

Weakly Supervised Set-Consistency Learning Improves Morphological Profiling of Single-Cell Images

Heming Yao* Phil Hanslovsky Jan-Christian Huetter
Burkhard Hoeckendorf David Richmond*
Biology Research | AI Development (BRAID), gCS, Genentech
* {yao.heming, richmond.david}@gene.com

Abstract

Optical Pooled Screening (OPS) is a powerful tool combining high-content microscopy with genetic engineering to investigate gene function in disease. The characterization of high-content images remains an active area of research and is currently undergoing rapid innovation through the application of self-supervised learning and vision transformers. In this study, we propose a set-level consistency learning algorithm, Set-DINO, that combines self-supervised learning with weak supervision to improve learned representations of perturbation effects in single-cell images. Our method leverages the replicate structure of OPS experiments (i.e., cells undergoing the same genetic perturbation, both within and across batches) as a form of weak supervision. We conduct extensive experiments on a large-scale OPS dataset with more than 5000 genetic perturbations, and demonstrate that Set-DINO helps mitigate the impact of confounders and encodes more biologically meaningful information. In particular, Set-DINO recalls known biological relationships with higher accuracy compared to commonly used methods for morphological profiling, suggesting that it can generate more reliable insights from drug target discovery campaigns leveraging OPS.

1. Introduction

High-content imaging combined with quantitative image analysis can be used to characterize cellular responses to genetic and chemical perturbations, and provides a powerful platform for target and drug discovery [4, 10, 34]. Despite the prevalence of this approach in the pharmaceutical industry, arrayed screening still suffers from limitations due to the high cost of scaling to large genetic and chemical libraries. Recently, Optical Pooled Screening (OPS), has been proposed as a cost-effective method for conducting high-content genetic screens at the whole-genome level [14, 28, 31, 35]. However, one caveat of pooled screens is

that cellular phenotypes are captured at the single-cell level, in contrast to arrayed screens, where fields of hundreds of cells receive the same treatment. This necessitates new research into methods for capturing cellular representations under perturbation that are robust to the high degree of noise and variability present in single-cell data. [5, 15, 31].

CellProfiler remains one of the most widely used tools for extracting expert-defined features (also referred to as “engineered” features) from high-content images [7, 32]. However, recent studies have focused on training deep learning algorithms, such as weakly supervised learning [3, 27], generative modeling [5] and self-supervised learning (SSL) [13, 23, 24, 31], to extract learned representations from high-content images. Among those methods, DINO [6] has emerged as a promising technique for extracting information-rich representations, and outperformed other approaches in a recent head-to-head comparison [13, 23].

While SSL frameworks are powerful feature extractors, they are unfortunately susceptible to learning unwanted confounding factors [13, 23]. Such factors can include plate-to-plate variation, well-position effect, and experimental conditions, all of which can influence image intensity, contrast, and texture. Despite efforts toward optimal experimental design, confounding factors remain a persistent challenge in high-content screening. Weakly supervised DINO [11, 17] was proposed to address the sensitivity of SSL to confounders by sampling image pairs across experimental batches and thus encouraging DINO to learn batch-invariant representations. This approach has been shown to improve the quality of learned representations on downstream biological tasks [11, 17, 18]; however, it has not yet been applied in the setting of single-cell images from optical pooled screens.

In this study, we develop an SSL framework explicitly for single-cell images. Inspired by [11, 17], we adopt a cross-batch sampling strategy in an attempt to learn representations that are invariant to confounders. However, we observed that DINO training with cross-batch sampling collapses due to the strong cell-to-cell variation exhibited

in single-cell data. To address this challenge, we combine cross-batch sampling with set-level representation to characterize cell populations undergoing a specific perturbation.

Our main contributions in this study include:

1. We propose weakly supervised set-consistency learning (Set-DINO), a novel representation learning framework designed specifically for single-cell images from optical pooled screens. To the best of our knowledge, this is the first time that set-level representation has been combined with DINO to facilitate self-supervised representation learning on noisy samples.
2. We apply Set-DINO to a large-scale OPS dataset of more than 5000 essential genes where it achieves a significantly better performance compared to both engineered features and the standard DINO framework. Through extensive experiments, we demonstrate that Set-DINO leverages the weak supervision provided by cell replicates to extract cell representations that are both less sensitive to confounding factors and contain more biologically meaningful information.

2. Related Work

2.1. Deep learning for high-content images

Deep learning has been extensively applied to the task of morphological profiling of high-content images. For example, weakly supervised learning using perturbation labels was applied to images from arrayed screening in [27], yielding improved performance at identifying treatments with the same mechanism of action (MoA) or genetic pathway, as compared to engineered features. However, leveraging perturbation labels as weak supervision may potentially result in learning spurious representations that falsely discriminate between different genetic perturbations with similar morphological effects. Not to mention the many genetic perturbations that have negligible effects on cell morphology.

Self-supervised consistency learning provides an elegant solution to this problem, because it doesn't assume "negative" relationships between samples from different perturbations. Furthermore, SSL and Vision Transformers (ViT) have recently achieved state-of-the-art performance in learning representations from natural images that can generalize to downstream tasks [6, 19]. Sivanandan *et al.* [31] and Doron *et al.* [13] applied similar methods to high-content images, and demonstrated that embeddings from DINO with ViT led to improved performance. Masked auto-encoders (MAE) have also been shown to outperform weakly supervised baselines in uncovering biological relationships [24]. Further, in [23], several SSL techniques including SimCLR, DINO, and MAE were compared and DINO embeddings achieved the best performance in terms of reproducibility and target prediction for compound per-

turbations.

