



Enhancing drug and cell line representations via contrastive learning for improved anti-cancer drug prioritization



Patrick J. Lawrence¹, Benjamin Burns² & Xia Ning^{1,2,3} ✉

Due to cancer's complex nature and variable response to therapy, precision oncology informed by omics sequence analysis has become the current standard of care. However, the amount of data produced for each patient makes it difficult to quickly identify the best treatment regimen. Moreover, limited data availability has hindered computational methods' abilities to learn patterns associated with effective drug-cell line pairs. In this work, we propose the use of contrastive learning to improve learned drug and cell line representations by preserving relationship structures associated with drug mechanisms of action and cell line cancer types. In addition to achieving enhanced performance relative to a state-of-the-art method, we find that classifiers using our learned representations exhibit a more balanced reliance on drug- and cell line-derived features when making predictions. This facilitates more personalized drug prioritizations that are informed by signals related to drug resistance.

Cancer, the leading cause of death worldwide, remains a challenge to treat due to its complex nature and variable response to therapy, even among patients with the same cancer type. Omics, which describes the collective and comprehensive analysis of biomolecular data, is being leveraged to conduct precision oncology. Transcriptomics, a subfield of omics that studies gene expression by quantifying relative levels of RNA molecules, has been an especially valuable tool. Oncologists use transcriptomics—comparing normal and tumor cells—to identify changes in gene expression. These alterations are used to pinpoint the molecular process(es) driving tumorigenesis in each patient, allowing clinicians to develop personalized treatment recommendations. However, because RNA sequencing (RNA-seq) measures the expression of more than 20,000 genes, manually evaluating each patient's data to determine the best treatment is neither scalable nor pragmatic.

Instead, machine learning models have been applied to omics data to predict cellular sensitivities to drug candidates. RefDNN¹, one state-of-the-art (SOTA) method, represents drugs via their structural similarity to a set of reference drugs. It then leverages these reference drugs to produce cell line

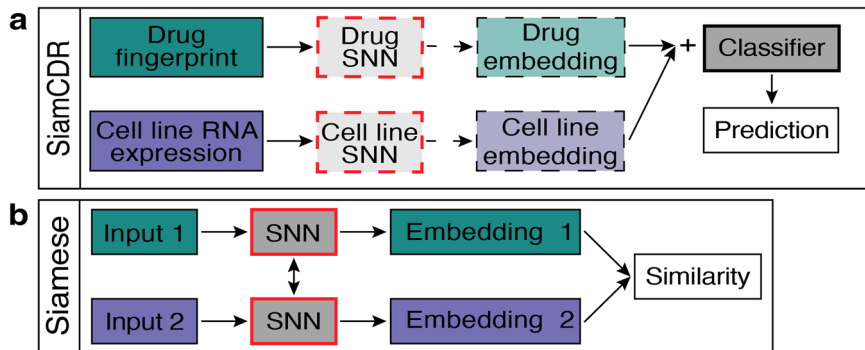
representations: each dimension is the output of an Elastic Net² model trained on transcriptomic data to predict a cell line's sensitivity to a distinct reference drug. During evaluation, RefDNN applies a dense neural network (DNN) to the Hadamard product of drug and cell line representations to predict cancer drug response (CDR). However, RefDNN is limited by data quality: it requires complete CDR data for all reference drug-cell line pairs during training. Other recent SOTA methods, such as SubCDR³, have explored the use of expert annotations genes, explicitly denoting the role each of the genes in their subset played in the tumorigenesis of each cancer type in their data set. This method also relies on the use of CDR values as 'side information'. While it improves on RefDNN, in that it can be trained on incomplete CDR data, SubCDR is only able to accurately predict CDR response for new combinations of cell lines and drugs seen during training. Furthermore, manually labeling genes by their role is both time-consuming and may introduce bias into the model.

DeepDSC⁴, another SOTA method, also leverages transcriptomic data to produce cell line representations. It does so via the latent embeddings of a pretrained autoencoder (AE). By encoding cell lines' transcriptomic profiles

¹Biomedical Informatics Department, The Ohio State University, 1800 Cannon Drive, Lincoln Tower 250, Columbus, OH 43210, USA. ²Computer Science and Engineering Department, The Ohio State University, 2015 Neil Avenue, Columbus, OH 43210, USA. ³Translational Data Analytics Institute, The Ohio State University, 1760 Neil Avenue, Columbus, OH 43210, USA.

✉ e-mail: ning.104@osu.edu

Fig. 1 | Model architectures. Depicts architectures for components proposed by this work: **a** Siamese neural network and **b** SiamCDR. For both, boxes with bold borders and a grey face denote trained components. The input pair in **a** is either a pair of drugs or pair of cell lines depending on if drug or cell line encoder is being trained. Dashed lines and box borders in **b** indicate optional components by variation. See respective sections in Methods for complete details.



into a lower-dimensional space, the AE captures key information while mitigating the risk of overfitting, which commonly occurs when training deep learning models on limited data. In DeepDSC, drugs are represented by Morgan fingerprints⁵. A dense neural network (DNN) is then applied to drug-cell line pairs to predict CDR. DeepDSC is more robust than RefDNN as it can be trained on incomplete data. However, its use of generic fingerprints may still hinder performance as they are not customized to predict CDR. Similarly, the cell line features derived from DeepDSC’s AE are optimized for transcriptomic profile reconstruction, meaning they, too, may not be relevant for CDR prediction.

As such, two major challenges in achieving precise CDR predictions are the creation of problem-specific embeddings and the integration of drug-cell line pairs, which are both complicated by the quality and availability of data. To address these issues, we propose a framework, denoted as SiamCDR, that leverages contrastive loss within a Siamese neural network⁶ (SNN) to enhance the expressiveness of drug and cell line representations, thereby improving Cancer Drug Response predictions (Fig. 1a). Specifically, our model learns to project drugs and cell lines to embedding spaces that encode the similarities of gene targets for drugs and cancer type for cell lines, respectively. This is guided by the intuition that drugs with similar targets will have similar effects. Additionally, drug efficacies among cells of the same cancer type should be more similar than drug efficacies among cells of different cancers. The benefit of using SNNs here is their, and other few-shot learning frameworks’, ability to learn such similarity relationships from a limited number of training instances for each group^{6–8}. Specifically, drugs are grouped by their mechanism of action (MOA); cell lines are grouped by their cancer type. Moreover, contrastive learning will ensure our method preserves similarity relationships tailored to predicting CDR.

Our experiments show that SiamCDR produces higher-quality and more personalized drug prioritizations than DeepDSC. In fact, a network analysis of genes whose expression is significantly correlated with SiamCDR’s docetaxel prioritization in breast cancer identified enriched pathways known to induce docetaxel-resistance. This suggests the SiamCDR can learn to modulate its recommendations from transcriptomic signals associated with drug efficacy. Finally, using SiamCDR, we identify multiple repurposing candidates for difficult-to-treat cancers.

Results

Evaluating model performance

High performing models will prioritize the most effective drugs for each cell line above ineffective drugs. To that end, we assess model performance using average $P_{cell}@k$ ($k = 1, 2, 3, 4, 5, 10$) and average $P_{cancer}@k$ ($k = 1, 2, 3, 4, 5$) defined by Eqs. (2) and (3), respectively. These measure the average proportion of highly effective drugs among a model’s top- k prioritized drugs. Specific details regarding how each metric is calculated can be found in the ‘Evaluation metrics’ subsection of Methods. We compare the performance of the top-performing hyperparameter set for each end-classifier evaluated within the context of our proposed architecture: SiamCDR_{LR}, SiamCDR_{RF}, and SiamCDR_{DNN}, against the DeepDSC baseline (see Methods section for details about each model). Table 1a and b

report the average $P_{cell}@k$ for cell lines with trained-on and novel cancers, respectively; Table 2a and b report the average $P_{cancer}@k$ for cell lines with trained-on and novel cancers, respectively. Tables 1 and 2 also report Bonferroni-corrected significance levels with respect to differences in performance relative to DeepDSC. Each SiamCDR model achieves improvements over DeepDSC for all reported metrics in Tables 1 and 2. However, the Bonferroni-corrected significance level of these improvements varies.

Table 1 | Evaluating model performance with respect to $P_{cell}@k$

a Trained-on cancer types						
Model	$P_{cell}@1$	$P_{cell}@2$	$P_{cell}@3$	$P_{cell}@4$	$P_{cell}@5$	$P_{cell}@10$
DeepDSC	0.7059	0.6000	0.5608	0.5392	0.5294	0.5031
SiamCDR _{LR}	0.9519	*0.9368	*0.9194	**0.8829	**0.8641	***0.8007
SiamCDR _{RF}	0.9647	*0.9196	*0.8915	**0.8657	**0.8486	***0.8185
SiamCDR _{DNN}	0.9490	*0.9137	*0.8941	*0.8598	**0.8353	***0.8405
b Novel cancer types						
Model	$P_{cell}@1$	$P_{cell}@2$	$P_{cell}@3$	$P_{cell}@4$	$P_{cell}@5$	$P_{cell}@10$
DeepDSC	0.7938	0.6846	0.6164	0.6023	0.5742	0.5463
SiamCDR _{LR}	*0.9508	0.9123	*0.9046	*0.8738	**0.8622	***0.8059
SiamCDR _{RF}	**0.9569	*0.9123	0.8892	*0.8554	*0.8363	***0.8090
SiamCDR _{DNN}	*0.9723	*0.9354	0.9036	*0.8792	**0.8591	***0.8239

Significance levels ($\alpha \leq 0.1, 0.05, 0.01$; indicated by *, **, ***, respectively) are determined from p -values obtained via Bonferroni correction ($n = 4$) of two-tailed t -tests comparing the performances of SiamCDR_{LR}, SiamCDR_{RF}, and SiamCDR_{DNN} against DeepDSC.

Table 2 | Evaluating model performance with respect to $P_{cancer}@k$

a Trained-on cancer types					
Model	$P_{cancer}@1$	$P_{cancer}@2$	$P_{cancer}@3$	$P_{cancer}@4$	$P_{cancer}@5$
DeepDSC	0.7182	0.6083	0.5640	0.5397	0.5347
SiamCDR _{LR}	0.9571	*0.9435	*0.9179	**0.8820	**0.8597
SiamCDR _{RF}	*0.9712	*0.9313	*0.8968	**0.8695	**0.8485
SiamCDR _{DNN}	0.9503	*0.9225	*0.9012	*0.8674	**0.8436
b Novel cancer types					
Model	$P_{cancer}@1$	$P_{cancer}@2$	$P_{cancer}@3$	$P_{cancer}@4$	$P_{cancer}@5$
DeepDSC	0.8157	0.7155	0.6447	0.6174	0.5793
SiamCDR _{LR}	*0.9621	0.9283	*0.9267	*0.8821	**0.8504
SiamCDR _{RF}	*0.9616	*0.9311	0.8840	0.8585	*0.8367
SiamCDR _{DNN}	*0.9794	*0.9547	0.9237	*0.8817	**0.8591

Significance levels ($\alpha \leq 0.1, 0.05, 0.01$; indicated by *, **, ***, respectively) are determined from p -values obtained via Bonferroni correction ($n = 4$) of two-tailed t -tests comparing the performances of SiamCDR_{LR}, SiamCDR_{RF}, and SiamCDR_{DNN} against DeepDSC. Precision@ k for each cancer is presented in Supplementary Table 6

Note, we do not observe significant differences in the performance of the three SiamCDR models. This indicates all three are equally proficient at recommending effective drugs to target the cells of trained-on and novel cancers.

