ACTIVE LEARNING TO DISCOVER PAIRWISE GENETIC INTERACTIONS VIA REPRESENTATION LEARNING

Moksh Jain[†1], **Alisandra Kaye Denton**[2], **Shawn T. Whitfield**[2], **Aniket Didolkar**[†1], **Berton Earnshaw**[2], **Jason Hartford**[2]
[1]Mila - Quebec AI Institute, Univérsité de Montréal [2]Valence Labs
[†]Work done during an internship at Valence Labs
moksh.jain@mila.quebec,jason.hartford@valencelabs.com

ABSTRACT

Embeddings of microscopy images from single gene knockouts can be used to infer biological interactions, but are limited to interactions that are revealed by single perturbations. If we want to detect effects that are only present in pairwise knockouts we need to (1) address the quadratic scaling of experimental costs, and (2) develop a method for detecting pairwise interactions. We present a set of theoretical independence assumptions under which the sum of embedding of single perturbations predicts the pairwise embeddings. Prediction failures then correspond to violations of these assumptions, and can be used to detect biological interactions. We used this prediction error as a reward in an active search algorithm and found that we can efficiently identify these instances of non-independence, and many of the selected pairs correspond to known biological interactions.

1 INTRODUCTION

Biologists use genetic perturbations to attempt to understand the mechanisms that govern the behaviour of cells (Jinek et al., 2012; Doudna and Charpentier, 2014). We can get a detailed view of the effects of these perturbations with cell painting assays (Bra, 2016; Fay et al., 2023; Chandrasekaran et al., 2023)—microscopy images of cells that have been stained with fluorescent dyes to expose different cellular components and organelles—that reveal the resulting change in the morphology of the cell. When embedded into a sufficiently expressive feature space, these images have been shown to be an effective tool for recovering known biological relationships (Moshkov et al., 2022; Sypetkowski et al., 2023; Kraus et al., 2023).

Given the cost of running perturbation experiments at scale, most prior work perturbs cells with single gene knockouts and infers potential relationships between genes by comparing the similarity between their resulting embeddings. For example, Kraus et al. compared the average embeddings of all single gene perturbations (after correcting for batch effects) and found that the tails of the distribution of cosine similarities between pairs of perturbations recovered a significant fraction of known relationships from biological databases. This works because proteins in the same biological pathway that have the same direction of effect (i.e. either up- or down-regulation), will tend to have a similar effects on the morphology of the cell, and have similar embeddings.

But there is a biological limit to what can be recovered like this: some relationships between genes only become apparent in double knockout experiments. For example, genes with *synthetic lethality* interactions (Nijman, 2011) result in the cell dying only if both genes are knocked out, but the cell remains healthy if either one of the two genes are knocked out. These kinds of interactions will not be detected from a single gene screen because, by definition, a single gene screen only ever knocks out one of the two genes at a time, so the associated embeddings will look like healthy cells.

We want methods that can find any pairwise interaction (i.e. not just those that result in cell death), without having to run all pairs of experiments. To achieve this, we need to address two challenges. First, how do we detect a biological interaction from the high dimensional output of a pairwise experiment[1]? And second, how do we avoid having to run all pairs of gene knockouts to recover all

[1]We will focus on assays that produce images, but our methods only depend on having an embedding space.

of these pairwise biological relationships. We address the first challenge by comparing the centered pairwise embedding from the final hidden layer of a classifier with the sum of individual embedding vectors. We present a set of theoretical results that show when this approach can be used to test the hypothesis that pairs of perturbations act independently on the cell. Empirically, we find that the norm of the difference between the pairwise and individual embeddings can be used to recover certain classes of known pairwise relationships such as pairs of apoptosis (programmed cell death) inhibitors.

To address the second challenge, we efficiently search the space of pairwise knockout experiments by treating this norm as a reward, thereby reducing the problem of finding interacting genes into an active matrix completion problem. This approach allows us to avoid the challenge of having to directly search the embedding space for interactions, and avoids the need to characterize uncertainty over the embedding function. Instead we directly model a posterior over the reward matrix and at every round select actions to trade off exploration and exploitation using Information Directed Sampling (IDS; Russo and Van Roy, 2016; Xu et al., 2022).

We tested this approach by collecting a benchmark dataset of all pairs of gene knockouts for 50 genes in HUVEC cells. We found that IDS discovered pairs of genes that result in large norm interactions significantly faster than random search, giving a 10% increase in the number of biological interactions that we were able to discover after 50 rounds of experimentation. The relationships that we detected were also complementary to that which would have been discovered using just single perturbations, and as a result, the two approaches can be combined to get a more detailed estimate of the relationships between genes from perturbation experiments.

In summary, we present an active learning approach for efficiently finding synthetic lethality-style interactions between genes that are only observable in pairwise assays. The approach uses the norm of the difference between the sum of single gene embeddings and a pairwise embedding as measure of the interaction between the two genes. We motivate this choice of score by deriving the theoretical conditions that show when a score of zero corresponds to no interaction.

2 Independence and Additivity of Representations

We assume that our images can be described by a set of (unknown) morphological features, $\{Z_i\}_{i=1}^n$ that describe the presence or absence of observable morphological phenotypes. For example, if a cell is experiencing apoptosis (programmed cell death), observable characteristics include chromatin compaction and nuclear fragmentation, the appearance of apoptotic bodies, and blebbing (protruding cell membrane). If a cell is recycling cellular material in response to resource limitations (macroautophagy a.k.a. autophagy), we observe the appearance of characteristic structures: phagophores and autophagosomes. These features are not assumed to be known, and are not mutually exclusive, but we do assume that they are conditionally independent, given the set of genes that have been knocked out[2]. Let δ denote an indicator vector that records which genes have been knocked out in a given experiment, with $\delta_i = 1$ if gene i is knocked out and 0 otherwise. Each of the features depends on whether or not a subset of genes have been knocked out for any given experiment.[3] For simplicity, we assume that the experiments are balanced so that $P(\delta_i) = P(\delta_j)$ for all i, j, but our approach could be generalized to unbalanced experiments by re-weighting. Taken together we have the following generative model,

$$Z_i = f_i(\delta, \epsilon_i), \quad P(\mathbf{Z}) = P(Z_1, \ldots, Z_n) = \prod_{i=1}^{n} P(Z_i) \quad X = g(Z_1, \ldots, Z_n, \epsilon_y) \tag{1}$$

where g is the observation function (in our case, a microscopy image), and X is the pixel level representation of the cell. We assume that g is a diffeomorphism: i.e. that it is bijective and both g and its inverse are continuously differentiable. The noise variables, ϵ_i, ϵ_y are assumed to be mutually independent. From this generative process, we observe the image, X, and δ, the experimental description of which genes were knocked out.

[2]This independence assumption should ideally be weakened to conditional independence given the underlying state of the cell. We leave this generalization to future work.

[3]These features were inspired by the *concepts* of (Wang et al., 2023) and play an analogous role although we use a very different embedding function and proof strategy.

