

---

# WISER: Weak Supervision and Supervised Representation Learning to Improve Drug Response Prediction in Cancer

---

Kumar Shubham<sup>1</sup> Aishwarya Jayagopal<sup>2</sup> Syed Mohammed Danish<sup>3</sup> Prathosh AP<sup>1</sup> Vaibhav Rajan<sup>2</sup>

## Abstract

Cancer, a leading cause of death globally, occurs due to genomic changes and manifests heterogeneously across patients. To advance research on personalized treatment strategies, the effectiveness of various drugs on cells derived from cancers (‘cell lines’) is experimentally determined in laboratory settings. Nevertheless, variations in the distribution of genomic data and drug responses between cell lines and humans arise due to biological and environmental differences. Moreover, while genomic profiles of many cancer patients are readily available, the scarcity of corresponding drug response data limits the ability to train machine learning models that can predict drug response in patients effectively. Recent cancer drug response prediction methods have largely followed the paradigm of unsupervised domain-invariant representation learning followed by a downstream drug response classification step. Introducing supervision in both stages is challenging due to heterogeneous patient response to drugs and limited drug response data. This paper addresses these challenges through a novel representation learning method in the first phase and weak supervision in the second. Experimental results on real patient data demonstrate the efficacy of our method WISER (Weak supervISion and supErvised Representation learning) over state-of-the-art alternatives on predicting personalized drug response. Our implementation is available at <https://github.com/kyrs/WISER>

## 1. Introduction

Cancer is a major cause of global morbidity and mortality (WHO, 2022). Cancer develops due to changes in our genome, which enable cancer cells to gain a selective advantage over healthy cells, resulting in uncontrolled proliferation as a cancerous tumour. Significant variability in treatment sensitivity and outcomes among patients makes cancer treatment difficult (Bedard et al., 2013). Hence, cancer care is transitioning from a ‘one-size-fits-all’ approach to a more personalized strategy guided by patient-specific genomic characteristics (Wahida et al., 2023).

To aid therapeutic development, there have been large-scale global efforts, e.g., through The Cancer Genome Atlas (TCGA) database (Hutter & Zenklusen, 2018), to catalog high-dimensional genomic information ( $\mathcal{X}$ ) of cancer patients. However, patient drug response data [ $\mathcal{Y}_t^{d_i}(\mathcal{X})$  for drug  $d_i$ ] is scarce due to limited number of patients, with only a few drugs administered on each patient (Sharifi-Noghabi et al., 2021). This has motivated researchers to explore preclinical datasets – e.g., *cell lines*, comprising cells extracted from patient cancers, which can be cloned in a way that the same genomic information is replicated across them. Such clones can be exposed to different drugs to obtain drug response information  $\mathcal{Y}_c^{d_i}(\mathcal{X})$  for multiple  $d_i$  on the same  $\mathcal{X}$ . This data is immensely useful and cannot be directly obtained from patients, who cannot be subjected to multiple drug regimens simultaneously. While such fine-grained drug response data is only available for a limited number of cell lines ( $\sim 1000$ ) and drugs, it provides a valuable starting point to build personalized drug response models based on genomic information.

However, previous studies have shown that such cell line-based response models do not accurately predict drug efficacy in patients due to several reasons (Seyhan, 2019). Cell line data ( $\mathcal{X}_c$ ) is more homogeneous than patient cancer cells ( $\mathcal{X}_t$ ) and the environments in which they reside are different. This results in differences in the distributions ( $P$ ) of genomic information across cell lines and patients ( $P(\mathcal{X}_c) \neq P(\mathcal{X}_t)$ ), and they can be considered as different domains (See Appendix B). Further, within the human body, in addition to the genomic structure, several other factors (e.g., the immune system) play a role in drug response. Thus,

---

<sup>1</sup>Department of Electrical Communication Engineering, Indian Institute of Science, Bengaluru, Karnataka, India <sup>2</sup>Department of Information Systems and Analytics, School of Computing, National University of Singapore, Singapore <sup>3</sup>Indian Institute of Technology, Patna, Bihar, India. Correspondence to: Kumar Shubham <shubhamkuma3@iisc.ac.in>.

*Proceedings of the 41<sup>st</sup> International Conference on Machine Learning*, Vienna, Austria. PMLR 235, 2024. Copyright 2024 by the author(s).

the drug response functions are different across cell lines and patients (i.e.,  $\mathcal{Y}_t^{d_i}(\cdot) \neq \mathcal{Y}_c^{d_i}(\cdot)$ ).

To address these challenges, several domain adaptation and transfer learning-based drug response models, that use a combination of cell line and patient data, have been developed. These methods generally consist of two stages: (1) an unsupervised representation learning phase where domain-invariant representations of genomic data are learned and, (2) a classification phase where these representations are used to train a drug response prediction model by categorizing responses as positive or negative based on the drug’s impact on inhibiting cancer growth. The classifier is trained using labeled data and used to predict drug response in patients.

Unsupervised representation learning approaches used by extant methods do not consider the drug response information ( $\mathcal{Y}_c^{d_i}(\mathcal{X}_c)$ ) associated with genomic profiles in cell lines, and hence do not distinguish between responders and non-responders to drugs. Supervised contrastive learning approaches (Barbano et al., 2022; Khosla et al., 2020; Graf et al., 2021; Hermans et al., 2017; Schroff et al., 2015; Lee et al., 2021) can address this by bringing the representations ( $\mathcal{Z}$ ) of data points with similar class labels closer together i.e.,  $\mathcal{Z}(\mathcal{X}_c^m) \sim \mathcal{Z}(\mathcal{X}_c^n)$  if  $\mathcal{Y}_c^{d_i}(\mathcal{X}_c^m) = \mathcal{Y}_c^{d_i}(\mathcal{X}_c^n)$ , emphasizing their shared characteristics over dissimilar classes. However, genomic profiles that respond to one drug may behave differently for another i.e, for drug  $d_i$ ,  $\mathcal{Y}_c^{d_i}(\mathcal{X}_c^m) = \mathcal{Y}_c^{d_i}(\mathcal{X}_c^n)$  but for drug  $d_k$ ,  $\mathcal{Y}_c^{d_k}(\mathcal{X}_c^m) \neq \mathcal{Y}_c^{d_k}(\mathcal{X}_c^n)$ . Hence, for drug  $d_i$ ,  $\mathcal{Z}(\mathcal{X}_c^m) \sim \mathcal{Z}(\mathcal{X}_c^n)$  but for drug  $d_k$ ,  $\mathcal{Z}(\mathcal{X}_c^m) \not\sim \mathcal{Z}(\mathcal{X}_c^n)$ . Further, limited patient data with documented drug response makes it difficult to find genomic profiles with similar efficacy across multiple drugs. These difficulties, in turn, limit the ability to use standard supervised contrastive learning methods to bring the representations of genomic profiles closer together. Our study addresses this challenge by learning a discrete representation per drug ( $\mathcal{R}$ ) (Van Den Oord et al., 2017) and representing each genomic profile as a weighted combination ( $\mathcal{Z} = \sum \mathcal{WR}$ ). To ensure that  $\mathcal{Z}$  is simultaneously reflective of the responses from multiple drugs, we increase the weights of drugs with positive response compared to those with negative response, through a supervised triplet loss (Hermans et al., 2017; Schroff et al., 2015; Barbano et al., 2022).

It is worth noting that while there is scarcity of labeled data in both domains, relatively abundant unlabeled patient data is available (See Appendix B). While prior studies have leveraged unlabeled patient data for learning domain-invariant representations, the training of drug response prediction classifiers has predominantly relied on the cell line dataset due to insufficient labeled response data for patients. Techniques like weak supervision can be employed to gener-

ate pseudo-labels for the abundant unlabeled data. However, naïvely using all pseudo-labeled samples does not improve performance (Lang et al., 2022; Shubham et al., 2023) (also seen in our results). In fact, there exists a trade-off between the noise introduced in the downstream classifier due to pseudo labels and the generalization it achieves when trained in a weak supervision setting (Lang et al., 2022). To address this, we introduce a subset selection step (Lang et al., 2022; Shubham et al., 2023), which to our knowledge is novel in this context and helps boost performance. We employ majority-vote-based weak supervision techniques (Ratner et al., 2017; Zhang et al., 2022) to create pseudo labels for patient genomic profiles without documented drug response, followed by a subset selection strategy (Muhlenbach et al., 2004). This subset is combined with labeled cell line data to train the drug response prediction classifier.

Our contributions can be summarized as follows:

- We design a new supervised domain-invariant representation learning approach which offers better distinction between drug responders and non-responders by addressing the challenges of limited sample size and heterogeneous drug response of genomic profiles.
- We propose a novel strategy that carefully selects a subset of least noisy pseudo-labeled patient data for classifier training on the domain-invariant representations.
- Using these techniques we propose a new method, called WISER, to estimate drug response for patients using unlabeled patient data and a small set of labeled cell line data.
- Our experiments on benchmark datasets demonstrate the superiority of WISER over state-of-the-art methods for drug response prediction, with improvements of up to 15.7% in AUROC.
- The most important features (genes) responsible for pseudo-labeling patient samples in the selected subsets correlates well with independent clinical evidence based on gene-drug interactions that impact patient survival, which further validates our pseudo-labeling approach.

## 2. Related Work

### 2.1. Drug Response Prediction

Prior literature on drug response prediction in patients has primarily focused on transfer learning (Pan & Yang, 2009). These approaches are useful when the target domain (patients) has limited samples, and a related source domain (cell lines) has more labeled samples. Transductive transfer

learning methods (Bousmalis et al., 2016; Sharifi-Noghabi et al., 2021; Sun & Saenko, 2016) use labeled source domain samples for drug response prediction but often build one model per drug and, thus, lack correlations across drugs. Inductive transfer learning methods (Sharifi-Noghabi et al., 2020; Ma et al., 2021) utilize few-shot and multi-task learning on available labeled patient samples but exhibit inferior performance to other approaches. Recent methods, like CODE-AE (He et al., 2022) learn shared representations using unlabeled genomic profiles from both domains.

Among the extant approaches, CODE-AE (He et al., 2022) has demonstrated superior predictive accuracy and robustness through extensive benchmark studies. CODE-AE is trained in two stages: (1) unsupervised pretraining of autoencoders to learn both domain-specific private and domain-invariant shared representations and (2) downstream drug response prediction based on the learned shared representations. A key shortcoming of this approach is that the representations learnt do not factor in the downstream drug response prediction task. Further, they do not utilize a large number of unlabeled patient genomic profiles in the downstream drug response prediction. Our proposed method, WISER, can handle these shortcomings and differs from CODE-AE in two aspects - (1) we incorporate drug response information of cell lines through supervised domain-invariant representation learning, and (2) we also utilize the available unlabeled patient genomic profiles through weak supervision techniques followed by subset selection.

