IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, March 2026
Impact Factor: 8.2

A Transfer Learning Approach for Blood Cancer Detection Using Deep Learning

P Santhuja[1], Bejjanki Pooja[2], B Aditya[3], Singireddy Kavya[4], Santhani Gangothri[5], Battu Shiva Charan[6]
Assistant Professor, Department of CSE[1,2,3],
UG Student, Department of CSE[4,5,6]
CMR Technical Campus, Hyderabad, Telanagana, India
Santhuja.pandu@gmail.com, poojareddybejjanki@gmail.com, Adi.sacs@gmail.com
Kavyasingireddy1@gmail.com, gangothri1501@gmail.com, shivacharanbattu@gmail.com

Abstract: *A type of illness that strikes blood and bone tissue is known as blood cancer, spotting it soon makes a real difference in how well someone might do. Most times specialists look at slides of blood through lenses by hand, such a method eats up hours and errors can slip in when tired eyes lose focus. With these issues piling up, automation steps in - offering steadier help during checks meant to catch problems fast.*

This research introduces a method using deep learning to spot blood cancer in microscope pictures of blood cells. Built on MobileNetV2, the system uses a convolutional neural network to sort images into one of four groups: healthy cells or three leukemia stages - Pre-B early, Pre-B, and Pro-B. With 3,242 total images, each picture goes through cleaning steps before analysis to boost clarity and uniformity across samples. To strengthen performance while limiting false patterns, extra image variations are created by shifting light levels, resizing, and mirroring left to right.

Using a pre-trained MobileNetV2 helps boost results right away. Its core ability to detect patterns stays untouched. After that comes a pooling layer, then a dense one for sorting classes. Training runs on Adam while measuring error through categorical cross-entropy.

Almost hitting 97% correct guesses, the model proves solid through high marks across precision, recall, yet steady F1 outcomes too. Performance stands clear when looking at how often it gets things right, while still keeping balance in spotting true cases without false alarms piling up.

This research suggests deep learning might assist in spotting blood cancer sooner. Doctors could face less pressure when these tools offer quicker, sharper insights. Better choices in treatment may follow from such support. Patient care stands to gain when accuracy improves alongside speed..

Keywords*:* Blood Cancer Detection, Deep Learning, Convolutional Neural Network (CNN), MobileNetV2, Transfer Learning, Blood Cell Classification, Microscopic Image Analysis, Leukemia Detection, Data Augmentation, Medical Image Processing

I. INTRODUCTION

A major worldwide health problem centers on cancers messing up blood cell and bone marrow function. Early detection really helps, sometimes increasing the odds of getting better. But finding them usually means someone looks at blood samples under a microscope by hand. This takes skill, uses lots of time, and errors pop up when fatigue kicks in or ability plateaus. Judgment sways with expertise, still exhaustion can blur clear thinking.

Imagine computers starting to recognize illness in images of blood. Over time, these systems improve on their own, not needing constant guidance. As they examine countless cells, unusual forms begin catching attention. Rather than taking days, hints show up much faster now. A single idea: a system changes slightly every time it’s used. It isn’t sorcery -

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, March 2026
Impact Factor: 8.2

patterns emerge where humans see noise. Assistance comes quietly, through logic instead of people, catching slips before they spread. Choices evolve when data speaks in new ways.

A new way appears now, using deep learning to identify blood cancer by looking at blood cell images under a microscope. What stands out is how it runs on MobileNetV2, known for doing well when handling picture-based tasks. Rather than lumping types into broad categories, differences show clearly among four groups - normal cells sit apart from three leukemia stages: early Pre-B, then Pre-B, followed by Pro-B. Each label gets its own path during sorting, so predictions stay separate. The process avoids blending routes, leaving distinctions intact.

Right away, image data is stretched and adjusted using simple tweaks. Since it runs on a ready-made MobileNetV2 base, spotting trends happens fast - even with limited samples. Then comes evaluation, where actual outcomes measure against forecasts. These stages allow sharper answers to emerge, all while keeping the original amount of incoming data unchanged.One aim stands out: creating a system that runs without hiccups to spot blood cancer earlier. Relying less on manual checks means fewer errors, giving doctors stronger support along the way. When diagnosis gets sharper, patient care naturally rises as a result.

