SINGLE-CELL MORPHOLOGICAL PROFILING REVEALS INSIGHTS INTO CELL DEATH

Benjamin Frey[1], **David Holmberg**[1], **Petter Byström**[2], **Ebba Bergman**[1], **Polina Georgiev**[1], **Martin Johansson**[1], **Patrick Hennig**[1], **Jonne Rietdijk**[1], **Dan Rosén**[1], **Jordi Carreras-Puigvert**[1,2], and **Ola Spjuth**[1,2]

[1]Department of Pharmaceutical Biosciences and Science for Life Laboratory, Uppsala University, Uppsala, Sweden.
[2]Phenaros Pharmaceuticals AB, Uppsala, Sweden.

ABSTRACT

Analysis of single-cell data has emerged as a powerful tool for studying biological processes and response to perturbations. However, its application in morphological profiling is less explored. In this study we profile six cell death subtypes induced by 50 small molecule drugs across six concentrations using the Cell Painting assay. We evaluate the performance of three feature extraction methods at single-cell and aggregated level and apply supervised and unsupervised analyses to uncover factors contributing to cell death mechanisms. Our results show that a bagged LightGBMXT model, trained on single-cell DeepProfiler achieved classification accuracy of 77.23%, with a top overall performance of 89.97% for corresponding aggregated profiles. Furthermore, self-supervised learning using the transformer-based DINO network revealed highly resolved and biologically meaningful subpopulations, shedding light on perturbation- and concentration-specific molecular targets and dose-dependant morphological changes. Our findings demonstrate the potential of studying phenomic data on single-cell level to enhance the characterization of cell death pathways, advancing our understanding of perturbation effects at a granular level.

1 Introduction

Single-cell data has become a cornerstone of modern biological research, providing unprecedented insights into cellular heterogeneity, immune responses, and disease mechanisms. Specifically in the field of transcriptomics, single-cell analysis has enabled the identification of subpopulations within heterogeneous samples, aided in uncovering the effects of perturbations, and advancing our understanding of tumor biology (Bunne et al., 2023; Lotfollahi et al., 2023; Peidli et al., 2024; Wu et al., 2024; Yao et al., 2022). Single-cell sequencing technologies have proven invaluable for exploring immune responses and elucidating disease mechanisms, particularly in the context of cancer biology (Olofsson et al., 2021). Despite the growing impact of single-cell sequencing, the application of high-throughput single-cell data for characterizing specifically cell death subtypes remains an open challenge.

While sequencing-based approaches are powerful, they are not capable of capturing phenotypic responses that can offer valuable insights into cellular processes such as cell death (Y. Chen et al., 2024). Image-based approaches provide a complementary perspective. Cell Painting is a high-content imaging method that characterizes cellular morphology through fluorescent labeling of eight distinct compartments, enabling researchers to visualize and analyze cellular features at scale (Bray et al., 2016). This assay has been widely adopted for its ability to enable high-throughput screening of millions of cells and thousands of compounds (Chandrasekaran et al., 2023; Kraus et al., 2024). By integrating morphological profiling with data-driven analysis, Cell Painting allows researchers to study how cells respond to perturbations at both population and individual cell levels. The technique has shown promise for identifying mechanisms of action (MOA) of compounds and assessing toxicity through feature extraction and subsequent downstream analyses (Caicedo et al., 2017; Rietdijk, Aggarwal, et al., 2022).

The analysis pipeline for Cell Painting data begins with a feature extraction step that leverages tools like CellProfiler (Carpenter et al., 2006), a dedicated software that extracts pre-defined features from cellular images. CellProfiler has

been instrumental in advancing morphological profiling by enabling researchers to extract detailed cellular characteristics (Harrison et al., 2023; Haslum et al., 2023; Rietdijk, Tampere, et al., 2021). However, the reliance on hand-crafted features has limitations, particularly as data scales and biological complexity increases. Recent advances in computer vision have introduced alternative approaches that leverage deep neural networks and optimized feature sets for automated feature extraction (Kim et al., 2024; Kraus et al., 2024; Stossi et al., 2024). These approaches include convolutional neural networks (CNN) (Tian et al., 2023), weakly-supervised techniques (Moshkov et al., 2024), and self-supervised models (Caron et al., 2021; He et al., 2021; Kraus et al., 2024).

Self-supervised vision models, in particular, have shown great promise in morphological profiling as they do not require extensive data labeling. One example, DINO employs two parallel networks to learn representations from different crops of the same image, enabling robust feature extraction. DINO has been successfully applied to Cell Painting data, outperforming existing methods in MOA prediction and toxicity estimation tasks (Doron et al., 2023; Gupta et al., 2024; Kim et al., 2024; Moshkov et al., 2024; Pfaendler et al., 2023). Similarly, Masked Autoencoders, which use pixel masking and encoder-decoder architectures, have demonstrated state-of-the-art performance on Cell Painting datasets. Notably, Recursion Pharmaceuticals recently released an open version of a Masked Autoencoder model trained on millions of Cell Painting images, a first foundation model for Cell Painting data (Kraus et al., 2024).

Despite these advancements, the majority of morphological profiling workflows where cells are perturbed using chemical or genetic perturbations aggregate features extracted from single-cell imaging into average profiles. By combining profiles from the same treatments and using mean/median cells as the main objects, outliers and data complexity can be reduced (Caicedo et al., 2017; Moshkov et al., 2024). While this approach improves data separation at the treatment level, it effectively sacrifices information about cellular heterogeneity and subpopulations at the single-cell level. This trade-off limits the deeper biological characterization of individual cells, where heterogeneity can reveal essential insights into responses to perturbations.

Applications where Cell Painting experiments have utilized single-cell data in the analysis remain relatively limited. Early efforts focused on tasks like protein co-localization analysis, classification, and the characterization of blood and immune cell types (Doron et al., 2023; Gupta et al., 2024; Moshkov et al., 2024; Pahl et al., 2023; Pfaendler et al., 2023). One specific application where single-cell analysis from Cell Painting has not been explored is cell death. However, the exploration of cell death mechanisms and cellular heterogeneity at the single-cell level has been largely unexplored. Schorpp et al. classified cell death subtypes using a CNN and larger image crops from four-channel Cell Painting experiments (Schorpp et al., 2023), covering two cell death subtypes (apoptosis, Ferroptosis, as well as non-treated cells) at seven inducing compounds per subtypes in five concentrations per compound. For their classification, they compared an end-to-end CNN with 512 x 512 crops as inputs with traditional machine learning models trained on 245-dimensional aggregated feature vectors.

In this manuscript we present a comprehensive exploration of six cell death subtypes induced by 50 compounds in MCF-7 breast cancer cells, each tested at six different concentrations using five-channel Cell Painting. In order to assess the relevance of single-cell morphological profiles for studying cell death subtypes, we apply unsupervised and supervised methods and compare the performance of three feature extraction approaches: CellProfiler (Stirling et al., 2021), DeepProfiler (Moshkov et al., 2024), and the self-supervised transformer-based DINO.

2 Methods and Materials

2.1 Compounds and cell death subtypes

50 compounds were selected where annotations spanned six distinct cell death subtypes: apoptosis, Necroptosis, Ferroptosis, Pyroptosis, as well as the non-regulated Autophagy and Immunogenic Cell Death (ICD). The compounds and cell death subtype annotations are shown in Figure 1 B; and with their associated targets in Supplementary Table 1. Apoptosis, as the most extensively studied form of cell death, is characterized via cell shrinkage followed by fragmentation and development of apoptotic bodies (Kerr et al., 1972). Necroptosis describes cell death defined by membrane rupture and subsequent release of cellular contents. Necroptotic processes are induced by tumour necrosis factor (TNF) leading to a RIPK and caspase 8 activation (Conrad et al., 2016). Ferroptosis is driven by metabolic processes, specifically involving fatty acyl moieties dependent on iron (Stockwell, 2022). It has been shown to be connected to different types of degenerative diseases involving the brain, heart and kidney. Pyroptosis presents a natural cellular defense mechanism induced by pathogen-associated molecular patterns (PAMPs) and damage-associated molecular patterns (DAMP) which trigger cell swelling, lysis and the release of pro-inflammatory cytokines, and shown to have cancer-related indications (Yu et al., 2021). Autophagy is a process driven by the lysosomal machinery, resulting in the self-degradation of cellular components (Shen et al., 2023). Immunogenic cell death, on the other hand, is initiated by immune cells, particularly T cells. This process involves alterations in the cell membrane structure, which

recruit highly specialized immune cells, such as CD8+ T cells, ultimately leading to the targeted destruction of the affected cell (Kroemer et al., 2013).

2.2 Cell culture

MCF7 human breast cancer cells (Sigma-Aldrich #86012803) were cultured to passage +24 and confirmed to be Mycoplasma-free using the MycoAlert kit (Lonza #LT07-218). Cells were maintained at 37°C, 5% CO_2, and 95% humidity in Dulbecco's Modified Eagle Medium (DMEM; Gibco #31885-023) supplemented with 10% (v/v) heat-inactivated fetal bovine serum (FBS; Gibco #10500064) and 1% (v/v) Penicillin-Streptomycin (Gibco #15140-122). At 80-90% confluency, cells were washed with Dulbecco's Phosphate-Buffered Saline (DPBS; Gibco #14190250), dissociated using TrypLE Select (Gibco #A1217701), and reseeded in fresh complete growth medium in T75 flasks (Thermo Scientific #156499).

2.3 Cell seeding and compound treatment

All test compounds were provided as DMSO solutions and diluted to stock concentrations using a 10-fold serial dilution. Phenotypic reference compounds that are known to induce a distinctive phenotype in a variety of cells (Willis et al., 2020), including Etoposide (Sigma-Aldrich #E1383), Fenbendazole (Sigma-Aldrich #F5396), and Staurosporine (Sigma-Aldrich #19-123-M), were purchased from Sigma-Aldrich and prepared as 10 mM stocks in DMSO (Sigma-Aldrich #D8418). Assay-ready plates containing test compounds, phenotypic reference compounds, and DMSO (negative controls) were prepared using an I-DOT (immediate drop-on-demand technology) liquid handler. Compounds were dispensed in precise volumes to achieve final treatment concentrations of 0.1, 0.3, 1, 3, 5, and 10 μM for test compounds; 3 μM for the positive controls Etoposide and Fenbendazole; 0.3 μM for the Staurosporine control; and 0.1% for DMSO.