The main hypothesis that our method aims to test is that, for a pair of genes, i and j, that influence Z_l we can partition the outcomes associated with each visual feature, Z_l, into the following *mutually exclusive* sets of events, each of which is a function of a single knockout. In particular,

Assumption 2.1. Assume that the realizations of each morphological feature, Z_l, can be partitioned into mutually exclusive events such that,

$$P(Z_l|\delta_i, \delta_j) = P(Z_l^i \cup Z_l^j \cup Z_l^n|\delta_i, \delta_j) = P(Z_l^i|\delta_i, \delta_j) + P(Z_l^j|\delta_i, \delta_j) + P(Z_l^n|\delta_i, \delta_j).$$

To understand this assumption by way of example, let Z_l denote the event that cell l died. Then, we have partitioned the set of states where the cell dies into those where knocking out gene i was responsible for cell l's death, Z_l^i, those where knocking out gene j was responsible for the cell's death, Z_l^j, and the remaining cell deaths that were caused by neither gene i nor gene j, Z_l^n. Additionally, assume the following conditional independencies hold:

Assumption 2.2. For any feature Z_l, we assume: (2.2.1) **No other effects.** For all i, j, $P(Z_l^i|\delta_i, \delta_j) = P(Z_l^i|\delta_i)$, and (2.2.2) **No off target effects.** For all i, j, $P(Z_l^n|\delta_i, \delta_j) = P(Z_l^n)$.

In our example, Assumption 2.2.1 is essentially a formal statement of *non*-synthetic lethality: if the probability that a cell died when knocking out gene i was influenced by whether or not gene j was knocked out, then we would have an example of a synthetic lethal relationship. Assumption 2.2.2 rules out interactions between i and j among the other causes of the morphological feature.

Our goal is to understand when embedding vectors can be used to detect interactions. Here we consider embedding vectors, $h(\cdot)$, that are constructed from the final hidden layer of a classifier. We assume that this classifier is trained optimally such that, $P(\delta_i|X) = \sigma(w_i^\top h(X))$, where $\sigma := \frac{\exp(x_i)}{\sum_j \exp(x_j)}$ is the softmax function, and that there are sufficiently diverse labels to ensure that this representation is identified up to a linear transformation; see Roeder et al. (2021) for details. Given a trained classifier, let $h_i := E[h(X)|\delta_i = 1] - E[h(X)|\delta = 0]$ be the average of the embeddings associated with a particular knockout centered around the control wells.

Theorem 2.3. *Assume the data generating process given in Equation 1, and that Assumption 2.1 and 2.2 hold. Then,*

$$h_i + h_j = h_{i,j},$$

i.e. the sum of the individual perturbation embeddings equals the pairwise embedding.

The proof is in Appendix B. This theorem shows that under the above assumptions, non-interaction between genes implies we can add single gene embeddings and predict the pairwise embedding. This is both useful for prediction—if we know these assumptions hold, we can predict the embedding that would be derived from an experiment without running it—and more importantly for our purposes, it implies that prediction failures will result from a violation of one of our assumptions. While prediction errors do not tell us which of the assumptions have failed, we know that independence assumed in Assumption 2.2 will be violated for synthetic-lethal-style interactions, and hence we can use this score to find candidates. Note that while we consider synthetic-lethality style interactions, the overall approach is applicable to other types of interactions.

3 Discovering Interacting Gene Pairs Efficiently

Through Theorem 2.3, we establish that we can use prediction error as an indicator of potential gene pair interactions. In practice, the goal is to identify gene-pair knockouts that induce large interactions in order to discover as many gene-gene relationships as possible within a fixed budget of experiments, as performing double gene knockouts on all pairs is unfeasible. We formalize this problem of selecting gene-pairs within the framework of *Adaptive Sampling for Discovery* (ASD; Xu et al., 2022). ASD adopts *Information Directed Sampling* (IDS; Russo and Van Roy, 2016) in a particular instantiation of the sleeping experts framework (Kanade et al., 2009), where each actions can be selected exactly once.

Concretely, let Δ denote the space of possible experiment designs, which in this case is the set of tuples $G \times G$ where G is the set of genes, each associated with some score $R : \Delta \to \mathbb{R}$. Following the previous sections, the reward for pair $(i, j) \in G \times G$ is defined as follows

$$R(i,j) \coloneqq (\mathbb{E}[h_{i,j} - h_i - h_j])^2 \quad \text{where } h_i := E[h(X)|\delta_i = 1] - E[h(X)|\delta = 0] \tag{2}$$

where the expectation is taken over a population of cells. The matrix is symmetric as the reward function is invariant to the order of the perturbations. We discuss further details of the algorithm in Appendix C.

Batching In high-throughput experimental screens, it is possible to run multiple experiments in parallel at the cost of a single experiment. So instead of a single action δ_t, we can select a set of actions $\{\delta_t^1, \ldots, \delta_t^b\}$ where b is the number of experiments we can run in parallel. However, as there are no efficient algorithms for combinatorial bandits in the discovery setting, we resort to a simple greedy scheme to select batches with ASD. Specifically, instead of picking the action which minimizes the information ratio, we pick b actions with the lowest information ratio. The overall approach is summarized in Algorithm 1

4 Empirical Results

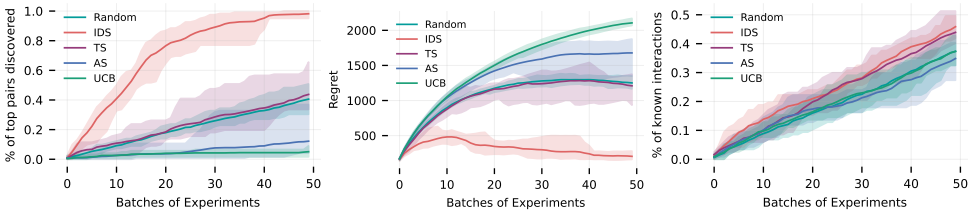

Figure 1: Empirical results on the benchmark task to acquire pairs defined from a selected set of 50 genes. We run experiments with embeddings from a DenseNet-based classifier. The solid lines represent the mean performance, whereas the shaded region represent all the runs (min-max). IDS outperforms all the baselines significantly in terms of top scoring pairs discovered as well as the regret. In terms of known interactions, IDS still outperforms the baselines by a small margin.

Setup We define a new benchmark task for in-silico validation of the approach. We consider a set of 50 genes selected with a bias towards genes with known gene-gene interactions. We then perform in vitro double gene knockout experiments on HUVEC cells with three CRISPR guides per gene. The experimental protocol used for collecting observations from the in vitro experiments is the same as Fay et al. (2023). We embed the images using a DenseNet-based classifier (Huang et al., 2017) trained on the `RxRx1` dataset (Sypetkowski et al., 2023). To ensure consistency, we follow the same data preprocessing procedures proposed by Sypetkowski et al. (2023) which consists of only singe gene interventions. The embeddings are averaged across all guides and replicates, resulting in 1,225 pairwise gene embeddings which form our "unknown" target $h_{i,j}$, and 50 single gene embeddings, h_i, which are assumed to be known at the start of the active learning experiment. The final reward matrix is illustrated in Figure 2.

Baselines In addition to a random policy, we consider as baselines Upper Confidence Bound (UCB), Thompson Sampling (TS) and Uncertainty Sampling (US). To instantiate UCB in the discovery setting, we use the uncertainty from the low-rank matrix posterior in place of the the counts used in the standard multi-armed bandit setting and mask the actions once they are selected. Similarly, in US we pick pairs based solely on the uncertainty of the posterior. TS is instantiated in the same way as IDS, but the pairs are selected to minimize only the instant regret.

