haematologica
Journal of the Ferrata Storti Foundation

Accurate automated diagnosis of B-acute lymphoblastic leukemia using deep learning and flow cytometry

by Sulov Chalise, Mikhail Roshal, Sophia Roshal, Jeeyeon Baik, Qi Gao, Anyi Li, Ahmet Dogan, Harini Veeraraghavan and Meng-Lei Zhu

Received: May 30, 2025.
Accepted: December 24, 2025.

Citation: Sulov Chalise, Mikhail Roshal, Sophia Roshal, Jeeyeon Baik, Qi Gao, Anyi Li, Ahmet Dogan, Harini Veeraraghavan and Meng-Lei Zhu. Accurate automated diagnosis of B-acute lymphoblastic leukemia using deep learning and flow cytometry.
Haematologica. 2026 Jan 8. doi: 10.3324/haematol.2025.288277 [Epub ahead of print]

Accurate automated diagnosis of B-acute lymphoblastic leukemia using deep learning and flow cytometry

Sulov Chalise[1], Mikhail Roshal[1], Sophia Roshal[2], Jeeyeon Baik[1], Qi Gao[1], Anyi Li[3], Ahmet Dogan[1], Harini Veeraraghavan[3]*, Meng-Lei Zhu[1]*

[1] Department of Pathology and Laboratory Medicine, Memorial Sloan Kettering Cancer Center, New York, NY, USA
[2] Carnegie Mellon University, Pittsburgh, PA, USA
[3] Department of Medical Physics, Memorial Sloan Kettering Cancer Center, New York, NY, USA

Running heads.
AI-based automatic flow cytometry analysis for B-ALL

Corresponding authors*
Harini Veeraraghavan, Department of Medical Physics, Memorial Sloan Kettering Cancer Center, New York, NY

Meng-Lei Zhu, Department of Pathology and Laboratory Medicine, Memorial Sloan Kettering Cancer Center, New York, NY

Disclosures
No relevant conflicts of interest to disclose.

Contributions
SC, MZ ,HV, MR, and AD designed the study. SC conducted data processing, model training, and performance evaluation. MR performed expert annotation of flow cytometry data. SR, JB, QG, and AL assisted in data collection and technical validation. SC, MR, AL, MZ, and HV wrote and reviewed the manuscript.

Data-sharing statement
The data that supports the findings of this study are available from the corresponding authors upon reasonable request.

Multi-parameter flow cytometry (MFC) is an essential ancillary technique used in B-lymphoblastic leukemia/lymphoma (B-ALL) diagnosis and monitoring to identify abnormal B precursor cell populations in patient specimens. Minimal (measurable) residual disease (MRD) analysis is essential for clinical management of B-ALL. [1] In experts' hands, MFC is a well-established, relatively inexpensive, and highly accurate technique for MRD with sensitivities as high as 2 in 10^6 cells. [2] Currently, accurate MFC result interpretation primarily depends on the expertise of highly trained personnel. Despite panel and instrument standardization, interpretive challenges remain a barrier to broader utilization. For example, a recent study involving expert, high-volume laboratories participating in Children's Oncology B-ALL trials revealed substantial variability in distinguishing normal immature B cells from leukemic counterparts. [3] Automated analysis approaches may, in theory, improve the situation, but have yet to show sufficient accuracy to be impactful in routine clinical practice.

In recent years, automated MFC analysis for cell population identification and characterization has been demonstrated in research studies.[4] CellCNN, a supervised deep-learning model that utilizes annotated data, can identify rare disease-associated cells using a small flow cytometry dataset.[5] Recent machine learning studies on B-ALL flow cytometry analysis, utilizing methods such as Gaussian mixture model–based pipelines, SVM-optimized radar plots, and deep neural networks, have developed high-performing tools that support expert interpretation and show strong concordance with manual gating. [6-8] However, comparisons across studies remain difficult due to differences in analytical focus, MRD inclusion, and evaluation metrics. Despite these advances, current systems have not yet achieved clinical-grade accuracy as standalone tools and are best suited for human-in-the-loop applications rather than fully automated assessment.[9]

To narrow the gap to the real-world clinical practice, we developed FlowARC (Flow analysis using Residual Convolutional network), a novel deep-learning approach for precise and automated B-ALL detection, including in MRD setting. FlowARC model employs a cascaded 3-stage architecture (Figure 1), a cell-level module (Audit stage), a

cell-ranking step (Reorder stage), and a sample-level module (Classify stage). In the Audit stage, individual cells within a flow cytometry sample are classified as normal or tumorous, with a leukemic probability score assigned to each cell. The Reorder stage ranks all cells by their leukemic probability and retains the top 7,500 most likely leukemic cells for further analysis. In the Classify stage, this subset is used by the sample-level module to determine whether the overall sample is tumorous or normal. Additionally, a separate quantification module estimates tumor burden, an essential feature for clinical reporting, particularly in the context of MRD.

FlowARC was trained and evaluated using a large retrospective clinical cohort of patients tested for B-ALL at our institution between 2015 and 2019, using an 8-color flow cytometry assay that followed standard diagnostic protocols. [10, 11] This assay included surface markers (CD20, CD34, CD10, CD33, CD58, CD45, CD19, CD38) along with forward scatter (FSC-H, FSC-A) and side scatter (SSC-H, SSC-A) parameters. The cohort comprised 1,681 flow cytometry samples from 333 patients (184 male, 149 female), including 1,149 B-ALL-positive and 532 B-ALL-negative cases, totaling approximately 85.7 million B cells. Most samples were derived from bone marrow (1,141 positive and 532 negative), with a small number from peripheral blood (n = 6 positive) and tissue (n = 2 positive). The mean age of patients was 32.8 years (SD 22.2, range 0-81 years). The Institutional Review Board at Memorial Sloan Kettering Cancer Center approved the study.

All FCS files from the cohort were processed using a standardized preprocessing pipeline including compensation and transformation using flowCore [12] (R package version 2.14.2) as well as automatic B-cells extraction by flowDensity (Supplementary Figure 1). [13] In negative cases, all extracted B-cells were labeled as normal. For a representative subset of 137 B-ALL positive bone marrow samples, abnormal immature B-cells were manually annotated by an expert hematopathologist (MR). These annotations were used to train the cell-level module and to generate synthetic datasets for training the sample-level and quantification modules. Synthetic data were generated to enhance the diversity of tumor cell populations across a wide range of tumor burdens,

including 25 to 500 B-cells (low MRD), 500 to 25,000 B-cells (high MRD), 25,000 to 50,000 B-cells (low tumor burden), and 50,000 to 500,000 B-cells (high tumor burden). In our synthetic sample generation, we set a minimum leukemic population of 25 cells, corresponding to a theoretical sensitivity of up to 10^{-5}, depending on the total number of cells analyzed, reflecting the real-world detection threshold of the underlying flow cytometry assay (Supplementary Figure 2). To avoid data leakage, patient-level splits were performed prior to synthetic data generation. Among negative cases, 65% were randomly allocated for training. Of the 137 manually annotated positive samples, 75% were used for training, while the remaining 25%, along with all unannotated positive cases, were reserved for testing. To evaluate model robustness, five independent instances of the cell-level, sample-level, and quantification modules were trained using different random weight initializations. Data preprocessing and model training (Supplementary Table 1) codes are available at https://github.com/MSK-Computational-HemePath/FlowARC.

