

# Identifying drug-resistant individual cells within tumors by semi-supervised domain adaptation

Kaishun Huang<sup>1</sup> and Hui Liu<sup>1\*</sup>

<sup>1</sup>College of Computer and Information Engineering, Nanjing Tech University, Nanjing, 211800 , Jiangsu, China.

\*Corresponding author(s). E-mail(s): [hliu@njtech.edu.cn](mailto:hliu@njtech.edu.cn);

## Abstract

The presence of pre-existing or acquired drug-resistant cells within the tumor often leads to tumor relapse and metastasis. Single-cell RNA sequencing (scRNA-seq) enables to elucidate the subtle differences in drug responsiveness among distinct cell subpopulations within tumors. A few methods have employed scRNA-seq data to predict the drug response of individual cells to date, but their performance is far from satisfactory. In this study, we propose SSDA4Drug, a semi-supervised few-shot transfer learning method for inferring drug-resistant cancer cells. SSDA4Drug extracts pharmacogenomic features from both bulk and single-cell transcriptomic data by utilizing semi-supervised adversarial domain adaptation. This allows us to transfer knowledge of drug sensitivity from bulk-level cell lines to single cells. We conduct extensive performance evaluation experiments across multiple independent scRNA-seq datasets, and demonstrate the state-of-the-art performance of SSDA4Drug. Remarkably, with only one or two labeled target-domain samples, SSDA4Drug significantly boosts the predictive performance of single-cell drug responses. Moreover, SSDA4Drug accurately recapitulates the temporally dynamic changes of drug responses during continuous drug exposure of tumor cells, and successfully identifies reversible drug-responsive states in lung cancer cells, which initially acquired resistance through drug exposure but later restore sensitivity induced by drug holiday. Also, our predicted drug responses consistently align with the developmental patterns of drug sensitivity observed along the evolutionary trajectory of oral squamous cell carcinoma cells. In addition, our derived SHAP values and integrated gradients effectively pinpoint the key genes involved in drug resistance in prostate cancer cells. These findings highlight the exceptional performance of our method in determining single-cell drug responses. This powerful tool holds the potential for identifying drug-resistant tumor cell subpopulations, paving the way for strides in precision medicine and novel drug development.

**Keywords:** Single-cell sequencing, drug response, semi-supervised learning, transfer learning, domain adaptation, adversarial training

## Introduction

Tumor heterogeneity often leads to substantial differences in patient responses to same drug treatment [1]. Despite remarkable progress in anti-cancer therapy, the presence of pre-existing or acquired drug-resistant cells within the tumor survive and proliferate, ultimately resulting in tumor relapse and metastasis [2]. The drug resistance is the leading cause of clinical failure of tumor therapy. Some scholars have developed bioinformatic methods to predict patients’ clinical response to drugs based on the bulk RNA-seq of cell lines already with large-scale drug sensitivity data, such as GDSC [3, 4], CCLE [5, 6] and CTRP [7]. For example, Chawla et al. introduce the Precily [8] framework that predicts drug response in patients and patient-derived xenografts (PDX) by integrating drug molecular structures and pathway enrichment scores. Meanwhile, He et al. introduce the context-aware deconfounding autoencoder (CODE-AE) [9], which transfer knowledge learned from cell lines to tumor tissues and deconfounds confounding factors involved in drug responses. Additionally, Ma et al. present a few-shot learning for drug response prediction that can translate from high-throughput screens to individual patients [10]. Despite the potential, bulk RNA-seq fails to discern variations in transcriptional programs and regulatory mechanisms within individual cells, thereby obscuring the intratumor heterogeneity. Consequently, previous models perform poorly on out-of-distribution samples, such as patients and individual cells.

The advancement of single-cell sequencing technology enables us to explore the complexity and variability among individual cells, offering us the opportunity to gain insights into the intratumor heterogeneity [11]. However, the currently available drug response data at single-cell level covers only a few cancer types and drugs due to cost and technical limitations [12], which pose challenges on the development of computational methods for predicting single-cell drug responses. Several studies pay attention to the knowledge transfer of drug responses from cell lines to single cells. The unsupervised domain adaptation technique, including maximum mean discrepancy [13–15], adversarial learning [16–19], or reconstruction learning [20, 21], has been used to align the data distribution between the source and target domains. For example, scDeal [22] integrates bulk RNA-seq and scRNA-seq data by aligning domains via maximum mean discrepancy (MMD) to predict single-cell drug responses. Following this idea, SCAD [23] employs adversarial domain adaptation to predict single-cell drug sensitivity. However, they can only extract domain-invariant features through unsupervised domain adaptation, but failed in determining the decision boundary of sensitivity and resistance in the target domain. Besides, Beyondcell [24] calculates an enrichment score to determine the therapeutic differences among cell populations and generate a prioritised sensitivity-based ranking for drug selection. scRank [25] employs a target-perturbed gene regulatory network to rank drug-responsive cell populations via *in*

*silico* drug perturbations using untreated single-cell transcriptomic data. Nevertheless, they either depends on human-designed score to rank drugs or prior knowledge to identify cell populations of interest.

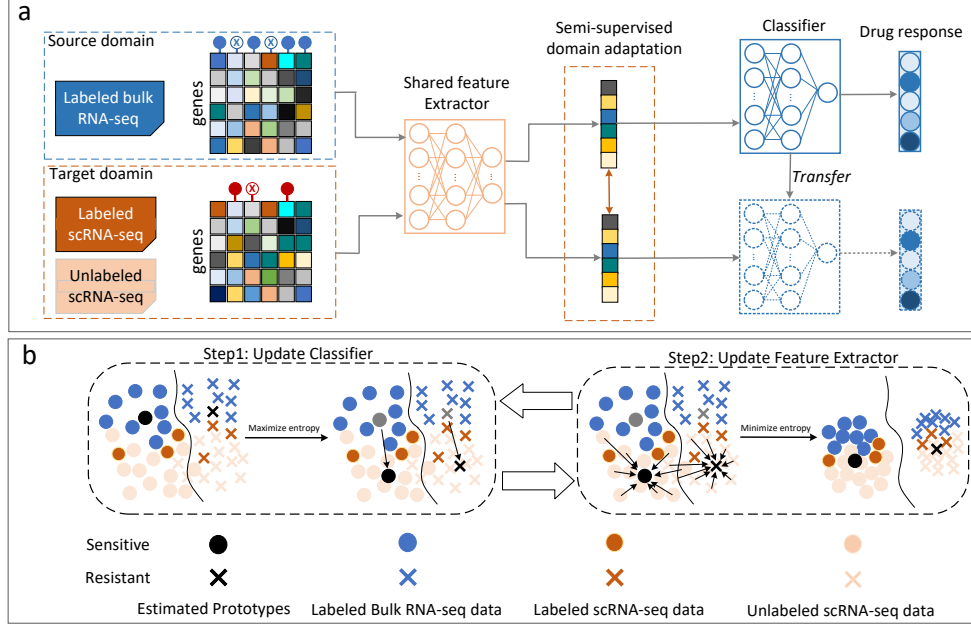
In recent years, semi-supervised domain adaptation methods achieve better performance by leveraging a small number of labeled samples in target domain [26–30]. Inspired by semi-supervised domain adaptation [31] and adversarial training [32], we introduce SSDA4Drug to transfer drug response from the source domain (bulk RNA-seq) to the target domain (scRNA-seq). By alternately maximizing the conditional entropy of unlabeled target samples with respect to the classifier and minimizing it with respect to the feature encoder, SSDA4Drug successfully reduces the distribution gap between source and target domains, and learns the discriminative features specific to the classification task, which significantly improves the performance in target domain. Remarkably, we found that with only a few labeled target-domain samples, SSDA4Drug effectively overcomes the limitations of unsupervised domain adaptation methods and greatly enhances drug response predictive performance at single-cell level. Our extensive experiments across multiple independent scRNA-seq datasets confirm that SSDA4Drug achieves superior performance compared to existing methods. Moreover, SSDA4Drug can not only effectively identifies the temporally dynamic changes in single-cell drug sensitivity during tumor cells exposure to drugs, but also accurately predict the developmental trend in drug sensitivity during cancer cell evolutionary trajectory. Finally, we explore the SHAP and integrated gradients of each input gene, and successfully uncover several genes relevant to drug responses.

## Results

### Semi-supervised few-shot transfer learning

SSDA4Drug is a deep learning framework that leverages semi-supervised domain adaptation techniques to translate drug response knowledge from source domain to the target domain (Fig. 1a; Supplementary Fig.7). In this study, we designate bulk RNA-seq data from cell lines as the source domain and scRNA-seq data as the target domain. Within the source domain, each cell line has been assigned response label to specific drug, whereas only a limited number of cells in the target domain carry such labels. Our model takes inspiration from adversarial perturbations [32] and minimax entropy techniques [31].

