UNIVERSITY OF TARTU
Institute of Computer Science
Computer Science Curriculum

Ali Zeynalli

Cell Cycle Phase Classification from Microscopy Images

Master's Thesis (30 ECTS)

Supervisor:
Dmytro Fishman, PhD

Tartu 2025

Contents

Cell Cycle Phase Classification from Microscopy Images

Abstract:

Accurate classification of cell cycle phases is essential for cancer research and drug discovery. While fluorescence microscopy provides high-contrast, biomarker-specific signals that support precise classification, it relies on staining protocols that limit scalability and compromise cell viability. In contrast, bright-field microscopy offers a label-free, cost-effective alternative but poses challenges due to its lower contrast. This study compares five computational strategies for cell cycle phase classification using fluorescence and bright-field microscopy: traditional feature-based classification, segmentation-based classification, mask-guided classification via segmentation, nuclei patch classification, and nuclei patch classification via segmentation. Results show that fluorescence images support near-perfect classification performance across all methods. For bright-field images, the highest balanced accuracy of 0.770 was achieved using a nuclei patch classification approach with a ResNet-50 backbone, followed closely by mask-guided classification. These findings highlight the potential of deep learning models for accurate cell cycle classification in bright-field microscopy, advancing the potential for scalable applications in biomedical research.

Keywords: Cell cycle phase classification, bright-field microscopy, deep learning

CERCS: P176 Artificial intelligence

Rakutsükli faaside klassifitseerimine mikroskoopia piltidelt.

Lühikokkuvõte:

Rakutsükli faaside täpne klassifitseerimine on oluline vähiuuringutes ja ravimite arendamises. Kuigi fluorestsentsmikroskoopia on täpne klassifitseerimismeetod kõrge kontrastsuse ja biomarkeripõhiste signaalide tõttu, tugineb see värvimisprotokollidele, mis piiravad meetodi edasiarengut ja võivad kahjustada rakkude elujõulisust. Seevastu heleda välja mikroskoopia on märgistusevaba ja kuluefektiivne alternatiiv, kuid siiski esineb väljakutseid tema madalama kontrastsuse tõttu. Käesolev uuring võrdleb viit arvutuslikku strateegiat rakutsükli faaside klassifitseerimiseks fluorestsents- ja heleda välja mikroskoopia abil: traditsiooniline tunnusepõhine klassifitseerimine, segmentatsioonipõhine klassifitseerimine, segmentatsioonimaski põhine klassifitseerimine, rakutuumaosade põhine klassifitseerimine ning rakutuumaosade põhine klassifitseerimine segmentatsiooni kaudu. Tulemused näitavad, et fluorestsentspildid võimaldavad peaaegu täiuslikku klassifitseerimist kõikide meetoditega. Heleda välja piltide puhul saavutati kõrgeimaks tasakaalustatud täpsuseks 0.770, milleks kasutati ResNet-50 põhist rakutuumaosade klassifitseerimismeetodit. Tulemuselt järgmine oli segmentatsioonimaski põhine klassifitseerimine. Need tulemused rõhutavad süvaõppemudelite potentsiaali rakutsükli erinevate faaside klassifitseerimiseks heleda välja mikroskoopia abil, mida saab potentsiaalselt rakendada biomeditsiinilises teadustöös.

Võtmesõnad: Rakutsükli faaside klassifitseerimine, heleda välja mikroskoopia, süvaõpe

CERCS: P176 Tehisintellekt

1. Introduction

Cells are the fundamental units of life. As organisms grow, they undergo a process known as cell division, which produces new cells. Approximately two trillion cells divide daily in the human body [1]. Cell division is vital for the development of organisms and plays a key role in various diseases, particularly cancer. Cancer originates from abnormalities in the cell cycle regulation, resulting in uncontrolled cell division. Therefore, accurate identification of cell cycle phases is crucial in cancer research and drug development [2]. For example, by monitoring changes in cell cycle phase distribution before and after treatment, researchers can identify molecular targets and improve anticancer therapies.

Drug discovery is a complex procedure typically involving high-throughput screening (HTS). This robust methodology uses thousands of micro-plates populated with millions of cultivated cells to evaluate vast libraries of candidate compounds. This approach enables researchers to assess compound effectiveness in diverse cell lines, varying chemical concentrations, and distinct stages of disease progression [3].

Despite the capability of biologists to visually discern cell cycle phases through fluorescent microscopy images, the manual annotation of such large datasets proves to be impractical due to the significant time, financial resources, and labour required. For example, annotating a single 3D microscopy image can take hours, making manual analysis impractical for large-scale studies [4]. Furthermore, manual annotation could be subjective, resulting in inconsistencies and unreliable outcomes. Given these challenges, there is an increasing demand for automated, objective, and reproducible methodologies to analyse cell cycle phases derived from microscopy images [5].

Fluorescence microscopy has traditionally been considered the gold standard in cellular biology, primarily due to its ability to illuminate specific cellular components using fluorescent biomarkers. However, this widely used technique comes with drawbacks. Fluorescence microscopy requires extensive sample preparation, which includes fixation and staining. Such procedures can significantly alter the physiological states of cells, potentially resulting in misleading interpretations regarding cellular behaviour. Furthermore, fluorescence microscopy encounters limitations such as photobleaching, limited availability of fluorescence channels, and higher operational costs.

In contrast, bright-field microscopy utilizes naturally transmitted light, eliminating the need for chemical staining or extensive preparation. Although bright-field imaging generally exhibits lower contrast and poses difficulties in precisely identifying cellular features, it offers distinct advantages. Bright-field microscopy is quicker, less expensive, does not compromise cell viability, and allows repeated time-series imaging of the same cell populations. Consequently, the automated analysis of bright-field images has excellent potential to facilitate the evaluation process, reduce costs, and improve overall drug discovery efficiency. Figure 2 illustrates representative examples from bright-field and fluorescence imaging modalities.

Despite the advantages of bright-field microscopy, limited research focuses on using bright-field microscopy to classify cell cycle phases, as most traditional approaches rely on fluorescence microscopy [6]. Few studies have addressed the challenges of cell cycle classification using

bright-field microscopy, revealing a gap in the literature. The development methodologies based on bright-field microscopy could offer a powerful and cost-effective approach for accurately classifying cell cycle phases, which could have important implications for cancer research and pharmaceutical development.

The primary aim of this thesis is to investigate the capabilities of deep learning models in classifying phases of the cell cycle from microscopy images. The specific objectives of this research are as follows:

- Generate reliable and robust ground-truth labels for cell cycle phases from fluorescence microscopy images, thereby facilitating accurate supervised training.
- Create a comprehensive deep learning pipeline for classifying cell cycle phases from microscopy images, exploring various classification strategies to improve predictive accuracy.
- Evaluate the applicability of models trained on fluorescent microscopy annotations in classifying bright-field microscopy images, potentially reducing the reliance on chemical staining.
- Identify limitations in existing workflows for microscopy image analysis

2. Background

Determining the cell cycle phase is crucial in cancer research and drug discovery, as insights into cell growth and division are essential for evaluating treatment effects and identifying therapeutic targets. In studies involving cancer cell lines, researchers often aim to assess how different compounds influence the distribution of cells across various cell cycle phases. Since these processes occur at the single-cell level, microscopy plays a pivotal role in visualising cellular structures and dynamics over time. This section provides an overview of the biological principles underlying the cell cycle and the relevant microscopy techniques used for analysis and discusses how machine learning and deep learning approaches have been employed to classify cell cycle phases from microscopy images.

2.1 Cell Cycle

The cell cycle describes the sequence of events that a cell undergoes as it grows and divides. The cycle is divided into two major stages: interphase and mitosis. Interphase includes the G1 (gap 1), S (synthesis), and G2 (gap 2) phases [6]. During G1, cells increase in size and prepare to duplicate their genetic material, known as DNA. If conditions are not favourable for division, cells may exit the cycle into a non-dividing resting state called G0 [7]. In the S phase, cells replicate their genetic material, resulting in a doubling of DNA content [6]. The G2 phase involves further growth and preparation for division. The final stage is the M phase (mitosis), during which the replicated genetic material is evenly distributed into two new cells [7]. Figure 1 illustrates the sequence of cell cycle phases.

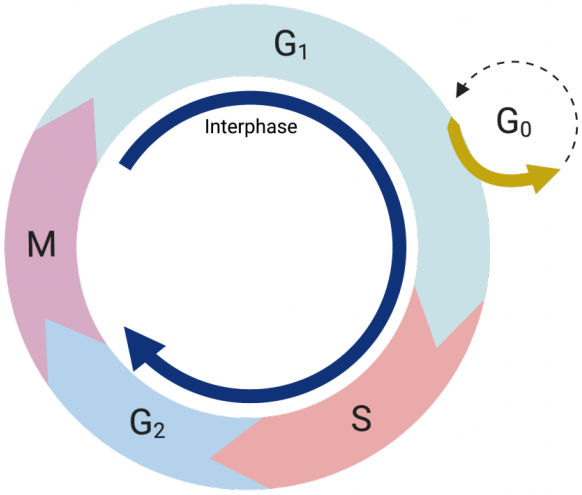

Figure 1. Illustration of the cell cycle phases, depicting the sequential progression through interphase (G1, S, G2), followed by mitosis (M phase). Cells in G1 may also enter a non-dividing resting state in G0. The diagram is adapted from OpenStax [8].

2.2 Microscopy Techniques for Cell Cycle Analysis

Microscopy is the fundamental tool for visualising and analysing cell cycle phases, offering high spatial and temporal resolution [9]. Various microscopy techniques are used in cell biology, each providing distinct advantages regarding resolution, contrast, and the ability to observe live or

fixed cells. Fixed cells are chemically treated to preserve their structure and cellular components, stabilising them for observation.

Among the classical approaches, bright-field microscopy uses transmitted white light, offering limited contrast for live, unstained cells, but is practical for assessing morphology and identifying cycle events in stained samples [10]. In contrast, fluorescence microscopy allows for the specific labelling of cellular components using fluorophores, which enables the quantification of DNA content and tracking of proteins involved in cell cycle regulation [11].

In this thesis, we utilise images acquired from both bright-field and fluorescence microscopy to classify cell cycle phases. Bright-field microscopy provides morphological context, while fluorescence microscopy provides molecular-level information through fluorescent labelling. Figure 2 illustrates representative examples from both imaging modalities. These imaging modalities enable manual observation of cellular morphology and generate large-scale image datasets, which are essential for computational analysis and the development of automated classification methods.

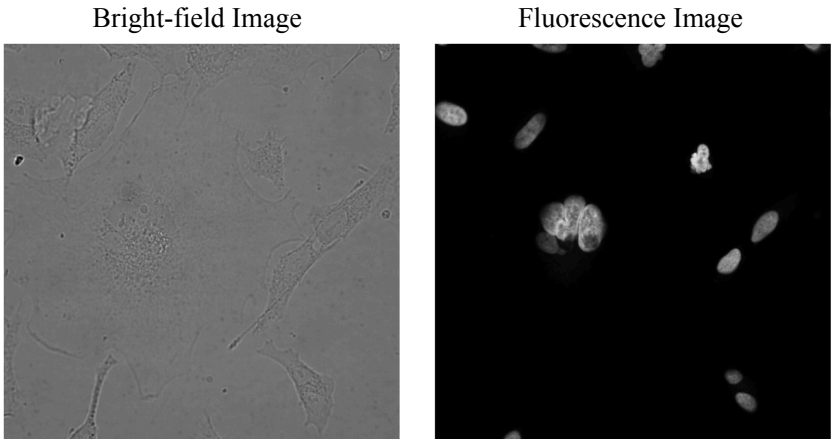

Figure 2. Illustration of two microscopy modalities used for cell cycle analysis. The left image shows a bright-field image, capturing overall cell morphology. The right image shows a fluorescence microscopy image, highlighting molecular structures labelled with a chemical targeting the nuclei of cells.

2.3 Experimental Determination of Cell Cycle Phase

A widely used approach for determining cell cycle phases involves quantifying cellular DNA content using fluorescent DNA-binding dyes [2]. As cells double their DNA during the S-phase, the total DNA content directly reflects the cell cycle stage: G0/G1 phase cells have 2N DNA content, S phase cells exhibit intermediate DNA content due to ongoing replication, and G2+M phase cells have 4N DNA content, where N refers to the amount of DNA in one set of chromosomes [2, 12]. Commonly used dyes include propidium iodide (PI), DAPI, and Hoechst. While PI and DAPI are typically used with fixed cells only, Hoechst can also be applied to live cells [2, 12].