In Table 1a, we observe no significant difference in performance between the models for $P_{\text{cell}}@1$. However, for all other values of k , we observe significant improvements in performance for all three SiamCDR models compared to DeepDSC. Notably, the significance level increases with k for all SiamCDR models. This indicates that, for trained-on cancers, SiamCDR models prioritize a greater number of effective drugs at the very top compared to DeepDSC, which suggests that SiamCDR is more robust than DeepDSC with respect to quality of its prioritizations. Identical trends are observed when performance is generalized to cancer type (P_{cancer}) in Table 2a, apart from SiamCDR_{RF} gaining significance at $k = 1$.

Table 1b reports significant improvements in $P_{\text{cell}}@1$ compared to DeepDSC for all SiamCDR models. However, SiamCDR_{LR} loses its significance for $k = 2$, while both SiamCDR_{RF} and SiamCDR_{DNN} do not achieve significant improvement for $k = 3$. When generalized to cancer type (P_{cancer}) in Table 2b, we observe substantial jumps in performance for DeepDSC, SiamCDR_{LR}, and SiamCDR_{DNN}. Conversely, SiamCDR_{RF} demonstrates more stable performance, which combined with the improved relative performance of DeepDSC, yields a loss of significance at $k = 4$. This suggests DeepDSC, SiamCDR_{LR}, and SiamCDR_{DNN} may all be more sensitive to cell line outliers than SiamCDR_{RF}. One potential explanation is these models have learned to identify broad-spectrum (or cancer-specific) anti-cancer drugs. In this case, presenting cell lines with resistance to these drugs would yield reduced performance. However, averaging P_{cell} by cancer type would reduce the impact each outlier has, thereby raising performance. If this were the case, that would indicate that SiamCDR_{RF} has a greater capacity for producing tailored cell line predictions compared to the other models.

To test whether DeepDSC, SiamCDR_{LR}, and SiamCDR_{DNN} models may, in fact, favor broad-spectrum anti-cancer drugs more compared to SiamCDR_{RF}, we visualize prediction trends for each drug-cell line pair. In Fig. 2, we plot continuous effective scores (y-axis)—denoted as CES and defined by Eq. (1)—against CDRs predicted by a) DeepDSC, b) SiamCDR_{RF}, c) SiamCDR_{LR}, and d) SiamCDR_{DNN} (x-axis). The binarization threshold for CDR labels is illustrated by horizontal, dashed lines. Pairs above this line are highly effective (top-10% with respect to CDR). To specifically assess drugs commonly predicted as effective, we filter the drugs to include only those prioritized among the top-5 for at least three cell lines by all four models. Each of the six identified drugs is highlighted with distinct shapes and colors in Fig. 2. Note that to qualitatively examine the relationship between the predicted score and CES, the scale of the x-axis in Fig. 2a has been adjusted; the range of predictions produced by DeepDSC (0.068 to 0.259) is much smaller than the range of the SiamCDR models' predictions (0.0 to 1.0).

We observe a positive association between CES and model predictions for each model. Considering only the six highlighted drugs, DeepDSC and SiamCDR_{RF} (Fig. 2a and b) produce scores with large drug-wise variance, suggesting their predictions are more informed by cell lines than either SiamCDR_{LR} or SiamCDR_{DNN}. Conversely, in Fig. 2c and d, we observe that SiamCDR_{LR} and SiamCDR_{DNN}'s predictions have small by-drug variance. This indicates that both have learned to identify broad-spectrum anti-cancer drugs. With respect to SiamCDR_{LR}, it is likely that the relationship(s) that determine(s) cell lines' individual drug responses are too complex to capture via logistic regression. Additionally, it is probable that there is insufficient data to train all SiamCDR_{DNN}'s parameters. This is supported by SiamCDR_{DNN}'s average training and validation loss curves, which suggest the model quickly overfits the training data (Supplementary Figure 1). As a result, both models have learned to identify broad-spectrum anti-cancer drugs as a way of minimizing loss during training. SiamCDR_{RF}'s ensemble approach allows decision trees to capture nonlinear interactions more effectively than LRs while being less prone to overfitting than DNNs, allowing it to provide unique predictions for each drug-cell line pair. This is supported by the average variance in SiamCDR_{LR}'s, SiamCDR_{RF}'s, and

SiamCDR_{DNN}'s predicted scores for any drug with a predicted score above 0.5 for at least one cell line. SiamCDR_{RF} achieves 5.11- and 2.72-times greater variance in its predictions than SiamCDR_{LR} and SiamCDR_{DNN}, respectively, indicating it may be a more suitable candidate for precision medicine applications.

Identifying drug- and cell line-derived feature importance to model predictions

We measure model feature importance (see Methods) to discern whether SiamCDR_{RF}'s tailored drug prioritizations result from a more balanced influence of drug- and cell line-derived features compared to DeepDSC, SiamCDR_{LR} or SiamCDR_{DNN}. In Fig. 3, we rank features by their importance to the predictions of a) DeepDSC, b) SiamCDR_{RF}, c) SiamCDR_{LR}, or d) SiamCDR_{DNN}. For each subplot, a maximum of the top-100, nonzero features are depicted. Relative feature importance is conveyed by bar height, and feature source—drug or cell line—is represented by color. The average importance for each source is denoted by a horizontal line of the source's respective color.

In Fig. 3a, we observe that DeepDSC's predictions are more heavily influenced by cell lines than drugs with 90% of the top-10 features being cell line-derived. Conversely, we find SiamCDR_{LR} and SiamCDR_{DNN} almost exclusively prefer drug-derived features (Fig. 3c and d, respectively). None of SiamCDR_{LR}'s top-100 features are derived from cell lines and SiamCDR_{DNN} preferred a single cell line-derived feature among its top-100 (46th). The lack of cell line influence yields non-personalized drug prioritization and explains the small by-drug variance in SiamCDR_{LR}'s and SiamCDR_{DNN}'s predictions observed in Fig. 2c and d.

On the other hand, in Fig. 2b, we observe a balanced influence of both drugs and cell lines on SiamCDR_{RF}'s predictions: among the top-10 features, four and six are drug- and cell line-derived, respectively. In addition, we observe that the average relative feature importance for drugs and for cell lines is more similar in magnitude for SiamCDR_{RF} than for any of the other models. This balanced integration of drug and cell line information when making predictions highlights SiamCDR_{RF}'s practical utility over both SiamCDR_{LR} and SiamCDR_{DNN} for precision medicine. Therefore, we consider only SiamCDR_{RF} for the remaining experiments.

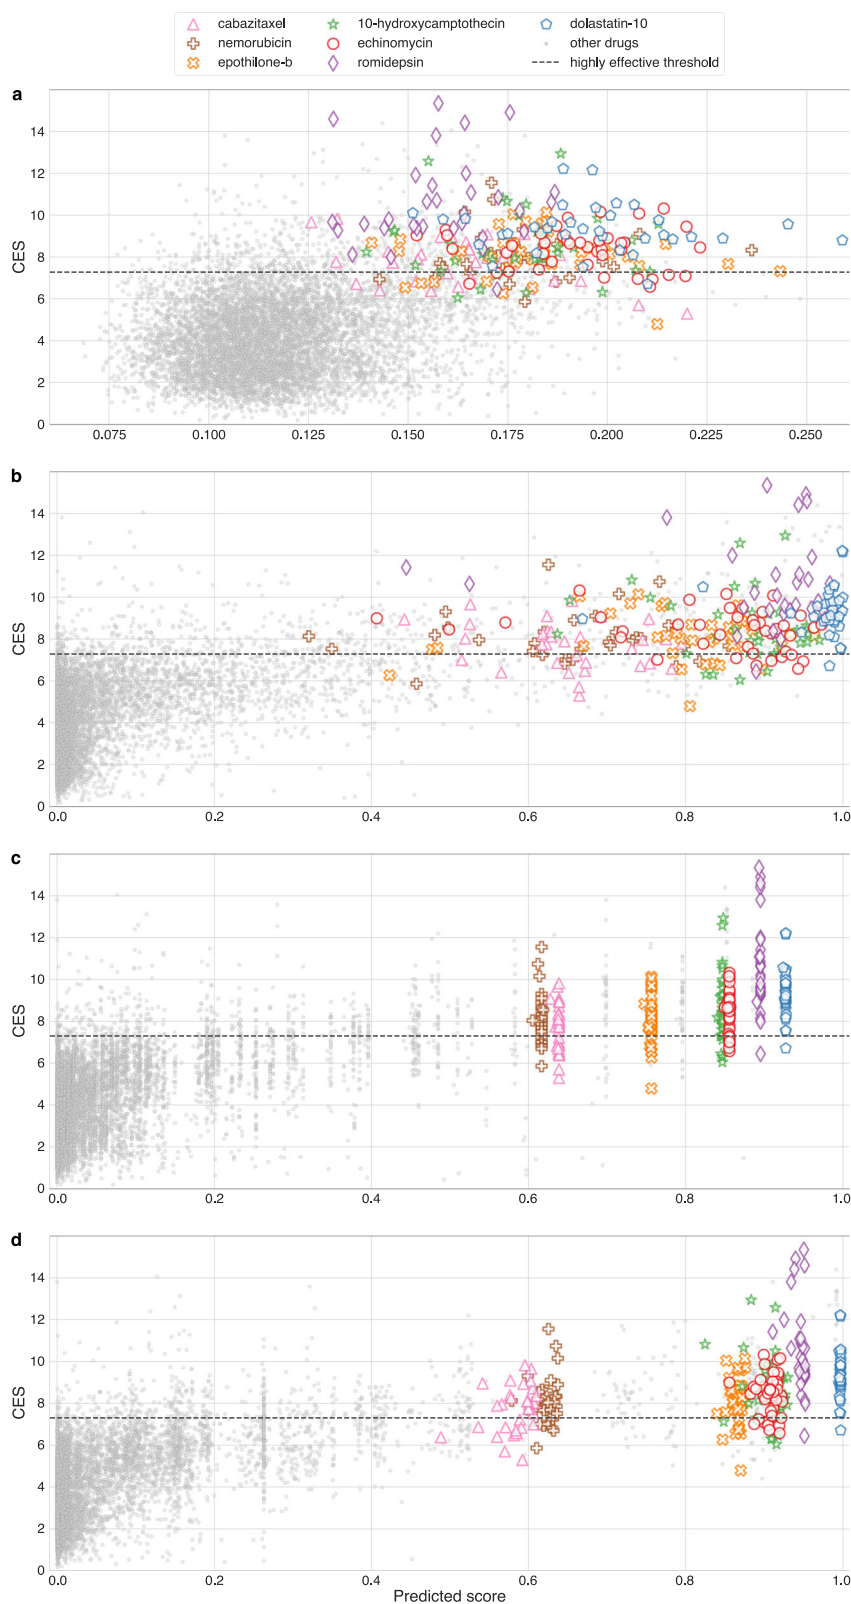
Comparing the expressiveness of DeepDSC's and SiamCDR_{RF}'s cell line representations

Both DeepDSC and SiamCDR_{RF} conduct representation learning to produce cell line embeddings (e_{ac} and e_{c} , respectively). However, e_{c} is learned via contrastive loss within an SNN framework, making it more likely to have captured task-specific information than e_{ac} . This is important as increasing the expressiveness of embeddings will enhance a model's ability to distinguish cell lines of different cancer types. The autoencoder framework that produces e_{ac} has been optimized for reconstruction, making it less likely to have captured the optimal information to predict CDR. To evaluate whether using e_{c} may better segregate cell lines by cancer than e_{ac} , we visualize the clustering of cell lines represented by DeepDSC's e_{ac} (Fig. 4a) and SiamCDR_{RF}'s e_{c} (Fig. 4b) with t-SNE plots. Only cancer types with more than 15 cell lines are shown.