Evaluation We evaluate each approach using three different metrics given a budget of 50 rounds, each with a batch size of 10 resulting in a total of 500 experiments covering 50% of all possible pairs of the gene set. First, we look at the fraction of the pairs with the top 5 percentile of scores recovered by the algorithm capturing the ability to explore the high scoring regions well. Next, we also evaluate the regret of each algorithm with respect to an optimal policy with access to the score matrix and acquires the highest scoring pairs at each round. Finally, we evaluate the number of known biological relations that appear in CORUM (Giurgiu et al., 2019), StringDB (Szklarczyk et al., 2021), Signor (Licata et al., 2020) and hu.MAP (Drew et al., 2021) to see how many each method is able to recover.

Results Figure 1 illustrates the empirical results. We observe that IDS discovers all the pairs with the top-5 percentile scores, whereas all the baselines barely recover half of the top pairs. Even in terms of the regret, IDS outperforms the baselines, with random performing the worst. This demonstrates that IDS is able to exploit the low-rank structure in the reward matrix effectively. In terms of the known interactions discovered, however, we observe that the performance of all methods is quite similar. IDS and TS, outperform all the baselines recovering $\sim$10% more known relations. The significant difference between the performance on the fraction of top pairs and the number of known relations prompts questions about the correlation between the prediction errors and the biological interactions.

5 Understanding the Biological Significance

The future applicability of active learning in this context is dependent on a high $R(i,j)$ score both providing information that 1) our independence assumptions have been violated in a biological meaningful way that is indicative of biological interactions, and 2) that this information is at least in part complementary to the squared cosine similarity, $C(i,j) \coloneqq \cos(w_i, w_j)^2$. Cosine similarity is an established metric for inferring biological relationships (Moshkov et al., 2022; Kraus et al., 2023), and is squared here to simplify the comparison between metrics. Figure 3 plots both metrics for all pairs of genes. We see 1) known relationships amount to the highest values for each metric, and 2) the metrics are complementary and not merely overlapping. Several functional categories are preferentially associated with variation in $R(i,j)$ or $C(i,j)$. For instance, autophagy and microtubule-annotated genes were found in a low-$R(i,j)$ regime with a number of samples showing higher $C(i,j)$; in contrast, p53 signaling, proteasome and apoptosis-related genes were found in a low $C(i,j)$ regime, with higher $R(i,j)$ (Figure 4).

We considered whether selecting well characterized gene pairs with elevated $R(i,j)$ would, from a biological perspective, be expected to violate our independent effects assumptions. Removing a presumably pleiotropic gene (e.g. PSMD1, proteasome), which we would expect to have multiple effects, on top of a gene buffering an essential pathway (e.g. BCL2L1, apoptosis), would violate our "No off-target effects" assumption (2.2.2); we might expect to see this gene pair in the high $R(i,j)$, low $C(i,j)$ regime and we do. Removing two genes known to buffer or compensate for each other (BCL2L1 and MCL1, apoptosis) would violate our assumption of "No other effects" (2.2.1), and again we see this gene pair in the high $R(i,j)$, low $C(i,j)$ regime.

We further looked at specific hypotheses for biological interactions. First, we expect that double knockouts of subunits in protein complexes would have an interaction where the phenotype of the double knockout is less severe than the sum of the parts. Taking $S(w) = \sum w \cdot \text{sign}(w_{i,j})$ as a proxy for severity, we find the severity of the double knockout to be very similar to that of the higher impact subunit knockout and less than the sum (Figure 5). Second, in known cases of synthetic lethality, where genes have a redundant function, we expect a more severe phenotype of the double knockout than predicted by the sum of individual knockouts. Here, existing knowledge is more limited and few high confidence literature examples were represented in the dataset. Of those that were, we did not observe the expected lower severity of the double prediction than the double knock out (Figure 6). We hypothesize that this pertains to an overall bias in the data where the average $\sum w_{i,j}$ of double perturbations was 0.76, but the average $\sum w_i + w_j$ notably higher at 1.03, making synthetic lethality interactions harder to detect. Despite this, and as $R(i,j)$ is affected by both total severity *and* direction in embedding space, a well characterized case of synthetic lethality, double knockout of BCL2L1 and EIF4E (Kuzuoglu-Ozturk et al., 2021), was in the highest 7% of $R(i,j)$ scores. Thus, $R(i,j)$ is already useful for detecting biological interactions, while future exploration of normalization and correcting for biases in scale are likely to result in further substantial improvement.

6 Discussion

This paper presented a method for efficiently finding failures of linear combinations of embeddings. We find that we can very efficiently find where linear predictions will fail and where they are likely to be good predictors of pairwise outcomes. This is immediately useful in that it implicitly implies a confidence score that allows us to predict when we can compositionally generalize to unseen pairs of perturbations by leveraging the additivity of embeddings. We also presented theory that described sufficient conditions on the data generating process that would imply this additivity. Violations of

this additivity correspond to failures of these assumption and hence can be used to detect potential causal interactions between genes. While many of the known interactions scored highly according to this metric, the overall correlation between known interactions and this metric was relatively low, suggesting that the test lacks specificity. We suspect that if we could find an alternative embedding function that did not require independence of the features, Z_i, this would reduce the number of false positives. Finally, it would be interesting to experimentally validate some of the high-scoring interactions using different assays to test whether these interactions are biologically meaningful.

REFERENCES

Cell painting, a high-content image-based assay for morphological profiling using multiplexed fluorescent dyes. *Nature Protocols*, 11(9):1757–1774, 2016.

P. Baillargeon, V. Fernandez-Vega, B. P. Sridharan, S. Brown, P. R. Griffin, H. Rosen, B. Cravatt, L. Scampavia, and T. P. Spicer. The scripps molecular screening center and translational research institute. *SLAS DISCOVERY: Advancing Life Sciences R&D*, 24(3):386–397, 2019.

M. Bereket and T. Karaletsos. Modelling cellular perturbations with the sparse additive mechanism shift variational autoencoder. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL https://openreview.net/forum?id=DzaCE00jGV.

E. Bingham, J. P. Chen, M. Jankowiak, F. Obermeyer, N. Pradhan, T. Karaletsos, R. Singh, P. A. Szerlip, P. Horsfall, and N. D. Goodman. Pyro: Deep universal probabilistic programming. *J. Mach. Learn. Res.*, 20:28:1–28:6, 2019. URL http://jmlr.org/papers/v20/18-403.html.

J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL http://github.com/google/jax.

S. N. Chandrasekaran, J. Ackerman, E. Alix, D. M. Ando, J. Arevalo, M. Bennion, N. Boisseau, A. Borowa, J. D. Boyd, L. Brino, P. J. Byrne, H. Ceulemans, C. Ch'ng, B. A. Cimini, D.-A. Clevert, N. Deflaux, J. G. Doench, T. Dorval, R. Doyonnas, V. Dragone, O. Engkvist, P. W. Faloon, B. Fritchman, F. Fuchs, S. Garg, T. J. Gilbert, D. Glazer, D. Gnutt, A. Goodale, J. Grignard, J. Guenther, Y. Han, Z. Hanifehlou, S. Hariharan, D. Hernandez, S. R. Horman, G. Hormel, M. Huntley, I. Icke, M. Iida, C. B. Jacob, S. Jaensch, J. Khetan, M. Kost-Alimova, T. Krawiec, D. Kuhn, C.-H. Lardeau, A. Lembke, F. Lin, K. D. Little, K. R. Lofstrom, S. Lotfi, D. J. Logan, Y. Luo, F. Madoux, P. A. M. Zapata, B. A. Marion, G. Martin, N. J. McCarthy, L. Mervin, L. Miller, H. Mohamed, T. Monteverde, E. Mouchet, B. Nicke, A. Ogier, A.-L. Ong, M. Osterland, M. Otrocka, P. J. Peeters, J. Pilling, S. Prechtl, C. Qian, K. Rataj, D. E. Root, S. K. Sakata, S. Scrace, H. Shimizu, D. Simon, P. Sommer, C. Spruiell, I. Sumia, S. E. Swalley, H. Terauchi, A. Thibaudeau, A. Unruh, J. V. de Waeter, M. V. Dyck, C. van Staden, M. Warchoł, E. Weisbart, A. Weiss, N. Wiest-Daessle, G. Williams, S. Yu, B. Zapiec, M. Żyła, S. Singh, and A. E. Carpenter. Jump cell painting dataset: morphological impact of 136,000 chemical and genetic perturbations. *bioRxiv*, 2023.