The total DNA content can be quantified at the single-cell level by measuring fluorescence intensity using flow cytometry or fluorescence microscopy. While flow cytometry offers high-

throughput analysis, fluorescence microscopy offers spatial resolution and image-based context, making it more suitable for microscopy-driven computational analysis [2]. In fluorescence microscopy, image analysis pipelines are used to segment nuclei and compute integrated fluorescence intensities for each cell, typically using Hoechst staining. These intensity values are then used to construct DNA content histograms, which typically show distinct peaks corresponding to different cell cycle phases.

Figure 3 illustrates an example of a DNA content histogram derived from fluorescence microscopy data, as presented by Chan et al. [13]. Their analysis used segmentation masks to identify individual nuclei and compute integrated pixel intensities for each cell in the Hoechst fluorescence channel. The resulting intensity values were $\log_2$-transformed and used to determine the centre of the 2N sub-population, corresponding to the G0/G1 phase. The intensity values were then normalised so that the 2N peak was set to the value of 1, and the 4N peak, corresponding to the G2+M phase, was set to the value of 2. This normalisation reflects the doubling of DNA content between the G0/G1 and G2+M phases. Based on these normalised values, cells were classified into five phases: sub-G1 (<0.75), G0/G1 (0.75–1.25), S (1.25–1.75), G2+M (1.75–2.5), and >4N (>2.5) [13].

As the dataset used in this thesis does not contain manual annotations for cell cycle phases, this histogram-based approach is employed as a baseline method for generating phase labels, which are used as ground truth labels.

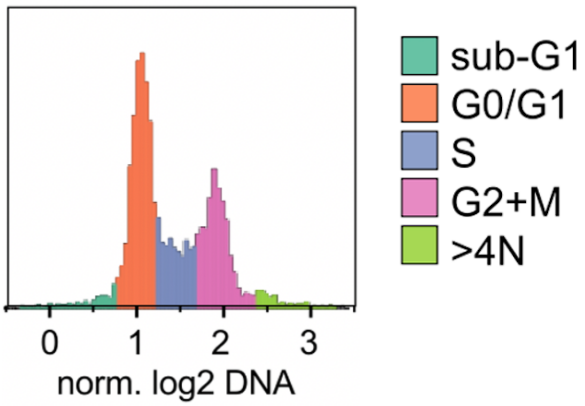

Figure 3. DNA content histogram is obtained from normalised and $\log_2$-transformed fluorescence intensity from individual Hoechst-stained nuclei. Distinct populations corresponding to sub-G1, G0/G1, S, and G2+M, >4N phases are visible. Figure adapted from Chan et al. [13].

Traditionally, cellular analysis of DNA content has relied on specific markers to identify the phases of the cell cycle [14]. These markers target the nucleus and display variations in expression based on the specific phase [15, 16]. While these experimental approaches have proven effective, they can introduce biases associated with the conditions under which they are conducted [17, 18]. This highlights the need for quantitative methodologies that are less invasive and provide a more objective classification of cell cycle phases across various biological and clinical samples. By implementing more standardised and unbiased techniques, researchers can improve the reliability of cell cycle analyses and enhance the comparability of findings across different studies.

2.4 Automated Methods for Cell Cycle Phase Classification

Given the limitations of traditional marker-based techniques, there is a demand for innovative solutions to identify and classify cell cycle phases in microscopy images accurately. We could develop straightforward and unbiased quantitative methods for cell classification based on intrinsic features by utilising commonly observed subcellular components as label-free markers.

To address these challenges, machine learning emerges as a promising computational strategy capable of extracting features associated with cell cycle phases without prior knowledge of which parameters to focus on [14]. Staining label-free techniques have also been developed to determine cell cycle phases, minimising the need for fluorescent probes and allowing cells to be studied in their native states [19, 20].

A growing literature highlights the utility of machine learning in conducting unbiased image analyses for various biological application [14]. This includes successful machine learning-based segmentation and classification across different cell types, such as yeast and mammalian cells [21–24]. Moreover, combining image-based classification with flow cytometry advances single-cell imaging cytometry [25–28], although classifying interphase phases — including G1, S, and G2 — remains challenging [25]. Therefore, the intersection of machine learning with image-based cell cycle phase classification presents a notable advancement in the field.

2.4.1 Feature-Based Classification

One practical way to classify cell cycle phases is to extract and analyse morphological and intensity-based features directly from microscopy images. These features, which often include nuclear size, shape, and fluorescence intensity, capture a broader range of cellular characteristics than DNA histogram-based methods, which focus solely on total DNA content. Consequently, feature-based classification can be applied to fluorescence images stained with dyes such as DAPI or Hoechst and label-free modalities like bright-field microscopy, expanding its utility across diverse experimental setups.

Recent studies underscore the effectiveness of this approach. For example, combining measurements of nuclear area and total fluorescence intensity can robustly distinguish G1/S from G2 phases, reflecting the doubling of DNA content during S phase [14]. Similarly, another study showed that combining both morphological and intensity-based metrics preserves the natural architecture of cells while still providing quantitative estimates of cell cycle stage [29].

Feature-based classification is also highly versatile. It can be performed on fixed cells stained with nuclear dyes or live cells, enabling time-lapse studies of how these features change over the cell cycle. Bright-field offers a cost-effective alternative to staining while reducing potential artefacts introduced by dyes or fluorescent proteins. Once these features are measured, an automated classification algorithm, such as a machine learning model, can assign each cell to a specific phase.

Machine learning (ML) is a branch of artificial intelligence that focuses on developing algorithms capable of learning from data and making predictions or decisions without being explicitly programmed [30]. ML has become a powerful tool for analysing complex biological data, particularly in image-rich domains like microscopy. With the increasing availability of high-

resolution imaging technologies and large-scale datasets, there is a growing need for automated and scalable methods to interpret biological images with high accuracy and reproducibility.

While feature-based classification offers a powerful alternative to methods relying solely on total DNA content, it still depends on manually selected features. A logical next step is to employ approaches that automatically identify the most relevant features directly from raw image data. This is where deep learning, a subset of more advanced ML techniques, comes into play—integrating feature extraction and classification into a single, end-to-end framework. The following subsection explores how deep learning can enhance image-based cell cycle classification.

2.4.2 Deep Learning-Based Classification

A major advancement in machine learning (ML) is deep learning (DL), a subfield that employs artificial neural networks with multiple layers to model complex, non-linear relationships in data. Convolutional neural networks (CNNs), a class of DL models designed specifically for image processing, have proven particularly effective at extracting spatial features from microscopy images [30, 31]. Unlike traditional ML approaches that depend on manually selected features, CNNs automatically learn task-relevant features during training, often resulting in superior performance and broader generalizability.

In microscopy-based research, ML and DL have transformed image analysis by enabling automated segmentation, classification, and object tracking across large datasets. These techniques are applied to various imaging modalities, including bright-field, fluorescence, and confocal microscopy, and they have become essential to high-throughput, quantitative analyses of cellular structures and dynamics [32, 33]. By minimizing the need for manual annotation and interpretation, ML-based pipelines significantly increase reproducibility, efficiency, and scalability in biomedical studies.

Cell cycle classification from microscopy images presents a complex challenge due to the morphological variability of cells across different phases and imaging modalities. Traditional approaches rely on manual feature extraction — such as measuring nuclear size or fluorescence intensity — to distinguish phases like G1, S, G2, and M [34]. These features are then used to train classical classifiers like support vector machines (SVMs) or decision trees [34]. More recent methods employ DL models, notably CNNs, that process raw image data end to end for classification [35]. For example, CNNs trained on DAPI-stained nuclei have achieved high accuracy in differentiating interphase from mitotic stages, in some cases surpassing human-level performance [36].

Furthermore, combining bright-field and fluorescence images as multi-channel inputs can boost classification results by providing morphological and molecular context [37]. Some studies have also incorporated temporal information from live-cell imaging to track phase transitions. Recurrent neural networks (RNNs), including long short-term memory (LSTM) architectures, model these time-dependent patterns to improve prediction consistency [38]. In parallel, explainable AI tools offer interpretability by highlighting which image regions most influenced a classification decision [39].

Despite these advances, several challenges remain. Challenges include dealing with class imbalance, variability in experimental conditions such as staining protocols, and the need to generalize findings across different cell types and imaging modalities. Effectively addressing these issues is essential for deploying ML models in real-world biomedical applications, ensuring robustness and interpretability.

3. Methods

This chapter presents our novel, data-driven methodology for classifying cell cycle phases, designed to address the limitations of traditional microscopy and deep learning techniques. We introduce a comprehensive pipeline for analysing microscopy images that minimises the need for manual annotations and expensive biomarkers. Our approach integrates data cleaning, labelling, preprocessing, classification strategy design, and evaluating both fluorescence and bright-field imaging modalities.

3.1 Dataset

This thesis utilises a dataset provided by Revvity through a collaborative agreement with the University of Tartu. The dataset comprises 1 968 microscopy images, each with a resolution of 1080×1080 pixels, capturing A549 human lung carcinoma cells. Nuclear staining was performed using Hoechst dye, which emits a DNA-dependent fluorescent signal captured in the fluorescence image. Each sample includes two bright-field images acquired at different focal planes, one fluorescence image highlighting the nuclei, and one instance segmentation mask generated using Revvity's Acapella software based on the fluorescence channel images (see Figure 4).

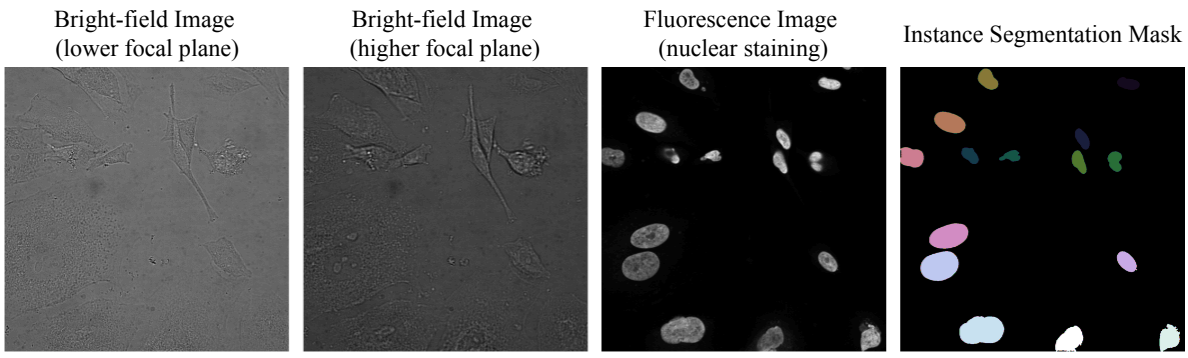

Figure 4. Example of a single sample from the dataset, comprising four images: two bright-field images acquired at different focal planes (left), one fluorescence image of nuclear staining (second from the right), and one instance segmentation mask (right), where each colour represents an individual nucleus.

3.2 Data Preprocessing

Effective training of deep learning models for microscopy image analysis requires high-quality, standardised input data. In this study, we implemented a comprehensive preprocessing pipeline designed to enhance image quality, isolate biologically meaningful regions, and structure the data to match the input requirements of our models. The pipeline includes manual removal of low-quality images, pixel intensity normalisation, dataset-level standardisation, and cropping of individual nuclei into patches for instance-level classification. Additionally, a feature extraction step was carried out by Revvity using their image analysis software, enabling us to perform feature classification based on quantified features of nuclear morphology and intensity.

3.2.1 Data Cleaning

During the initial exploration and visual inspection of the dataset, we identified several low-quality images. These images were affected by artefacts, blurriness, or partial focus — often

capturing the edge of the well instead of the cell population. Figure 5 illustrates representative examples of the excluded images.