We also calculate e_{c} and e_{ac} 's intra-group similarity and inter-group separability, where groups are defined as cancers. Use of e_{c} reduces intra-cancer similarity by -8.85% compared to e_{ac} , with mean inter-group similarities of 0.91 and 1.00. This indicates each cancer's cell lines are not as tightly clustered when using e_{c} versus e_{ac} . Interestingly, a mean intra-group similarity of 1.00 suggests e_{ac} may not be able to differentiate between cell lines of the same cancer and may consider them identical. In terms of inter-cancer type separability, e_{c} achieves a significant improvement of 77.03% compared to e_{ac} (1.77 vs 1.00; $p < 0.0001$). Importantly, a separability score of 1.0 conveys that cancers cannot be differentiated when using e_{ac} . The inability of e_{ac} to discern cells of the same and different cancer(s) is supported by the lack of well-defined clusters in Fig. 4a.

Conversely, e_{c} 's inter-cancer type separability implies that it can better distinguish between cell lines of distinct cancers. This is supported by Fig. 4b,

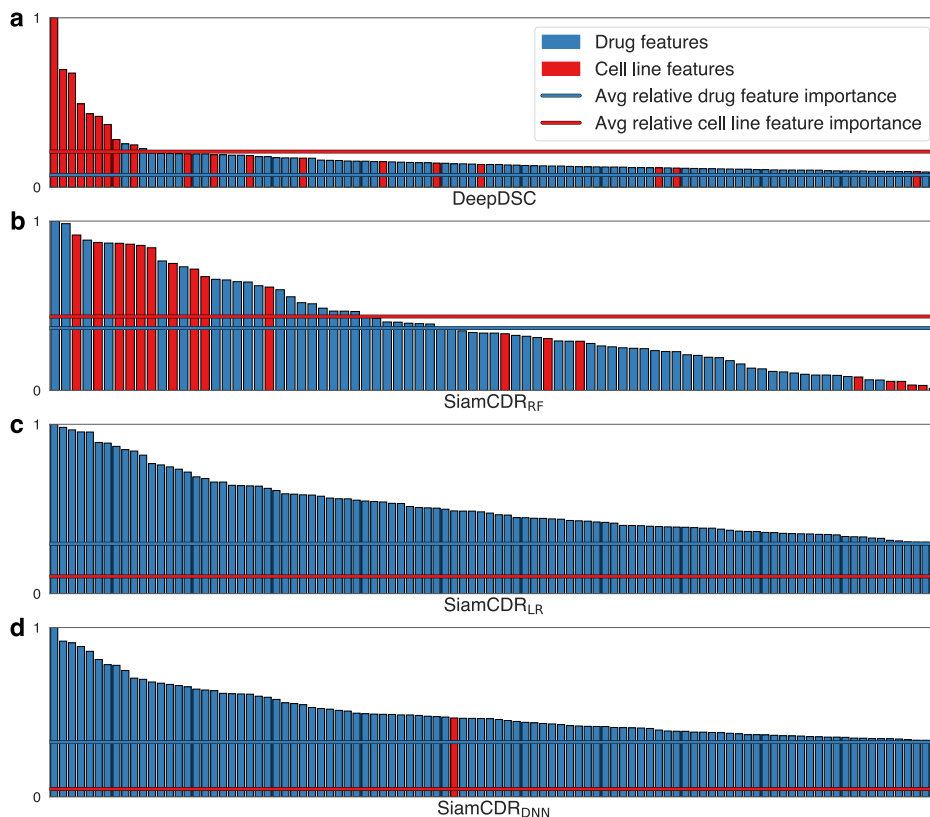
Fig. 2 | Comparing effective score to model predictions. Plots the relationship between CES (continuous effective score) and scores predicted by **a** DeepDSC, **b** SiamCDR_{RF}, **c** SiamCDR_{LR}, and **d** SiamCDR_{DNN}. Each point represents a drug-cell line pair. Drug-cell line pairs containing drugs recommended in the top-3 for at least 3 cell lines by all 4 models are highlighted with distinct colors and shapes (see legend). Note the scale of predicted score in **a** is different than **b**, **c**. This was done to allow the general trend to be visualized.



which illustrates that SiamCDR_{RF}'s e_c separates into distinct clustering structures. For example, all cell lines of lung cancer, leukemia, lymphoma, myeloma, and skin cancer, respectively, exist in single, well-defined clusters for their respective cancer type (circled with dotted lines). The distance of these clusters relative to others corresponds to a limited overlap in tumorigenesis for these cancers. Other cancer types, including bone cancer,

sarcoma, brain cancer, and neuroblastoma, exist within well-defined clusters that are geometrically close to other cancer clusters. This is highlighted in *cluster i* in Fig. 4b (magnified in Fig. 4c). The proximity of brain cancer and neuroblastoma is unsurprising given both involve tumorigenesis within the nervous system. Likewise, the proximity of bone cancer and sarcoma is expected as sarcomas often originate in bone tissue. Previous evidence also

Fig. 3 | Scaled drug- and cell line-derived feature importance for model predictions. Minmax scaled feature importance (>0.01) for the top-100, non-zero features is plotted in descending order along the x-axis for **a** DeepDSC, **b** SiamCDR_{RF}, **c** SiamCDR_{LR}, and **d** SiamCDR_{DNN}. In each plot, the average relative feature importance for both drug- and cell line-derived features is plotted with horizontal lines.



demonstrates a close relationship between some forms of sarcoma and neuroblastomas⁹. *Cluster ii*, magnified in Fig. 4d, represents another multi-cancer grouping (breast cancer, rhabdoid, and kidney). The close proximity of rhabdoid tumors and kidney cancers in the embedding space corresponds to the fact that rhabdoid tumors often originate in the kidneys, indicating that the embeddings accurately capture cancer relations¹⁰. Additionally, the proximity between kidney and breast cancers may indicate that embedding geometry may even capture a cell line's driving mutation(s) (e.g., the risk for both breast and kidney cancer is significantly elevated by *PTEN* mutations¹¹). Finally, we observe in *cluster iii* of Fig. 4b (magnified in Fig. 4e) that the embeddings of colorectal, gastric, bile, and pancreatic cancer cell lines are in close geometric proximity to one another, which may loosely represent gastrointestinal cancers. Each set of multi-cancer groupings further demonstrates the capacity of SiamCDR's framework to embed highly nuanced relationships, thereby producing embeddings (e_c) with high expressiveness that are fine-tuned to predict CDR than DeepDSC's embeddings (e_{ae}).

Comparing the expressiveness of SiamCDR_{RF}'s and DeepDSC's drug representations

The drug representations (e_d) produced by SiamCDR's framework should also capture more task-specific information than the Morgan fingerprints (f) employed by DeepDSC. This is because the fingerprints are generic, sparse vectors produced from predefined heuristics. To assess the expressiveness of each representation scheme, we use t-SNE plots to visualize how drugs are clustered when using either f (Fig. 4f) or e_d (Fig. 4g). Only MOAs with more than 10 drugs present in the pretraining data are highlighted. The included MOAs and their drug counts are reported in Supplementary Table 8.

We also calculate intra-MOA similarity and inter-MOA separability for both f and e_d , finding that the use of e_d significantly improves intra-MOA similarity and inter-MOA separability compared to f by 89.50% (0.77 and 0.40) and 60.82% (1.92 and 1.19), respectively. The significance level for each is less than 0.0001. Interestingly, this improvement demonstrates that

preserving drug relations based on shared gene target captures rich semantic information related to MOA. Moreover, the improvement in both metrics garnered by e_d indicates e_d has higher expressiveness than f .

This is further illustrated by the t-SNE plots. We observe, in Fig. 4g, multiple well-defined clusters corresponding to distinct MOAs. This contrasts with largely unsegregated MOAs in Fig. 4f. Interestingly, in Fig. 4g, *cluster iv* (magnified in Fig. 4h) is comprised of both phosphatidylinositol-3-kinase (PI3K) inhibitors and mammalian target of rapamycin (mTOR) inhibitors and is highly separated from all other MOAs. Both PI3K and mTOR belong to a signaling pathway that controls cell growth and survival¹². Because genes comprising these pathways may be closely linked, the clustering of drugs by pathways may result from overlapping genetic targets of drugs spanning these MOA.

Additionally, some MOAs, such as protein synthesis inhibitors (PSI), do not segregate well and span multiple clusters. Protein synthesis represents a broad category modulated by many genes, which, like genetic pathways, may overlap with other MOAs. For example, in *cluster v* of Fig. 4g, we observe three PSIs near tubulin polymerization inhibitors (TPIs). *Cluster v* is magnified in Fig. 4i. One such PSI, brefeldin-a (BFA), targets *ARF1*, which controls protein secretion and coordinates tubulin polymerization^{13,14}. Thus, it is not unexpected to observe BFA closely associated with TPIs. In another example, CUDC-907 (fimepinostat), classified as a PI3K inhibitor, is found among HDAC inhibitors in *cluster vi* (magnified in Fig. 4j). HDACs facilitate histone modification, modulating gene accessibility for transcription to control gene expression¹⁵. This reclassification is supported by recent evidence investigating its dual-inhibitory properties¹⁶. The work concluded that CUDC-907's HDAC inhibition elicits more significant changes to gene expression than those driven by its PI3K inhibition. As such, the drug embeddings produced by SiamCDR_{RF} may improve MOA classifications, a conclusion supported by the classification of embeddings produced by SiamCDR_{RF} for novel drugs. We embed two drugs, YM-201636 and BNC-105, which are missing gene target data and were excluded from pretraining but have reported MOAs (PI3K inhibitor and TPI). We find that both embeddings place the drugs

