Learning Perturbation-specific Cell Representations for Prediction of Transcriptional Response across Cellular Contexts

Gal Keinan[1], Karen Sayal[*1, 2], Alon Gonen[*1], Jiang Zhu[1], Lena Granovsky[1], and Jeremy England[1]

[1]Artificial Intelligence and Machine Learning, GSK
[2]Department of Oncology, Royal Berkshire NHS Foundation Trust, UK

Abstract

Pre-clinical and clinical scientific endeavours provide complementary perspectives on fundamental biological processes of translational value. Harnessing the information value on the totality of such knowledge requires novel approaches to integrate across the breadth of experimental space. High-throughput screens are often the first step on this bridge to the patient. However, their representative capacity to encompass all the cellular contexts encountered in a patient are often limited due to experimental constraints. Thus, we present PerturbX, a new deep learning model to predict transcriptional responses to chemical or genetic perturbations in unobserved cellular contexts, and to uncover interpretable factors of variation associated with the predicted response. PerturbX can be applied in both an unimodal or multimodal setting. We believe that our proposed approach has the ability to inform novel biomarker discovery and contribute to a redefinition of the drug development pipeline.

1 Introduction

The opportunities in precision oncology are expanding at an increasing pace. One important driving factor has been the explosion in diverse pre-clinical and clinical datasets. Key initiatives include the Cancer Dependency Map (DepMap) project (Cancer Cell Line Encyclopedia & Genomics of Drug Sensitivity in Cancer, 2015), The Cancer Genome Atlas (TCGA) program (Weinstein et al., 2013) and the Human Cell Atlas (HCA) (Rozenblatt-Rosen et al., 2017).

An unresolved question remains as to how best to harness and integrate the insights gained by such distinct but inter-related endeavours. One component will involve the identification and validation of signatures to map between the pre-clinical and clinical settings. It should be a bidirectional relationship. Insights gained in the pre-clinical domain should be of value in the clinical domain and vice versa.

In such a light, information-rich phenotypes provide opportunities to capture a detailed picture of the cellular consequences of chemical or genetic perturbations which, in specific contexts, may have a downstream effect on cellular fitness. Gene expression profiles in particular are a robust and informative phenotypic measure of cellular responses to perturbations. They can be applied to a wide range of applications, including advancing our understanding of gene regulatory networks

[*]These authors contributed equally to this work

and the mechanism of action of screened compounds. Furthermore, single cell level cellular profiles enable interrogation of phenotypic heterogeneity at a level which has been hitherto inaccessible.

High-throughput screens are utilized to profile transcriptional states across multiple cellular contexts under different perturbations (McFarland et al., 2020). However, their representative capacity is still limited relative to the vast combinatorial landscape of all cell type - perturbation pairs. Therefore, a computational model which learns from observed experimental data to predict outcomes in unobserved cellular contexts is of immense value.

We propose PerturbX, a novel deep learning model trained to predict the transcriptional effect of a given perturbation in a range of cellular contexts (e.g. different cell types), in which the perturbational response has not been experimentally observed. PerturbX is based on an encoder - decoder architecture that learns a mapping between the unperturbed state representation of cells to the transcriptional effect of a given perturbation. Furthermore, we show that the model is able to uncover interpretable factors of variation within the unperturbed state which are associated with the observed transcriptional response patterns. These "biomarkers" of response could ultimately be mapped between the pre-clinical and clinical settings and thereby advance our understanding of the representative capacity of cell models for patient-based in vivo biology.

In the setting presented in this paper, the unperturbed state is represented by the unperturbed gene expression profiles of cell lines, and the transcriptional effect is represented by the vector of differential expression. However, the proposed approach is flexible and can be applied in a multimodal fashion. Gene expression profiles at the input could be easily expanded by incorporating mutational profiles as one example. Such considerations would be guided by individual project-specific requirements.

Our main contributions are:

1. We introduce PerturbX and demonstrate its ability to successfully predict transcriptional responses to perturbations in unseen cellular contexts using data readily available in the public domain.

2. We show that PerturbX can learn biologically meaningful and interpretable representations of cell types.

3. We propose a method for identifying the predictive features, or biomarkers, captured by PerturbX.

2 Related Work

Several methods have been recently proposed to address the challenge of predicting the transcriptional response of cells to a range of perturbations. These methods are predominantly based on variational approaches.

scGen (Lotfollahi et al., 2019) utilises a conditional variational autoencoder (CVAE) to predict the post-perturbation single cell distribution in unobserved cell types. trVAE (Lotfollahi et al., 2019) also uses a CVAE with the addition of Maximum Mean Discrepancy regularisation to the decoder network. Compositional perturbation autoencoder (CPA) (Lotfollahi et al., 2021) aims to predict transcriptional responses across different cell types, doses and combinations. It learns disentangled representations of these attributes in order to then generate predictions from unseen combinations of the observed attributes. However, CPA does not generalise to novel cell types or novel perturbations. ChemCPA (Hetzel et al., 2022) is a recent extension of CPA which can generalise predictions to unseen perturbations. sVAE+ (Lopez et al., 2022) is a variant of the sparse VAE model (Lachapelle & Lacoste-Julien, 2022) which assumes that each perturbation changes a small number of latent variables. It exploits the sparse mechanism shift assumption to learn causally interpretable disentangled representations.

In contrast to previous works, we suggest a non-generative approach to the problem of post-perturbation transcriptional prediction. We formulate this as a regression problem in which we aim to learn a mapping from an unperturbed state representation of a specific cellular context to the corresponding transcriptional effect under a given perturbation. Such a mapping allows our model to generalise to unseen cellular contexts and to discover predictive signatures of response which can ultimately be carried over to and utilised in clinical settings. Furthermore, this choice encourages latent proximity between cell types which respond similarly, rather than cell types with similar unperturbed expression. Consequently, the resulting latent representation is simpler and more relevant to the task at hand. [1]

3 PerturbX

PerturbX is a deep learning model trained to predict the transcriptional effect of a perturbation across cellular contexts by learning a *perturbation-specific* cell type representation. This representation captures the complex transcriptional similarities between different cell types in response to a specific compound or genetic perturbation, and is inferred from features of the unperturbed state of a cell. The proposed method allows for generalisation on experimentally unobserved cell types.