We utilize a shared denoising autoencoder to extract the low-dimensional representations for both source and target domains. A fully-connected network functions as a classifier to determine drug response labels. To effectively exploit the unlabeled target samples for extracting discriminative features, the class prototypes (representatives of each class) are trained to gravitate towards target samples by maximizing the conditional entropy of unlabeled target examples with respect to the classifier, while the feature extractor is updated to minimize the entropy to make them better clustered around their respective prototypes (Fig. 1b). By optimizing the minimax loss on the conditional entropy, coupled with adversarial perturbations to bolster the local smoothness, SSDA4Drug not only reduces the distribution gap between domains but also learns discriminative features specific to the classification task.



**Fig. 1:** Overview of model architecture and illustration of semi-supervised domain adaptation. (a) Extraction of discriminative and domain-invariant features via semi-supervised domain adaptation to predict drug response for target-domain samples. (b) Illustrative diagram of semi-supervised domain adaptation based on minimax entropy. The feature extractor and classifier are trained by adversarial learning to fully exploiting unlabeled target samples to derive better features for classification tasks.

## SSDA4Drug achieved high performance in predicting single-cell drug responses

To evaluate the predictive power of our method on single-cell drug responses, we firstly tested it on the datasets released by SCAD [23]. The datasets included ten single-cell drug response data of five distinct cell lines treated by 9 chemotherapy and targeted drugs (see Table 1 for details). We considered two different gene sets as model input: "all" included all genes shared by both source and target domains, "tp4k" included the top 4,000 genes exhibiting highly variable expression levels in the target domain. To balance the proportion of sensitive and resistant cell lines in the source domain, a weighted-sampling algorithm [33] was employed to randomly selected samples for minority class based on the assigned weights.

For each dataset listed in Table 1, we built a model and evaluated its performance separately. It is worth noting that only 3 cells were randomly selected as labeled target samples, while the remaining served as unlabeled samples for semi-supervised model training and performance evaluation as well. To avoid deviation induced by randomness, we repeated the process 50 times and calculated the mean AUC and mean AUPR as performance metrics. For comparison to SCAD, we used "tp4k" as



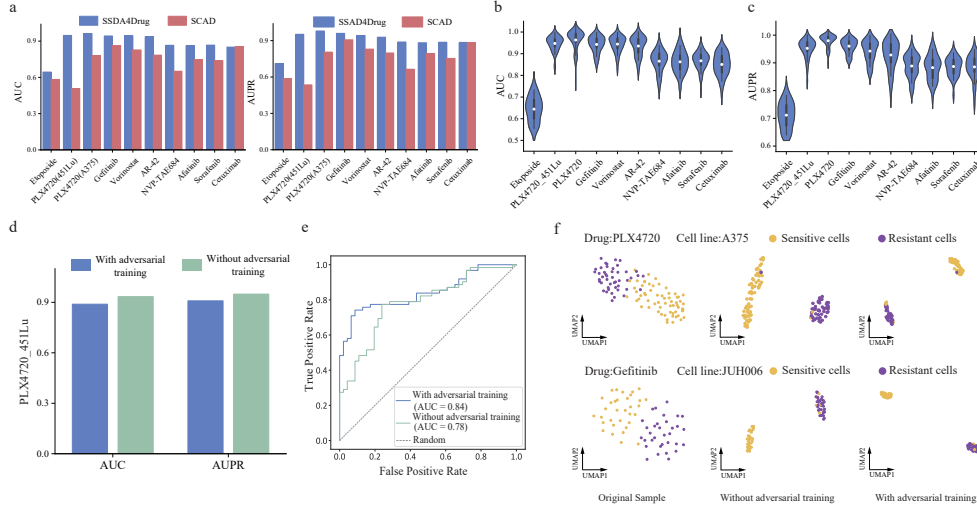
**Table 1:** Details of the single-cell drug response dataset released by SCAD [23]

Drug	Dataset	Cell line	#res.	#sen.	#gene	Cancer type
Etoposide	GSE149215	PC9	764	629	2668	Lung cancer
PLX4720	GSE108383	451Lu	81	74	2856	Melanoma
PLX4720	GSE108383	A375	46	63	2873	Melanoma
Vorinostat	CCL	JUH006	33	33	2933	Oral squamous cell carcinomas
Gefitinib	CCL	JUH006	33	33	2933	Oral squamous cell carcinomas
AR-42	CCL	JUH006	33	33	2933	Oral squamous cell carcinomas
NVP-TAE684	CCL	SCC47	60	60	2876	Oral squamous cell carcinomas
Afatinib	CCL	SCC47	60	60	2876	Oral squamous cell carcinomas
Sorafenib	CCL	SCC47	60	60	2876	Oral squamous cell carcinomas
Cetuximab	CCL	SCC47	60	60	2876	Oral squamous cell carcinomas

the input genes. The results showed that our model achieved significantly higher AUC and AUPR values than SCAD in predicting single-cell responses for all drugs, except Cetuximab (Fig. 2a). Particularly, for PLX4720, our method boosted the AUC by 43.9% and 18.3% compared to SCAD on the 451Lu and A375 melanoma cell lines, respectively. In the case of Cetuximab, our AUC was marginally lower than SCAD’s by just 0.4%. We also illustrated the distribution of AUC and AUPR values across the 50 repetitions for each dataset (Fig. 2b,c), and their low variances (Supplementary Table 1 and Supplementary Data 1) underscored the consistency and reliability of our model.

To assess our model’s generalization capability, we evaluated its performance across different cell lines treated with the same drug. Specifically, we trained the model on the 451Lu cell line treated with PLX4720 and tested it on the A375 cell line exposed to the same drug. For this analysis, we utilized "all" 11,355 common genes shared by both cell line datasets as model inputs. Only 3 451Lu cells were randomly selected as labeled target-domain samples, while the remaining cells served as unlabeled data for model training. Meanwhile, we investigated whether the adversarial perturbation could enhance model generalizability. First, for 451Lu single-cell responses, our model exhibited comparable performance with and without adversarial training (Fig. 2d). Remarkably, when tested on A375 cells, the model trained with adversarial perturbations achieved a significantly higher AUC value of 0.84 compared to the model without adversarial training, which attained 0.78 AUC value (Fig. 2e). These results not only demonstrate our model’s superiority in predicting single-cell drug responses across multiple datasets but also highlight its excellent generalization ability across different cell lines.

To further validate the beneficial effects of adversarial training on model robustness, we applied Uniform Manifold Approximation and Projection (UMAP) dimensionality reduction [34] to visualize the latent representations of A375 and JUH006 cells treated by PLX4720 and Gefitinib, respectively. The UMAP plots showed that adversarial training led to tighter clusters of intra-class samples from the target domain in the latent space, while inter-class samples were more widely separated (Fig. 2f and Supplementary Fig.1). This observation showed that the adversarial perturbations refined the decision boundary, making our model less sensitive to slight variations in input data and significantly enhancing its generalization performance.



**Fig. 2:** Evaluation of SSDA4Drug performance in predicting single-cell drug responses. (a) Average AUC and AUPR values of SSDA4Drug and SCAD on ten scRNA-seq datasets spanning 9 distinct drugs and 5 cell lines. (b-c) Violin plots of AUC and AUPR values of SSDA4Drug achieved by 50 random repetitions on each scRNA-seq dataset. (d) Barplots of the AUC and AUPR values of SSDA4Drug with and without adversarial training for 451Lu cells subjected to PLX4720 treatment. (e) ROC curves and AUC values for SSDA4Drug trained on 451Lu cell line while tested on A375 cell line, with and without adversarial training. (f) UMAP plots of the latent representations obtained by SSDA4Drug with and without adversarial training for the A375 and JUH006 cells treated by PLX4720 and Gefitinib, respectively.