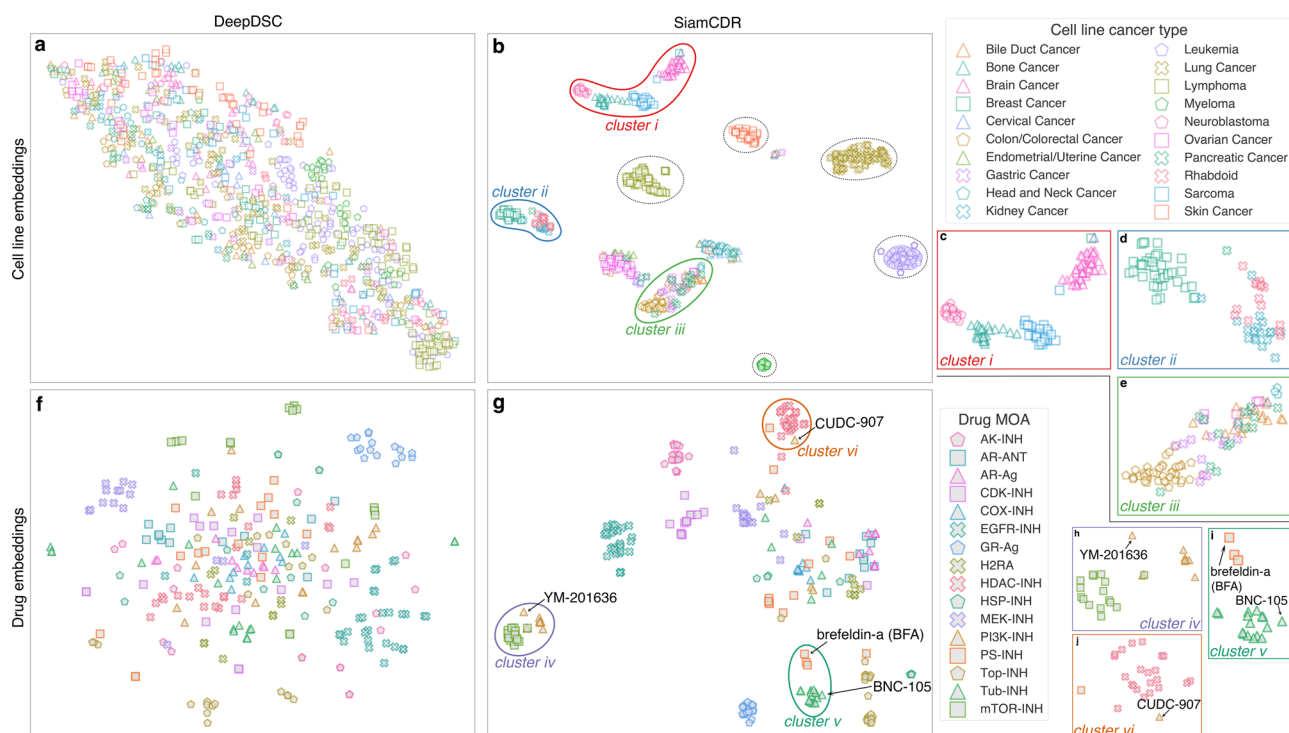


Fig. 4 | t-SNE plots for cell line and drug feature representations. t-SNE plots were from the embeddings produced by both **a** DeepDSC's autoencoder and **b** SiamCDR's cell line encoder for cancers with at least 15 cells. Each cancer is represented by a distinct color/shape combination. In **b**, clusters of single cancer types are highlighted with dotted, black lines; clusters discussed in the Results are highlighted and labeled with distinct colored lines. *Clusters i, ii, and iii* in **b** are magnified in **c, d, and e**, respectively. We also produce t-SNE plots for drugs with MOAs with at least 10 drugs in the pre-training data using either **f** 256-bit Morgan fingerprints (DeepDSC)

or **g** SiamCDR_{RF}'s drug encoder embeddings. Each MOA is represented by a unique color/shape combination. Clusters discussed in the Results section are highlighted in **g** and labeled with distinct colored lines. Clusters *iv, v, and vi* in **g** are magnified in **h, i, and j**, respectively. The axis scales for all plots have been adjusted to best fit the data. PS protein synthesis, TOP topoisomerase, H2RA histamine receptor antagonist, TP tubulin polymerization, AR adrenergic receptor, COX cyclooxygenase, GR glucocorticoid receptor, AK aurora kinase, -INH inhibitor, -A antagonist, -Ag agonist.

within the clusters of their respective MOAs (*clusters v and vi*, respectively). Altogether, this evidence supports the use of our proposed framework for learning biologically relevant drug representations.

Evaluating FDA-approved drug prioritizations

Because prioritization quality can be difficult to ascertain when evaluating novel candidates, we first compare SiamCDR_{RF}'s and DeepDSC's ability to prioritize FDA-approved drugs. If a model can highly prioritize known/approved treatments, it is likely that other highly prioritized candidates may also be effective. Specifically, we examine how highly each model prioritizes the subset of drugs FDA-approved to treat each cell line based on its cancer type compared to all 369 drugs with at least one reported indication. We leverage 805 unseen cell lines for which there was also at least one FDA-approved drug among our data. We score all combinations of these drugs and cell lines using both DeepDSC and SiamCDR_{RF}. For each cell line, both its top-prioritized FDA-approved therapy and the average priority of its FDA-approved drug(s) are presented in Supplementary Data 1. In Table 3, we report by-cancer summaries of these results and include significance levels comparing SiamCDR_{RF}'s and DeepDSC's mean and max prioritizations of FDA-approved treatments; SiamCDR_{RF} achieves significant improvements in the average and max prioritization for 40 and 35% of the evaluated cancers, respectively. Despite not achieving improvements across the majority of evaluated cancers, for 80% of cancers, SiamCDR_{RF}'s top prioritization is a unique FDA-approved drug. Conversely, DeepDSC exhibits limited uniqueness among its top-prioritized drugs; for 50% of the evaluated cancers either docetaxel or doxorubicin is DeepDSC's top-prioritized FDA-approved therapy. This implies DeepDSC has learned to prioritize a few broad-spectrum anti-cancer drugs rather than personalized candidates, which may be driven by an inability to discern individual cells

after they have been embedded as illustrated in Fig. 4a. This notion is further supported by the standard deviation (std) in the priority of each cancer's top-prioritized FDA-approved therapies for each model. The mean std of SiamCDR_{RF}'s 24 top-prioritized drugs is 17.32, while for DeepDSC's 13 drugs it is 1.36. This indicates that while DeepDSC's predictions do vary by cell line (Fig. 2a), the priority-order of drugs does not change. The variance in by-cell line variance in score may indicate DeepDSC has learned to identify cell lines that have fewer effective drugs. However, the lack of individualized drug prioritization suggests that, like SiamCDR_{DNN}, DeepDSC may have overfit to the data.

Case Study 1 – Exploring docetaxel prioritization in breast cancer

To better understand what SiamCDR_{RF} may be leveraging to tailor its prioritizations, we explore SiamCDR_{RF}'s prioritization of docetaxel for breast cancer. This pair was chosen as it follows a bimodal distribution—the model prioritizes docetaxel very highly for some breast cancers and lowly for others. We measure the correlation in expression for each of the 463 genes with SiamCDR_{RF}'s prioritization of docetaxel across all breast cancer cells to ascertain whether transcriptomic differences in breast cancer cells have influenced docetaxel's priority. In Supplementary Data 2, we report 40 genes whose expression is significantly correlated (magnitude > 0.35; significance ≤ 0.1). We use STRING¹⁷ to perform network analysis on the isolated genes and observe significant enrichment of the PI3K-Akt (19 genes), MAPK (11 genes), and RAS (12 genes) signaling pathways. This aligns with documented evidence of transcriptomic changes in *Akt* pathways inducing docetaxel resistance^{18–21}. Moreover, this suggests SiamCDR_{RF} has learned to leverage signals associated with transcriptomic modulation of CDR to tailor its predictions, making it a valuable tool for precision oncology.

Table 3 | DeepDSC and SiamCDR_{RF}'s by-cancer prioritizations of FDA-approved drugs

Cancer	Cell line count	Drug count	Mean priority			Max priority			Highest prioritized drug (%)	
			DSC	CDR	α	DSC	CDR	α	DSC	CDR
Bladder	17	3	22.19	98.92	****	5.65	20.71	****	doxorubicin (100.0)	valrubicin (64.7)
Breast	43	17	133.14	163.68	****	2.00	2.72	****	docetaxel (62.8)	gemcitabine (100.0)
Cervical	18	1	10.00	104.94	****	10.00	104.94	****	topotecan (100.0)	topotecan (100.0)
Colorectal	50	9	183.50	238.18	****	10.98	16.1	****	irinotecan (100.0)	SN-38 (100.0)
E/U	22	1	222.27	332.27	****	222.27	332.27	****	MTX (100.0)	MTX (100.0)
Gastric	39	3	144.81	159.10	**	2.13	25.54	****	docetaxel (100.0)	docetaxel (100.0)
HNC	37	2	168.86	33.27	****	2.11	25.86	****	docetaxel (100.0)	docetaxel (86.5)
Kidney	33	7	84.71	204.76	****	5.64	8.70	****	doxorubicin (100.0)	temsirolimus (84.8)
Leukemia	104	33	190.98	169.65	****	5.62	4.47	****	doxorubicin (100.0)	vincristine (89.4)
Liver	9	1	118.33	290.11	****	118.33	290.11	****	sorafenib (100.0)	sorafenib (100.0)
Lung	118	19	134.55	169.89	****	2.00	2.58	****	docetaxel (91.5)	gemcitabine (99.2)
Lymphoma	83	12	176.66	103.43	****	5.14	3.82	****	doxorubicin (65.1)	romidepsin (100.0)
Myeloma	30	9	143.19	73.22	****	5.67	6.20	*	doxorubicin (100.0)	ixazomib (90.0)
Neuroblastoma	29	1	271.28	67.10	****	271.28	67.10	****	cytoxan (100.0)	cytoxan (100.0)
Ovarian	38	7	176.88	78.73	****	2.79	2.71	****	paclitaxel (100.0)	gemcitabine (100.0)
Pancreatic	24	7	193.88	226.23	****	21.96	2.54	****	everolimus (100.0)	gemcitabine (100.0)
Prostate	10	5	77.78	35.66	****	2.40	2.80	****	docetaxel (100.0)	cabazitaxel (100.0)
Sarcoma	35	2	101.10	140.20	****	58.97	30.86	****	vinblastine (100.0)	vinblastine (100.0)
Skin	51	6	235.68	174.33	****	112.96	39.82	****	vindesine (100.0)	dabrafenib (100.0)
Thyroid	7	5	205.01	286.19	****	89.53	204.73	****	lenvatinib (100.0)	cinacalcet (60.0)

The average mean and max priority of indicated drugs across cell lines with a given cancer (Supplementary Data 1) are presented. Bolded cancers indicate those for which SiamCDR_{RF} achieves higher mean prioritizations than DeepDSC. Significance levels ($\alpha \leq 0.1, 0.05, 0.01, 0.001$; indicated by *, **, ***, and **** respectively) are determined from *p*-values obtained from a two-tailed, independent *t*-test and denote the significance of SiamCDR_{RF}'s prioritizations (CDR) compared to DeepDSC's prioritizations (DSC). (%) in highest prioritized drug indicates the % of cell lines prioritizing the listed drug highest among their FDA-approved indications. Cancer and drug abbreviations – CC Colon/Colorectal, E/U Endometrial/Uterine, HNC Head and Neck, MTX methotrexate, Cytoxan cyclophosphamide.

Case study 2 – Examining repurposing candidates for difficult-to-treat cancers

Lastly, after demonstrating high-performance and personalized prioritizations, we evaluate SiamCDR_{RF}'s capacity to repurposing candidates for cancer treatment. We consider four cancers: two *trained-on* (bladder, BC; and head and neck, HNC) and two *novel* (gastric, GC; and prostate, PC). These were selected via Table 3 from the subset of cancers with fewer than 10 approved therapies and an average max priority for prioritized FDA-approved therapies less than 30. Because SiamCDR_{RF} consistently prioritizes FDA-approved therapies highly, the other highly prioritized candidates may also be effective. Moreover, the limited number of approved therapies implies these cancers are difficult to treat.

