where $C = |\mathcal{G}|$, followed by an MLP that produces per-gene logits $\boldsymbol{\ell} = \mathrm{MLP}_\theta(\mathbf{E}_\theta) \in \mathbb{R}^C$. We use a differentiable relaxation of sampling without replacement (Xie & Ermon, 2019):

$$r_i = \ell_i + G_i, \quad G_i = -\log(-\log u_i), \;\; u_i \sim \mathcal{U}(0,1), \tag{9}$$

$$\boldsymbol{\alpha}^1 = \mathbf{r}, \qquad \mathbf{q}^j = \mathrm{softmax}\big(\boldsymbol{\alpha}^j/\tau\big), \tag{10}$$

$$\boldsymbol{\alpha}^{j+1} = \boldsymbol{\alpha}^j + \log\big(1 - \mathbf{q}^j\big), \quad j = 1, \ldots, K, \tag{11}$$

$$\mathbf{s}_{\text{soft}} = \min\Big(\textstyle\sum_{j=1}^{K} \mathbf{q}^j, \; \mathbf{1} \Big) \in [0,1]^C, \tag{12}$$

The clipped relaxation $\mathbf{s}_{\text{soft}}$ need not sum exactly to K. Let $\mathbf{s}_{\text{hard}} = \mathrm{hardTopK}(\mathbf{r})$ be the binary mask with exactly K selected genes. Standard selection uses $\mathbf{s}_{\text{ST}} = \mathbf{s}_{\text{soft}} + \mathrm{sg}(\mathbf{s}_{\text{hard}} - \mathbf{s}_{\text{soft}})$, where sg stops gradients: the forward mask is hard, but its gradient is soft. Soft forward instead uses $\mathbf{s}_{\text{soft}}$. During selection, mask weights multiply direction-specific predicted response vectors, not the numerical $0/1/2$ flag codes; the resulting response is mapped to the full gene space. Final evaluation uses hard top-K sets from the logits.

Transcriptomic-change cap. To limit the magnitude of predicted perturbation responses, we cap each response relative to the measured disease–control shift $\mathbf{a} = \boldsymbol{\mu}_{\text{ctrl}} - \boldsymbol{\mu}_{\text{dis}}$:

$$\mathrm{cap}_\kappa(\boldsymbol{\delta}) = \boldsymbol{\delta} \cdot \min\big(1, \; \kappa\|\mathbf{a}\|_2 / \|\boldsymbol{\delta}\|_2\big). \tag{13}$$

We define $\mathrm{cap}_\kappa(\mathbf{0}) = \mathbf{0}$. At $\kappa = 1$, the disease–control shift bounds the response norm; the cap never enlarges a response.

Algorithm 1 CellReprogrammer target selection (per disease × cell type).

Require: expression $\mathbf{X}$, donor labels $\mathbf{y}$, fine-tuned surrogate g_ϕ, intervention size K, steps T, seeds $\mathcal{R}$

form donor pseudobulks $\mathbf{p}_d = |\mathcal{C}_d|^{-1} \sum_{i \in \mathcal{C}_d} \mathbf{x}_i$
augment classifier training with random-K perturbation pseudobulks assigned the original label y_d
train disease classifier $(\mathbf{w}, b)$ and determine perturbation directions $\mathbf{e}$
for seed $r \in \mathcal{R}$ **do**
    initialize sampler $\theta = (\mathbf{E}_\theta, \mathrm{MLP}_\theta)$ and temperature τ
    for step $t = 1, \ldots, T$ **do**
        sample $\mathbf{r}$ and $\mathbf{s}_{\mathrm{soft}}$ (Eqs. 9–12)
        $\mathbf{s}_{\mathrm{ST}} \leftarrow \mathbf{s}_{\mathrm{soft}} + \mathrm{sg}(\mathrm{hardTopK}(\mathbf{r}) - \mathbf{s}_{\mathrm{soft}})$
        evaluate $\mathcal{L}_{\mathrm{c}}(\mathbf{s}_{\mathrm{ST}})$ (Eq. 6) with frozen g_ϕ, h_ψ; update θ by Adam and anneal τ
    end for
    $S^{(r)} = \text{top-}K(\boldsymbol{\ell})$
end for
return K genes with the highest selection frequencies; report S_{core} (Eq. 17)