## 2.2. Weak Supervision and Subset Selection

Weak supervision techniques (Ratner et al., 2016) are designed to address the challenge of limited data size. They leverage information from various sources (*label functions*), such as data from different domains (Mazzetto et al., 2021; Zhang et al., 2022), to generate cost-effective but noisy labels for unlabeled data. To further enhance the accuracy of the estimation process, confident predictions from different sources of pseudo labels are systematically combined, through weighing or voting schemes (Ratner et al., 2016; Dawid & Skene, 1979; Fu et al., 2020). For a smaller set of label functions, a majority vote-based scheme (Ratner et al., 2017) outperforms weighing techniques.

In addition, recent works (Lang et al., 2022; Shubham et al., 2023), in weak supervision, have shown that a subset of original data can generate optimal results compared to the use of the entire pseudo-labeled dataset. In fact, there exists a trade-off between the generalization achieved by the classifier and the noise introduced by the pseudo labels. Previous studies on subset selection have primarily concentrated on natural language tasks and employed pre-trained word embeddings (Kenton & Toutanova, 2019). However, the application of these on cancer research remains unexplored.

## 3. Proposed Method

### 3.1. Problem Formulation and Solution Overview

**Problem Definition:** Let us assume that there are  $\mathcal{N}_c$  labeled samples of genomic profiles associated with cell lines ( $\mathcal{G}_{cell}$ ) and  $\mathcal{N}_t$  unlabeled samples of genomic profiles from patients ( $\mathcal{G}_{patient}$ ). In this work, although we focus on gene expression profiles, our method can also be applied to other omics data types, such as mutations. In general,  $\mathcal{N}_c \ll \mathcal{N}_t$ . Let  $\{d_1, d_2 \dots d_n\}$  be the set of  $n$  drugs with documented drug response for  $\mathcal{G}_{cell}$  and  $\mathcal{Y}_c^{d_i}(\mathcal{X}_c^j)$  be the corresponding response of a drug ( $d_i$ ) to a genomic profile  $\mathcal{X}_c^j \in \mathcal{G}_{cell}$  and  $\mathcal{Y}_t^{d_i}(\mathcal{X}_t^m)$  be the drug response for patients  $\mathcal{X}_t^m \in \mathcal{G}_{patient}$ . Note that  $\mathcal{Y}_c^{d_i}(\mathcal{X}_c^j) \in \{-1, 0, 1\}$  where 1 indicates a positive response of a genomic profile  $\mathcal{X}_c^j$  to drug  $d_i$ , 0 indicates a negative response to the drug  $d_i$  and -1 represents that the response data is not available. The main objective of our work is to use the labeled cell line data ( $\mathcal{G}_{cell}, \mathcal{Y}_c^{d_i}$ ) for  $n$  drugs  $\{d_1, d_2 \dots d_n\}$  and the unlabeled patient genomic profile ( $\mathcal{G}_{patient}$ ) to estimate drug response for patients ( $\mathcal{Y}_t^{d_i}$ ). Further details about both domains are provided in Appendix B.

**Solution Overview:** Here, we describe the overview of our method comprising four major stages as depicted in Fig. 1.

**Stage 1: Representation Learning** In the first stage, we learn representations that are invariant between patient and cell line domains. Specifically, we learn discrete latent representations for individual drugs. The desired domain-invariant representation  $\mathcal{Z}$  is generated through a weighted combination of these drug representations.

**Stage 2: Weak Supervision** To incorporate the unlabeled patient genomic profiles in the training of the downstream drug response prediction model, we train multiple classifiers (*label functions*) using labeled cell line data and the domain invariant representation ( $\mathcal{Z}$ ). These label functions are then used to predict labels for the unlabeled patient dataset. The confident predictions from all label functions are combined based on majority-vote to assign the pseudo labels.

**Stage 3: Subset Selection** In this stage, we propose to utilize a subset of genomic profiles with confident predictions as indicated by the label functions. We employ cut statistics (Muhlenbach et al., 2004) in conjunction with the domain-invariant representation ( $\mathcal{Z}$ ) to select a subset of least noisy samples.

**Stage 4: Drug Response Prediction** We combine the subset of patient genomic profiles and associated pseudo labels, chosen after subset selection in Stage 3, with the labeled cell line genomic profiles to train a downstream drug response prediction classifier. This classifier can be used to infer drug responses in new patients.

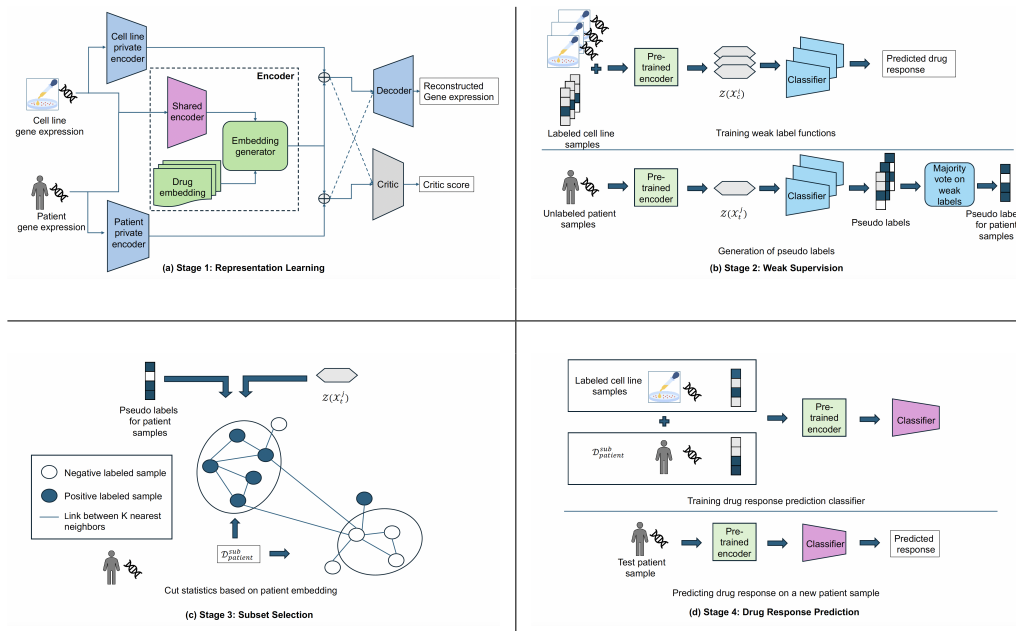


Figure 1. This diagram outlines WISER’s comprehensive training process, divided into four key phases. First, in the Representation Learning phase, a domain-invariant representation ( $\mathcal{Z}$ ) is learned between cell line and patient genomic profiles using a shared encoder and private encoding scheme. Next, in the Weak Supervision phase, multiple label functions are trained using labeled genomic profiles of cell lines to assign pseudo labels to unlabeled patient genomic profiles. Following that, in the Subset Selection phase, pseudo labels and the domain-invariant representation ( $\mathcal{Z}$ ) are used to select a subset of patient genomic profiles ( $\mathcal{D}_{patient}^{sub}$ ) and associated pseudo labels based on the consistency of the labels among nearest neighbors. Finally, in the Drug Response Prediction phase, the selected subset, along with labeled genomic profiles from cell lines, is utilized for downstream classifier training and predicting drug responses among patients.

### 3.2. Representation Learning

Genomic profile data collected from cell lines and patients exhibit distributional shifts owing to multiple confounding factors (He et al., 2022). This can cause a model trained using cell line data to not generalize to patients. In line with previous work, we address this using a private and shared encoder scheme, where a shared encoder ( $\mathcal{C}_S$ ) captures a domain invariant representation between the two domains while a private encoder ( $\mathcal{C}_P$ ) captures domain specific information. However, He et al. (2022) do not consider the drug response information ( $\mathcal{Y}_c^{d_i}(\mathcal{X}_c)$ ) during representation learning. We address this by representing the genomic profile ( $\mathcal{Z}$ ) as a weighted combination of drug embedding ( $\mathcal{R}$ ) (Eq. 1, Eq. 3) and used a triplet loss to learn these weights based on the drug efficacy results (Eq. 4).

In line with discrete representation learning methods (Lee et al., 2021), we leverage information on how a specific drug responds to a genomic profile, to generate a drug-specific discrete latent representation ( $\mathcal{R} = \{\mathcal{R}_{d_1}, \mathcal{R}_{d_2} \dots \mathcal{R}_{d_n}\}$ ). Similarly, inspired by contextual attention maps (Graves et al., 2014; Bahdanau et al., 2014), we combine the discrete representations of drugs ( $\mathcal{R}$ ) and the shared representation of genomic profiles ( $\mathcal{C}_S$ ) to form a new representation of the given genomic profile ( $\mathcal{Z}$ ). This new representation is a weighted sum of drug embeddings ( $\mathcal{R}$ ), with weights

( $\mathcal{W}$ ) indicating the efficacy of the different drugs on a given genomic profile. To obtain  $\mathcal{W}$ , we calculate the cosine similarity ( $sim(\cdot)$ ) between  $\mathcal{R}$  and  $\mathcal{C}_S$ .