FlowARC demonstrated high performance in differentiating abnormal B-cell precursors from normal cells at the cell level, achieving an area under the receiver operating characteristic curve (AUROC) of 0.995 (95% CI: 0.994–0.995). The corresponding confusion matrix included 10,028,906 true normal cells, 2,057,835 true tumor cells, 77,621 false normal cells (0.6%), and 341,424 false tumor cells (2.7%) (Figure 2C). At the sample level, reflecting clinical diagnoses, FlowARC continued to show excellent results. It achieved an AUROC of 0.994 (95% CI: 0.990–0.998) on the synthetic test set and maintained similarly high performance on real-world patient samples with an AUROC of 0.991 (95% CI: 0.987–0.995), accuracy of 97.5%, sensitivity of 96.4%, and specificity of 98.7% (Figure 2A). The sample-level confusion matrix comprised 225 true normal, 245 true tumor, 4 false normal (0.8%), and 6 false tumor cases (1.6%) (Figure 2D). Compared with the CellCNN model trained and tested on identical preprocessed data, FlowARC significantly outperformed CellCNN (AUROC: 0.995 vs. 0.869; $p < 10^{-11}$ by DeLong's test) (Figure 2B). Tumor burden quantification model was performed on real-world patient samples, and demonstrated strong agreement with expert-assessed tumor content, achieving an R^2 of 0.879 (95% CI: 0.877–0.881), a slope of 0.963 (95%

CI: 0.960–0.966), and an intercept of 0.111 (95% CI: 0.095–0.127), indicating that the model's estimates closely matched the true measured values (Figure 2C).

Beyond accurate B-ALL detection, FlowARC provides explainable outputs and visual tools to aid clinical interpretation. Using the SHAP (Shapley Additive ExPlanations) algorithm, we identified key markers contributing to predictions, with low CD38**,** high forward scatter (FSC-A), and increased CD58 emerging as the strongest predictors of tumor cells, consistent with known disease phenotypes (Figure 2D). These SHAP values help elucidate the phenotypic profiles underlying the model's abnormal cell predictions.

FlowARC's cell-level predictions can be visualized using standard bi-parametric flow cytometry plots. In true tumor cases, high-probability cells form distinct clusters across multiple markers (Figure 3A, bottom), whereas in true negatives or false positives, misclassified cells appear scattered without clustering (Figure 3A, top; 3B, top). In false negatives, abnormal cells may still form clusters resembling true positives (Figure 3B, bottom). These visualization features support pathologist-in-the-loop verification, improving diagnostic accuracy and interpretability. The same approach is effective for low-cell-count samples (<7,500 B-cells), where true tumors form compact clusters and true negatives remain diffuse (Figure 3C), extending FlowARC's utility even in challenging cases.

In summary, we developed the FlowARC model for automated detection of B-ALL, achieving over 0.99 AUROC in real-world clinical cases (Figure 1E). Its novel cascaded architecture, combining cell-level and sample-level modules, overcomes the limitations of previous methods and enables clinical-grade detection of lymphoblastic lymphoma in both diagnostic and MRD samples. The model achieved performance comparable to expert pathologists, with superior reproducibility and consistency, whereas human performance is influenced by inter-observer variability with reported concordance of 74–93%.[3, 9, 14] Misclassifications primarily occurred in samples with very low tumor burden (<0.01%) or overlapping aberrant and regenerative immunophenotypes, which are

scenarios that challenge both AI and human reviewers, while cell-level visualization of prediction scores helped identify and correct some of these borderline cases.

Although FlowARC was developed and validated using data from our institution, its architecture and training strategy are broadly adaptable. The model is panel-specific, but the training framework is universally applicable and can be retrained on any clinical flow cytometry panel with limited locally annotated data. Model performance aligns with the analytical capability of the target clinical assay, facilitating cross-institutional consistency in diagnostic quality. Moreover, the same architecture can be adapted to other leukemia types, such as AML and T-ALL, by retraining on disease-specific panels and ensuring sufficient positive cases to account for immunophenotypic variability.

FlowARC addresses the extensive data requirements of deep learning by using synthetic samples generated from real clinical data, which closely mirror real-world cellular variability. This approach enables effective training with a limited number of curated cases. An ablation study showed that approximately 200 annotated cases (100 normal and 100 abnormal) are sufficient to achieve clinically acceptable performance, with only a modest decline from the full model. Using locally representative cases, ideally verified by an expert hematopathologist, further ensures FlowARC's practicality and accessibility for clinical laboratories integrating AI-driven diagnostics into routine workflows.

The resulting automation enabled by FlowARC holds considerable potential for optimizing diagnostic and MRD evaluation workflows. It streamlines key analytical processes, resulting in notably shorter diagnostic turnaround times, improved workflow efficiency, and reduced variability from manual interpretation. These enhancements in diagnostic accuracy and speed are likely to lead to more informed clinical decisions and improved patient outcomes.

Reference:

1. Borowitz MJ, Wood BL, Devidas M, et al. Prognostic significance of minimal residual disease in high risk B-ALL: a report from Children's Oncology Group study AALL0232. Blood. 2015;126(8):964-971.
2. Tembhare PR, Subramanian Pg PG, Ghogale S, et al. A High-Sensitivity 10-Color Flow Cytometric Minimal Residual Disease Assay in B-Lymphoblastic Leukemia/Lymphoma Can Easily Achieve the Sensitivity of 2-in-10(6) and Is Superior to Standard Minimal Residual Disease Assay: A Study of 622 Patients. Cytometry B Clin Cytom. 2020;98(1):57-67.
3. Keeney M, Wood BL, Hedley BD, et al. A QA Program for MRD Testing Demonstrates That Systematic Education Can Reduce Discordance Among Experienced Interpreters. Cytometry B Clin Cytom. 2018;94(2):239-249.
4. Hu Z, Bhattacharya S, Butte AJ. Application of Machine Learning for Cytometry Data. Front Immunol. 2021;12:787574.
5. Arvaniti E, Claassen M. Sensitive detection of rare disease-associated cell subsets via representation learning. Nat Commun. 2017;8:14825.
6. Reiter M, Diem M, Schumich A, et al. Automated Flow Cytometric MRD Assessment in Childhood Acute B- Lymphoblastic Leukemia Using Supervised Machine Learning. Cytometry A. 2019;95(9):966-975.
7. Shopsowitz KE, Liu L, Setiadi A, et al. Machine learning optimized multiparameter radar plots for B-cell acute lymphoblastic leukemia minimal residual disease analysis. Cytometry B Clin Cytom. 2022;102(5):342-352.
8. Seheult JN, Otteson GE, Timm MM, et al. Artificial Intelligence Accelerates the Interpretation of Measurable Residual B Lymphoblastic Leukemia by Flow Cytometry. Blood Adv. 2025 Sep 18. doi:10.1182/bloodadvances.2025016126. [Epub ahead of print]
9. Verbeek MWC, Rodriguez BS, Sedek L, et al. Minimal residual disease assessment in B-cell precursor acute lymphoblastic leukemia by semi-automated identification of normal hematopoietic cells: A EuroFlow study. Cytometry B Clin Cytom. 2024;106(4):252-263.
10. Gao Q, Liu Y, Aypar U, et al. Highly sensitive single tube B-lymphoblastic leukemia/lymphoma minimal/measurable residual disease test robust to surface antigen directed therapy. Cytometry B Clin Cytom. 2023;104(4):279-293.
11. Geyer MB, Ritchie EK, Rao AV, et al. Pediatric-inspired chemotherapy incorporating pegaspargase is safe and results in high rates of minimal residual disease negativity in adults up to age 60 with Philadelphia chromosome-negative acute lymphoblastic leukemia. Haematologica. 2021;106(8):2086-2094.
12. Hahne F, LeMeur N, Brinkman RR, et al. flowCore: a Bioconductor package for high throughput flow cytometry. BMC Bioinformatics. 2009;10:106.
13. Malek M, Taghiyar MJ, Chong L, et al. flowDensity: reproducing manual gating of flow cytometry data by automated density-based cell population identification. Bioinformatics. 2015;31(4):606-607.
14. Maurer-Granofszky M, Schumich A, Buldini B, et al. An Extensive Quality Control and Quality Assurance (QC/QA) Program Significantly Improves Inter-Laboratory Concordance Rates of Flow-Cytometric Minimal Residual Disease Assessment in Acute

Lymphoblastic Leukemia: An I-BFM-FLOW-Network Report. Cancers (Basel). 2021;13(23):6148.