Drug-realizability score. For each candidate gene g, let $\sigma_g \in \{+1, -1\}$ denote the required intervention direction and let u_g and w_g denote the numbers of brain-penetrant LINCS L1000 compound signatures (Subramanian et al., 2017) that significantly increase and decrease its expression, respectively. We define

$$R_g = \max\Big(0,\ \sigma_g \frac{u_g - w_g}{u_g + w_g}\Big), \qquad R_g = 0 \text{ if } u_g + w_g = 0, \tag{14}$$

The corresponding optimization penalty is

$$\mathcal{L}_{\mathrm{drug}} = \frac{1}{K} \sum_{g \in \mathcal{G}} s_g (1 - R_g). \tag{15}$$

Target-toxicity score. We estimate inhibition-associated target risk using DepMap knockout lethality (Broad DepMap, 2024) and gnomAD loss-of-function intolerance (Karczewski et al., 2020). For gene g, let c_g be its Chronos gene-effect value; the DepMap component is $d_g = \min(1, \max(0, -c_g))$, and the gnomAD component is $q_g = \max(0, 1 - \mathrm{LOEUF}_g)$. We use $b_g = (d_g + q_g)/2$ when both are available and the available component otherwise. For optimization, $t_g = b_g$ for inhibited targets and $t_g = 0$ for activated targets, giving

$$\mathcal{L}_{\mathrm{tox}} = \frac{1}{K} \sum_{g \in \mathcal{G}} s_g t_g. \tag{16}$$

Activated targets receive no loss-of-function penalty because these sources characterize inhibition or loss of function; this does not establish their safety under activation.

Multi-seed stability selection. Because the optimization objective is non-convex, selected gene sets may vary across optimization seeds. We therefore run the optimizer over seeds $\mathcal{R}$ while holding the data and predictive models fixed. For the resulting sets $\{S^{(r)}\}_{r \in \mathcal{R}}$, we define

$$\mathrm{freq}(g) = \frac{1}{|\mathcal{R}|} \sum_{r \in \mathcal{R}} \mathbb{1}\Big[g \in S^{(r)}\Big], \qquad S_{\mathrm{core}} = \{g : \mathrm{freq}(g) \geq \gamma\}, \qquad \gamma = 0.8, \tag{17}$$

For final full-cohort summaries, runs whose hard selections do not lower mean predicted disease probability are excluded. We report the K genes with the highest selection frequencies and use S_{core} as a stability diagnostic. Leave-one-donor-out evaluation performs selection independently within each training fold.

Pseudocode and complexity. Algorithm 1 summarizes target selection. Let $C = |\mathcal{G}|$, T be the number of optimization steps, and $R = |\mathcal{R}|$ the number of seeds. Each relaxed top-K sample makes K passes over C logits, costing $O(KC)$ time and $O(KC)$ automatic-differentiation memory. Across seeds, this costs $O(RTKC)$ for the relaxation alone, avoiding exhaustive enumeration of candidate sets. Full runtime also includes logit computation and gradients through the frozen surrogate.

A.4 Implementation and hyperparameters

Disease classifier. Donor pseudobulks average measured expression across sampled nuclei from each donor. The intervention evaluator is L_2-regularized logistic regression ($C = 0.5$), refit after excluding each evaluated disease donor. Four random-perturbation pseudobulks per training donor retain its label. The separate five-fold

Table B.1: **Perturb-seq sources across the five screened cell types.** IN and OLG were evaluated for cell-type selection but excluded from the main intervention benchmark.