For each of these cancers' cell lines, we obtain SiamCDR_{RF}'s prioritizations across all 1119 drugs. Then, for each cancer, we examine the literature for anti-cancer evidence associated with any drug prioritized more

highly, on average, than the top-prioritized FDA-approved drug recorded in Table 3. We present the drug candidates with positive evidence in Table 4. MOAs, gene targets, and indications are obtained from the Broad Institute's drug repurposing hub²². The first row for each cancer in Table 4 is its top-prioritized FDA-approved drug. For cancers whose highest prioritized FDA-approved drug is above 50, only candidates prioritized among the top 50 are considered.

For BC, we found 13/45 (28.9%) drugs prioritized above its top-prioritized FDA-approved drug, valrubicin, with anti-BC evidence. The mean priority of the anti-BC candidates was 18.8 compared to valrubicin's mean priority of 70.6. Notably, dolastatin-10 (dol-10) (mean priority: 5.0) was approved to treat BC after the LINCS and DRH data were published. In addition, there are several drugs with published in vivo anti-BC evidence. Cabazitaxel was found to increase objective response in muscle-invasive bladder cancer by over 2 times the current gold standard (26% to 57%)

Table 4 | Highly prioritized drugs as candidates for repurposing

Cancer type: FDA-approved targeted therapies	Drug	Avg Rank	MOA	Gene targets	Indication	Evidence
Bladder: atezolizumab, avelumab, enfortumab, erdafitinib, nivolumab, pembrolizumab, sacituzumab govitecan-hziy	valrubicin	70.6	DNA-INH TOP-INH	TOP2A	A (BC)	-
	triptolide	2.1	RNA-POL-INH	CYP2C19 RELA	P3 (KD)	26,27
	HOPT	4.8	TOP-INH	TOP1	preclinical	26
	dol-10	5.0	TP-INH	TUBB	P2 (L / Ly)	
	romidepsin	11.5	HDAC-INH	CYP2B6 CYP2C19 CYP3A5 HDAC[1-9]	A (Ly)	55,56
	YM-155	12.1	SUR-INH	BIRC5	P2 (PC, M, Ly)	57
	cabazitaxel	13.4	miTUB-INH	CYP2C8 CYP3A5 SLC01B3 TUBA1[A-B] TUBA3[D-E] TUBA4A TUBB, TUBB[1,6,8] TUBB2[A-B] TUBB4A	A (PC)	23
	gemcitabine	15.3	RR-INH	CMPK1 RRM[1-2] TYMS	A (BrC, NSCLC, OC, PaC)	27,58
	GZD824	20.0	BAK-INH	ABL1 BCR	preclinical	25
	OTS167	25.6	MELZK-INH	MELK	P1 (L)	59
	Oligpraz	26.1	NRF2 activator	ANG CYP2B6 NFE2L2	P3 (NAFLD)	24
	CJDC-907	30.0	PI3K-INH	HDAC2 PIK3R1	P3 (Ly)	60
	temsirolimus	37.2	mTOR-INH	MTOR PTEN	A (RCC)	61
	ixazomib	42.0	Ptsm-INH		A (My)	62
	HNC: cetuximab, nivolumab, pembrolizumab	docetaxel	85.1	TP-INH	BLC2 MAPI2, 4, T] NR112 TUBA1[A-C] TUBA3[C-E] TUBA4A TUBB TUBB2[A-B] TUBB[3, 6, 8] TUBB4[A-B]	A (BrC, NSCLC, PC, GC, HNC)
YM-155		8.9	SUR-INH	BIRC5	P2 (PC, M, Ly)	34
BGT226		23.6	PI3K-INH	MTOR PIK3CA PIK3CB PIK3CG	P1/2 (BrC)	35
OTS167		26.3	MELZK-INH	MELK	P1 (L)	63
temsirolimus		33.4	mTOR-INH	MTOR PTEN	A (RCC)	33
ninlaro		37.3	Ptsm-INH		A (My)	32
litronesib		37.4	KLSP-INH	KIF11	P2 (BrC)	64
SN-38		42.5	TOP-INH	TOP1	A (CC)	36
poziotinib		45.9	EGFR-INH	EGFR ERBB[2, 4]	P2 (LC)	30
docetaxel		79.4	TP-INH	BLC2 MAPI2, 4, T] NR112 TUBA1[A-C] TUBA3[C-E] TUBA4A TUBB TUBB2[A-B] TUBB[3, 6, 8] TUBB4[A-B]	A (BrC, NSCLC, PC, GC, HNC)	-
triptolide		2.3	RNA-POL-INH	CYP2C19 RELA	P3 (KD)	28
EXM		7.1	TOP-INH	TOP1	P3 (PaC)	38
romidepsin		9.3	HDAC-INH	CYP2B6 CYP2C19 CYP3A5 HDAC[1-9]	A (Ly)	40
YM-155		11.7	SUR-INH	BIRC5	P2 (PC, M, Ly)	41
alvespimycin		13.1	HSP-INH	HSP90AA1	P2 (BrC)	65
cabazitaxel	15.6	miTUB-INH	CYP2C8 CYP3A5 SLC01B3 TUBA1[A-B] TUBA3[D-E] TUBA4A TUBB, TUBB[1,6,8] TUBB2[A-B] TUBB4A	A (PC)	37	
BGT226	21.3	PI3K-INH	MTOR PIK3CA PIK3CB PIK3CG	P1/2 (BrC)		
GZD824	22.0	BAK-INH	ABL1 BCR	preclinical	66	
Gastric: fam-trastuzumab, nivolumab, pembrolizumab, ramucirumab, trastuzumab						

Table 4 (continued) | Highly prioritized drugs as candidates for repurposing
Cancer type: FDA-approved targeted therapies

Drug	Avg Rank	MOA	Gene targets	Indication	Evidence
cabazitaxel	17.7	miTUB-INH	CYP2C8 CYP3A5 SLC01B3 TUBA1[A-B] TUBA3[D-E] TUBA4A TUBB, TUBB[1,6,8] TUBB2[A-B] TUBB4A	A (PC)	-
dol-10	2.1	TP-INH	TUBB	P2 (L / Ly)	44
triptolide	3.3	RNA-POL-INH	CYP2C19 RELA	P3 (KD)	29
HCPT	6.3	TOP-INH	TOP1	preclinical	43
EXM	8.0	TOP-INH	TOP1	P3 (PaC)	67
alvespimycin	8.8	HSP-INH	HSP90AA1	P2 (BrC)	39
YM-155	9.1	SUR-INH	BIRC5	P2 (PC, M, Ly)	
camptothecin	13.6	TOP-INH	TOP1	P3 (CC)	42
gemcitabine	16.1	RR-INH	CMPK1 RRM[1-2] TYMS	A (BrC, NSCLC, OC, PaC)	68

Prostate: abiraterone acetate, apalutamide, cabazitaxel, darolutamide, enzalutamide, lutetium Lu 177 vipivotide tetraxetan, olaparib, talazoparib tosylate, radium 223 dichloride, rucaparib camsylate

Abbreviations by column – **Drug** HCPT 10-hydroxycamptothecin, dol-10 dolastatin-10, EXM exatecan-mesylate, nintaro ixazomib citrate, MOA -INH inhibitor, -A antagonist, -Ag agonist, -SA stabilizing agent, TOP topoisomerase, RNA-POL/RNA polymerase, TP tubulin polymerization, HSP heat shock protein, MEK mitogen-activated protein kinase, BCWS bacterial cell wall synthesis, PI3K phosphatidylinositol-3 kinase, mTUB microtubule, RR ribonucleotide reductase, BAK Bcr-Abl kinase, EF estrogen receptor, GF glucocorticoid receptor, PR progesterone receptor, K2SP kinesin-like spindle protein, MELZK maternal embryonic leucine zipper kinase, Ptsm proteasome; **Gene targets** Multiple genes that are identical except for the last character are collapsed using square brackets. Within the brackets, the final letters or numbers are either comma delimited or indicated with ranges. For example, HOX1[-3, 6] would indicate that the drug targets HOX1, HOX2, HOX3, and HOX6; **Indication** BC bladder cancer, BrC breast cancer, CC colorectal cancer, HNC head and neck cancer, KD kidney disease, L leukemia, LC lung cancer, Ly lymphoma, M melanoma, NAFLD non-alcoholic fatty liver disease, NSCLC non-small cell lung carcinoma, OC ovarian cancer, PaC pancreatic cancer, PC prostate cancer, RCC renal cell carcinoma, P clinical trial phase, A FDA-approved.

during a phase II clinical trial²³. Oltipraz reduced bladder carcinogenesis by detoxifying a bladder-specific carcinogen in mice²⁴. In another mouse model, GZD824 displayed activity against FGFR1-mutant BCs, which are especially difficult to treat²⁵. Gemcitabine and 10-hydroxycamptothecin (HCPT), two drugs with demonstrated anti-BC activity, both became more effective when administered in combination with triptolide^{26,27}. Similarly, SiamCDR_{RF} also highly prioritized triptolide to treat both GC and PC: in vitro, triptolide reportedly enhanced apoptotic activity of other anti-cancer drugs when used to pretreat GC cells²⁸ and inhibited PC cell growth²⁹.

For HNC, 8/44 (18.2%) with higher priority than the highest prioritized FDA-approved drug, docetaxel, have anti-HNC evidence. The mean priority of these drugs was 31.9 compared to docetaxel’s priority of 85.1. In clinical trials, poziotinib, litronesib, ninlaro, and temsirolimus each positively affected HNC progression³⁰⁻³³. Additionally, previously published in vitro and in vivo studies present evidence of anti-HNC activity for YM-155, BGT226, and SN-38³⁴⁻³⁶.

For GC, 17.8% of drugs (8/45) with higher priority than the highest prioritized FDA-approved drug, docetaxel, have evidence of anti-GC activity (mean priority of 12.8 and 79.4, respectively). Specifically, alvespimycin, cabazitaxel, and exatecan-mesylate (ExM) each achieved positive results in clinical trials³⁷⁻³⁹. Furthermore, three additional drugs, including triptolide, exhibited positive, pre-clinical evidence: romidepsin and YM-155^{40,41}. Another highly prioritized drug, BGT226, lacks published evidence exploring its anti-GC activity; however, one of its gene targets, PIK3CA, is a common oncogene known to stimulate GC tumorigenesis. As such, further investigation into this candidate is warranted.

Finally, for 43.8% of the drugs (8/16) with higher priority than cabazitaxel, PC’s top-prioritized FDA-approved drug have documented evidence of anti-BC activity. The mean priority of these drugs is 8.4 compared to cabazitaxel’s mean priority of 17.7. Notably, YM-155 was recently approved to treat PC. Alvespimycin demonstrated the capacity to achieve complete responses in patients with PC during phase I clinical trials³⁹. One in vitro study cited ExM as the most potent TOP-INH against PC cells. HCPT, dol-10, and camptothecin each showcased positive pre-clinical evidence as well⁴²⁻⁴⁴.

Discussion

Due to variable drug responses among patients with the same cancer type, optimizing cancer treatment for each patient remains a challenge. Computational methods have been proposed to predict CDR, but their performance is limited by data availability and modeling strategies. To address this, we propose SiamCDR_{RF}, which uses SNNs to pretrain drug and cell line encoders to produce embeddings with high expressiveness. SNNs excel in scenarios with limited data availability as they focus on learning the (dis) similarity between training instances, which enables fine-grained differences, relevant to predicting CDR, to be captured. SiamCDR_{RF} uses an RF model to predict CDR from the learned embeddings. Its balance of simplicity, to prevent overfitting, with an ability to capture complex, non-linear relationships, affords the model enhanced precision in its pan-cancer drug prioritizations over LR and DNN classifiers.