Figure legends:

Figure 1: **FlowARC Architecture**
FlowARC, a novel and multi-module deep-learning-based model, follows a cascaded approach for detecting tumor presence in flow cytometry data. **(A)** **Flow cytometry data acquisition and preprocessing pipeline**: Raw flow cytometry samples undergo preprocessing to extract B-cell populations based on forward/side scatter detectors and fluorescent markers. **(B)** **Cell-level module training (Audit stage)**: A small subset of cells is manually annotated by pathologists as normal or tumor cells. These annotations train the cell-level neural network to classify individual cells and assign tumor probability scores. **(C) Probability-based cell reordering (Reorder stage)**: Within each flow cytometry sample, cells are automatically reordered based on their tumor probability scores from the cell-level module, ranking cells from highest to lowest tumor likelihood. This allows sample truncation from hundreds of thousands of cells to the most relevant thousands for efficient processing. **(D)** **Sample-level classification (Classify stage)**: The reordered and truncated flow samples, labeled at the sample level (normal vs. tumor), train the sample-level module to classify entire flow cytometry samples as normal or tumorous. (**E) Summarization of Key Innovations and Clinical Impact of the FlowARC Model.**

Figure 2: **Performance of FlowARC**
Comprehensive performance evaluation of FlowARC for tumor detection in flow cytometry samples. **(A)** Receiver operating characteristic (ROC) curve with 95% confidence interval for FlowARC on patient test cases. The inset shows a magnified view of the high-performance region (AUC = 0.991 ± 0.006). **(B)** Comparison of area under the ROC curve (AUC) between CellCNN and FlowARC models with 95% confidence intervals. **(C)** Tumor burden was calculated by combining information from two complementary parts of the model. For samples with fewer than 7,500 tumor cells, the quantification module was used to estimate tumor content. In contrast, samples exceeding this threshold were analyzed using the cell-level module. Tumor burden quantification showing correlation between FlowARC-predicted and true tumor cell populations. Each point represents a patient case, color-coded by percentage of tumor cells in white blood cells (WBC) as determined by institutional case reports. The dashed line indicates perfect prediction (R^2 = 0.879 ± 0.002, slope = 0.963 ± 0.003, intercept = 0.111 ± 0.016). Cases marked "Neg" represent tumor-negative samples. **(D)** SHAP (SHapley Additive exPlanations) feature importance plots for tumor cell prediction (top) and normal cell prediction (bottom). Features (cell surface markers) are ranked by decreasing importance, with individual SHAP values shown as points colored by raw feature value (blue = low, pink = high).

Figure 3: **Flow Cytometry Visualization of B-ALL Detection by FlowARC Cell-Level Module**
Direct visualization of FlowARC's cell-by-cell tumor probability assignments across representative patient cases. Each row represents a single patient case

displayed across four biaxial plots: CD45 vs CD20, CD34 vs CD38, CD10 vs CD38, and t-SNE dimensionality reduction (components 1 and 2). Individual cells are color-coded by FlowARC-predicted tumor likelihood (purple = high tumor probability, orange = normal cells, as indicated by color bar). **(A)** Correctly classified cases: True negative (top row) shows uniform orange coloring indicating correct identification of normal B-cells; true positive (bottom row) displays distinct purple tumor cell populations clearly separated from normal cells. **(B)** Misclassified cases: False positive (top row) shows scattered likely-tumor cells that do not form tight clusters in any biaxial plots, indicating misclassification of the case as abnormal; false negative (bottom row) demonstrates likely-tumor cells that, although few, form tight clusters in CD34 vs CD38 and t-SNE plots, indicating model misclassification as normal. **(C)** Cases with insufficient B-cell counts (<7,500 cells) excluded from sample-level analysis: normal case (top row) does not show clusters of likely-tumor cells, while the abnormal case (bottom row) displays a distinct purple tumor cluster as expected. Despite being unable to process these cases through the complete FlowARC pipeline due to low B-cell numbers, the cell-level module visualization remains highly informative for identifying tumor populations even in these low cell count samples. These visualizations demonstrate FlowARC's ability to identify phenotypically distinct tumor populations while highlighting edge cases and current limitations.