Cell type	Source study	Accession	Perturbation
EXN	Tian et al. (2021)	GSE152988	CRISPRi/a
EXN, IN	Zheng et al. (2024)	GSE249416	in vivo loss of function
EXN, IN, MG, AST, OLG	Jin et al. (2020)	GSE157977	in vivo Cas9 knockout
MG	Dräger et al. (2022)	GSE178317	CRISPRi/a
MG	Cardona et al. (2025)	GSE289721	CRISPRi
AST	Leng et al. (2022)	GSE182308	CRISPRi

classification screen is reported in Table B.2. A predefined set of sex-associated marker genes (Appendix B) and genes unmeasured in a cohort are excluded from intervention candidates; the marker mask is also applied to classifier features. Other X/Y-chromosome genes remain eligible.

Perturbation surrogate. scGPT is fine-tuned per cell type with the residual response and Eq. (8) ($\gamma = 5$, $\lambda_0 = 0.1$, $\lambda_1 = 1$, $\lambda_2 = 0.075$). Training uses Adam (lr 10^{-4}), batch size 32, 25 epochs, dropout 0.2, gradient clipping 1.0, at most 120 control inputs per condition, and an $85/15$ target-level split. The supervised target is the measured perturbation-minus-control pseudobulk effect.

Combination optimizer. The learned gene embeddings and MLP parameterize Gumbel-top-K selection. We use Adam (lr 5×10^{-3}), 150 steps, and temperature $2.0 \to 0.1$, with $K = 3$ unless stated. Penalty weights $(\lambda_{\text{drug}}, \lambda_{\text{tox}}) = (0.5, 1)$ are fixed across contexts. Standard optimization uses hard forward and soft backward selection; soft forward changes only the forward mask. Both evaluate hard top-K sets. Directions follow the disease–control pseudobulk difference. Final full-cohort sets use 100 seeds (Appendix A.3); benchmark selections are fold-specific.

B Data Details

Patient cohorts (classification and combination optimization). We use 2,289,891 nuclei from 331 donors across SCZ, BPD, ASD, MDD, and PTSD, counting shared MDD/PTSD controls once. Cohort control/disease donor totals and the seven-cell-type coverage are given in Table 1. Donor labels encode diagnostic status.

Perturb-seq (response modeling). The EXN, MG, and AST screens comprise 422 cell-type–target pairs across 263,209 cells, including controls: EXN 284 targets/78,371 cells; MG 76/84,093; AST 62/100,745. IN and OLG screens support the cell-type selection analysis in Table 1 but are not used for the main intervention benchmark. Table B.1 identifies all five screened cell types and their source studies. Mouse in vivo responses were mapped to human gene orthologs.

Counts and shared controls. Table 1 retains donor–cell-type groups with at least eight nuclei. MDD and PTSD controls are counted once overall but appear in both disease columns.

Candidate pools. The screen-grounded pool $\mathcal{G}_{\text{scr}}$ is the set of Perturb-seq-assayed target genes for each cell type (284/76/62 for EXN/MG/AST), each with a directly measured perturbation response. The expanded pool $\mathcal{G}_{3\text{k}}$ is a 3,000-gene pool per cell type: the union of all screened target genes and the top patient-derived highly variable genes, computed on genes common to the patient cohort, Perturb-seq screens, and surrogate vocabulary after excluding the following predefined sex-associated markers: *XIST*, *TSIX*, *RPS4Y1*, *RPS4Y2*, *DDX3Y*, *UTY*, *USP9Y*, *KDM5D*, *EIF1AY*, *NLGN4Y*, *TXLNGY*, *ZFY*, *TMSB4Y*, *PRKY*, *TTTY14*, *TTTY10*, and *UTY1*. The perturbation surrogate predicts a per-gene perturbation delta for *every* gene in $\mathcal{G}_{3\text{k}}$; only the $\mathcal{G}_{\text{scr}}$ subset carries a measured delta to validate against (the extrapolation caveat of Appendix A.2).