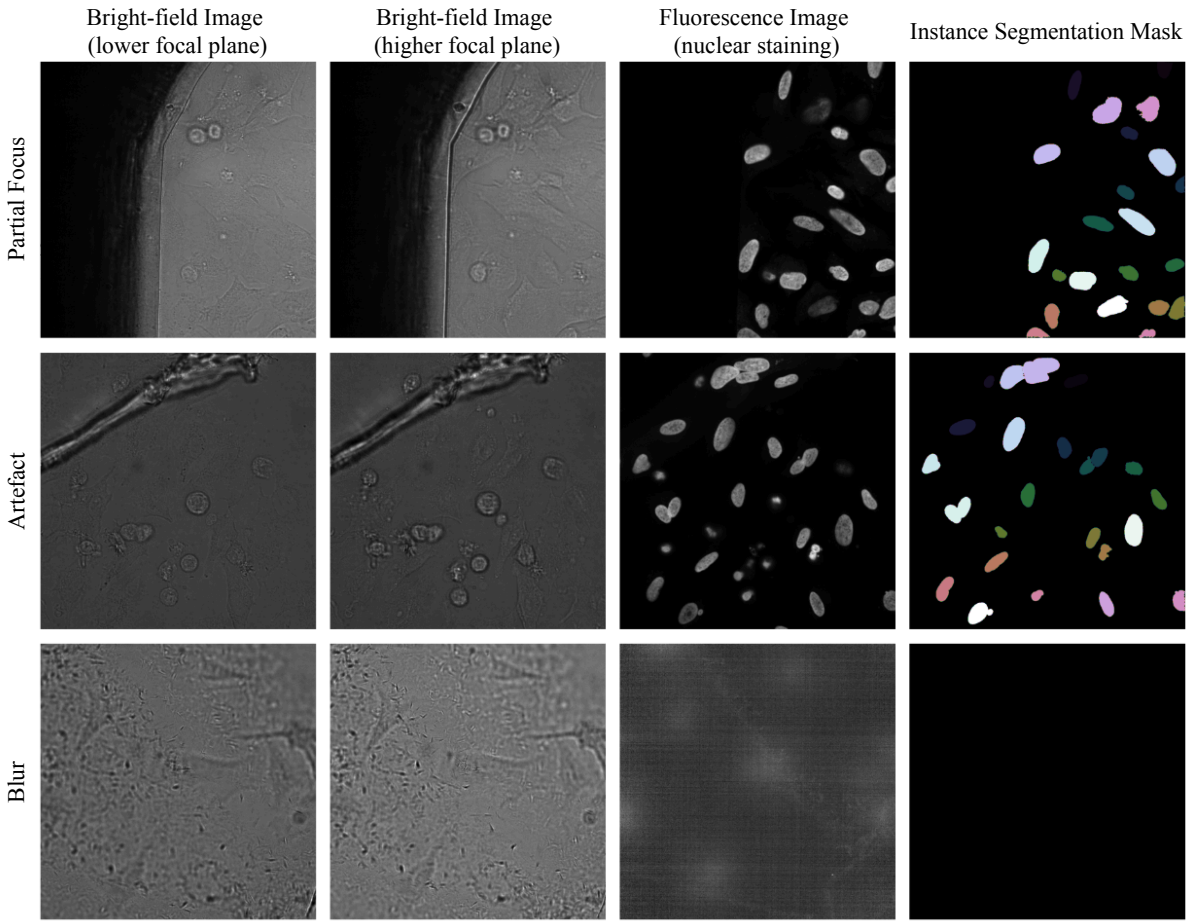

Figure 5. Examples of excluded low-quality images: partially focused images, capturing the edge of the well rather than the central cell population (top row); images with artefacts, such as non-biological structures (middle row); blurry images with loss of clarity (bottom row).

Defective images may contain noise or artefacts that do not accurately represent the true biological features of the cells. This can result in incorrect measurements and cell cycle phase misclassification, ultimately compromising the validity of model predictions. For example, if noise is interpreted as a biological signal, it can lead to false positives or false negatives in model predictions. This problem is well documented in the literature. Ali et al. [40] demonstrated that artefacts in microscopy images degrade both segmentation and classification accuracy. Their method, ArtSeg, which uses weak supervision to segment artefact regions, showed that explicitly excluding these regions improves downstream analysis. They reported higher pixel-wise and object-wise F1 scores when artefacts were filtered out, highlighting the importance of rigorous image quality control.

To address this issue, we performed a careful manual review of the entire dataset and removed 275 low-quality images. The final dataset used for this study consists of 1 693 images. By removing these low-quality images, we ensure that the dataset accurately reflects the actual cellular characteristics, thereby enhancing the reliability and validity of our analysis.

3.2.2 Data Normalisation and Standardisation

Standardisation and normalisation of input data are essential for achieving stable training, faster convergence, and improved generalisation in deep learning models [41, 42]. In microscopy image analysis, raw pixel intensity values can vary significantly due to differences in imaging conditions, sensor sensitivity, and background illumination [43]. Normalisation ensures a consistent dynamic range across samples. Standardisation improves learning by centring the data distribution and scaling input values to have zero mean and unit variance.

All microscopy images in the dataset were stored in 16-bit TIFF format. To normalise the intensity range, each pixel value in an image was divided by 65 535, the maximum representable value in 16-bit images. This operation scaled all pixel intensities to the interval $[0, 1]$. Following normalisation, the dataset was split into training, validation, and test subsets using a 60%, 20%, and 20% split, respectively. Standardisation was then applied by computing the mean and standard deviation of pixel intensities from the training set only. Computing these statistics exclusively on the training data prevents information leakage into validation and test sets and ensures that the model's input scaling reflects only what it has seen during training, thereby preserving the integrity of downstream evaluation. These statistics were subsequently used to standardise all three subsets by subtracting the mean and dividing by the standard deviation.

This preprocessing strategy combines global normalisation with dataset-level standardisation. This strategy was selected over alternatives such as per-image min-max normalisation. Although per-image min-max scaling is commonly used to reduce variation across images, it is not suitable in this context. In cell cycle phase classification, pixel intensity levels often encode biologically meaningful information that reflects the structural or functional state of cells. These intensity levels are also critical in generating ground-truth phase labels. Applying min-max scaling independently to each image would eliminate inter-sample intensity differences by rescaling every image to the same range. This would discard valuable discriminative features that are important for phase classification. In contrast, our approach preserves biologically relevant intensity variation while ensuring consistent input scaling across the dataset. Furthermore, computing standardisation parameters solely from the training set avoids information leakage, thereby supporting fair model evaluation and improved generalisation.

3.2.3 Data Cropping into Patches

An additional preprocessing step was implemented for classification strategies that required input in the form of individual nucleus patches instead of full-size images. To facilitate this, instance segmentation masks were used to localise each nucleus. For each segmented instance, a bounding box was computed by identifying the minimum and maximum pixel coordinates that corresponded to the mask region. These bounding boxes accurately enclosed the spatial extent of each nucleus, allowing for precise localisation within the image. The extracted bounding boxes were then utilised to crop localised regions around each nucleus from the corresponding fluorescence and bright-field images.

To provide high-quality and complete representations, nuclei located at the edges of the image were excluded from cropping. These instances often result in partially cropped nuclei, where important morphological features may be missing or truncated. Such incomplete representations

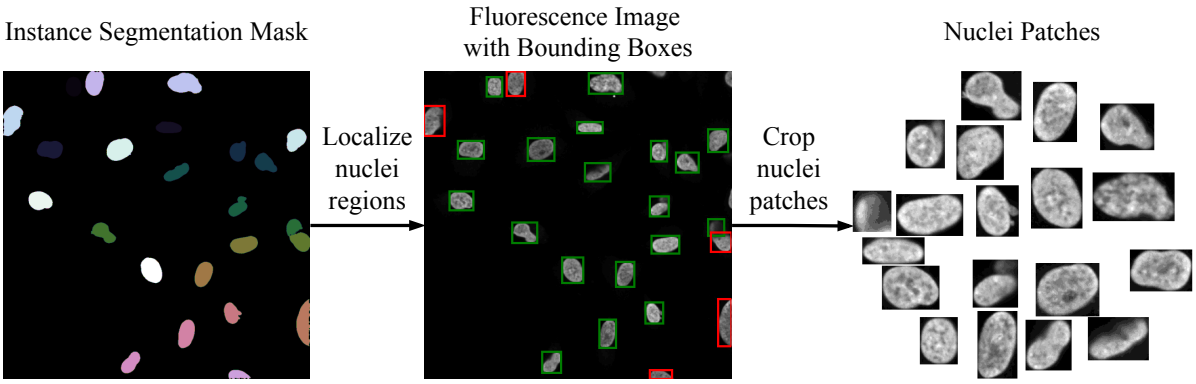

Figure 6. Illustration of the nuclei patch extraction process. Bounding boxes are derived from the instance segmentation mask (left), with green boxes indicating valid nuclei selected for cropping and red boxes marking excluded border instances. These bounding boxes are used to localise corresponding regions in the fluorescence image (middle), from which individual nuclei patches are extracted (right). The same strategy was applied to the bright-field images.

can introduce ambiguity, potentially misleading the classifier by providing insufficient or distorted visual information. For example, a cropped nucleus missing part of its boundary might resemble an earlier or later cell cycle phase, thereby increasing the risk of misclassification. Figure 6 illustrates how instance segmentation masks are applied to extract individual crops from fluorescence images and t bounding boxes selected for cropping are indicated in green, while instances at the borders that were excluded are marked in red.

3.2.4 Feature Extraction

In addition to image-based preprocessing, we implemented a feature extraction step to support the feature-based classification approach. This process was carried out by Revvity using their proprietary software, which automatically computed a diverse set of morphological and intensity-based descriptors for each nucleus. Features were extracted from both fluorescence and bright-field microscopy images, capturing various aspects of nuclear morphology, including area, size, shape, and integrated intensities. The extraction pipeline generated two feature tables depending on the imaging modality. For fluorescence images, 775 features were computed per nucleus. For bright field images captured at two different focal planes, a total of 1 448 features were derived by aggregating statistics across both planes. The output was stored in CSV format, where each row corresponds to a single nucleus and contains a comprehensive feature vector along with its image association. Figure 7 illustrates the feature table generated by Revvity's software. In the feature-based classification approach, this structured tabular representation was used as input for traditional machine learning models.

3.3 Data Annotation

The dataset did not contain ground truth labels for cell cycle phases, which are essential for supervised learning tasks such as classification. Since this study aims to classify the cell cycle phase of individual nuclei, it was necessary to generate these labels as part of the data annotation process. Manual annotation of microscopy datasets at the single-cell level is impractical due to the large volume of data. A single microscopy image can contain tens to hundreds of nuclei, and

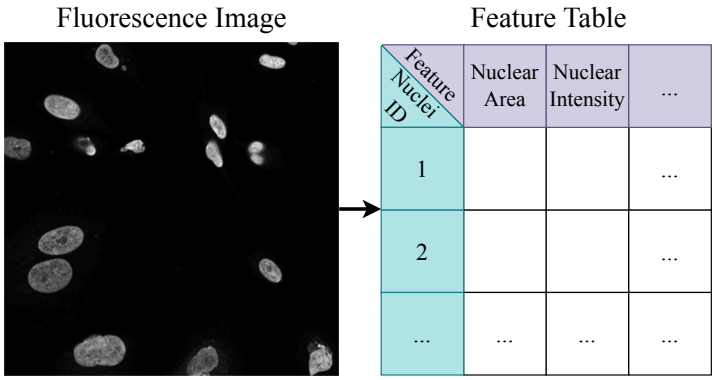

Figure 7. Feature table generation using Revvity's software. For each image, the software produces a table where each row corresponds to a nucleus, characterized by morphological and intensity-based metrics extracted from fluorescence or brightfield microscopy images.

large-scale experiments often consist of thousands of such images. Annotating each nucleus manually would require substantial expert time and effort, result in high costs, and introduce potential subjectivity and inconsistency. To address the absence of cell cycle annotations, we adopted a semi-automated labelling strategy based on DNA content estimation from fluorescence intensity.

3.3.1 Generating Labels

We followed the methodology outlined in the data analysis section of Chan et al. [13], as described earlier in Section 2.3 and illustrated in Figure 3. This method uses the fluorescence intensity of DNA-binding stains as a proxy for DNA content, which correlates with cell cycle progression. By constructing DNA content histograms and identifying distinct peaks corresponding to DNA content levels characteristic of each cell cycle phase, nuclei were automatically assigned to one of the major phases (G1, S, or G2/M) using predefined intensity thresholds. The main steps of our labelling procedure are summarised below.

1. Instance segmentation masks were used to localise individual nuclei in each image.
2. For each nucleus, the integrated pixel intensity was calculated from the corresponding fluorescence image. These intensity values were then $\log_2$-transformed to reduce skewness, stabilise variance, and facilitate robust peak detection. This transformation is commonly used in DNA content analysis, where measured intensities reflect approximately doubling DNA content between cell cycle phases.
3. DNA content histograms were generated across all cells in each well (48 wells in total). These histograms were used to identify the intensity peaks corresponding to the G0/G1 and G2+M phases. The detected peak positions were averaged across all wells to obtain robust reference values.
4. All fluorescence intensity values were then normalised using these averaged peak centres, such that the G0/G1 peak was centred at 1 and the G2+M peak at 2.
5. Cells were classified into five phases based on normalised intensity using fixed thresholds:

- **sub-G1:** < 0.75
- **G0/G1:** 0.75–1.25
- **S:** 1.25–1.75
- **G2+M:** 1.75–2.50
- **>4N:** > 2.50

Figure 8 illustrates the complete data annotation pipeline used in this study. This semi-automated approach enables reliable and scalable annotation while maintaining biological interpretability.