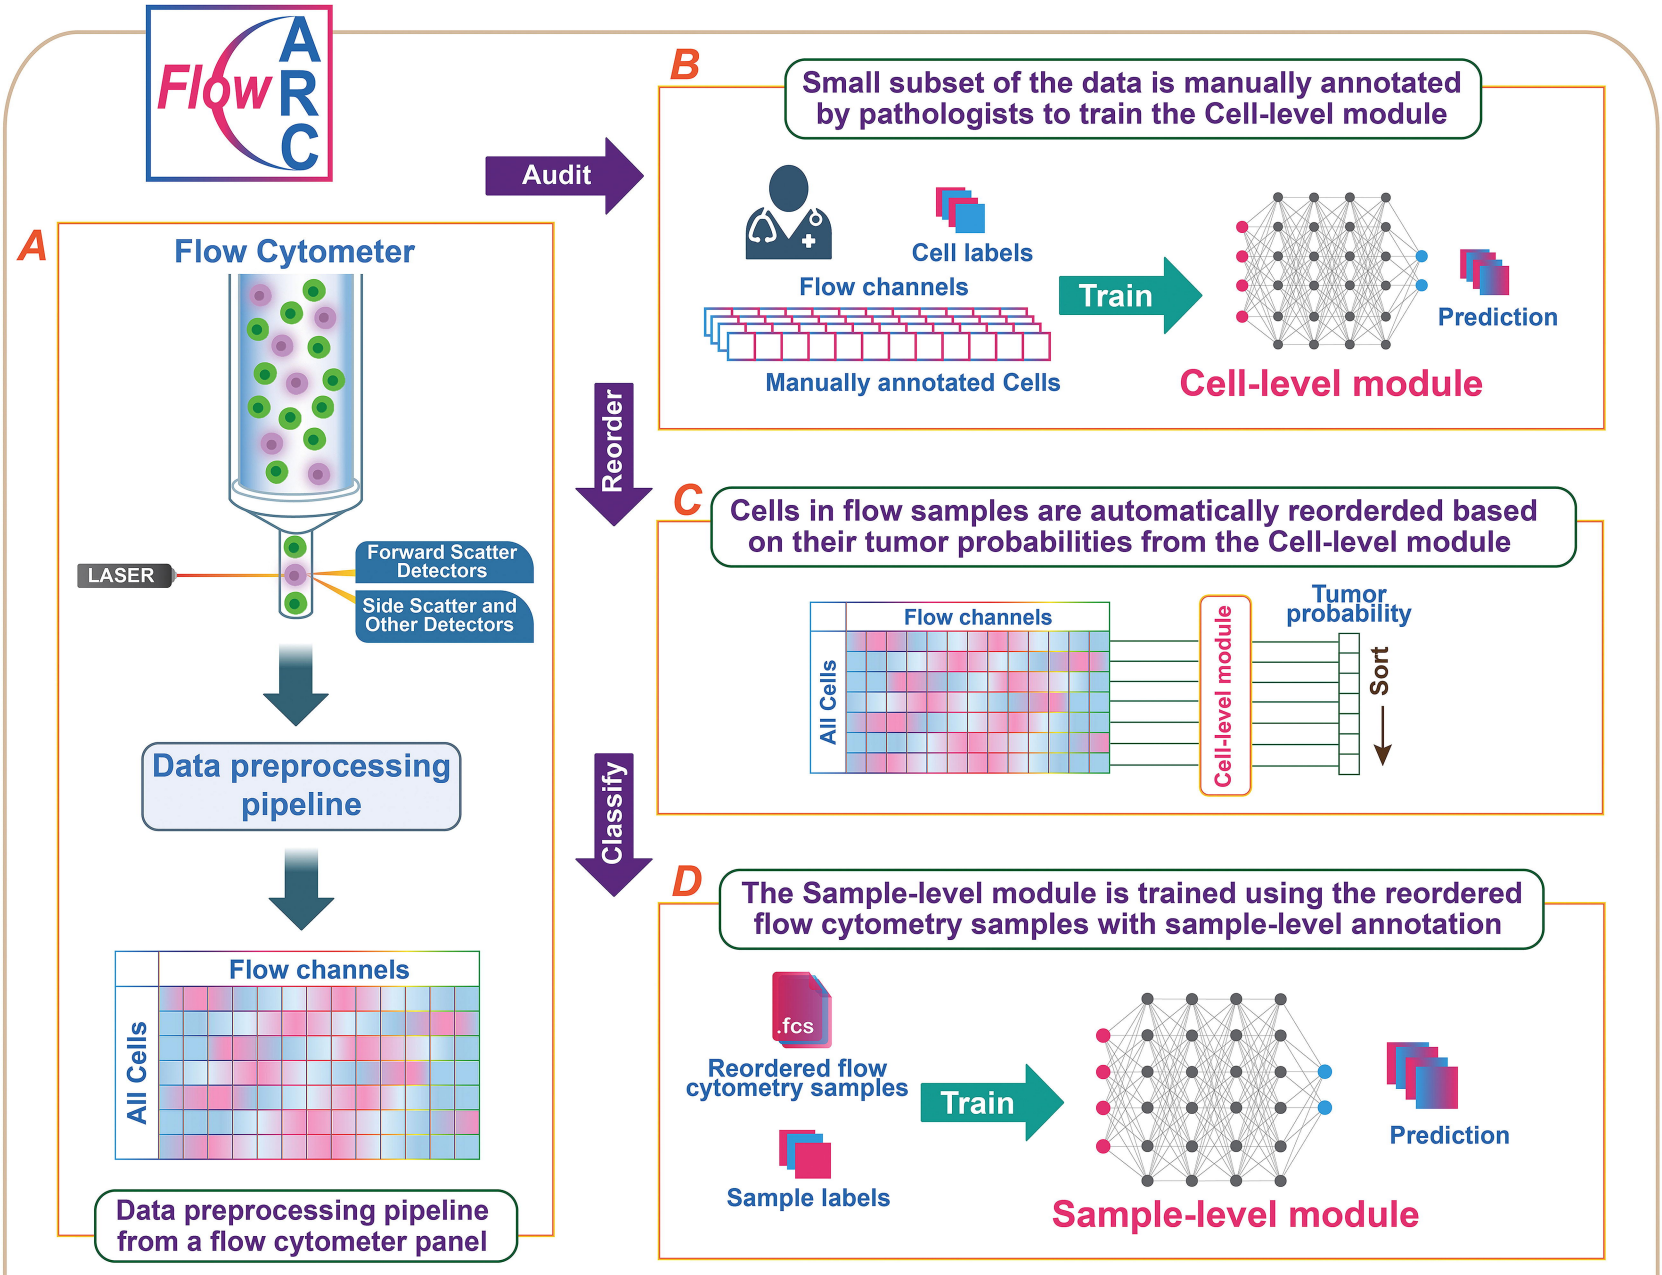

E

Aspect	FlowARC vs. prior approaches	Key impact
Architecture	Three-stage cascade (cell → ranked subset → sample) instead of single-stage model	Greater accuracy; lower noise
Data focus	Selects top 7,500 high-probability abnormal cells, not all events	Better sensitivity for MRD cases
Training	Synthetic MRD samples from real data	Reduces need for extensive manual labeling
Interpretability	SHAP marker attribution and probability mapping	Transparent, biologically consistent results
Clinical utility	Integrated tumor burden estimate; visual review on standard plots	Enables MRD reporting and pathologist verification

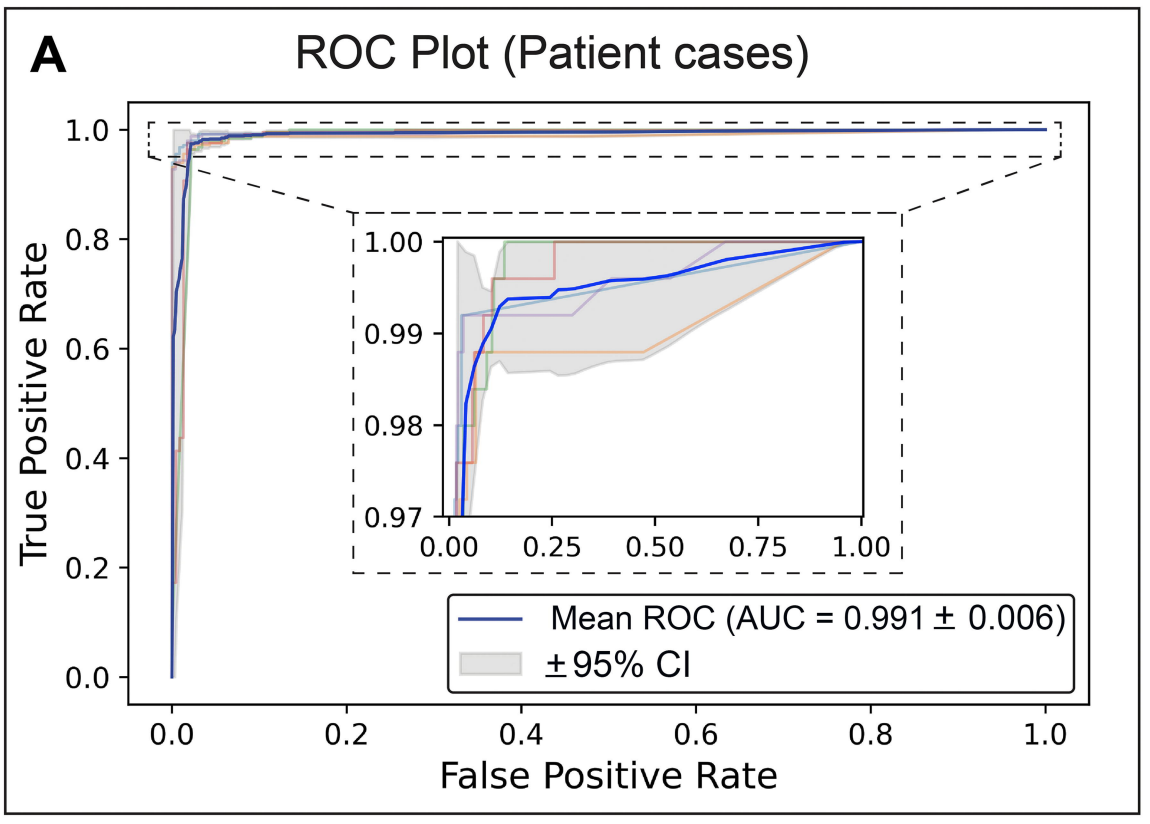
A
ROC Plot (Patient cases)
True Positive Rate
False Positive Rate
Mean ROC (AUC = 0.991 ± 0.006)
±95% CI