The model architecture is based on an encoder-decoder network. However, unlike classical autoencoders, we do not aim to reconstruct the input but rather predict a high-dimensional response vector.

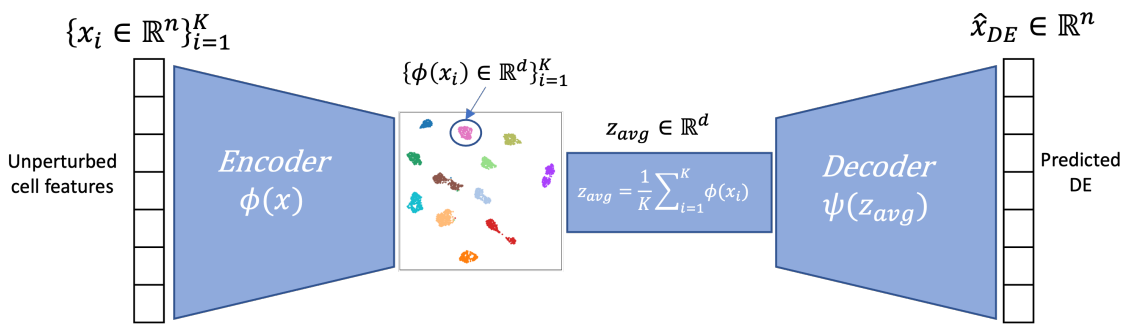

Figure 1: Schematic of the PerturbX model architecture. Given K single cells of a specific cell type as input, the model's encoder maps each of them into the latent space. The K different embeddings are averaged to summarize the latent representation for that cell type. The averaged representation is then mapped by the decoder into a prediction of the differential expression.

The input to the model is a set $\{x_i \in \mathbb{R}^n\}_{i=1}^K$ of K *single cells* from a specific cell type. Each x_i is a n-dimensional vector representing the cellular context of the cells in their unperturbed state. A natural choice for this representation is the unperturbed expression profile of n genes in the cell which can be obtained on the single-cell level. This could be expanded to include additional modalities, such as the mutation status or copy number variations (CNV) for each gene. However, it is expected that these latter modalities would be homogenous between single cells arising from the same cell type. The output of the model is the predicted differential expression (DE) vector, $\hat{x}_{DE} \in \mathbb{R}^n$ which indicates the transcriptional effect on that cell type. For simplicity, we define this target vector to be the difference in mean expression values between the perturbed and control (e.g. DMSO-treated) single cell populations, $x_{DE} = (\overline{X}_{perturbed} - \overline{X}_{DMSO})$. We note one could

[1] We note that some of the methods mentioned above, such as scGen, can only infer the perturbed distribution of cell-types whose unperturbed distribution has been observed during training. Our method is more flexible in the sense that it requires access to the unperturbed distribution of gene expression only at inference time.

alternatively utilize the second moment and normalize this vector by the sample covariance matrix, e.g. $x_{DE} = \Sigma^{-1}(\overline{X}_{perturbed} - \overline{X}_{DMSO})$, assuming that $\Sigma_{perturbed} = \Sigma_{DMSO} = \Sigma$ in a manner similar to Linear Discriminant Analysis.

We formulate this as a set prediction problem (see section 3.1) where each single cell x_i in the set is transformed into its lower dimensional representation, $z_i = \phi(x_i)$. For every cell type, the set $\{z_i\}_{i=1}^{K}$ is then averaged to produce the latent vector $z_{avg} \in \mathbb{R}^d$ which summarizes the latent representation for that specific cell type. Finally, z_{avg} is propagated through the decoder to produce the predicted differential expression vector, $\widehat{x}_{DE} \in \mathbb{R}^n$.

The choice of defining the target variable to be the differential expression rather than the perturbed expression itself is motivated by the fact that perturbation effects are usually manifested as fluctuations around the unperturbed expression in a limited number of genes. Consequently, to predict the full perturbed transcriptome, the latent representation has to preserve most of the information on x_i and will mostly reflect similarities in the unperturbed transcriptome. By focusing on the differential expression, the latent representation is encouraged to learn a low-dimensional representation of the perturbation effect and discard all information not relevant for this task.

To train our model, we minimize the following loss function:

$$L = \frac{1}{C}\sum_{c=1}^{C}\left(w^T \cdot \left\|x_{(c)}^{DE} - \widehat{x}_{(c)}^{DE})\right\|^2 + \ \beta \cdot \sum_{j=1}^{K_c}\left[\|z_j\|^2\right]\right)$$

where C is the number of different cell types in the data and $x_{(c)}^{DE}$ and K_c are respectively the differential expression vector and the number of single cells in the unperturbed population of cell type c. β is the regularisation strength, controlling the flexibility of the latent representation to avoid model overfitting and w is a gene weighting vector that reflects the variability of each gene in response to the given perturbation. In our experiments we set $w_i = 1$ if gene i is in the top k most highly variable genes, otherwise $w_i = 0.01$. The choice of k is data-dependent with typical values in our experiments being between $50 - 200$. Alternatively, a continuous weighting which is proportional to the variability could also be explored.

To maintain constant set sizes and induce stochasticity during training, we use a bootstrapped sample of single cells for each cell type in each training step. For a certain cell type A, let X^A be the set of unperturbed single cells of type A, then the model takes in the bootstrapped sample $X_{bs} \subseteq X^A$ and is trained to reconstruct the target response vector $x_{DE} = \overline{X}_{perturbed} - \overline{X}_{bs}$, i.e., the difference between the mean expression of the perturbed cell population and the bootstrapped DMSO treated sample.