2.2. Removing nuisance in morphological profiles

Batch effect is a central challenge in learning biological meaningful representations from high-content imaging data. Numerous batch correction techniques have been developed and successfully applied to mitigate unwanted variations in morphological profiles [13, 18, 23]. However, those methods may not fully benefit deep learning approaches, as they are typically applied as a post-processing step. Sypetkowski *et al.* applied adaptive batch normalization to normalize features during training, using statistics from individual batches [33]. Their proposed method mitigated batch effect and helped the model generalize to unseen batches. Inspired by this work, we explore an image normalization method based on image statistics from control cells.

Furthermore, numerous methods leverage the replicate structure of high-content screening data to learn invariance to batch effects. For example, Cross-Zamirski *et al.* proposed a weakly supervised DINO model (WS-DINO) to incorporate the treatment labels in arrayed screening, and demonstrated improved performance in MoA prediction [11]. Similarly, cross-domain consistency learning was proposed by Haslum *et al.* with additional loss terms to force the model to disregard batch-specific signals [17]. Those methods follow a rationale similar to supervised contrastive learning [22], where weak labels improve the robustness and informativeness of learned representations. Our study further validates the effectiveness of training SSL algorithms with weak labels, and extends this finding to single-cell images, requiring the use of a set-level loss. Our results indicate that while both weakly supervised learning and SSL approaches have their respective limitations, their combination can improve representations for high-content imaging data.

2.3. Modeling population heterogeneity

When analyzing high-content imaging data, profiles of biological replicates are typically aggregated to represent the average or median response of a cell population to each perturbation [2]. Existing research also suggests that higher-order statistics such as the dispersion and covariance of features may provide additional information and improve performance on downstream tasks [29]. By leveraging a set-level implementation, we benefit from a smoother, population-level loss, while retaining single-cell level profiles.

Deep Sets is a popular technique that offers a general framework for extracting representations from sets of objects [36]. The concept of set-level representation has been successfully integrated with SimCLR to improve unsupervised meta-learning performance on natural images [25].

Dijk *et al.* applied Deep Sets to pre-extracted single-cell profiles from CellProfiler, learning an aggregation function that down-weights noisy cells [12]. The aggregation function was trained using weakly supervised contrastive learning, and the resulting profiles prioritize biological signals over batch effects. In this study, instead of using hand-crafted features, we explore the possibility of combining set-level representation with end-to-end training of weakly supervised DINO to facilitate representation learning directly from raw images.

3. Methods

3.1. Dataset

We use a publicly released large-scale OPS dataset profiling CRISPR knockout of 5072 essential genes on cultured human cells [15]. Four guide RNA (sgRNA) sequences were used per gene target, and an additional 250 non-targeting sgRNAs were used as negative controls. The entire sgRNA library was delivered to a pool of cells, and the experiment was replicated across 46 wells from 8 plates (each plate has at most 6 wells).

In total, the dataset contains around 32 million cells, with a median of 6,000 cells per gene perturbation. The dataset was released with raw 4-channel images (stained for DNA, DNA damage, F-actin, and tubulin), metadata including the sgRNA that each cell received, and precomputed morphological features. The released features are normalized by the median and median absolute deviation of non-targeting controls (NTCs) within the same well.

For the model training and evaluation, we divide the data from 8 plates into a training set (6 plates with 28 wells), a validation set (1 plate with 6 wells), and a test set (2 plates with 12 wells). The model is trained on the training set and the checkpoints and hyper-parameters are selected on the validation set. The test set is exclusively used to evaluate the model’s performance and generalizability on unseen data. In this study, we regard each well as one experimental *batch*. The median number of cells with the same sgRNA in each batch is 29.

3.2. Image preprocessing

We follow the established practice for preprocessing of high-content images [1, 20, 23]. Images are flat-field corrected, and intensity values are clipped at the 0.1 and 99.9 percentiles, and then linearly re-scaled to [0, 1]. Single-cell images are then cropped using the cell centroids provided with the released metadata, and applying a 96-pixel-by-96-pixel bounding box around each cell.

We evaluate two methods for normalizing the single-cell images. The first method is image-wise z-score normalization (referred to as *z-score*), a common method where pixel intensities are image-wise and channel-wise normalized by

z-score for each single-cell image [20]. The second method is image normalization using statistics of the NTCs from the corresponding batch (referred to as *NTC z-score*). In this approach, each single-cell image is normalized by the channel-wise mean and standard deviation of pixel intensities from all NTCs in the corresponding batch. We compare the performance of both normalization methods in our results section.

3.3. Set-DINO framework

Similar to the standard DINO [6], Set-DINO consists of a student branch and a teacher branch (Figure 1). In the standard DINO framework, a single image with different augmentations is fed into the student and teacher branches and the model is trained to maximize the similarity between the embeddings from the two branches.

In Set-DINO, we sample a set of n single-cell images from cells receiving perturbation p in batch b : $X_{p,b} = \{x_{p,b}^1, \dots, x_{p,b}^n\}$. The i th tensor $x_{p,b}^i \in \mathbb{R}^{C,H,W}$ represents a multi-channel image of one cell receiving perturbation p from batch b , where C, H, W denote the number of channels, height, and width of the image, respectively. Similarly, we sample a second set of n single-cell images from cells receiving the same perturbation p from a different batch b' : $X_{p,b'} = \{x_{p,b'}^1, \dots, x_{p,b'}^n\}$.

The image sets $X_{p,b}$ and $X_{p,b'}$ are fed into the student network ϕ_s and teacher network ϕ_t , respectively, to generate single-cell latent embeddings. Then, an aggregation layer Λ aggregates the image-level embeddings to set-level embeddings:

$$\pi_{p,b} = \Lambda(\phi_s(x_{p,b}^1), \dots, \phi_s(x_{p,b}^n)), \quad (1)$$

$$\pi_{p,b'} = \Lambda(\phi_t(x_{p,b'}^1), \dots, \phi_t(x_{p,b'}^n)). \quad (2)$$

The embeddings $\pi_{p,b}$ and $\pi_{p,b'}$ represent a consensus of the cell populations receiving perturbation p in batches b , and b' , respectively. They are used as a pair of views whose similarity is optimized during training of the student network ϕ_s :

$$\mathcal{L} = H(\gamma(\pi_{p,b}), \gamma(\pi_{p,b'})), \quad (3)$$

where γ is a multi-layer perceptron (MLP), and H is the cross-entropy loss.