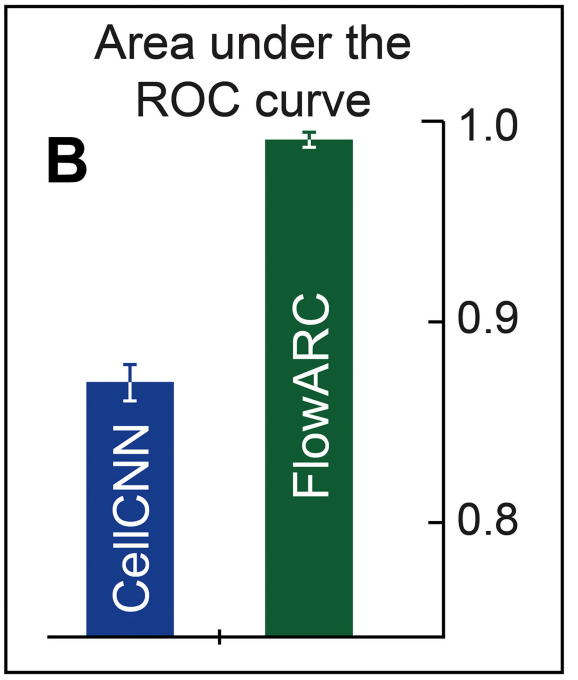
Area under the ROC curve
B
CellCNN
FlowARC

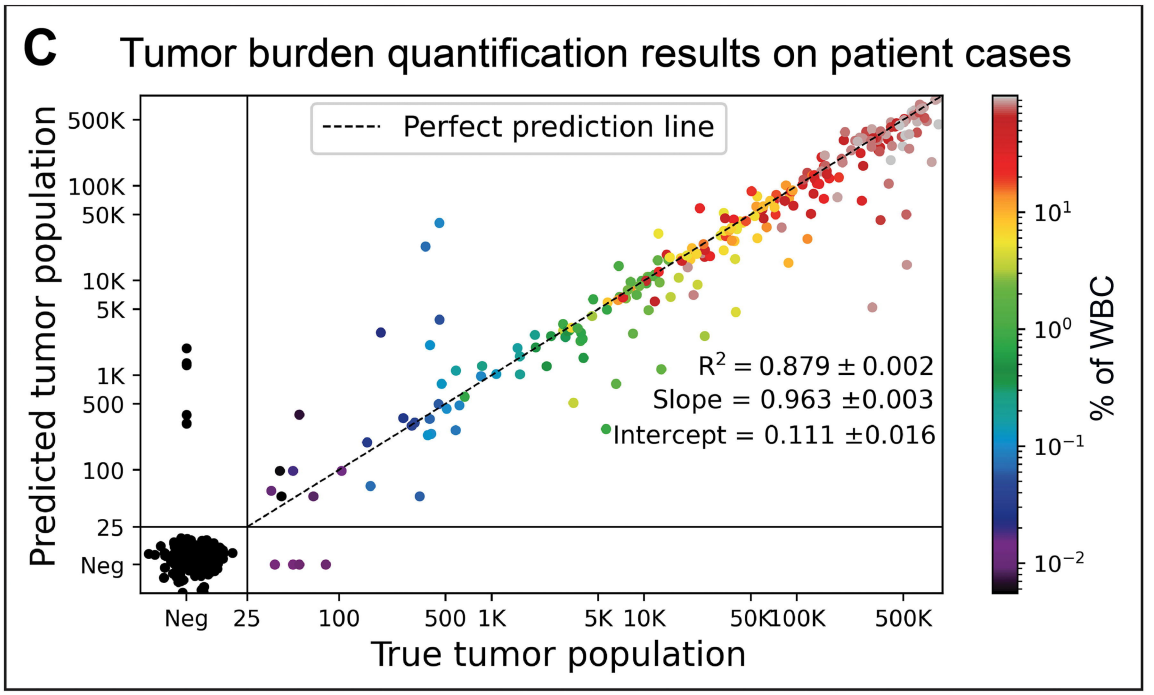
C
Tumor burden quantification results on patient cases
Perfect prediction line
$R^2 = 0.879 \pm 0.002$
Slope = 0.963 ±0.003
Intercept = 0.111 ±0.016
Predicted tumor population
True tumor population
% of WBC

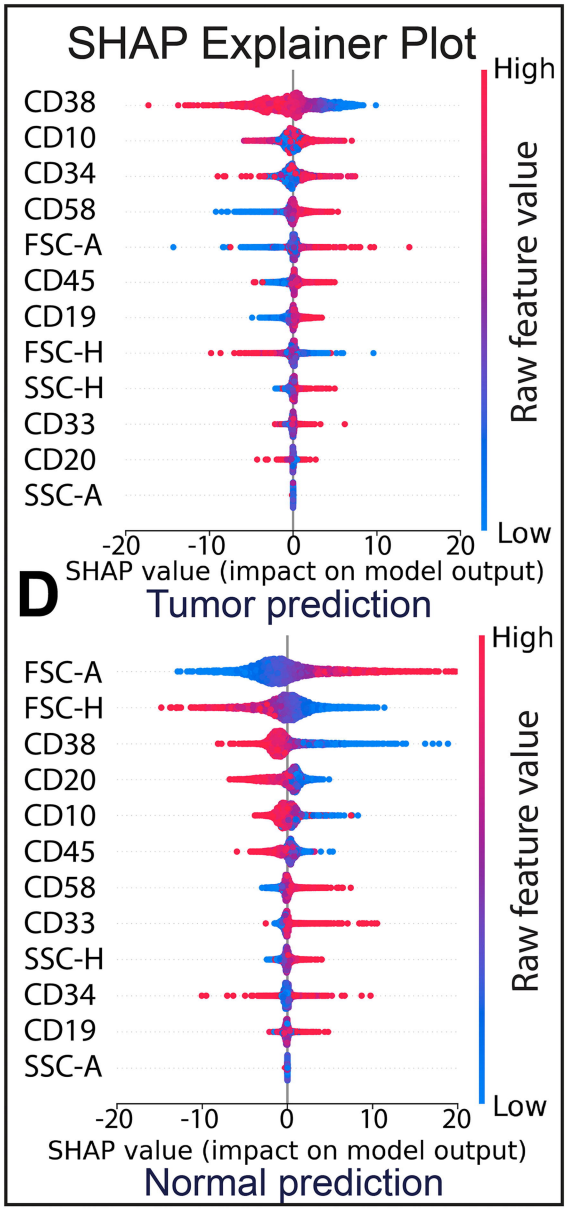
SHAP Explainer Plot
CD38
CD10
CD34
CD58
FSC-A
CD45
CD19
FSC-H
SSC-H
CD33
CD20
SSC-A
High
Low
Raw feature value
SHAP value (impact on model output)
D
Tumor prediction
Normal prediction

