

FlowCyt: A Comparative Study of Deep Learning Approaches for Multi-Class Classification in Flow Cytometry Benchmarking

Lorenzo Bini *

University of Geneva, Switzerland

LORENZO.BINI@UNIGE.CH

Fatemeh Nassajian Mojarad *

University of Geneva, Switzerland

FATEMEH.NASSAJIAN@UNIGE.CH

Margarita Liarou

University of Geneva, Switzerland

MARGARITA.LIAROU@UNIGE.CH

Thomas Matthes

Geneva University Hospital, Switzerland

THOMAS.MATTHES@HCUGE.CH

Stéphane Marchand-Maillet

University of Geneva, Switzerland

STEPHANE.MARCHAND-MAILLET@UNIGE.CH

Abstract

This paper presents FlowCyt, the first comprehensive benchmark for multi-class single-cell classification in flow cytometry data. The dataset comprises bone marrow samples from 30 patients, with each cell characterized by twelve markers. Ground truth labels identify five hematological cell types: T lymphocytes, B lymphocytes, Monocytes, Mast cells, and Hematopoietic Stem/Progenitor Cells (HSPCs). Experiments utilize supervised inductive learning and semi-supervised transductive learning on up to 1 million cells per patient. Baseline methods include Gaussian Mixture Models, XGBoost, Random Forests, Deep Neural Networks, and Graph Neural Networks (GNNs). GNNs demonstrate superior performance by exploiting spatial relationships in graph-encoded data. The benchmark allows standardized evaluation of clinically relevant classification tasks, along with exploratory analyses to gain insights into hematological cell phenotypes. This represents the first public flow cytometry benchmark with a richly annotated, heterogeneous dataset. It will empower the development and rigorous assessment of novel methodologies for single-cell analysis.

Data and Code Availability This paper uses the flow cytometry dataset provided by our University Hospital, which is available in our GitHub benchmark’s repository along with our code¹.

Institutional Review Board (IRB) This study was reviewed and approved by the Institutional Review Board of our University Hospital.

1. Introduction

In the constantly changing field of analyzing hematologic cell populations, multi-class classification remains a challenging task, requiring innovative diagnostic strategies capable of understanding the complexities of diverse cellular manifestations. Among various diagnostic tools, flow cytometry has become a key element in clinical practices, providing quick insights into cell populations and enabling the identification and characterization of abnormal cell types. However, the diagnostic potential of flow cytometry is counterbalanced by the formidable challenges posed by the complexity and heterogeneity of hematologic populations, requiring a shift in analytical methodologies.

We endeavor to offer the first public benchmark for flow cytometry data, providing automated solutions for hematological relevant tasks,

* These authors contributed equally.

1. [FlowCyt Webpage](#)

such as multi-class single-cell classification. It is noteworthy that this is the first available benchmark for flow cytometry data, equipped with ground truth, baseline models and advanced deep learning methods. This benchmark aims to provide a standardized platform for evaluating the performance of various algorithms and models, facilitating the development of novel, more efficient, and accurate diagnostic strategies.

1.1. Flow Cytometry

Flow cytometry (FC) is a laboratory technique used to detect and measure physical and chemical characteristics of population of cells or particles in a solution. In medicine, particularly in hematology and immunology, flow cytometry is mainly applied to characterize and count types of white blood cells in the evaluation of infectious diseases, autoimmune disorders, immunodeficiencies, or in the diagnosis of blood cancers such as leukemias or lymphomas.

Typically, samples obtained from the blood or bone marrow of patients are thereby labeled with antibodies that recognize specific surface proteins expressed by the cell populations present. The antibodies are coupled to fluorochromes so that when they are passed before a laser beam light is absorbed and then emitted in a band of wavelengths. Tens of thousands of cells can be quickly examined by modern flow cytometers in a matter of seconds and the data gathered subsequently processed by a computer. Highly complex data sets are thereby generated with data from up to several million cells and information about the presence or not of between 15 to 50 different cell proteins. Despite originating in the 1960s (Fulwyler, 1965; Gray et al., 1975), the fundamental structure of flow cytometry has undergone minimal alteration. Nevertheless, ongoing enhancements to both flow cytometers and fluorescent dyes have substantially accelerated cell analysis and expanded the range of detectable protein markers. In the 2000s, cytometry by time of flight (CyTOF), also referred to as mass spectrometry, was invented (Bandura et al., 2009; Bendall et al., 2011). CyTOF can identify isotope peaks without notable spectrum overlap with the use of heavy metal isotope coupled antibodies. This enables the simultaneous profiling of over 50 protein markers. In this work, we fo-

cus on providing and describing flow cytometry data, along with baseline performance of automated analysis technique.

1.2. Manual FC Data Analysis

Traditional manual gating employs a series of two-dimensional plots for data visualization and employs hierarchical gates to identify cell populations (Salama et al., 2022). It has the advantage that it can integrate prior knowledge into cytometry data analysis, encompassing the functionality of protein markers and the developmental correlation among cell populations. However, it remains a challenge to analyse the cytometry data due to its large cell numbers, high dimensionality, human bias and heterogeneity between datasets.