W. Chen, Y. Wang, and Y. Yuan. Combinatorial multi-armed bandit: General framework and applications. In *International conference on machine learning*, pages 151–159. PMLR, 2013.

R. Combes, M. S. Talebi Mazraeh Shahi, A. Proutiere, et al. Combinatorial bandits revisited. *Advances in neural information processing systems*, 28, 2015.

Z. Dai, Q. P. Nguyen, S. S. Tay, D. Urano, R. Leong, B. K. H. Low, and P. Jaillet. Batch bayesian optimization for replicable experimental design. *arXiv preprint arXiv:2311.01195*, 2023.

T. Desautels, A. Krause, and J. W. Burdick. Parallelizing exploration-exploitation tradeoffs in gaussian process bandit optimization. *Journal of Machine Learning Research*, 15:3873–3923, 2014.

J. A. Doudna and E. Charpentier. The new frontier of genome engineering with crispr-cas9. *Science*, 346(6213):1258096, 2014.

A. Dove. High-throughput screening goes to school. *Nature Methods*, 4(6):523–532, 2007.

K. Drew, J. B. Wallingford, and E. M. Marcotte. hu. map 2.0: integration of over 15,000 proteomic experiments builds a global compendium of human multiprotein assemblies. *Molecular systems biology*, 17(5):e10016, 2021.

M. M. Fay, O. Kraus, M. Victors, L. Arumugam, K. Vuggumudi, J. Urbanik, K. Hansen, S. Celik, N. Cernek, G. Jagannathan, et al. Rxrx3: Phenomics map of biology. *bioRxiv*, pages 2023–02, 2023.

A. Foster. *Variational, Monte Carlo and policy-based approaches to Bayesian experimental design*. PhD thesis, University of Oxford, 2021.

R. Garnett, Y. Krishnamurthy, X. Xiong, J. Schneider, and R. Mann. Bayesian optimal active search and surveying. In *Proceedings of the 29th International Coference on International Conference on Machine Learning*, pages 843–850, 2012.

M. Giurgiu, J. Reinhard, B. Brauner, I. Dunger-Kaltenbach, G. Fobo, G. Frishman, C. Montrone, and A. Ruepp. Corum: the comprehensive resource of mammalian protein complexes—2019. *Nucleic acids research*, 47(D1):D559–D563, 2019.

G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.

K. Huang, R. Lopez, J.-C. Hutter, T. Kudo, A. Rios, and A. Regev. Sequential optimal experimental design of perturbation screens guided by multi-modal priors. *bioRxiv*, pages 2023–12, 2023.

M. Jinek, K. Chylinski, I. Fonfara, M. Hauer, J. A. Doudna, and E. Charpentier. A programmable dual-rna–guided dna endonuclease in adaptive bacterial immunity. *science*, 337(6096):816–821, 2012.

V. Kanade, H. B. McMahan, and B. Bryan. Sleeping experts and bandits with stochastic action availability and adversarial rewards. In *Artificial Intelligence and Statistics*, pages 272–279. PMLR, 2009.

A. Karbasi, V. Mirrokni, and M. Shadravan. Parallelizing thompson sampling. *Advances in Neural Information Processing Systems*, 34:10535–10548, 2021.

O. Kraus, K. Kenyon-Dean, S. Saberian, M. Fallah, P. McLean, J. Leung, V. Sharma, A. Khan, J. Balakrishnan, S. Celik, et al. Masked autoencoders are scalable learners of cellular morphology. *arXiv preprint arXiv:2309.16064*, 2023.

D. Kuzuoglu-Ozturk, Z. Hu, M. Rama, E. Devericks, J. Weiss, G. G. Chiang, S. T. Worland, S. E. Brenner, H. Goodarzi, L. A. Gilbert, et al. Revealing molecular pathways for cancer cell fitness through a genetic screen of the cancer translatome. *Cell reports*, 35(13), 2021.

T. Lattimore and C. Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.

L. Licata, P. Lo Surdo, M. Iannuccelli, A. Palma, E. Micarelli, L. Perfetto, D. Peluso, A. Calderone, L. Castagnoli, and G. Cesareni. Signor 2.0, the signaling network open resource 2.0: 2019 update. *Nucleic acids research*, 48(D1):D504–D510, 2020.

D. V. Lindley. On a measure of the information provided by an experiment. *The Annals of Mathematical Statistics*, 27(4):986–1005, 1956.

R. Lopez, N. Tagasovska, S. Ra, K. Cho, J. Pritchard, and A. Regev. Learning causal representations of single cells via sparse mechanism shift modeling. In M. van der Schaar, C. Zhang, and D. Janzing, editors, *Proceedings of the Second Conference on Causal Learning and Reasoning*, volume 213 of *Proceedings of Machine Learning Research*, pages 662–691. PMLR, 11–14 Apr 2023. URL `https://proceedings.mlr.press/v213/lopez23a.html`.

M. Lotfollahi, A. Klimovskaia Susmelj, C. De Donno, L. Hetzel, Y. Ji, I. L. Ibarra, S. R. Srivatsan, M. Naghipourfar, R. M. Daza, B. Martin, J. Shendure, J. L. McFaline-Figueroa, P. Boyeau, F. A. Wolf, N. Yakubova, S. Günnemann, C. Trapnell, D. Lopez-Paz, and F. J. Theis. Predicting cellular responses to complex perturbations in high-throughput screens. *Molecular Systems Biology*, 19(6):e11517, 2023. doi: https://doi.org/10.15252/msb.202211517. URL `https://www.embopress.org/doi/abs/10.15252/msb.202211517`.

C. Lyle, A. Mehrjou, P. Notin, A. Jesson, S. Bauer, Y. Gal, and P. Schwab. Discobax discovery of optimal intervention sets in genomic experiment design. In *International Conference on Machine Learning*, pages 23170–23189. PMLR, 2023.

A. Mehrjou, A. Soleymani, A. Jesson, P. Notin, Y. Gal, S. Bauer, and P. Schwab. Genedisco: A benchmark for experimental design in drug discovery. In *International Conference on Learning Representations*, 2021.

N. Moshkov, M. Bornholdt, S. Benoit, M. Smith, C. McQuin, A. Goodman, R. A. Senft, Y. Han, M. Babadi, P. Horvath, et al. Learning representations for image-based profiling of perturbations. *Biorxiv*, pages 2022–08, 2022.