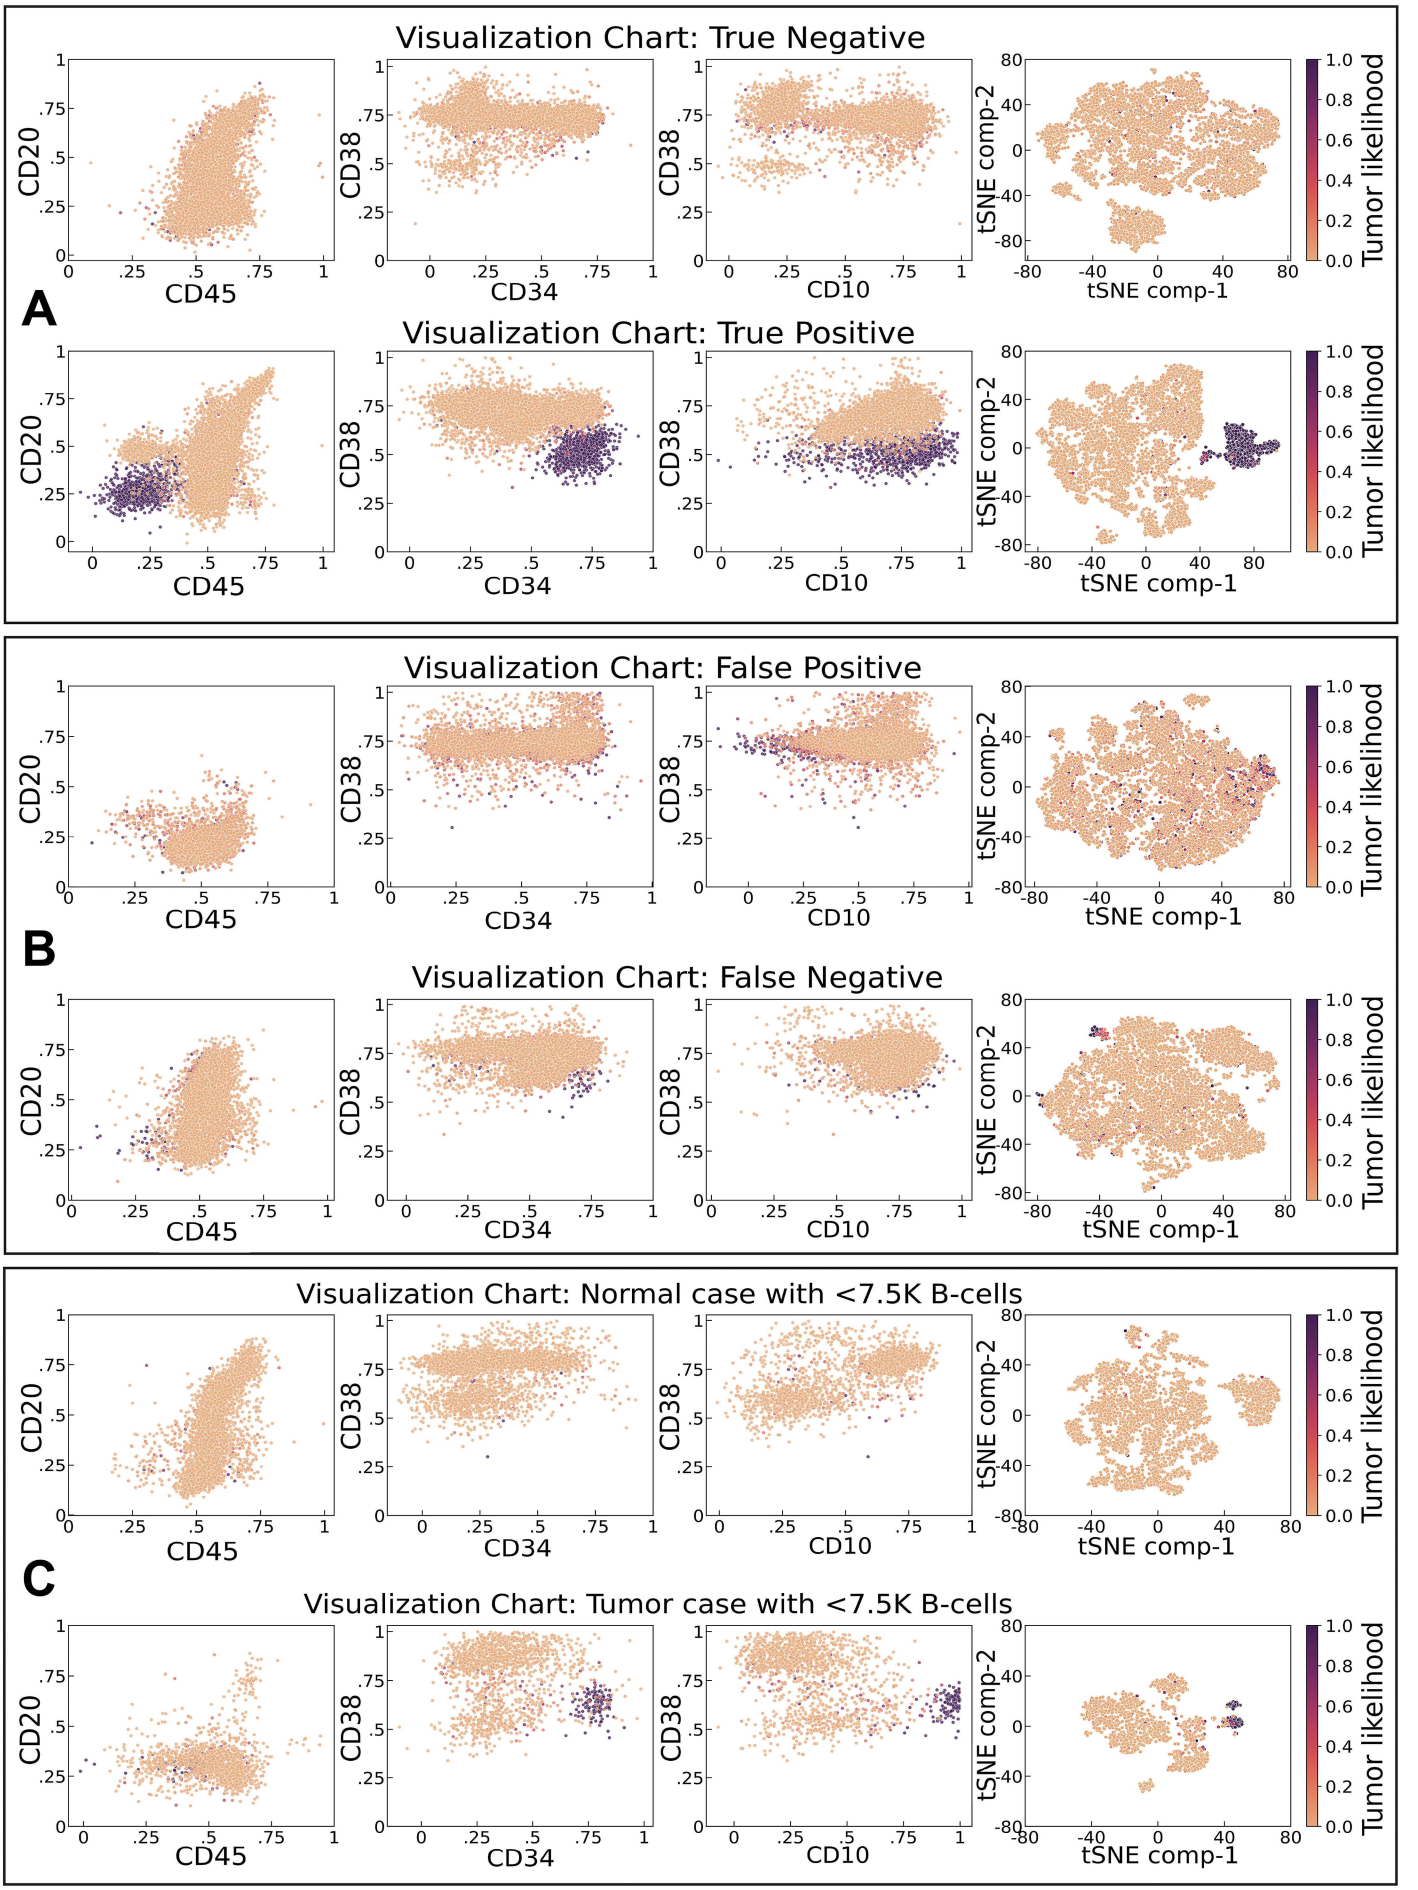
A
Visualization Chart: True Negative
Visualization Chart: True Positive
B
Visualization Chart: False Positive
Visualization Chart: False Negative
C
Visualization Chart: Normal case with <7.5K B-cells
Visualization Chart: Tumor case with <7.5K B-cells
CD20
CD45
CD38
CD34
CD38
CD10
tSNE comp-2
tSNE comp-1
Tumor likelihood