Special software has been developed which allow cytometrists to analyse this multi-dimensional data, to focus on specific cell populations, and to establish the phenotypic profile of a cell or cell population. This manual analysis is time-consuming, laborious, needs a high expertise, but still remains subjective and user-dependent.

Methods are therefore urgently needed which allow the analysis of this data in an efficient, rapid and precise way.

1.3. Related Work

There already exist numerous mathematical and statistical tools to analyze cytometric multi-dimensional data sets such as cell population identification (Hu et al., 2018; Mosmann et al., 2014; Finak et al., 2009; Aghaeepour et al., 2011; Van Gassen et al., 2015; Dorfman et al., 2016), data visualization (Amir et al., 2013; Becht et al., 2019; Qiu et al., 2011), sample classification (Bruggner et al., 2014; Van Gassen et al., 2016; Arvaniti and Claassen, 2017; Hu et al., 2020), batch normalization (Schuyler et al., 2019; Van Gassen et al., 2020), quality control (Monaco et al., 2016) and trajectory inference (Dai et al., 2021). These tools employ a wide range of computational methods which range from rule-based algorithms to machine learning models.

2. Dataset Description

In this section, we present an overview of Flow-Cyt, which is the first benchmark for single-cell multi-class classification using flow cytometry data. It utilizes a curated dataset comprising samples from 30 different patients.

2.1. Biological Description

Tabular data, especially in the field of biological data analysis, has a distinct and complicated structure. Each row in the table represents a single cell, which is characterized by several parameters like size, granularity, and fluorescence intensity. These parameters collectively constitute a high-dimensional feature space that captures the intricate biological variations among the cells. This data’s inherent structure often contains underlying relationships and dependencies that are not immediately apparent.

In this research, we begin by collecting raw data from the bone marrow of patients using a cytometer. The cytometer assesses the physical and chemical features of cells in a fluid as it passes through a laser beam. The surface molecules are fluorescently labeled and then activated by the laser to emit light at different wavelengths. Detectors collect the emitted light and transform it into electric signals. The raw data obtained from the cytometer are saved in Flow Cytometry Standard (FCS) files. Each FCS file comprises multi-dimensional data that corresponds to thousands of cells, with each cell characterized by various parameters such as size, granularity, and fluorescence intensity.

Thirty bone marrow (BM) samples were obtained from patients who underwent a flow cytometric analysis for diagnostic purposes. All samples were processed by the Diagnostics laboratory of the University Hospital and no malignant disease was detected. Samples were stained with the following antibodies: CD14=FITC, CD19=PE, CD13=ECD, CD13=PerCP5.5, CD34=PC7, CD117=APC, CD7=APC700, CD16=APC750, HLA-DR=PB and CD45=KO. Between 250’000 and 1’000’000 cells were acquired from each sample on a 10-color Navios cytometer (Beckman&Coulter) and analyzed manually with the KALUZA software (B&C’).

After the exclusion of cell debris, dead cells, and doublets, cell counts were established for the following five different cell populations: B lymphocytes (CD19pos), T lymphocytes (CD7pos), monocytes (CD14pos and CD33pos), mast cells (CD117 strongly pos), and HSPCs (hematopoietic stem and progenitor population, CD34pos). Therefore, we characterized in our dataset the following populations:

1. T Lymphocytes: These are a specific type of white blood cell that plays a pivotal role in the adaptive immune system, which is the response that involves the activation of immune cells to fight infection. They are responsible for directly killing infected host cells, activating other immune cells, producing cytokines, and regulating the immune response.
2. B Lymphocytes: These cells are significant contributors to the adaptive immune system. They are responsible for producing antibodies against antigens, which are substances that the immune system recognizes as foreign. Each mature B cell is programmed to make one specific antibody. When a B cell encounters its triggering antigen, it gives rise to many large cells known as plasma cells, each of which is essentially an antibody factory.
3. Monocytes: These are a type of white blood cell and a part of the innate immune system. They play a vital role in the body’s defense against infections and other foreign invaders. Monocytes circulate in the bloodstream, and when they migrate into tissues, they differentiate into macrophages or dendritic cells, which are capable of engulfing and digesting pathogens and apoptotic cells.
4. Mast cells: These are a type of immune cell that plays a crucial role in the body’s response to allergies and certain infections. Mast cells are found in most tissues, but especially in areas close to the external environment, such as the skin and mucous membranes. They contain granules filled with potent chemicals, including histamine, which they release in response to contact with an allergen. This release triggers inflammation, which can lead to allergic reactions.

5. Hematopoietic stem and progenitor cells (HSPC): These are a type of stem cell found in the bone marrow and cord blood. They have the unique ability to give rise to all other types of blood cells, including red blood cells, platelets, and all types of white blood cells. This makes them crucial for maintaining the body’s blood supply and immune system.

Each sample in the dataset includes ground truth labels indicating the corresponding hematological cell type. Researchers can use these labels to train and evaluate classification models aimed at discerning different cell types. Moreover, the dataset enables exploratory analyses such as clustering and dimensionality reduction, allowing comprehensive investigations into the heterogeneity and dynamics of hematological cell populations.

2.2. Technical Description

Table 1 displays the features that we utilized for all the analysis and tasks. All files have been gated by hematologists according to the list of cell types (labels) provided below.