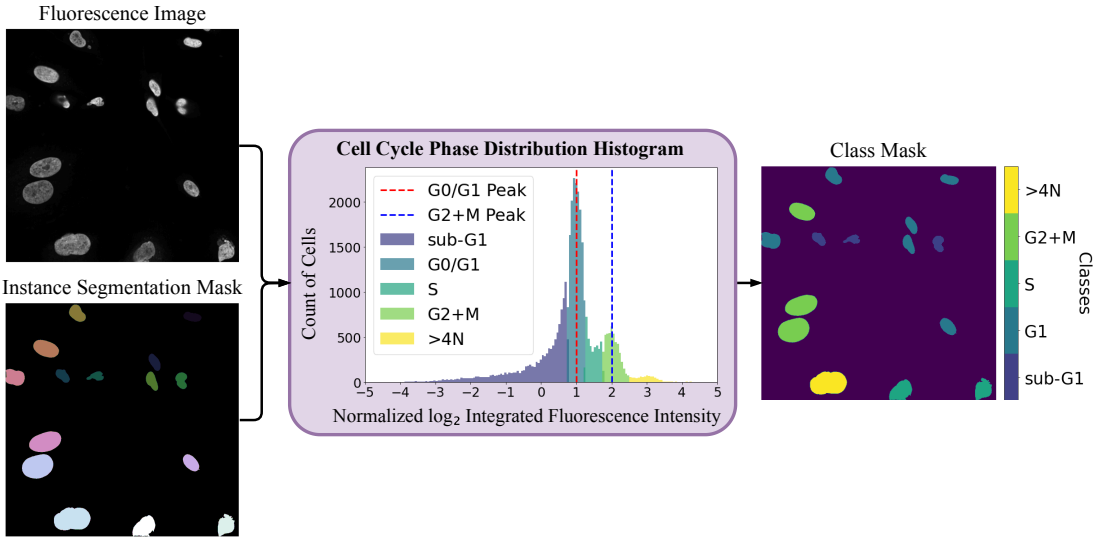

Figure 8. Overview of the data annotation pipeline. The fluorescence image and instance segmentation mask are used to extract integrated fluorescence intensities for each nucleus. These intensities are $\log_2$-transformed, normalised, and classified into cell cycle phases based on DNA content thresholds. As a result, a class mask is generated, where each nucleus is assigned to one of the five cell cycle phases.

3.3.2 Generating Confident Labels

While the histogram-based approach with fixed thresholds provides initial labels, nuclei with intensity values near class boundaries may be ambiguous and prone to misclassification. To improve label reliability, we applied an additional filtering step. Within each cell cycle class, only nuclei with integrated fluorescence intensities between the 12.5th and 87.5th percentiles were retained. This percentile-based filtering excludes borderline cases and yields a more robust set of confidently labelled instances. Figure 9 illustrates the filtering process, showing the initial class mask, the filtered class distribution, and the resulting confident class mask.

3.4 Classification Strategies

Accurate and automated classification of cell cycle phases is essential to overcome the limitations of traditional experimental approaches, which are expensive, time-consuming and labour-intensive. As discussed earlier, conventional methods typically rely on molecular markers that are effective. However, these markers are not always scalable or standardised across studies [17, 18]. To address these challenges, we explored computational strategies capable of performing cell cycle phase classification directly from microscopy data. Specifically, we investigated

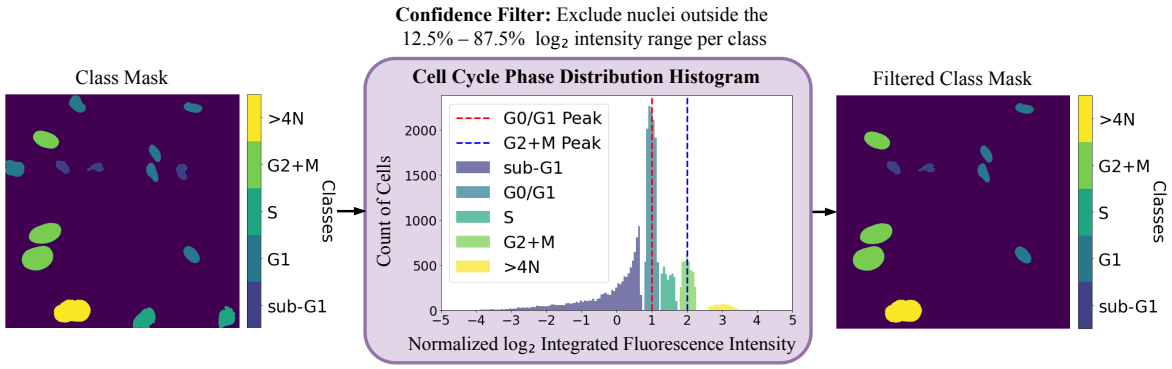

Figure 9. Workflow for generating confident labels. Initial class mask with nuclei assigned to five cell cycle phases (left); filtered class distribution histogram after excluding nuclei outside the 12.5%–87.5% $\log_2$ fluorescence intensity range (middle); filtered class mask retaining only confidently classified nuclei (right).

feature-based classification using software-extracted cellular features, and deep learning-based classification, which leverages raw image data to automatically learn discriminative features. These strategies enables evaluation of both traditional machine learning methods and modern end-to-end learning frameworks within the context of microscopy-based cell cycle analysis.

3.4.1 Feature-Based Classification

The feature-based classification approach utilizes tabular data generated from the feature extraction process detailed in Section 3.2.4. These features represent the morphological and intensity-based characteristics of individual nuclei obtained through fluorescence and bright-field microscopy images. Once extracted, these cellular features serve as inputs for machine learning models to train models that predict the cell cycle phase of each nucleus. Ground truth labels were generated using the annotation pipeline outlined in Section 3.3. By combining the cellular features with their corresponding cell cycle labels, we enabled supervised learning. An overview of the feature-based classification pipeline is presented in Figure 10.

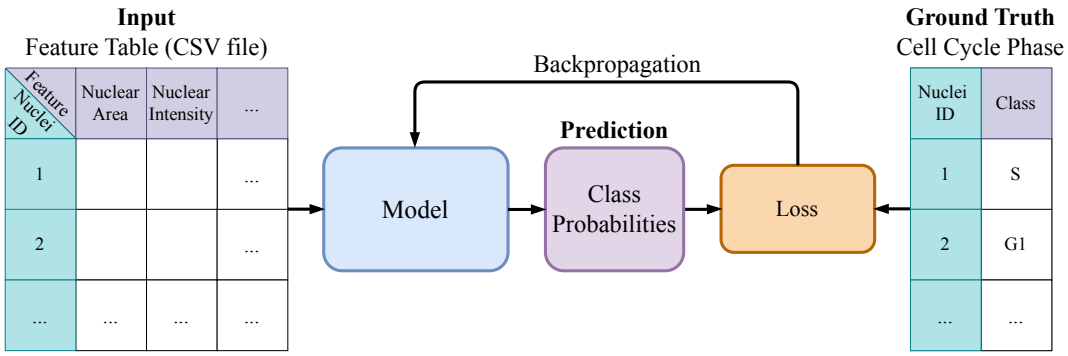

Figure 10. Overview of the feature-based classification approach. Extracted features and ground truth labels are used to train a machine learning model for cell cycle phase prediction.

We employed two ensemble learning models, Random Forest and Extreme Gradient Boosting (XGBoost), to perform classification. Both models are suitable for structured data, capable of

modelling non-linear relationships, and robust against moderate noise. Their interpretability and ability to compute feature importance make them valuable for understanding biological systems.

In Random Forest model, entropy was used as the split criterion during tree construction. Entropy quantifies uncertainty in class distribution and helps identify the most informative feature splits, guiding trees toward purer subsets. The entropy H at a node with C classes is defined as:

$$H = -\sum_{c=1}^{C} p_c \log_2(p_c) \tag{1}$$

where p_c is the proportion of samples of class c.

In XGBoost model, multi-class logarithmic loss was used as the evaluation function. This loss is also known as categorical cross-entropy, penalises incorrect predictions based on the confidence of the model's output probabilities. The multi-class logarithmic loss function is defined as

$$\mathcal{L}_{\text{mlog}} = -\frac{1}{N}\sum_{i=1}^{N}\sum_{c=1}^{C} y_{i,c} \log(\hat{y}_{i,c}) \tag{2}$$

where N is the number of samples, C is the number of classes, $y_{i,c}$ is the true label indicator, and $\hat{y}_{i,c}$ is the predicted probability.

Feature-based classification provides a straightforward and interpretable framework for cell cycle analysis. However, its reliance on predefined features may limit its ability to capture the full complexity of nuclear morphology associated with different cell cycle stages, which deep learning models can learn from raw image data.

3.4.2 Classification via Segmentation

In this study, we implemented a segmentation-based approach to classify cell cycle phases by formulating the task as multi-class semantic segmentation. This method was initially introduced by Chan et al. [13]. A deep neural network is trained to predict pixel-wise class labels using microscopy images as input and class masks as supervision. An overview of the segmentation-based classification pipeline is shown in Figure 11, which presents the workflow using a fluorescence image as an example. The same approach was also applied to bright-field images in a parallel evaluation.

To optimise model performance, we employed a composite loss function that combines Cross-Entropy Loss and Dice Loss. Cross-Entropy Loss penalises incorrect pixel-wise class predictions and is widely used in multi-class classification problems. The Cross-Entropy Loss $\mathcal{L}_{\text{CE}}$ is defined as:

$$\mathcal{L}_{\text{CE}} = -\sum_{c=1}^{C} y_c \log(\hat{y}_c) \tag{3}$$

where C is the number of classes, y_c is the one-hot encoded ground truth for class c, and $\hat{y}_c$ is the predicted probability for class c.

Dice Loss measures the spatial overlap between the predicted and ground-truth segmentations, making it particularly effective in the presence of class imbalance. The Dice Loss $\mathcal{L}_{\text{Dice}}$ is given

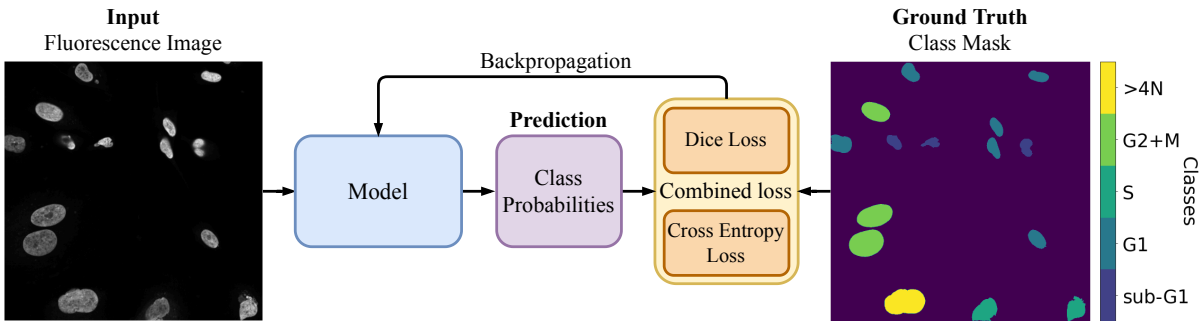

Figure 11. Classification via segmentation. A deep neural network receives a microscopy image (fluorescence or bright-field) as input and produces a pixel-wise segmentation map, assigning each pixel to one of five biological classes: sub-G1, G0/G1, S, G2+M, or >4N. The model is trained using a combination of cross-entropy and Dice loss.

by:

$$\mathcal{L}_{\text{Dice}} = 1 - \frac{2\sum_i y_i \hat{y}_i}{\sum_i y_i + \sum_i \hat{y}_i} \tag{4}$$

where y_i and $\hat{y}_i$ represent the ground truth and predicted segmentation values at pixel i, respectively.

The total loss is computed as a weighted sum::

$$\mathcal{L}\text{Total} = \alpha \cdot \mathcal{L}\text{CE} + \beta \cdot \mathcal{L}_{\text{Dice}} \tag{5}$$

with $\alpha = 0.5$ and $\beta = 0.5$ to balance pixel-level classification accuracy and region-level overlap.

3.4.3 Mask-Guided Classification via Segmentation

To improve consistency in cell cycle phase predictions, we extended the baseline semantic segmentation approach by incorporating instance-level information using instance segmentation masks. The aim of this approach is to encourage the model to produce consistent predictions for entire nuclei, rather than relying purely on pixel-wise accuracy. This is particularly important in microscopy settings, where partial misclassification within a single nucleus may lead to biologically misleading interpretations.