Similar to the standard DINO framework, a stop-gradient (SG) operator is applied on the teacher network ϕ_t and the parameters in the teacher network are updated with an exponential moving average (EMA) of the student parameters. We use a ViT backbone for ϕ_s and ϕ_t , and a 3-layer MLP for γ . The aggregation function Λ can be any function that is invariant to permutations such as feature statistics [25], Deep Sets [36], or self-attention layers [26]. In this study, we use the arithmetic mean because of its simplicity and effectiveness [25].

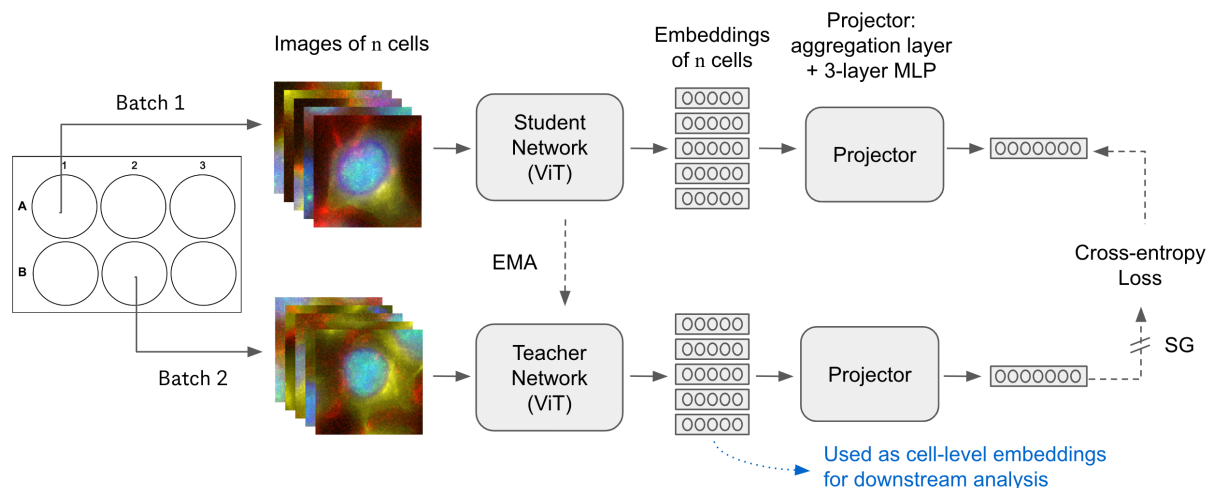


Figure 1. **Overview of the Set-DINO framework.** The inputs are two sets of single-cell images undergoing the same perturbation in different batches. Each 4-channel image is processed individually by the Vision Transformer (ViT) to generate a set of embeddings. The projector consists of an aggregation layer, followed by three fully-connected layers. The resulting consensus embeddings from the student and teacher branches are used to calculate the cross-entropy loss to train the model. After the model is trained, the single-cell image embeddings from ViT are used as the cell-level morphological features. SG: stop-gradient, EMA: exponential moving average.

The motivation of the set-level representation is to better characterize the cell population within a specific condition, while retaining cell-to-cell variation. Moreover, we found that the set-level representation is necessary for stabilizing DINO training when applying a cross-batch sampling strategy, due to the large degree of variation in single-cell images.

The cross-batch sampling strategy can be regarded as a form of data augmentation using biological replicates. Compared to common image augmentation techniques such as rotations and Gaussian blur, the utilization of cell replicates from different batches serves as a more biologically meaningful form of augmentation. However, despite having received the same genetic perturbation, single cells sampled from different batches may demonstrate very different morphology due to variations in cell states, cell cycle, and batch-level technical variations. This is especially true considering that the perturbation effects from many genetic perturbations can be extremely subtle [27]. Moreover, due to the varying effectiveness of the guide RNA, some cells may “escape” the perturbation, and exhibit a morphology similar to NTCs. Consequently, contrasting two single-cell images sampled from different batches may make the network insensitive to both batch effects and biologically meaningful information. Our experiments in this study demonstrate that without set-level aggregation, the model collapses, as excessively strong data augmentation forces the model to extract very general features that are identical for all cells in the dataset. To address this problem, we create views from each experimental batch using a set of cells receiving the same perturbation. The two sets of cells are assumed to

contain similar distributions in cell states.

In this study, we explore different cell sampling strategies. Given that every gene target in our dataset has four sgRNAs, we compare sampling cells receiving the same sgRNA versus those receiving perturbation of the same gene target. As an ablation study, we also experiment with sampling sets of cells from the same experimental batch.

3.4. Implementation

In Set-DINO, the data loader samples cellular images based on perturbation labels. To build one mini-batch, we initially select N_P perturbations, followed by sampling a pair of batches for each perturbation. Subsequently, for all cells in a given batch b with perturbation p , we randomly sample n cells to build the image set $X_{p,b}$. We experiment with $n \in \{1, 4, 8, 16\}$, and maintain $N_P \times n = 512$ to maximize GPU utilization. Each epoch consists of 50k mini-batches.

We use ViT-small/16 as the backbone and set the hidden dimension of the MLP to 2048. The model is trained with an Adam optimizer for 300 epochs. We follow the same warm-up and cosine schedule for learning rate and weight decay as in the standard DINO framework [6], with a base learning rate of 0.04. The teacher temperature is set to 0.01. Eight local crops are used for each single-cell image. The Set-DINO framework is implemented in PyTorch and distributed over 2 GPUs. With $n = 16$, the model training takes 12 days.