Table 1: Flow cytometry data markers.

Class	Marker	Description
0	FS INT	Forward Scatter (FSC) - Cell’s size
1	SS INT	Side Scatter (SSC) - Cell’s granularity
2	CD14-FITC	Cluster of Differentiation 14 - Antigen
3	CD19-PE	Cluster of Differentiation 19 - Antigen
4	CD13-ECD	Cluster of Differentiation 13 - Antigen
5	CD33-PC5.5	Cluster of Differentiation 33 - Antigen
6	CD34-PC7	Cluster of Differentiation 34 - Antigen
7	CD117-APC	Cluster of Differentiation 117 - Antigen
8	CD7-APC700	Cluster of Differentiation 7 - Antigen
9	CD16-APC750	Cluster of Differentiation 16 - Antigen
10	HLA-PB	Human Leukocyte Antigen
11	CD45-KO	Cluster of Differentiation 45 - Antigen

Data and labels: Flow cytometry data is primarily tabular data. Each sample is represented by a $N \times D$ matrix of values. It is therefore equivalent to a D -dimensional point cloud that cytometrists visualize via 2D projections. The result of such operations is the definition and quantification of different cell populations according to their phenotype.

The dataset we propose contains samples from 30 patients; each sample with an average of $\bar{N} = 500'000$ cells (see supplementary material for the detailed count) and $D = 12$ dimensions (markers, according to Table 1). At that stage, data has not

undergone any other normalization than the one provided by the flow cytometry apparatus.

While we acknowledge that our study utilizes a dataset with 12 markers, this number is well within the range of typical flow cytometry panels used in clinical practice. Many previous studies have successfully demonstrated the utility of machine learning techniques on similar or even lower dimensional flow data. For instance, [Arvaniti and Claassen \(2017\)](#) applied representation learning on a 10-marker dataset to detect rare disease populations. [Bruggner et al. \(2014\)](#) identified stratifying signatures across 13 surface markers, while [Finak et al. \(2009\)](#) proposed a mixture model-based approach for cell population identification on a datasets with 8-12 markers, showcasing the effectiveness of their method in accurately identifying cell populations across diverse marker panels. Moreover, the key challenge in FC analysis does not solely stem from high dimensionality but rather from the inherent heterogeneity, complexity, and rarity of cell populations of interest amidst a majority of irrelevant cells. Our dataset, comprising over 1 million cells per patient, with 5 and 6 distinct and unbalanced cell types, accurately reflects these real-world challenges faced by cytometrists.

Files: Files have been anonymized and stored in “Flow Cytometry Standard” (FCS)² format that is generically produced by the flow cytometer. Readers for such format include the **FlowCal** Python library³ or the **flowCore** package in the larger **bioconductor** R library⁴.

Subsequently data files were cleaned manually with the help of the KALUZA software so that cell debris and dead cells were eliminated. The populations of interest were then defined manually and stored as separate files for each patient. For each patient, we obtained therefore seven data files:

- **Case_A.fcs** which contained the total cell population,
- **Case_0.fcs** (first class) contained only T lymphocytes,
- **Case_N.fcs** (second class) contained only B lymphocytes,

2. [International Society for Advancement of Cytometry](#)
 3. [Python FlowCal library](#)
 4. [R FlowCore package](#)

- **Case.G.fcs** (third class) contained only monocytes,
- **Case.P.fcs** (fourth class) contained only mast cells,
- **Case.K.fcs** (fifth class) contained only hematopoietic (HSPCs) cells,
- **Case.B.fcs** (sixth class) contained all the remaining cells denoted as $\{B\} = \{A\} \setminus \{O, N, G, P, K\}$.

Labels were kept aside and stored in basic “Comma Separated Value” (CSV) format that is easy to parse in any language. This offers the possibility to add new labels to the same data. The correspondence between cells (lines of FCS files) and their labels (lines in CSV files) is simply made by line numbering.

Overall, the data repository includes files for a total volume of around 1.3 GB. The data is distributed in our convenient group repository⁵.

3. Structure of the Benchmark

FlowCyt aims to facilitate the flow cytometry process undertaken by hematologists by promoting reproducibility and extensibility in single-cell classification research. The benchmarking framework and dataset are openly accessible, allowing researchers to explore novel methodologies and algorithms while reproducing existing results. By providing a standardized evaluation framework and a richly annotated dataset, FlowCyt empowers researchers to push the boundaries of single-cell analysis in hematological/immunological flow cytometry data.

3.1. Classification Task:

The main focus of our benchmark is to classify different cell types present in the heterogeneous samples. To make a fair comparison and evaluation of the performance of different baseline methods, we have provided various metrics, such as accuracy, precision, recall, and F1 score. Since it is a multi-class problem, these metrics are essential for ensuring robustness and comparability of the results. Moreover, being the dataset

strongly imbalanced, we have computed the ratio of the correctly predicted labels for each class in each patient. These measures will assist researchers in precisely evaluating and comparing the classification models.

3.1.1. SUB-POPULATION CLASSIFICATION:

We offer a specialized evaluation that focuses only on the classification of cells that belong to the $\{O, N, G, P, K\}$ population. This analysis allows researchers to examine the effectiveness of classification techniques in differentiating the variations within this specific subset of hematological cells. By isolating this distinct cell population, researchers can better understand the capacity of classification models to accurately identify and classify cells that are clinically relevant.