Similar to the baseline segmentation approach, the model receives microscopy image as input and produces a pixel-wise segmentation map indicating the predicted cell cycle phase for each pixel. In this extended framework, we integrate the instance segmentation mask directly into the model's input by using it as an additional channel alongside the original image. This channel input explicitly provides object-level boundaries, helping the network differentiate between adjacent nuclei and maintain internal consistency within each instance. During training, ground truth class masks are used for supervision, and the loss function is adapted to operate at the instance level. Specifically, predictions are aggregated across all pixels within each segmented nucleus, and the averaged prediction is compared against the majority ground truth class for that instance using instance-level cross-entropy loss. This design encourages the model to produce consistent predictions for each nucleus as a whole. An overview of the mask-guided classification via segmentation pipeline is shown in Figure 12. Although a fluorescence image is shown in the figure, the same method was also applied to bright-field images during evaluation.

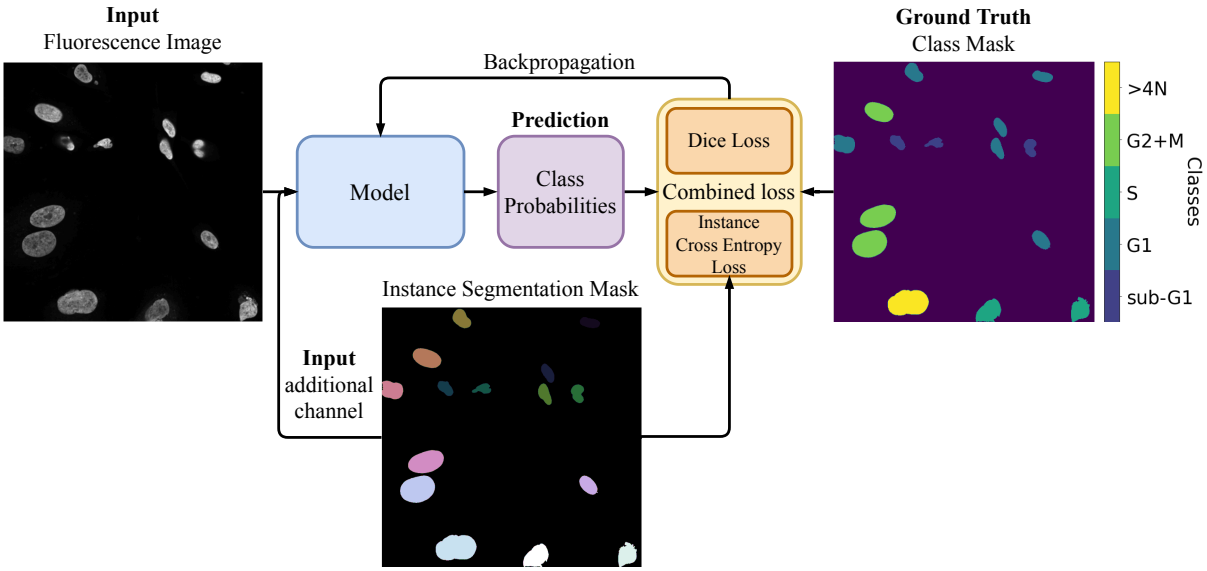

Figure 12. Overview of the mask-guided classification via segmentation approach. A deep neural network receives a three-channel input comprising a microscopy image (fluorescence or bright-field) and the corresponding instance segmentation mask. The model outputs a pixel-wise cell cycle phase map. During training, instance-level cross-entropy loss is used to enforce prediction consistency across each nucleus.

In this mask-guided approach, the standard Cross-Entropy Loss is modified to operate at the instance level. Specifically, for each nucleus identified in the instance segmentation mask, we isolate all corresponding pixels and compute the majority ground-truth class label across those pixels. The predicted class scores for that nucleus are then averaged across all pixels within the instance to produce a single prediction vector. This aggregated prediction is compared to the instance-level ground truth using standard Cross-Entropy.

The instance-aware Cross-Entropy Loss for a single instance k is defined as:

$$\mathcal{L}_{\text{CE}}^{(k)} = -\sum_{c=1}^{C} y_c^{(k)} \log(\hat{y}_c^{(k)}) \tag{6}$$

where C is the number of classes, $y_c^{(k)}$ is a one-hot encoded vector indicating the majority ground-truth label for instance k, and $\hat{y}_c^{(k)}$ is the softmax-normalised average prediction across all pixels in instance k.

Let $\mathcal{K}$ denote the set of all valid nucleus instances in the batch. The total instance-level Cross-Entropy Loss is computed by averaging over all instances:

$$\mathcal{L}_{\text{CE-Instance}} = \frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} \mathcal{L}_{\text{CE}}^{(k)} \tag{7}$$

The final training objective remains a weighted combination of Dice Loss and instance-aware Cross-Entropy Loss:

$$\mathcal{L}_{\text{Total}} = \alpha \cdot \mathcal{L}_{\text{CE-Instance}} + \beta \cdot \mathcal{L}_{\text{Dice}} \tag{8}$$

where α and β are weighting coefficients. As in the baseline formulation, we set both $\alpha = 0.5$ and $\beta = 0.5$ to equally balance pixel-level overlap and instance-level consistency. This formulation penalises the model more strongly when it misclassifies an entire nucleus and encourages coherence across all pixels within an object. It is particularly beneficial in dense cellular environments where pixel-wise predictions may be noisy or inconsistent.

3.4.4 Nuclei Patch Classification

In addition to segmentation-based strategies, we implemented a traditional image classification approach by approaching the task as a nucleus-level prediction problem. Individual nuclei were cropped into patches using the preprocessing pipeline described in Section 3.2.3. These nucleus-centered patches were then used as inputs to convolutional neural networks trained to predict the corresponding cell cycle phase.

Separate classification models were trained for fluorescence and bright-field image modalities. In both cases, the network received a cropped nucleus patch as input and produced a probability distribution over the five target classes: sub-G1, G0/G1, S, G2+M, and >4N. Ground-truth labels were assigned based on the annotation strategy described earlier. An overview of the patch-based classification pipeline is illustrated in Figure 13.

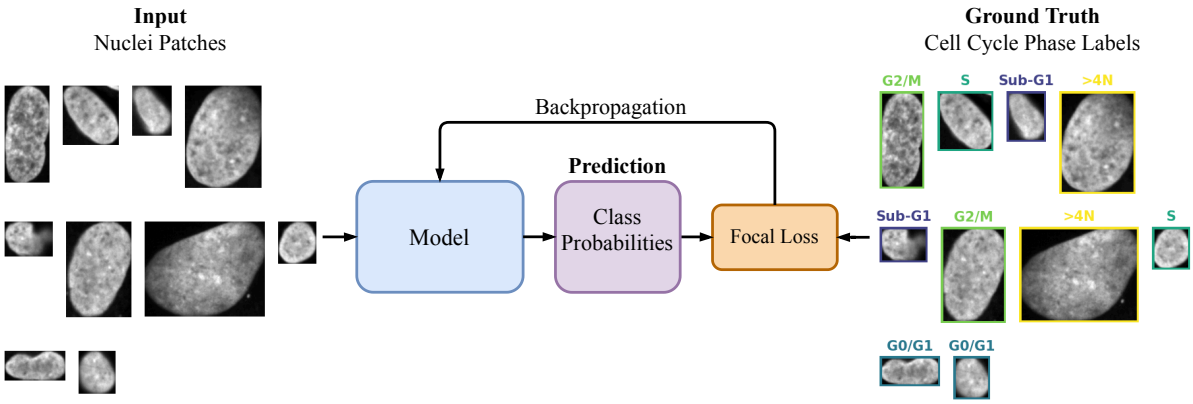

Figure 13. Nuclei patch classification. Cropped nucleus patches are extracted microcopy images (fluorescence or bright-field) and fed into a classification model. The model outputs a probability distribution over five cell cycle phases: sub-G1, G0/G1, S, G2+M, and >4N. Focal Loss is used during training to address class imbalance and focus learning on difficult cases.

To train the models, we used Focal Loss function. Although Cross-Entropy Loss is widely used for multi-class classification, it proved insufficient in our case due to the class imbalance present in the dataset—particularly the under-representation of the S and >4N phases. Focal Loss mitigates this issue by dynamically scaling the contribution of each training sample based on the confidence of its prediction, placing greater emphasis on hard-to-classify and minority-class examples [44]. This makes it especially effective for biological datasets, where the distribution of cell states may be naturally skewed. The Focal Loss is defined as:

$$\mathcal{L}_{\text{Focal}} = -\alpha_t (1 - \hat{y}_t)^{\gamma} \log(\hat{y}_t) \tag{9}$$

where $\hat{y}_t$ is the predicted probability for the true class t, α_t is a class-specific weighting factor to address imbalance, γ is the focusing parameter that adjusts the contribution of well-classified vs.

misclassified examples. In our implementation, we adapted this formulation using PyTorch's built-in Cross-Entropy Loss with class weights, and applied Focal Loss as:

$$\mathcal{L}_{\text{Focal}} = (1 - p_t)^{\gamma} \cdot \mathcal{L}_{\text{CE}} \tag{10}$$

where $p_t = \exp(-\mathcal{L}_{\text{CE}})$ corresponds to the predicted probability for the correct class. We used $\gamma = 2.0$ and incorporated class weights to reflect the label distribution. This configuration led to more stable training and improved generalisation across under-represented classes.

3.4.5 Nuclei Patch Classification via Segmentation

This classification strategy evaluates whether segmentation-based models can benefit from localized input by applying them to cropped nuclei patches instead of full-size microscopy images. The goal is to determine whether limiting the spatial context to individual nuclei can improve classification performance by reducing interference from neighboring cells and background noise.

As in the full-image segmentation approach, the model is trained to assign a cell cycle phase to each pixel within the input image. Each input consists of a cropped nucleus patch extracted from either fluorescence or bright-field microscopy data. Ground truth labels for supervision are provided by the corresponding region in the class mask.

To maintain consistency with the full-image segmentation strategy, we used the same loss formulation, combining pixel-wise Cross-Entropy Loss and Dice Loss. This design allows the model to balance pixel-level classification accuracy with spatial coherence in the predicted segmentation. Localising the task to the patch level enables the network to focus more precisely on the morphology and internal texture of individual nuclei, which may lead to improved discrimination between subtle morphological changes across cell cycle phases.

An overview of this patch-based segmentation pipeline is illustrated in Figure 14. While the example shows a fluorescence nucleus patch, the same approach was also applied to bright-field images.

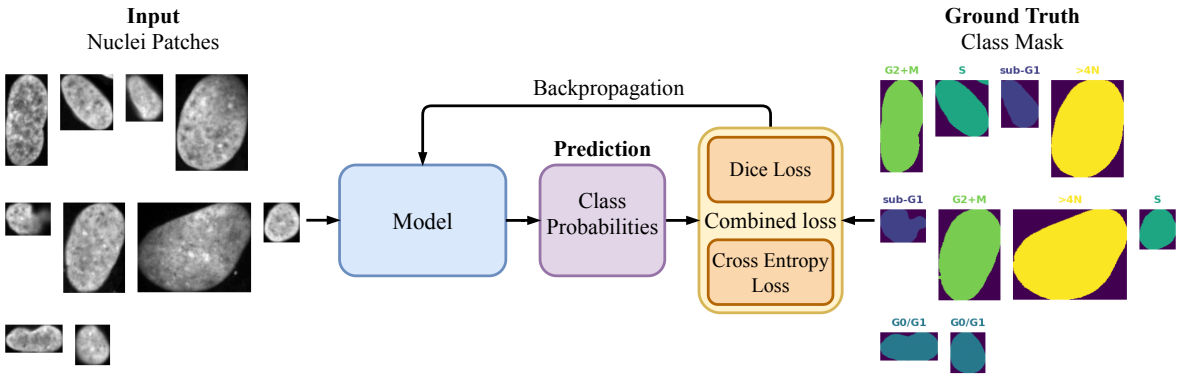

Figure 14. Nuclei patch classification via segmentation. A cropped nucleus patch from microscopy images (fluorescence or bright-field) is fed into a semantic segmentation model, which assigns each pixel to a specific cell cycle phase. The network is trained using a combined loss function (Cross-Entropy and Dice Loss) with class masks as ground truth.

3.5 Model Selection

To support the five classification strategies explored in this study — feature-based classification, classification via segmentation, mask-guided classification via segmentation, nuclei patch classification, and nuclei patch classification via segmentation — we employed a diverse set of models tailored to the specific requirements of each task.

For the feature-based classification strategy, we used two ensemble learning models: Random Forest and XGBoost. These models are particularly suitable for structured, tabular data derived from cell morphology measurements. Random Forest offers strong performance with minimal tuning and is robust to overfitting [45]. XGBoost enhances this by introducing gradient boosting and regularisation, making it highly effective in datasets with noisy or imbalanced classes [46].