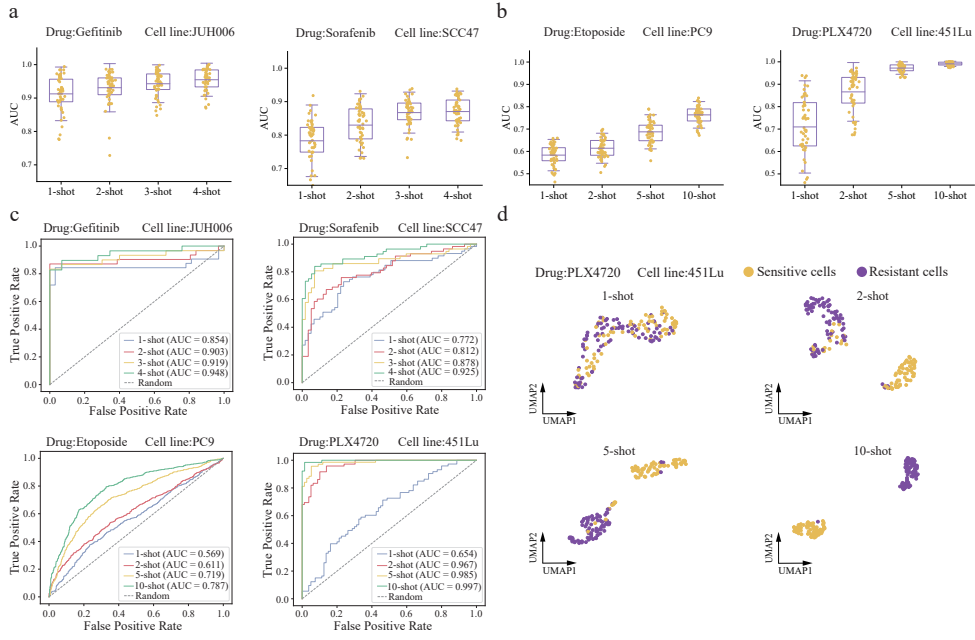
## Few-shot brought great performance enhancement

For systematic evaluation of few-shot impact on performance, we gradually increased the number of labeled target samples included in the training set. Using the datasets in Table 1, we employed distinct increment in labeled target samples, due to differing numbers of single cells across various drugs. For Gefitinib and Sorafenib, we progressively increased the labeled cells from 1 to 4. For Etoposide and PLX4720-451Lu, we used 1, 2, 5, and 10 shots for model training, respectively. In each few-shot task, the designated number of labeled target samples were randomly selected for the training set, with the remainder used to assess model performance. To mitigate randomness-induced deviations, we repeated each few-shot task 50 times.

The results revealed a significant enhancement in performance for all four drugs as the number of labeled target samples increased (Fig. 3 and Supplementary Fig.2). Specifically, for Gefitinib and Sorafenib, the mean AUC values started at 0.912 and 0.782 for 1-shot and increased by 4.6% and 11.2% at 4-shot, respectively (Fig. 3a). In the case of Etoposide and PLX4720, their performance at 10-shot gained a substantial boost of 31% and 39% in mean AUC values compared to 1-shot scenario, respectively (Fig. 3b, Supplementary Data 2). A close-up view of the ROC curves and AUC values

revealed the great performance improvement brought by the shot increment on the four drug-induced single-cell response datasets (Fig. 3c). Notably, the variance in AUC for PLX4720-451Lu dataset decreased greatly with the increasing number of shots, indicating improved model reliability.

To visualize these effects, we utilized UMAP plots to depict the embeddings of target samples from the PLX4720-451Lu dataset, derived from models trained with varying shot numbers (Fig. 3d). Evidently, as the number of shots increases, the intra-class single-cell embeddings clustered closer together, while the inter-class samples moved further apart. These findings not only confirmed the benefits of few-shot learning in predicting single-cell drug responses, but also demonstrated that our model's discriminatory power and reliability boosted with the increasing number of few-shot samples.



**Fig. 3:** Performance enhancement of SSDA4Drug brought by increasing number of shots in target domain. (a) Boxplot of AUC values on Gefitinib-JUH006 and Sorafenib-SCC47 scRNA-seq datasets by SSDA4Drug trained using 1, 2, 3, and 4 labeled target-domain shots. (b) Boxplots of AUC values on Etoposide-PC9 and PLX4720-451Lu scRNA-seq datasets, obtained by SSDA4Drug trained using 1, 2, 5, and 10 labeled target-domain shots. (c) ROC curves and AUC values on four different scRNA-seq datasets achieved by SSDA4Drug trained with varying number of target-domain shots, respectively. (d) UMAP plots of the latent representations for PLX4720-451Lu scRNA-seq data achieved by SSDA4Drug trained with different number of target-domain shots.

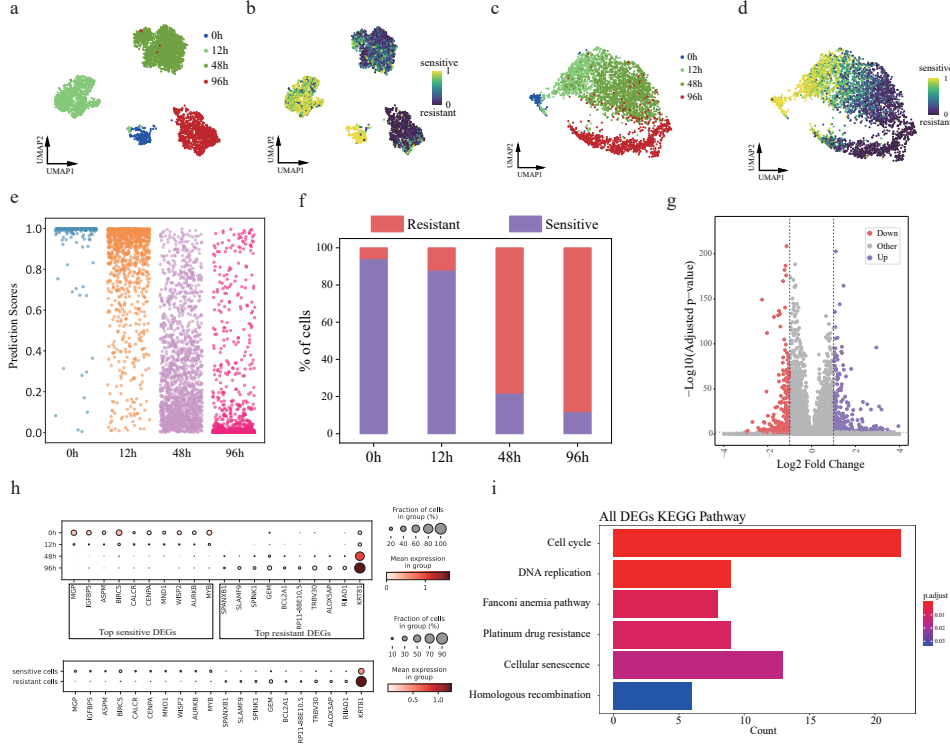
## SSDA4Drug captured dynamic state transition toward drug resistance

Previous studies have shown that cancer cells often evolve rapidly in response to therapeutic pressure, and eventually develop resistance to specific drug [35]. To further validate the performance of SSDA4Drug, we assessed its ability to predict temporally dynamic changes of drug response for individual cells during persistent chemotherapy drug exposure. For this purpose, we used the dataset published by Ben-David et al. [36], which encompassed 7,440 single-cell clones derived from MCF7 cell line. These cells were exposed to 500nM Bortezomib, a 26S proteasome inhibitor, for 2 days. After filtering out low-quality data, 3,740 cells were retained for subsequent analysis. They underwent single-cell sequencing at various time points: before treatment (t0,  $n=160$ ), 12 hours after drug exposure (t12,  $n=994$ ), 48 hours after exposure (t48,  $n=1623$ ), and 96 hours after drug washout and recovery (t96,  $n=963$ ) (Fig. 4a). The UMAP plots revealed distinct clustering of the transcriptional profiles of these cells into four groups, implying rapid transcriptional state transitions induced by drug exposure.

We tested whether SSDA4Drug could accurately predict the temporal variations of MCF7 cells in response to Bortezomib based on the scRNA-seq data. To this end, we assigned 20% of the cells at t0 ( $n=32$ ) as sensitive and 20% of the cells at t96 ( $n=193$ ) as resistant. Then, we randomly selected only 3 cells from each category as labeled target samples, with the remaining serving as unlabeled target samples for model training. All cells were used to verify SSDA4Drug’s ability to capture the temporal changes of drug efficacy at different time points. We analyzed the predicted scores of individual cells, and revealed that almost all MCF7 cells were sensitive to Bortezomib before treatment. However, as drug exposure duration increased, the number of sensitive cells gradually decreased while the number of resistant cells increased. At 96 hours post-treatment, the cell population contained the highest proportion of resistant cells (Fig. 4b,e). The UMAP plots showed that the learned embeddings of the cells diverged gradually from sensitivity to resistance induced by drug treatments (Fig. 4c,d). Our results were highly consistent with the findings of a previous study [37].

We were intrigued by the molecular mechanism that underlies the transition of cellular state from sensitive to resistant. For this purpose, we categorized all cells as either resistant or sensitive based on the predicted scores using a threshold 0.5. This categorization revealed a gradual increase in the proportion of resistant cells over time (Fig. 4f). Subsequently, we conducted differential expression analysis ( $\log_{2}FC > 1$  and  $p\text{-adj} < 0.05$ ) between the two groups of cells with assigned response labels, and identified 385 drug-sensitive genes and 255 drug-resistant genes (Fig. 4g). Among them, we selected the top 10 overexpressed genes from both sensitive and resistant DEGs, and visualized their respective cell proportions and gene expression levels at each time point (Fig. 4h). Two distinct gene expression patterns emerged when comparing t0 and t12 to t48 and t96. Among the top 10 drug-resistant DEGs, a few genes have been reported to be associated to drug resistance (Supplementary Table 2). For instance, BCL2A1 has been shown that its high expression in tumor samples correlates with in vivo chemoresistance [38]. SPINK1 overexpression correlates with Hepatocellular Carcinoma treatment Resistance [39]. Finally, we performed KEGG functional enrichment analysis on the differentially expressed genes and found that the overexpressed

genes were enriched in cell proliferation-related signaling pathways, such as cell cycle and DNA replication (Fig. 4i). These experimental results demonstrated the capacity of our method to predict the dynamics of cellular responses to drug treatment across multiple time points.