3.1.2. TOTAL POPULATION CLASSIFICATION:

To evaluate the learning robustness and generalization capabilities of the classification models, FlowCyt includes a classification task that involves the entire population A for each patient. By evaluating the models with the inherent complexity and diversity present within population A , researchers can determine how well the classification methodologies can work in real-world clinical scenarios.

3.2. Other Tasks:

The benchmark we’ve created for flow cytometry data analysis goes beyond classification and can be used for various analytical tasks, such as clustering, dimensionality reduction, trajectory inference, and anomaly detection. The dataset can be explored by researchers to gain insights into cellular types, making it a versatile resource for researchers working on FCS data. FlowCyt is therefore a resource that we aim to enrich with new data, labels, and task definitions (see Section 5.1).

It will also de facto get enriched by the community via its usage.

4. Experiments

We now discuss the details of our experimental design, which includes both inductive (IL) and transductive learning (TL) tasks. By combining

⁵. [FlowCyt Webpage](#)

these two learning methodologies, we aim to utilize their strengths to conduct a comprehensive analysis of our data. This approach also ensures that our results are more reliable and robust.

4.1. Techniques

To solve our tasks, which are mainly related to multi-class classification, we utilized well-known classification techniques as our baseline. All of these baselines were tested in the inductive learning framework, which was the basis for our medical evaluation, as depicted in Figure 1 and fully described in Section 4.1.2. Additionally, we tested the best-performing models in a transductive learning scenario to further evaluate the robustness of the learning process. This last scenario should be considered as an extension of the classification (inductive) task, where the reliability of the models was proven under semi-supervised learning, rather than full-supervised one.

4.1.1. BASELINES

FlowCyt offers a variety of methods for single-cell classification, ranging from traditional techniques like Gaussian Mixture Models (GMM), XGBoost (XGB), and Random Forest (RF) to more advanced methods like Deep Neural Networks (DNNs) and Graph Neural Networks (GNNs). GNNs stand out in terms of performance, demonstrating the ability of graph learning techniques to identify patterns in a cell’s neighborhood and their interactions. This enhances our understanding of hematological situations by utilizing spatial information and network structures.

To use these graph learning methods, FlowCyt provides the option to convert tabular data to graph-structured data that can be fed into the various proposed GNNs. We achieved this by treating each cell as a node in the graph space, with \mathcal{N}_i representing the set of k -nearest neighbors for node i (where $k = 7$) using the L_2 metric in the feature space. For each node, we built all possible pairs of its neighbors and added them to the edge set of the graph, resulting in a fully connected graph for each patient-graph instance. After experimenting with different values, we decided to use $k = 7$ to strike a balance between making the graph too sparse (which could miss

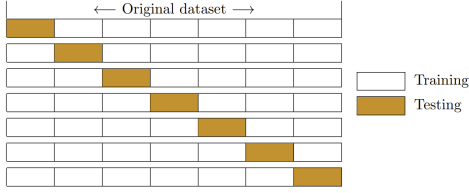
important connections) and too dense (which could include irrelevant connections). This approach enabled us to capture both the local structure (individual neighbors) and higher-order relationships (combinations of neighbors) in the data, helping us identify the most relevant relationships without adding noise. Moreover, constructing the graph via the k -nearest neighbor is intrinsically local and does not force us to make any assumptions about the distribution of the data, making it a good choice for datasets like ours where the full distribution of the data could be unknown.

4.1.2. INDUCTIVE LEARNING


To accomplish this specific task, we used the dataset from our cohort of 30 patients. To ensure a reliable evaluation methodology, we employed a randomized approach. We carried out a 7-fold test procedure and we selected 10% of each training fold uniformly at random as a validation set to evaluate the predicted labels \hat{y}_i by our model. This approach means that out of 30 patients, 4 or 5 are for testing, 2 or 3 for validation, and the rest for training. It’s important to note that the testing graphs are completely unseen during training.

We trained our model and adjusted the hyperparameters on the validation set to assess performance on the test set. We ran our models five times with five different seed initialization to ensure the robustness of our results.

Transductive Learning Motivated by the best results produced in Section 4.1.2, we further proved the effectiveness of the models by using the transductive learning methodology outlined in Yang et al. (2016), where during the training phase the models were given access to the feature vectors of all nodes. While the concept of transductive learning has been explored in various domains, its application to flow cytometry data is relatively novel, and there is little literature explicitly discussing its implications in this field. Hu et al. (2020) and Zhang et al. (2019) proposed a robust and interpretable end-to-end deep learning model for cytometry data and scRNA-Seq analysis, incorporating elements of self-supervised and semi-supervised learning. They discuss the advantages of leveraging unlabeled data to improve model generalization and



(a) 7-fold test procedure diagram



(b) Train/Val splitting

Figure 1: The dataset was split into seven segments for analysis. Six of these segments were used as training data, while one was reserved as test data. Additionally, 10% of each training segment was randomly selected to create a validation set.

robustness, which is crucial in the context of heterogeneous and high-dimensional cytometry data to reduce the need for extensive manual gating and annotation.