The scores over different drugs are further normalized using a softmax function with an inverse temperature ( $\Delta$ ) to generate the weight  $\mathcal{W}$ . A weighted combination of  $\mathcal{R}$  using  $\mathcal{W}$  is used to generate  $\mathcal{Z}$ , as given in Eq. 1.

$$\mathcal{Z}(\mathcal{X}) = \sum_{i=1}^n \mathcal{W}_i(\mathcal{X}) \mathcal{R}_{d_i}$$

$$\mathcal{W}_i(\mathcal{X}) = \frac{\exp(\Delta * sim(\mathcal{C}_S(\mathcal{X}), \mathcal{R}_{d_i}))}{\sum_{j=1}^n \exp(\Delta * sim(\mathcal{C}_S(\mathcal{X}), \mathcal{R}_{d_j}))}$$

$$sim(\mathcal{C}_S(\mathcal{X}), \mathcal{R}_{d_i}) = \frac{\mathcal{C}_S(\mathcal{X})^T \mathcal{R}_{d_i}}{\|\mathcal{C}_S(\mathcal{X})\| \|\mathcal{R}_{d_i}\|} \quad (1)$$

For the training of our encoder models ( $\mathcal{C}_s, \mathcal{C}_p$ ) we concatenate ( $\oplus$ ) the weighted representation ( $\mathcal{Z}$  from Eq.1) and the private representation ( $\mathcal{C}_P = \{\mathcal{C}_P^t, \mathcal{C}_P^c\}$ ) of a genomic profile before passing it through a shared decoder ( $D$ ) for reconstruction ( $l_1$ ) in both the domains ( $\mathcal{X}_c^i \in \mathcal{G}_{cell}, \mathcal{X}_t^j \in \mathcal{G}_{patient}$ ), with the following reconstruction loss:

$$\mathbf{l}_{recon} = \frac{\sum_{i=1}^{\mathcal{N}_c} \mathbf{l}_1(\mathcal{X}_c^i, \mathcal{C}_P^c)}{\mathcal{N}_c} + \frac{\sum_{j=1}^{\mathcal{N}_t} \mathbf{l}_1(\mathcal{X}_t^j, \mathcal{C}_P^t)}{\mathcal{N}_t}$$

$$\text{Where, } \mathbf{l}_1(\mathcal{X}, \mathcal{C}_P) = \|D(\mathcal{Z}(\mathcal{X}) \oplus \mathcal{C}_P(\mathcal{X})) - \mathcal{X}\|^2 \quad (2)$$

To ensure that generated embedding ( $\mathcal{Z}$ ) and the private embedding ( $\mathcal{C}_P$ ) do not capture redundant information, we introduce an orthogonal loss (He et al., 2022) between these two embeddings as:  $\mathbf{l}_{ortho} = \|\mathcal{Z}(\mathcal{X}_c)^T \mathcal{C}_P^c(\mathcal{X}_c)\|^2 + \|\mathcal{Z}(\mathcal{X}_t)^T \mathcal{C}_P^t(\mathcal{X}_t)\|^2$ .

We further use the embedding loss  $\mathbf{l}_{embed}$  (Van Den Oord et al., 2017) to ensure that the generated embedding ( $\mathcal{Z}$ ) and the encoded genomic profiles ( $\mathcal{C}_S$ ) are closer to each other for both cell lines and patients. Eq. 3 illustrates this where  $sg(\cdot)$  denotes the stop gradient operator.

$$\mathbf{l}_{embed} = \frac{\sum_{i=1}^{\mathcal{N}_c} \mathbf{l}(\mathcal{X}_c^i)}{\mathcal{N}_c} + \frac{\sum_{j=1}^{\mathcal{N}_t} \mathbf{l}(\mathcal{X}_t^j)}{\mathcal{N}_t}$$

$$\mathbf{l}(\mathcal{X}) = \|\mathcal{Z}(\mathcal{X}) - sg(\mathcal{C}_S(\mathcal{X}))\|^2 + \|\mathcal{Z}(\mathcal{X}) - sg(\mathcal{Z}(\mathcal{X})) - \mathcal{C}_S(\mathcal{X})\|^2 \quad (3)$$

To ensure that the learnt representations reflect the drug efficacy on labeled genomic profiles, we rely on supervised triplet loss (Hermans et al., 2017; Schroff et al., 2015), which has a direct correspondence to modern supervised contrastive loss (Barbano et al., 2022). Triplet loss minimizes the distance between an anchor and positive labeled samples while maximizing the distance from negative labeled samples. In our formulation we use the cosine distance ( $dis(\cdot) = 1 - sim(\cdot)$  in Eq. 1), with the drug representation ( $\mathcal{R}$ ) as the anchor. The goal is to minimize ( $\mathbf{l}_{cns}$ ) the average distance of this anchor from the genomic representation with positive efficacy ( $\mathbf{s}^+$ ) and maximize its distance from the genomic representation with negative efficacy ( $\mathbf{s}^-$ ), (Eq. 4) where  $\mathbf{1}(\mathcal{Y}_c^{d_j}(\mathcal{X}_j^i) = 1)$  is an indicator function capturing the positive efficacy of drug ( $d_j$ ) on genomic profile ( $\mathcal{X}_j^i$ ) and  $\mathbf{1}(\mathcal{Y}_c^{d_j}(\mathcal{X}_j^i) = 0)$  captures the negative efficacy;  $\delta$  is the minimum offset between  $\mathbf{s}^+$  and  $\mathbf{s}^-$ .

$$\mathbf{l}_{cns} = \max(\mathbf{s}^+ - \mathbf{s}^- + \delta, 0)$$

$$\mathbf{s}^+ = \frac{\sum_{i=1}^{\mathcal{N}_c} \sum_{j=1}^n \mathbf{1}(\mathcal{Y}_c^{d_j}(\mathcal{X}_c^i) = 1) dis(\mathcal{C}_s(\mathcal{X}_c^i), \mathcal{R}_{d_j})}{\sum_{i=1}^{\mathcal{N}_c} \sum_{j=1}^n \mathbf{1}(\mathcal{Y}_c^{d_j}(\mathcal{X}_c^i) = 1)}$$

$$\mathbf{s}^- = \frac{\sum_{i=1}^{\mathcal{N}_c} \sum_{j=1}^n \mathbf{1}(\mathcal{Y}_c^{d_j}(\mathcal{X}_c^i) = 0) dis(\mathcal{C}_s(\mathcal{X}_c^i), \mathcal{R}_{d_j})}{\sum_{i=1}^{\mathcal{N}_c} \sum_{j=1}^n \mathbf{1}(\mathcal{Y}_c^{d_j}(\mathcal{X}_c^i) = 0)} \quad (4)$$

### 3.2.1. DOMAIN ADAPTATION

For efficient generalization of the downstream model, the generated representation should be invariant to the domain. Recent works (He et al., 2022; Ganin et al., 2016), have tried to generate such a representation using adversarial networks. Within this framework, a separate critic network is trained to distinguish between the embeddings from the two domains, while an encoder tries to generate indistinguishable embeddings for the critic. This additional training step ensures that as the training proceeds, an equilibrium is reached where the embedding is invariant for the critic network ( $\mathcal{F}$ ). In our work, we have used the Wasserstein GAN (WSGAN) (Arjovsky et al., 2017) with a gradient penalty-based adversarial loss (Gulrajani et al., 2017) to train our critic network (Eq. 5). The critic network takes as input a concatenation ( $\hat{\mathcal{C}}$ ) of generated embedding and private representation from both the domains.  $\mathbf{l}_{critic}$  tries to minimize the difference between the mean critic scores for patients  $\mathcal{F}(\hat{\mathcal{C}}(\mathcal{X}_t^j, \mathcal{C}_P^t))$  and cell lines  $\mathcal{F}(\hat{\mathcal{C}}(\mathcal{X}_c^i, \mathcal{C}_P^c))$ . In contrast, the patient representations are learnt to obtain a higher critic score ( $\mathbf{l}_{gen}$ ). A gradient penalty term is added, which encourages the gradient of the critic to have a norm close to 1 to maintain Lipschitz continuity (Arjovsky et al., 2017). These gradients are calculated on linear interpolate of input representation from both the domains ( $\mathcal{L}$ ), where  $\mathcal{L} = \epsilon \hat{\mathcal{C}}(\mathcal{X}_c, \mathcal{C}_P^c) + (1 - \epsilon) \hat{\mathcal{C}}(\mathcal{X}_t, \mathcal{C}_P^t)$  and  $\epsilon \sim U(0, 1)$ . Mathematically, the aforementioned loss functions are defined as follows:

$$\mathbf{l}_{critic} = \frac{1}{\mathcal{N}_t} \sum_{j=1}^{\mathcal{N}_t} \mathcal{F}(\hat{\mathcal{C}}(\mathcal{X}_t^j, \mathcal{C}_P^t)) - \frac{1}{\mathcal{N}_c} \sum_{i=1}^{\mathcal{N}_c} \mathcal{F}(\hat{\mathcal{C}}(\mathcal{X}_c^i, \mathcal{C}_P^c)) + \lambda (\|\nabla_{\mathcal{L}} \mathcal{F}(\mathcal{L})\| - 1)^2$$

$$\mathbf{l}_{gen} = -\frac{1}{\mathcal{N}_t} \sum_{i=1}^{\mathcal{N}_t} \mathcal{F}(\hat{\mathcal{C}}(\mathcal{X}_t^i, \mathcal{C}_P^t))$$

$$\hat{\mathcal{C}}(\mathcal{X}, \mathcal{C}_P) = \mathcal{Z}(\mathcal{X}) \oplus \mathcal{C}_P(\mathcal{X}) \quad (5)$$

The complete training occurs in two stages - first where the model is trained only using the loss ( $\mathbf{l}_{pl} = \mathbf{l}_{recon} + \mathbf{l}_{cns} + \mathbf{l}_{embed} + \mathbf{l}_{ortho}$ ) for a few epochs and later using  $\mathbf{l}_{total} = \mathbf{l}_{pl} + \mathbf{l}_{gen}$ , and  $\mathbf{l}_{critic}$  for the critic network.