3.1 Single-cell RNA-seq and Set Prediction

In contrast to previously published work, we do not seek to predict the expression profiles of the entire population of perturbed single cells. Instead, we propose to focus on the general effect of the perturbation on each specific cell type. We observe that attempts in predicting the distribution of single cells may sometimes miss the average effect of the perturbation and is mostly accurate for genes with simple unimodal distributions. By simplifying the target variable and focusing on the average effect, we are able to distill the information that is associated with different forms of response and achieve more accurate predictions in unobserved cellular contexts.

As described above, we form our post-perturbation predictions from the single cell distribution of the observed unperturbed state in the following way: denote $X^A = \{x \mid x \textit{ in unperturbed population of cell type } A\}$, and let x_{DE}^A be the differential expression vector indicating the perturbation response of cell type A, then the function learned by the model $f : \mathbb{R}^{|X^A| \times n} \to \mathbb{R}^n$ should satisfy $f\left(X^A\right) = x_{DE}^A$.

Clearly, this is a set prediction problem and so f should be permutation-invariant with respect to the set X^A. Thus, we follow the structure from Zaheer et al. (2017) in which $f = \psi(\sum_{x \in X^A} \phi(x))$

and adapt it to the encoder-decoder structure such that $\phi(\cdot)$ is parameterized by the encoder network, the aggregation is done over the latent embeddings and $\psi(\cdot)$ is parameterized by the decoder network.

3.2 Model Explainability

A key aspect of the perturbation response prediction is the ability to explain the predicted outcome. This can be done by identifying the factors of variation in the input which contribute the most to the predicted outcome. If the unperturbed state representation consists of the gene expression profiles, this corresponds to identifying the genes whose up- or down-regulation has a strong effect on the model's response prediction, or equivalently, on the latent embedding of the cell types.

Providing explanations for predictions made by deep networks is not a trivial task. Further complexity is added by the high-dimensional space of the network outputs. To address these difficulties, we make a few simplifications. First, we consider the latent representation, Z as the output variable which we aim to explain. Z is trained to capture a low-dimensional representation of the response. Changes in Z reflect axes of variation in the predicted differential expression. Second, we further simplify the output variable by decomposing the latent representation into its principal components in order to identify the features that affect the outcome in each separate axis of variation.

To identify the input features that contribute to the variability in a given axis, we train an auxiliary interpretable regression model to recapitulate the encoder's outputs along that axis. More formally, let $X \in \mathbb{R}^{M \times n}$ be the set of all single cells in our training data where M is the number of cells and n is the number of genes in the transcriptome. If $Z = \phi(X) \in \mathbb{R}^{M \times d}$ denotes the latent embeddings of the single cells and $w_k \in \mathbb{R}^d$ denotes the k'th principal component of Z, then the auxiliary regression model is trained to approximate the function $X \rightarrow \phi(X) \cdot w_k$. It is important to note that by inspection of the data variability along a chosen principal component, one can associate a subset of input features with specific changes in the response. The proposed decomposition into principal components loses the global structure of Z. However, it serves as a pragmatic compromise between model explainability and a rich latent representation required for predictive capacity. Explainability is essential to ensure PerturbX will be widely adopted by biological and clinical domain experts.

4 Experiments

To demonstrate the performance of PerturbX, we used the pooled single cell RNA sequencing data generated by McFarland et al. (2020). A novel method, MIX-Seq, was developed by McFarland et al. (2020) to enable multiplexed single cell post-perturbational transcriptional pooling on a diverse range of 25 – 100 cell lines. Gene expression profiles were assessed at 6- or 24-hours post-perturbation. All experiments included Dimethyl Sulfoxide (DMSO) treatment as a negative control.

For detailed exploration, we excluded perturbations which induced homogenous responses across all cell lines since the predictive task in such cases is trivial. We also excluded perturbations for which the single-cell population in each condition was too small to reflect their underlying distribution. After filtering according to these exclusion criteria, the three most suitable perturbations to take forward for investigation were selected: idasanutlin, trametinib and dabrafenib. The transcriptional responses to trametinib and dabrafenib are highly correlated. Trametinib was selected as the representative example.

Idasanutlin and trametinib share similar high-level biological features. They are both small molecule inhibitors whose response is closely linked to established signaling pathways. Both inhibitors are expected to demonstrate heterogenous perturbational responses in different cellular contexts. Idasanutlin and trametinib are in clinical trial development or established in routine clinical practice. Therefore, they are representative of the challenges faced in translational drug development. In all conducted experiments, we benchmarked the predictive performance of PerturbX

to scGen since scGen allows for predicting the transcriptional outcomes in unobserved cell lines.

4.1 Predicting the Response to idasanutlin

Idasanutlin blocks the interaction of MDM2 with p53 (Ding et al., 2013). p53 is an established tumor suppressor (Baker et al., 1989). In its native context, MDM2 binds to p53 resulting in the enzymatic degradation of p53 (Momand et al., 1992). Idasanutlin binds directly to MDM2, blocks the MDM2-p53 interaction and thereby restores the tumour suppressive properties of p53 (Ding et al., 2013). *TP53* is the gene encoding the p53 protein. Many cancer cell lines have inactivating mutations in *TP53* which prevent them from responding to idasanutlin (Michaelis et al., 2011).

Figure 2 shows an UMAP embedding of the gene expression profiles colour-coded by (a) cell type and (b) idasanutlin and DMSO perturbations. The embeddings distinguish cell lines according to response to idasanutlin. Non-responsive cell lines show overlap between the DMSO and idasanutlin-treated clusters whilst responsive cell lines show cluster separation.