Each compound plate included 20 or 21 test compounds at 6 concentrations and 2 technical replicates, 3 phenotypic reference compounds at 1 concentration and 2 technical replicates, and 14 wells of DMSO. The distribution of treatments and controls on the compound plates was optimized using PLAID (Plate Layouts using Artificial Intelligence Design) (Francisco Rodríguez et al., 2023), a flexible constraint-programming model that ensures a robust microplate layout.

Prior to cell seeding on the assay ready plates, the compounds pre-spotted onto the assay plates were dissolved by adding 20 μL of complete growth medium per well using a Biotek Multiflo FX dispenser. The plates were then placed on a shaker at 150 rpm at room temperature for one hour to ensure complete dissolution of the compounds. Cells from passage 24 were counted and seeded on top of the plates with dissolved compounds at a density of 1100 cells per well by dispensing 20 μL/well of cell suspension with Biotek Multiflo FX dispenser, resulting in a total final volume of 40 μl/well. The seeded plates were left to stand at room temperature for 20 minutes before being incubated at 37°C in a 5% CO_2 atmosphere for 48 hours.

2.4 Cell Painting assay

Cell painting experiments were conducted according to the protocol described Bray et al., 2016 with minor modifications. The fluorescent labeling process was performed in an automated laboratory setup, which included a UR10 robotic arm for handling plates, a Biotek MultiFlo FX for dispensing solutions, and a BlueWasher (Blue Cat Bio) for washing steps. In brief, MCF7 cells were cultured in 384-well plates and treated with compounds at various doses for 48 hours. Following treatment, cells were live-cell-labeled with MitoTracker Deep Red (Invitrogen) and subsequently fixed with 4% paraformaldehyde (Histolab) for 20 minutes. After fixation, cells were permeabilized with 0.1% Triton X-100 (Cytiva) for 20 minutes and stained with a solution containing Hoechst 33342, Wheat Germ Agglutinin, Phalloidin, SYTO 14, and Concanavalin A (all from Thermo Fisher Scientific) for 20 minutes. Finally, plates were washed, sealed, and stored at 4°C, protected from light until image acquisition.

2.5 Imaging and image pre-processing

Fluorescence microscopy was performed using a high-throughput ImageXpress Micro XLS (Molecular Devices), capturing nine sites per well across five fluorescence channels. Excitation and emission wavelengths were optimized for each dye. Image processing and analysis were conducted with CellProfiler and CellPose 2.0 (Stringer et al., 2021), including quality control, illumination correction, segmentation, and feature extraction. A total of 107,405 images were collected, resulting in 2TB of data. Segmentation results were saved as binary masks and later used as input for feature extraction (CellProfiler), image cropping (DeepProfiler) and training (DINO).

2.6 CellProfiler

A feature extraction pipeline using Cell Profiler version 4.3.4 (available at `https://doi.org/10.17044/scilifelab.21378906.v2` and described in (Harrison et al., 2023)) was run and resulted in a 2,160-dimensional feature vector for each cell. These features capture fluorescence intensity, location, granularity, and other measurements of the imaged cells, grouped into cell, nucleus, and cytoplasm features. Before feature extraction, an illumination correction function was calculated and applied to the images.

2.7 DeepProfiler

DeepProfiler is a convolutional neural network trained in a weakly supervised setting, as described by (Moshkov et al., 2024). In this study, we utilized the pre-trained model CellPainting_CNN with weights from (Moshkov et al., 2024) for inference without additional training. Features were extracted as 672-dimensional vectors from the block_6a_-activation layer of the EfficientNetB0 architecture, generating one feature vector per single-cell crop. To ensure accurate representation, we used nucleus-centered image crops with a resolution of 150 × 150 pixels. The cropping size was selected to fully encapsulate 95% of all cells. Illumination correction was applied during the DeepProfiler feature extraction process, using a down_scale_factor of 8 and a median_filter_size of 48. Importantly, segmentation masks were omitted during inference to retain contextual neighborhood information surrounding the cells.

2.8 DINO

Self-distillation with no labels (DINO) falls under the self-supervised learning paradigm to learn representation from images without the need for explicit label annotation (Caron et al., 2021). The approach combines identical student and teacher networks to learn image features through self-distillation. We trained DINO with a ViT-small network, a patch size of 16 pixels, and five input channels. We adapted the implementation by (Doron et al., 2023), to allow for multi-channel inputs. The network was trained from scratch with a subset of our Cell Death dataset consisting of >185,000 single-cell images, comprising all compounds at all concentrations. As input, we used the segmented five-channel site-level images, center-cropped around the nuclei, and stacked the images into a (150 x 150 x 5) image for each cell. Pixel intensities were normalized to [0, 1] and images resized to 224 x 224 pixels. Using the same augmentations as (Doron et al., 2023) (random channel dropping, random brightness and contrast, artificial warping), we trained the DINO network for 200 epochs with a learning rate of 0.0005. Training was performed on a single NVIDIA GTX 3090 with a batch size of 20 for 3 days. We included randomly sampled DMSO cells to allow the network to train on as much variance in morphologies as possible. Inference was performed on the teacher network, resulting in 384-dimensional single-cell feature vectors. To evaluate the model, we inferred 625,000 randomly sampled cells from our cell death dataset.

2.9 Data post-processing

To ensure comparability, we applied the same post-processing and analysis pipeline to all three feature sets. In the following, we will also refer to the extracted feature vectors as "profiles". Feature vectors were first standardized to the negative control (DMSO) by subtracting the mean and dividing by the standard deviation of the DMSO cells per plate. For our aggregated analysis, we calculated the median profile per site. Feature selection was performed by removing missing value features, highly correlated (>99%), low variance (<0.01%) and blacklisted features as defined in the pycytominer package (Way et al., 2022). Additionally, features with >25% constant values across cells were removed. This resulted in just over 1,400 features for CellProfiler, 384 features for DINO, and 672 for DeepProfiler, respectively.

2.10 Energy distance

Energy distance (e-distance) (Peidli et al., 2024) was originally implemented for single-cell sequencing data as a metric for perturbation strength. We calculated e-distances per compound concentration set using single-cell profiles and the first 50 principal components.

Let $x_1, ...x_N \in C_1$ and $y_1, ...y_M \in C_2$ be single-cell profiles, where *N*, *M* represent the respective sample size and C_T be our treatment spaces. We further define:

$$\begin{aligned} \delta_{XY} &= \frac{1}{NM} \sum_{i=1}^{N} \sum_{j=1}^{M} \|x_i - y_j\| \\ \sigma_X &= \frac{1}{N^2} \sum_{i=1}^{N} \sum_{j=1}^{N} \|x_i - x_j\| \end{aligned} \tag{1}$$

The energy distance is then defined as:

$$E(X, Y) := 2\delta_{XY} - \sigma_X - \sigma_Y \tag{2}$$

Using the e-distance as a test statistic, we performed a permutation test with 50,000 permutations to calculate a p-value for each compound concentration in relation to the negative control. We removed all compound concentrations with an adjusted p-value > 0.05.

2.11 Grit score

Grit score measures the perturbation strength and reproducibility of a given treatment group. We calculated grit scores on a compound-concentration level using aggregated profiles. We filtered compound concentrations with a grit score below 1.5.

Let $d_1, ...d_N \in N$ be morphological profiles of a control group and $t \in C_T$ a target profile of given treatment *T* with *M* replicate profiles $r_1, ..., r_M$. Then define $D := \{d_1, ..., d_M\}$ where $d_i = corr(t, c_i)$ and $S := \{s_1, ..., s_M\}$ where $s_i = corr(t, r_i)$. Let further μ_D, σ_D be the respective statistics of *D*. Then the grit score *G* is defined as:

$$\begin{aligned} g_i &= \frac{s_i - \mu_D}{\sigma_D} \\ G &= \frac{1}{M} \sum_{i=1}^{M} g_i \end{aligned} \tag{3}$$

where *M* represents the number of replicates.

Both e-distance and grit score were calculated with respect to the negative control, which was then omitted for further analysis. Reference compounds in Supplementary Table 2 that were not annotated with specific cell death subtypes were omitted.

2.12 Visualisation and clustering

We used the scanpy (Wolf et al., 2018) library to calculate PCA, neighborhood graphs, and UMAP embeddings. We trained our UMAP on single-cell profiles using PAGA for initialization and cosine similarity as the distance metric. For visualization purposes, cells extracted with CellProfiler and DeepProfiler were sampled to 10% in a stratified manner. For DINO results, we used the full data set post-filtering as it matched the sampled CellProfiler and DeepProfiler sets in size.

2.13 Supervised model training

The extracted profiles from CellProfiler, DeepProfiler, and DINO were used to train a classifier both on single-cell profiles and the aggregated profiles. First, profiles were split into a training and testing set (70/30), stratified over Cell Death subtypes. As our dataset contained a limited number of compounds per Cell Death subtype, we implemented a split with no overlap between seeding wells. Next, a set of models was trained using the AutoGluon Tabular framework (Erickson et al., 2020). The model selection included models from the XGBoost, CatBoost, Extra Trees, Random Forest, and FastAI NN families (Breiman, 2001; T. Chen et al., 2016; Dorogush et al., 2018; Korsunsky et al., 2019). AutoGluon Tabular runs an automated stacking, combination, cross validation and hyper-parameter tuning of models and reports a ranked list of trained models. We ran AutoGluon Tabular on the "high_quality" setting without custom models and removed the KNN model family. The training time was set to 24 hours. The two top-performing models were selected for reporting ensuring overlap between the training scenarios. In addition to single-cell and aggregate classifications, we implemented a third strategy ('Majority vote') where we first predicted cell death subtypes on single-cell level followed by a majority vote over predicted labels per site. A prediction was assigned as correct if and only if the majority-voted label matched the true label. To evaluate our classification model, we reported F1, Recall, Precision as well as Accuracy. All metrics were reported on macro and micro (weighted) level. Additionally, we generated confusion matrices for all models.

Table 1: Number of samples used for aggregated and Single-Cell classification models across Feature Extractors.