B.1 Disease classifier performance

Table B.2 reports five-fold donor AUROC on cell-type-specific top-3,000 variance genes. This classification screen differs from the augmented leave-one-disease-donor-out classifier used for intervention scoring (Appendix A.4; scoring metrics in Appendix D.1).

C Baseline Details

Baselines select within the same $\mathcal{G}_{3\text{k}}$ universe and are scored through the identical magnitude-capped donor pipeline. The measured-effect linear ranking has measured responses only for $\mathcal{G}_{\text{scr}} \subset \mathcal{G}_{3\text{k}}$ and therefore selects only from that subset; no effect is imputed for unassayed targets. Other baselines may rank the expanded pool. Selected genes receive disease–control-derived intervention directions.

Table B.2: **Disease classification across seven cell types.** Donor AUROC, mean±SD across five folds; cell-type-specific top-3000 variance genes. Bold: intervention-benchmark cell types.

Cell type	ASD	BPD	SCZ	MDD	PTSD	mean
EXN	$0.75_{\pm 0.11}$	$0.76_{\pm 0.08}$	$0.64_{\pm 0.17}$	$0.68_{\pm 0.11}$	$0.60_{\pm 0.21}$	0.69
MG	$0.60_{\pm 0.16}$	$0.58_{\pm 0.18}$	$0.56_{\pm 0.09}$	$0.76_{\pm 0.15}$	$0.59_{\pm 0.22}$	0.62
AST	$0.63_{\pm 0.15}$	$0.68_{\pm 0.20}$	$0.63_{\pm 0.20}$	$0.74_{\pm 0.15}$	$0.58_{\pm 0.18}$	0.65
IN	$0.70_{\pm 0.12}$	$0.71_{\pm 0.11}$	$0.64_{\pm 0.15}$	$0.79_{\pm 0.09}$	$0.66_{\pm 0.18}$	0.70
OLG	$0.68_{\pm 0.10}$	$0.57_{\pm 0.12}$	$0.53_{\pm 0.11}$	$0.68_{\pm 0.19}$	$0.66_{\pm 0.20}$	0.62
OPC	$0.63_{\pm 0.12}$	$0.69_{\pm 0.14}$	$0.66_{\pm 0.18}$	$0.73_{\pm 0.17}$	$0.62_{\pm 0.16}$	0.66
END	$0.58_{\pm 0.16}$	$0.48_{\pm 0.13}$	$0.69_{\pm 0.12}$	$0.61_{\pm 0.21}$	$0.77_{\pm 0.20}$	0.63

Subscripts give fold standard deviations; Table 1 summarizes cell-type means.

C.1 Association-based ranking

Each ranking method assigns a per-gene score and takes the top K. *DE* scores $|\bar{\mathbf{p}}^{\text{dis}} - \bar{\mathbf{p}}^{\text{ctrl}}|$, the absolute donor pseudobulk difference. *clf-w* scores $|w_g|$, the magnitude of the disease classifier weight. *net-prox* scores the negative closest-distance from each candidate to the disease module in the STRING interaction network (Guney et al., 2016), the disease module being the curated DisGeNET genes; it uses topology only, with no surrogate. *risk* ranks the curated DisGeNET disease genes themselves (SFARI for autism). The interaction network uses STRING v12.0 human physical links at combined score ≥ 400. Autism-associated genes come from the SFARI Gene human-gene export downloaded in June 2026.