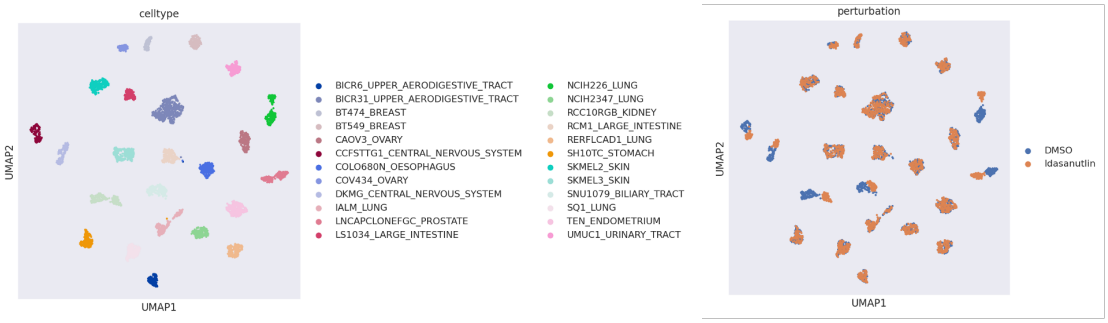

(a) (b)

Figure 2: The 2D UMAP embeddings enable visualisation of differences in transcriptional response to idasanutlin amongst a panel of 24 cell lines according to (a) cell type and (b) perturbation.

The biology of p53 is complex and has many intricacies which can be challenging to functionally deconvolve. We hypothesize PerturbX can integrate such intricacies into the trained latent representations of the unperturbed cell lines in order to establish the model's predictive capacity for the post-perturbation setting.

We hold out three different cell lines, SNU1079, DKMG and BICR31, and train the model on all other cell lines. SNU1079 and DKMG are *TP53* wild-type (WT) and show significant transcriptional responses to idasanutlin. In contrast, BICR31 is *TP53* mutated and shows no transcriptional effect. Figure 3 shows a scatterplot of the true vs predicted differential expression profiles of response to idasanutlin generated by PerturbX (top) and scGen (bottom) in the three unobserved cell lines. It demonstrates the model's ability to distinguish responsive and non-responsive cell lines according to *TP53* mutational status and to make accurate predictions of the transcriptional effect. scGen by design averages the transcriptional effect in the latent space over all observed cell lines. In contrast, PerturbX identifies the specific cell lines which are expected to respond most similarly to the predicted unobserved cell line.

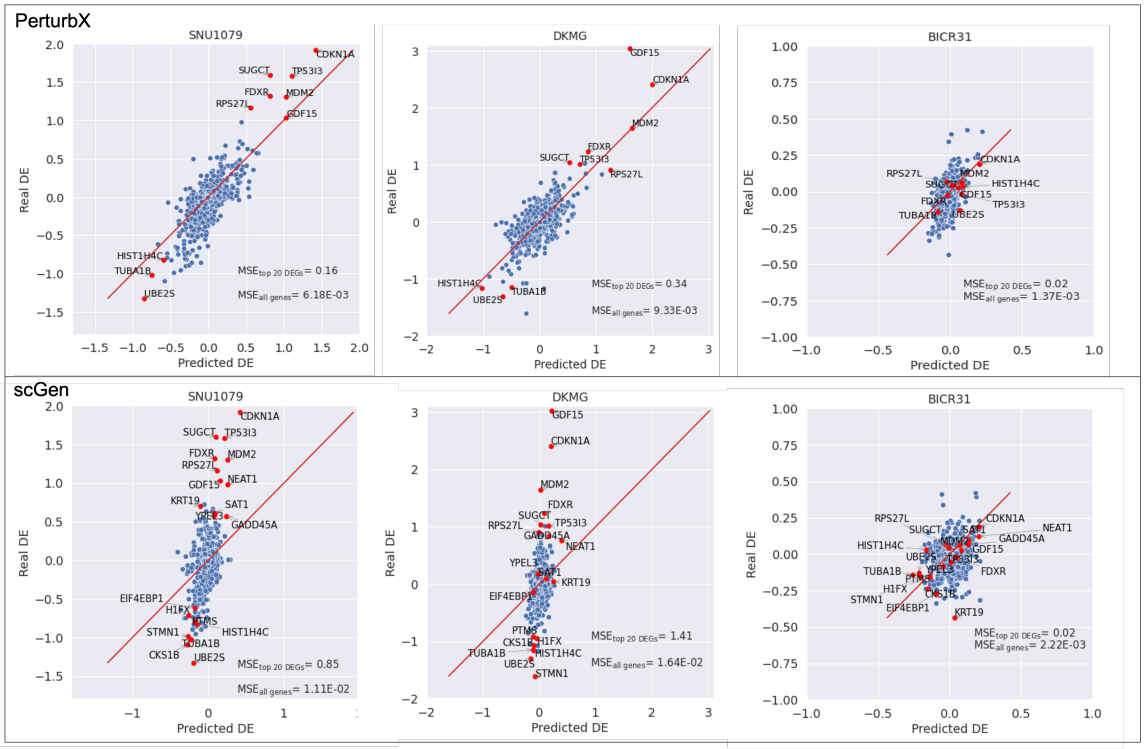

Figure 3: PerturbX successfully predicts gene expression in cell lines which do (SNU1079, DKMG) and do not (BICR31) respond to idasanutlin. Predictive performance of Perturbx (top) is compared to scGen (bottom). In all examples, PerturbX outperforms scGen as shown by consistently lower MSE. The top 20 differentially-expressed genes (DEGs) were chosen based on the response across all cell lines. The data points highlighted in red indicate the top 10 DEGs.