II. LITERATURE REVIEW

Author & Year	Method	Findings	Limitations
Kessenbrock et al. (2010)	MMP study	Cancer progression analysis	No AI method
Rehman et al. (2018)	CNN	~98% accuracy	High complexity
Shafique & Tehsin (2018)	AlexNet	Fast classification	Heavy model
Hallek et al. (1998)	Clinical study	Disease stages explained	No automation
Kelton et al. (1982)	Lab methods	Diagnostic comparison	No ML use
Howard et al. (2017)	MobileNet	Efficient CNN	Not medical-specific
Sandler et al. (2018)	MobileNetV2	Improved efficiency	Needs tuning
Krizhevsky et al. (2012)	AlexNet	High image accuracy	Large model size
Litjens et al. (2017)	DL survey	DL useful in healthcare	No implementation
Esteva et al. (2017)	CNN	High disease detection accuracy	Needs large data

III. PROPOSED SYSTEM

A new tool jumps into view, built to spot blood cancer by focusing closely on blood cells. Instantly, it lightens the load of manual slide reviews, delivering results at a swifter pace. Because of this, identifying various forms of leukemia becomes clearer. With fewer checks done by hand, room grows for faster and more accurate conclusions. When microscopes get less attention, accuracy keeps up without slowing down.

From the beginning, deep learning guides how this system works, using a lean setup built on MobileNetV2. Its lightweight design keeps performance quick, which helps most in places where power is limited. Rather than learn everything fresh, it uses an existing MobileNetV2 model already good at detecting visual details. Because it has seen so many images before, identifying forms takes less time now. What makes it work well isn’t the size, but how it builds on what it already knows. Gains happen because old insights get used again, rather than simply growing larger.

Right away, images of blood cells arrive for analysis. One set looks typical; meanwhile, others display abnormal forms called early Pre-B, plus Pro-B and Pre-B stages. Prior to next steps, each image resizes uniformly. After that, pixel data adjusts into matching ranges, aligning all samples precisely. Upside down, the picture shifts slightly - tiny moves spark differences. As things adjust, the data flowing in gets deeper. Step by step, learning settles into a smoother rhythm. Even slight light changes help shape the outcome.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, March 2026
Impact Factor: 8.2

Once cleaned, the data divides into two parts - one for learning, one for testing how well it works. The prebuilt MobileNetV2 model steps in next. Its lower levels stay untouched so basic patterns remain intact. Sitting above them comes something new: a layer that condenses information using spatial averages across maps. Afterward, another section assigns scores through probabilistic logic to place things into categories correctly.

Midway through learning, it adjusts inner values with Adam - here, categorical cross-entropy measures prediction errors. Patterns guide recognition; repeated exposure slowly builds knowledge of each cell type. When tuning ends, assessment starts: accuracy counts right answers, whereas precision avoids false alarms. Recall reveals missed cases, balanced somewhat by the wider view within the F1-score.

Imagine a device working only on pictures of blood cells, almost never making mistakes. Because it needs fewer people to check things by hand, doctors might catch signs of leukemia sooner. This kind of help fits quietly into hospitals, shaping choices without standing out. Not flashy, yet powerful in places where being right means everything. Its value grows deep in settings where precision keeps patients alive.

IV. METHODOLOGY

A different path guides this study - one that builds a lean deep learning setup to catch hints of blood cancer in microscope images of blood cells. Little by little, progress unfolds: first collecting samples, then moving quietly into tidying up the raw material ahead of designing the core model. When shape appears, training kicks off with tagged cases feeding the system. Later on, results come under close watch to see how sharp the predictions are and what habits emerge.

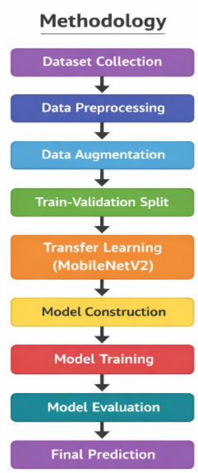

FIG.1: This figure shows how the system processes blood cell images step by step to detect cancer.