We make the Set-DINO framework and the checkpoint of a trained model publicly available. ¹

¹<https://github.com/Genentech/set-dino>

Table 1. **Evaluation of batch-level gene profiles and consensus gene profiles.** We compare the performance of profiles from the Set-DINO framework with the standard DINO framework and engineered features from [15]. “Set-DINO - sgRNA” means that the teacher and student views are sampled from cells with the same guide RNA, while “Set-DINO - gene target” means that the views are sampled from cells with the same gene target (multiple guide RNAs). “KNN@k=5” refers to the accuracy of the k-nearest neighbor classifier when $k = 5$. All values are displayed in percentages. Best values are highlighted in bold.

	Batch Effect		Reproducibility		Biological Recall			
	↓ KNN @k=5	↓ GC @k=5	↑ KNN @k=5	↑ mAP	CORUM		CORUM (curated)	
					↑ Recall @5%	↑ Recall @10%	↑ Recall @5%	↑ Recall @10%
Engineered features	19.2	31.8	2.58	1.62	25.9	33.0	27.9	35.6
DINO (z-score)	12.6	10.1	0.68	0.41	21.1	28.4	23.5	31.3
DINO (NTC z-score)	21.7	34.5	1.16	0.64	25.4	33.7	28.4	36.9
Set-DINO (z-score) - sgRNA	22.5	42.8	5.48	3.55	27.3	46.2	33.8	43.9
Set-DINO (NTC z-score) - sgRNA	20.2	34.4	5.71	3.71	28.6	37.9	35.0	45.3
Set-DINO (z-score) - gene target	19.2	32.1	6.18	4.10	28.7	37.3	35.2	45.5
Set-DINO (NTC z-score) - gene target	17.3	23.4	6.87	4.51	29.5	38.3	36.1	46.9

3.5. Representation levels

After model training, single-cell embeddings from ViT are extracted, processed and then aggregated into multiple levels for further evaluation [2].

Single-cell profiles: For Set-DINO and standard DINO models, the embeddings of the class token from the last four ViT layers are used as cell-level features. As a baseline, we also use the engineered features released by [15]. The raw features are normalized using batch-wise normalization based on the median and median absolute deviation of features from NTCs, aiding in data alignment across different batches and mitigating batch effects [2]. Experimentally, we find that engineered features benefit from Principal Component Analysis (PCA), but learned features do not. As a result, we apply PCA to engineered features (after normalization) with a cutoff of 95% variance.

Batch-level gene profiles: Single-cell profiles of cells with the perturbations of the same gene target from the same batch are aggregated by an arithmetic mean operation into batch-level gene profiles. Batch-level gene profiles represent the cell population with a specific genetic perturbation from a specific batch.

Consensus gene profiles: Batch-level gene profiles are batch-wise centered on the means of features from NTCs and subsequently aggregated across batches to generate consensus gene profiles. Consensus gene profiles represent the average morphological changes resulting from individual genetic perturbations. These embeddings can be used to infer gene functions and gene-gene relationships.

3.6. Evaluation protocols

We employ multiple metrics to evaluate our learned representations on the basis of reproducibility, batch effect and

biological recall.

Reproducibility: After feature processing and aggregation, we first evaluate the reproducibility of the batch-level gene profiles, following previous evaluation approaches [9, 23]. Specifically, a graph is constructed where the nodes are batch-level gene profiles and the edge weights are given by the cosine distance between every pair of nodes. Then, we compute the average precision for the task of predicting the genetic perturbation of each node from the genetic perturbations of its distance-ranked neighbors, which measures the ability to retrieve the profiles of the same genetic perturbations from different batches against the background of all other perturbations.

Following this, the mean average precision (mAP) is calculated across all nodes. Also, we calculate the k-nearest neighbors (KNN) classification accuracy on perturbations when $k = 5$. A high mAP and accuracy indicate that batch-level gene profiles from the same genetic perturbation are clustered and dissimilar to other genetic perturbations as well as NTCs. Given that many perturbations exhibit negligible effects and those cells post-perturbation display very similar profiles to NTCs, the absolute values of mAP and accuracy are not expected to be high.

Batch Effect: To evaluate the batch effect, we calculate the KNN classification accuracy on experimental batches using the same graph described above ($k = 5$). In addition, we compute the Graph Connectivity (GC) [23] on the KNN graph. To calculate GC, subgraphs are constructed by retaining only nodes from a certain batch. GC is then defined as the average ratio of the number of nodes in the largest connected component and the total number of nodes in the subgraph. Low batch prediction accuracy and GC values indicate that the embeddings of cells from different experimental batches are well mixed (i.e., low batch effect).