To investigate how PerturbX can capture a biologically representative latent embedding, we next examined the 2D UMAP plots of the trained latent embedding Z. The UMAP plots are shown in Figure 4a. It confirms that cell lines are mapped to different regions of the Z space according to *TP53* status. *TP53* mutant cell lines in the training and test data cluster together and in a region distinct compared to *TP53* wild-type cell lines. Therefore, the latent embeddings cluster according to important biological axes. Furthermore, the difference in the predicted differential expression for

SNU1079 and DKMG, together with the UMAP embedding, reveal the model's ability to capture an even more granular heterogeneity in response within the group of TP53 WT cell lines.

We next determined the contribution of the input features (gene list) on the predicted outcome, as described in section 3.2. We trained an auxiliary Random Forest Regressor (Breiman, 2001) to approximate the function $X \rightarrow \phi(X) \cdot w_1$, where w_1 is the first principal component of Z. Feature importance was assessed by the Shapley (SHAP) values of the auxiliary regression model. The list of top 10 genes is shown in Figure 4b. 7 out of the 10 genes are known *TP53* targets: *CCL2* (Tang et al., 2012), *RRAD* (Zhang et al., 2014), *RPS27L* (He & Sun, 2007), *TPM1* (Tang et al., 2018), *NUPR1* (Lopez et al., 2015) , *C1QL1* (Mei et al., 2008) and *MDM2* (Momand et al., 1992).

Furthermore NDUFA4L2, a component of the electron transport chain complex I subunit (Meng et al., 2019), and p53 share a common regulator HIF-1α, a transcription factor which accumulates under moderate and severe hypoxia (Meng et al., 2019; Chen et al., 2003).

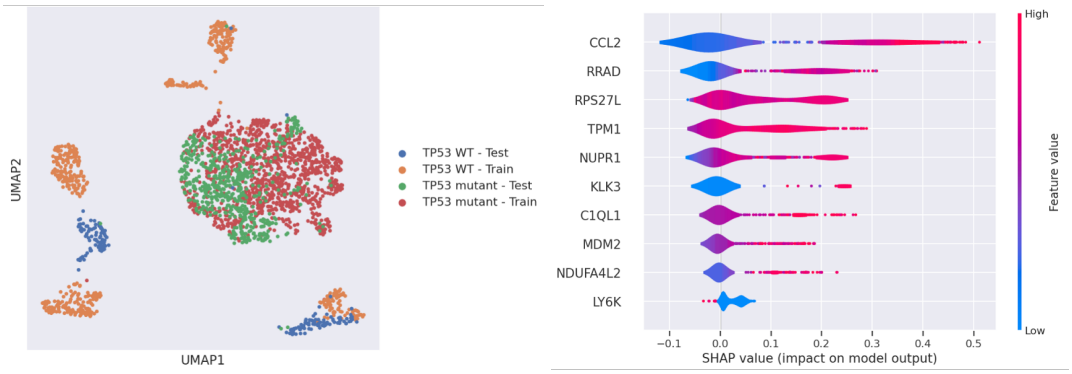

(a) (b)

Figure 4: (a) UMAP visualisation of the trained embedding Z confirms cell line clustering according to TP53 mutational status. (b) Genes contributing to the predicted expression upon perturbation of p53 signaling are enriched for known p53 targets.

4.2 Predicting the Response to Trametinib

We next evaluated PerturbX under compound treatment with trametinib on a larger RNA-seq dataset involving 97 cell lines. Trametinib is a reversible allosteric inhibitor of mitogen-activated extracellular signal regulated kinase 1 (MEK1) and MEK2 (Wright & McCormack, 2013). MEK1 and MEK2 are closely related kinases. They are a component of the Ras/Raf/MEK/ERK extracellular signaling pathway which is involved in driving cell proliferation and survival (Ohren et al., 2004).

BRAF mutations result in dependency on MEK signaling (Solit et al., 2006). Trametinib has been most extensively explored in the context of the BRAF V600E/V600K mutational background (Solit et al., 2006). At a clinical level, trametinib is approved for use in combination with dabrafenib in patients with advanced melanoma or non-small cell lung cancer with a BRAF V600 mutation (Robert et al., 2019; Planchard et al., 2017).

Compared to idasanutlin, the transcriptional response profiles to trametinib are significantly more diverse and heterogenous, with several genes showing pronounced effect in a small subset of cell lines. Such heterogeneity in response to trametinib treatment makes it significantly more challenging for the model to capture the exact differential expression profile in an unobserved cell line. In spite of such difficulty, PerturbX is still able to effectively capture the diverse response patterns and learn a meaningful representation of the different cell lines.

We adopted the same experimental pipeline as utilised in our investigations of idasanutlin. We selected three cell lines with diverse response profiles for our hold out set: UACC257, NCIH2347 and LNCAP clone FGC. UACC257, a melanoma cell line with a BRAF V600E mutation and NCIH2347, a non-small cell lung cancer cell line with a KRAS L19F mutation demonstrated a marked transcriptional response to trametinib. In contrast, LNCAP, a prostate adenocarcinoma cell line which is *BRAF* and *KRAS* WT, showed minimal response to trametinib.

The predictive performance of PerturbX compared to scGen in the hold out set is shown in Figure 5a. In spite of the limited number of cell lines sensitive to trametinib, PerturbX could reliably distinguish between the responsive vs non-responsive setting. The predicted differential response is in accordance with the published literature on MEK targeted therapies. In the trametinib-responsive cases, Perturbx outperformed scGen as confirmed by a lower MSE.