C.2 Perturbation-based ranking

scGO ranks single-gene perturbations using disease-classifier reversal. Linear ranking uses measured single-gene perturbation effects only for screened genes. Response-based baselines use predictions from GEARS, scGen, CPA, and the single-cell models listed in Table 2. *CMap* scores signature reversal (Lamb et al., 2006), $-\text{corr}\big(\boldsymbol{\delta}_g,\ \bar{\mathbf{p}}^{\text{dis}} - \bar{\mathbf{p}}^{\text{ctrl}}\big)$ over the differentially expressed genes, so a gene scores highly when its perturbation delta points opposite to the disease signature. *GEARS-select* scores $-\mathbf{D}_g^{\text{GEARS}} \cdot \mathbf{w}$, the GEARS-predicted perturbation effect projected onto the classifier, so the most negative projection ranks first. The *foundation-model* rankings (scGPT, scFoundation, Geneformer) are zero-shot in-silico perturbation: the disease and control cell means are embedded, the disease→control axis is formed in embedding space, each candidate is perturbed in its disease–control-derived direction (inhibition to zero or deletion; activation to the control level or overexpression), the perturbed mean is re-embedded, and genes are ranked by the resulting shift toward the control state. These models provide rankings; the combination optimizer does not reuse their pretrained weights.

C.3 Combinatorial selection

PDGrapher. PDGrapher (Gonzalez et al., 2025) predicts intervention targets from pre- and post-intervention expression, an intervention label, and a STRING graph. We adapted it to disease-state transitions using disease-donor pseudobulks as pre-intervention states and surrogate-predicted signed single-gene responses to form proxy post-intervention states; each label identifies the perturbed gene. Training uses up to 24,000 proxy triples per disease–cell-type context, with 25 epochs and one graph layer. At inference, disease and control donor means define the requested transition, and the perturbation-discovery output ranks targets. Because baseline models are trained on full cohorts before held-out classifier scoring, this comparison does not provide fully donor-held-out target selection or validation against measured disease–treatment responses.

Combinatorial search. Simulated annealing, a genetic algorithm, random search, and greedy selection optimize the same set-level objective directly over K-gene sets. The stochastic searches share a fixed candidate-evaluation budget (160 per restart); we report means over eight restarts.

D Supplementary Evaluation

D.1 Evaluation metrics

We provide detailed definitions of the effectiveness, feasibility, toxicity, and perturbation-response fidelity metrics used in the main text.

Effectiveness. For each held-out donor, $\Delta\ell = \text{logit}\, h_\psi(\tilde{\mathbf{p}}_d) - \text{logit}\, h_\psi(\mathbf{p}_d)$ and $\Delta\ell_{\text{adj}} = \Delta\ell - \Delta\ell_{\text{rand}}$. More negative values indicate stronger predicted effects relative to matched random perturbations. Means are computed over donors within each disease–cell-type setting, then equally over settings. Random controls are matched separately to each intervention size and transcriptomic-change cap using 25 antithetic direction pairs.

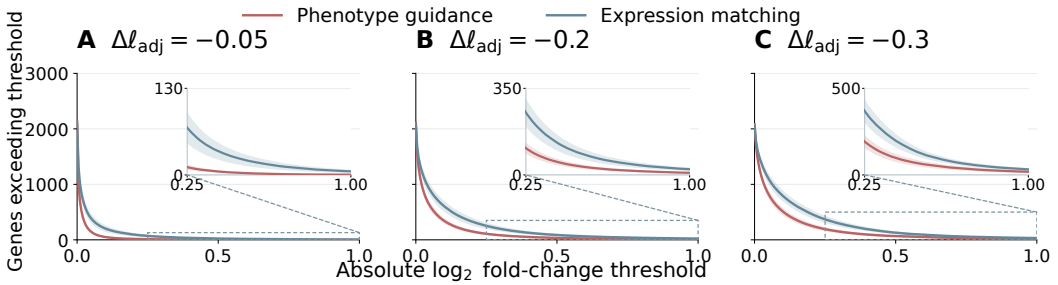

Figure D.1: **Downstream changes at additional matched effects.** (A–C) Gene counts above absolute $\log_2$ fold-change thresholds at $\Delta\ell_{\text{adj}} = -0.05, -0.2, -0.3$; shading, SE across settings. Dashed outlines mark the enlarged regions.

Feasibility. Intervention size is K; druggability is the fraction of selected targets with direction-matched LINCS support ($R_g > 0$; Appendix A.3). The fraction is averaged over donors, then equally over settings. It measures expression support, not direct drug binding.