In FlowCyt, our primary goal was to classify the status of cells, distinguishing between the five different classes, while also evaluating the reliability of the classification performances in a semi-supervised setup utilizing up to one million cells per patient. To achieve this, we conducted five different seed iterations using our models. We randomly assigned 10% of the nodes for validation and 10% for testing while keeping the classes’ ratio balanced, the same as in the full dataset. Since our model is strongly imbalanced, we used a weighted negative log-likelihood loss function. Then we randomly masked the labels of 50% of the training, validation, and test set. We can therefore think of our problem as having a large graph where half of the graph is unlabeled, and our goal is to predict the labels based on the ones that we have.

4.2. Results

In the following sections, we provide a detailed account of our experimental setup. We present the classification results for both sub and total cell populations (see Section 3.1) in the inductive scenario. Additionally, we present the TL results for the previous best models.

4.2.1. SUB-POPULATION IL

We present here the results of the classification task as discussed in Section 3.1.1. We measure the performance in terms of accuracy, precision, recall, and F1 score as reported in Table 2. We compare state-of-the-art methods for tab-

Table 2: Average metrics across all 30 patients on sub-population. All the results are averaged with ± 0.01 , and the names are shortened.

Model	Accuracy	Precision	Recall	F1-Score
GAT	0.97	0.99	0.98	0.98
SAGE	0.98	0.97	0.98	0.98
GCN	0.96	0.96	0.95	0.95
DNN	0.86	0.88	0.84	0.80
RF	0.94	0.91	0.83	0.88
XGB	0.97	0.96	0.98	0.97
GMM	0.47	0.77	0.36	0.50

ular data classification, such as DNNs, XGB, RF, and GMM, with popular GNNs like GraphSAGE (SAGE) (Hamilton et al., 2017), Graph Convolutional Networks (GCNs) (Kipf and Welling, 2016), and Graph Attention Networks (GATs) (Veličković et al., 2017).

We calculate these metrics for each class i , using the following formula (where TP stands for True Positive, FN for False Negative, and FP for False Positive):

$$\begin{aligned}
 \text{Precision}_i &= \frac{\text{TP}_i}{\text{TP}_i + \text{FP}_i}, \\
 \text{Recall}_i &= \frac{\text{TP}_i}{\text{TP}_i + \text{FN}_i}, \\
 \text{F1 Score}_i &= 2 \times \frac{\text{Precision}_i \times \text{Recall}_i}{\text{Precision}_i + \text{Recall}_i}.
 \end{aligned} \tag{1}$$

We have similar equations for $i = 1, 2, 3, 4, 5$. Encouraged by the predictive ability of some of the models, we take a deeper look, and in Table 3

and Table 4 we show the (correct) predicted label across patients and cell types.

Table 3: Comparison of the average correct ratios across all 30 patients on the sub-population. Results are averaged with ± 0.01 .

Type	GAT	DNN	RF	XGB	GMM
T Cells	99.18	96.71	93.25	97.02	63.71
B Cells	90.53	73.62	90.60	82.63	51.00
Monocytes	97.89	90.28	90.38	93.01	25.06
Mast	98.28	88.78	1.74	49.28	30.00
HSPC	92.36	83.68	84.07	88.76	30.76

Table 4: Cont’d of Table 3 on the sup-population classification task.

Type	GAT	SAGE	GCN
T Cells	99.10	98.18	99.12
B Cells	90.53	86.86	84.22
Monocytes	97.19	93.71	96.55
Mast Cells	98.28	99.50	97.24
HSPC	92.36	93.25	92.14

4.2.2. A-POPULATION IL

We summarise in Table 5 the results on the total cell population (Section 3.1.2) for each patient, and we show in Table 6 of Appendix A the hyperparameters used throughout all the experiments. It is important to note that some models were not able to scale up on larger populations of up to one million nodes while maintaining a good ratio of correctly predicted labels. This indicates that only graph models, which are based on graph-encoded data, can capture the cell relations that are crucial for the classification task. While we only mentioned here GAT as a graph model, please refer to the code repository for all the results obtained using various GNNs related to Table 5. Moreover, moved by the extremely accurate results provided by GAT, we also interpret the choices made by the model, analyzing which markers (features) were considered more important than others, as shown in Figure 2 where the top 10 are shown. As we can see, the features that are given the most importance (at least above 55%) are 2,1,5,0 and 6, therefore we




Figure 2: Feature importance as highlighted by the model. For the sake of simplicity, we remind the reader to match the labels with the corresponding Table 1 for future explanations.

break down the choice of the model for each one. In order of importance, we have:

- **CD14-FITC.** The first important feature demonstrates the expression level of CD14. CD14 is a cell surface receptor that binds to lipopolysaccharide (LPS), which is a component of the bacterial cell wall. CD14 is mainly expressed by monocytes, macrophages, and activated granulocytes. It plays a role in the immune response against bacterial infections. A high importance of this feature may imply that our model can distinguish between cells that are involved in innate immunity and those that are not, detecting the presence of bacterial infection in the sample.
- **Side Scatter (SSC)-Cell’s granularity.** This feature denotes the level of light scatter at a 90-degree angle in relation to the laser beam. SSC reflects the internal complexity or granularity of the cell, encompassing features like the presence of granules, nuclei, or other organelles. A high importance of this feature may suggest that either our model can differentiate between cells that have different degrees of complexity, such as lymphocytes, monocytes, and granulocytes, or our model can identify cells that have abnormal granularity, such as blast cells or malignant cells.