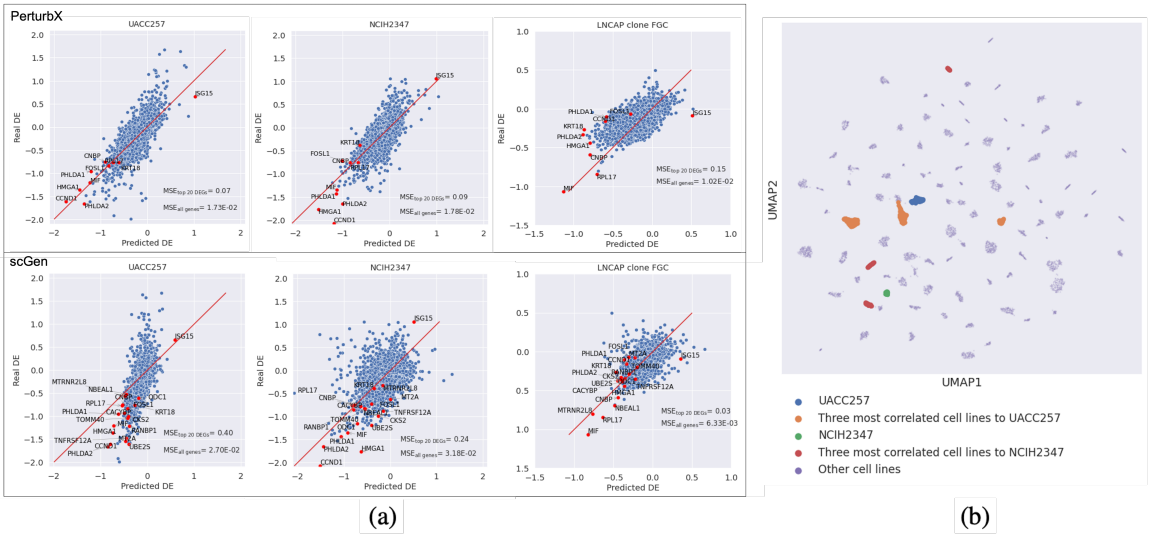

Figure 5: (a) PerturbX successfully predicts gene expression in cell lines which do (UACC257 and NCIH2347) and do not (LNCAP clone FGC) respond to trametinib. Perturbx (top) outperforms scGen (bottom) in the responsive cell lines. The LNCAP cell line is correctly predicted to have a weak response. However, the prediction is noisier compared to scGen. Top 20 DEGs were chosen based on the response across all cell lines. The data points highlighted in red indicate the top 10 DEGs. (b) PerturbX successfully clusters unobserved cell lines according to their response to trametinib.

The diversity in response to trametinib naturally lends itself to a more complex latent representation. Such a representation should be flexible enough to reflect the heterogenous responses to trametinib whilst also remaining tightly regularised to avoid model overfitting. As a qualitative evaluation of the learned representation, we explored the UMAP embedding of the latent space and focused on the trametinib responsive cell lines from the hold out set, UACC257 and NCIH2347. Cell lines from the training set were ranked according to their similarity in response to UACC257 and NCIH2347. Similarity was measured by the Pearson correlation coefficient of the differential expression vectors. We see that the unobserved cell lines are mapped close to two out of the three most similar cell lines in terms of the transcriptional effect, demonstrating that PerturbX can effectively capture the features in the unperturbed state which indicate the similarity in response.

5 Discussion

We present PerturbX, a deep learning model which can receive as input the unperturbed state representation of a given cellular context and predict the transcriptional effect in response to a certain perturbation. We demonstrate the ability of PerturbX to capture a low dimensional representation of the "response space" in which cell types are embedded close together if they exhibit similar transcriptional responses to a given perturbation and to generalise to unobserved cellular contexts. Furthermore, PerturbX identifies genetic features from the unperturbed state which most significantly contribute to the predicted output in a manner aligned with a priori domain-specific knowledge. Significantly, this combined ability showcases the unique potential of PerturbX to inform novel biomarker discovery whilst remaining congruous to the relevant underlying biology.

The architecture presented in this paper allows expansion of the representation of the unperturbed state from gene expression profiles to additional modalities, such as mutation status or CNV. However, future work is needed to evaluate the contribution of such multimodal learning. We hypothesize that it may improve response predictions by maximizing the information content of the learned representation, potentially by integrating information from orthogonal biological axes which are often inter-dependent and exhibit complex correlation relationships. Multimodal learning may also address some of the challenges arising from technical biases specific to the unimodal space. Expansion into multimodal learning will likely require larger datasets for model training and potential model adaptations.

In summary, the future opportunities and challenges offered by PerturbX sits at the evolving interface between ML and biology. It signals the future of ML-guided experimental design which will bridge across traditional interdisciplinary boundaries.

References

S J Baker, E R Fearon, J M Nigro, S R Hamilton, A C Preisinger, J M Jessup, P vanTuinen, D H Ledbetter, D F Barker, Y Nakamura, R White, and B Vogelstein. Chromosome 17 deletions and p53 gene mutations in colorectal carcinomas. *Science*, 244(4901):217–221, Apr 1989. ISSN 0036-8075 (Print); 0036-8075 (Linking). doi: 10.1126/science.2649981.

Leo Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001. doi: 10.1023/A:1010933404324. URL `https://doi.org/10.1023/A:1010933404324`.

Consortium Cancer Cell Line Encyclopedia and Consortium Genomics of Drug Sensitivity in Cancer. Pharmacogenomic agreement between two cancer cell line data sets. *Nature*, 528(7580):84–87, Dec 2015. ISSN 1476-4687 (Electronic); 0028-0836 (Print); 0028-0836 (Linking). doi: 10.1038/nature15736.

Delin Chen, Muyang Li, Jianyuan Luo, and Wei Gu. Direct interactions between hif-1 alpha and mdm2 modulate p53 function. *J Biol Chem*, 278(16):13595–13598, Apr 2003. ISSN 0021-9258 (Print); 0021-9258 (Linking). doi: 10.1074/jbc.C200694200.