Furthermore, we find SiamCDR_{RF} achieves significant improvements in performance compared to the current SOTA, DeepDSC. SiamCDR_{RF} also identifies FDA-approved therapies and recommends drug repurposing candidates with reasonable success. Notably, via pathway analysis of genes significantly correlated with SiamCDR_{RF}’s drug prioritization, we find significant enrichment of pathways known to induce resistance. This implies the proposed model has learned to tailor its predictions based on transcriptomic signals of drug resistance. Finally, we present 19 drug repurposing candidates for the treatment of either BC, HNC, GC, or PC. Eight of these drugs are unique to an individual cancer, further demonstrating SiamCDR_{RF}’s tailored prioritization.

We also demonstrate in Supplementary Note 2 that SiamCDR_{RF} outperforms SubCDR³ despite the method’s inclusion of additional hard-coded annotations, indicating that a data-driven, unsupervised approach to learning embeddings is better suited and more efficient at capturing relevant

task-specific patterns. Its ability to identify effective drug candidates with high precision for both novel cell lines and cancer types sets SiamCDR_{RF} apart from previous SOTA methods. Moreover, the highly personalized predictions obtained by SiamCDR_{RF} as well as their interpretable nature demonstrates SiamCDR_{RF}'s viability as a tool with the potential to be leveraged by clinicians to advance precision medicine's standard of care for all cancers.

Despite its extremely promising performance, SiamCDR_{RF}, like other deep learning models, would likely benefit from additional data. Drug representations may be further enhanced by leveraging inhibition/activation information while pretraining Enc_d . As observed in Fig. 4d, by using only gene target to determine drug similarity, agonists and antagonists are grouped together in the embedding space. While outside the scope of the current work, in the future, it would also be worthwhile to evaluate common molecular representation techniques to ascertain the optimal method for producing learned drug embeddings within our contrastive learning framework. Cell line representations may also be improved with the use of additional omics types as each omics type possesses distinct information that may further enhance the model's predictive power. Finally, producing pan-cancer drug recommendations is a difficult problem as drug sensitivities vary both within and across cancer types. Given that drug sensitivities among cells of the same cancer are generally more similar than drug sensitivities of cells of different cancer types, a mixture of experts may improve SiamCDR_{RF}'s recommendations by enabling each 'expert' model to focus only on predicting drug sensitivities for cells grouped by cancer type.

Methods

Drug-cell line pairs

Cancer drug response (CDR) data were obtained from the Broad Institute (PRISM Repurposing 19Q4's secondary screen)³¹. In total, this data set contains information on 701,004 drug-cell line pairs. Data were collected from chemical-perturbation viability screens for serial dilutions of 1,448 small-molecule drugs against 480 cell lines. Dose-response curves were fit to the viability screens and used to predict drug response-related metrics. The CDR was calculated from features of the dose-response curve via a custom effective score that will be introduced by Eq. (1) (see "Effective score"). The score was calculated from AUC, lower limit, and IC₅₀ values. As such, pairs missing these values were excluded. As lower limit is the minimum cell viability a drug can achieve, pairs with predicted lower limits < 0 were excluded. Additionally, we excluded drug-cell line pairs with R² < 0.7 to ensure the dose-response curves were well fit to the data. We further processed data by removing duplicate pairs with the following process: (1) in accordance with PRISM documentation, we retained cell lines with the MTS010 screen ID as these are higher quality screens; and (2) if none of the pair's duplicates were from the MTS010 screen, the duplicate with the highest R² was retained. We then removed (1) previously withdrawn drugs and (2) cell lines matching any of the following criteria: (a) less than 1% of the cell line's screened drugs were highly effective ($CES \geq 7.2734$; see "Effective score"); (b) the cell line's cancer type was unknown; and (c) the cell line's RNA-seq was unavailable (see "Cell line gene expression"). These criteria ensure that, for each cell line, a minimum number of effective drugs were observed to learn to predict CDR and that any drug identified as high-priority candidates would be actionable. After removing pairs with missing or low-quality data, we were left with 67,838 drug-cell line pairs from 1,105 drugs and 419 cell lines. Drugs and cell lines are initially represented as 256-bit Morgan fingerprints (radius = 3)^{5,45} and RNA expression of a subset of cancer genes (see "Cell line gene expression") and are denoted by vectors \mathbf{f} and \mathbf{g} , respectively.

Effective score

IC₅₀, or the half-maximal inhibitory concentration, represents the concentration of a drug needed to reach 50% cell viability and partially conveys drug effectiveness. However, this measure only evaluates drug potency, completely ignoring both drug efficacy and the minimum viability a drug can achieve. Two drugs with similar IC₅₀ would be considered similar even if

one reached 5% viability and the other, 49%. To better balance potency and efficacy, we propose a custom effective score (CES):

$$CES = \log \left(\frac{AUC + \text{lower limit} + IC_{50}}{2 \times AUC \times \text{lower limit} \times IC_{50}} \right). \quad (1)$$

AUC for dose-dependent curves conveys a drug's cumulative effect across treatment concentrations by incorporating both potency and effectiveness. Specifically, lower AUC values indicate higher cellular sensitivity to treatment, even at low concentrations. The lower limit is the minimum possible viability a drug can elicit. Including this term will give greater importance to drugs that are more likely to completely eradicate all cells of a cancer. The score is binarized—threshold $\approx 7.27(\mu + 1.28\sigma)$ —such that those with effective scores in the top 10% are labeled 1 and 0 otherwise. Using these labels allows models to be trained to identify only the most effective drugs for each cell line.

Cell line gene expression

We produce cell line representations by leveraging gene expression data obtained from the Broad Institute's Cancer Cell Line Encyclopedia (CCLE) data set (version: DepMap Public 22Q2)⁴⁶. Gene expression was leveraged as it is both information-rich and the most widely available of any omics data type, making our model and its predictions more accessible. The expression of 18,964 protein-coding genes was measured across 1406 cell lines (473 of which were evaluated for their CDR; see "Drug-cell line pairs"). Using KEGG's documented cancer pathways^{47–49}, we considered the expression for only 463 cancer-related genes (Supplementary Note 1). We excluded any remaining cell lines, not evaluated for CDR data if 1) their cancer types had fewer than 10 cell lines with expression data or 2) their listed cancer type was either 'Unknown' or 'Noncancerous'. This yielded 864 cell lines, which we used to pretrain DeepDSC's autoencoder⁴ and SiamCDR's cell line Siamese neural network (SNN) (see "Learning cell line and drug representations"). The distribution of cancers among the pretraining data is presented in Supplementary Table 1a.

The 419 cell lines with CDR data were used for model training, validation, and testing. Cancers evaluated on fewer than 15 cell lines had all their cell lines (67) reserved as a test set; we denote this subset of the data: *novel*. These were withheld to facilitate an evaluation of model generalizability on cancer types unseen during training. Supplementary Table 1b reports the distribution of cell lines by cancer among the novel cancer test set. The remaining cell lines underwent random training/test splitting. To preserve the distribution of cancer across each split, random sampling was conducted by cancer. Specifically, 15% of the cell lines from each cancer were reserved for testing model generalizability on cell lines with cancer types observed during training. This subset is denoted *trained-on* cancers. The remaining cell lines were randomly assigned to one of five folds, stratified by cancer type, to be used for 5-fold cross-validation. The distribution of cell lines by cancer in the training folds and trained-on test set is reported in Supplementary Table 1c.

Baseline method

We evaluate our proposed framework by comparing the performance of models trained from it against a SOTA method: DeepDSC⁴. We select this method for its good performance relative to other published methods and its use of both gene expression and drugs as input. As such, improvements in performance garnered by models produced by our proposed framework will not be due to different data used. We do not compare against RefDNN, despite its use of transcriptomics, as it requires a reference drug set for which the response is known for all cell lines. There are no drugs for which this is the case in our CDR data set, nor in real applications is this likely to be the case. Other models such as DeepCDR⁵⁰, and DeepDR⁵¹ may achieve comparable performance to DeepDSC; however, these also utilize multi-omics data sets. Doing so provides the models with additional information about cell lines that is unavailable when using a single omics data type. This precludes fair comparison of methods as it becomes impossible to ascertain

whether differences in performance arise from methodological differences or the use of distinct information.

Overview of SiamCDR framework

Fig. 1a highlights the distinct components comprising the proposed SiamCDR framework. Solid boxes are consistent across each variation of the framework, while boxes with dashed borders represent aspects evaluated during hyperparameter tuning. Drug-cell line pairs are input to the framework. Each drug is initially represented by a vector \mathbf{f} ; each cell line is represented by a vector \mathbf{g} . Depending on the model, one or both of the drugs and cell lines are projected to an embedding space using a DNN encoder (Enc) pretrained by a Siamese neural network (SNN) (see “Learning cell line and drug representations” and Fig. 1b). We denote Enc_d and Enc_c as the drug and cell line encoders, respectively, with Enc_d and Enc_c projecting \mathbf{f} to \mathbf{e}_d and \mathbf{g} to \mathbf{e}_c , respectively. Here, \mathbf{e}_d and \mathbf{e}_c are vectors that represent drug and cell line embeddings, respectively. Embedding drugs and cell lines in this way enhances the framework’s ability to capture the most salient information for CDR predictions, imparting higher expressiveness to the embeddings produced. Drug and cell line representations are then combined and input to an end classifier (see “SiamCDR’s end classifier”) which is used to predict the relative CDR for each drug-cell line pair.

Learning cell line and drug representations

Our proposed SiamCDR framework leverages a SNN⁶ to pretrain Enc_d and Enc_c to encode drugs or cell lines via contrastive learning, respectively. The SNN structure is illustrated in Fig. 1b. During training, an Enc will be applied to a pair of inputs, producing an embedding for each. Note the Enc ’s weights are shared when embedding each input. The SNN then applies a sigmoid activation function to the Euclidean distance between the embeddings, producing a probability that the inputs are from different groups. Ground-truth labels indicate if the inputs are from different groups (1) or the same group (0). Binary cross entropy loss is optimized such that Enc preserves intra- and inter-group similarity relationships within the embedding space. By doing this, members of the same group are mapped close together and those of different groups are mapped further apart, allowing us to explicitly capture task-specific information.

For our context, we employ nearly identical training structures when pretraining Enc_d and Enc_c to embed drugs and cell lines. The only differences are the input and how the ground-truth labels are determined. Drugs are grouped by their gene targets; cell lines are grouped by cancer types. Enc_d , which produce \mathbf{e}_d , is pretrained on the subset of drugs in the drug-cell line pair data with at least one reported gene target. Enc_c , which produce \mathbf{e}_c , is pretrained on 864 cells not evaluated for CDR. Hyperparameters were tuned via an exhaustive search; the best hyperparameters for both Enc_d and Enc_c are indicated in Supplementary Table 2.

Note that the quality of embeddings produced using the SNN framework was evaluated against those produced using momentum contrastive learning⁵². We found that, for our context, SNNs produced embeddings with 5-times greater separation while retaining comparable cohesion. The complete comparison is presented in Supplementary Note 3.