Cell Death Subtype (Label)	CellProfiler		DeepProfiler		DINO	
	Aggregated	**Single Cells**	**Aggregated**	**Single Cells**	**Aggregated**	**Single Cells**
Apoptosis (0)	853	173,437	860	268,678	1,097	37,691
Autophagy (1)	2,375	535,027	2,371	362,324	1,162	38,311
Ferroptosis (2)	1,181	282,551	1,078	157,970	1,098	39,128
Immunogenic Cell Death (3)	598	262,684	843	263,173	845	29,226
Necroptosis (4)	617	251,172	848	99,668	1,002	36,264
Pyroptosis (5)	515	112,237	535	82,212	476	17,500
Total	6,139	1,617,108	6,535	1,234,025	5,680	198,120

3 Results

3.1 Measuring perturbation strength using e-distance enables robust cell filtering

We calculated grit scores and e-distances for sets of single cells with a count of at least 200 cells, treated with the same compound concentration. We used profiles of Dimethyl Sulfoxide (DMSO, solvent control) treated cells from our experiments as a negative control. For all subsequent analysis, DMSO cells were omitted.

Grit calculations resulted in 54.04% of sets kept for CellProfiler features, 53.43% for DeepProfiler features, and 51.56% for DINO features. Notably, four compound concentration sets extracted with CellProfiler showed significantly higher grit scores than in DeepProfiler. Further inspection revealed that these sets did not appear in DeepProfiler features as post-processing of the cells resulted in lower cell counts and subsequent removal (Table 1). The Venn diagram (Figure 1C) shows that 133 sets resulted in a grit > 1.5 for all feature extractors, with 20 sets only in DINO and CellProfiler features. Inspection of the distributions of grit scores further reveals that DINO features result in a significant proportion of sets with grit scores > 4. This indicates strong phenotypical changes when compared to DMSO (Supplementary Figure S2).

Etest and e-distance results revealed similar distributions in distance and significance results between CellProfiler, DeepProfiler and DINO features (Supplementary Figure S2). For CellProfiler, higher absolute distances were observed with a maximum e-distance of 11,590 as compared to 3,090 in DeepProfiler and 420.93 for DINO (Supplementary Figure S1). Due to the non-interpretability of the distances, filtering was performed using results from the corresponding etest. Results show 74.90% of sets as significant for DeepProfiler features, 66.54% significant for CellProfiler features and 71.00% for DINO. Here, 155 common sets were observed between the three feature sets (Figure 1C). It is noteworthy that the number of sets differs between grit and e-distance calculations, as the metrics have different underlying assumptions. Therefore, the intersection of significant etest and grit-scores >1.5 was used as a cut-off for further analysis.

As shown, e-distance effectively filters phenotypically relevant sets across our feature sets, showing consistent distributions and overlapping significant results for CellProfiler, DeepProfiler, and DINO features. The intersection of significant etest results with grit scores >1.5 provides a robust selection criterion, refining the focus to sets with the strongest phenotypic changes and enhancing the reliability of our following analysis tasks.

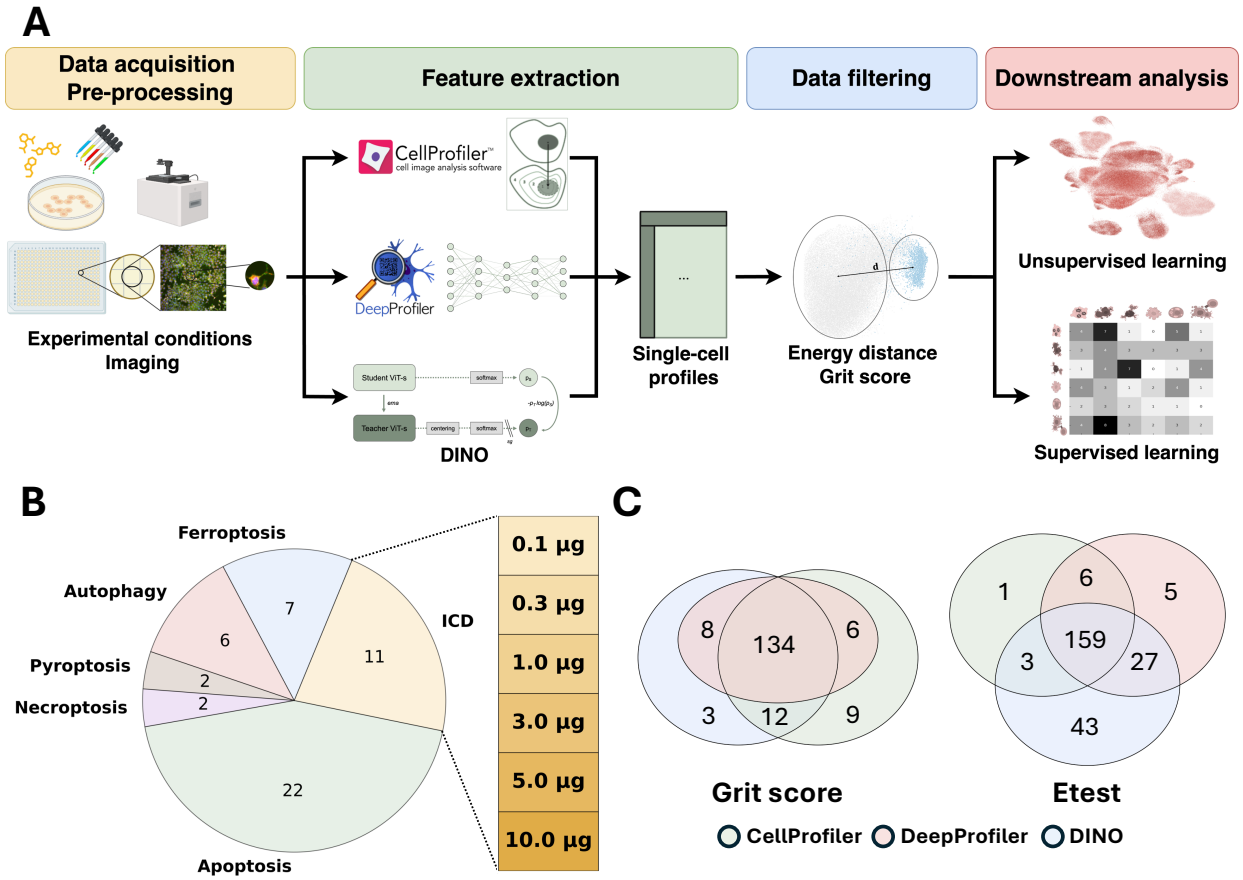

Figure 1: Overview of experimental and analysis pipeline and filtering outcomes. (A) Schematic of the workflow used in this work. Experimental conditions are established, followed by imaging and feature extraction using CellProfiler, DeepProfiler, or DINO. Each extracted feature set undergoes separate data filtering and downstream analysis, including metric calculations, unsupervised and supervised learning. (B) Distribution of compound counts across different of cell death subtypes. Bar chart indicates concentrations used for the different compounds. The same concentrations were used in each compound. (C) Venn diagram depicting number of significant compound concentration sets between CellProfiler, DeepProfiler and DINO. Left diagram shows filtered grit sets, right diagram shows significant etest sets.

3.2 DINO features learn higher resolution morphological patterns on single-cell level

In order to characterize the selected cell death subtypes from a single-cell perspective, we used different feature extraction methods applied to single-cell data as well as aggregated (median cell) data.

We first compared the feature sets based on lower-dimensional UMAP embeddings. Our results show stronger separation of cell death subtypes at the single-cell level using DeepProfiler features compared to CellProfiler (Figure 2A). A clear separation of the Pyroptosis class is observable in both the single-cell UMAP along with the marginal kernel density estimations, as well as in the corresponding correlation matrix (Figure 2C). For CellProfiler features, we observe a positive correlation of profiles from the Pyroptosis, Ferroptosis, and Autophagy classes. For DeepProfiler features, we notice positive correlations between profiles from the Ferroptosis and Autophagy classes, but a clear separation from the Pyroptosis class. Further, a clear cluster formation of the Apoptosis class can be observed in the DeepProfiler embedding. We identified the compound FK866 to be responsible for the formation of this cluster as indicated in Supplementary Figure S3B middle. When comparing the results with the DINO embedding, more noticeable cluster formation on a single-cell level as well as in the correlation maps can be found (Figure 2A and C DINO). Clusters appear to be more dense as indicated by the marginal KDEs in the single-cell UMAP as well as by highly correlated profiles in the corresponding heatmaps (Figure 2). When taking into account the compounds, we observed a clear pattern; individual compounds corresponding to the observed clusters in morphological space (Supplementary Figure S3B). Density estimations of the cell death subtypes in our UMAP spaces revealed an overall more uniform distribution

of cells in both CellProfiler and DeepProfiler features (Supplementary Figure S4A and B), while the subtypes in our DINO embedding space form multiple sub-clusters (Supplementary Figure S4C).

Subsequent analysis of median profiles revealed a clearer separation of cell death subtypes and specific compounds in all three feature sets (Figure 1B and Supplementary Figure S3A). In all feature sets, we observed a region of strong overlap between the cell death subtypes and identified these regions as groups with lower grit and e-distance scores and less separable (Figure 2B). Overall, the differences between feature sets on an aggregated level are less apparent when compared to single-cell level.

One noticeable observation in the DeepProfiler features is the clearly separated cluster in the UMAP embedding (Figure 2A.) The cluster was identified to consist of compounds from all cell death subtypes without any distinct pattern indicative of its formation (Figure 2 DeepProfiler and Supplementary Figure S14). Analysis of the corresponding microscopy images using our *CellViewer* tool did not reveal any patterns in the images that would explain this cluster (Supplementary Figure S14C). Further analysis of the original feature intensities revealed a distinct pattern in all cells originating from that cluster (Supplementary Figure S14B), but due to the non-interpretability of the features, we were limited in our understanding of this anomaly. As a removal of the blob did not result in a change of the remaining morphological space, we kept the points in for all further analysis.