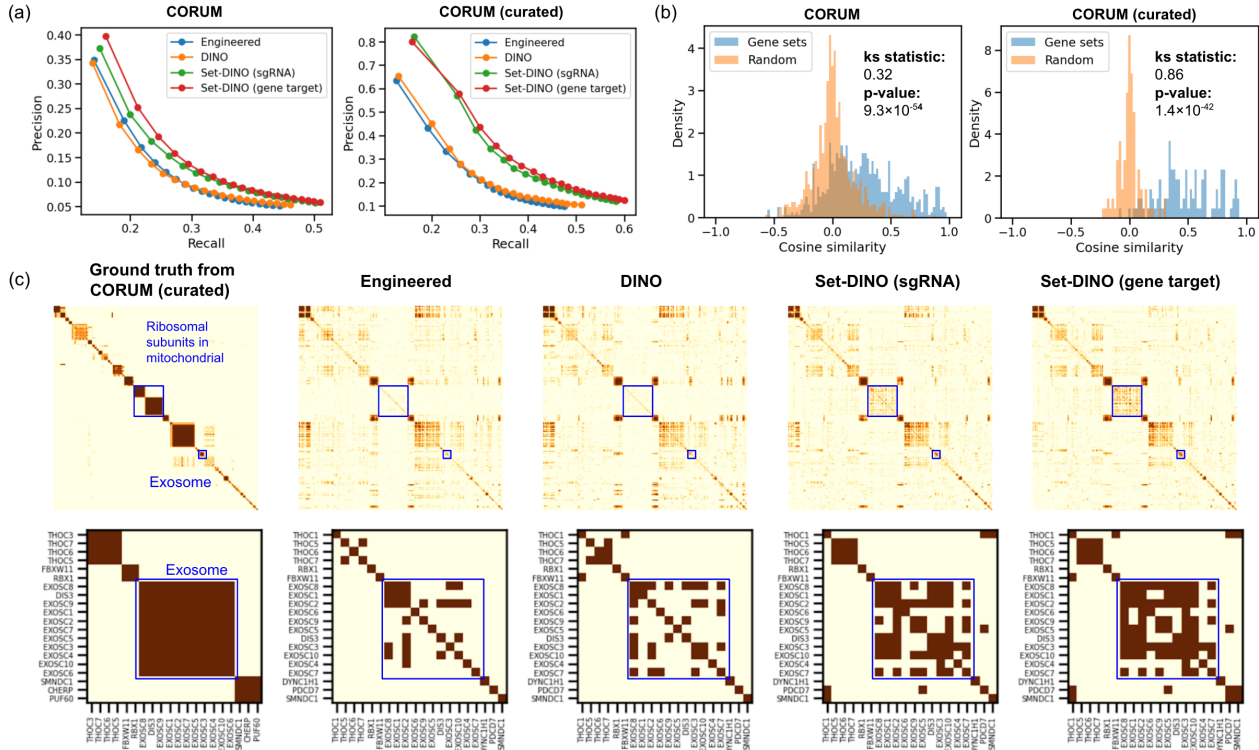


Figure 2. **Evaluation of consensus gene profiles in identifying gene-gene relationships.** (a) Precision-recall curves for predicting gene-gene relationships according to CORUM and curated CORUM. The curves are drawn by varying cutoff percentiles from top 1% to top 20%. (b) Distribution of cosine similarities between gene pairs. “Gene Sets” contains gene pairs that are involved in the same protein complex, while “Random” contains randomly sampled gene pairs. (c) The first row contains the full adjacency matrices of the ground truth graph (leftmost column) and prediction graphs (other columns). The second row provides a zoomed-in view of the adjacency matrix focused around the exosome. Although self-edges are included in the visualization, they are not considered when calculating biological recall and precision. All DINO and Set-DINO models shown in this figure were trained with NTC z-score normalization.

Biological Recall: Finally, we evaluate the biological information in consensus gene profiles by measuring how well they can infer gene-gene relationships [8]. The biological “ground truth” is built from the CORUM database [16], a public collection of manually curated mammalian protein complexes. A ground truth graph is built by connecting every pair of genes in the same protein complex. A prediction graph is constructed by connecting every pair of genes whose cosine similarity between the morphological profiles exceeds a certain percentile of the pairwise similarity distribution. The prediction graph is compared with the ground truth graph, and the recall of the gene-gene relationships is calculated using the top 5% percentile and top 10% percentile [8] as cutoffs. In addition, a precision-recall curve is calculated for further evaluation.

In CORUM, some protein complexes significantly overlap with others, potentially causing the involved genes to dominate the ground truth graph. To avoid this bias, we utilized a curated CORUM database from [15], which includes only protein complexes with limited overlap with

other complexes. The full CORUM contains gene-gene relationships from 1263 genes that are perturbed in our OPS dataset, and the curated CORUM includes a subset of 538 genes.

4. Results and Discussion

4.1. Set-DINO achieves superior performance compared to existing methods

Table 1 shows our results on reproducibility and batch effect on batch-level gene profiles, as well as the performance of gene-gene relationship inference using consensus gene profiles. The reproducibility metrics and batch effect metrics should be considered together to assess the quality of the learned representation. The former evaluates the replicate consistency and how well the model captures the perturbation effect, while the latter evaluates the model’s resistance to batch-level confounding factors.

These results show that the standard DINO model with NTC z-score yields a performance similar to engineered

Table 2. **Ablation study on cell sampling strategies.** We compare the performance of three sampling strategies with different numbers of cells, n . In the *same cell(s)* strategy the teacher and student views are built from the same set of cells. When $n = 1$, this is equivalent to the standard DINO framework. In the *within-batch* strategy the views are built from two sets of cells from the same batch with the same perturbation (same guide). In the *cross-batch* strategy the views are built from two sets of cells from different batches with the same perturbation (same guide). All values in this table are displayed in percentages. Values are omitted for models trained with the *cross-batch* strategy and $n \in 1, 4$ because their training collapsed. All DINO and Set-DINO models shown in this table were trained with NTC z-score normalization on input images. “KNN@k=5” refers to the accuracy of the k-nearest neighbor classifier when $k = 5$. Best values are highlighted in bold.