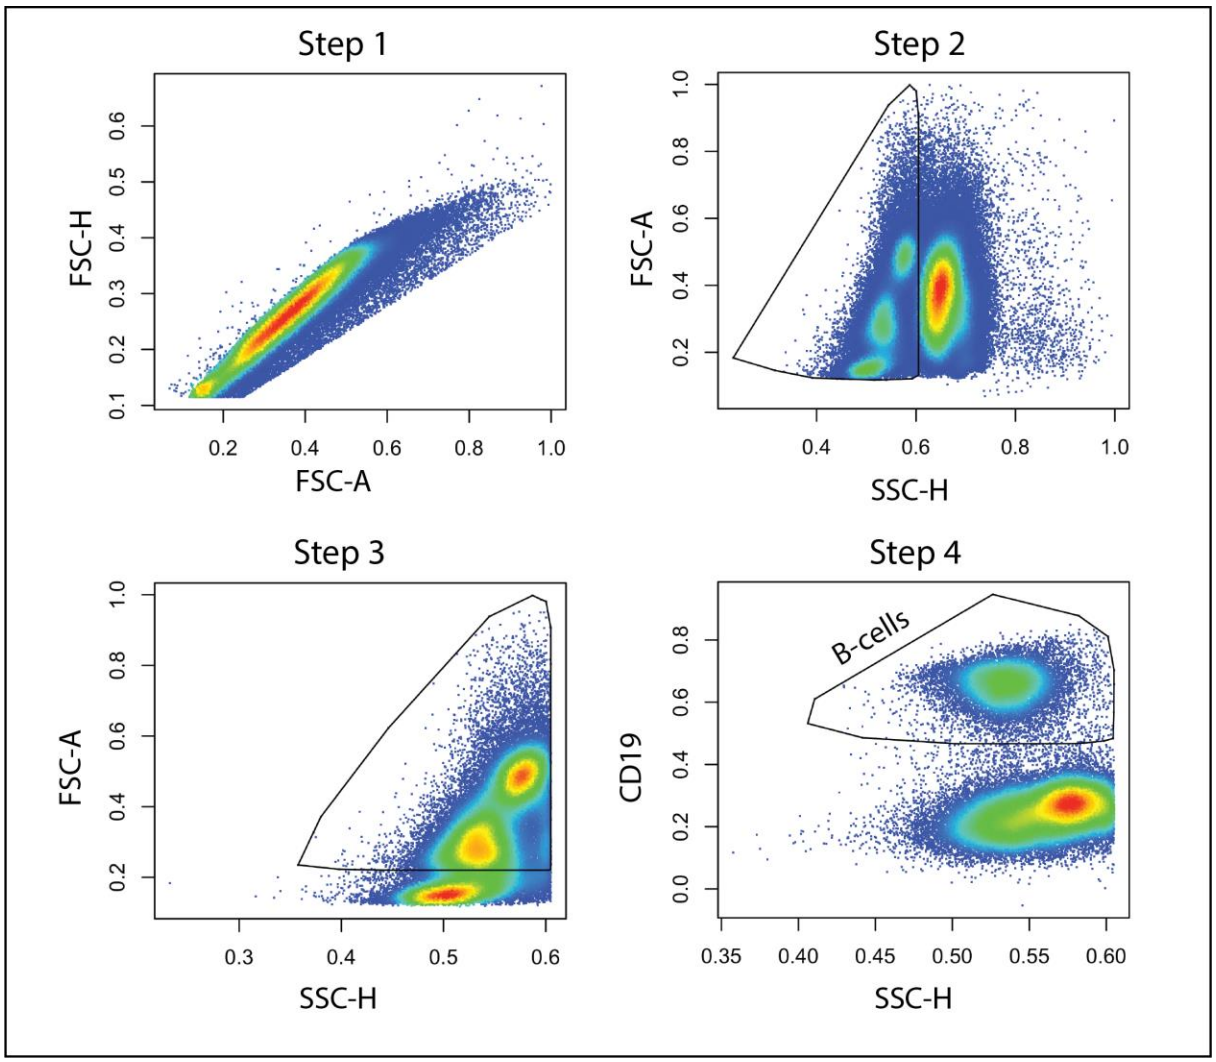

Supplementary Figure 1

Automated B-cell extraction via flowDenstiy

A four-step automated gating strategy for CD19+ B-cell isolation from flow cytometry data using flowDensity. Step 1: Quality control filtering removes margin events and doublets. Margin events are excluded using empirically defined FSC and SSC minimum/maximum thresholds, while doublets are identified and removed based on FSC-A/FSC-H ratio exceeding median plus four standard deviations. Steps 2-3: Sequential gating isolates viable mononuclear cell populations using side scatter parameters (Step 2: SSC-H <0.605; Step 3: SSC-A >0.22). Step 4: CD19+ B-cells are identified from the mononuclear population using density-based clustering with a CD19 expression threshold >0.3. All threshold values were optimized by comparing automated gating results to manual expert annotations across a validation set of 15 randomly selected patient cases, ensuring concordance between automated and manual B-cell populations. This standardized approach enables consistent, reproducible B-cell extraction across large-scale flow cytometry datasets without manual intervention.