**Fig. 4:** SSDA4Drug accurately predict the temporal dynamics of 3,740 MCF7 cells in response to Bortezomib at four time points. (a-b) UMAP plots of the scRNA-seq data all cells exposed to Bortezomib colored by time points and predicted sensitivity scores, respectively. (c-d) UMAP plots of the latent representations obtained by SSDA4Drug for all cells colored by time points and SSDA4Drug predicted scores, respectively. (e) Scatter plots of SSDA4Drug's predicted scores for all cells at four time point. (f) Percentage of sensitive and resistant cells predicted by SSDA4Drug at each time point. (g) Volcano plot for differential expression analysis between sensitive and resistant cell groups predicted by SSDA4Drug. (h) Upper and lower plots show the cell fraction and normalized expression levels of top 10 differentially expressed genes across four time points and predicted cell groups, respectively. (i) KEGG enrichment analysis for the differentially expressed genes between sensitive and resistant cells predicted by SSDA4Drug.

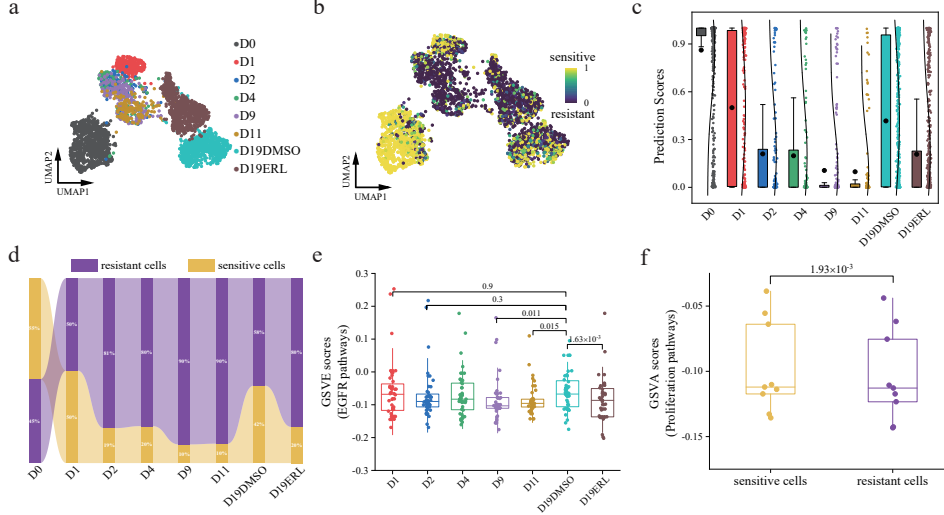
## SSDA4Drug identified reversible drug response states of lung cancer cells

To further evaluate whether SSDA4Drug can identify reversible drug response states of cancer cells, we applied it to the dataset published by Aissa et al. [40]. This dataset comprised the Drop-seq data from lung cancer PC9 cells cultured in vitro, including untreated cells, erlotinib-treated cells, and cells underwent a drug holiday (Fig. 5a). After filtering out low-quality data, 3,053 cells were retained for analysis. Specifically, These cells covered various time points during Erlotinib treatment: before treatment (D0,  $n=765$ ), after 1 day (D1,  $n=236$ ), 2 days (D2,  $n=148$ ), 4 days (D4,  $n=100$ ), 9 days (D9,  $n=228$ ), and 11 days (D11,  $n=143$ ) of drug exposure, as well as the cells followed by re-treatment for 2 days after a 6-day drug holiday (D19ERL,  $n=726$ ) and cells followed an 8-day drug holiday (D19DMSO,  $n=707$ ). The number of living cells decreased from D1 to D4 during drug exposure and then increased thereafter, indicating a transition from Erlotinib sensitivity to resistance in PC9 cells (Supplementary Fig.3). The cell subpopulation that has developed resistance survived and proliferated, thereby leading to growth in cell numbers even during drug exposure. In contrast, the D19DMSO cells regained sensitivity to Erlotinib to a large extent, implying drug holiday successfully reversed acquired drug resistance [40]. In fact, previous studies have suggested the possibility of temporary reversal of drug resistance, exemplified by the “drug holiday” phenomenon observed during clinical therapy by EGFR inhibitors [41].

Our objective was to evaluate SSDA4Drug’s predictive ability regarding the reversible drug response states and explore the underlying biological mechanism. We labeled the cells before treatment (D0) as sensitive, and the cells received 9 days (D9) and 11 days (D11) treatments as resistant ones. These cells were used as labeled target-domain samples to train the model, which was then used to predict the drug response states of the cells at other time points. Consistent with the actual observations, the numbers of sensitive cells predicted by SSDA4Drug declined from D1 to D11 during drug exposure (Fig. 5b,c), while the cells predicted to be resistant increased in number. As expected, the predicted proportion of sensitive cells in D19DMSO significantly increased after drug holiday (Supplementary Fig.4). According to the predicted scores (higher score means higher sensitivity), we categorized individual cells as either sensitive or resistant using a predicted value threshold of 0.5. As a result, the proportion of sensitive cells was found to gradually decrease but increased again after the drug holiday (Fig. 5d). Specifically, D19DMSO cells showed a 22.8% increase in sensitive cells compared to D2 cells, and a 22.3% increase compared to D19REL cells re-treated after the drug holiday. These findings demonstrate a high degree of consistency between our predictions and the actual situation.

Considering that Erlotinib is an inhibitor of EGFR pathway, we employed the GSVA tool [42] to assess the activity levels of EGFR-related signaling pathways (Supplementary Data 5) in each cell. The mean GSVA scores across all cells at various time points served as indicators of pathway activity levels. The results revealed a correlation between the GSVA scores for EGFR pathway and the proportion of drug-sensitive cells, elucidating the dynamic shifts in cellular sensitivity to the drug treatment (Fig. 5c,e).

Additionally, drug-sensitive cells exhibited higher GSVA scores in cell proliferation-related signaling pathways compared to drug-resistant cells (Fig. 5f). The findings was consistent with previous study indicating that actively proliferating cells exhibit high sensitivity to specific anticancer drugs [8]. These collective insights demonstrated that our method actually captured pharmacogenomic features related to drug response from expression profiles, rather than just memorizing the samples, thereby yielding to strong generalizability.



**Fig. 5:** SSDA4Drug successfully identified reversible drug response states of PC9 cells subjected to Erlotinib treatment. (a-b) UMAP plot of scRNA-seq data of PC9 cells exposed to Erlotinib, colored by eight time points and predicted sensitivity scores, respectively. (c) Boxplots of SSDA4Drug's predicted scores for PC9 cells at eight time points. (d) Area plot of the percentage of sensitive and resistant cells predicted by SSDA4Drug at each time point. (e) Boxplots of GSVA scores regarding the activity levels of EGFR-related pathways. P-values were obtained from the paired-sample t-test. (f) Boxplot of GSVA scores regarding the activity levels of proliferation-related pathways in the sensitive and resistant cell groups. P-values were obtained from the paired-sample t-test.

## SSDA4Drug predicted development of drug resistance along with evolutionary trajectory

We seek to explore the evolutionary trajectory of cancer cells and its intricate linkage to the development of drug resistance. When subjected to the selective pressure of drug therapy, cancer cells undergo genetic alterations and phenotypic shifts, enabling them to withstand the lethal effects induced by therapeutic agents. This adaptation



poses significant challenges to the treatment efficacy. Therefore, it is important to identify the pivot point where cancer cells transition from being drug-sensitive to drug-resistant during their evolutionary journey. This motivated us to explore the correlation between drug response state predicted by SSDA4Drug and the evolutionary trajectory of tumor cells.