- **CD33-PC5.5.** This feature indicates the expression level of CD33, a cell surface receptor that belongs to the sialic acid-binding immunoglobulin-like lectin (Siglec) family. CD33 is expressed by myeloid cells, such as monocytes, macrophages, granulocytes, and mast cells, and it modulates the immune response by inhibiting the activation of these cells. A high importance of this feature may suggest that our model can distinguish between myeloid and non-myeloid cells, or that our model can detect the expression of CD33 as a marker for certain types of leukemia, such as acute myeloid leukemia (AML) or chronic myelomonocytic leukemia (CMML).
- **Forward scatter (FSC)-Cell’s size.** This feature indicates the level of light scatter along the path of the laser beam. FSC is proportional to the diameter or surface area of the cell, and it can be used to discriminate cells by size. Placing emphasis on this feature may suggest that our model can differentiate between cells that have different sizes, such as small lymphocytes and large monocytes.
- **CD34-PC7.** This feature indicates the expression level of CD34, which is a cell surface glycoprotein that belongs to the sialomucin family. CD34 is expressed by HSPCs, and functions as a cell-cell adhesion molecule. The importance on this feature suggests that the model can identify the presence of HSPCs in the sample.

Table 5: Comparison of the average correct ratios across all 30 patients on the A-population. Results are averaged and rounded with ± 0.01 .

Type	GAT	DNN	RF	XGB	GMM
T Cells	97.91	91.86	62.67	81.45	72.04
B Cells	95.17	56.60	-	6.33	-
Monocytes	98.03	89.31	-	72.63	68.74
Mast	97.40	47.09	-	15.54	-
HSPC	90.79	32.79	-	24.89	-
Others	87.77	94.82	91.07	80.03	-

4.2.3. POPULATIONS ON TL

In this section, we discuss the results obtained for the sub-population on TL set-up (see Section

3.1.1). We have evaluated our metrics and found that the F1 score is around $\{0.94 \pm 0.01\}$ for GAT, GCN, and SAGE, which confirms the robustness of graph models and graph encoded data even for semi-supervised approach with large graphs. The results for the employed models have been obtained using those hyperparameters to fit the constraints of 1 GPU with 11GB memory (RTX 2080 Ti), as shown in Table 7 of Appendix A. To highlight the strength of graph models, Figure 3 shows the t-SNE (Van der Maaten and Hinton, 2008) embeddings visualization of GAT for one random patient (patient 23 has been randomly picked). Despite the difficulties due to the mask setup that is typical of transductive learning, we can see how the graph model has clustered the five classes into its final embedding layer. To view the results for the sub-population under the TL framework, please refer to our code repository.




Figure 3: t-SNE projection for one random patient, for the transductive learning task.

5. Discussion of FlowCyt

FlowCyt provides a valuable resource for researchers in the field of hematological cell populations. The dataset is extensively annotated and captures the complexity and heterogeneity of these cell populations. The high-dimensional tabular data allows for various analytical tasks, such as dimensionality reduction, clustering, and trajectory analysis, to be performed beyond classification. By utilizing their domain knowledge, researchers can explore the biological relation-

ships and hierarchical structures among the cell types based on the provided labels.

The feasibility of developing automated solutions to assist cytometrists has been demonstrated through multi-class classification experiments throughout Section 4. The effectiveness of graph neural networks in modeling spatial patterns and inter-cell dependencies via graph encoding has been highlighted by the results obtained. This benchmark represents the first public flow cytometry benchmark that empowers rigorous assessment of novel methodologies. It will facilitate research that pushes the boundaries and clinical translation of automated single-cell analysis to supplement traditional manual examination. Through collaborative efforts, this versatile dataset may uncover additional tasks and biological insights, thus contributing to advancing the field of hematological cell populations (Bini et al., 2024).

5.1. Future Research

This benchmark is a solid foundation for expanding research in flow cytometry through diverse data and advanced analytical tasks. To better represent real clinical scenarios and improve model robustness, it is important to increase our patient cohort and include patients with hematologic/immunologic diseases, such as leukemias (e.g. Acute Myeloid Leukemia-AML). Obtaining longitudinal samples that track patients over treatment courses could also reveal informative dynamics and could help to understand how cell populations evolve in response to therapies.

Advanced tasks should be defined in collaboration with hematologists/immunologists to ensure clinical relevance. There are opportunities to develop methodologies that support rather than replace manual analysis, providing a supplemental perspective. For example, automated strategies for quality control, across-patient normalization, batch effect correction, and data visualization could aid cytometrists. Beyond improving existing practices, innovative new paradigms could extract novel biological insights.

Moreover, there are numerous learning paradigms for flow cytometry: hierarchical relationships between hematopoietic cells could be modeled (Nassajian Mojarrad et al., 2024), aligning with developmental pathways to iden-

tify pathological deviations from healthy cell distribution patterns. Hybrid approaches could combine unsupervised representation learning with supervised diagnosis.