	Batch Effect		Reproducibility		Biological Recall			
	↓ KNN @k=5	↓ GC @k=5	↑ KNN @k=5	↑ mAP	CORUM		CORUM (curated)	
					↑ Recall @5%	↑ Recall @10%	↑ Recall @5%	↑ Recall @10%
Engineered features	19.2	31.8	2.58	1.62	25.9	33.0	27.9	35.6
Same cell(s), n=1	21.7	34.5	1.16	0.64	25.4	33.7	28.4	36.9
Same cell(s), n=4	24.2	45.8	1.00	0.54	22.9	30.9	27.3	36.1
Within-batch, n=1	47.1	75.7	0.97	0.53	21.9	30.5	26.3	36.1
Within-batch, n=4	44.2	71.8	0.47	0.28	19.8	29.0	24.9	34.3
Cross-batch, n=1,4	—	—	—	—	—	—	—	—
Cross-batch, n=8	20.0	36.2	5.83	3.78	27.7	36.4	34.7	45.1
Cross-batch, n=16	20.2	34.4	5.71	3.71	28.6	37.9	35.0	45.3

features in predicting gene-gene relationships. Notably, Set-DINO significantly outperforms both of these on the gene-relationship task using both CORUM and curated CORUM. With the optimal Set-DINO gene profiles, the recall at top 5% cutoff increases by 8.2% (29.4% relative increase) on curated CORUM compared to engineered features. Additionally, the reproducibility metrics for Set-DINO profiles markedly surpass those of both DINO and engineered features. This boost suggests that Set-DINO’s weakly supervised training encourages the model to learn more biologically meaningful representations.

Interestingly, within the self-supervised setting, the inference performance of gene-gene relationships correlates with our reproducibility metric. This suggests that combining weak supervision on perturbation labels with SSL leads to morphological features that encode biologically meaningful information more effectively.

Furthermore, comparing the guide-level (sgRNA) and gene-level variants of Set-DINO, we observe that building views from all cells with the same gene target leads to a lower batch effect, and a relative increase of over 20% on reproducibility metrics, and a slight increase in predicting gene-gene relationships.

Finally, we observe a consistent improvement in all performance metrics when using NTC z-score normalization. For standard DINO with z-scoring, we note that although its batch effect metrics are low, it also exhibits low replicate consistency, which indicates that this model has less capability to extract distinguishable features. Finally, our batch effect metrics only consider batch-level confounding

factors. Therefore, it is possible that the embeddings from DINO with z-score are dominated by other nuisance factors, such as cell positions within the well.

4.2. Set-DINO encodes biologically meaningful information

We further illustrate the biological information encoded in consensus gene profiles in Figure 2. Figure 2a compares recall-precision curves from different methods using cut-off ranges from top 1% to top 20%, focusing on the high-precision regime that is required for target discovery. We note that while standard DINO and engineered features exhibit similar performance, Set-DINO achieves a substantial performance boost, especially on curated CORUM.

We also observe that gene pairs with known relationships tend to have higher cosine similarity than randomly sampled gene pairs (Figure 2b). The KS-statistic between the two distributions is 0.32 on CORUM, and 0.86 on curated CORUM.

Figure 2c presents the adjacency matrices from the ground truth graph on curated CORUM, as well as the predicted graphs from different methods. We find that consensus gene profiles from Set-DINO models achieve higher recall in multiple protein complexes. Notably, the gene-gene relationships in some protein complexes are almost entirely missed in engineered and DINO profiles but captured by Set-DINO. Two examples of this are “ribosomal subunits in mitochondria”, which are critical for mitochondrial translation [30], and exosomes, which play an integral role in cell-cell communication [21]. Both of these structures are

highlighted in Figure 2c.

4.3. Ablation studies on Set-DINO framework

We perform a series of ablation studies to evaluate the effect of weak supervision in self-supervised learning of cellular embeddings. See Table 2 for a summary of these results.

Our first observation is that model training tends to collapse if we apply the proposed cross-batch sampling strategy with a small value of n . This phenomenon is likely due to the substantial differences in cell state distribution and the general batch-level distribution shift exhibited by the two sets of cells. Increasing the value of n aids in stabilizing the model during training. Based on our results, increasing beyond $n = 8$ has a negligible impact on performance.

In addition to cross-batch sampling, we propose two alternative sampling strategies. The first strategy is the “same cell” approach, in which the identical sets of cells are used for both teacher and student branches. When $n = 1$, this is equivalent to the standard DINO framework. A comparison of the results from $n = 1$ and $n = 4$ reveals that utilizing a set of images leads to higher batch effects and lower reproducibility metrics. The recall in gene-gene relation prediction is also lower. These results are likely due to the averaging done in the aggregation layer, which may reduce the effect of image augmentation.

The second “within-batch” strategy involves building teacher and student views by sampling different cells with the same perturbation from the same batch. In this case, the differences between the two views mainly arise from the variation in cell state distributions due to random sampling. The results indicate that the profiles from within-batch sampling exhibit markedly worse batch effects as well as lower reproducibility metrics and gene-gene relationship prediction performance. This suggests that in this scenario, the model extracts primarily batch-related information to ensure consistency between the teacher and student branches.

For additional validation, we perform a PCA on the batch-level gene profiles, and calculate the reproducibility and batch effect metrics on an increasing number of Principal Components (PCs) by order. This analysis provides insight into the amount of perturbation-specific and batch-specific signals contained in the PCs with the highest variances. Figure 3 shows that the batch-specific signals dominate the high-ranked PCs of gene profiles from the model trained using the “within-batch” strategy, while perturbation-specific signals dominate the high-ranked PCs from the “cross-batch” trained model. In addition, according to the metrics we observed on the validation set during model training for within-batch sampling (not shown), batch-specific signals begin to dominate at an early stage in the training process.

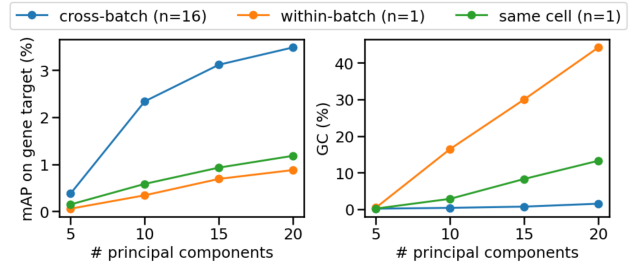


Figure 3. **Performance analysis on an increasing number of principal components.** Principal component analysis (PCA) is performed on batch-level gene profiles. Reproducibility and batch effect metrics on an increasing number of principal components are measured.