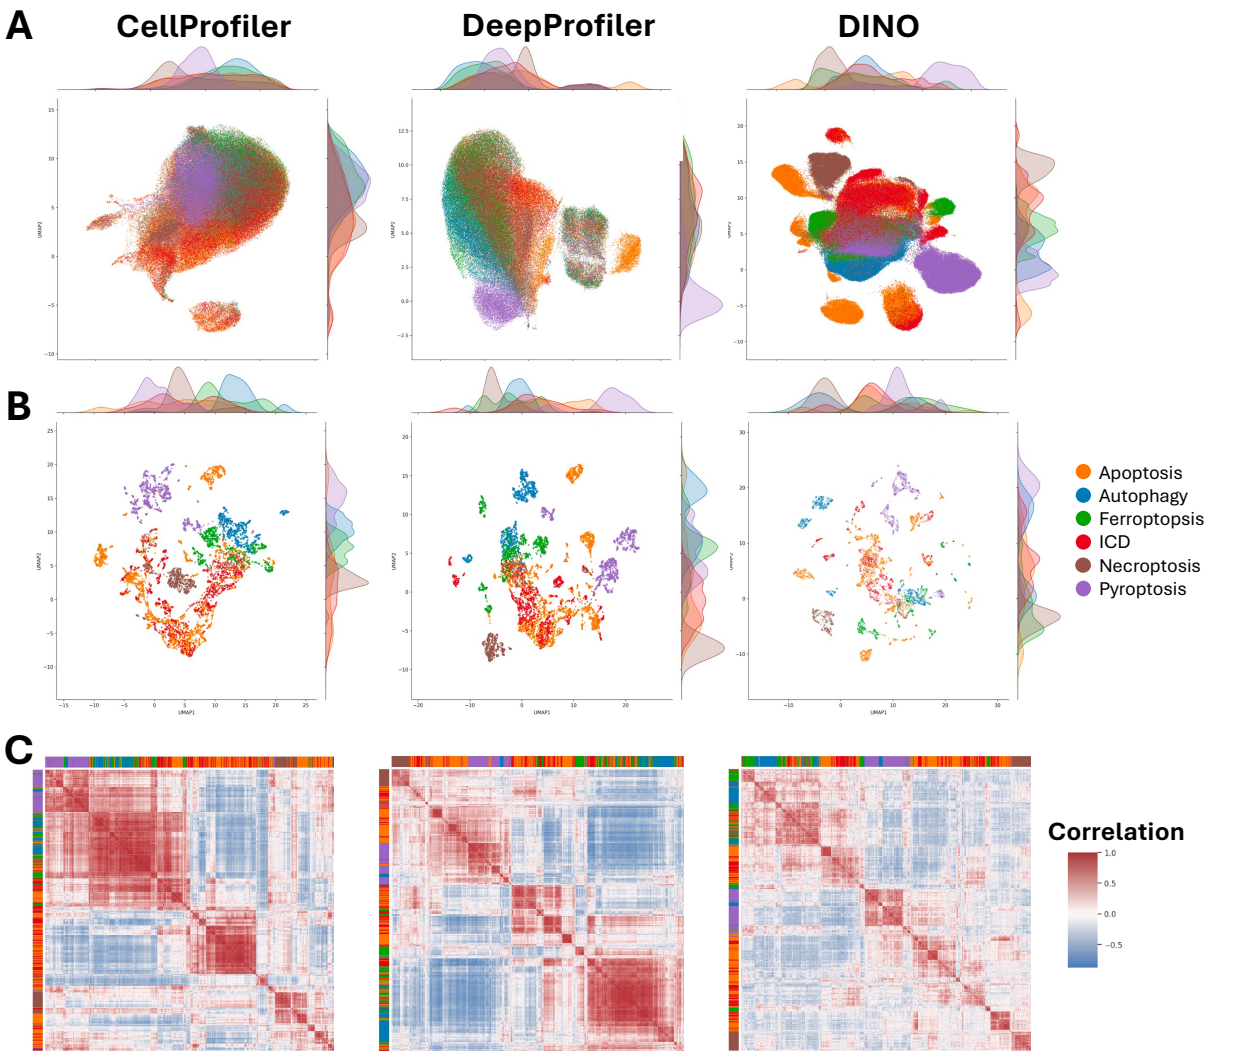

Figure 2: Comparison of morphological profiling across feature extraction methods. (A) UMAP embeddings of cell morphologies extracted using CellProfiler, DeepProfiler, and DINO. Each point represents a single-cell, colored by its cell death subtype. Marginal density plots along the axes display the distribution of cells for each subtype, highlighting clustering patterns unique to each method. (B) UMAP embeddings of aggregated cell morphologies extracted using CellProfiler, DeepProfiler, and DINO. Each point represents a aggregated profile, colored by its cell death subtype. Marginal density plots along the axes display the distribution of cells for each subtype, highlighting clustering patterns unique to each method. (C) Correlation heatmaps comparing pairwise feature similarities among cell treatments. Rows and columns correspond to site-level aggregated profiles, with colors representing correlation values (red for positive and blue for negative correlations). The distribution of cell death subtypes is indicated by colored bars along the heatmap axes.

3.3 Single-cell density analysis reveals biological heterogeneity within apoptotic cells

Building on the previous results, we sought to explore the heterogeneity within our cell death subtypes. For this, we focused on the DINO-derived features. We selected apoptosis for further analysis due to its high variability across compounds, providing a case for studying the cellular heterogeneity.

To investigate, a new UMAP embedding was trained on the single-cell DINO features of the apoptotic cells. The embedding revealed distinct cluster formations within the morphological space (Figure 3 and Supplementary Figure S7). To further understand the emergence of these clusters, we utilized our *CellViewer* tool to examine different regions within the UMAP. Representative cells from each region were identified using a k-medoid algorithm, and their corresponding cell images were annotated directly onto the UMAP (Figure 3A). The analysis revealed that specific regions in the morphological space were associated with distinct cellular morphologies. Notably, cells in the region highlighted in the middle exhibited morphological traits typically associated with apoptosis, including shrinkage,

isolation, nuclear fragmentation, and vesicle formation. Interestingly, this region contained cells originating from a diverse array of compounds (Supplementary Figure S7). Outward from this central apoptotic region, more distinct morphological patterns emerged, highlighting morphological heterogeneity within the apoptosis subtype.

Given this observed heterogeneity, we next investigated its potential origins. As outlined, cells in this experiment were treated with a wide range of compounds known to induce apoptosis. To elucidate the relationship between morphology and molecular mechanisms, compounds were next annotated with their respective targets and which were mapped onto the UMAP. This revealed a connection between distinct regions of the UMAP and specific molecular targets (Figure 3B and Supplementary Figure S8C). For instance, cells in the bottom-right region (Figure 3A) were predominantly treated with compounds targeting Topoisomerases I/II, DNA polymerases, or intercalating DNA agents (Figure 3B). The morphological characteristics in this region reflect these molecular mechanisms specifically by changes in the nucleus channel (indicated by blue color).

Analogous analyses using DeepProfiler and CellProfiler features demonstrated similar trends in morphological subspaces, albeit with less pronounced cluster formation (Supplementary Figure S8A and B). Overall, DINO-derived features provided the clearest differentiation of sets of cells while maintaining the most homogeneity among distinct regions in the morphological space. These findings underscore the ability of single-cell morphological features to reveal fine-grained heterogeneity within cell death subtypes and highlight the superior performance of DINO features in distinguishing regions connected to a molecular target.