Another promising direction is single-cell trajectory inference (TI), where apart from providing insights into phenotypes of different blood cell types, FlowCyt’s dataset can also be used for better understanding the developmental trajectories of hematopoietic precursors. The use of our dataset is not limited to supervised classification problems and can be extended to benchmarking different trajectory inference methods, which so far have been reviewed for transcriptomics (Saelens et al., 2019) and not flow cytometry data. In such a case, access to cell type annotations will be valuable to evaluate the quality of the reconstructed trajectories and compare the performance of different TI methods suitable for non-time resolved data. Moreover, our cohort data will allow us to establish a healthy trajectory model that will serve in the future as a reference to comprehend in more detail the differences between healthy and malignant hematopoiesis.

6. Conclusion

Overall, this benchmark enables diverse analytical approaches from conventional supervised learning to emerging paradigms in flow cytometry research. Collaborative efforts to generate new labeled datasets, formulate meaningful clinical questions, and develop inventive computational methodologies will catalyze impactful discoveries in hematology and immunology. By fostering rigorous yet creative investigations leveraging its versatility, this benchmark aims to accelerate the advancement of automated single-cell analysis.

Acknowledgments

The Swiss National Science Foundation partially funds this work under grant number 207509 "Structural Intrinsic Dimensionality".

References

N. Aghaeepour, R. Nikolic, H.H. Hoos, and R.R. Brinkman. Rapid cell population identification

- in flow cytometry data. *Cytom Part J Int Soc Anal Cytol*, 79:6–13, 2011.
- E.A.D. Amir, K.L. Davis, M.D. Tadmor, E.F. Simonds, J.H. Levine, S.C. Bendall, D.K. Shenfeld, S. Krishnaswamy, G.P. Nolan, and D. Pe’er. visne enables visualization of high dimensional single-cell data and reveals phenotypic heterogeneity of leukemia. *Nature biotechnology*, 31(6):545–552, June 2013.
- E. Arvaniti and M. Claassen. Sensitive detection of rare disease-associated cell subsets via representation learning. *nature communications*, 8(1):14825, April 2017.
- D.R. Bandura, V.I. Baranov, O.I. Ornatsky, A. Antonov, R. Kinach, X. Lou, S. Pavlov, S. Vorobiev, J.E. Dick, and S.D. Tanner. Mass cytometry: technique for real time single cell multitarget immunoassay based on inductively coupled plasma time-of-flight mass spectrometry. *Analytical chemistry*, 81(16):6813–6822, August 2009.
- E. Becht, L. McInnes, J. Healy, C.A. Dutertre, I.W. Kwok, L.G. Ng, F. Ginhoux, and E.W. Newell. Dimensionality reduction for visualizing single-cell data using umap. *Nature biotechnology*, 37(1):38–44, January 2019.
- S.C. Bendall, E.F. Simonds, P. Qiu, E.A.D. Amir, P.O. Krutzik, R. Finck, R.V. Bruggner, R. Melamed, A. Trejo, O.I. Ornatsky, and R.S. Balderas. Single-cell mass cytometry of differential immune and drug responses across a human hematopoietic continuum. *Science*, 332(6030):687–696, May 2011.
- L. Bini, F. Nassajian Mojarrad, T. Matthes, and S. Marchand-Maillet. Hemagraph: Breaking barriers in hematologic single cell classification with graph attention. *arXiv preprint, arXiv:2402.18611*, 2024.
- R.V. Bruggner, B. Bodenmiller, D.L. Dill, R.J. Tibshirani, and G.P. Nolan. Automated identification of stratifying signatures in cellular subpopulations. *Proceedings of the National Academy of Sciences*, 111(26):E2770–E2777, July 2014.
- Y. Dai, A. Xu, J. Li, L. Wu, S. Yu, J. Chen, W. Zhao, X.J. Sun, and J. Huang. Cytotree: an r/bioconductor package for analysis and visualization of flow and mass cytometry data. *BMC bioinformatics*, 22(1):1–20, 2021.
- D.M. Dorfman, C.D. LaPlante, and B. Li. Flock cluster analysis of plasma cell flow cytometry data predicts bone marrow involvement by plasma cell neoplasia. *Leukemia Research*, 48:40–45, September 2016.
- G. Finak, A. Bashashati, R. Brinkman, and R. Gottardo. Merging mixture components for cell population identification in flow cytometry. *Advances in bioinformatics*, 2009, 2009.
- M.J. Fulwyler. Electronic separation of biological cells by volume. *Science*, 150(3698):910–911, November 1965.
- J.W. Gray, A.V. Carrano, L.L. Steinmetz, M.A. Van Dilla, D.H. Moore 2nd, B.H. Mayall, and M.L. Mendelsohn. Chromosome measurement and sorting by flow systems. *Proceedings of the National Academy of Sciences*, 72(4):1231–1234, April 1975.
- W. Hamilton, Z. Ying, and J. Leskovec. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30, 2017.
- Z. Hu, C. Jujjavarapu, J.J. Hughey, S. Andorf, H.C. Lee, P.F. Gherardini, M.H. Spitzer, C.G. Thomas, J. Campbell, P. Dunn, and J. Wiser. Metacyto: a tool for automated meta-analysis of mass and flow cytometry data. *Cell reports*, 24(5):1377–1388, July 2018.
- Z. Hu, A. Tang, J. Singh, S. Bhattacharya, and A.J. Butte. A robust and interpretable end-to-end deep learning model for cytometry data. *Proceedings of the National Academy of Sciences*, 117(35):21373–21380, September 2020.
- T.N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- G. Monaco, H. Chen, M. Poidinger, J. Chen, J.P. de Magalhães, and A. Larbi. flowai: automatic and interactive anomaly discerning tools for flow cytometry data. *Bioinformatics*, 32(16):2473–2480, August 2016.