SiamCDR’s end classifiers

We evaluate the following classifiers: DNN, logistic regression (LR), random forest (RF). DNNs are extremely popular because of their flexibility and capacity to learn complex, nonlinear patterns. However, this capacity comes at the cost of many learnable parameters, increasing the risk of overfitting when training data is limited. On the other hand, LR models are simple models that make predictions based on linear combinations of features. Thus, while they benefit from having fewer trainable parameters than DNN, making them less prone to overfitting, they fail to capture non-linear patterns. Finally, RF models are ensemble methods comprised of multiple decision trees with each tree trained on a bootstrap-sampled subset of the training data. Its ensemble nature, which relies on consensus among trees, provides greater stability to predictions. Furthermore, the risk of overfitting can be mitigated by controlling tree depth and the minimum number of

observations in each split. The complexity of the entire RF model is limited—compared to DNNs—by the simplicity of individual trees. However, by making a series of nested, linear decision boundaries, these trees can capture more complex relationships than LR models. Hyperparameter tuning options for RF and DNN classifiers are reported in Supplementary Tables 3a and b, respectively.

Model training

Each DL method (Enc_d , Enc_c , and SiamCDR_{DNN}) was trained using an ADAM optimizer. Additionally, to mitigate the risk of overfitting we implement dropout and early stopping. The learning rate is also exponentially decayed during training. The specific hyperparameters used to train Enc_d , Enc_c , and SiamCDR_{DNN} are reported in Supplementary Tables 2b and 3b. To improve SiamCDR_{RF}’s ability to generalize during inference, we evaluated the number of estimators and the minimum samples required for splits (Supplementary Table 3a).

Model selection

We evaluate the performance of all combinations of the overall framework discussed above. This includes each combination of drug (\mathbf{f} or \mathbf{e}_d) and cell line (\mathbf{g} or \mathbf{e}_c) embeddings. End classifiers are applied to the concatenated drug-cell line pairs. From the results presented in Supplementary Tables 4 and 5, we select the top-performing architecture for each end classifier with respect to $P_{cell}@k$ and $P_{cancer}@k$ for trained-on cancers. The best model for the DNN, RF, and LR end classifiers is denoted SiamCDR_{DNN}, SiamCDR_{RF}, and SiamCDR_{LR}, respectively. All three models leverage learned cell line representation (\mathbf{e}_c). SiamCDR_{RF} represents drugs with learned embeddings (\mathbf{e}_d), while SiamCDR_{DNN} and SiamCDR_{LR} obtain their best performance when using Morgan fingerprints (\mathbf{f}).

Evaluation metrics

We evaluate model performance via precision@k for both cell lines and cancers (Eqs. (2) and (3), respectively), which measures the proportion of highly effective drugs among a model’s top-k prioritized drugs. Each metric provides a different level of granularity with which to evaluate the precision of model prioritizations. Specifically, $P_{cell}@k$ reports the average precision@k across all cell lines; $P_{cancer}@k$ reduces the impact of outliers by mean pooling $P_{cell}@k$ for each cancer type. For each model, we obtain each cell line’s prioritization order by sorting drug candidates by their predicted CDR scores in descending order. The predicted score conveys the probability that a drug is effective against the given cell line. The closer a drug’s score is to 1, the more likely a model believes the drug to be effective. Drugs prioritized near the top are considered more promising candidates for a given cell line than lower-prioritized drugs.

$$P_{cell}@k(i) = \frac{D_i(k) \cap E_i}{D_i(k)} \tag{2}$$

$$P_{cancer}@k(j) = \frac{\sum_{i=1}^{N_j} P_{cell}@k(i)}{N_j} \tag{3}$$

In Eq. (2), $D_i(k)$ is the set of top-k prioritized drugs for the i th cell line and E_i is the set of all effective drugs for that cell line. In Eq. (3), N_j is the set of cell lines of the j th cancer. In both cases, higher $P_{cell}@k$ and $P_{cancer}@k$ scores indicate a model is better able to prioritize effective drugs at the very top.

Statistical analysis

Significance of the performance of top-3 best SiamCDR models (SiamCDR_{LR}, SiamCDR_{RF}, and SiamCDR_{DNN}) compared to the DeepDSC baseline across 5-fold cross-validation was obtained via Bonferroni multiple-hypothesis corrected, pairwise, two-tailed, independent t-tests.

Feature importance

We measure feature importance to discern their individual contributions to model predictions. Higher importance reflects greater contribution. In our context, we anticipate high-performance models to leverage both drug and cell line information when making predictions, ranking features derived from both drugs and cell lines highly with respect to their importance. Random forest models (e.g., SiamCDR_{RF}) provide feature importance as a native attribute. For DNNs (e.g., SiamCDR_{DNN} or DeepDSC), we estimate individual feature importance as the mean magnitude of SHAP⁵³ values calculated across 5000 randomly sampled training examples. Finally, we estimate feature importance for logistic regression classifiers (e.g., SiamCDR_{LR}) by the magnitude of their coefficients. However, we first evaluate coefficient stability as unstable coefficients serve as poor indicators of feature importance. Coefficient stability reflects consistency in model feature preferences and can be assessed by measuring the variance in coefficients obtained across multiple training folds. High stability is indicated by a small variance relative to the average magnitude of coefficients and vice versa. We find the mean-variance of SiamCDR_{LR}'s coefficients to be two orders of magnitude smaller than the mean value of SiamCDR_{LR}'s coefficients (7.99×10^{-3} and 0.38, respectively), suggesting that SiamCDR_{LR}'s coefficients can be reliably used to estimate feature importance.

Clustering

We use the t-SNE method⁵⁴ (Euclidean distance and perplexity = 30) to project the drug and cell line embedding spaces into two-dimensional spaces. This facilitates a qualitative evaluation of the expressiveness of embeddings used by DeepDSC and SiamCDR_{RF}. Proximity within the projected spaces is positively associated with embedding similarity. As such, we expect embeddings with high expressiveness to produce well-defined clusters comprised of drugs or cell lines of the same (or similar) MOA or cancer, respectively.

Embedding expressiveness

We assess the expressiveness of our learned drug and cell line representations using intra-group similarity and inter-group separability. Groups are defined as MOAs for drugs and cancer types for cell lines. In other words, these metrics evaluate how well drug representations cluster by MOA and how well cell line representations cluster by cancer type. We define intra-group similarity as the average pairwise cosine similarity between all members of a group and inter-group similarity as the average pairwise cosine similarities of each group member to all non-group members. Both range from -1 to 1 with more positive values denoting increased similarity and more negative values associated with increased dissimilarity. We gauge the distinctiveness of a group's cluster from those of other groups via the ratio between its intra- and inter-group similarity, which we denote 'inter-group separability'. Higher inter-group separability indicates higher degrees of separation between groups. In our context, embeddings with high expressiveness should be capable of discerning drugs of different MOAs and cell lines of different cancer types. As such, they will, ideally, attain high values with respect to both metrics.

Reporting summary

Further information on research design is available in the Nature Research Reporting Summary linked to this article.

Data availability

All of the data used in this work is open source and can be freely obtained from the Broad Institute. Both the processed CCLE⁴⁶ gene expression data (CCLE_expression.csv) and the cell line meta data (sample_info.csv) can be accessed through the DepMap portal (<https://www.depmap.org>; version: DepMap22Q2 Public). The CDR data (filename: secondary-screen-replicate-collapsed-logfold-change.csv) is also available through the DepMap portal (version: PRISM Repurposing 19Q4). The code used to process the data is freely available on GitHub (See **Code availability**).

Code availability

All of the code and instructions to replicate our work is freely available here: <https://github.com/ninglab/SiamCDR>. Additionally, the source and versions of our work's primary software, hardware, and data dependencies are presented in Supplementary Table 9.

Received: 30 October 2023; Accepted: 22 March 2024;

Published online: 18 May 2024

References

- Choi, J., Park, S. & Ahn, J. RefDNN: a reference drug based neural network for more accurate prediction of anticancer drug resistance. *Sci. Rep.* **10**, 1861 (2020).
- Zou, H. & Hastie, T. Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Ser. B: Stat. Methodol.* **67**, 301–320 (2005).
- Liu, X. & Zhang, W. A subcomponent-guided deep learning method for interpretable cancer drug response prediction. *PLoS Comput. Biol.* **19**, e1011382 (2023).
- Li, M. et al. DeepDSC: A deep learning method to predict drug sensitivity of cancer cell lines. *IEEE/ACM Trans. Comput. Biol. Bioinform.* **18**, 575–582 (2021).
- Morgan, H. L. The generation of a unique machine description for chemical structures—a technique developed at chemical abstracts service. *J. Chem. Doc.* **5**, 107–113 (1965).
- Koch, G., Zemel, R. & Salakhutdinov, R. Siamese Neural Networks for One-shot Image Recognition. *ICML Deep Learning Workshop* **2** (2015)
- Hoffer, E. & Ailon, N. Deep Metric Learning Using Triplet Network. in *Similarity-Based Pattern Recognition* (eds. Feragen, A., Pelillo, M. & Loog, M.) 84–92 https://doi.org/10.1007/978-3-319-24261-3_7 (Springer International Publishing, Cham, 2015).
- Wang, J. et al. Learning Fine-Grained Image Similarity with Deep Ranking. in 1386–1393 <https://doi.org/10.1109/CVPR.2014.180> (IEEE Computer Society, 2014).
- Cavazzana, A. O., Miser, J. S., Jefferson, J. & Triche, T. J. Experimental evidence for a neural origin of Ewing's sarcoma of bone. *Am. J. Pathol.* **127**, 507–518 (1987).
- Geller, J. I., Roth, J. J. & Biegel, J. A. Biology and treatment of Rhabdoid Tumor. *Crit. Rev. Oncog.* **20**, 199–216 (2015).
- Fusco, N. et al. PTEN alterations and their role in cancer management: are we making headway on precision medicine? *Genes* **11**, 719 (2020).
- Peng, Y., Wang, Y., Zhou, C., Mei, W. & Zeng, C. PI3K/Akt/mTOR pathway and its role in cancer therapeutics: are we making headway? *Front. Oncol.* **12**, 819128 (2022).
- Misumi, Y. et al. Novel blockade by brefeldin A of intracellular transport of secretory proteins in cultured rat hepatocytes. *J. Biol. Chem.* **261**, 11398–11403 (1986).
- Bellouze, S. et al. Golgi fragmentation in pmn mice is due to a defective ARF1/TBCE cross-talk that coordinates COPI vesicle formation and tubulin polymerization. *Hum. Mol. Genet.* **23**, 5961–5975 (2014).
- Seto, E. & Yoshida, M. Erasers of histone acetylation: the histone Deacetylase enzymes. *Cold Spring Harb. Perspect. Biol.* **6**, a018713 (2014).
- Liao, W. et al. Therapeutic Potential of CUDC-907 (Fimepinostat) for Hepatocarcinoma Treatment Revealed by Tumor Spheroids-Based Drug Screening. *Front. Pharmacol.* **12**, 658197 (2021).
- Szklarczyk, D. et al. The STRING database in 2021: customizable protein-protein networks, and functional characterization of user-uploaded gene/measurement sets. *Nucleic Acids Res.* **49**, D605–D612 (2021).
- Huang, P. et al. A comprehensive RNA study to identify circRNA and miRNA biomarkers for Docetaxel resistance in breast cancer. *Front. Oncol.* **11**, 669270 (2021).