All segmentation-based strategies such as classification via segmentation, mask-guided classification via segmentation, and nuclei patch classification via segmentation, were implemented using the U-Net architecture [47]. Designed specifically for biomedical segmentation tasks, U-Net features an encoder-decoder structure with skip connections that facilitate precise, pixel-level predictions. For these segmentation models, we employed ResNet [48] and EfficientNet [49] as encoder backbones. Both were pre-trained on the ImageNet dataset [50] to leverage transfer learning and accelerate training. ResNet uses residual blocks to support deep learning without degradation, while EfficientNet balances network depth, width, and resolution to achieve high accuracy with lower parameter counts.

To ensure a broad comparison of model performance, we selected three different sizes from each backbone architecture, covering a range from lightweight to more complex designs. Specifically, we used ResNet-18, ResNet-50, and ResNet-152 from the ResNet family, as well as EfficientNet-B0, EfficientNet-B4, and EfficientNet-B7 from the EfficientNet family. This setup allowed us evaluation of small, medium, and large models across the implemented classification strategies.

For the patch-based image classification strategy, we used the same set of ResNet and EfficientNet architectures as standalone classifiers. These convolutional neural networks are highly effective for visual tasks and allow for efficient learning of local morphological features in cropped nucleus patches. Table 1 summarises the models and backbone configurations used in each classification strategy.

Table 1. Summary of models used for each classification strategy. All deep learning backbones were pre-trained on ImageNet weights.

Classification Strategy	**Model and Architecture**
Feature-Based Classification	XGBoost, Random Forest
Classification via Segmentation	U-Net with ResNet-18, 50, 152 and EfficientNet-B0, B4, B7
Mask-Guided Classification via Segmentation	
Nuclei Patch Classification via Segmentation	
Nuclei Patch Classification	ResNet-18, 50, 152 and EfficientNet-B0, B4, B7

3.6 Training Pipeline

Our training pipeline was designed to support both traditional machine learning models and deep learning-based classification strategies used in this study. Traditional models such as Random Forest and XGBoost were implemented using the scikit-learn and XGBoost libraries, respectively. Deep learning models were developed using the PyTorch framework, with additional support from utility libraries, including Segmentation Models PyTorch for U-Net-based segmentation tasks. Backbone architectures such as ResNet and EfficientNet were employed in two contexts: as encoders within U-Net for segmentation tasks, and as standalone models in image classification strategies. This modular and flexible pipeline enabled consistent experimentation across all classification strategies while ensuring reproducibility and scalability.

3.6.1 Data Preparation and Augmentation

The same data augmentation protocol was applied across all deep learning experiments to enhance generalisation and mitigate overfitting. Augmentations were implemented using the Albumentations library [51], chosen for its efficiency and compatibility with PyTorch-based workflows. The applied transformations included horizontal flips, vertical flips, and 90-degree rotations. These augmentations were selected to increase the model's robustness to orientation and positional variance — an essential requirement in microscopy image analysis, where nuclei can appear in arbitrary orientations. Horizontal and vertical flips help the model learn reflection-invariant representations, while random rotations simulate multiple viewing angles and promote rotational invariance. A deterministic random seed was used to ensure reproducibility of training results. Collectively, this augmentation strategy helps the model maintain high performance in the presence of varied spatial arrangements and potential domain shifts.

Besides standard augmentations, all classification strategies required resizing of input images to comply with the fixed-size requirements of deep learning models. Images larger than the target resolution were resized down using bilinear interpolation for microscopy images and nearest-neighbour interpolation for segmentation masks, ensuring spatial integrity while maintaining input consistency across modalities.

An additional resizing step was necessary for classification strategies that operate on cropped nuclei patches. Since nuclei vary in size and morphology, some patches were smaller than the target input size. In these cases, zero-padding was used to preserve morphological fidelity instead of applying interpolation. This padding approach ensured that small nuclei were not artificially enlarged or blurred, while still conforming to the input shape requirements of the network.

This combination of augmentation, resizing, and padding formed a unified preprocessing strategy that promoted consistent, high-quality input data for robust and generalisable model training.

3.6.2 Training Details

All deep learning models were trained for 50 epochs, which was empirically determined to be sufficient for convergence while balancing computational efficiency. The batch size was selected as the largest value that could fit into the GPU memory for each task. Learning rates were tuned for each model class using an initial grid search and validation accuracy. For model optimisation, we used the Adam optimiser [52], which offers adaptive learning rate updates and performs well

in computer vision tasks. To improve convergence and adapt the learning rate during training, we employed a ReduceLROnPlateau scheduler, which monitors validation loss and reduces the learning rate by a factor of 0.1 if no improvement is observed for five consecutive epochs. This setup contributed to more stable and efficient training dynamics. Table 2 provides a summary of training configurations.

Table 2. Training configurations for each classification strategy. Asterisk (*) denotes fluorescence images, and double asterisk (**) denotes bright-field images (both focal planes).

Classification Strategy	Input Size	Batch Size	Loss Function	Learning Rate
Feature-Based Classification	775* or 1448** features	–	Entropy or Multiclass Log	–
Classification via Segmentation	512×512	8	Dice + Cross-Entropy	1e-4
Mask-Guided Classification via Segmentation	512×512	8	Dice + Instance Cross-Entropy	1e-4
Nuclei Patch Classification via Segmentation	224×224	32	Dice + Cross-Entropy	1e-4
Nuclei Patch Classification	224×224	32	Focal Loss	5e-5

3.6.3 Evaluation Metrics and Testing

Balanced accuracy was used as the primary evaluation metric for all classification strategies. This metric is particularly effective in the presence of class imbalance, as it computes the average recall across all classes. By giving equal importance to each class, balanced accuracy ensures that rare classes contribute fairly to the overall performance, which is essential given the uneven distribution of cell cycle phases in our dataset.

To address the ambiguity in the ground truth labels, generated using the threshold-based method, we also report the top-3 balanced accuracy. This metric evaluates whether the correct class is among the three most probable predictions made by the model, providing a more flexible assessment in scenarios where class boundaries may be uncertain.

Model predictions were evaluated on a held-out test set. Instance-level metrics were computed using instance segmentation masks to ensure precise alignment between predictions and ground-truth annotations. All evaluations were performed on the subset of nuclei with confident labels, as defined by the filtering strategy described in Section 9.

3.6.4 Computational Resources

All deep learning models were trained on the High-Performance Computing (HPC) cluster at the University of Tartu, which provides access to NVIDIA Tesla V100 GPUs with 32 GB VRAM and NVIDIA Tesla A100 GPUs with 40 GB VRAM [53]. Training parameters were selected to maximise computational efficiency without compromising model convergence and performance.

Traditional machine learning models, including Random Forest and XGBoost, were trained on CPU resources, as their implementations are optimised for tabular data and do not require GPU acceleration. This setup ensured efficient training for both deep learning and machine learning components of the pipeline.

3.7 Writing Assistance

To support the scientific writing and technical presentation of this thesis, we used ChatGPT [54] and Grammarly [55]. ChatGPT assisted in structuring content, refining technical explanations, and generating LaTeX code for tables and figures. Grammarly was employed to enhance grammar, clarity, and readability throughout the text. Together, these tools contributed to the overall clarity and professionalism of the final manuscript.

4. Experiments and Results

This section summarizes the performance of the classification strategies presented in Section 3, evaluated on both fluorescence and bright-field microscopy images. The experiments aim to assess and compare the effectiveness of different models in accurately predicting cell cycle phases, with particular emphasis on bright-field images, which lack biomarker-specific signals and exhibit lower contrast. A comparative analysis of all strategies is conducted to identify the most accurate and generalisable pipeline for cell cycle phase classification in label-free imaging modality.

4.1 Feature-Based Classification

Table 3. Balanced accuracy of XGBoost and Random Forest models on fluorescence and bright-field microscopy images. Models were trained on extracted morphological and intensity features and evaluated for their ability to classify cell cycle phases. The best score is highlighted in **bold**, the second-best score is underlined.

Model	Fluorescence	Bright-field
XGBoost	**0.996**	**0.736**
Random Forest	0.875	0.689

We employed two ensemble learning algorithms for the feature-based classification strategy, namely Random Forest and XGBoost. Both models were trained on tabular data composed of morphological and intensity-derived features extracted from individual nuclei to predict their cell cycle phase. The Random Forest model was optimized through a grid search over key hyperparameters. The best-performing configuration included a maximum tree depth of 20, a minimum of five samples required to split a node, and 200 estimators. The entropy criterion was used to evaluate information gain during tree construction. In comparison, XGBoost was trained using its default hyperparameters, with multi-class logarithmic loss as the evaluation metric. This configuration showed strong performance, and no further tuning was conducted.

Table 3 presents the balanced accuracy achieved by both models across fluorescence and bright-field microscopy data. On fluorescence images, both models performed well. XGBoost reached a near-perfect accuracy of 0.996, while Random Forest achieved 0.875. On bright-field images, accuracy decreased for both models. Nevertheless, XGBoost maintained a higher accuracy of 0.736, demonstrating better adaptability to bright-field features than Random Forest, which achieved 0.689.

4.2 Classification via Segmentation

We employed the U-Net model to conduct multiclass pixel-wise semantic segmentation for classifying cell cycle phases. To systematically investigate the impact of encoder capacity on segmentation performance, we chose backbone architectures from two convolutional neural network families, ResNet and EfficientNet. We selected three representative variants from each family, providing a spectrum of models that range from lightweight to high-capacity. This approach enables a thorough evaluation of networks with differing depths and parameter counts across both fluorescence and bright-field imaging modalities.

Table 4. Balanced accuracy of U-Net models with different backbone architectures of EfficientNet and ResNet, evaluated on fluorescence and bright-field microscopy images. Each model was trained to perform semantic segmentation for cell cycle phase classification. The best score is highlighted in **bold**, the second-best score is underlined.

Backbone Model	Fluorescence	Bright-field
EfficientNet-B0	0.943	0.599
EfficientNet-B4	0.946	0.614
EfficientNet-B7	0.946	0.622
ResNet-18	0.954	0.614
ResNet-50	0.952	**0.652**
ResNet-152	**0.962**	0.625

Table 4 presents the balanced accuracy scores obtained by each model. On fluorescence images, all U-Net configurations achieved strong performance, with the ResNet-152 backbone reaching the highest accuracy of 0.962. In contrast, performance on bright-field data was lower for all models. The best-performing backbone, ResNet-50, achieved a balanced accuracy of 0.652, underscoring the challenge of applying pixel-level segmentation in low-contrast bright-field images. These results suggest that while U-Net models are highly effective in fluorescence images, their generalisability to bright-field microscopy remains limited.

4.3 Mask-Guided Classification via Segmentation

Table 5. Balanced accuracy of U-Net models with different backbone architectures of EfficientNet and ResNet, evaluated on fluorescence and bright-field microscopy images. Each model was trained using instance-level supervision to perform mask-guided classification via segmentation. The best score is highlighted in **bold**, the second-best score is underlined.

Backbone Model	Fluorescence	Bright-field
EfficientNet-B0	0.956	0.728
EfficientNet-B4	0.968	0.723
EfficientNet-B7	0.960	0.720
ResNet-18	0.979	0.730
ResNet-50	**0.981**	0.732
ResNet-152	0.979	**0.749**

Table 5 presents a summary of the balanced accuracy achieved through mask-guided classification via segmentation, utilizing instance segmentation masks for both input and loss supervision. The results demonstrate strong performance on fluorescence images across all models, with U-Net utilizing a ResNet-50 backbone achieving the highest score of 0.981, closely succeeded by the ResNet-152 model.

In contrast, the accuracy for bright-field images was lower than that for fluorescence images; however, it consistently surpassed the results obtained from baseline segmentation method in

Section 4.2. The U-Net model with a ResNet-152 backbone achieved a balanced accuracy of 0.749. Notably, the performance trend indicates that ResNet backbones generally outperformed their EfficientNet variants in this approach. Moreover, the integration of instance-level supervision enhanced the consistency of predictions across individual nuclei in both imaging modalities.

4.4 Nuclei Patch Classification

Table 6. Balanced accuracy of patch-based classification models applied to fluorescence and bright-field microscopy images. Cropped nucleus patches were used as input to convolutional neural networks with ResNet and EfficientNet backbones. The best score is highlighted in **bold**, and the second-best is underlined.

Model	Fluorescence	Bright-field
EfficientNet-B0	0.999	0.725
EfficientNet-B4	0.986	0.717
EfficientNet-B7	**1.000**	0.697
ResNet-18	0.998	0.768
ResNet-50	0.999	**0.770**
ResNet-152	**1.000**	0.756

Table 6 presents the balanced accuracy of patch-based classification models trained to predict cell cycle phases from cropped nucleus regions. On fluorescence images, all models achieved very high accuracies, with ResNet-152 and EfficientNet-B7 reaching perfect scores of 1.000.