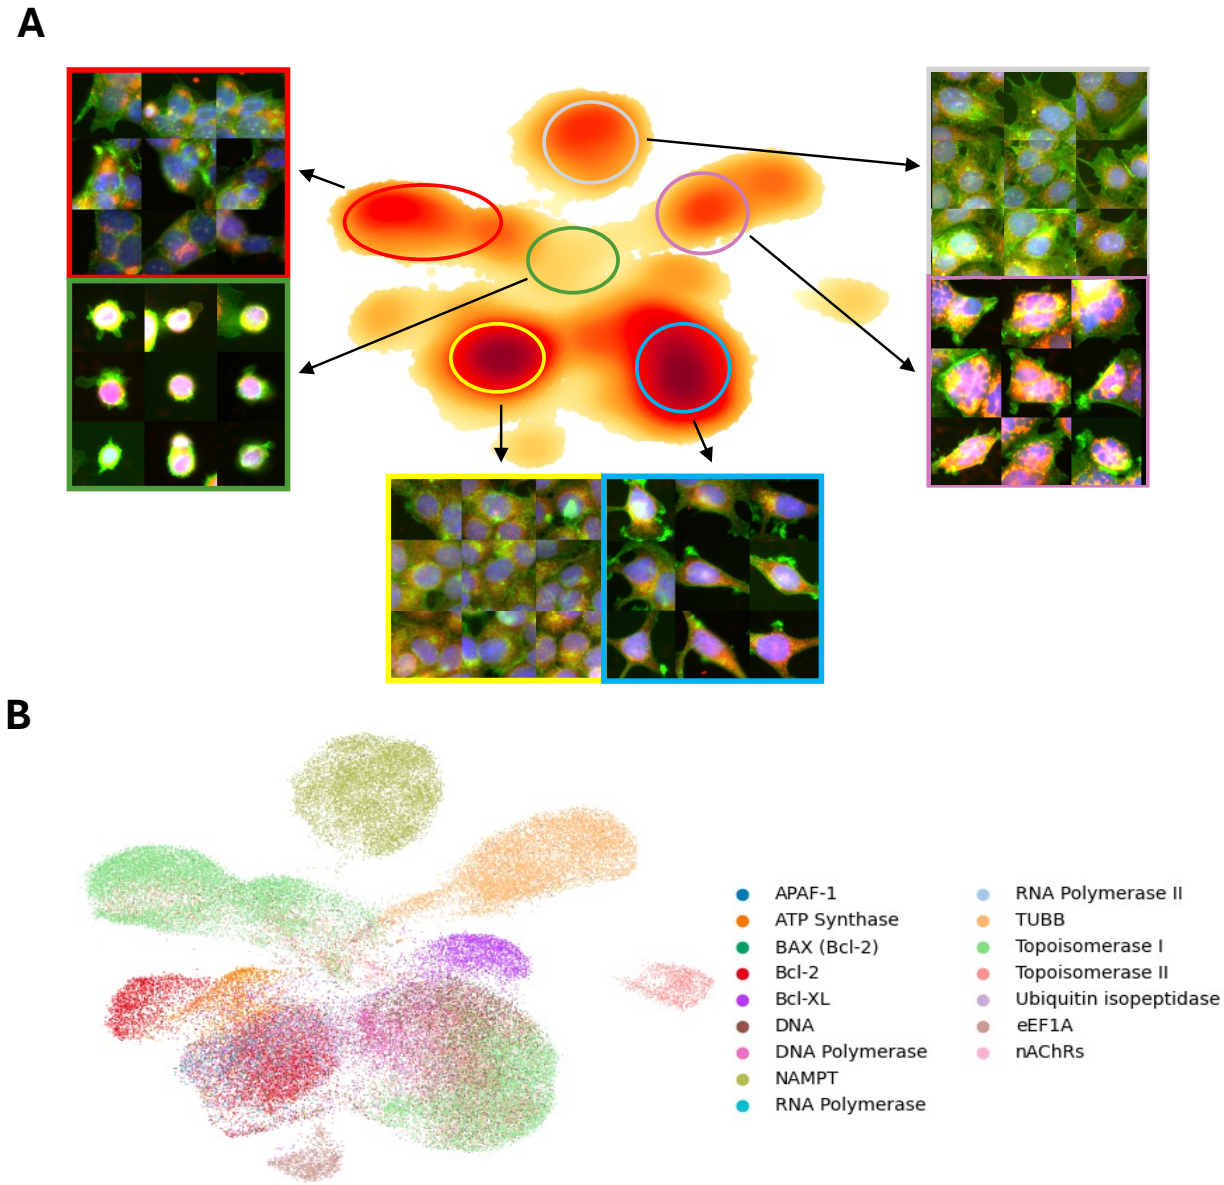

Figure 3: Density mapping of cells from the apoptosis subtype on DINO embedding. (A) UMAP showing the distribution of cell morphologies across treatments in the apoptosis class, with warmer colors (yellow to red) indicating higher densities. Colored circles highlight specific clusters linked to distinct phenotypes. Representative cell images adjacent to the UMAP are overlays of the Hoechst (nucleus, blue), Mitotracker (red/yellow) and Phalloidin (Cytoskeleton/ actin, green) channels. (B) Single-cell UMAP embedding of the apoptosis class colored by their annotated targets. Specific target, compound matching can be found in Supplementary Table 1.

3.4 Single-cell analysis reveal dose-dependent morphological effects

We next investigated how compound concentration influences morphological changes on single-cell level. We projected chemical perturbations at varying concentrations onto the morphological space of the corresponding cell death subtypes. As a showcase, we focused on the apoptosis class, which contained the highest abundance of compounds. We further selected three compounds (Cladribine, SN 38, and Topotecan) to examine their dose-dependent morphologies. We were able to detect measurable shifts in morphology within the morphological space, driven by compound dosage (Figure 4). Notably, SN 38 exhibited the most prominent morphological changes, with cells shifting from a cluster in the bottom right to a region in the top left of the apoptosis-specific morphological space, based on DINO features (Figure 4B). We used the *CellViewer* tool to identify key representative cells across concentrations. Representative cells highlighted this transition, showing a progression from triangular-shaped cells with large nuclei at the lowest concentration (0.1 μg) to elongated cells at the highest concentration (5.0 μg) as well as cytoskeletal re-organization. A similar dose-dependent effect was observed for Cladribine (Figure 4C), where cells at the lowest concentration resembled untreated MCF-7 cells. With increasing concentrations, cells exhibited a larger nucleus, stronger mitochondrial signals, and a lack of neighboring cells. However, not all compounds demonstrated pronounced changes. For instance, Topotecan showed minimal morphological alterations at lower concentrations, with notable shifts only emerging at the highest tested concentration.

To further explore how we could use dose-response effects in our understanding of cell death, we examined features derived from the CellProfiler suite. Specific cellular features have previously been used to investigate effects relevant to cell death and the cell cycle, such as the intensity of the nucleus (Mueller et al., 2006). To demonstrate its utility, we analyzed a range of compounds and their corresponding feature distributions in the integrated DNA intensity across concentrations (Supplementary Figure S15). Our findings revealed clear shifts in the intensity as compound concentrations are increased, indicating a potential change in cell cycle composition. These results underscore how morphological features, particularly those derived from CellProfiler, can provide quantitative insights into dose-dependent cellular responses, specifically leveraging information provided by single cells.

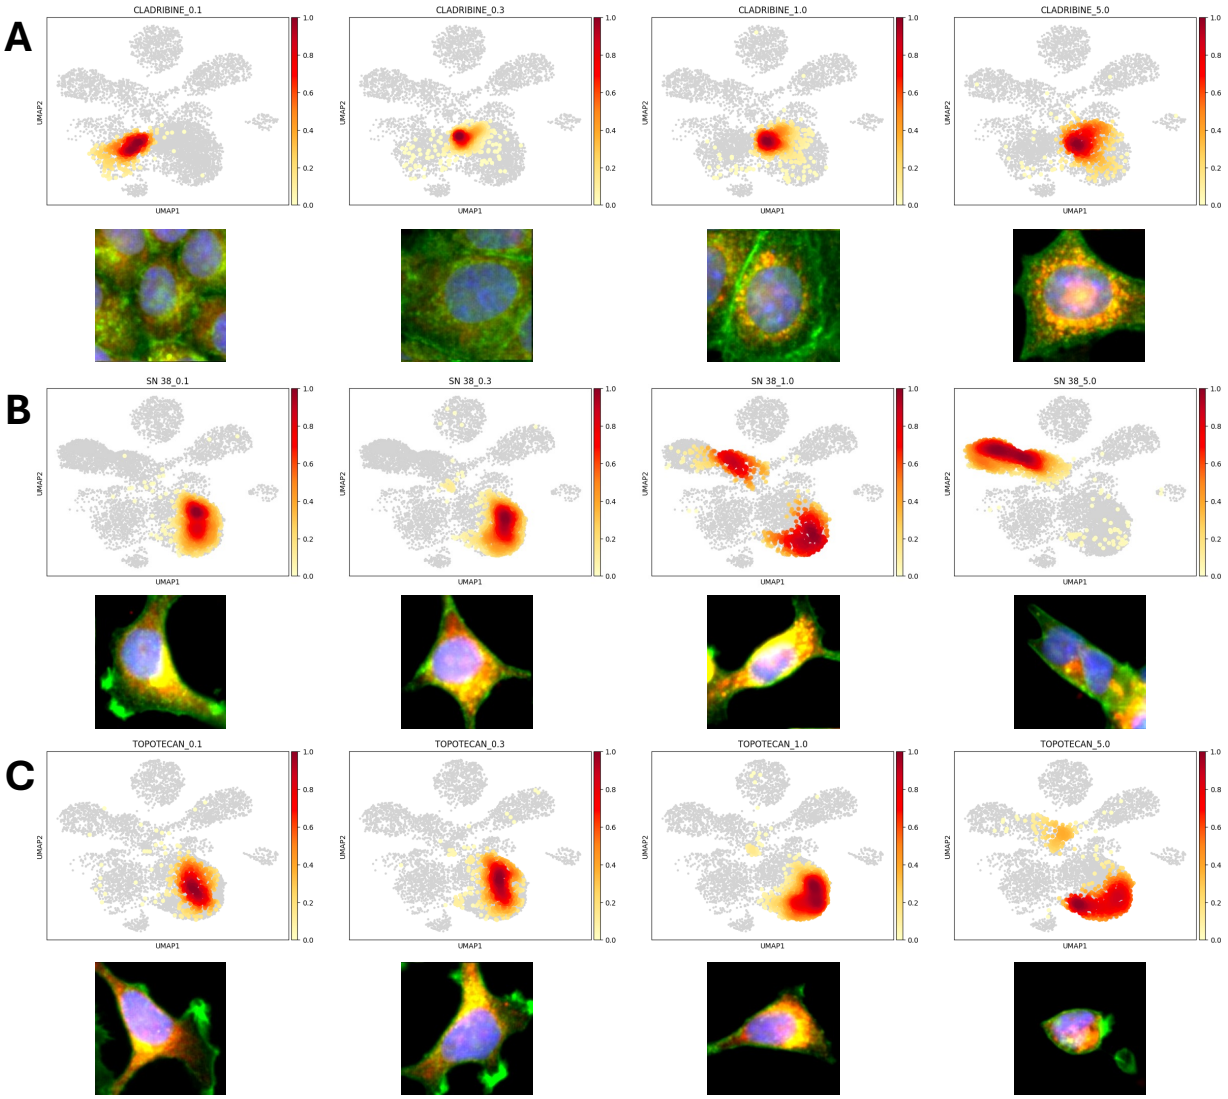

Figure 4: UMAP representation of single-cell morphological profiles treated with varying concentrations of compounds, accompanied by representative cell images for the apoptosis class in DINO features. Each row shows UMAP projections of cells treated with increasing concentrations (0.1, 0.3, 1.0, 3.0, and 5.0 µM). UMAP plots highlight the distribution of treated cells (colored) relative to the remaining cells in the apoptosis. Representative cell images below each UMAP depict an overlay of nuclei (blue), cytoskeleton (green), and mitochondria (red/yellow). Cells were identified with *CellViewer*. (A) UMAP projections and representative images of cells treated with Cladribine. (B) UMAP projections and representative images of cells treated with SN 38. (C) UMAP projections and representative images of cells treated with Topotecan.

3.5 Aggregate profiles outperform single-cells in cell death classification

After identifying distinct cell death subtypes through unsupervised analysis, we proceeded to evaluate whether single-cell morphological profiles could be utilized to accurately classify cell death subtypes. We assessed their performance via single-cell, aggregated, and majority-voted strategies to identify the most effective approaches for this classification task.

At the single-cell level, our analysis demonstrated that a bagged LightGBMXT model trained on DeepProfiler features achieved the highest classification performance, with a macro F1 score of 77.30% (Table 2). Comparatively, the top-performing model trained on CellProfiler features achieved a macro F1 score of 77.04%, showing a slight decrease in overall accuracy but comparable F1 performance. Interestingly, models utilizing DINO features exhibited significantly lower performance at the single-cell level. The highest macro F1 achieved for DINO features 48.78%, representing a performance drop of up to 32.74% compared to the LightGBMXT model trained on DeepProfiler features. When evaluating the classification performance on aggregated profiles, the results revealed improved metrics across all feature

extraction methods. The LightGBMXT model trained on DeepProfiler features attained a macro F1 score of 90.00%, with a comparable performance to the model trained on CellProfiler features, which achieved a macro F1 score of 91.49% (Table 2). In contrast, DINO features continued to under perform, with the NeuralNet_FastAI model achieving the highest macro F1 of 60.80% and the LightGBMXT model following closely with an accuracy of macro F1 of 59.70% (Table 2).