S. M. Nijman. Synthetic lethality: General principles, utility and detection using genetic screens in human cells. *FEBS Letters*, 585(1):1–6, 2011. ISSN 0014-5793.

A. Pacchiano, J. Lee, and E. Brunskill. Experiment planning with function approximation. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023a.

A. Pacchiano, D. Wulsin, R. A. Barton, and L. Voloch. Neural design for genetic perturbation experiments. In *The Eleventh International Conference on Learning Representations*, 2023b.

D. Phan, N. Pradhan, and M. Jankowiak. Composable effects for flexible and accelerated probabilistic programming in numpyro. *arXiv preprint arXiv:1912.11554*, 2019.

T. Rainforth, R. Cornish, H. Yang, A. Warrington, and F. Wood. On nesting monte carlo estimators. In *International Conference on Machine Learning*, pages 4267–4276. PMLR, 2018.

T. Rainforth, A. Foster, D. R. Ivanova, and F. B. Smith. Modern bayesian experimental design. *arXiv preprint arXiv:2302.14545*, 2023.

R. Ranganath, S. Gerrish, and D. Blei. Black box variational inference. In *Artificial intelligence and statistics*, pages 814–822. PMLR, 2014.

H. Robbins. Some aspects of the sequential design of experiments. 1952.

G. Roeder, L. Metz, and D. Kingma. On linear identifiability of learned representations. In M. Meila and T. Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 9030–9039. PMLR, 18–24 Jul 2021. URL https://proceedings.mlr.press/v139/roeder21a.html.

D. Russo and B. Van Roy. An information-theoretic analysis of thompson sampling. *The Journal of Machine Learning Research*, 17(1):2442–2471, 2016.

E. G. Ryan, C. C. Drovandi, J. M. McGree, and A. N. Pettitt. A review of modern computational algorithms for bayesian optimal design. *International Statistical Review*, 84(1):128–154, 2016.

P. Sebastiani and H. P. Wynn. Maximum entropy sampling and optimal bayesian experimental design. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 62(1):145–157, 2000.

N. Srinivas, A. Krause, S. Kakade, and M. Seeger. Gaussian process optimization in the bandit setting: no regret and experimental design. In *Proceedings of the 27th International Conference on International Conference on Machine Learning*, pages 1015–1022, 2010.

M. Sypetkowski, M. Rezanejad, S. Saberian, O. Kraus, J. Urbanik, J. Taylor, B. Mabey, M. Victors, J. Yosinski, A. R. Sereshkeh, et al. Rxrx1: A dataset for evaluating experimental batch correction methods. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4284–4293, 2023.

D. Szklarczyk, A. L. Gable, K. C. Nastou, D. Lyon, R. Kirsch, S. Pyysalo, N. T. Doncheva, M. Legeay, T. Fang, P. Bork, et al. The string database in 2021: customizable protein–protein networks, and functional characterization of user-uploaded gene/measurement sets. *Nucleic acids research*, 49 (D1):D605–D612, 2021.

W. R. Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3-4):285–294, 1933.

Z. Wang, L. Gui, J. Negrea, and V. Veitch. Concept algebra for (score-based) text-controlled generative models. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL https://openreview.net/forum?id=SGlrCuwdsB.

D. Wingate and T. Weber. Automated variational inference in probabilistic programming. *arXiv preprint arXiv:1301.1299*, 2013.

Z. Xu, E. Shim, A. Tewari, and P. Zimmerman. Adaptive sampling for discovery. *Advances in Neural Information Processing Systems*, 35:1114–1126, 2022.

D. Xun, R. Wang, X. Zhang, and Y. Wang. Microsnoop: A generalized tool for unbiased representation of diverse microscopy images. *bioRxiv*, 2023. doi: 10.1101/2023.02.25.530004. URL https://www.biorxiv.org/content/early/2023/05/06/2023.02.25.530004.

A. Zanette, K. Dong, J. N. Lee, and E. Brunskill. Design of experiments for stochastic contextual linear bandits. *Advances in Neural Information Processing Systems*, 34:22720–22731, 2021.

A BACKGROUND

Representation learning for gene knockouts Our work builds on the classifier-based approach to embedding cells presented in Sypetkowski et al. (2023). More recently, (Kraus et al., 2023; Xun et al., 2023) showed that you can get improved embeddings for detecting biological relationships with cosine similarity by using masked autoencoders. There have also been a number of papers that have attempted to learn disentangled representations of cells (e.g. Lotfollahi et al., 2023; Bereket and Karaletsos, 2023; Lopez et al., 2023), mostly from expression data. Our work takes inspiration from Wang et al. (2023)'s concept algebra framework that analyses the properties of the score function as a representation; we use the final hidden layer of a classifier which naturally leads to a different set of assumptions and proof strategy.

Experimental design Experimental capacity for in-vivo experiments is often limited due to the complexity and cost associated with running each experiment. As a consequence, designing of experiments that acquire the most information about the system efficiently is an important and well-studied problem. The problem can be formalized as sequential Bayesian optimal experimental design (BOED; Ryan et al., 2016; Foster, 2021; Rainforth et al., 2023), where the goal to design experiments $x \in \mathcal{X}$ with outcomes $y \in \mathcal{Y}$ governed by a generative process $y \sim p(y \mid \gamma, x)$ with parameters γ. The experiments are performed sequentially $(x_1, \ldots, x_T)$, with the objective of maximizing a measure of utility: the information gain (Lindley, 1956; Sebastiani and Wynn, 2000).

$$x_t^\star = \arg \max_{x \in \{\mathcal{X} \setminus D_x^t\}} MI(\gamma; \{x, y\} \mid D^t) \tag{3}$$

where $D_i = \{(x_i, y_i)\}_{i=1}^{t-1}$ denotes the set of experiments performed till step $t-1$. Solving this optimization problem is challenging in general as even computing the mutual information requires estimating the posterior $p(\gamma \mid D_i)$ and knowledge of the generative process $p(y \mid \gamma, x)$ and nested integral over y and γ which can suffer from poor convergence rates when estimated from samples using a NMC estimator (Rainforth et al., 2018). This complexity is compounded further when x is high-dimensional. With the availability of high-throughput screenings (Dove, 2007; Baillargeon et al., 2019), it is possible to run a batch of experiments in parallel at roughly the same cost. The problem then becomes selecting a sequence of *batches* of experiments instead of sequence of experiments. This also adds further complexity to the experimental design problem as one has to search over a combinatorial space of sets of experiments to be selected at each step (Dai et al., 2023). An instantiation of the experimental design framework is active search (Garnett et al., 2012) where the goal is to pick informative experiments in the presence of binary feedback.

Bandits The framework of multi-armed bandits (Thompson, 1933; Robbins, 1952; Lattimore and Szepesvári, 2020) encapsulates the exploration-exploitation dilemma that is central in a wide variety of decision-making problems. The most general instantiation of the problem involves learning a policy π which maps a history of observations to a distribution over a set of possible actions $\mathcal{A}$ where each action $a \in \mathcal{A}$ is associated with an unknown potentially stochastic reward $f(a)$ such that after T actions sampled from the policy $(a_1, \ldots, a_T)$ the regret $R_T = T \max_{a \in \mathcal{A}} \mathbb{E}[f(a)] - \sum_{i=1}^{T} f(a_i)$ is minimized. There is a rich connection between bandits and the experimental design problem (Zanette et al., 2021; Pacchiano et al., 2023a). For instance the regret bounds for bandit optimization in some settings are characterized in terms of the information gain (Srinivas et al., 2010). While the batch setting, where the policy selects a batch of candidates in each round, is not studied extensively in the context of bandits, there are parallelized version of the UCB (Desautels et al., 2014) and Thompson Sampling (Karbasi et al., 2021) algorithms. This problem of selecting batches of actions can also be viewed as an instantiation of *combinatorial* bandits (Chen et al., 2013; Combes et al., 2015).