- T.R. Mosmann, I. Naim, J. Rebhahn, S. Datta, J.S. Cavanaugh, Weaver J.M., and G. Sharma. Swift-scalable clustering for automated identification of rare cell populations in large, high-dimensional flow cytometry datasets, part 2: Biological evaluation. *Cytom Part J Int Soc Anal Cytol*, 85:422–433, February 2014.
- F. Nassajian Mojarrad, L. Bini, T. Matthes, and S. Marchand-Maillet. Why attention graphs are all we need: Pioneering hierarchical classification of hematologic cell populations with leukograph. *arXiv preprint, arXiv:2402.18610*, 2024.
- P. Qiu, E.F. Simonds, S.C. Bendall, K.D. Gibbs Jr, R.V. Bruggner, M.D. Linderman, K. Sachs, G.P. Nolan, and S.K. Plevritis. Extracting a cellular hierarchy from high-dimensional cytometry data with spade. *Nature biotechnology*, 29(10):886–891, October 2011.
- W. Saelens, R. Cannoodt, H. Todorov, and Y. Saeys. A comparison of single-cell trajectory inference methods. *Nature biotechnology*, 37(5):547–554, 2019.
- M.E. Salama, G.E. Otteson, J.J. Camp, J.N. Seheult, D. Jevremovic, D.R. Holmes III, H. Olteanu, and M. Shi. Artificial intelligence enhances diagnostic flow cytometry workflow in the detection of minimal residual disease of chronic lymphocytic leukemia. *Cancers*, 14(10):2537, May 2022.
- R.P. Schuyler, C. Jackson, J.E. Garcia-Perez, R.M. Baxter, S. Ogolla, R. Rochford, D. Ghosh, P. Rudra, and E.W. Hsieh. Minimizing batch effects in mass cytometry data. *Frontiers in immunology*, 10:2367, October 2019.
- L. Van der Maaten and G. Hinton. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11), 2008.
- S. Van Gassen, B. Callebaut, M.J. Van Helden, B.N. Lambrecht, P. Demeester, T. Dhaene, and Y. Saeys. Flowsom: Using self-organizing maps for visualization and interpretation of cytometry data. *Cytom Part J Int Soc Anal Cytol*, 87:636–645, 2015.
- S. Van Gassen, C. Vens, T. Dhaene, B.N. Lambrecht, and Y. Saeys. Floremi: Flow density survival regression using minimal feature redundancy. *Cytom Part J Int Soc Anal Cytol*, 89:22–29, 2016.
- S. Van Gassen, B. Gaudilliere, M.S. Angst, Y. Saeys, and N. Aghaeepour. Cytonorm: a normalization algorithm for cytometry data. *Cytometry Part A*, 97(3):268–278, March 2020.
- P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Lio, and Y. Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.
- Z. Yang, W. Cohen, and R. Salakhudinov. Revisiting semi-supervised learning with graph embeddings. In *International conference on machine learning*, pages 40–48. PMLR, 2016.
- Z. Zhang, D. Luo, X. Zhong, J.H. Choi, Y. Ma, S. Wang, E. Mahrt, W. Guo, E.W. Stawiski, Z. Modrusan, and S. Seshagiri. SCINA: a semi-supervised subtyping algorithm of single cells and bulk samples. *Genes*, 10(7):531, 2019.

Appendix A. Hyperparameters

Tables 7 and Table 6 below show the default hyperparameters used in all the experiments:

Table 6: Default hyperparameters used in all inductive experiments. For XGBoost, Random Forest, and Gaussian Mixture Model we used the sklearn package as shown in our code repository.

Params	GAT	SAGE	GCN	DNN
Layers	3	3	3	3
Hidden-ch	64	64	64	256
Att-heads	8	-	-	-
Optimizer	Adam	Adam	Adam	Adam
Lr-sched	[0.01,1e-7]	[0.01,1e-7]	[0.01,1e-7]	[0.01,1e-7]
Weigh-decay	0.0005	0.005	0.005	0.005
Dropout	0.3	0.4	0.4	0.5
Epochs	1000	1000	1000	1000
Early-stop	40	50	50	40

Table 7: Default hyperparameters used in all transductive experiments to fit our memory GPU constraints.

Parameter	GAT	SAGE	GCN
Layers	3	3	3
Hidden-ch	8	16	16
Att-heads	8	-	-
Optimizer	Adam	Adam	Adam
Lr scheduler	[0.01,1e-7]	[0.01,1e-7]	-
Weight decay	0.0005	0.005	0.005
Dropout	0.3	0.5	0.5
Train epochs	1000	1000	1000
Early stopping	40	50	50