### 3.3. Weak Supervision

Once we learn the domain invariant representations, they are subsequently employed to generate pseudo labels for the unlabeled genomic profile of patients. For this task, we partition the labeled cell line data into  $\mathcal{O}$  distinct subsets ( $\mathcal{D}_{cell}^i$   $i \in 1 \dots \mathcal{O}$ , where  $\mathcal{D}_{cell}^i \subset \mathcal{G}_{cell}$ ) and train a classifier ( $\mathcal{M}_i$ ) using their representations ( $\mathcal{Z}$ ). Each

individual classifier acts as a label function in our weak supervision framework and is utilized to infer the probability of drug response prediction for the genomic profile of patients ( $\mathcal{P}_i(y|\mathcal{X}_t^j)$ , where  $\mathcal{X}_t^j \in \mathcal{G}_{patient}$ ). The model assigns a label  $\hat{y} = 1$ , when the predicted drug response probability exceeds a threshold  $t^+$  and  $\hat{y} = 0$ , when the probability falls below a threshold  $t^-$ . For all intermediate probabilities where the confidence in model predictions is low, it abstains from assigning any class and labels the sample as -1 (Eq 6).

$$\mathcal{P}_i(y|\mathcal{X}_t^j) = \mathcal{M}_i(\mathcal{Z}(\mathcal{X}_t^j)) \text{ s.t. } i \in \{1 \dots \mathcal{O}\}, \mathcal{X}_t^j \in \mathcal{G}_{patient}$$

$$\hat{y}_i^j = \begin{cases} 1, & \text{if } \mathcal{P}_i(y|\mathcal{X}_t^j) > t^+ \\ 0, & \text{if } \mathcal{P}_i(y|\mathcal{X}_t^j) < t^- \\ -1, & \text{otherwise} \end{cases} \quad (6)$$

Samples with atleast one valid prediction (not abstained) from the label functions are used subsequently. The final pseudo label ( $\mathbf{y}_t^j$ ) for a given patient genomic profile ( $\mathcal{X}_t^j$ ) is decided by a majority vote across all non abstained predictions ( $\hat{y}_i^j$ ). The details are in Eq. 7, where ( $\mathbf{1}(\hat{y}_i^j = 1)$ ) and ( $\mathbf{1}(\hat{y}_i^j = 0)$ ) are indicator functions.

$$\mathbf{y}_t^j = \begin{cases} 1, & \text{if } \sum_{i=1}^{\mathcal{O}} \mathbf{1}(\hat{y}_i^j = 1) > \sum_{i=1}^{\mathcal{O}} \mathbf{1}(\hat{y}_i^j = 0) \\ 0, & \text{if } \sum_{i=1}^{\mathcal{O}} \mathbf{1}(\hat{y}_i^j = 1) \leq \sum_{i=1}^{\mathcal{O}} \mathbf{1}(\hat{y}_i^j = 0) \end{cases} \quad (7)$$

### 3.4. Subset selection and Drug Response Prediction

Once the pseudo labels have been assigned to the non-abstained patient genomic profiles, they can be directly used in conjunction with the labeled cell line data for the training of the drug response prediction classifier. However, recent works (Lang et al., 2022; Shubham et al., 2023) have shown that in a weak supervision setting, a complete set of non-abstained samples generates sub-optimal performance whereas considering a subset, improves performance.

In our work, we use cut statistics (Muhlenbach et al., 2004) to select a subset of the non-abstained dataset ( $\mathcal{V}$ ) by using the domain invariant representation ( $\mathcal{Z}$ ) and the pseudo labels ( $\mathbf{y}_t$ ) assigned to them. Each data sample ( $\mathcal{X}_t^i, \mathbf{y}_t^i$ ) where ( $\mathcal{X}_t^i \in \mathcal{V}$ ) is assigned a normalized Z score ( $z_i$ ) as explained below. For each patient ( $\mathcal{X}_t^i \in \mathcal{G}_{patient}$ ), we first find the nearest neighbors  $\text{NN}(\mathcal{X}_t^i) = \{\mathcal{X}_t^l : \text{where } (\mathcal{X}_t^l, \mathcal{X}_t^i) \text{ are } K \text{ nearest neighbors based on } L2 \text{ distance between } \mathcal{Z}(\mathcal{X}_t^l), \mathcal{Z}(\mathcal{X}_t^i)\}$ . A graph ( $G = (\mathcal{V}, \mathcal{E})$ ) is created with the number of nodes equal to the number of non-abstained patient genomic profile ( $\mathcal{V}$ ) and edges ( $\mathcal{E}$ ) defined as the nearest neighbor for each sample ( $\text{NN}(\mathcal{X}_t^i), \mathcal{X}_t^i \in \mathcal{V}$ ). For every edge in the graph a weight ( $w_{i,j}$ ) is assigned, so that samples with similar representation ( $\mathcal{Z}$ ) has higher weight compared to dissimilar ones i.e.,

$w_{i,j} = (1 + \|\mathcal{Z}(\mathcal{X}_t^i) - \mathcal{Z}(\mathcal{X}_t^j)\|)^{-1}$  where  $\mathcal{X}_t^j \in \text{NN}(\mathcal{X}_t^i)$ . In general a set of data points (sub-graph) with similar representation (higher  $w_{i,j}$ ) but sharing different pseudo labels are considered to be noisy and should not be considered for downstream training (Muhlenbach et al., 2004). Under given assumption, each sample  $\mathcal{X}_t^i$  is assigned a score  $\mathcal{J}^i$ , a sum of weights of samples sharing different class labels ( $\mathbf{1}(\mathbf{y}_t^i \neq \mathbf{y}_t^j)$ ) among the nearest neighbor. Further, under a null hypothesis of independent assignment of class labels with probability  $\mathcal{P}(\mathbf{y}_t)$  a Z-score ( $z_i$ ) is calculated for  $\mathcal{J}^i$  using the mean ( $\mu_i$ ) and variance ( $\sigma_i$ ) calculated according to Muhlenbach et al. (2004).  $\mathcal{P}(\mathbf{y}_t)$  is approximated by the bin counts of both positive and negative classes amongst the non-abstained samples. A smaller  $z_i$  signifies the consistency of class labels amongst the nearest neighbors and is an indicator of less noisy pseudo labels. In our work, the non-abstained patient data is sorted based on  $z_i$ , (Eq. 8) and the top  $b\%$  (also referred to as budget) is used to obtain  $\mathcal{D}_{patient}^{sub}$  which is then used in conjunction with labeled cell line data to train the final classifier for drug response prediction.

$$z_i = \frac{\mathcal{J}_i - \mu_i}{\sigma_i}$$

$$\mathcal{J}_i = \sum_{j \in \text{NN}(\mathcal{X}_t^i)} w_{i,j} \mathbf{1}(\mathbf{y}_t^i \neq \mathbf{y}_t^j)$$

$$\mu_i = [1 - \mathcal{P}(\mathbf{y}_t^i)] \sum_{j \in \text{NN}(\mathcal{X}_t^i)} w_{i,j}$$

$$\sigma_i^2 = \mathcal{P}(\mathbf{y}_t^i) [1 - \mathcal{P}(\mathbf{y}_t^i)] \sum_{j \in \text{NN}(\mathcal{X}_t^i)} w_{i,j}^2 \quad (8)$$

Algorithm 1 (In the appendix) describes the complete procedure of our method called WISER (Weak supervISion and supErvised Representation learning).

## 4. Experiments

### 4.1. Experiment Settings

We evaluate the proposed method in **Four** experimental settings - (1) Drug response prediction: In this task we compare different baselines by training a binary classifier to predict efficacy of a given drug on patients, (2) understanding the medical relevance of weak supervision and subset selection techniques in this context, (3) ablation study of the proposed method to compare the performance of the model with and without weak supervision and (4) measuring sensitivity of subset size on classification performance.

**Data** We have used the cancer cell lines and patient genomic profiles (comprising gene expression data from 1426 genes) as in CODE-AE (He et al., 2022). 677 labeled cancer cell line samples, from DepMap portal (Ghandi et al., 2019),

Table 1. Performance comparison of predicted patient response using AUROC and AUPRC metrics of our proposed method (WISER). Data related to clinical relapse is used for all the evaluations. The result is noted in the form of (mean / std) where the score has been obtained over five fold cross validation. The best performer among all baselines is reported in bold, while the predictions that were not meaningful are denoted by ‘-’. On an average, our method outperforms others baselines on all the drugs for at least one metric. The best performer is highlighted in bold.

Methods	5-Fluorouracil		Temozolomide		Sorafenib		Gemcitabine		Cisplatin	
	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC
<b>WISER</b>	0.715/0.036	<b>0.741/0.023</b>	<b>0.760/0.006</b>	<b>0.786/0.019</b>	<b>0.727/0.007</b>	0.728/0.024	<b>0.649/0.037</b>	<b>0.752/0.002</b>	<b>0.851/0.007</b>	<b>0.861/0.020</b>
CODE-AE	<b>0.868/0.030</b>	0.740/0.006	0.751/0.017	0.762/0.001	0.631/0.020	0.705/0.062	0.594/0.016	0.751/0.006	0.652/0.071	0.743/0.011
DAE	0.591/0.066	0.573/0.066	0.685/0.013	0.668/0.105	0.485/0.053	0.613/0.046	0.530/0.036	0.511/0.048	0.522/0.087	0.581/0.096
CORAL	0.578/0.015	0.651/0.135	0.675/0.020	0.654/0.020	0.491/0.023	0.616/0.048	0.597/0.030	0.544/0.037	0.617/0.072	0.617/0.124
VELODROME	0.598/0.054	0.403/0.000	0.701/0.028	0.668/0.000	0.505/0.029	<b>0.749/0.000</b>	0.547/0.030	0.434/0.000	0.583/0.029	0.442/0.000
ENET	0.435/0.092	0.454/0.070	-	-	-	-	-	-	0.637/0.076	0.623/0.045
TCRP	0.596/0.080	0.546/0.073	0.675/0.009	0.662/0.012	0.441/0.053	0.521/0.054	0.462/0.057	0.502/0.055	0.414/0.048	0.432/0.037
MLP	0.569/0.050	0.599/0.042	0.646/0.022	0.624/0.038	0.444/0.035	0.501/0.035	0.467/0.036	0.498/0.049	0.459/0.070	0.496/0.070
DSN-DANN	0.635/0.065	0.596/0.101	0.683/0.015	0.690/0.040	0.533/0.050	0.628/0.069	0.555/0.070	0.582/0.044	0.585/0.103	0.608/0.133
VAEN	0.633/0.157	0.585/0.100	0.648/0.035	0.632/0.162	0.600/0.021	0.668/0.112	0.526/0.087	0.618/0.223	0.694/0.049	0.698/0.065
COXEN	0.336/0.000	0.403/0.000	0.726/0.000	0.668/0.000	0.639/0.000	<b>0.749/0.000</b>	0.378/0.000	0.434/0.000	0.393/0.000	0.442/0.000
COXRF	0.562/0.070	0.598/0.063	0.388/0.080	0.451/0.031	0.418/0.072	0.505/0.044	0.506/0.078	0.506/0.037	0.554/0.074	0.564/0.065
CELLIGNER	0.536/0.060	0.531/0.024	-	-	-	-	0.575/0.029	0.529/0.053	0.497/0.042	0.550/0.033
ADAE	0.68/0.040	0.725/0.036	0.707/0.010	0.757/0.003	0.540/0.092	0.678/0.040	0.499/0.093	0.691/0.123	0.633/0.165	0.755/0.080
DSN-MMD	0.678/0.074	0.674/0.103	0.712/0.031	0.759/0.051	0.515/0.036	0.582/0.090	0.465/0.041	0.491/0.069	0.650/0.023	0.605/0.067

and 9808 unsupervised patient samples from TCGA (Hutter & Zenklusen, 2018) were used. 179 samples of labeled TCGA genomic profiles were used for evaluation. Drug response in cell lines was based on z-score calculated on Area Under the Dose Response Curve (AUDRC) scores. Cell lines with a z-score less than 0 were considered positive respondents and greater than 0 as negative respondents to the drug. For patients, the assessment relied on cancer relapse time post-chemotherapy, categorizing values greater than the median as positive respondents and those less than median as negative respondents. The specifics of data preprocessing and related details are available in He et al. (2022). A set of 20 drugs present in both DepMap and TCGA ( $\mathcal{D} = \{d_1, d_2, \dots, d_{20}\}$ ), were considered for the experiment. Details of drugs are provided in Appendix C. Due to the limited number of labeled patient genomic profiles, the evaluation was done only on 5-Fluorouracil (Fu), Temozolomide (Tem), Sorafenib (Sor), Gemcitabine (Gem) and Cisplatin (Cis), with drug responses available in atleast 20 patients.