Qingjie Ding, Zhuming Zhang, Jin-Jun Liu, Nan Jiang, Jing Zhang, Tina M Ross, Xin-Jie Chu, David Bartkovitz, Frank Podlaski, Cheryl Janson, Christian Tovar, Zoran M Filipovic, Brian Higgins, Kelli Glenn, Kathryn Packman, Lyubomir T Vassilev, and Bradford Graves. Discovery of rg7388, a potent and selective p53-mdm2 inhibitor in clinical development. *J Med Chem*, 56(14): 5979–5983, Jul 2013. ISSN 1520-4804 (Electronic); 0022-2623 (Linking). doi: 10.1021/jm400487c.

H He and Y Sun. Ribosomal protein s27l is a direct p53 target that regulates apoptosis. *Oncogene*, 26(19):2707–2716, Apr 2007. ISSN 0950-9232 (Print); 0950-9232 (Linking). doi: 10.1038/sj.onc.1210073.

Leon Hetzel, Simon Böhm, Niki Kilbertus, Stephan Günnemann, Mohammad Lotfollahi, and Fabian Theis. Predicting Cellular Responses to Novel Drug Perturbations at a Single-Cell Resolution. *arXiv e-prints*, art. arXiv:2204.13545, April 2022. doi: 10.48550/arXiv.2204.13545.

Sébastien Lachapelle and Simon Lacoste-Julien. Partial Disentanglement via Mechanism Sparsity. *arXiv e-prints*, art. arXiv:2207.07732, July 2022. doi: 10.48550/arXiv.2207.07732.

Maria Belen Lopez, Maria Noé Garcia, Daniel Grasso, Jennifer Bintz, Maria Inés Molejon, Gabriel Velez, Gwen Lomberk, Jose Luis Neira, Raul Urrutia, and Juan Iovanna. Functional characterization of nupr1l, a novel p53-regulated isoform of the high-mobility group (hmg)-related protumoral protein nupr1. *J Cell Physiol*, 230(12):2936–2950, Dec 2015. ISSN 1097-4652 (Electronic); 0021-9541 (Print); 0021-9541 (Linking). doi: 10.1002/jcp.25022.

Romain Lopez, Nataša Tagasovska, Stephen Ra, Kyunghyn Cho, Jonathan K. Pritchard, and Aviv Regev. Learning Causal Representations of Single Cells via Sparse Mechanism Shift Modeling. *arXiv e-prints*, art. arXiv:2211.03553, November 2022. doi: 10.48550/arXiv.2211.03553.

Mohammad Lotfollahi, Mohsen Naghipourfar, Fabian J. Theis, and F. Alexander Wolf. Conditional out-of-sample generation for unpaired data using trVAE. *arXiv e-prints*, art. arXiv:1910.01791, October 2019. doi: 10.48550/arXiv.1910.01791.

Mohammad Lotfollahi, F Alexander Wolf, and Fabian J Theis. scgen predicts single-cell perturbation responses. *Nat Methods*, 16(8):715–721, Aug 2019. ISSN 1548-7105 (Electronic); 1548-7091 (Linking). doi: 10.1038/s41592-019-0494-8.

Mohammad Lotfollahi, Anna Klimovskaia Susmelj, Carlo De Donno, Yuge Ji, Ignacio L. Ibarra, F. Alexander Wolf, Nafissa Yakubova, Fabian J. Theis, and David Lopez-Paz. Learning interpretable cellular responses to complex perturbations in high-throughput screens. *bioRxiv*, 2021. doi: 10.1101/2021.04.14.439903. URL `https://www.biorxiv.org/content/early/2021/05/18/2021.04.14.439903`.

James M McFarland, Brenton R Paolella, Allison Warren, Kathryn Geiger-Schuller, Tsukasa Shibue, Michael Rothberg, Olena Kuksenko, William N Colgan, Andrew Jones, Emily Chambers, Danielle Dionne, Samantha Bender, Brian M Wolpin, Mahmoud Ghandi, Itay Tirosh, Orit Rozenblatt-Rosen, Jennifer A Roth, Todd R Golub, Aviv Regev, Andrew J Aguirre, Francisca Vazquez, and Aviad Tsherniak. Multiplexed single-cell transcriptional response profiling to define cancer vulnerabilities and therapeutic mechanism of action. *Nat Commun*, 11(1):4296, Aug 2020. ISSN 2041-1723 (Electronic); 2041-1723 (Linking). doi: 10.1038/s41467-020-17440-w.

Jie Mei, Qi-Ya Zhang, Zhi Li, Shuo Lin, and Jian-Fang Gui. C1q-like inhibits p53-mediated apoptosis and controls normal hematopoiesis during zebrafish embryogenesis. *Dev Biol*, 319(2):273–284, Jul 2008. ISSN 1095-564X (Electronic); 0012-1606 (Linking). doi: 10.1016/j.ydbio.2008.04.022.

Lifei Meng, Xuhui Yang, Xiao Xie, and Mingsong Wang. Mitochondrial ndufa4l2 protein promotes the vitality of lung cancer cells by repressing oxidative stress. *Thorac Cancer*, 10(4):676–685, Apr 2019. ISSN 1759-7714 (Electronic); 1759-7706 (Print); 1759-7706 (Linking). doi: 10.1111/1759-7714.12984.

M Michaelis, F Rothweiler, S Barth, J Cinatl, M van Rikxoort, N Löschmann, Y Voges, R Breitling, A von Deimling, F Rödel, K Weber, B Fehse, E Mack, T Stiewe, H W Doerr, D Speidel, and J Jr Cinatl. Adaptation of cancer cells from different entities to the mdm2 inhibitor nutlin-3 results in the emergence of p53-mutated multi-drug-resistant cancer cells. *Cell Death Dis*, 2(12):e243, Dec 2011. ISSN 2041-4889 (Electronic). doi: 10.1038/cddis.2011.129.