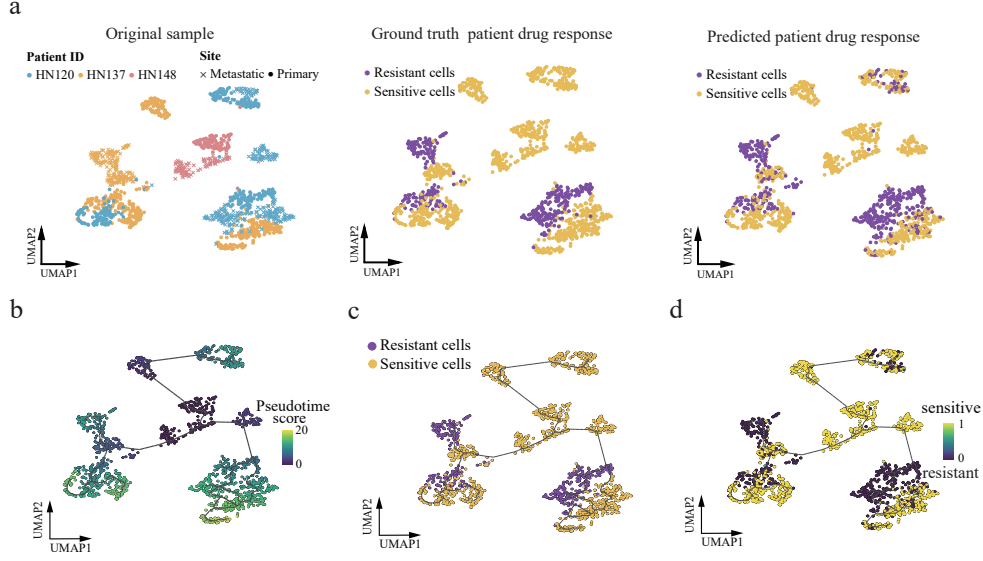
To this end, we conducted analysis on the GEO dataset (accession number GSE117872) [43]. This dataset comprised the scRNA-seq data of 1,302 cells collected from three patients with oral squamous cell carcinomas (OSCC). Each cell in the dataset has been labeled as either sensitive or resistant to Cisplatin. We randomly selected 50 cells as labeled target samples to train the model, and subsequently assessed its performance on the remaining cells. The results showed that 88% of the cells were correctly predicted, with an AUC score of 0.927 and an AUPR score of 0.978. Among them, 98% of the predicted resistant cells and 84.6% of the predicted sensitive cells matched their ground-truth labels (Fig. 6a). Next, we performed pseudotime analysis on the scRNA-seq data using Monocle3 [44]. The result reflected the inferred evolutionary trajectory based on transcriptional profiles of these cancer cells (Fig. 6b). By depicting the actual drug response state of the cells along the trajectory, we could clearly observe the progression of cancer cells from the sensitive state to resistant state (Fig. 6c). As expected, our comparative analysis between the predicted drug resistance progression and the ground-truth data confirmed that the development of drug resistance is tightly coupled with the evolutionary trajectory of cancer cells (Fig. 6c,d). These findings suggested that our method successfully identified the transcriptional pattern of cancer cells shifting from sensitivity to resistance along the evolutionary trajectory.

### SSDA4Drug uncovered key genes relevant to drug response

We were interested in exploring the ability of SSDA4Drug to uncover critical genes determining the single-cell drug response. So, we applied it to another single-cell drug response dataset (GEO accession number GSE140440) [45]. This dataset comprised the scRNA-seq profiles of 324 cells derived from two prostate cancer cell lines, PC3 and DU145. Each cell has been assigned the label either sensitive or resistant to Docetaxel (Fig. 7a). We randomly selected three cells as labeled target-domain samples to train the model. For performance evaluation on this dataset, we found that 96% of the cells were correctly predicted for their sensitivity or resistance to Docetaxel, with an AUC value of 0.97. Based on the predicted scores, we classified each cell as either resistant or sensitive using the threshold of 0.5, and found that the predicted drug responses aligned almost perfectly with the ground-truth drug responses.

Following the proof of performance, we employed the SHapley Additive exPlanation (SHAP) [46] value to uncover the key genes involved in drug response. Therefore, we calculated the mean absolute SHAP values for each gene in sensitive and resistant cell groups, which represented their average impact on the drug response prediction (Fig. 7b). To further illustrate the contributions of the genes to resistance and sensitivity, we generated beeswarm plots for the resistant cell group (Fig. 7c) and the sensitive cell group (Supplementary Fig.6). By examining two representative Docetaxel-resistant genes, GPSM3 and HLA-DQB1 (Fig. 7d), we observed their

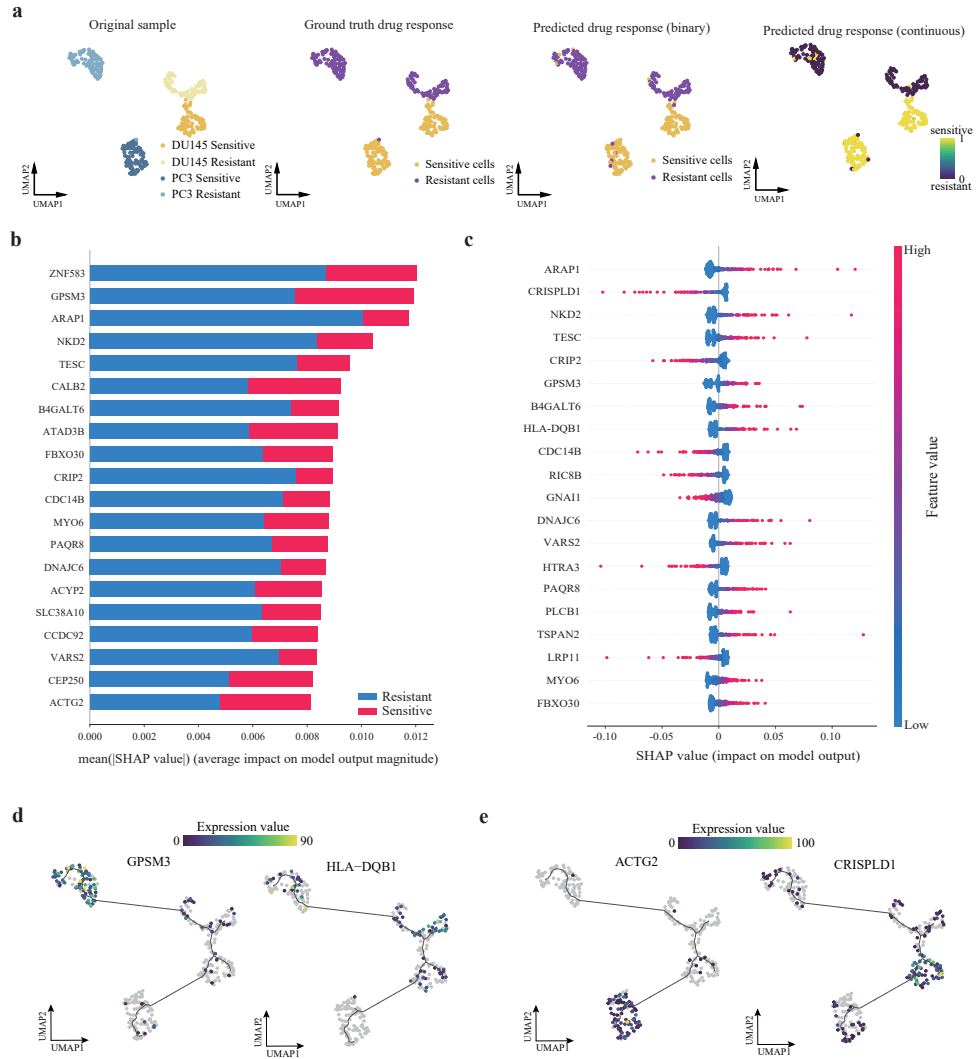




**Fig. 6:** SSDA4Drug predicted drug resistance progression highly coupled with the evolutionary trajectory of cancer cells when subjected to the selective pressure of Cisplatin treatment. (a) UMAP plots of the scRNA-seq data of OSCC cells colored by sample sources, ground-truth drug response labels, binarized drug response labels based on predicted scores, respectively. (b-d) Pseudotime trajectory of the OSCC cells in which the cell was colored by pseudotime scores, ground-truth drug response labels, and predicted drug response probability, respectively.

increased expression levels closely coupled with the transition of cells towards drug-resistant phenotype (Fig. 7a,d). Conversely, two representative Docetaxel-sensitive genes, ACTG2 and CRISPLD1, demonstrated an opposite expression pattern (Fig. 7e). Previous study has reported GPSM3 as a potential immunotherapy target in low-grade Gliomas [47], while ACTG2 was shown to control the growth of small intestinal neuroendocrine tumor (SI-NET) cells [48].

In addition, we leveraged the integrated gradients (IG) [49] to uncover key genes. To this end, We computed the integrated gradient score for each gene in individual cells, and then run differential expressed analysis (sensitive vs resistant) based on their IG values (Supplementary Data 4). We visualized the top 10 genes with the highest IG values in both resistant and sensitive cells (Supplementary Fig.5), from which we found distinct patterns between the two cell populations. These findings demonstrated our model's interpretability for identifying the genes that contributed most significantly to drug responses.



**Fig. 7:** Discovery of key genes involved in the drug response of PCa cells to Docetaxel treatment. (a) From left to right: UMAPS plots of cells colored by tissue sources, ground-truth labels, binarized drug response labels and predicted scores, respectively. (b) Mean absolute SHAP values of top 20 genes contributing significantly to the drug response state of PCa cells. (c) SHAP beeswarm plot of top 20 genes contributing to drug resistance in PCa cells. (d) Pseudotime trajectory plots for two representative drug-resistant genes, GPSM3 and HLA-DQB1, colored by their respective gene expression levels. (e) Pseudotime trajectory plots of two representative drug-sensitive genes, ACTG2 and CRISPLD1, colored by their respective gene expression levels.