Recognizing that single-cell profiles are rarely analyzed in isolation but are part of broader distributions of data points, we explored an alternative strategy to enhance classification performance. We computed predictions from single-cell classifiers and aggregated these predictions using a majority-vote strategy at the site level. This approach significantly improved the performance of our top-performing single-cell model. The LightGBMXT model trained on DeepProfiler features saw a performance boost of 9.48%, achieving an accuracy of 86.60% and a macro F1 score of 86.98% (Table 2). Similar performance improvements were observed across all remaining models, demonstrating the effectiveness of this aggregation strategy. However, despite the observed gains, none of the majority-voted models surpassed the top-performing models trained on aggregated profiles.

To further investigate model performance, we evaluated the confusion matrices for our classification models. Across the models from CellProfiler and DeepProfiler, we observed a large fraction of misclassification between the ICD class (label 3) and apoptosis class (label 0) (Supplementary Figures S11-S13). Analysis of our compounds revealed structural homology between a few compounds in these classes; Etoposide (apoptosis) and Teniposide (ICD), as well as Doxorubicin/ Epirubicin (apoptosis) and Daunorubicin (ICD). We noticed strong similarities among their induced morphological effects, which could explain the misclassification of these specific classes (Supplementary Figure S4). This trend was not observed in features extracted with DINO. Instead, we observed overall higher misclassification rates among classes, which reflect our previous findings.

In an attempt to validate these results, we visualised the distribution of cell death classes in the UMAP spaces through densities. Here, we observed homogeneity among classes for both CellProfiler and DeepProfiler features (Supplementary Figure S4A and B) , while DINO features exhibited significantly stronger spread (Supplementary Figure S4C). We point out the overlap of the apoptosis and ICD classes in the UMAP space for DeepProfiler and CellProfiler, further hinting at our theory of similar compounds between the two classes (Supplementary Figure S4A and B).

Table 2: Comparison of Performance Metrics across Feature Sets with Standard, Majority Vote, and Aggregated Results. Each column depicts Accuracy, Precision, Recall, and F1 scores (all macro). Bold indicates the highest score per metric across models and strategies.

Model	**Strategy**	Accuracy	Precision	Recall	F1
CellProfiler (CP)					
LightGBMXT_{L2}	Single-cell	0.7476	0.7666	0.7780	0.7703
	Majority vote	0.8137	0.8398	0.8227	0.8443
	Aggregated	0.8954	**0.9335**	**0.9001**	**0.9149**
$\text{NeuralNetFastAI}_{L1}$	Single-cell	0.7220	0.7574	0.7456	0.7500
	Majority vote	0.8111	0.8487	0.8296	0.8375
	Aggregated	0.7865	0.8265	0.8028	0.8019
DeepProfiler (DP)					
LightGBMXT_{L2}	Single-cell	0.7723	0.7840	0.7660	0.7730
	Majority vote	0.8660	0.9040	0.8554	0.8698
	Aggregated	**0.8997**	0.9260	0.8966	0.9000
$\text{NeuralNetFastAI}_{L1}$	Single-cell	0.7712	0.7800	0.7662	0.7710
	Majority vote	0.8652	0.9000	0.8553	0.8685
	Aggregated	0.7986	0.8452	0.8017	0.8091
DINO					
LightGBMXT_{L2}	Single-cell	0.4889	0.4900	0.4857	0.4878
	Majority vote	0.5876	0.5959	0.5799	0.5853
	Aggregated	0.6020	0.5970	0.5987	0.5970
$\text{NeuralNetFastAI}_{L1}$	Single-cell	0.4449	0.4510	0.4472	0.4469
	Majority vote	0.5750	0.5883	0.5642	0.5714
	Aggregated	0.6100	0.6073	0.6117	0.6080

4 Discussion

Analyzing data on a single-cell level enables studying heterogeneous cellular behaviors, identifying rare subpopulations, and uncovering responses to perturbations that would otherwise be masked in aggregated analyses. In this work, we investigated how single-cell data from Cell Painting experiments can be leveraged in the data analysis, demonstrated on a novel dataset specifically designed to study cell death subtypes. We compared three feature extraction approaches on varying aggregation levels and investigated which factors contribute to different types of cell death for supervised and unsupervised downstream analysis tasks.

Removal of perturbations without meaningful morphological changes is a crucial step in any profiling analysis. Specifically, when working on a single-cell level, minimizing noise in the form of unaffected cells becomes a crucial part of the analysis. The current version of the grit score is impractical to compute at the single-cell level for larger datasets, restricting its applicability in high-throughput experiments. For instance, on our cluster with 256 GB of RAM, grit scores could not be reliably computed for over 15,000 cells in one go, greatly curtailing its applicability in such experiments. To complement grit scores, we evaluated e-distance, a computationally more efficient metric suitable for large-scale single-cell datasets (Peidli et al., 2024). Our results showed that e-distance and grit scores produced comparable filtering outcomes, suggesting that either metric, or a combination, could be effectively used to identify treatments inducing significant morphological changes (Figure 1C). The choice of filtering metric ultimately depends on the experimental context and availability of computational resources, with e-distance offering scalability for single-cell data while grit scores provide robustness at the aggregated level.

Unsupervised analysis illustrated that single-cell profiles offer a powerful means to study how different types of cell death emerge, revealing both subpopulation dynamics and distinct morphological characteristics (Figures 2 and 3). A notable observation across all feature extractors was the non-homogeneity of cell death subtypes. Using the apoptosis class as an example, we were able to identify distinct cellular populations within the morphological space, each associated with specific morphological traits (Figure 3). Indeed, apoptosis encompasses a wide range of intrinsic mechanisms (Dehghan et al., 2023; Kerr et al., 1972)), leading to diverse morphological effects, such as the formation of apoptotic bodies, decrease in cellular volume, enlargement of the ER, and membrane blebbing or blistering, among others (Y. Chen et al., 2024; Kerr et al., 1972), which may explain the diversity of phenotypes that we identify.

To validate the identified apoptosis subpopulations, we mapped compound targets to these regions in the morphological space. Strikingly, this revealed a clear correlation between the subpopulations and specific target annotations, suggesting that distinct molecular mechanisms drive the observed morphological diversity (Figure 3B). Compounds targeting similar pathways or molecular components clustered together within the morphological embedding, providing an additional layer of biological interpretation to the observed heterogeneity. Moreover, we demonstrated that dose-dependent changes in morphology further contribute to heterogeneity, as compounds induce distinct morphological effects only at specific concentration levels (Figure 4 and Supplementary Figures S5 and S6) and certain cell cycle stages (Supplementary Figure S15).

The choice of feature extractor plays a crucial role in shaping the observed morphological patterns. Deep learning-based feature extraction methods, such as DeepProfiler and DINO, leverage fundamentally different architectures to learn features from imaging data. DeepProfiler employs a weakly supervised approach explicitly trained to differentiate treatments (Moshkov et al., 2024), while DINO uses a self-supervised paradigm to learn representations by minimizing a loss between teacher and student networks, focusing on local and global features (Caron et al., 2021). Our results indicate that DINO tends to capture more intricate features, as evidenced by dense clustering of similar treatments in both UMAP embeddings and correlation maps (Figures 2, 3 and Supplementary Figure S7).

As deep learning based features are inherently difficult to interpret biologically, it becomes necessary to use additional approaches like our mapping of underlying imaging data to validate their relevance. To further extend on this, explainability remains a critical consideration for deep learning approaches, especially in biological and medical applications (Tanevski et al., 2022). While tools like CellProfiler provide interpretable features, neural network-based methods often require techniques like GradCAM (Selvaraju et al., 2019) or visualization of attention heads in vision transformers (Caron et al., 2021) to understand what the model is focusing on during training and inference. In our evaluation of attention maps, we observed that the DINO network captured biologically meaningful aspects of the cells, emphasizing the validity of deep learning models for extracting features. We observed how the network learned features correlated to specific channels and subcellular features (Supplementary Figures S9 and S10). However, we point to anomalies, such as the unexplained cluster in the DeepProfiler features (Supplementary Figure S14), which suggest limitations in deep-learning approaches. We hypothesized that this cluster could originate from differences in the training set (using U2OS and A549 cell lines) compared to the MCF-7 cells we used for inference here in our study (Moshkov et al., 2024). This underscores the importance of domain-specific fine-tuning, when using larger pre-trained models and careful evaluation of obtained results.

Despite the strengths of single-cell profiles in uncovering subpopulation dynamics and morphological heterogeneity, several limitations must be acknowledged. While our supervised classification tasks achieved macro F1 scores of up to 77%, our single-cell classification models consistently underperformed compared to majority-voting and aggregated profile approaches. This performance gap may be attributed to an oversimplification of underlying population heterogeneity when relying solely on single-cell data and treating them as individual data points. In an attempt to consider treatments more as populations, we implemented a simple majority voting, reaching similar performance levels compared to the aggregated (Table 2). Recently, more advanced modeling approaches such as optimal transport, archetype analysis, and deep sets (Baran et al., 2019; Bunne et al., 2023; Dijk et al., 2024; Dong et al., 2023) have emerged, offering valuable insights into perturbation effects and cellular responses at the same time. These methods hold promise for future addressing the challenges associated with single-cell analysis by taking underlying distributions and their changes into consideration.

In addition to the aggregation levels, significant classification performance differences were observed between CellProfiler/DeepProfiler and DINO models, with disparities of up to 30% in accuracy (Table 2). One possible factor is the difference in single-cell dataset sizes used for training (Table 1), as larger datasets have been shown to improve performance by better approximating the underlying data distributions (Althnian et al., 2021; Dosovitskiy et al., 2021). However, since aggregated results across feature extractors were more consistent, other factors likely play a more dominant role. DINO's self-supervised approach, which emphasizes capturing intricate morphological features, leads to clustering of morphologically similar groups with a higher resolution (Figures 2 and 3). This results in greater within-class heterogeneity, making classification tasks more challenging. In contrast, DeepProfiler's weakly supervised training optimizes for treatment differentiation (Moshkov et al., 2024), reducing within-class variability. Similarly, CellProfiler's hand-crafted features, while less flexible, offer greater consistency across classes. We believe that the inability of DINO models to generalize effectively stems from an emphasis on subtle morphological differences that increase within-class variability. These findings underscore the importance of considering both training data size and feature extraction methodology when designing models for morphological profiling. They further highlight the potential trade-offs between capturing detailed biological variability and achieving state-of-the-art performance in classification tasks. Future works could explore strategies to enhance DINO's generalization capabilities, such as incorporating additional pre-training data, fine-tuning on class-specific labels, or using hybrid approaches that combine supervised and self-supervised learning paradigms or contrastive learning frameworks with other data modalities as seen in (Fradkin et al., 2024; Sanchez-Fernandez et al., 2023).