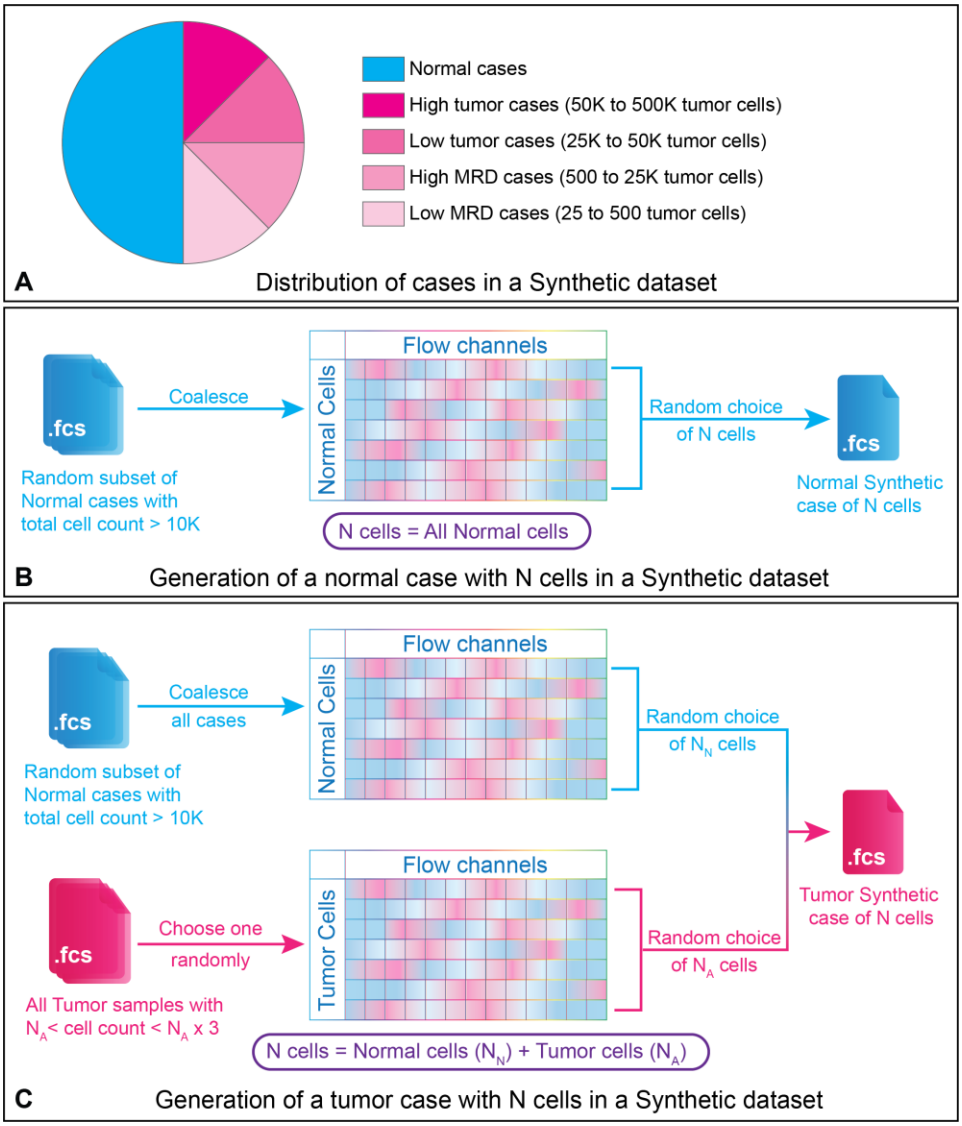

Supplementary Figure 2

Synthetic Dataset Generation Strategy for Sample-Level Model Training

Construction of balanced synthetic flow cytometry datasets for training, validation, and testing of the sample-level module. (A) Distribution of cases in synthetic datasets: Each dataset maintains a 50:50 balance between normal and tumor-positive cases, with tumor cases further stratified into four burden categories: low MRD (25-500 tumor cells), high MRD (500-25,000 tumor cells), low tumor burden (25,000-50,000 tumor cells), and high tumor burden (50,000-500,000 tumor cells). A minimum of 25 tumor cells was required for all positive cases. (B) Normal case generation workflow: Normal cases are synthesized by: (1) randomly selecting multiple normal patient samples with combined cell count >10,000, (2) coalescing these samples, and (3) randomly sampling N cells (ranging from 10,000-50,000) to create individual synthetic normal cases. This approach captures the natural variability in cell counts and phenotypes observed in clinical samples. (C) Tumor case generation workflow: Synthetic tumor cases combine normal and tumor cell populations by: (1) generating a normal B-cell background (NN cells) using the process from panel B, (2) selecting a tumor sample containing NA < cell count < 3× NA tumor cells (where NA is the desired tumor cell count), and (3) randomly sampling NA tumor cells to combine with the normal population, yielding cases with N = NN + NA total cells. This strategy generated 15,000 training, 6,000 validation, and 9,000 testing synthetic samples from annotated patient data, ensuring robust model development across the full spectrum of tumor burdens encountered clinically, from MRD to overt disease.

Module	Input	Output	Key Implementation & Hyperparameters
Cell-level Module	(1 × 12) Each cell with 12 cell markers	(2 × 1) probability for normal and tumor class	Random weight initialization; Loss function: cross-entropy + Orthogonal Projection Loss (gamma=0.5, alpha=1); Trained for 100 epochs, batch size =800, Stochastic gradient descent optimizer with momentum=0.9, weight decay=1e-3; learning rate =0.01 decaying 10% at epochs (5,10,20,30,50,70); class imbalance handled via imbalanced sampler; hyperparams selected by AUROC metric
Sample-level Module	(3 × 7,500 × 14) Matrix of top 7,500 cells after reordering and added cell-level probability for normal and tumor + FFT real & imaginary of that matrix	(2 × 1) probability for normal vs tumor class	Random weight initialization; Loss function: cross-entropy + Orthogonal Projection Loss (gamma=0.5, alpha=1); Trained for 100 epochs, batch size =10, Stochastic gradient descent optimizer with momentum=0.9, weight decay=1e-3; learning rate =0.01 decaying 10% at epochs (5,10,20,30,50,70); FFT added for noise robustness; hyperparams selected by AUROC metric
Quantification Module	Same as sample-level module with the value at [0,0,0] replaced with the pre-truncation cell count for that sample	(1 × 1) value between 0 and 1 multiplied by input size to get absolute tumor count	Same as sample-level module but with Huber loss (delta=0.007); trained only on synthetic B-ALL positive samples; hyperparams selected by R^2 metric

Supplementary Table 1

FlowARC Module Architecture and Training Parameters

Summary of neural network architectures and hyperparameters for the three FlowARC modules.Cell-level Module:Processes individual cells (1×12 feature vector) using 1D ResNet-18 (15) to output binary tumor/normal probabilities (2×1). Trained with combined cross-entropy and Orthogonal Projection Loss (γ=0.5, α=1) to handle class imbalance via imbalanced sampling. Sample-level Module:Analyzes top 7,500 reordered cells as a 3×7,500×14 tensor (cells × markers × probability scores + FFT transformations) using 2D ResNet-101 to classify entire samples as normal or tumor. FFT components enhance noise robustness.Quantification Module:Employs the same architecture as the sample-level module but replaces probability scores with pre-truncation cell counts and uses Huber loss (δ=0.007) for regression to estimate absolute tumor cell numbers. All modules utilize stochastic gradient descent with momentum (0.9), weight decay (10^{-3}), and learning rate scheduling (initial 0.01, 10% decay at epochs 5, 10, 20, 30, 50, 70). Hyperparameters were optimized using AUROC for classification modules and R^2 for the quantification module on validation datasets. The quantification module was trained exclusively on synthetic B-ALL positive samples to ensure accurate tumor burden estimation.