## Discussion

Tumor heterogeneity often leads to significant differences in the response to the same drug treatment among different cancer cells. The advancement of single-cell sequencing technology has opened up new avenues to identify pre-existing and acquired drug-resistant tumor cell subpopulations. However, the currently available single-cell drug responses covers only a few cancer types and drugs, due to cost and technical limitations. Therefore, the computational methods with strong predictive power is a practical solution to determine single-cell drug response. However, due to insufficient data for model training, the performance of current predictive models is far from satisfactory. In this paper, we utilized the semi-supervised domain adaptation to transfer the pharmacogenomic knowledge from bulk-level to single-cell level, establishing a more effective method for building predictive model with superior performance. The experimental results demonstrated that our method possessed sufficient generalizability and notably outperformed previous methods based on unsupervised domain adaptation.

Although this study primarily used the gene expression profiles for feature extraction, the flexibility of the SSDA4Drug framework allows for potential integration with multi-omics data, including but not limited to genomic mutations, copy number variations, epigenetic data. Furthermore, despite of one model per drug, we believe that a more valuable research direction in the future would undoubtedly be to build a framework capable of predicting unseen drugs. This would avoid the tedium and limitations of training separate models for each drug.

One limitation of our model is the rely on a small number of single cells with response labels, which are not currently always available. However, as the cost of single-cell sequencing technology decreases, the scRNA-seq data of clinical tissues from recurrent or drug-resistant tumors is bound to become more common, effectively addressing the issue of insufficient data. Meanwhile, we can continuously update the cell lines used for model training to cover more tumor types and drugs. Finally, with the growth of bulk drug response data of specific cancer type, we can attempt to train models using specific cancer cell lines instead of pan-cancer cell lines, thereby yielding better predictive performance of single-cell drug response for specific cancer types of interest.

## Conclusion

In this paper, we proposed a semi-supervised deep transfer learning method for predicting drug responses of individual tumor cells. We have demonstrated that incorporation of a few labeled single-cell samples into the domain adaptation not only aligned the distribution of source and target domains, but also learned a better classification decision boundary. Our extensive experiments have shown that our model outperformed current state-of-the-art methods [23], but also captured the dynamic changes in drug sensitivity of tumor cell populations during drug exposure, and reversals of response state induced by drug holidays. Moreover, we explored the model interpretability for uncovering key genes that play important role in cellular response to drug treatments.

## Materials and Methods

### Data source and preprocessing

#### Bulk drug sensitivity

The bulk drug sensitivity data was retrieved from the Genomics of Drug Sensitivity in Cancer (GDSC) database [3, 4], which provided drug sensitivities of 1074 cancer cell lines upon 226 drugs. Drug sensitivity was evaluated using half maximal inhibitory concentration (IC50) and area under the dose-response curve (AUC) measurements. We assigned drug response label to each cell line, by using the method similar to CODE-AE [9]. For a drug of interest, we ranked the IC50 values of all cell lines tested against this drug, and then categorized them as sensitive or resistant. The binarization threshold was set as the average IC50 value across all cell lines tested against the drug. The sensitive cell lines were labeled as 1, while the resistant ones were labeled as 0. Together with the drug sensitivity data, we retrieved the matched bulk RNA-seq data (RMA-normalized basal expression profiles) of the cell lines from GDSC.

#### Single-cell drug response

For comprehensive performance evaluation, we collected a number of scRNA-seq datasets where cells were derived from different sources, including cell lines and patient tumor tissues. The scRNA-seq data and single-cell drug response labels were obtained from the National Center for Biotechnology Information’s (NCBI) Gene Expression Omnibus (GEO) and Broad Institute (see Data availability).

### Overview of SSDA4Drug model

Given the abundant drug response labels in the source domain (cell lines) and a limited number of labels in the target domain (single cells), we proposed the SSDA4Drug model designed for single-cell drug response prediction. Inspired by minimax entropy [31], our model integrates the semi-supervised domain adaptation and adversarial training techniques, aiming to learn robust decision boundaries for classification tasks and effectively align domain distributions simultaneously (Fig. 1a).

Our model consists of two components: a autoencoder and a classifier. We define the feature extractor  $\mathcal{F}$  as the encoder plus the network layers except for the last linear layer of the classifier, while the last linear layer of the classifier as  $\mathcal{C}$  (Supplementary Fig.7). We train  $\mathcal{F}$  and  $\mathcal{C}$  using labeled source-domain samples and a small number of labeled target-domain samples, and evaluate the prototypes of the sensitive and resistant cells. Since the classifier is primarily trained on labeled samples, the estimated prototypes may be biased towards the source samples due to scarcity of labeled target samples. Let the conditional entropy represent the similarity between the estimated prototypes and target features, we firstly train  $\mathcal{C}$  to maximize the entropy between the prototypes and the unlabeled target features. This would pull the prototypes toward the center of the unlabeled target features. Next, we train  $\mathcal{F}$  to minimize the entropy so that the unlabeled target features are clustered around the prototypes. To summarize, the minimax entropy process is actually the adversarial learning between  $\mathcal{F}$

and  $\mathcal{C}$ . Repeating such entropy maximization and entropy minimization process yields discriminative features. Meanwhile, we utilize adversarial perturbations to enhance the local smoothness of the model. The adversarial samples regarding the data points near the decision boundary could clarify the decision boundaries between drug-sensitive and drug-resistant cells.

## Denoising autoencoder for feature extraction

We utilize a common denoising autoencoder to extract a low-dimensional representation for both the source and target domains, taking as input the expression profiles of labeled and unlabeled samples. Let  $D_s = \{(\mathbf{x}_i^s, y_i^s)\}_{i=1}^{N_s}$  represent  $N_s$  labeled source samples, where  $\mathbf{x}_i$  represents the expression profile and  $y_i = \{1, 0\}$  represents the corresponding label (sensitive or resistant) of  $i$ -th example upon specific drug treatment. Similarly,  $D_t = \{(\mathbf{x}_i^t, y_i^t)\}_{i=1}^{N_t}$  represents  $N_t$  labeled target samples, and  $D_u = \{(\mathbf{x}_i^u)\}_{i=1}^{N_u}$  represents  $N_u$  unlabeled target samples ( $N_t \ll N_u$ ).

To learn robust features, we introduce random noise generated from a binomial distribution to each gene expression value, so that we get the noisy expression profile  $\mathbf{x}'$ . The encoder maps the noisy gene expression profiles to a low-dimensional latent representation  $z$ , while the decoder strives to reconstruct the original noise-free data from this latent representation. To minimize the reconstruction error, we employ the mean squared error (MSE) loss function as below

$$\mathcal{L}_{reco} = \sum_{i=1}^{N_s+N_t+N_u} \mathbb{E}_{\mathbf{x}_i' \sim q_D(\mathbf{x}_i'|\mathbf{x}_i)} [MSE(\mathbf{x}_i, D(E(\mathbf{x}_i')))] \quad (1)$$

where  $q_D(\mathbf{x}_i'|\mathbf{x}_i)$  represents the probability distribution from which the noise drawn.

## Drug response predictor

Suppose that the final linear layer of the classifier, i.e.  $\mathcal{C}$ , consists of weight vectors  $\mathbf{W}_C \in \mathbb{R}^{d \times c}$ , which maps the features  $\mathcal{F}(\mathbf{x})$  extracted by previous layers to binary class probabilities. Since previous studies have demonstrated that normalization can enhance model stability and performance, we normalize the features and feed them into the last linear layer  $\mathcal{C}$ , followed by a softmax layer with a temperature parameter  $T$ . As a result, the prediction score for each class is defined as  $\mathbf{p}(y|\mathbf{x}) = \sigma(\frac{1}{T} \frac{\mathbf{W}_c^T \mathcal{F}(\mathbf{x})}{\|\mathcal{F}(\mathbf{x})\|})$ , where  $\sigma$  indicates a softmax function. In our study, the sensitive and resistant prototypes are required to update their positions based on maximizing and minimizing entropy, we thus refer to the weight vectors  $\mathbf{W}_C$  as the class prototypes.