Design of Gene Knockout Experiments Due to the profound potential impact of genetic interventions, there has been considerable work in the recent years to develop algorithms for design gene knockout experiments. Typically the problem is studied in the context of discovering single gene-knockouts (Mehrjou et al., 2021) which result in a particular phenotypic effect of interest. A variety of methods including bandits (Pacchiano et al., 2023b) and traditional experimental design (Lyle et al., 2023) approaches have been studied in this context. Further, approaches which aim to learn predictors for the effects of unseen gene knockouts typically operate on RNA-seq data (Huang et al., 2023). RNA-Seq data can have several advantages but also disadvantages. While PerturbSeq scales

well for double perturbations, a major drawback is that perturbations are stochastic and we can only observe perturbations which survive long enough to be measured.

B PROOF OF THEOREM 2.3

Theorem 2.3 *Assume the data generating process given in Equation 1, and that Assumption 2.1 and 2.2 hold.*

$$h_i + h_j = h_{i,j}.$$

That is, the sum of the centered individual perturbation embeddings equals the centered pairwise embedding.

Proof. First note that under our assumptions, by the change of variable formula, we can rewrite the conditional density of the images, X, for any pair of gene knockouts, δ_i, δ_j, as,

$$\begin{aligned} P(X = x|\delta_i, \delta_j) &= P(g^{-1}(X = x)|\delta_i, \delta_j) \left|\det\left(\frac{\partial g^{-1}}{\partial y}\right)\right| = \prod_i P(Z_i|\delta_i, \delta_j)P(\epsilon) \left|\det\left(\frac{\partial g^{-1}}{\partial y}\right)\right| \\ &= P(Z_l^i|\delta_i)P(Z_l^j|\delta_j)P(Z_l^n)\prod_{i\neq j} P(Z_i)P(\epsilon)|\det(\frac{\partial g}{\partial y})| \end{aligned}$$

Now consider the following single and double knockouts as a ratio with the control where $\mathbf{x} = 0$ (i.e. $\delta_i = 0$ for all i),

$$\begin{aligned} \frac{P(X = x|\delta_i = 1)}{P(X = x|\delta = 0)} &= \frac{[P(Z_l^i|\delta_i = 1) + P(Z_l^j|\delta_j = 0) + P(Z_l^n)]\prod_{i\neq j} P(Z_i)P(\epsilon)|\det(\frac{\partial g}{\partial y})|}{[P(Z_l^i|\delta_i = 0) + P(Z_l^j|\delta_j = 0) + P(Z_l^n)]\prod_{i\neq j} P(Z_i)P(\epsilon)|\det(\frac{\partial g}{\partial y})|} \\ &= \frac{P(Z_l^i|\delta_i = 1) + P(Z_l^j) + P(Z_l^n)}{P(Z_l^i) + P(Z_l^j) + P(Z_l^n)} \\ \frac{P(X = x|\delta_j = 1)}{P(X = x|\delta = 0)} &= \frac{P(Z_l^i) + P(Z_l^j|\delta_j = 1) + P(Z_l^n)}{P(Z_l^i) + P(Z_l^j) + P(Z_l^n)} \\ \frac{P(X = x|\delta_i, \delta_j = 1)}{P(X = x|\delta = 0)} &= \frac{P(Z_l^i|\delta_i = 1) + P(Z_l^j|\delta_j = 1) + P(Z_l^n)}{P(Z_l^i) + P(Z_l^j) + P(Z_l^n)} \\ \frac{P(X = x|\delta = 0)}{P(X = x|\delta = 0)} &= \frac{P(Z_l^i) + P(Z_l^j) + P(Z_l^n)}{P(Z_l^i) + P(Z_l^j) + P(Z_l^n)} = 1 \end{aligned}$$

Now, notice that,

$$\begin{aligned} &\frac{P(X = x|\delta_i = 1)}{P(X = x|\delta = 0)} + \frac{P(X = x|\delta_j = 1)}{P(X = x|\delta = 0)} - \frac{P(X = x|\delta_i = 1, \delta_j = 1)}{P(X = x|\delta = 0)} - \frac{P(X = x|\delta_i = 0, \delta_j = 0)}{P(X = x|\delta = 0)} \\ =&\frac{P(Z_l^i|\delta_i = 1) + P(Z_l^j) + P(Z_l^n) + P(Z_l^i) + P(Z_l^j|\delta_j = 1) + P(Z_l^n)}{P(Z_l^i) + P(Z_l^j) + P(Z_l^n)} - \\ &\frac{P(Z_l^i|\delta_i = 1) + P(Z_l^j|\delta_j = 1) + P(Z_l^n) + P(Z_l^i) + P(Z_l^j) + P(Z_l^n)}{P(Z_l^i) + P(Z_l^j) + P(Z_l^n)} = 0 \end{aligned} \tag{4}$$

Now consider,

$$\begin{aligned} &h_i + h_j - h_0 - h_{i,j} \\ &= \mathbb{E}\left[\left[P(X|\delta_i = 1) + P(X|\delta_j = 1) - P(X|\delta = 0) - P(X|\delta_i = 1, \delta_j = 1)\right]h(X)\right] \\ &= \mathbb{E}\left[\underbrace{\left[\frac{P(X|\delta_i = 1)}{P(X|\delta = 0)} + \frac{P(X|\delta_j = 1)}{P(X|\delta = 0)} - \frac{P(X|\delta = 0)}{P(X|\delta = 0)} - \frac{P(X|\delta_i = 1, \delta_j = 1)}{P(X|\delta = 0)}\right]}_{=0} P(X|\delta = 0)h(X)\right] \\ &= 0 \end{aligned}$$

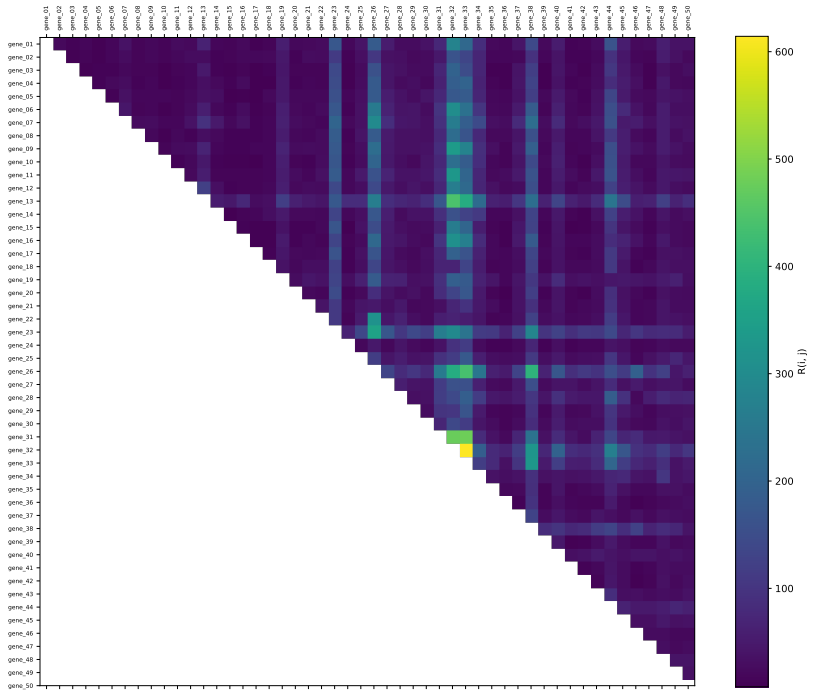

Figure 2: Reward matrix consisting of the prediction error for all possible gene pairs from the selected set of genes.