Target toxicity. We average b_g (Appendix A.3) over selected inhibited targets with an available score and report it on a 0–100 scale. Targets missing both sources are excluded; if none are scored, the reported value is zero. This risk score is not a probability of cell death.

Downstream toxicity. We report the fraction of predicted post-perturbation values outside the gene-wise 1st–99th percentile interval of training-donor pseudobulks, including pre-existing outliers.

Figure 3E independently maps each outward-better utility u to radius $0.25 + 0.75u$. The ranges are effect $+0.2$ to -0.6, druggability 0–100%, target-risk score 100–0, and downstream toxicity 7–4%; size utility is $(20 - K)/19$, or zero if the effectiveness threshold is not reached. Polygon area is not a combined score.

Perturbation-response fidelity. We evaluate the surrogate on target genes held out from fine-tuning. For a perturbation targeting gene g, the predicted and observed expression changes relative to control are

$$\widehat{\boldsymbol{\Delta}}_g = \overline{g_\phi(\mathbf{x}, \mathbf{f}_g) - g_\phi(\mathbf{x}, \mathbf{0})}, \qquad \boldsymbol{\Delta}_g = \bar{\mathbf{p}}^{\,g} - \bar{\mathbf{p}}^{\text{ctrl}}, \tag{18}$$

where $\mathbf{f}_g \in \{0, 1, 2\}^m$ specifies the perturbation, the first average is over control cells $\mathbf{x}$, and $\bar{\mathbf{p}}^{\,g}$ and $\bar{\mathbf{p}}^{\text{ctrl}}$ are the mean observed profiles of perturbed and control cells. The predicted change is the difference between surrogate outputs with and without the perturbation flag, as in the combination optimizer (Eq. (5)). We measure fidelity over the 20 genes with the largest observed changes:

$$\rho_\Delta = \text{corr}\big(\widehat{\boldsymbol{\Delta}}_g[\mathcal{T}_g],\ \boldsymbol{\Delta}_g[\mathcal{T}_g]\big), \tag{19}$$

where $\mathcal{T}_g$ denotes these response genes. Direction accuracy (Acc.) is the mean fraction of $\mathcal{T}_g$ genes with matching predicted and observed signs across held-out targets. For this metric, the predicted change is $\overline{g_\phi(\mathbf{x}, \mathbf{f}_g) - \mathbf{x}}$; ρ_Δ instead uses the flag-on-minus-flag-off contrast above. Control profiles are pooled across sources, so these estimates do not establish source-matched response-amplitude fidelity.

D.2 Matched-effect analysis

Phenotype guidance at matched effect. For Figure 4A, we reselect $K = 3$ targets in each held-out fold using either the disease-classifier objective or squared distance to the training-control mean. Both arms share the surrogate, candidate pool, seed, 150 Adam steps, and target penalties; each objective is normalized by its initial training-donor mean. After selection at $\kappa = 1$, we solve for a cap in $[0, 4]$ yielding $\Delta\ell_{\text{adj}} = -0.1$ with matched random controls. Both arms reach this level in $314/510$ folds across all 15 settings.

For each donor, we calculate $\log_2[(\bar{x}_{\text{pert}} + 0.1)/(\bar{x}_{\text{orig}} + 0.1)]$ from mean CP10k expression. The common downstream background excludes either arm's targets and genes with training-cell mean below 0.1 CP10k. Gene counts above each absolute fold-change threshold are averaged within settings and then equally across settings, with setting-level SE.

Additional matched-effect levels. Figure D.1 extends the $\Delta\ell_{\text{adj}} = -0.1$ comparison in Figure 4A to -0.05, -0.2, and -0.3. Both arms reach these levels in 356, 244, and 190 paired folds, covering 15, 14, and 13 settings, respectively. Coverage differs between levels; the expression filter, downstream background, and pseudocount are shared.