From a computational perspective, the scalability of single-cell morphological profiling remains a challenge. High-throughput analyses require efficient pre-processing and filtering strategies to ensure reliable downstream analyses. In this study, deep learning-based feature extraction methods not only outperformed traditional tools like CellProfiler but also required fewer features to achieve similar or superior results. For example, while CellProfiler generates 1,400 hand-crafted features, DeepProfiler and DINO used only 672 and 384 features, respectively, to deliver comparable or better outcomes. Additionally, inference of features was obtainable in substantially less time with 2 hours per plate in DeepProfiler and 1 hour for DINO, both running on a single NVIDIA GTX 3090. As CellProfiler is as of now not capable of running feature extraction on GPUs, the inference times reached > 4 hours per plate. Overall, these computational cost reductions significantly alleviate computational bottlenecks, making deep learning frameworks suitable for large-scale, high-throughput applications. A recently published approach can be found with the SPACe framework, which combines efficient data processing with interpretable and reduced feature sets, when compared to standard CellProfiler pipeline (Stossi et al., 2024).

Single-cell morphological data can offer invaluable insights into cellular heterogeneity and subpopulation dynamics, complementing aggregated profiles for exploratory and hypothesis-driven analyses. However, we showed that for its effective use, rigorous methodological choices are required, including feature extraction strategies and filtering metrics, tailored to the specific biological context. By integrating computationally efficient and scalable analysis methods, single-cell morphological profiling can be further refined to address its current limitations and unlock its full potential for studying complex cellular processes.

5 Authors contributions

BF: Design, methodology, analysis, visualization, interpretation, manuscript writing; DH: Methodology, analysis; PB: Methodology, analysis; EB: Methodology, interpretation; PG: Cell Painting experiments; MJ: Cell Painting experiments; PH: Methodology, interpretation; DR: Software development; JR: Cell Painting experiments, interpretation; JCP: Design, conceptualization, interpretation, manuscript writing, supervision; OS: Design, conceptualization, interpretation, manuscript writing, supervision; All of the authors reviewed, edited, and contributed to discussions on the manuscript and approved the final version of the manuscript.

6 Funding

O.S. acknowledges funding from the Swedish Research Council (Grants 2020-03731, 2020-01865, 2024-04576, and 2024-03566), FORMAS (Grant 2022-00940), Swedish Cancer Foundation (22 2412 Pj 03 H), and Horizon Europe (Grant Agreements 101057014 (PARC) and 101057442 (REMEDI4ALL)).

7 Acknowledgements

Compounds were provided by Beactica Therapeutics. The computations were enabled by resources provided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS), partially funded by the Swedish Research Council through grant agreement no. 2022-06725, resources provided by Uppsala University at UPPMAX, and the Berzelius resource provided by the Knut and Alice Wallenberg Foundation at the National Supercomputer Centre.

8 Conflict of interest

J.C.P. and O.S. declare ownership in Phenaros Pharmaceuticals.

9 Data and Code availability

The code to run the analysis and generate the figures can be found on Github (https://github.com/pharmbio/sc_-celldeath). The feature profiles, metrics and splits to run the analysis and supervised training can be found on FigShare: https://doi.org/10.17044/scilifelab.28202864.

References

Althnian, Alhanoof, Duaa AlSaeed, Heyam Al-Baity, Amani Samha, Alanoud Bin Dris, Najla Alzakari, Afnan Abou Elwafa, and Heba Kurdi (Jan. 2021). "Impact of Dataset Size on Classification Performance: An Empirical Evaluation in the Medical Domain". *Applied Sciences*, 11: 2, p. 796.

Baran, Yael, Akhiad Bercovich, Arnau Sebe-Pedros, Yaniv Lubling, Amir Giladi, Elad Chomsky, Zohar Meir, Michael Hoichman, Aviezer Lifshitz, and Amos Tanay (Oct. 2019). "MetaCell: analysis of single-cell RNA-seq data using K-nn graph partitions". *Genome Biology*, 20: 1, p. 206.

Bray, Mark-Anthony, Shantanu Singh, Han Han, Chadwick T. Davis, Blake Borgeson, Cathy Hartland, Maria Kost-Alimova, Sigrun M. Gustafsdottir, Christopher C. Gibson, and Anne E. Carpenter (Sept. 2016). "Cell Painting, a high-content image-based assay for morphological profiling using multiplexed fluorescent dyes". *Nature Protocols*, 11: 9, pp. 1757–1774.

Breiman, Leo (Oct. 2001). "Random Forests". *Machine Learning*, 45: 1, pp. 5–32.

Bunne, Charlotte, Stefan G. Stark, Gabriele Gut, Jacobo Sarabia del Castillo, Mitch Levesque, Kjong-Van Lehmann, Lucas Pelkmans, Andreas Krause, and Gunnar Rätsch (Nov. 2023). "Learning single-cell perturbation responses using neural optimal transport". *Nature Methods*, 20: 11, pp. 1759–1768.

Caicedo, Juan C, Sam Cooper, Florian Heigwer, Scott Warchal, Peng Qiu, Csaba Molnar, Aliaksei S Vasilevich, Joseph D Barry, Harmanjit Singh Bansal, Oren Kraus, et al. (Sept. 2017). "Data-analysis strategies for image-based cell profiling". *Nature Methods*, 14: 9, pp. 849–863.

Caron, Mathilde, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin (May 2021). *Emerging Properties in Self-Supervised Vision Transformers*. arXiv, 2021. DOI: 10.48550/arXiv.2104.14294.

Carpenter, Anne E., 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. (Oct. 2006). "CellProfiler: image analysis software for identifying and quantifying cell phenotypes". *Genome Biology*, 7: 10, R100.

Chandrasekaran, Srinivas Niranj, Jeanelle Ackerman, Eric Alix, D. Michael Ando, John Arevalo, Melissa Bennion, Nicolas Boisseau, Adriana Borowa, Justin D. Boyd, Laurent Brino, et al. (Mar. 2023). *JUMP Cell Painting dataset: morphological impact of 136,000 chemical and genetic perturbations*. bioRxiv, 2023. DOI: 10.1101/2023.03.23.534023.

Chen, Tianqi and Carlos Guestrin (Aug. 2016). "XGBoost: A Scalable Tree Boosting System" *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 785–794.

Chen, Yao, Xiaohua Li, Minfeng Yang, and Song-Bai Liu (May 2024). "Research progress on morphology and mechanism of programmed cell death". *Cell Death & Disease*, 15: 5, pp. 1–13.

Conrad, Marcus, José Pedro Friedmann Angeli, Peter Vandenabeele, and Brent R. Stockwell (May 2016). "Regulated necrosis: disease relevance and therapeutic opportunities". *Nature Reviews Drug Discovery*, 15: 5, pp. 348–366.

Dehghan, Sepehr, Nasim Kheshtchin, Shaghayegh Hassannezhad, and Maryam Soleimani (Dec. 2023). "Cell death classification: A new insight based on molecular mechanisms". *Experimental Cell Research*, 433: 2, p. 113860.

Dijk, Robert van, John Arevalo, Mehrtash Babadi, Anne E. Carpenter, and Shantanu Singh (Nov. 11, 2024). "Capturing cell heterogeneity in representations of cell populations for image-based profiling using contrastive learning". *PLOS Computational Biology*, 20: 11, e1012547.

Dong, Mingze, Bao Wang, Jessica Wei, Antonio H. de O. Fonseca, Curtis J. Perry, Alexander Frey, Feriel Ouerghi, Ellen F. Foxman, Jeffrey J. Ishizuka, Rahul M. Dhodapkar, et al. (Nov. 2023). "Causal identification of single-cell experimental perturbation effects with CINEMA-OT". *Nature Methods*, 20: 11, pp. 1769–1779.

Dorogush, Anna Veronika, Vasily Ershov, and Andrey Gulin (Oct. 2018). *CatBoost: gradient boosting with categorical features support*. arXiv, 2018. DOI: 10.48550/arXiv.1810.11363.

Doron, Michael, Théo Moutakanni, Zitong S. Chen, Nikita Moshkov, Mathilde Caron, Hugo Touvron, Piotr Bojanowski, Wolfgang M. Pernice, and Juan C. Caicedo (June 2023). *Unbiased single-cell morphology with self-supervised vision transformers*. bioRxiv, 2023. DOI: 10.1101/2023.06.16.545359.

Dosovitskiy, Alexey, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. (June 2021). *An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale*. arXiv, 2021. DOI: 10.48550/arXiv.2010.11929.

Erickson, Nick, Jonas Mueller, Alexander Shirkov, Hang Zhang, Pedro Larroy, Mu Li, and Alexander Smola (Mar. 2020). *AutoGluon-Tabular: Robust and Accurate AutoML for Structured Data*. arXiv, 2020. DOI: 10.48550/arXiv.2003.06505.

Fradkin, Philip, Puria Azadi, Karush Suri, Frederik Wenkel, Ali Bashashati, Maciej Sypetkowski, and Dominique Beaini (Sept. 2024). *How Molecules Impact Cells: Unlocking Contrastive PhenoMolecular Retrieval*. arXiv, 2024. DOI: 10.48550/arXiv.2409.08302.

Francisco Rodríguez, María Andreína, Jordi Carreras Puigvert, and Ola Spjuth (Dec. 2023). "Designing microplate layouts using artificial intelligence". *Artificial Intelligence in the Life Sciences*, 3: p. 100073.

Gupta, Ankit, Zoe Wefers, Konstantin Kahnert, Jan N. Hansen, Will Leineweber, Anthony Cesnik, Dan Lu, Ulrika Axelsson, Frederic Ballllosera Navarro, Theofanis Karaletsos, et al. (Dec. 2024). *SubCell: Vision foundation models for microscopy capture single-cell biology*. bioRxiv, 2024. DOI: 10.1101/2024.12.06.627299.