**Model Configuration** The encoder and decoder networks, used in representation learning, consist of two linear layers of the neural network. The hidden units associated with the encoder and decoder are (512, 256) and (256, 512) dimensions respectively. Both networks use ReLU based activation units. The critic network and the classifier (used for weak supervision and downstream drug response prediction) consist of two layers of neural network with (64, 32) dimensions of hidden unit with ReLU activation for the first layer. The critic network uses linear layer as final activation, while the classifier uses a sigmoid layer. Same architecture has been used for all the baseline methods for fair comparison. Further details about training and hyper parameter tuning is provided in Appendix C.

**Baselines** We have compared our method with CODE-AE (He et al., 2022), VAEN (Jia et al., 2021) and DAE (Vincent et al., 2008). Further the proposed method is compared with domain adaptation techniques like Celligner (Warren et al., 2021), Velodrome (Sharifi-Noghabi et al., 2021), Deep CORAL (Sun & Saenko, 2016) and DSN (MMD and DANN variant) (Bousmalis et al., 2016). Recent methods like ADAE (Dincer et al., 2020), COXEN + Random Forest (COXRF) and COXEN (Zhu et al., 2020) were also included for comparison. To compare with algorithms which do not use representation learning, the results of TCRP (Ma et al., 2021), MLP (Sakellaropoulos et al., 2019) and ElasticNet (Kuenzi et al., 2020) were also included.

**Metrics** For comparison, we have used area under the receiver operating characteristics (AUROC) and area under the precision-recall curve (AUPRC) scores (He et al., 2022). The classifier used for drug response prediction was trained using 5-fold stratified validation data of cell line and tested on patient data from TCGA.

## 4.2. Results

### 4.2.1. DRUG RESPONSE PREDICTION

Table 1 shows a performance comparison of our method with other baselines. Our method (WISER) exhibits superior performance in terms of AUROC scores for Cisplatin, Temozolomide, Gemcitabine, and Sorafenib, surpassing baselines by 15.7%, 0.9%, 5.2% and 8.8% respectively while for AUPRC score it shows an enhancement of 0.1%, 2.4%, 0.1% and 10.6% for 5-Fluorouracil, Temozolomide, Gemcitabine and Cisplatin respectively. Comparison with other traditional methods, self supervised learning based losses and another dataset are provided in Appendix E, Appendix G and Appendix H respectively.




Figure 2. Ablation on weak supervision and sensitivity test on subset size over the performance of the model.

Table 2. Experiment to examine the medical relevance of weak supervision and subset selection. In the given experiment, the set of genes with significant drug-gene interaction (P-val<0.05) for the survival of patients with cancer, from GDISC, is compared with the genes considered relevant by weak supervision and subset selection. The precision and recall between the two sets is reported.

Drug	Fu	Tem	Sor	Gem	Cis
Gene (P-val<0.05)	418	706	-	526	831
Gene (P-val≥0.05)	521	456	-	473	143
Precision	0.419	0.609	-	0.499	0.860
Recall	0.459	0.500	-	0.464	0.503

### 4.2.2. MEDICAL RELEVANCE OF THE METHOD

We next examine the medical relevance of the pseudo labels and the subset selected for the downstream prediction, generated using the best hyperparameters. We identify the genes most relevant in the generation of these pseudo labels, through the feature selection procedure of Extra-Trees Classifier (Alfian et al., 2022). This is done by fitting an Extra-Trees Classifier model on the patient genomic profiles (from the selected subset) and their pseudo labels, and selecting genes with top 50% feature importance. We compare the selected genes against the GDISC database (Spainhour et al., 2017), which has independently identified genes associated with chemotherapy response in TCGA. The authors provide information for 22 drugs, however, details for Sor were unavailable. The resulting set of significant genes and the corresponding overlap are highlighted in Table 2. Drugs with relevant information were evaluated on two metrics (1) **Precision**: This measures the ratio of genes marked as significant by GDISC among all the genes selected by the Extra-Trees Classifier. (2) **Recall**: This assesses the ratio of genes selected by the Extra-Trees Classifier among the entire set of genes marked as significant by the GDISC database for a given drug. Cisplatin, Temozolomide, Gemcitabine, 5-Fluorouracil achieve a precision of 0.860, 0.609, 0.499, 0.419 respectively. Similarly, the recall achieved by these drugs are 0.503, 0.500, 0.464, 0.459. This score correlates with the performance of the drugs in Table 1, where a higher precision and recall generate better AUROC and AUPRC, thus suggesting the faithfulness (Alvarez Melis & Jaakkola, 2018) of our explanations in terms of gene importances.

### 4.2.3. ABLATION STUDIES

We conducted an ablation test on the effect of weak supervision and subset selection, by directly using the representa-

tions of labeled cell line samples for the downstream drug response prediction. The results (Figure 2) were compared for the best hyperparameter configuration of each drug. The results indicate that weak supervision and subset selection (WISER) improve AUROC by 4.58% and AUPRC by 3.4% on average. Further ablation on the proposed loss functions are provided in Appendix F.

### 4.2.4. SENSITIVITY ANALYSIS

Since the ablation study indicates the importance of weak supervision and subset selection, we next examine the impact of the subset budget (b) on the overall performance. This test was performed by varying b while maintaining the optimal configuration for the remaining parameters. Figure 2 summarizes the result of the experiment. For AUROC, the subset selection setting generated better results for 5-Fluorouracil (b=40%), Cisplatin (b=20%) and Gemcitabine (b=20%) than the complete non-abstained dataset (b=100%). An improvement of 10.2%, 2.4% and 7.7% are seen in these 3 drugs respectively. For AUPRC, the subset selection setting generated better results for all drugs other than Temozolomide, with budget b set to 20%, 10%, 10%, 80% for Cisplatin, Gemcitabine, Sorafenib and 5-Fluorouracil respectively. An improvement of 0.8%, 28.2%, 1.8% and 6.9% was observed for these drugs respectively. It can be seen that using subset selection leads to optimal performance compared to the complete non-abstained dataset. Sensitivity analysis on other hyperparameters are provided in Appendix D.

## 5. Conclusion

Recent cancer drug response prediction methods have largely followed the paradigm of unsupervised domain-invariant representation learning followed by a downstream drug response classification step. Although supervised training could improve performance, doing so was limited by the heterogeneity in patient responses across drugs and limited availability of labeled patient data. Our approach, addresses these challenges by modeling genomic profiles as a combination of discrete drug representations, reflective of heterogeneous drug responses. We also use weak supervision and subset selection on unlabeled patient genomic profiles to improve generalization of the classifier. WISER demonstrates improved drug response prediction for several clinically sig-



nificant anti-cancer drugs. To the best of our knowledge, our method is the first to use domain-invariant representation for subset selection with weak supervision, and can be applied to similar settings with large unlabeled datasets. However, the performance of our method is limited by the available labeled dataset and the set of drugs considered for discrete representation learning. Future work can explore further improvements of our approach through other sources of distant supervision, e.g., through knowledge graphs.

## Impact Statement

This research seeks to enhance the effectiveness of personalized cancer treatment by integrating laboratory data and patient information, thereby bridging gaps between research and real-world outcomes. The study tackles the scarcity of labeled patient data through the use of weak supervision techniques, aiming to contribute to the improvement of reliable and accessible personalized cancer treatments.

## Acknowledgment

This research/project is supported (in part for setting up the GPU compute) by the Indian Institute of Science through a start-up grant and by the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG-100E-2023-116). Prathosh AP is supported by the Infosys Foundation Young investigator award. Kumar Shubham is supported by the Kotak IISc AI-ML Centre (KIAC) fellowship. Aishwarya Jayagopal is supported by the National University of Singapore Research Scholarship.

We want to thank Pranav Sastry for his feedback and assistance in the project.

## References

- Alfian, G., Syafrudin, M., Fahrurrozi, I., Fitriyani, N. L., Atmaji, F. T. D., Widodo, T., Bahiyah, N., Benes, F., and Rhee, J. Predicting breast cancer from risk factors using svm and extra-trees-based feature selection method. *Computers*, 11(9):136, 2022.
- Alsagaf, I., Buchan, D., and Wan, C. Improving cell type identification with gaussian noise-augmented single-cell rna-seq contrastive learning. *Briefings in Functional Genomics*, pp. elad059, 2024.
- Alvarez Melis, D. and Jaakkola, T. Towards robust interpretability with self-explaining neural networks. *Advances in neural information processing systems*, 31, 2018.
- Arjovsky, M., Chintala, S., and Bottou, L. Wasserstein generative adversarial networks. In *International conference on machine learning*, pp. 214–223. PMLR, 2017.
- Bahdanau, D., Cho, K., and Bengio, Y. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- Barbano, C. A., Dufumier, B., Tartaglione, E., Grangetto, M., and Gori, P. Unbiased supervised contrastive learning. *arXiv preprint arXiv:2211.05568*, 2022.
- Bedard, P. L., Hansen, A. R., Ratain, M. J., and Siu, L. L. Tumour heterogeneity in the clinic. *Nature*, 501(7467): 355–364, 2013.
- Bousmalis, K., Trigeorgis, G., Silberman, N., Krishnan, D., and Erhan, D. Domain separation networks. *Advances in neural information processing systems*, 29, 2016.
- Bruna, A., Rueda, O. M., Greenwood, W., Batra, A. S., Callari, M., Batra, R. N., Pogrebniak, K., Sandoval, J., Cassidy, J. W., Tufegdzcic-Vidakovic, A., et al. A biobank of breast cancer explants with preserved intra-tumor heterogeneity to screen anticancer compounds. *Cell*, 167(1): 260–274, 2016.
- Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pp. 1597–1607. PMLR, 2020.
- Dawid, A. P. and Skene, A. M. Maximum likelihood estimation of observer error-rates using the em algorithm. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 28(1):20–28, 1979.
- Dincer, A. B., Janizek, J. D., and Lee, S.-I. Adversarial deconfounding autoencoder for learning robust gene expression embeddings. *Bioinformatics*, 36(Supplement\_2): i573–i582, 2020.
- Fu, D., Chen, M., Sala, F., Hooper, S., Fatahalian, K., and Ré, C. Fast and three-rious: Speeding up weak supervision with triplet methods. In *ICML*, pp. 3280–3291. PMLR, 2020.
- Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., and Lempitsky, V. Domain-adversarial training of neural networks. *The journal of machine learning research*, 17(1):2096–2030, 2016.
- 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.
- Graf, F., Hofer, C., Niethammer, M., and Kwitt, R. Dissecting supervised contrastive learning. In *International Conference on Machine Learning*, pp. 3821–3830. PMLR, 2021.