5. Conclusion

In this study, we propose the Set-DINO model with a cross-batch sampling strategy that combines weak supervision and self-supervised learning to obtain better single-cell representations for cell morphology images. Our results demonstrate that the proposed framework outperforms established baselines using engineered features and the standard DINO framework in extracting morphological profiles of single-cell images from a held-out test set. We conduct ablation studies to confirm that both set-level representation and cross-batch sampling are critical to achieving success.

Additional evaluation based on prior biological knowledge reveals that the consensus gene profiles learned by Set-DINO significantly improve the prediction of gene-gene relationships. Thus, we anticipate that the proposed framework could benefit future target discovery and drug discovery research. Furthermore, while this study focuses on single-cell imaging data from optical pooled screens, Set-DINO may also be applicable to other single-cell datasets containing weak labels, as well as single-cell crops from arrayed cell painting datasets.

One limitation of this study is that a simple arithmetic average is used in the aggregation layer to fuse the latent embeddings from single-cell images to population-level representations. Previous studies have explored more sophisticated aggregation methods for computing set-level representations [25, 26, 36]. Future work will investigate which strategy is optimal for feature aggregation.

Acknowledgments

We would like to thank Avtar Singh and Luke Funk for productive discussions, and their support with the OPS dataset.

References

- [1] D Michael Ando, Cory Y McLean, and Marc Berndl. Improving phenotypic measurements in high-content imaging

- screens. *BioRxiv*, page 161422, 2017. 3
- [2] Juan C Caicedo, Sam Cooper, Florian Heigwer, Scott Warchal, Peng Qiu, Csaba Molnar, Aliaksei S Vasilevich, Joseph D Barry, Harmanjit Singh Bansal, Oren Kraus, et al. Data-analysis strategies for image-based cell profiling. *Nature methods*, 14(9):849–863, 2017. 2, 5
- [3] Juan C Caicedo, Claire McQuin, Allen Goodman, Shantanu Singh, and Anne E Carpenter. Weakly supervised learning of single-cell feature embeddings. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9309–9318, 2018. 1
- [4] Juan C Caicedo, Shantanu Singh, and Anne E Carpenter. Applications in image-based profiling of perturbations. *Current opinion in biotechnology*, 39:134–142, 2016. 1
- [5] Rebecca J Carlson, Michael D Leiken, Alina Guna, Nir Hacohen, and Paul C Blainey. A genome-wide optical pooled screen reveals regulators of cellular antiviral responses. *Proceedings of the National Academy of Sciences*, 120(16):e2210623120, 2023. 1
- [6] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021. 1, 2, 3, 4
- [7] Anne E Carpenter, Thouis R Jones, Michael R Lamprecht, Colin Clarke, In Han Kang, Ola Friman, David A Guertin, Joo Han Chang, Robert A Lindquist, Jason Moffat, et al. Cellprofiler: image analysis software for identifying and quantifying cell phenotypes. *Genome biology*, 7:1–11, 2006. 1
- [8] Safiye Celik, Jan-Christian Hütter, Sandra Melo Carlos, Nathan H Lazar, Rahul Mohan, Conor Tillinghast, Tommaso Biancalani, Marta Fay, Berton A Earnshaw, and Imran S Haque. Biological cartography: Building and benchmarking representations of life. *Biorxiv*, pages 2022–12, 2022. 6
- [9] Srinivas Niranj Chandrasekaran, Beth A Cimini, Amy Goodale, Lisa Miller, Maria Kost-Alimova, Nasim Jamali, John G Doench, Briana Fritchman, Adam Skepner, Michelle Melanson, et al. Three million images and morphological profiles of cells treated with matched chemical and genetic perturbations. *Biorxiv*, pages 2022–01, 2022. 5
- [10] Michael J Cox, Steffen Jaensch, Jelle Van de Waeter, Laure Cougnaud, Daan Seynaeve, Soulaïman Benalla, Seong Joo Koo, Ilse Van Den Wyngaert, Jean-Marc Neefs, Dmitry Malkov, et al. Tales of 1,008 small molecules: phenomic profiling through live-cell imaging in a panel of reporter cell lines. *Scientific reports*, 10(1):13262, 2020. 1
- [11] Jan Oscar Cross-Zamirski, Guy Williams, Elizabeth Mouchet, Carola-Bibiane Schönlieb, Riku Turkki, and Yin-hai Wang. Self-supervised learning of phenotypic representations from cell images with weak labels. *arXiv preprint arXiv:2209.07819*, 2022. 1, 2
- [12] Robert Van Dijk, John Arevalo, Shantanu Singh, and Anne E Carpenter. Learning representations of cell populations for image-based profiling using contrastive learning. In *NeurIPS 2022 Workshop on Learning Meaningful Representations of Life*, 2022. 3
- [13] Michael Doron, Théo Moutakanni, Zitong S Chen, Nikita Moshkov, Mathilde Caron, Hugo Touvron, Piotr Bojanowski, Wolfgang M Pernice, and Juan C Caicedo. Unbiased single-cell morphology with self-supervised vision transformers. *bioRxiv*, pages 2023–06, 2023. 1, 2
- [14] David Feldman, Avtar Singh, Jonathan L Schmid-Burgk, Rebecca J Carlson, Anja Mezger, Anthony J Garrity, Feng Zhang, and Paul C Blainey. Optical pooled screens in human cells. *Cell*, 179(3):787–799, 2019. 1
- [15] Luke Funk, Kuan-Chung Su, Jimmy Ly, David Feldman, Avtar Singh, Brittanica Moodie, Paul C Blainey, and Iain M Cheeseman. The phenotypic landscape of essential human genes. *Cell*, 185(24):4634–4653, 2022. 1, 3, 5, 6
- [16] Madalina Giurgiu, Julian Reinhard, Barbara Brauner, Irmtraud Dunger-Kaltenbach, Gisela Fobo, Goar Frishman, Corinna Montrone, and Andreas Ruepp. Corum: the comprehensive resource of mammalian protein complexes—2019. *Nucleic acids research*, 47(D1):D559–D563, 2019. 6
- [17] Johan Fredin Haslum, Christos Matsoukas, Karl-Johan Leuchowius, Erik Müllers, and Kevin Smith. Metadata-guided consistency learning for high content images. In *Medical Imaging with Deep Learning*, pages 918–936. PMLR, 2024. 1, 2
- [18] Johan Fredin Haslum, Christos Matsoukas, Karl-Johan Leuchowius, and Kevin Smith. Bridging generalization gaps in high content imaging through online self-supervised domain adaptation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 7738–7747, 2024. 1, 2
- [19] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16000–16009, 2022. 2
- [20] Rens Janssens, Xian Zhang, Audrey Kauffmann, Antoine de Weck, and Eric Y Durand. Fully unsupervised deep mode of action learning for phenotyping high-content cellular images. *Bioinformatics*, 37(23):4548–4555, 2021. 3
- [21] Raghu Kalluri and Valerie S LeBleu. The biology, function, and biomedical applications of exosomes. *Science*, 367(6478):eaau6977, 2020. 7
- [22] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *Advances in neural information processing systems*, 33:18661–18673, 2020. 2
- [23] Vladislav Kim, Nikolaos Adaloglou, Marc Osterland, Flavio M Morelli, Marah Halawa, Tim König, David Gnutt, and Paula A Marin Zapata. Self-supervision advances morphological profiling by unlocking powerful image representations. *BioRxiv*, pages 2023–04, 2023. 1, 2, 3, 5
- [24] Oren Kraus, Kian Kenyon-Dean, Saber Saberian, Maryam Fallah, Peter McLean, Jess Leung, Vasudev Sharma, Ayla Khan, Jia Balakrishnan, Safiye Celik, et al. Masked autoencoders are scalable learners of cellular morphology. *arXiv preprint arXiv:2309.16064*, 2023. 1, 2
- [25] Dong Bok Lee, Seanie Lee, Joonho Ko, Kenji Kawaguchi, Juho Lee, and Sung Ju Hwang. Self-supervised set representation learning for unsupervised meta-learning. *arXiv*