Model performance remained lower for bright-field images compared to fluorescence images. However, ResNet-50 has achieved the highest balanced accuracy at 0.770, compared to the other classification approaches. Across all configurations, ResNet backbones outperformed their EfficientNet backbones on bright-field images.

This strategy proved highly effective across both modalities, particularly for bright-field images. Patch-based models could extract relevant structural information by focusing classification on localised nuclear regions, highlighting their robustness and practical utility for cell cycle phase classification.

4.5 Nuclei Patch Classification via Segmentation

In the experiment assessing various segmentation models on cropped nucleus patches, different architectures were utilised to evaluate their effectiveness across fluorescence and bright-field images. For fluorescence images, the models demonstrated strong performance overall, with EfficientNet-B0 standing out by achieving the highest accuracy of 0.974. The high performance of these models in this modality indicates their ability to extract meaningful features from the fluorescence images, which tend to be richer in contrast and detail.

Conversely, the performance on bright-field images was generally lower across all models, compared to fluorescence images. The ResNet-50 backbone stood out, achieving the best performance with an accuracy of 0.737. This indicates that while bright-field imaging presents

Table 7. Balanced accuracy of patch-based segmentation models trained on cropped nucleus regions from fluorescence and bright-field images. Models were trained to assign pixel-wise labels corresponding to cell cycle phases. The best score is highlighted in **bold**, and the second-best is underlined.

Model	Fluorescence	Bright-field
EfficientNet-B0	**0.974**	0.717
EfficientNet-B4	0.963	0.723
EfficientNet-B7	0.962	0.732
ResNet-18	0.970	0.729
ResNet-50	0.973	**0.737**
ResNet-152	0.952	0.725

challenges, ResNet-50's architecture still allows it to maintain a competitive edge. EfficientNet-B7 and ResNet-18 also showed promising results, although they did not match the accuracy of ResNet-50 in this case. The consistent performance of ResNet-50 across both fluorescence and bright-field modalities demonstrates its robustness and generalisation capabilities.

4.6 Comparison of Approaches

Table 8. Top-3 balanced accuracy scores for each classification strategy across fluorescence and bright-field images. The best score is highlighted in **bold**, and the second-best is underlined.

Classification Strategy	Fluorescence			Bright-field		
	Top-1	Top-2	Top-3	Top-1	Top-2	Top-3
Feature-Based Classification	0.996	0.999	**1.000**	0.736	0.944	0.996
Classification via Segmentation	0.962	0.979	0.988	0.652	0.833	0.907
Mask-Guided Classification via Segmentation	0.981	0.993	0.994	0.749	0.929	0.964
Nuclei Patch Classification	**1.000**	**1.000**	**1.000**	**0.770**	**0.965**	**0.997**
Nuclei Patch Classification via Segmentation	0.974	**1.000**	**1.000**	0.737	0.934	0.975

Table 8 compares the best-performing models from each classification strategy using top-1, top-2, and top-3 balanced accuracy on fluorescence and bright-field microscopy data. The top-3 accuracy illustrates how often the correct class is included among their top three predictions.

All classification approaches demonstrated strong classification performance on fluorescence images, with Nuclei Patch Classification achieving perfect scores across all metrics. These results demonstrate that each strategy can extract discriminative features for accurate phase classification from fluorescence images.

In the more challenging bright-field images, performance varied immensely among the methods. Nuclei Patch Classification achieved the highest top-1 balanced accuracy of 0.770, closely followed by Mask-Guided Classification via Segmentation at 0.749, and Nuclei Patch Classification via Segmentation at 0.737. These results indicate that strategies targeting localised

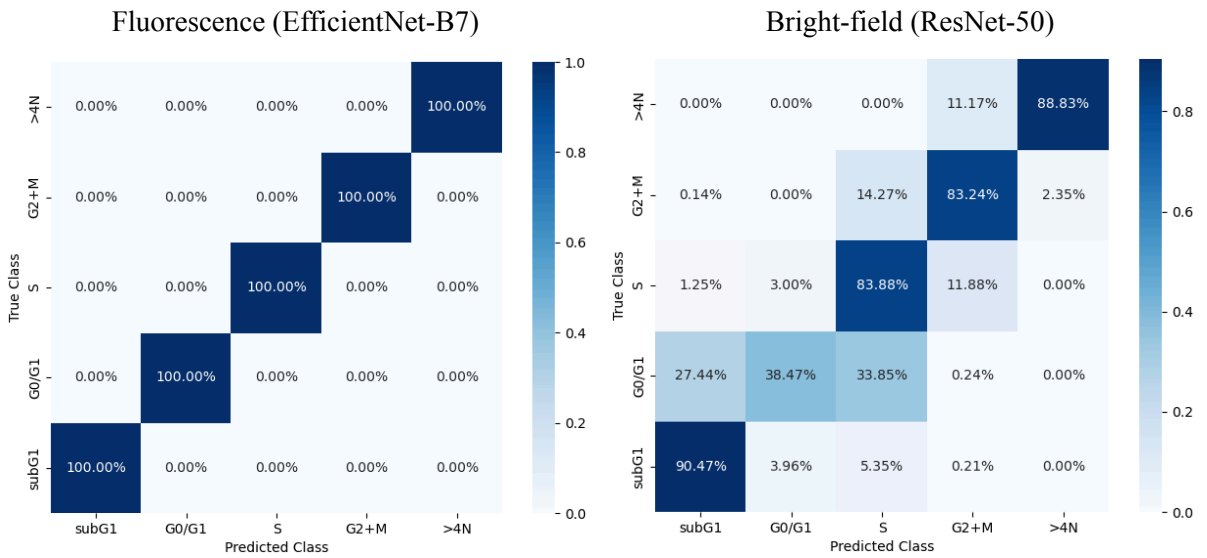

Figure 15. Confusion matrices of the best-performing Nuclei Patch Classification models for each imaging modality. EfficientNet-B7 achieved top performance on fluorescence images, while ResNet-50 performed best on bright-field data.

nuclear regions are particularly effective for cell cycle phase prediction from bright-field images, likely due to their capacity to isolate and emphasise relevant morphological details.

Figure 15 illustrates the confusion matrices for the top-performing Nuclei Patch Classification models applied to fluorescence and bright-field images, providing insight into model performance. The matrices indicate that most misclassifications occur between adjacent cell cycle phases, particularly among subG1, G1/G0, and S phases in bright-field images. This can be attributed to our ground truth labels being generated through a semi-automated pipeline using thresholding, which allows for potential misclassifications within these adjacent classes. Understanding these patterns is crucial for enhancing the accuracy of our annotations, thereby producing more reliable ground truth data. This will allow models to be trained on robust and reliable ground truth, resulting in better overall performance.

Notably, mask-guided classification via the segmentation approach consistently outperformed the baseline segmentation approach across all metrics. This underscores the advantage of enforcing instance-level consistency, which enhances robustness in low-contrast images by promoting coherent predictions across each nucleus.

5. Discussion

This study evaluated various strategies for classifying cell cycle phases from microscopy images, with a particular focus on bright-field imaging as a cost-effective and label-free alternative to fluorescence imaging. We compared traditional feature-based approaches with deep learning-based pipelines, analyzing segmentation and classification models across different input modalities and architectures.

The results demonstrate that deep learning strategies outperform traditional machine learning models when applied to raw image data, especially in the challenging bright-field modality. Among the tested approaches, the nuclei patch classification strategy utilizing ResNet-50 achieved the highest balanced accuracy on bright-field images, reaching 0.770. This was closely followed by the mask-guided classification through the segmentation U-Net model with ResNet-152, which achieved an accuracy of 0.749. These findings indicate that patch-level models can effectively learn local morphological cues, even in lower-contrast imaging conditions.

Fluorescence-based models consistently achieved near-perfect accuracy across all methodologies, validating the proposed architectures and loss formulations. However, as anticipated, bright-field images presented greater challenges due to the lack of biomarker-specific contrast. Despite this, several strategies demonstrated good performance on bright-field data, especially when guided by instance segmentation or patch-localized input. This suggests that nuclear morphology remains a sufficiently informative feature under label-free conditions when paired with robust models.

Interestingly, the best-performing models were not always the largest. While deeper architectures like ResNet-152 and EfficientNet-B7 often led to strong fluorescence results, mid-sized models such as ResNet-50 offered the best trade-off between performance and generalisability on bright-field data.

5.1 Limitations

One limitation of this thesis is its reliance on a single dataset, which defines the scope and experimental context of the study. While this approach allows for controlled comparisons and methodological consistency, it may restrict the broader applicability of the findings. To strengthen validation and gain deeper insights into the robustness and generalizability of the proposed models, it would be beneficial to extend the analysis to include a more diverse range of datasets, encompassing multiple cell lines, imaging modalities, and acquisition conditions.

Another limitation comes from the computational resources needed to train and evaluate complex models like U-Net. Microcopy images tend to be large, and when we have thousands of images in our datasets, it becomes challenging to meet the computational resource requirements. To address this, we have been resizing the images and using smaller batch sizes. However, training more advanced models with full-size images could significantly increase performance, but this may not be feasible in all research environments, potentially limiting the broader adoption of these advanced models.

5.2 Further Research

A promising direction for future research involves the integration three-dimensional (3D) spatial contextual information. While this thesis focused on two-dimensional (2D) microscopy images, many biological processes — such as cell cycle progression — occur in a continuous spatial and temporal context that 2D imaging alone cannot fully capture. In particular, 3D volumetric imaging offers a more comprehensive view of cellular morphology and organization across multiple focal planes, providing richer information for classification tasks. Extending the current framework to incorporate these modalities through 3D convolutional networks or spatiotemporal models could lead to more accurate and generalizable classification, especially in complex or heterogeneous cellular environments.

The mask-guided classification through segmentation strategy presented in this study provides a promising framework for enhancing image analysis. By enforcing instance-level consistency during training, this approach effectively reduces the noise often encountered in pixel-wise segmentation, particularly in low-contrast images. There is significant potential for future research to build on this concept by integrating attention mechanisms to prioritize biologically relevant regions better. Moreover, combining instance-aware training with temporal tracking could further improve the accuracy of cell cycle phase classification.

6. Conclusion

This thesis explored and compared multiple computational strategies for classifying cell cycle phases from microscopy images, focusing on bright-field imaging as a cost-effective and label-free alternative to fluorescence-based methods. Five classification pipelines were evaluated, ranging from traditional feature-based methods to advanced deep learning models incorporating segmentation, mask guidance, and patch-level inputs. Fluorescence images served as a performance benchmark, with all approaches achieving high accuracy due to the strong fluorescence signal provided by biomarker-specific contrasts. However, bright-field images presented greater challenges due to their lower contrast and lack of explicit labels. Among all strategies, nuclei patch classification using ResNet-50 emerged as the most effective on bright-field images, achieving a top-1 balanced accuracy of 0.770. These results highlight the effectiveness of localized and instance-aware models in bright-field images and offer a promising direction for scalable, automated cell cycle analysis in biomedical research.

7. Acknowledgements

I would like to express my gratitude to my supervisor, Dmytro Fishman, for his expert guidance, constant support, and insightful feedback throughout this journey. I also wish to thank the members of the Biomedical Computer Vision Lab for their collaborative spirit, helpful discussions, and encouragement.

References

[1] Cell Division - Mitosis and Meiosis | Ask A Biologist. https://askabiologist.asu.edu/cell-division.

[2] Ligasová A., Frydrych I., and Koberna K. Basic Methods of Cell Cycle Analysis. International Journal of Molecular Sciences 24.4 (2023), p. 3674. DOI: 10.3390/ijms24043674.

[3] Auld D. S., Coassin P. A., Coussens N. P., Hensley P., Klumpp-Thomas C., Michael S., Sittampalam G. S., Trask O. J., Wagner B. K., Weidner J. R., et al. Microplate selection and recommended practices in high-throughput screening and quantitative biology. Assay Guidance Manual (2020).

[4] Sonneck J., Zhao S., and Chen J. On the risk of manual annotations in 3D confocal microscopy image segmentation. Proceedings of the IEEE/CVF International Conference on Computer Vision. 2023, pp. 3894–3902.

[5] Hayatigolkhatmi K., Soriani C., Soda E., Ceccacci E., El Menna O., Peri S., Negrelli I., Bertolini G., Franchi G. M., Carbone R., et al. Automated workflow for the cell cycle analysis of (non-) adherent cells using a machine learning approach. eLife 13 (2024), RP94689. DOI: 10.7554/eLife.94689.3.

[6] He Y. R., He S., Kandel M. E., Lee Y. J., Hu C., Sobh N., Anastasio M. A., and Popescu G. Cell cycle stage classification using phase imaging with computational specificity. ACS photonics 9.4 (2022), pp. 1264–1273. DOI: 10.1021/acsphotonics.1c01779.