- Graves, A., Wayne, G., and Danihelka, I. Neural Turing machines. arXiv preprint arXiv:1410.5401, 2014.
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. C. Improved training of Wasserstein GANs. Advances in neural information processing systems, 30, 2017.
- He, D., Liu, Q., Wu, Y., and 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.
- Hermans, A., Beyer, L., and Leibe, B. In defense of the triplet loss for person re-identification. arXiv preprint arXiv:1703.07737, 2017.
- Hinton, G. E. and Zemel, R. Autoencoders, minimum description length and Helmholtz free energy. Advances in neural information processing systems, 6, 1993.
- Hutter, C. and Zenklus, J. C. The cancer genome atlas: creating lasting value beyond its data. Cell, 173(2):283–285, 2018.
- Jia, P., Hu, R., Pei, G., Dai, Y., Wang, Y.-Y., and Zhao, Z. Deep generative neural network for accurate drug response imputation. Nature communications, 12(1):1740, 2021.
- Kenton, J. D. M.-W. C. and Toutanova, L. K. Bert: Pre-training of deep bidirectional transformers for language understanding. In Proceedings of NAAACL-HLT, volume 1, pp. 2, 2019.
- Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C., and Krishnan, D. Supervised contrastive learning. Advances in neural information processing systems, 33:18661–18673, 2020.
- Kingma, D. P. and Welling, M. Auto-encoding variational Bayes. arXiv preprint arXiv:1312.6114, 2013.
- Kuenzi, B. M., Park, J., Fong, S. H., Sanchez, K. S., Lee, J., Kreisberg, J. F., Ma, J., and Ideker, T. Predicting drug response and synergy using a deep learning model of human cancer cells. Cancer cell, 38(5):672–684, 2020.
- Lang, H., Vijayaraghavan, A., and Sontag, D. Training subset selection for weak supervision. Advances in Neural Information Processing Systems, 35:16023–16036, 2022.
- Lee, H. H., Tang, Y., Yang, Q., Yu, X., Bao, S., Landman, B. A., and Huo, Y. Attention-guided supervised contrastive learning for semantic segmentation. arXiv preprint arXiv:2106.01596, 2021.
- 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.
- Mazzeo, A., Cousins, C., Sam, D., Bach, S. H., and Upfal, E. Adversarial multi class learning under weak supervision with performance guarantees. In International Conference on Machine Learning, pp. 7534–7543. PMLR, 2021.
- Muhlenbach, F., Lallich, S., and Zighed, D. A. Identifying and handling mislabelled instances. Journal of Intelligent Information Systems, 22(1):89–109, 2004.
- Pan, S. J. and Yang, Q. A survey on transfer learning. IEEE Transactions on knowledge and data engineering, 22(10):1345–1359, 2009.
- Ratner, A., Bach, S. H., Ehrenberg, H., Fries, J., Wu, S., and Ré, C. Snorkel: Rapid training data creation with weak supervision. In Proceedings of the VLDB Endowment. International Conference on Very Large Data Bases, volume 11, pp. 269. NIH Public Access, 2017.
- Ratner, A. J., De Sa, C. M., Wu, S., Selsam, D., and Ré, C. Data programming: Creating large training sets, quickly. NeurIPS, 29, 2016.
- Sakellaropoulos, T., Vougas, K., Narang, S., Koinis, F., Kotsinas, A., Polyzos, A., Moss, T. J., Piha-Paul, S., Zhou, H., Kardala, E., et al. A deep learning framework for predicting response to therapy in cancer. Cell reports, 29(11):3367–3373, 2019.
- Schroff, F., Kalenichenko, D., and Philbin, J. Facenet: A unified embedding for face recognition and clustering. In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 815–823, 2015.
- Seyhan, A. A. Lost in translation: the valley of death across preclinical and clinical divide—identification of problems and overcoming obstacles. Translational Medicine Communications, 4(1):1–19, 2019.
- Sharifi-Noghabi, H., Peng, S., Zolotareva, O., Collins, C. C., and Ester, M. Aitl: adversarial inductive transfer learning with input and output space adaptation for pharmacogenomics. Bioinformatics, 36(Supplement\_1):i380–i388, 2020.
- Sharifi-Noghabi, H., Harjandi, P. A., Zolotareva, O., Collins, C. C., and Ester, M. Out-of-distribution generalization from labelled and unlabelled gene expression data for drug response prediction. Nature Machine Intelligence, 3(11):962–972, 2021.

- Shubham, K., Sastry, P., and AP, P. Fusing conditional sub-modular gan and programmatic weak supervision. arXiv preprint arXiv:2312.10366, 2023.
- Spainhour, J. C. G., Lim, J., and Qiu, P. Gdisc: a web portal for integrative analysis of gene–drug interaction for survival in cancer. Bioinformatics, 33(9):1426–1428, 2017.
- Sun, B. and Saenko, K. Deep coral: Correlation alignment for deep domain adaptation. In Computer Vision–ECCV 2016 Workshops: Amsterdam, The Netherlands, October 8-10 and 15-16, 2016, Proceedings, Part III 14, pp. 443–450. Springer, 2016.
- Van Den Oord, A., Vinyals, O., et al. Neural discrete representation learning. Advances in neural information processing systems, 30, 2017.
- Vincent, P., Larochelle, H., Bengio, Y., and Manzagol, P.-A. Extracting and composing robust features with denoising autoencoders. In Proceedings of the 25th international conference on Machine learning, pp. 1096–1103, 2008.
- Wahida, A., Buschhorn, L., Fröhling, S., Jost, P. J., Schneeweiss, A., Lichter, P., and Kurzrock, R. The coming decade in precision oncology: six riddles. Nature Reviews Cancer, 23(1):43–54, 2023.
- Warren, A., Chen, Y., Jones, A., Shibue, T., Hahn, W. C., Boehm, J. S., Vazquez, F., Tsherniak, A., and McFarland, J. M. Global computational alignment of tumor and cell line transcriptional profiles. Nature Communications, 12(1):22, 2021.
- WHO. Cancer. <https://www.who.int/news-room/fact-sheets/detail/cancer>, February 2022.
- Zbontar, J., Jing, L., Misra, I., LeCun, Y., and Deny, S. Barlow twins: Self-supervised learning via redundancy reduction. In International conference on machine learning, pp. 12310–12320. PMLR, 2021.
- Zhang, J., Hsieh, C.-Y., Yu, Y., Zhang, C., and Ratner, A. A survey on programmatic weak supervision. arXiv preprint arXiv:2202.05433, 2022.
- Zhu, Y., Brettin, T., Evrard, Y. A., Xia, F., Partin, A., Shukla, M., Yoo, H., Doroshov, J. H., and Stevens, R. L. Enhanced co-expression extrapolation (coxen) gene selection method for building anti-cancer drug response prediction models. Genes, 11(9):1070, 2020.

## A. Appendix

### Algorithm 1 WISER: Weak supervISion and supERvised Representation learning to improve drug response prediction in cancer

**Require:** Genomic profile for cell line ( $\mathcal{X}_c$ ), genomic profile for patients ( $\mathcal{X}_t$ ), epoch for initial training ( $\mathcal{P}_i$ ), epoch for domain adaptation based training ( $\mathcal{P}_d$ ), epoch for critic training ( $\mathcal{P}_c$ ), batch size ( $\mathcal{B}$ ), weak supervision thresholds  $t^+$  and  $t^-$ , number of chunks  $\mathcal{O}$  of cell line data for training, subset selection budget  $b$  and nearest neighbor size  $K$ .

- 1: **##Representation Learning**
- 2: **for** epoch in  $[0 \dots \mathcal{P}_i]$  : **do**
- 3:   Sample batch of cell line and patient genomic data from the dataloader without replacement.  $\{\mathcal{X}_c^{(i)}\}_{i=0}^{\mathcal{B}} \{\mathcal{X}_t^{(j)}\}_{j=0}^{\mathcal{B}}$
- 4:   Train shared encoder ( $\mathcal{C}_S$ ), private encoder ( $\mathcal{C}_P^c, \mathcal{C}_P^t$ ), discrete embedding ( $\mathcal{R}$ ), and decoder ( $\mathcal{D}$ ) with sampled batch using ( $\mathcal{L}_{pt}$ ) loss.
- 5: **end for**
- 6:  $\mathcal{N} = 0$
- 7: **for** epoch in  $[0 \dots \mathcal{P}_d]$  : **do**
- 8:   Sample batch of cell line and patient genomic data from the dataloader without replacement.  $\{\mathcal{X}_c^{(i)}\}_{i=0}^{\mathcal{B}} \{\mathcal{X}_t^{(j)}\}_{j=0}^{\mathcal{B}}$
- 9:   Train the critic network ( $\mathcal{F}$ ) with ( $\mathcal{L}_{critic}$ ) loss.
- 10:    $\mathcal{N} += 1$
- 11:   **if**  $\mathcal{N} \% \mathcal{P}_c == 0$  **then**
- 12:     Sample batch of genomic data  $\{\mathcal{X}_c^{(i)}\}_{i=0}^{\mathcal{B}} \{\mathcal{X}_t^{(j)}\}_{j=0}^{\mathcal{B}}$
- 13:     Train  $\mathcal{C}_S, \mathcal{C}_P^c, \mathcal{C}_P^t, \mathcal{R}, \mathcal{D}$  with sampled batch using ( $\mathcal{L}_{total}$ )
- 14:     **end if**
- 15:   **end for**
- 16:   Use the representation ( $\mathcal{Z}$ ) generated by shared encoder and drug-based embeddings ( $\mathcal{C}_S$  and  $\mathcal{R}$ ).
- 17: **## Weak Supervision**
- 18: **for**  $i$  in  $[1 \dots \mathcal{O}]$ : **do**
- 19:   Train a classifier  $\mathcal{M}_i$  using  $D_{cell}^i$  where  $D_{cell}^i \subset G_{cell}$ .
- 20:   Infer  $\mathcal{P}_i(\mathbf{y} | \mathcal{X}_t^j)$  using the trained classifier  $\mathcal{M}_i$ , where  $\mathcal{X}_t^j \in G_{patient}$ .
- 21: **end for**
- 22:   Label samples based on  $t^+$  and  $t^-$  (Eq. 6).
- 23:   Assign the final pseudo label ( $\mathbf{y}_i^j$ ), based on Majority Voting strategy, for non-abstained samples (Eq. 7).
- 24: **## Subset Selection and Drug Response Prediction**
- 25:   Calculate  $\mathbf{z}_i$  for non-abstained patient samples as in Eq. 8, sort by  $\mathbf{z}_i$  and choose top  $b$  % as the subset.
- 26:   Use the patient genomic profiles associated with this subset, along with their pseudo labels, in conjunction with  $\mathcal{X}_c$  to train a drug response prediction classifier.