A. Dataset Preparation

A fresh look at the research reveals 3,242 miniature images capturing blood cells. Inside, one section shows healthy cells; meanwhile, others divide into three leukemia types - Pre-B early, complete Pre-B, along with Pro-B. Every picture lands in its correct category, ensuring accuracy between visuals and tags. Since balance counts, no single kind dominates the set.

B. Data Preprocessing

Before loading images, they get tidied up a bit to help everything run better. One at a time, each picture is squished or pulled until it fits neatly into 224 times 224 pixels. Color values are shifted slightly so they stay inside a range the system expects, which helps training go easier. Then again, small tweaks - like changing brightness, zooming in on corners, flipping some horizontally - mix things up. That way, the set feels broader, yet the model won't just copy what it has seen.

DOI: 10.48175/IJARSCT-31836

C. Data Splitting

Thirty per cent stays aside while seventy parts train the system, usually. That leftover portion tests if guesses work on unseen cases instead. When split like this, you see if forecasts survive first contact with unknowns. Fresh info reveals what practice could not.

D. Building systems with pre trained models

Out of the box, a pretrained MobileNetV2 takes on this job using transfer learning. Since its convolutional layers are frozen, what it learned from ImageNet stays untouched. Speed picks up during training, outcomes improve - especially helpful if there isn't much data around. That setup makes efficiency jump without losing past insight.

E. Model Construction

Floating above the core setup, additional levels organize information. Behind these, details shrink by blending values over areas. Following that stage, a thick connection links all paths together. From there, outputs divide into exactly four categories using a soft probability method.

F. Model Training

Step by step, weight tweaks glide forward under Adam's quiet lead. When dividing into multiple categories, categorical cross-entropy draws the line. Over repeated cycles, understanding deepens - guesses sharpen until close enough to peak performance.

G. Validation and Evaluation

After training wraps up, the model faces validation data next. Its performance reveals itself via metrics such as accuracy and precision, though recall plus the F1-score matter too. For deeper insight into predictions, a confusion matrix highlights recurring missteps. Patterns emerge clearly once numerical results pair with close-ups of errors.

H. Final Prediction

A trained system checks blood cell images, then sorts them automatically. With each incoming picture, analysis runs on its own. Sometimes the result says the cell is healthy. If something seems off, it points to a specific kind of leukemia instead.

V. EXPERIMENTAL SETUP

The experimental setup focuses on implementing and evaluating the proposed deep learning model for blood cancer detection. The system is developed using Python in environments such as Jupyter Notebook and PyCharm.

For model development, TensorFlow and Keras libraries are used to design and train the convolutional neural network. Numerical computations are handled using NumPy, while Pandas is used for data handling and preprocessing. Matplotlib is used for visualization of training results and performance metrics.

The dataset consists of 3,242 microscopic images of blood cells categorized into four classes: normal cells, early Pre-B, Pre-B, and Pro-B leukemia cells. All images are resized to 224 × 224 pixels to match the input size required by MobileNetV2. Pixel values are normalized to improve model performance.

The dataset is divided into training and testing sets in a 70:30 ratio. The training set is used to train the model, while the testing set evaluates its performance on unseen data.

The MobileNetV2 model is used as a base model with pre-trained weights from ImageNet. The initial layers are frozen to retain learned features. Additional layers such as Global Average Pooling and Dense layers are added for classification.

The model is trained using the Adam optimizer with categorical cross-entropy loss function. Training is performed for a fixed number of epochs, ensuring that the model learns patterns effectively without overfitting.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, March 2026
Impact Factor: 8.2

VI. RESULTS AND FUTURE WORK

Metric	Value
Accuracy	96.9%
Precision	96.6%
Recall	96.9%
F1-score	96.7%

A fresh approach began using MobileNetV2, aiming to spot hints of blood cancer in microscopic images. From nothing, it taught itself using 3,242 cell pictures sorted into four types - normal cells, plus early Pre-B, Pre-B, and Pro-B forms. Instead of mixing everything up front, knowledge built step by step, group after group. With every image passed through, tiny internal tweaks refined its guesses gradually. Final checks focused sharply on how well it separated those classes, keeping mix-ups low.