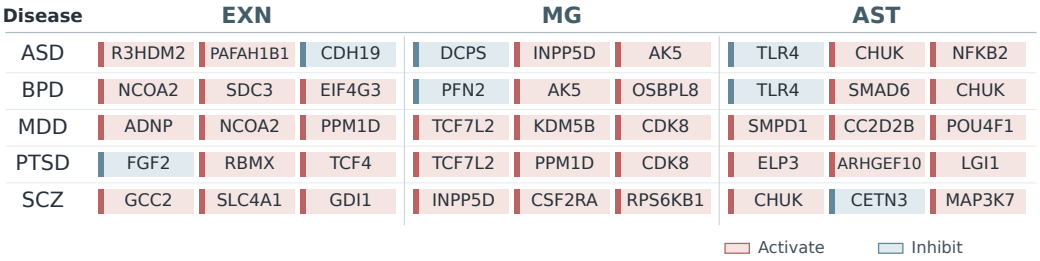

Figure D.2: **Selected targets across 15 disease–cell-type contexts.** Three most frequent targets per context at $K = 3$; red, activate; blue, inhibit.

D.3 Context specificity

Figure 2E uses full-cohort gene sets and classifiers. For a set transferred from context i to j, let $R_{i\to j} = -\Delta\ell_{\mathrm{adj}}(i \to j)/G_j$, where G_j is the recipient disease–control logit gap. Transfer shortfall is $100(R_{j\to j} - R_{i\to j})$ percentage points; positive values favour recipient-specific selection.

Downstream responses. Figure 2F,G uses the three most frequently selected targets and their recorded directions in each full-cohort context, with $\kappa = 1$. Donor-level fold changes use mean CP10k expression and a 0.1 pseudocount; targets and genes with unperturbed mean below 0.1 CP10k are excluded. For panel F, we rank the absolute values of mean signed gene responses and compute Spearman correlations over genes eligible in all settings. Diagonal entries compare different diseases within a cell type, never a context with itself.

For panel G, each disease cohort supplies its top 100 control-derived markers per cell type. Ten markers per cell type are chosen by frequency across the five cohorts, then mean specificity. Colors show within-set rank percentiles of mean absolute donor responses; gray denotes unavailable or excluded genes. Marker enrichment among each donor's top 100 absolute responses uses log odds ratios with a 0.5 correction, averaged within and then equally across diseases. These descriptive responses include the cell-type-specific surrogate and do not isolate the optimizer's contribution.

D.4 Biological support

Target functions and directions. Figure D.2 shows the three most frequent targets and their proposed directions in each context. Three of five astrocyte sets include *TLR4* or *CHUK*/IKKα; scutellarin suppresses astrocytic TLR4/NF-κB signaling in a mouse model (Lu et al., 2021). Microglial *CDK8* and *PFN2* have reported roles in differentiation and phagocytosis (Dräger et al., 2022). The proposed directions and combinations remain experimental hypotheses.

Pathway enrichment and network connectivity. Astrocyte sets are enriched for NF-κB/TLR signaling (OR $= 47.7$, Fisher $p = 7 \times 10^{-6}$) and inflammatory genes (OR $= 5.0$, $p = 0.008$). With the druggability penalty, transcription factors, kinases, and epigenetic regulators are also enriched (OR $= 10.9$, 4.7, and 5.7). Four of 15 sets show above-chance STRING connectivity.

GO enrichment of targets and downstream responses. We tested Figure D.2 targets against the 3,000-gene candidate pool and each setting's 100 largest absolute non-target downstream responses against eligible genes. One-sided hypergeometric tests covered GO terms with 10–500 background genes; Benjamini–Hochberg correction covered setting–term tests separately for targets and responses. Figure 2H shows the union of each cell type's top three target terms by significant-disease count (ties: smaller q).

Shared pathway profiles. Figure 2J compares downstream GO fold-enrichment vectors across terms shared by all five diseases (715 in EXN, 799 in MG, 870 in AST), using Spearman correlations for the ten disease pairs. Overlapping terms and absolute responses preclude conclusions about matching response directions.