## B. Distinction Between the Two Domains and Details about Dataset

Table 3 provides details of the two domains under consideration in our study. The cell line domain is notable for its abundant labeled responses to diverse drugs, whereas the patient data predominantly comprises unlabeled samples. For our experiments, we selected 20 drugs which were administered in both patients and cell lines. To evaluate our approach on patients, we considered 5 drugs with a documented response in at least 20 patients (Table 4).

Table 3. Details about the two domains in cancer drug response prediction.

Domains	Unlabeled data	Labeled data	Drug response label	Number of drugs with response	Number of drugs selected in our experiments
Cell line	1305	686	Evaluated using Z-score computed on AUDRC scores. (1) Z-score less than 0 considered as positive respondents. (2) Z-score greater than 0 considered as negative respondents.	449	20
Patients	9808	179	Cancer relapse time post-chemotherapy (1) Values greater than the median considered positive respondents. (2) Values less than the median considered negative respondents.	78	5

Table 4. Distribution of testing dataset.

Drug	5-Fluorouracil	Temozolomide	Gemcitabine	Cisplatin	Sorafenib
TCGA samples	21	46	46	40	26

### C. Training Details and Hyperparameter

The training of the model happens in four stages as mentioned in Algorithm 1. For representation learning, a grid search was performed on the initial training epoch ( $\mathcal{P}_i$ ), domain adaptation epoch ( $\mathcal{P}_d$ ) and inverse temperature value ( $\Delta$ ). The value considered for the experiments were  $[50, 100, 300]$ ,  $[1000, 2000, 2500, 3000]$  and  $[0.001, 0.1, 1, 2, 2.5, 10, 100]$  respectively. A set of 20 drugs were used for representation learning ( $\mathcal{R}$ ) namely 5-Fluorouracil, Gemcitabine, Temozolomide, Cisplatin, Sorafenib, Sunitinib, Doxorubicin, Tamoxifen, Paclitaxel, Carmustine, Cetuximab, Methotrexate, Topotecan, Erlotinib, Irinotecan, Bicalutamide, Temozolomide, Oxaliplatin, Docetaxel, Etoposide. For weak supervision, 5 label functions were trained on 5 different chunks ( $\mathcal{O}$ ) of labeled cell line dataset. The number of chunk for training was decided based on previous works (He et al., 2022; Ratner et al., 2017), considering the limited number of labeled cell line data and optimal performance of majority vote for less than 10 (Lfs). The value of ( $t^+$  and  $t^-$ ) were determined based on grid search over  $[(0.7, 0.3), (0.55, 0.49), (0.51, 0.46)]$  respectively. Similar experiments were performed on the median score of the predicted probabilities of label functions. For subset selection,  $K=20$  was considered in line with previous work (Lang et al., 2022). The optimal value of subset size ( $b$ ) was determined by a grid search over  $[0.2, 0.4, 0.5, 0.6, 0.8, 1]$ . All the experiments were done on NVIDIA A6000 Graphic card with 20 core and 160 GB memory. Our code is available at <https://github.com/kyrs/WISER>.

### D. Experiment Without Weak Supervision

Table 5. Performance comparison of predicted patient response using AUROC and AUPRC metrics of our proposed method without weak supervision (WISER(w\o WS)) with other transfer learning based approaches. Data related to clinical relapse is used for all the evaluations. The result is noted in the form of (mean / std) where the score has been obtained over five fold cross validation. The best performer among all baselines is reported in bold, while the predictions that were not meaningful are denoted by '-'. On an average, our method performs the best on 3 out of 5 drugs for atleast one metric. The best performer is highlighted in bold.

Methods	5-Fluorouracil		Temozolomide		Sorafenib		Gemcitabine		Cisplatin	
	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC
WISER- w\o WS	0.687/0.029	0.688/0.047	0.694/0.020	0.725/0.014	<b>0.683/0.021</b>	0.725/0.094	<b>0.617/0.013</b>	<b>0.752/0.005</b>	<b>0.792/0.038</b>	<b>0.808/0.035</b>
CODE-AE	<b>0.868/0.030</b>	<b>0.740/0.006</b>	<b>0.751/0.017</b>	<b>0.762/0.001</b>	0.631/0.020	0.705/0.062	0.594/0.016	0.751/0.006	0.652/0.071	0.743/0.011
TCRP	0.596/0.080	0.546/0.073	0.675/0.009	0.662/0.012	0.441/0.053	0.521/0.054	0.462/0.057	0.502/0.055	0.414/0.048	0.432/0.037
CORAL	0.578/0.015	0.651/0.135	0.675/0.020	0.654/0.020	0.491/0.023	0.616/0.048	0.597/0.030	0.544/0.037	0.617/0.072	0.617/0.124
VELODROME	0.598/0.054	0.403/0.000	0.701/0.028	0.668/0.000	0.505/0.029	<b>0.749/0.000</b>	0.547/0.030	0.434/0.000	0.583/0.029	0.442/0.000

To assess the impact of incorporating labeled drug response data for representation learning, we conducted a separate experiment with our method, excluding weak supervision and subset selection. The results were compared with other transfer learning-based approaches, and the findings are presented in Table 5. Our method (WISER w\o WS) demonstrated superior performance in three out of five drugs, showcasing improved AUROC and AUPRC metrics. Specifically, it exhibited gains of approximately 14%, 2.3%, and 5.2% in AUROC for Cisplatin, Gemcitabine, and Sorafenib, respectively. Additionally, there were gains of 6.5% and 0.1% in AUPRC for Cisplatin and Gemcitabine, respectively.

#### D.1. Sensitivity Test on Hyperparameter

To analyze the impact of different hyperparameters on representation learning, we have done a sensitivity analysis of initial training epoch ( $\mathcal{P}_i$ ), adversarial training ( $\mathcal{P}_d$ ) and temperature ( $\Delta$ ) on AUROC and AUPRC performance while using best configuration for other hyper parameters.

##### D.1.1. INITIAL TRAINING EPOCH

Table 6 and Table 7 shows the impact of initial training epoch on the performance of the model. As per the result the Cisplatin, Gemcitabine, Temozolomide, Sorafenib, and 5-Fluorouracil achieves best AUROC score for 50, 300, 300, 50, 50 epoch respectively and best AUPRC score for 50, 100, 300, 100, 100 respectively. In general, different drugs performs differently for this hyper parameter, where training for less number of epochs is favourable for Cisplatin while training for

more iteration is favoured in Temozolomide.

Table 6. Sensitivity analysis of the initial training epoch ( $\mathcal{P}_i$ ) on AUROC scores.

Drug / Epoch	50	100	300
Fu	<b>0.687/0.027</b>	0.642/0.043	0.650/0.034
Sor	<b>0.683/0.019</b>	0.636/0.047	0.441/0.043
Tem	0.558/0.004	0.573/0.024	<b>0.694/0.017</b>
Gem	0.499/0.097	0.440/0.120	<b>0.617/0.011</b>
Cis	<b>0.792/0.034</b>	0.527/0.045	0.529/0.120

Table 7. Sensitivity analysis of the initial training epoch ( $\mathcal{P}_i$ ) on AUPRC scores.

Drug / Epoch	50	100	300
Fu	0.626/0.069	<b>0.688/0.042</b>	0.575/0.022
Sor	0.431/0.100	<b>0.725/0.085</b>	0.566/0.065
Tem	0.547/0.029	0.575/0.024	<b>0.725/0.014</b>
Gem	0.75/0.0	<b>0.752/0.005</b>	0.75/0.0
Cis	<b>0.808/0.031</b>	0.575/0.057	0.608/0.081

#### D.1.2. ADVERSARIAL TRAINING EPOCH

Next we analyze the impact of adversarial training epoch on the performance of the model. Table 8 and Table 9 shows the result of given experiment. Where, Cisplatin, Gemcitabine, Temozolomide, Sorafenib, and 5-Fluorouracil achieves best AUROC score for 2500, 2000, 1000, 2000, 2500 respectively and best AUPRC score for 2000, 1000,1000, 2000, 2000 respectively. In general, we see an impact of domain adaptation on the performance of the model, where in drugs like Cisplatin best results are generated for larger number of training epochs.

Table 8. Sensitivity analysis of the domain adversarial training epoch ( $\mathcal{P}_d$ ) on AUROC scores.

Drug / Epoch	1000	2000	2500
Fu	0.509/0.063	0.680/0.013	<b>0.687/0.026</b>
Sor	0.452/0.061	<b>0.683/0.019</b>	0.586/0.027
Tem	<b>0.694/0.018</b>	0.681/0.010	0.664/0.031
Gem	0.518/0.082	<b>0.617/0.011</b>	0.517/0.111
Cis	0.391/0.011	0.654/0.085	<b>0.792/0.033</b>

Table 9. Sensitivity analysis of the domain adversarial training epoch ( $\mathcal{P}_d$ ) on AUPRC scores.