The first equality follows by definition of the embeddings, the second uses the fact that the density ratio in Equation 4 sums to zero. Given that this expression evaluates to zero, by rearranging the terms we get that,

$$h_i + h_j = h_0 + h_{i,j}$$

as required.

□

C EXPERIMENTAL DETAILS

We denote the set of possible distributions defined over Δ by $\mathcal{D}(\Delta)$, the history of actions and their corresponding rewards until round t by $H_t = ((\delta_k, R(\delta_k)))_{k=1}^{t-1}$, and Δ_t denotes the set of available designs at round t. The policy π_{IDS} is defined as a map from H_t to $\mathcal{D}(\Delta_t)$. As we consider a discrete space of actions, each element of $\mathcal{D}(\Delta_t)$ is a vector which represents the distribution over available actions. We define the instant regret of taking an action $\Delta_t(\delta) = \mathbb{E}_{f\sim p(R|H_t)}[\max_{\delta' \in \Delta_t} f(\delta') - f(\delta)]$,

Algorithm 1 ASD for Designing Gene Pair Knockouts

```
Initialize H_1 = {}
for t = 1, ..., T do
    Estimate posterior p(R | H_t)
    Compute information ratio Ψ_t
    Pick batch (δ_1, ..., δ_b) greedily which minimize Ψ_t
    Perform experiments and compute R(x)
    Update H_{t+1} = H_t ∪ {(δ_t^1, R(δ_t^1)), ..., (δ_t^b, R(δ_t^b))}
end for
```

where the expectation is over samples from the posterior over the reward function f given the history of observations H_t. Additionally, we define $g(x) = MI(\delta^*_{t,1}, \ldots, \delta^*_{t,T-t+1}; R(\delta_t)|H_t, \delta_t = x)$ which represents the information gain about the top $T - t + 1$ unselected actions. The IDS policy at round t can be then be computed by minimizing the *information ratio*

$$\pi_{\text{IDS}} \in \arg \min_{\pi \in \mathcal{D}(\Delta_t)} \Psi_{\pi,t} := \frac{(\Delta_t^\top \pi)^\lambda}{g_t^\top \pi} \tag{5}$$

λ controls the tradeoff between lower instant regret (exploitation) and higher information gain (exploration). Following Russo and Van Roy (2016); Xu et al. (2022) we rely on an approximate algorithmic choice, replacing g_t with the conditional variance $v_t(x) = \text{Var}_t(\mathbb{E}[R(\delta_t)|\delta^*_1, \delta_t = x])$, as it is lower bound on the information gain $g_t(x) \geq v_t(x)$. We adopt the low-rank matrix with a prior of row and column spaces being sampled from a standard Gaussian. To obtain samples from the posterior distribution over the low-rank (where m is the rank) reward matrix, we use stochastic variational inference (Wingate and Weber, 2013; Ranganath et al., 2014), and specifically the implementation from `numpyo` (Bradbury et al., 2018; Bingham et al., 2019; Phan et al., 2019). We train the variational posterior for 5000 epochs with a learning rate of 0.01 using the Adam optimizer. We then generate k samples from the posterior.

For each algorithm, we run a sweep over all a set of hyperparameters. We then pick the best hyperparameters and run the experiment over 10 different seeds to get the final results. For IDS, we tune $m \in \{3, 5, 7, 10, 12\}$, $\lambda \in \{2, 3, 4, 5\}$ and $k \in \{500, 750, 1000, 1500\}$. For TS and US, we tune $m \in \{3, 5, 7, 10, 12\}$ and $k \in \{500, 750, 1000, 1500\}$. For UCB we tune $m \in \{3, 5, 7, 10, 12\}$, $\beta \in \{0.01, 0.1, 0.2, 0.5, 1, 2, 5\}$ and $k \in \{500, 750, 1000, 1500\}$, where β controls the exploration.

D Additional Figures

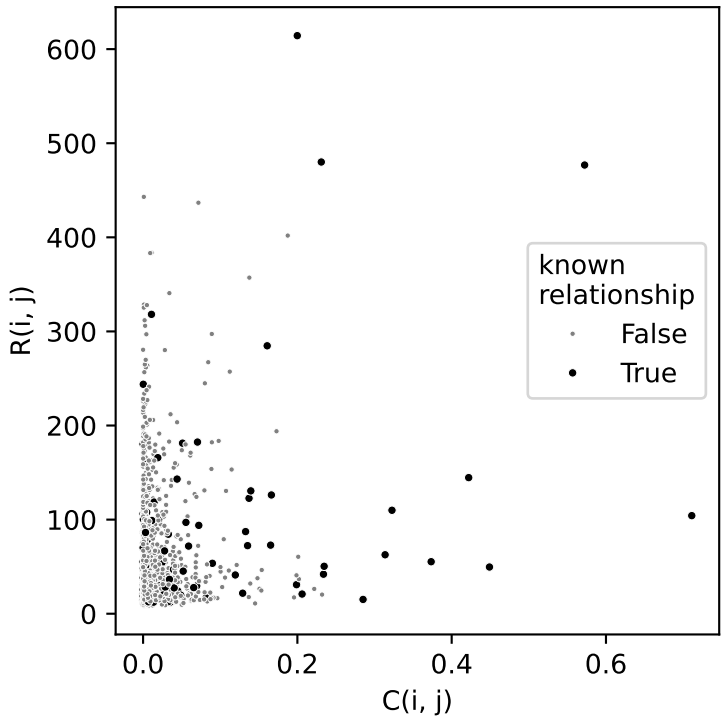

Figure 3: Contrasts the established cosine similarity with the l2 norm. Both metrics show gene pairs with known relationships from the hu.Map, Corum, Signor, or StringDB among their highest values.