We firstly train  $\mathcal{F}$  and  $\mathcal{C}$  using  $D_s$  and  $D_t$  so that the labeled examples can correctly classified. The standard cross-entropy is used to evaluate the classification loss:

$$\mathcal{L}_{cls} = \mathbb{E}_{(\mathbf{x}, y) \in D_s, D_t} CE(\mathbf{p}(y|\mathbf{x}), y) \quad (2)$$

## Adversarial perturbation

The adversarial training [50] introduces subtle perturbations within the vicinity of the input data. The perturbations are oriented towards the direction of loss gradient ascent and are typically generated under  $l_2$  norm constraints. This training strategy demands that the model not only minimizes the empirical risk but also minimizes the adversarial loss, making the model less sensitive to slight changes in the input. Formally, the adversarial loss is defined as below:

$$\begin{aligned} \mathcal{L}_{adv} &= \mathbb{E}_{(\mathbf{x}, y) \in \mathcal{D}_s, \mathcal{D}_t} D[q(y), p(y | \mathbf{x} + r_{adv})], \\ \text{where } r_{adv} &\approx \epsilon \frac{g}{\|g\|_2}, g = \nabla_{\mathbf{x}} D[q(y), p(y | \mathbf{x})]. \end{aligned} \quad (3)$$

Note that  $D[\cdot, \cdot]$  represents the Kullback-Leibler divergence that measures the difference between two distributions,  $p(y|\mathbf{x})$  denotes the probability distribution of the predicted label  $y$  given input  $\mathbf{x}$ ,  $q(y)$  is the true distribution of label which is generally approximated by one-hot vector of  $y$ ,  $g$  is the gradient that can be effectively computed using backpropagation, and  $\epsilon$  is the magnitude of the adversarial perturbation. This perturbation strives to maximize the divergence between actual and predicted label distributions by choosing the direction of gradient ascent.

## Semi-supervised domain adaptation

The minimax entropy method [31] is adopted for semi-supervised domain adaptation, which maximizes the entropy of unlabeled target samples to optimize  $\mathcal{C}$ , and minimizes the entropy with respect to  $\mathcal{F}$  to aggregate features. For the unlabeled samples, we assume the existence of a domain-invariant prototype in both the sensitive and resistant categories, representing the central points for two domains. Initially,  $\mathcal{F}$  and  $\mathcal{C}$  are trained using the labeled source samples and a small number of labeled target samples, and then the initial drug-sensitive and drug-resistant prototypes are evaluated. We utilize the unlabeled target samples to train  $\mathcal{C}$ , maximizing the entropy between the prototypes and the features of unlabeled target samples. This allows us to update the prototype, and bring them closer to the unlabeled target samples. Next, to extract discriminative feature, we train  $\mathcal{F}$  to reduce the entropy between the prototypes and the unlabeled target features, thereby the features are clustered around one of the two prototypes in the latent space. The entropy is defined as below:

$$H = -\mathbb{E}_{(\mathbf{x}, y) \in \mathcal{D}_u} \sum p(y | \mathbf{x}) \log p(y | \mathbf{x}) \quad (4)$$

Therefore, the minimax entropy stands as an adversarial training process, which maximizes the entropy to find the prototypes that can best represent the unlabeled target features and then strives to minimize entropy to cluster the unlabeled features based on these prototypes. In this study we utilize the gradient reversal layer to implement the adversarial learning. Denote by  $\mathcal{L}_{mme}$  the loss of the minimax entropy-based adversarial learning, we have

$$\mathcal{L}_{mme} = \arg \max_{\Theta} \min_{\Phi} H \quad (5)$$

in which  $\Theta$  and  $\Phi$  are the parameters of classifier  $\mathcal{C}$  and feature extractor  $\mathcal{F}$ . As a result, the total loss function of our SSDA4Drug is defined as

$$\mathcal{L} = \mathcal{L}_{cls} + \mathcal{L}_{adv} + \mathcal{L}_{reco} + \lambda \mathcal{L}_{mme} \quad (6)$$

in which  $\lambda$  is a hyperparameter used for the trade-off between the minimax entropy loss and the classification task loss on the labeled samples.

## Identification of key genes

We first consider the Integrated Gradients [49] to identify key genes. Integrated Gradients provides a quantitative assessment of the importance of each input feature in the decision-making, enhancing the interpretability of our model. In our practice, the integrated gradient (IG) scores were calculated using the “IntegratedGradients” function in the Python Captum library [51], taking as inputs the expression matrix, trained model and output labels. It output the IG matrices with the same shape as the input expression matrix, its rows represent genes and columns represent cells and elements are the corresponding IG values [22]. To identify genes with significantly higher IG values in sensitive (or resistant) cells, we performed differential expression analysis using the “rank\_genes\_groups” function. We considered the genes with  $\log_{2}FC > 1$  and  $p\text{-adj} < 0.05$  as key genes.

We also leverage the SHAP values to uncover key genes involved in drug response. SHAP quantifies how much an input feature contributes to pushing the model’s prediction away from the base value, which is typically defined as the average output over the training dataset. We use the DeepExplainer from the SHAP [46] library to compute the SHAP values for each cell. To visualize the SHAP values, the “summary\_plot” function is used to plot the distribution of SHAP values for top 20 genes contributing significantly to final prediction, i.e. resistance or sensitivity. The “beeswarm\_plot” function is used to plot the distribution of SHAP values within a specific group of cancer cells, such as drug-resistant cells.

## Pseudotime analysis

Monocle3 was used to analyze the evolutionary trajectory of different cell subgroups, using the gene-cell matrix in the UMI count scale extracted from the Seurat subset. The UMAP tool was used to reduce the dimension. After dimensionality reduction and cell sorting, evolutionary trajectories were inferred using default parameters of Monocle3.

## References

- [1] Vitale, I., Shema, E., Loi, S., Galluzzi, L.: Intratumoral heterogeneity in cancer progression and response to immunotherapy. *Nature medicine* **27**(2), 212–224 (2021)

- [2] Cree, I.A., Charlton, P.: Molecular chess? hallmarks of anti-cancer drug resistance. *BMC cancer* **17**, 1–8 (2017)
- [3] Yang, W., Soares, J., Greninger, P., Edelman, E.J., Lightfoot, H., Forbes, S., Bindal, N., Beare, D., Smith, J.A., Thompson, I.R., *et al.*: Genomics of drug sensitivity in cancer (gdsc): a resource for therapeutic biomarker discovery in cancer cells. *Nucleic acids research* **41**(D1), 955–961 (2012)
- [4] Iorio, F., Knijnenburg, T.A., Vis, D.J., Bignell, G.R., Menden, M.P., Schubert, M., Aben, N., Gonçalves, E., Barthorpe, S., Lightfoot, H., *et al.*: A landscape of pharmacogenomic interactions in cancer. *Cell* **166**(3), 740–754 (2016)
- [5] Li, H., Ning, S., Ghandi, M., Kryukov, G.V., Gopal, S., Deik, A., Souza, A., Pierce, K., Keskula, P., Hernandez, D., *et al.*: The landscape of cancer cell line metabolism. *Nature medicine* **25**(5), 850–860 (2019)
- [6] Ghandi, M., Huang, F.W., Jané-Valbuena, J., Kryukov, G.V., Lo, C.C., McDonald III, E.R., Barretina, J., Gelfand, E.T., Bielski, C.M., Li, H., *et al.*: Next-generation characterization of the cancer cell line encyclopedia. *Nature* **569**(7757), 503–508 (2019)
- [7] Rees, M.G., Seashore-Ludlow, B., Cheah, J.H., Adams, D.J., Price, E.V., Gill, S., Javaid, S., Coletti, M.E., Jones, V.L., Bodycombe, N.E., *et al.*: Correlating chemical sensitivity and basal gene expression reveals mechanism of action. *Nature chemical biology* **12**(2), 109–116 (2016)
- [8] Chawla, S., Rockstroh, A., Lehman, M., Ratther, E., Jain, A., Anand, A., Gupta, A., Bhattacharya, N., Poonia, S., Rai, P., *et al.*: Gene expression based inference of cancer drug sensitivity. *Nature communications* **13**(1), 5680 (2022)
- [9] He, D., Liu, Q., Wu, Y., Xie, L.: A context-aware deconfounding autoencoder for robust prediction of personalized clinical drug response from cell-line compound screening. *Nature Machine Intelligence* **4**(10), 879–892 (2022)
- [10] Ma, J., Fong, S.H., Luo, Y., Bakkenist, C.J., Shen, J.P., Mourragui, S., Wessels, L.F., Hafner, M., Sharan, R., Peng, J., *et al.*: Few-shot learning creates predictive models of drug response that translate from high-throughput screens to individual patients. *Nature Cancer* **2**(2), 233–244 (2021)
- [11] Han, Y., Wang, D., Peng, L., Huang, T., He, X., Wang, J., Ou, C.: Single-cell sequencing: a promising approach for uncovering the mechanisms of tumor metastasis. *Journal of hematology & oncology* **15**(1), 59 (2022)
- [12] Wei, W., Shin, Y.S., Xue, M., Matsutani, T., Masui, K., Yang, H., Ikegami, S., Gu, Y., Herrmann, K., Johnson, D., *et al.*: Single-cell phosphoproteomics resolves adaptive signaling dynamics and informs targeted combination therapy in glioblastoma. *Cancer cell* **29**(4), 563–573 (2016)