[7] Matthews H. K., Bertoli C., and Bruin R. A. de. Cell cycle control in cancer. Nature reviews Molecular cell biology 23.1 (2022), pp. 74–88.

[8] OpenStax. The Cell Cycle. Accessed: 2025-03-21. Salt Lake Community College Pressbooks. 2022. https://slcc.pressbooks.pub/collegebiology1/chapter/the-cell-cycle/.

[9] Murphy D. B. and Davidson M. W. Fundamentals of Light Microscopy and Electronic Imaging. Wiley-Blackwell, 2012.

[10] Ren J., Liu X., and Zhang Z. Brightfield microscopy-based cell cycle staging in fixed samples. Journal of Microscopy 282 (2021), pp. 123–130.

[11] Alberts B. et al. Molecular Biology of the Cell. 6th. Garland Science, 2015.

[12] Eastman A. E. and Guo S. The palette of techniques for cell cycle analysis. FEBS letters 594.13 (2020), pp. 2084–2098. DOI: 10.1002/1873-3468.13842.

[13] Chan G. K. Y., Kleinheinz T. L., Peterson D., and Moffat J. G. A simple high-content cell cycle assay reveals frequent discrepancies between cell number and ATP and MTS proliferation assays. PloS one 8.5 (2013), e63583.

[14] Nagao Y., Sakamoto M., Chinen T., Okada Y., and Takao D. Robust classification of cell cycle phase and biological feature extraction by image-based deep learning. Molecular biology of the cell 31.13 (2020), pp. 1346–1354. DOI: 10.1091/mbc.E20-03-0187.

[15] Sakaue-Sawano A., Kurokawa H., Morimura T., Hanyu A., Hama H., Osawa H., Kashiwagi S., Fukami K., Miyata T., Miyoshi H., et al. Visualizing spatiotemporal dynamics of multicellular cell-cycle progression. Cell 132.3 (2008), pp. 487–498. DOI: 10.1016/j.cell.2007.12.033.

[16] Gookin S., Min M., Phadke H., Chung M., Moser J., Miller I., Carter D., and Spencer S. L. A map of protein dynamics during cell-cycle progression and cell-cycle exit. PLoS biology 15.9 (2017), e2003268. DOI: 10.1371/journal.pbio.2003268.

[17] Sobecki M., Mrouj K., Colinge J., Gerbe F., Jay P., Krasinska L., Dulic V., and Fisher D. Cell-cycle regulation accounts for variability in Ki-67 expression levels. Cancer research 77.10 (2017), pp. 2722–2734. DOI: 10.1158/0008-5472.CAN-16-0707.

[18] Miller I., Min M., Yang C., Tian C., Gookin S., Carter D., and Spencer S. L. Ki67 is a graded rather than a binary marker of proliferation versus quiescence. Cell reports 24.5 (2018), pp. 1105–1112. DOI: 10.1016/j.celrep.2018.06.110.

[19] Doan M., Vorobjev I., Rees P., Filby A., Wolkenhauer O., Goldfeld A. E., Lieberman J., Barteneva N., Carpenter A. E., and Hennig H. Diagnostic potential of imaging flow cytometry. Trends in biotechnology 36.7 (2018), pp. 649–652. DOI: 10.1016/j.tibtech.2017.12.008.

[20] Blasi T., Hennig H., Summers H. D., Theis F. J., Cerveira J., Patterson J. O., Davies D., Filby A., Carpenter A. E., and Rees P. Label-free cell cycle analysis for high-throughput imaging flow cytometry. Nature communications 7.1 (2016), p. 10256. DOI: 10.1038/ncomms10256.

[21] Chong Y. T., Koh J. L., Friesen H., Duffy S. K., Cox M. J., Moses A., Moffat J., Boone C., and Andrews B. J. Yeast proteome dynamics from single cell imaging and automated analysis. Cell 161.6 (2015), pp. 1413–1424. DOI: 10.1016/j.cell.2015.04.051.

[22] Xu Y.-Y., Yang F., Zhang Y., and Shen H.-B. Bioimaging-based detection of mislocalized proteins in human cancers by semi-supervised learning. Bioinformatics 31.7 (2015), pp. 1111–1119. DOI: 10.1093/bioinformatics/btu772.

[23] Pärnamaa T. and Parts L. Accurate classification of protein subcellular localization from high-throughput microscopy images using deep learning. G3: Genes, Genomes, Genetics 7.5 (2017), pp. 1385–1392. DOI: 10.1534/g3.116.033654.

[24] Zhang Y., Xie Y., Liu W., Deng W., Peng D., Wang C., Xu H., Ruan C., Deng Y., Guo Y., et al. DeepPhagy: a deep learning framework for quantitatively measuring autophagy activity in Saccharomyces cerevisiae. Autophagy 16.4 (2020), pp. 626–640. DOI: 10.1080/15548627.2019.1632622.

[25] Eulenberg P., Köhler N., Blasi T., Filby A., Carpenter A. E., Rees P., Theis F. J., and Wolf F. A. Reconstructing cell cycle and disease progression using deep learning. Nature communications 8.1 (2017), p. 463. DOI: 10.1038/s41467-017-00623-3.

[26] Han-Shu F., Ming-Fei L., and Jing S. New methods for cell cycle analysis. Chinese Journal of Analytical Chemistry 47.9 (2019), pp. 1293–1301. DOI: 10.1016/S1872-2040(19)61186-2.

[27] Gupta A., Harrison P. J., Wieslander H., Pielawski N., Kartasalo K., Partel G., Solorzano L., Suveer A., Klemm A. H., Spjuth O., et al. Deep learning in image cytometry: a review. Cytometry Part A 95.4 (2019), pp. 366–380. DOI: 10.1002/cyto.a.23701.

[28] Meng N., Lam E. Y., Tsia K. K., and So H. K.-H. Large-scale multi-class image-based cell classification with deep learning. IEEE journal of biomedical and health informatics 23.5 (2018), pp. 2091–2098. DOI: 10.1109/JBHI.2018.2878878.

[29] Ferro A., Mestre T., Carneiro P., Sahumbaiev I., Seruca R., and Sanches J. M. Blue intensity matters for cell cycle profiling in fluorescence DAPI-stained images. Laboratory Investigation 97.5 (2017), pp. 615–625. DOI: 10.1038/labinvest.2017.13.

[30] LeCun Y., Bengio Y., and Hinton G. Deep Learning. Nature 521.7553 (2015), pp. 436–444. DOI: 10.1038/nature14539.

[31] Litjens G., Kooi T., Bejnordi B. E., Setio A. A. A., Ciompi F., Ghafoorian M., Laak J. A. van der, Ginneken B. van, and Sánchez C. I. A Survey on Deep Learning in Medical Image Analysis. Medical Image Analysis 42 (2017), pp. 60–88. DOI: 10.1016/j.media.2017.07.005.

[32] Caicedo J. C., Goodman A., Karhohs K. W., Cimini B. A., Ackerman J., Haghighi M., Heng C., Becker T., Doan M., McQuin C., et al. Nucleus Segmentation Across Imaging Experiments: The 2018 Data Science Bowl. Nature Methods 16.12 (2019), pp. 1247–1253. DOI: 10.1038/s41592-019-0612-7.

[33] Pawlowski N., Caicedo J. C., Singh S., Carpenter A. E., and Storkey A. Automated Analysis of High-Content Microscopy Data with Deep Learning. Methods 96 (2016), pp. 12–25. DOI: 10.1016/j.ymeth.2015.12.015.

[34] Zhou X., Taramelli R., Pedrali-Noy G., and Bianchi T. High-Content Screening Using Multivariate Image Analysis: A Review. Molecular BioSystems 8.3 (2012), pp. 744–755. DOI: 10.1039/C2MB05370B.

[35] Paul R., Achim A., Lorenz P., and Borau C. Automatic Cell Cycle Phase Classification Using Deep Convolutional Neural Networks. Scientific Reports 11.1 (2021), pp. 1–12. DOI: 10.1038/s41598-021-86221-0.

[36] Liu Q., Hu Q., Li J., and Gao Y. Deep Learning Enables Accurate Clustering with Batch Effect Removal in Single-Cell RNA-Seq Analysis. Nature Communications 10.1 (2019), pp. 1–14. DOI: 10.1038/s41467-019-13056-x.

[37] Zhang H., Wang J., Li X., and Chen M. Multi-Modal Learning for Robust Cell Cycle Classification. Bioinformatics 37.Supplement_1 (2021), pp. i153–i161. DOI: 10.1093/bioinformatics/btab290.

[38] Gao Y., Wang Z., Zhang Z., and Wu X. DeepCell: Automated Analysis of Cell Dynamics and Function Using Deep Learning. Nature Methods 15.11 (2018), pp. 1013–1020. DOI: 10.1038/s41592-018-0179-z.

[39] Holzinger A., Biemann C., Pattichis C. S., and Kell D. B. What Do We Need to Build Explainable AI Systems for the Medical Domain? arXiv (2017). DOI: 10.48550/arXiv.1712.09923.

[40] Ali M. A., Hollo K., Laasfeld T., Torp J., Tahk M.-J., Rinken A., Palo K., Parts L., and Fishman D. ArtSeg—Artifact segmentation and removal in brightfield cell microscopy images without manual pixel-level annotations. Scientific Reports 12.1 (2022), p. 11404.

[41] Shao J., Hu K., Wang C., Xue X., and Raj B. Is normalization indispensable for training deep neural network? Advances in Neural Information Processing Systems 33 (2020), pp. 13434–13444.

[42] Bjorck N., Gomes C. P., Selman B., and Weinberger K. Q. Understanding batch normalization. Advances in neural information processing systems 31 (2018).

[43] Seoni S., Shahini A., Meiburger K. M., Marzola F., Rotunno G., Acharya U. R., Molinari F., and Salvi M. All you need is data preparation: A systematic review of image

harmonization techniques in Multi-center/device studies for medical support systems. Computer Methods and Programs in Biomedicine (2024), p. 108200.
[44] Lin T.-Y., Goyal P., Girshick R., He K., and Dollár P. Focal loss for dense object detection. Proceedings of the IEEE international conference on computer vision. 2017, pp. 2980–2988.
[45] Breiman L. Random forests. Machine learning 45.1 (2001), pp. 5–32.
[46] Chen T. and Guestrin C. XGBoost: A scalable tree boosting system. Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mi 2016, pp. 785–794.
[47] Ronneberger O., Fischer P., and Brox T. U-net: Convolutional networks for biomedical image segmentation. International Conference on Medical image computing and computer-assisted intervention. Springer. 2015, pp. 234–241.
[48] He K., Zhang X., Ren S., and Sun J. Deep residual learning for image recognition. Proceedings of the IEEE conference on computer vision and pattern recognition. 2016, pp. 770–778.
[49] Tan M. and Le Q. EfficientNet: Rethinking model scaling for convolutional neural networks. International conference on machine learning. PMLR. 2019, pp. 6105–6114.
[50] Deng J., Dong W., Socher R., Li L.-J., Li K., and Fei-Fei L. ImageNet: A large-scale hierarchical image database. 2009 IEEE conference on computer vision and pattern recognition (2009), pp. 248–255.
[51] Buslaev A., Iglovikov V., Khvedchenya E., Parinov A., Druzhinin A., and Kalinin A. A. Albumentations: Fast and Flexible Image Augmentations. Information 11.2 (2020), p. 125.
[52] Kingma D. P. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014).
[53] University of Tartu. High-Performance Computing Cluster. https://hpc.ut.ee. 2024.
[54] OpenAI. ChatGPT (Mar 14 version). https://chat.openai.com. 2023.
[55] Inc. G. Grammarly Writing Assistant. https://www.grammarly.com. 2023.

License

Non-exclusive licence to reproduce the thesis and make the thesis public

I, Ali Zeynalli,
(author´s name)

1. grant the University of Tartu a free permit (non-exclusive licence) to

   reproduce, for the purpose of preservation, including for adding to the digital archives of the University of Tartu until the expiry of the term of copyright, my thesis

   Cell Cycle Phase Classification from Microscopy Images,
   (title of thesis)

   supervised by Dmytro Fishman;

   (supervisor´s name)

2. grant the University of Tartu a permit to make the thesis specified in point 1 available to the public via the web environment of the University of Tartu, including via the digital archives, under the Creative Commons licence CC BY NC ND 4.0, which allows, by giving appropriate credit to the author, to reproduce, distribute the work and communicate it to the public, and prohibits the creation of derivative works and any commercial use of the work until the expiry of the term of copyright;

3. am aware of the fact that the author retains the rights specified in points 1 and 2;

4. confirm that granting the non-exclusive licence does not infringe other persons' intellectual property rights or rights arising from the personal data protection legislation.

Ali Zeynalli
15/05/2025