- preprint arXiv:2310.06511*, 2023. [2](#), [3](#), [8](#)
- [26] Juho Lee, Yoonho Lee, Jungtaek Kim, Adam R Kosiorek, Seungjin Choi, and Yee Whye Teh. Set transformer. 2018. [3](#), [8](#)
- [27] Nikita Moshkov, Michael Bornholdt, Santiago Benoit, Matthew Smith, Claire McQuin, Allen Goodman, Rebecca A Senft, Yu Han, Mehrtash Babadi, Peter Horvath, et al. Learning representations for image-based profiling of perturbations. *Nature Communications*, 15(1):1594, 2024. [1](#), [2](#), [4](#)
- [28] Meraj Ramezani, Julia Bauman, Avtar Singh, Erin Weisbart, John Yong, Maria Lozada, Gregory P. Way, Sanam L. Kavari, Celeste Diaz, Marzieh Haghighi, Thiago M. Batista, Joaquín Pérez-Schindler, Melina Claussnitzer, Shantanu Singh, Beth A. Cimini, Paul C. Blainey, Anne E. Carpenter, Calvin H. Jan, and James T. Neal. A genome-wide atlas of human cell morphology. *bioRxiv*, 2023. [1](#)
- [29] Mohammad H Rohban, Hamdah S Abbasi, Shantanu Singh, and Anne E Carpenter. Capturing single-cell heterogeneity via data fusion improves image-based profiling. *Nature communications*, 10(1):2082, 2019. [2](#)
- [30] Juan Sastre, Federico V Pallardó, José García de la Asunción, and José Viña. Mitochondria, oxidative stress and aging. *Free radical research*, 32(3):189–198, 2000. [7](#)
- [31] Srinivasan Sivanandan, Bobby Leitmann, Eric Lubeck, Mohammad Muneeb Sultan, Panagiotis Stanitsas, Navpreet Ranu, Alexis Ewer, Jordan E Mancuso, Zachary F Phillips, Albert Kim, et al. A pooled cell painting crispr screening platform enables de novo inference of gene function by self-supervised deep learning. *bioRxiv*, pages 2023–08, 2023. [1](#), [2](#)
- [32] David R Stirling, Madison J Swain-Bowden, Alice M Lucas, Anne E Carpenter, Beth A Cimini, and Allen Goodman. Cellprofiler 4: improvements in speed, utility and usability. *BMC bioinformatics*, 22:1–11, 2021. [1](#)
- [33] Maciej Sypetkowski, Morteza Rezanejad, Saber Saberian, Oren Kraus, John Urbanik, James Taylor, Ben Mabey, Mason Victors, Jason Yosinski, Alborz Rezazadeh Sereshkeh, et al. Rrx1: A dataset for evaluating experimental batch correction methods. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4284–4293, 2023. [2](#)
- [34] Mojca Mattiazi Usaj, Erin B Styles, Adrian J Verster, Helena Friesen, Charles Boone, and Brenda J Andrews. High-content screening for quantitative cell biology. *Trends in cell biology*, 26(8):598–611, 2016. [1](#)
- [35] Russell T Walton, Avtar Singh, and Paul C Blainey. Pooled genetic screens with image-based profiling. *Molecular Systems Biology*, 18(11):e10768, 2022. [1](#)
- [36] Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Russ R Salakhutdinov, and Alexander J Smola. Deep sets. *Advances in neural information processing systems*, 30, 2017. [2](#), [3](#), [8](#)