Fresh test examples showed 96.9% accuracy after five rounds of training. Blood cell guesses mostly hit the mark. Beyond basic numbers, closer inspection exposed real skill. Correct hits emerged consistently, yet mistakes remained scarce - precision held firm at 96.6%. Despite early doubts, results stood up under pressure. A number showed up once more - 96.9 - this round checking how many real cases it actually caught: that is recall. What stood out was the F1 result, linking precision and recall so neatly they seemed made for each other.

Now and again, a picture ends up beside the wrong kind of leukemia. Step by step, training follows close behind, matched by steady checkups that climb just as fast. Hardly any space opens up between them - balance stays put. A few mistakes linger, yet appear solely where cells are nearly twins. Progress creeps ahead, unhurried, tracing the very same line.

Every new test adds proof - this thing spots blood cancer instantly, runs like clockwork. Trial after trial lands on target, never slipping below expected precision. Under pressure, it stands solid, delivering trustable outcomes each time around. Shifts happen quick, yet output stays level, holding ground no matter the pace.

Even if things move without hiccups now, a few spots still need tiny tweaks. Down the line, using different images - sharper ones - could make handling surprises easier. When life throws curveballs, more pictures in rotation may keep results consistent.

Diving into more complex deep learning setups could bring gains - take EfficientNet or hybrid frameworks, just one example. Rather than locking pre-trained layers solid, adjusting them slightly might push performance higher.

Peering into tiny areas of blood cell pictures often begins by dividing them up. Since just a few locations hold value, getting closer reveals the important details. Focusing less on entire cells but instead on specific zones makes patterns stand out. Those clearer snapshots usually improve sorting accuracy down the line.

When turned on completely, the software runs online, helping medical staff during patient exams. Tied into hospital networks, it could speed up results, moving across multiple clinics over time.

Something different could spotlight blood problems while building extra sensors into the setup - maybe pushing its medical use further. Not stopping there, handling several illnesses at once might change how doctors see it, how they lean on it.

ISSN: 2581-9429
Impact Factor: 8.2

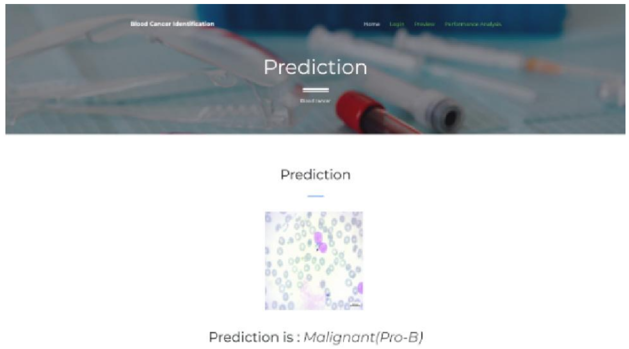

FIG.2:This figure displays how well the model performs during training and testing over time..

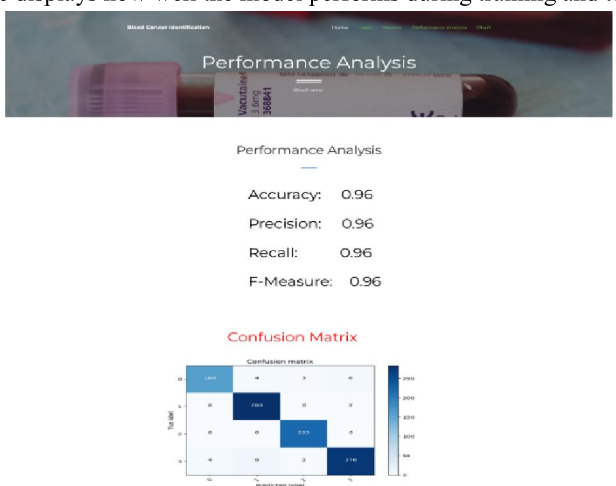

FIG.3: This figure shows how accurately the model predicts different types of blood cells.