19. Dey, G., Bharti, R., Das, A. K., Sen, R. & Mandal, M. Resensitization of Akt induced docetaxel resistance in breast cancer by 'Iturin A' a Lipopeptide molecule from marine bacteria *Bacillus megaterium*. *Sci. Rep.* **7**, 17324 (2017).
20. Vinod, B. S. et al. Resveratrol chemosensitizes HER-2-overexpressing breast cancer cells to docetaxel chemoresistance by inhibiting docetaxel-mediated activation of HER-2-Akt axis. *Cell Death Discov.* **1**, 1–9 (2015).
21. Qin, Y. et al. Autophagy and cancer drug resistance in dialogue: Pre-clinical and clinical evidence. *Cancer Lett.* **570**, 216307 (2023).
22. Corsello, S. M. et al. The drug repurposing Hub: a next-generation drug library and information resource. *Nat. Med.* **23**, 405–408 (2017).
23. Challapalli, A. et al. A single-arm Phase II trial of Neoadjuvant Cabazitaxel and Cisplatin Chemotherapy for muscle-invasive transitional cell carcinoma of the urinary bladder. *Clin. Genitourin. Cancer* **19**, 325–332 (2021).
24. Iida, K. et al. Nrf2 is essential for the chemopreventive efficacy of Oltipraz against urinary bladder carcinogenesis. *Cancer Res.* **64**, 6424–6431 (2004).
25. Jiang, K. et al. GZD824 overcomes FGFR1-V561F/M mutant resistance in vitro and in vivo. *Cancer Med.* **10**, 4874–4884 (2021).
26. Wang, T. et al. Synergistic antitumor effects of triptolide plus 10-hydroxycamptothecin on bladder cancer. *Biomed. Pharmacother.* **115**, 108899 (2019).
27. Yang, Y. et al. Synergistic antitumor effects of triptolide plus gemcitabine in bladder cancer. *Biomed. Pharmacother.* **106**, 1307–1316 (2018).
28. Yuan, W. et al. The antigastric cancer effect of triptolide is associated with H19/NF- κ B/FLIP axis. *Front. Pharm.* **13**, 918588 (2022).
29. Huang, W. et al. Triptolide inhibits the proliferation of prostate cancer cells and down-regulates SUMO-specific protease 1 expression. *PLoS One* **7**, e37693 (2012).
30. Lee, J. H. et al. A phase II study of poziotinib in patients with recurrent and/or metastatic head and neck squamous cell carcinoma. *Cancer Med.* **10**, 7012–7020 (2021).
31. Corsello, S. M. et al. Discovering the anticancer potential of non-oncology drugs by systematic viability profiling. *Nat. Cancer* **1**, 235–248 (2020).
32. Smith, D. C. et al. Phase 1 study of ixazomib, an investigational proteasome inhibitor, in advanced non-hematologic malignancies. *Invest N. Drugs* **33**, 652–663 (2015).
33. Dunn, L. A. et al. A phase II study of temsirolimus added to low-dose weekly carboplatin and paclitaxel for patients with recurrent and/or metastatic (R/M) head and neck squamous cell carcinoma (HNSCC). *Ann. Oncol.* **28**, 2533–2538 (2017).
34. Kumar, B. et al. YM155 reverses cisplatin resistance in head and neck cancer by decreasing cytoplasmic survivin levels. *Mol. Cancer Ther.* **11**, 1988–1998 (2012).
35. Chang, K.-Y. et al. Novel phosphoinositide 3-kinase/mTOR dual inhibitor, NVP-BGT226, displays potent growth-inhibitory activity against human head and neck cancer cells in vitro and in vivo. *Clin. Cancer Res.* **17**, 7116–7126 (2011).
36. Lee, Y.-M. et al. SN-38, an active metabolite of irinotecan, enhances anti-PD-1 treatment efficacy in head and neck squamous cell carcinoma. *J. Pathol.* **259**, 428–440 (2023).
37. Shah, M. A. et al. Multicenter phase II study of Cabazitaxel in advanced gastroesophageal cancer: Association of HER2 Expression and M2-like Tumor Associated Macrophages with Patient Outcome. *Clin. Cancer Res.* **26**, 4756–4766 (2020).
38. Ajani, J. A. et al. A phase II clinical and pharmacokinetic study of intravenous exatecan mesylate (DX-8951f) in patients with untreated metastatic gastric cancer. *Invest N. Drugs* **23**, 479–484 (2005).
39. Pacey, S. et al. A Phase I study of the Heat Shock Protein 90 inhibitor alvespimycin (17-DMAG) given intravenously to patients with advanced solid tumors. *Clin. Cancer Res.* **17**, 1561–1570 (2011).
40. Hui, K. F., Yeung, P. L. & Chiang, A. K. S. Induction of MAPK- and ROS-dependent autophagy and apoptosis in gastric carcinoma by combination of romidepsin and bortezomib. *Oncotarget* **7**, 4454–4467 (2015).
41. Cheng, X. J. et al. Survivin inhibitor YM155 suppresses gastric cancer xenograft growth in mice without affecting normal tissues. *Oncotarget* **7**, 7096–7109 (2016).
42. Zhao, Y., Zheng, Y., Chen, X., Du, R. & Yan, Z. Camptothecin derivatives induce apoptosis and inhibit proliferation of prostate cancer PC-3M cells through downregulation of PI3K/Akt signaling pathway. *Phytochem. Lett.* **46**, 79–89 (2021).
43. Gu, R. & Zhang, Q. Effects of low-dose bufalin combined with hydroxycamptothecin on human castration-resistant prostate cancer xenografts in nude mice. *Exp. Ther. Med.* **22**, 1–9 (2021).
44. Gao, G., Wang, Y., Hua, H., Li, D. & Tang, C. Marine antitumor Peptide Dolastatin 10: Biological activity, structural modification and synthetic chemistry. *Mar. Drugs* **19**, 363 (2021).
45. Landrum, G. et al. Rdkit/rdkit: 2022_03_4 (Q1 2022) Release. Release_2022_03_4. <https://doi.org/10.5281/zenodo.6798971>.
46. Next-generation characterization of the Cancer Cell Line Encyclopedia | Nature. <https://www.nature.com/articles/s41586-019-1186-3>.
47. Kanehisa, M. & Goto, S. KEGG: kyoto encyclopedia of genes and genomes. *Nucleic Acids Res.* **28**, 27–30 (2000).
48. Kanehisa, M. Toward understanding the origin and evolution of cellular organisms. *Protein Sci.* **28**, 1947–1951 (2019).
49. Kanehisa, M., Furumichi, M., Sato, Y., Kawashima, M. & Ishiguro-Watanabe, M. KEGG for taxonomy-based analysis of pathways and genomes. *Nucleic Acids Res.* **51**, D587–D592 (2023).
50. Liu, Q., Hu, Z., Jiang, R. & Zhou, M. DeepCDR: a hybrid graph convolutional network for predicting cancer drug response. *Bioinformatics* **36**, i911–i918 (2020).
51. Chiu, Y.-C. et al. Predicting drug response of tumors from integrated genomic profiles by deep neural networks. *BMC Med. Genomics* **12**, 18 (2019).
52. Chen, X., Xie, S. & He, K. An Empirical Study of Training Self-Supervised Vision Transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision, ICCV*, 9640–9649 (2021).
53. Lundberg, S. M. & Lee, S.-I. A Unified Approach to Interpreting Model Predictions. In *Advances in Neural Information Processing Systems* vol. 30 (Curran Associates, Inc., 2017).
54. van der Maaten, L. & Hinton, G. Visualizing data using t-SNE. *J. Mach. Learn. Res.* **9**, 2579–2605 (2008).
55. Pailas, S. et al. The Histone Deacetylase Inhibitor Romidepsin spares normal tissues while acting as an effective radiosensitizer in bladder tumors in vivo. *Int J. Radiat. Oncol. Biol. Phys.* **107**, 212–221 (2020).
56. Okubo, K., Miyai, K., Kato, K., Asano, T. & Sato, A. Simvastatin-romidepsin combination kills bladder cancer cells synergistically. *Transl. Oncol.* **14**, 101154 (2021).
57. Cui, X. et al. NF- κ B suppresses apoptosis and promotes bladder cancer cell proliferation by upregulating survivin expression in vitro and in vivo. *Sci. Rep.* **7**, 40723 (2017).
58. Messing, E. M. et al. Effect of intravesical instillation of gemcitabine vs saline immediately following resection of suspected low-grade non-muscle-invasive bladder cancer on tumor recurrence: SWOG S0337 randomized clinical trial. *JAMA* **319**, 1880–1888 (2018).
59. Chen, S. et al. Inhibition of MELK produces potential anti-tumour effects in bladder cancer by inducing G1/S cell cycle arrest via the ATM/CHK2/p53 pathway. *J. Cell Mol. Med* **24**, 1804–1821 (2020).
60. Ho, J.-N., Jeon, J. S., Kim, D. H., Ryu, H. & Lee, S. CUDC-907 suppresses epithelial-mesenchymal transition, migration and invasion in a 3D spheroid model of bladder cancer. *Oncol. Rep.* **49**, 1–9 (2023).

61. Pulido, M. et al. Safety and efficacy of temsirolimus as second line treatment for patients with recurrent bladder cancer. *BMC Cancer* **18**, 194 (2018).
62. Sato, A., Asano, T., Okubo, K., Isono, M. & Asano, T. Ritonavir and ixazomib kill bladder cancer cells by causing ubiquitinated protein accumulation. *Cancer Sci.* **108**, 1194–1202 (2017).
63. Ren, L., Deng, B., Saloura, V., Park, J.-H. & Nakamura, Y. MELK inhibition targets cancer stem cells through downregulation of SOX2 expression in head and neck cancer cells. *Oncol. Rep.* **41**, 2540–2548 (2019).
64. A Study in Ovarian, Non-Small Cell Lung, Prostate, Colorectal, Gastroesophageal Cancers, and Squamous Cell Carcinoma of the Head and Neck - Study Results - ClinicalTrials.gov. <https://clinicaltrials.gov/ct2/show/results/NCT01059643>.
65. Kim, J. G. et al. HSP90 inhibitor 17-DMAG exerts anticancer effects against gastric cancer cells principally by altering oxidant-antioxidant balance. *Oncotarget* **8**, 56473–56489 (2017).
66. Liu, X. et al. Preclinical development of HQP1351, a multikinase inhibitor targeting a broad spectrum of mutant KIT kinases, for the treatment of imatinib-resistant gastrointestinal stromal tumors. *Cell & Bioscience* **9**, 88 (2019).
67. Jo, U. et al. TOP1-DNA trapping by exatecan and combination therapy with ATR inhibitor. *Mol. Cancer Ther.* **21**, 1090–1102 (2022).
68. Correia, C. et al. Development of potent CPP6–gemcitabine conjugates against human prostate cancer cell line (PC-3). *RSC. Med. Chem.* **11**, 268–273 (2020).

Acknowledgements

This project was made possible, in part, by support from the National Science Foundation grant no. IIS-2133650 (X.N.) and the National Institute of Aging grant no. 5R01AG071470-03 (X.N., P.L.). Any opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of these funding agencies.

Author contributions

Conceptualization, P.L. and X.N.; methodology, P.L. and X.N.; software, P.L. and B.B.; validation, P.L., B.B., and X.N.; formal analysis, P.L. and X.N.;

investigation, P.L. and X.N.; resources, X.N.; data curation, P.L. and B.B.; writing – original draft, P.L.; writing – review & editing, P.L. and X.N.; visualization, P.L.; supervision, X.N.; project administration, X.N.; funding acquisition, X.N.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41698-024-00589-8>.

Correspondence and requests for materials should be addressed to Xia Ning.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2024