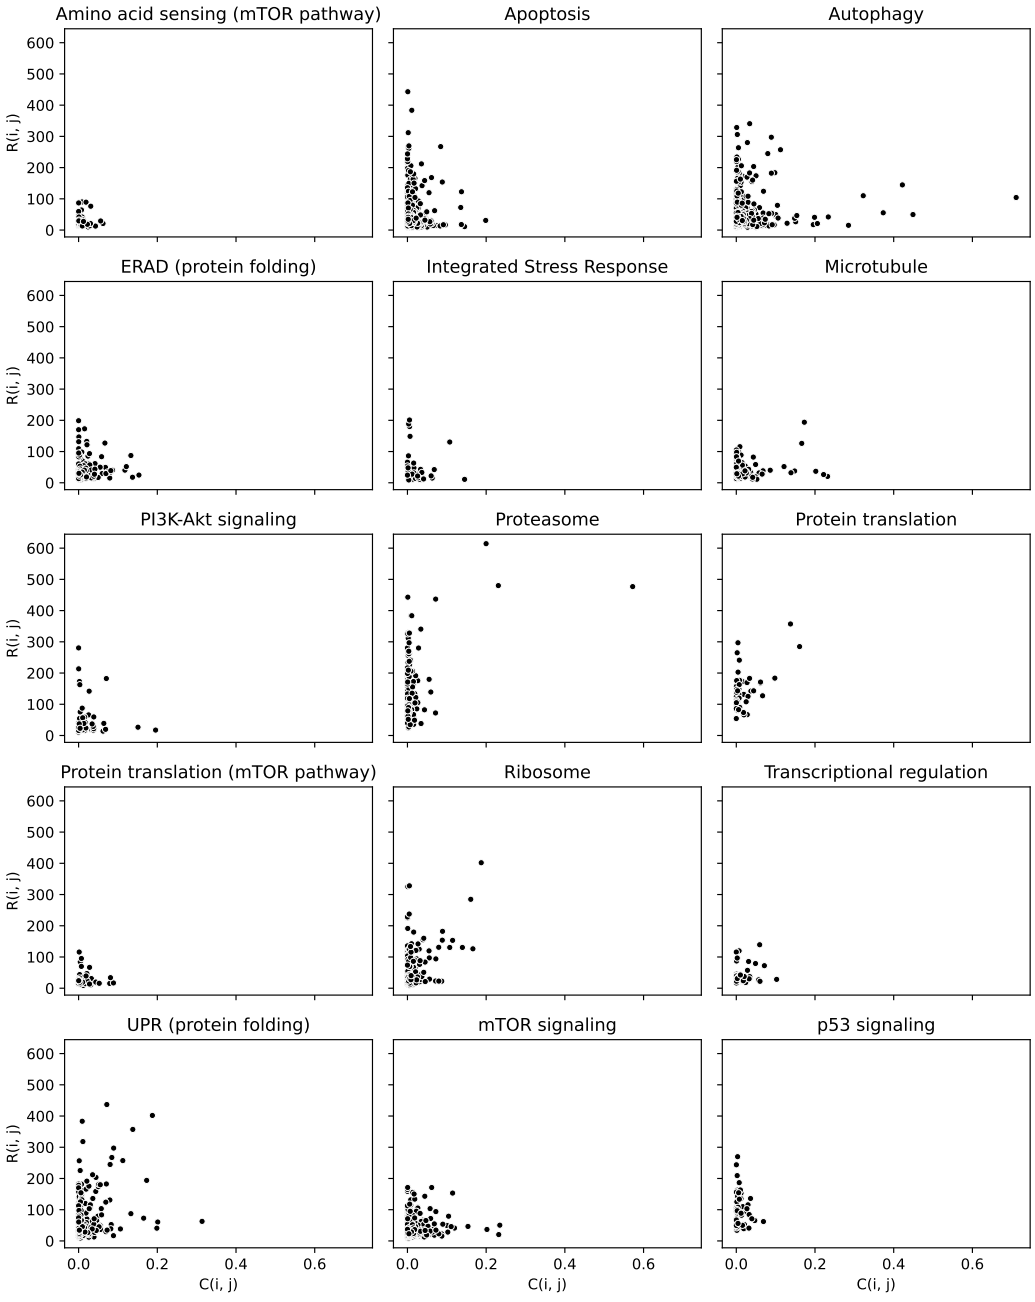

Figure 4: Breaks down the relationship between l2 norm and cosine similarity by the biological pathway of *either* targeted gene.

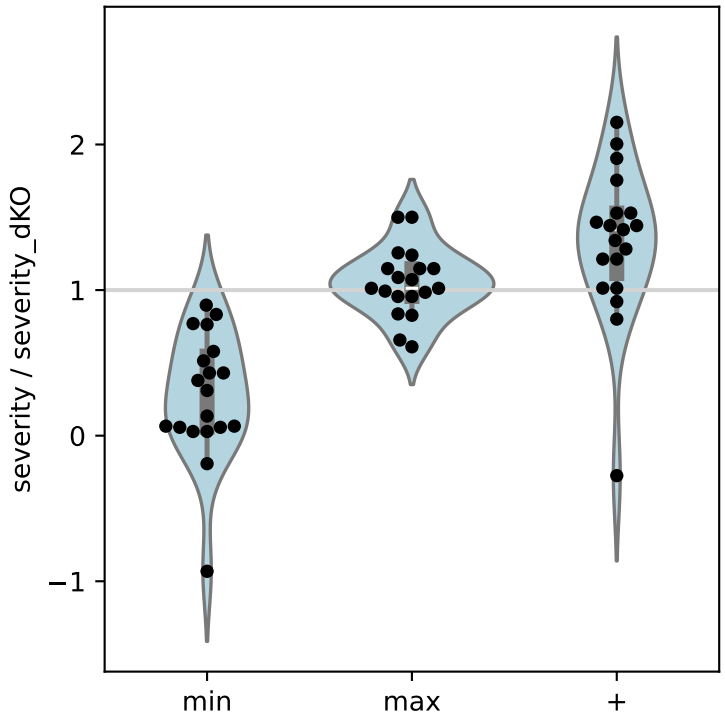

Figure 5: Shows the phenotype severity proxy, S, for the lower severity (min), higher severity (max), and sum (+) of individual embeddings. Values are normalized to the embedding of the double knockout. All gene pairs plotted here have an 'in complex' relationship in Corum.

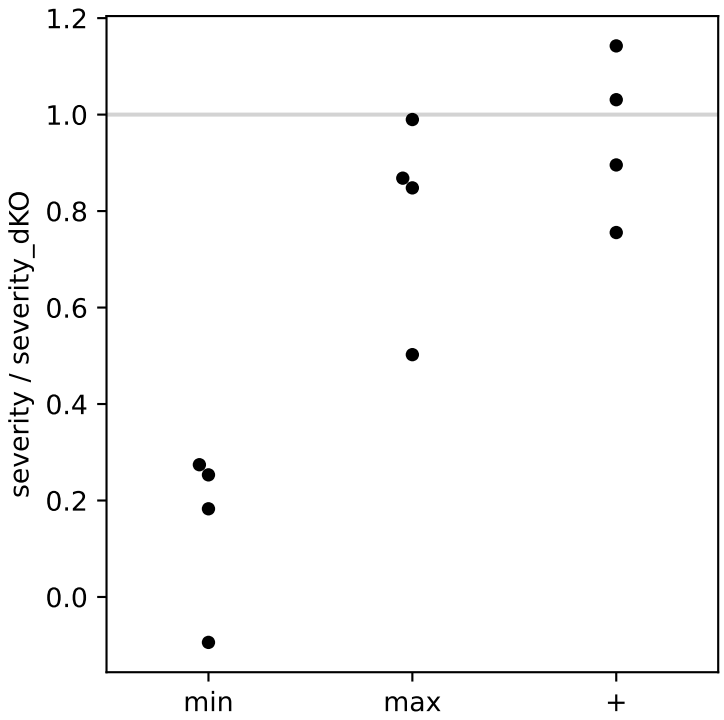

Figure 6: Shows the phenotype severity proxy, S, for the lower severity (min), higher severity (max) individual embeddings, and for the sum (+) of the individual embeddings. Values are normalized to the embedding of the double knockout. All gene pairs plotted here were manually identified as having a candidate synthetic lethality relationship by literature search.