Drug / Epoch	1000	2000	2500
Fu	0.589/0.108	<b>0.688/0.042</b>	0.656/0.0348
Sor	0.523/0.113	<b>0.725/0.085</b>	0.565/0.101
Tem	<b>0.725/0.013</b>	0.683/0.035	0.659/0.048
Gem	<b>0.750/0.000</b>	-	<b>0.750/0.000</b>
Cis	0.525/0.016	0.710/0.058	<b>0.808/0.031</b>

#### D.1.3. INVERSE TEMPERATURE

We have further conducted a sensitivity test to analyze the impact of the inverse temperature ( $\Delta$ ) on the performance of the model. Table 10 and Table 11 summarizes the result of the given experiment. Based on the presented outcome, we found an influence of inverse temperature on the performance of all the drugs. Inverse temperature controls the weights associated with the drug embeddings ( $\mathcal{R}$ ). For 5-Fluorouracil and Temozolomide best AUROC and AUPRC scores are generated for

$\Delta = 10$ , Cisplatin generates best results for  $\Delta = 2.5$  while for Gemcitabine, Sorafenib best AUPRC score were generated for  $\Delta = 0.001$  and AUROC score for 0.1 and 2.5 respectively.

Table 10. Sensitivity analysis of the Inverse Temperature ( $\Delta$ ) on AUROC scores.

Drug / Inv temp	0.01	0.1	1	2	2.5	10
Fu	0.487/0.094	0.483/0.231	0.500/0.079	0.486/0.031	0.478/0.091	<b>0.687/0.026</b>
Tem	0.492/0.07	0.534/0.103	0.366/0.016	0.445/0.035	0.489/0.071	<b>0.694/0.017</b>
Gem	0.465/0.035	<b>0.617/0.011</b>	0.390/0.015	0.324/0.02	0.330/0.014	0.365/0.016
Sor	0.483/0.141	0.456/0.127	0.473/0.045	0.638/0.06	<b>0.638/0.051</b>	0.586/0.027
Cis	0.526/0.057	0.558/0.120	0.541/0.047	0.502/0.043	<b>0.791/0.033</b>	0.485/0.15

Table 11. Sensitivity analysis of the Inverse Temperature ( $\Delta$ ) on AUPRC scores.

Drug / Inv temp	0.01	0.1	1	2	2.5	10
Fu	0.601/0.09	0.542/0.177	0.469/0.060	0.499/0.057	0.440/0.0491	<b>0.688/0.042</b>
Tem	0.561/0.079	0.533/0.098	0.416/0.017	0.490/0.019	0.486/0.066	<b>0.725/0.013</b>
Gem	<b>0.670/0.086</b>	0.481/0.072	0.572/0.078	0.424/0.027	0.402/0.115	0.431/0.01
Sor	<b>0.724/0.084</b>	0.609/0.162	0.479/0.024	0.553/0.083	0.526/0.042	0.554/0.021
Cis	0.603/0.054	0.598/0.120	0.664/0.034	0.626/0.047	<b>0.808/0.031</b>	0.501/0.021

### E. Analysis of Different Components

We have further analyzed the importance of different components in our method. For this, we have compared WISER with (1) WISER- w\o WS : A derivative of our work with only supervised discrete representation learning module and does not use weak supervision. (2) Code-AE : Code-AE is the closest baseline to our work which does not use weak supervision and supervised discrete representation learning. (3) Next we compare our method with other representation learning based methods without domain adaptation module i.e., Variational autoencoder (Kingma & Welling, 2013) (VAE), autoencoder (Hinton & Zemel, 1993) (AE). (4) Finally we compare our method with random forest a standard model without neural network (RF). As per the result WISER performs optimal for all the drugs on atleast one metric.

Table 12. Performance comparison of predicted patient response using AUROC and AUPRC metrics of our proposed method (WISER). Data related to clinical relapse is used for all the evaluations. The result is noted in the form of (mean / std) where the score has been obtained over five fold cross validation. On an average, our method outperforms others baselines on all the drugs for at least one metric. The best performer is highlighted in bold.

Methods	5-Fluorouracil		Temozolomide		Sorafenib		Gemcitabine		Cisplatin	
	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC
WISER	0.715/0.036	<b>0.741/0.023</b>	<b>0.760/0.006</b>	<b>0.786/0.019</b>	<b>0.727/0.007</b>	<b>0.728/0.024</b>	<b>0.649/0.037</b>	<b>0.752/0.002</b>	<b>0.851/0.007</b>	<b>0.861/0.020</b>
WISER- w\o WS	0.687/0.029	0.688/0.047	0.694/0.020	0.725/0.014	0.683/0.021	0.725/0.094	0.617/0.013	0.752/0.005	0.792/0.038	0.808/0.035
CODE-AE	<b>0.868/0.030</b>	0.740/0.006	0.751/0.017	0.762/0.001	0.631/0.020	0.705/0.062	0.594/0.016	0.751/0.006	0.652/0.071	0.743/0.011
VAE	0.636/0.032	0.616/0.067	0.671/0.023	0.688/0.020	0.472/0.023	0.554/0.024	0.514/0.090	0.484/0.048	0.552/0.103	0.631/0.034
AE	0.636/0.019	0.576/0.046	0.659/0.048	0.610/0.030	0.528/0.061	0.597/0.101	0.553/0.029	0.553/0.050	0.623/0.042	0.607/0.067
RF	0.565/0.100	0.595/0.099	0.632/0.03	0.619/0.048	0.366/0.131	0.482/0.103	0.452/0.026	0.480/0.023	0.470/0.062	0.473/0.044

### F. Ablation for Discrete Representation

We have extended our analysis by performing ablation studies to assess the significance of the proposed loss functions used in training of domain invariant representation(Z). For this experiment we successively removed triplet loss  $l_{cns}$  and discrete representation loss  $l_{embed}$  while using the best hyperparameter for the model. The experiments were performed without weak supervision and subset selection strategy in the downstream drug response prediction task, so as to understand the effects of the loss terms on Z in isolation (indicated by 'Wiser(Z)'). The results (mean/std. over 5-fold cross validation) of this experiment are provided in Table 13.

Table 13. An ablation study to investigate the impact of different loss functions on AUROC and AUPRC metrics for discrete embedding. Data related to clinical relapse is used for all the evaluations. The result is noted in the form of (mean / std) where the score has been obtained over five fold cross validation. On an average, our method outperforms others baselines on all the drugs for at least one metric. The best performer is highlighted in **bold**.

Methods	5-Fluorouracil		Temozolomide		Sorafenib		Gemcitabine		Cisplatin	
	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC
WISER(Z)	<b>0.687/0.029</b>	0.688/0.047	<b>0.694/0.020</b>	<b>0.725/0.014</b>	<b>0.683/0.021</b>	<b>0.725/0.094</b>	<b>0.617/0.013</b>	<b>0.752/0.005</b>	<b>0.792/0.038</b>	<b>0.808/0.035</b>
WISER(Z) w/o $l_{cns}$	0.593/0.041	<b>0.745/0.047</b>	0.664/0.013	0.675/0.024	0.568/0.027	0.603/0.082	0.510/0.079	0.551/0.114	0.736/0.017	0.754/0.021
WISER(Z) w/o $\{l_{cns}, l_{embed}\}$	0.545/0.032	0.483/0.023	0.617/0.018	0.601/0.027	0.487/0.024	0.540/0.045	0.450/0.050	0.422/0.026	0.417/0.058	0.520/0.013

### G. Comparison with Self Supervised Learning Methods

We have further compared our method with other self supervised learning approaches. Based on prior literature (Alsaggaf et al., 2024), we use Gaussian noise based perturbation of the genomic samples for the domain-invariant representations. On the augmented data, we apply various SSL methods like SimCLR (Chen et al., 2020) and Barlow Twins (Zbontar et al., 2021) on CODE-AE. The results (mean/std. over 5-fold cross validation) is provided in Table 14.

Table 14. Performance comparison of predicted patient response using AUROC and AUPRC metrics of our proposed method (WISER) with other self supervised learning approaches. Data related to clinical relapse is used for all the evaluations. The result is noted in the form of (mean / std) where the score has been obtained over five fold cross validation. On an average, our method outperforms others baselines on all the drugs for at least one metric. The best performer is highlighted in **bold**.

Methods	5-Fluorouracil		Temozolomide		Sorafenib		Gemcitabine		Cisplatin	
	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC
WISER	0.715/0.036	0.741/0.023	<b>0.760/0.006</b>	<b>0.786/0.019</b>	<b>0.727/0.007</b>	<b>0.728/0.024</b>	<b>0.649/0.037</b>	<b>0.752/0.002</b>	<b>0.851/0.007</b>	<b>0.861/0.020</b>
CODE-AE	<b>0.868/0.030</b>	0.740/0.006	0.751/0.017	0.762/0.001	0.631/0.020	0.705/0.062	0.594/0.016	0.751/0.006	0.652/0.071	0.743/0.011
CODE-AE + SimCLR	0.663/0.051	0.699/0.129	0.707/0.007	0.733/0.024	0.479/0.024	0.610/0.027	0.490/0.029	0.609/0.027	0.469/0.05	0.518/0.070
CODE-AE + Barlow Twins	0.747/0.029	<b>0.767/0.048</b>	0.680/0.015	0.681/0.036	0.569/0.029	0.576/0.033	0.621/0.052	0.621/0.036	0.670/0.116	0.702/0.028

### H. Generalization on Other Drugs and Datasets

To establish the generalizability of proposed method in two aspects - (1) on unknown drugs (unseen during representation learning) and (2) on a different dataset, we conducted similar experiments on the PDTTC breast cancer dataset (Bruna et al., 2016) (32 samples per drug), on drugs unused in TCGA. The results(mean/std. over 5-fold cross validation) are shown in Table 15.

Table 15. Performance comparison of predicted response using AUROC and AUPRC metrics of our proposed method (WISER) on PDTTC dataset. The result is noted in the form of (mean / std) where the score has been obtained over five fold cross validation. On an average, our method outperforms others baselines on two out of three drugs. The best performer is highlighted in **bold**.

Methods	Az628		Gefitinib		Axitinib	
	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC
WISER	<b>0.792/0.069</b>	<b>0.789/0.029</b>	<b>0.700/0.025</b>	<b>0.793/0.031</b>	<b>0.864/0.011</b>	0.836/0.037
CODE-AE	0.754/0.097	0.679/0.148	0.613/0.037	0.778/0.009	0.840/0.049	0.762/0.033
Velodrome	0.513/0.015	0.625/0.064	0.446/0.091	0.495/0.052	0.786/0.041	<b>0.841/0.014</b>