- [13] Xiao, N., Zhang, L.: Dynamic weighted learning for unsupervised domain adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 15242–15251 (2021)
- [14] Chen, C., Chen, Z., Jiang, B., Jin, X.: Joint domain alignment and discriminative feature learning for unsupervised deep domain adaptation. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 33, pp. 3296–3303 (2019)
- [15] Long, M., Zhu, H., Wang, J., Jordan, M.I.: Deep transfer learning with joint adaptation networks. In: International Conference on Machine Learning, pp. 2208–2217 (2017). PMLR
- [16] Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., March, M., Lempitsky, V.: Domain-adversarial training of neural networks. *Journal of machine learning research* **17**(59), 1–35 (2016)
- [17] Long, M., Cao, Z., Wang, J., Jordan, M.I.: Conditional adversarial domain adaptation. *Advances in neural information processing systems* **31** (2018)
- [18] Tzeng, E., Hoffman, J., Saenko, K., Darrell, T.: Adversarial discriminative domain adaptation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 7167–7176 (2017)
- [19] Chen, L., Chen, H., Wei, Z., Jin, X., Tan, X., Jin, Y., Chen, E.: Reusing the task-specific classifier as a discriminator: Discriminator-free adversarial domain adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 7181–7190 (2022)
- [20] Zhuang, F., Cheng, X., Luo, P., Pan, S.J., He, Q.: Supervised representation learning: Transfer learning with deep autoencoders. In: Twenty-fourth International Joint Conference on Artificial Intelligence (2015)
- [21] Zheng, H., Fu, J., Mei, T., Luo, J.: Learning multi-attention convolutional neural network for fine-grained image recognition. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 5209–5217 (2017)
- [22] Chen, J., Wang, X., Ma, A., Wang, Q.-E., Liu, B., Li, L., Xu, D., Ma, Q.: Deep transfer learning of cancer drug responses by integrating bulk and single-cell rna-seq data. *Nature Communications* **13**(1), 6494 (2022)
- [23] Zheng, Z., Chen, J., Chen, X., Huang, L., Xie, W., Lin, Q., Li, X., Wong, K.-C.: Enabling single-cell drug response annotations from bulk rna-seq using scad. *Advanced Science* **10**(11), 2204113 (2023)
- [24] Fustero-Torre, C., Jiménez-Santos, M.J., García-Martín, S., Carretero-Puche, C., García-Jimeno, L., Ivanchuk, V., Di Domenico, T., Gómez-López, G., Al-Shahrour, F.: Beyondcell: targeting cancer therapeutic heterogeneity in single-cell

rna-seq data. *Genome Medicine* **13**(1), 187 (2021)

- [25] Li, C., Shao, X., Zhang, S., Wang, Y., Jin, K., Yang, P., Lu, X., Fan, X., Wang, Y.: scrank infers drug-responsive cell types from untreated scrna-seq data using a target-perturbed gene regulatory network. *Cell Reports Medicine* (2024)
- [26] Kim, T., Kim, C.: Attract, perturb, and explore: Learning a feature alignment network for semi-supervised domain adaptation. In: *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIV 16*, pp. 591–607 (2020). Springer
- [27] Li, J., Li, G., Shi, Y., Yu, Y.: Cross-domain adaptive clustering for semi-supervised domain adaptation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2505–2514 (2021)
- [28] Jiang, P., Wu, A., Han, Y., Shao, Y., Qi, M., Li, B.: Bidirectional adversarial training for semi-supervised domain adaptation. In: *IJCAI*, pp. 934–940 (2020)
- [29] Li, B., Wang, Y., Zhang, S., Li, D., Keutzer, K., Darrell, T., Zhao, H.: Learning invariant representations and risks for semi-supervised domain adaptation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1104–1113 (2021)
- [30] Singh, A.: Clda: Contrastive learning for semi-supervised domain adaptation. *Advances in Neural Information Processing Systems* **34**, 5089–5101 (2021)
- [31] Saito, K., Kim, D., Sclaroff, S., Darrell, T., Saenko, K.: Semi-supervised domain adaptation via minimax entropy. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 8050–8058 (2019)
- [32] Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572* (2014)
- [33] Efraimidis, P.S., Spirakis, P.G.: Weighted random sampling with a reservoir. *Information processing letters* **97**(5), 181–185 (2006)
- [34] McInnes, L., Healy, J., Melville, J.: Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426* (2018)
- [35] Mansoori, B., Mohammadi, A., Davudian, S., Shirjang, S., Baradaran, B.: The different mechanisms of cancer drug resistance: a brief review. *Advanced pharmaceutical bulletin* **7**(3), 339 (2017)
- [36] Ben-David, U., Siranosian, B., Ha, G., Tang, H., Oren, Y., Hinohara, K., Stratthdee, C.A., Dempster, J., Lyons, N.J., Burns, R., *et al.*: Genetic and transcriptional evolution alters cancer cell line drug response. *Nature* **560**(7718), 325–330 (2018)

- [37] Pellecchia, S., Viscido, G., Franchini, M., Gambardella, G.: Predicting drug response from single-cell expression profiles of tumours. *BMC medicine* **21**(1), 476 (2023)
- [38] Vogler, M.: Bcl2a1: the underdog in the bcl2 family. *Cell Death & Differentiation* **19**(1), 67–74 (2012)
- [39] Yang, C., Guo, L., Du, J., Zhang, Q., Zhang, L.: Spink1 overexpression correlates with hepatocellular carcinoma treatment resistance revealed by single cell rna-sequencing and spatial transcriptomics. *Biomolecules* **14**(3), 265 (2024)
- [40] Aissa, A.F., Islam, A.B., Ariss, M.M., Go, C.C., Rader, A.E., Conrardy, R.D., Gajda, A.M., Rubio-Perez, C., Valyi-Nagy, K., Pasquinelli, M., *et al.*: Single-cell transcriptional changes associated with drug tolerance and response to combination therapies in cancer. *Nature communications* **12**(1), 1628 (2021)
- [41] Song, T., Yu, W., Wu, S.-X.: Subsequent treatment choices for patients with acquired resistance to egfr-tkis in non-small cell lung cancer: restore after a drug holiday or switch to another egfr-tki? *Asian Pacific Journal of Cancer Prevention* **15**(1), 205–213 (2014)
- [42] Hänzelmann, S., Castelo, R., Guinney, J.: Gsva: gene set variation analysis for microarray and rna-seq data. *BMC bioinformatics* **14**, 1–15 (2013)
- [43] Sharma, A., Cao, E.Y., Kumar, V., Zhang, X., Leong, H.S., Wong, A.M.L., Ramakrishnan, N., Hakimullah, M., Teo, H.M.V., Chong, F.T., *et al.*: Longitudinal single-cell rna sequencing of patient-derived primary cells reveals drug-induced infidelity in stem cell hierarchy. *Nature communications* **9**(1), 4931 (2018)
- [44] Cao, J., Spielmann, M., Qiu, X., Huang, X., Ibrahim, D.M., Hill, A.J., Zhang, F., Mundlos, S., Christiansen, L., Steemers, F.J., *et al.*: The single-cell transcriptional landscape of mammalian organogenesis. *Nature* **566**(7745), 496–502 (2019)
- [45] Schnepf, P.M., Shelley, G., Dai, J., Wakim, N., Jiang, H., Mizokami, A., Keller, E.T.: Single-cell transcriptomics analysis identifies nuclear protein 1 as a regulator of docetaxel resistance in prostate cancer cells. *Molecular Cancer Research* **18**(9), 1290–1301 (2020)
- [46] Lundberg, S.M., Lee, S.-I.: A unified approach to interpreting model predictions. *Advances in neural information processing systems* **30** (2017)
- [47] Wang, M., Jia, J., Cui, Y., Peng, Y., Jiang, Y.: Molecular and clinical characterization of a novel prognostic and immunologic biomarker gpm3 in low-grade gliomas. *Brain Sciences* **11**(11), 1529 (2021)
- [48] Edfeldt, K., Hellman, P., Westin, G., Stalberg, P.: A plausible role for actin gamma smooth muscle 2 (actg2) in small intestinal neuroendocrine tumorigenesis.

BMC Endocrine Disorders **16**, 1–8 (2016)

- [49] Sundararajan, M., Taly, A., Yan, Q.: Axiomatic attribution for deep networks. In: International Conference on Machine Learning, pp. 3319–3328 (2017). PMLR
- [50] Miyato, T., Dai, A.M., Goodfellow, I.: Adversarial training methods for semi-supervised text classification. arXiv preprint arXiv:1605.07725 (2016)
- [51] Kokhlikyan, N., Miglani, V., Martin, M., Wang, E., Alsallakh, B., Reynolds, J., Melnikov, A., Kliushkina, N., Araya, C., Yan, S., et al.: Captum: A unified and generic model interpretability library for pytorch. arXiv preprint arXiv:2009.07896 (2020)