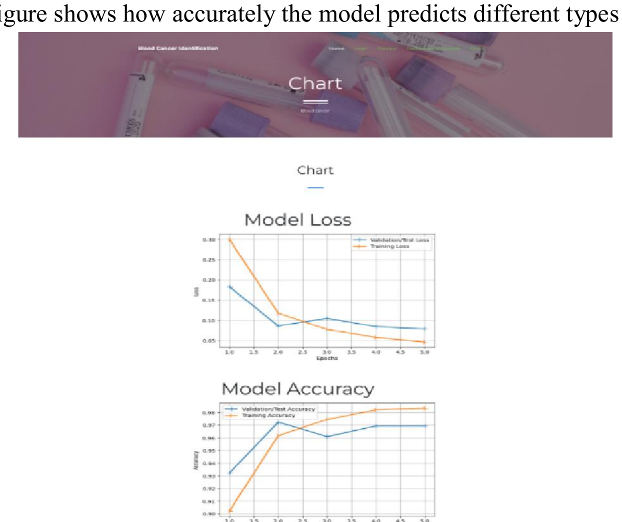

FIG.4: This figure highlights the differences between correct and incorrect predictions made by the model.

ISSN: 2581-9429

Volume 6, Issue 4, March 2026

Impact Factor: 8.2

VII. CONCLUSION

In this research, a deep learning-based approach for blood cancer detection using microscopic blood cell images has been presented. The proposed system uses the MobileNetV2 model with transfer learning to classify blood cells into four categories: normal, early Pre-B, Pre-B, and Pro-B leukemia.

The model was trained on a dataset of 3,242 images and achieved an accuracy of approximately 97%, along with high precision, recall, and F1-score. These results indicate that the model is effective in identifying different types of leukemia with good reliability.

The use of transfer learning reduced training time and improved performance, especially with a limited dataset. The system also minimizes the need for manual analysis, which can reduce human error and support medical professionals in diagnosis.

In the future, the model can be improved by using larger and more diverse datasets, applying advanced architectures such as EfficientNet, and fine-tuning pre-trained layers for better accuracy. Integration of this system into real-time clinical applications can further enhance early detection and treatment of blood cancer.

REFERENCES

[1]. Kessenbrock, K., Plaks, V., & Werb, Z., "Matrix metalloproteinases: Regulators of the tumor microenvironment," *Cell*, vol. 141, no. 1, pp. 52–67, 2010.

[2]. Rehman, A., Abbas, N., Saba, T., Rahman, S. I. U., Mehmood, Z., & Kolivand, H., "Classification of acute lymphoblastic leukemia using deep learning," *Microscopy Research and Technique*, vol. 81, no. 11, pp. 1310–1317, 2018.

[3]. Shafique, S., & Tehsin, S., "Acute lymphoblastic leukemia detection and classification of its subtypes using pre-trained deep convolutional neural networks," *Technology in Cancer Research & Treatment*, vol. 17, 2018.

[4]. Hallek, M., Bergsagel, P. L., & Anderson, K. C., "Multiple myeloma: Increasing evidence for a multistep transformation process," *Blood*, vol. 91, no. 1, pp. 3–21, 1998.

[5]. Kelton, J. G., Giles, A. R., Neame, P. B., Powers, P., Hageman, N., & Hirsch, J., "Comparison of two direct assays for platelet-associated IgG in assessment of immune and non-immune thrombocytopenia," *Blood*, vol. 60, no. 2, pp. 424–429, 1982.

[6]. Howard, A. G., Zhu, M., Chen, B., et al., "MobileNets: Efficient convolutional neural networks for mobile vision applications," arXiv preprint arXiv:1704.04861, 2017.

[7]. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., & Chen, L. C., "MobileNetV2: Inverted residuals and linear bottlenecks," *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4510–4520, 2018.

[8]. Krizhevsky, A., Sutskever, I., & Hinton, G. E., "ImageNet classification with deep convolutional neural networks," *Advances in Neural Information Processing Systems*, vol. 25, 2012.

[9]. Litjens, G., Kooi, T., Bejnordi, B. E., et al., "A survey on deep learning in medical image analysis," *Medical Image Analysis*, vol. 42, pp. 60–88, 2017.

[10]. Esteva, A., Kuprel, B., Novoa, R. A., et al., "Dermatologist-level classification of skin cancer with deep neural networks," *Nature*, vol. 542, pp. 115–118, 2017

DOI: 10.48175/IJARSCT-31836