J Momand, G P Zambetti, D C Olson, D George, and A J Levine. The mdm-2 oncogene product forms a complex with the p53 protein and inhibits p53-mediated transactivation. *Cell*, 69(7):1237–1245, Jun 1992. ISSN 0092-8674 (Print); 0092-8674 (Linking). doi: 10.1016/0092-8674(92)90644-r.

Jeffrey F Ohren, Huifen Chen, Alexander Pavlovsky, Christopher Whitehead, Erli Zhang, Peter Kuffa, Chunhong Yan, Patrick McConnell, Cindy Spessard, Craig Banotai, W Thomas Mueller, Amy Delaney, Charles Omer, Judith Sebolt-Leopold, David T Dudley, Iris K Leung, Cathlin Flamme, Joseph Warmus, Michael Kaufman, Stephen Barrett, Haile Tecle, and Charles A Hasemann. Structures of human map kinase kinase 1 (mek1) and mek2 describe novel noncompetitive kinase inhibition. *Nature Structural & Molecular Biology*, 11(12):1192–1197, 2004. doi: 10.1038/nsmb859. URL `https://doi.org/10.1038/nsmb859`.

David Planchard, Egbert F Smit, Harry J M Groen, Julien Mazieres, Benjamin Besse, Åslaug Helland, Vanessa Giannone, Anthony M Jr D'Amelio, Pingkuan Zhang, Bijoyesh Mookerjee, and Bruce E Johnson. Dabrafenib plus trametinib in patients with previously untreated braf(v600e)-mutant metastatic non-small-cell lung cancer: an open-label, phase 2 trial. *Lancet Oncol*, 18 (10):1307–1316, Oct 2017. ISSN 1474-5488 (Electronic); 1470-2045 (Linking). doi: 10.1016/S1470-2045(17)30679-4.

Caroline Robert, Jean J. Grob, Daniil Stroyakovskiy, Boguslawa Karaszewska, Axel Hauschild, Evgeny Levchenko, Vanna Chiarion Sileni, Jacob Schachter, Claus Garbe, Igor Bondarenko, Helen Gogas, Mario Mandalá, John B.A.G. Haanen, Celeste Lebbé, Andrzej Mackiewicz, Piotr Rutkowski, Paul D. Nathan, Antoni Ribas, Michael A. Davies, Keith T. Flaherty, Paul Burgess, Monique Tan, Eduard Gasal, Maurizio Voi, Dirk Schadendorf, and Georgina V. Long. Five-year outcomes with dabrafenib plus trametinib in metastatic melanoma. *New England Journal of Medicine*, 381(7):626–636, 2019. doi: 10.1056/NEJMoa1904059. URL `https://doi.org/10.1056/NEJMoa1904059`. PMID: 31166680.

Orit Rozenblatt-Rosen, Michael J T Stubbington, Aviv Regev, and Sarah A Teichmann. The human cell atlas: from vision to reality. *Nature*, 550(7677):451–453, Oct 2017. ISSN 1476-4687 (Electronic); 0028-0836 (Linking). doi: 10.1038/550451a.

David B Solit, Levi A Garraway, Christine A Pratilas, Ayana Sawai, Gad Getz, Andrea Basso, Qing Ye, Jose M Lobo, Yuhong She, Iman Osman, Todd R Golub, Judith Sebolt-Leopold, William R Sellers, and Neal Rosen. Braf mutation predicts sensitivity to mek inhibition. *Nature*, 439(7074): 358–362, Jan 2006. ISSN 1476-4687 (Electronic); 0028-0836 (Print); 0028-0836 (Linking). doi: 10.1038/nature04304.

Chao Tang, Jin Wang, Qi Wei, Yi-Peng Du, He-Ping Qiu, Chao Yang, and Yu-Chuan Hou. Tropomyosin-1 promotes cancer cell apoptosis via the p53-mediated mitochondrial pathway in renal cell carcinoma. *Oncol Lett*, 15(5):7060–7068, May 2018. ISSN 1792-1074 (Print); 1792-1082 (Electronic); 1792-1074 (Linking). doi: 10.3892/ol.2018.8204.

X Tang, M Asano, A O'Reilly, A Farquhar, Y Yang, and S Amar. p53 is an important regulator of ccl2 gene expression. *Curr Mol Med*, 12(8):929–943, Sep 2012. ISSN 1875-5666 (Electronic); 1566-5240 (Print); 1566-5240 (Linking). doi: 10.2174/156652412802480844.

John N Weinstein, Eric A Collisson, Gordon B Mills, Kenna R Mills Shaw, Brad A Ozenberger, Kyle Ellrott, Ilya Shmulevich, Chris Sander, and Joshua M Stuart. The cancer genome atlas pan-cancer analysis project. *Nat Genet*, 45(10):1113–1120, Oct 2013. ISSN 1546-1718 (Electronic); 1061-4036 (Print); 1061-4036 (Linking). doi: 10.1038/ng.2764.

Cameron J. M. Wright and Paul L. McCormack. Trametinib: First global approval. *Drugs*, 73(11):1245–1254, 2013. doi: 10.1007/s40265-013-0096-1. URL `https://doi.org/10.1007/s40265-013-0096-1`.

Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Ruslan Salakhutdinov, and Alexander Smola. Deep Sets. *arXiv e-prints*, art. arXiv:1703.06114, March 2017. doi: 10.48550/arXiv.1703.06114.

Cen Zhang, Juan Liu, Rui Wu, Yingjian Liang, Meihua Lin, Jia Liu, Chang S Chan, Wenwei Hu, and Zhaohui Feng. Tumor suppressor p53 negatively regulates glycolysis stimulated by hypoxia through its target rrad. *Oncotarget*, 5(14):5535–5546, Jul 2014. ISSN 1949-2553 (Electronic); 1949-2553 (Linking). doi: 10.18632/oncotarget.2137.