Harrison, Philip John, Ankit Gupta, Jonne Rietdijk, Håkan Wieslander, Jordi Carreras-Puigvert, Polina Georgiev, Carolina Wählby, Ola Spjuth, and Ida-Maria Sintorn (July 2023). "Evaluating the utility of brightfield image data for mechanism of action prediction". *PLOS Computational Biology*, 19: 7, e1011323.

Haslum, Johan Fredin, Charles Lardeau, Johan Karlsson, Riku Turkki, Karl-Johan Leuchowius, Kevin Smith, and Erik Müllers (Apr. 2023). *Cell Painting-based bioactivity prediction boosts high-throughput screening hit-rates and compound diversity*. bioRxiv, 2023. DOI: 10.1101/2023.04.03.535328.

He, Kaiming, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick (Dec. 2021). *Masked Autoencoders Are Scalable Vision Learners*. arXiv, 2021. DOI: 10.48550/arXiv.2111.06377.

Kerr, J. F. R., A. H. Wyllie, and A. R. Currie (Aug. 1972). "Apoptosis: A Basic Biological Phenomenon with Wideranging Implications in Tissue Kinetics". *British Journal of Cancer*, 26: 4, pp. 239–257.

Kim, Vladislav, Nikolaos Adaloglou, Marc Osterland, Flavio M. Morelli, Marah Halawa, Tim König, David Gnutt, and Paula A. Marin Zapata (Jan. 2024). *Self-supervision advances morphological profiling by unlocking powerful image representations*. bioRxiv, 2024. DOI: 10.1101/2023.04.28.538691.

Korsunsky, Ilya, Nghia Millard, Jean Fan, Kamil Slowikowski, Fan Zhang, Kevin Wei, Yuriy Baglaenko, Michael Brenner, Po-ru Loh, and Soumya Raychaudhuri (Dec. 2019). "Fast, sensitive and accurate integration of single-cell data with Harmony". *Nature Methods*, 16: 12, pp. 1289–1296.

Kraus, Oren, Kian Kenyon-Dean, Saber Saberian, Maryam Fallah, Peter McLean, Jess Leung, Vasudev Sharma, Ayla Khan, Jia Balakrishnan, Safiye Celik, et al. (Apr. 2024). *Masked Autoencoders for Microscopy are Scalable Learners of Cellular Biology*. arXiv, 2024. DOI: 10.48550/arXiv.2404.10242.

Kroemer, Guido, Lorenzo Galluzzi, Oliver Kepp, and Laurence Zitvogel (Mar. 2013). "Immunogenic Cell Death in Cancer Therapy". *Annual Review of Immunology*, 31: Volume 31, 2013, pp. 51–72.

Lotfollahi, Mohammad, Anna Klimovskaia Susmelj, Carlo De Donno, Leon Hetzel, Yuge Ji, Ignacio L Ibarra, Sanjay R Srivatsan, Mohsen Naghipourfar, Riza M Daza, Beth Martin, et al. (June 2023). "Predicting cellular responses to complex perturbations in high-throughput screens". *Molecular Systems Biology*, 19: 6, e11517.

Moshkov, Nikita, Michael Bornholdt, Santiago Benoit, Matthew Smith, Claire McQuin, Allen Goodman, Rebecca A. Senft, Yu Han, Mehrtash Babadi, Peter Horvath, et al. (Feb. 2024). "Learning representations for image-based profiling of perturbations". *Nature Communications*, 15: 1, p. 1594.

Mueller, Sandra, Marcus Schittenhelm, Friedemann Honecker, Elke Malenke, Kirsten Lauber, Sebastian Wesselborg, Joerg T. Hartmann, Carsten Bokemeyer, and Frank Mayer (Aug. 2006). "Cell-cycle progression and response of germ cell tumors to cisplatin in vitro". *International Journal of Oncology*, 29: 2, pp. 471–479.

Olofsson, Karl, Valentina Carannante, Madoka Takai, Björn Önfelt, and Martin Wiklund (Aug. 2021). "Single cell organization and cell cycle characterization of DNA stained multicellular tumor spheroids". *Scientific Reports*, 11: p. 17076.

Pahl, Axel, Beate Schölermann, Philipp Lampe, Marion Rusch, Mark Dow, Christian Hedberg, Adam Nelson, Sonja Sievers, Herbert Waldmann, and Slava Ziegler (June 2023). "Morphological subprofile analysis for bioactivity annotation of small molecules". *Cell Chemical Biology*.

Peidli, Stefan, Tessa D. Green, Ciyue Shen, Torsten Gross, Joseph Min, Samuele Garda, Bo Yuan, Linus J. Schumacher, Jake P. Taylor-King, Debora S. Marks, et al. (Jan. 2024). "scPerturb: harmonized single-cell perturbation data". *Nature Methods*, pp. 1–10.

Pfaendler, Ramon, Jacob Hanimann, Sohyon Lee, and Berend Snijder (Jan. 2023). *Self-supervised vision transformers accurately decode cellular state heterogeneity*. bioRxiv, 2023. DOI: 10.1101/2023.01.16.524226.

Rietdijk, Jonne, Tanya Aggarwal, Polina Georgieva, Maris Lapins, Jordi Carreras-Puigvert, and Ola Spjuth (Aug. 2022). "Morphological profiling of environmental chemicals enables efficient and untargeted exploration of combination effects". *Science of The Total Environment*, 832: p. 155058.

Rietdijk, Jonne, Marianna Tampere, Aleksandra Pettke, Polina Georgiev, Maris Lapins, Ulrika Warpman-Berglund, Ola Spjuth, Marjo-Riitta Puumalainen, and Jordi Carreras-Puigvert (Dec. 2021). "A phenomics approach for antiviral drug discovery". *BMC Biology*, 19: 1, p. 156.

Sanchez-Fernandez, Ana, Elisabeth Rumetshofer, Sepp Hochreiter, and Günter Klambauer (Nov. 2023). "CLOOME: contrastive learning unlocks bioimaging databases for queries with chemical structures". *Nature Communications*, 14: 1, p. 7339.

Schorpp, Kenji, Alaa Bessadok, Aidin Biibosunov, Ina Rothenaigner, Stefanie Strasser, Tingying Peng, and Kamyar Hadian (July 2023). "CellDeathPred: a deep learning framework for ferroptosis and apoptosis prediction based on cell painting". *Cell Death Discovery*, 9: 1, pp. 1–11.

Selvaraju, Ramprasaath R., Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra (Dec. 2019). *Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization*. arXiv, 2019. DOI: 10.48550/arXiv.1610.02391.

Shen, Sikou, Yina Shao, and Chenghua Li (Aug. 2023). "Different types of cell death and their shift in shaping disease". *Cell Death Discovery*, 9: 1, pp. 1–12.

Stirling, David R., Madison J. Swain-Bowden, Alice M. Lucas, Anne E. Carpenter, Beth A. Cimini, and Allen Goodman (Dec. 2021). "CellProfiler 4: improvements in speed, utility and usability". *BMC Bioinformatics*, 22: 1, p. 433.

Stockwell, Brent R. (July 2022). "Ferroptosis turns 10: Emerging mechanisms, physiological functions, and therapeutic applications". *Cell*, 185: 14, pp. 2401–2421.

Stossi, Fabio, Pankaj K. Singh, Michela Marini, Kazem Safari, Adam T. Szafran, Alejandra Rivera Tostado, Christopher D. Candler, Maureen G. Mancini, Elina A. Mosa, Michael J. Bolt, et al. (Nov. 2024). "SPACe: an open-source, single-cell analysis of Cell Painting data". *Nature Communications*, 15: 1, p. 10170.

Stringer, Carsen, Tim Wang, Michalis Michaelos, and Marius Pachitariu (Jan. 2021). "Cellpose: a generalist algorithm for cellular segmentation". *Nature Methods*, 18: 1, pp. 100–106.

Tanevski, Jovan, Ricardo Omar Ramirez Flores, Attila Gabor, Denis Schapiro, and Julio Saez-Rodriguez (Apr. 2022). "Explainable multiview framework for dissecting spatial relationships from highly multiplexed data". *Genome Biology*, 23: 1, p. 97.

Tian, Guangyan, Philip J Harrison, Akshai P Sreenivasan, Jordi Carreras-Puigvert, and Ola Spjuth (Dec. 2023). "Combining molecular and cell painting image data for mechanism of action prediction". *Artificial Intelligence in the Life Sciences*, 3: p. 100060.

Way et al., Gregor (2022). *Pycytominer: Data processing functions for profiling perturbations*. 2022. URL: https://github.com/cytomining/pycytominer.

Willis, Clinton, Johanna Nyffeler, and Joshua Harrill (Aug. 2020). "Phenotypic Profiling of Reference Chemicals across Biologically Diverse Cell Types Using the Cell Painting Assay". *SLAS DISCOVERY: Advancing the Science of Drug Discovery*, 25: 7, pp. 755–769.

Wolf, F. Alexander, Philipp Angerer, and Fabian J. Theis (Feb. 2018). "SCANPY: large-scale single-cell gene expression data analysis". *Genome Biology*, 19: 1, p. 15.

Wu, Yueshu, Jun Yang, Gang Xu, Xiaolin Chen, and Xiaochen Qu (Apr. 2024). "Integrated analysis of single-cell and bulk RNA sequencing data reveals prognostic characteristics of lysosome-dependent cell death-related genes in osteosarcoma". *BMC Genomics*, 25: 1, p. 379.

Yao, Wenchao, Xuxu Liu, Yuanhang He, Maolan Tian, Shixin Lu, Qiang Wang, Yi Zheng, Zhenyi Lv, Chenjun Hao, Dongbo Xue, et al. (Dec. 2022). "ScRNA-seq and bulk RNA-seq reveal the characteristics of ferroptosis and establish a risk signature in cholangiocarcinoma". *Molecular Therapy - Oncolytics*, 27: pp. 48–60.

Yu, Pian, Xu Zhang, Nian Liu, Ling Tang, Cong Peng, and Xiang Chen (Mar. 2021). “Pyroptosis: mechanisms and diseases”. *Signal Transduction and Targeted Therapy*, 6: 1, pp. 1–21.