



OPEN

OVision A raspberry Pi powered portable low cost medical device framework for cancer diagnosis

Samaira Mehta^{1,2}✉

Cancer remains a major global health challenge, with significant disparities in access to advanced diagnostic and prognostic technologies, especially in resource-constrained settings. Existing medical treatments and devices for cancer diagnosis are often prohibitively expensive, limiting their reach and impact. Pathologists' scarcity exacerbates cancer diagnosis accuracy, elevating mortality risks. To address these critical issues, this study presents OVision - a low cost, deep learning-powered framework developed to assist in histopathological diagnosis. The key objective is to leverage the portable, low-power computing Raspberry Pi. By designing standalone devices that eliminate the need for internet connectivity and high-end infrastructure, we can dramatically reduce costs while maintaining accuracy. As a proof of concept, the study demonstrated the viability of this framework through a compact, self-contained device capable of accurately detecting ovarian cancer subtypes with 95% accuracy, on par with traditional methods, while costing a small fraction of the price. This portable, off-grid solution has immense potential to improve access to precision cancer diagnostics, especially in underserved regions of the world that lack the resources to deploy expensive, infrastructure-heavy medical technologies. In addition, by classifying each tile, the tool can provide percentages of each histologic subtype detected within the slide. This capability enhances the diagnostic precision, offering a detailed overview of the heterogeneity within each tissue sample, helps in understanding the complexity of histologic subtypes and tailoring personalized treatment plans. In conclusion, this work proposes a transformative model for developing affordable, accessible medical devices that can bring advanced healthcare benefits to all, laying the foundation for a more equitable, inclusive future of precision medicine.

Keywords Cancer, Diagnostics, Deep learning, Histopathology, Raspberry Pi, Precision medicine, Ovarian cancer

Cancer remains a one of the biggest global health challenges, claiming an estimated 10 million lives in 2020 alone¹. Despite significant advancements in diagnostics and treatment, access to these vital resources remains severely constrained^{2–4} in many regions, exacerbating the burden of the disease. Compounding this issue is the reliance on pathologists whose expertise is pivotal in accurate cancer diagnosis; however, a shortage⁵ of these specialists and slow turnaround time of slide evaluation can delay treatment and/or negatively impact patient well-being.

Recognizing these challenges, research teams and commercial entities are leveraging AI to develop tools aimed at assisting pathologists⁶. While these innovations are promising, the high cost associated with such medical devices presents a barrier to widespread deployment, particularly in resource-limited settings.

One malignancy that embodies these challenges is ovarian cancer. With a 5-year survival rate⁷ of just 49.4%, ovarian cancer exemplifies the pressing need for accurate and timely detection to improve diagnosis. Ovarian cancer is characterized by significant histological heterogeneity, which poses challenges for accurate diagnosis and treatment. This heterogeneity is not only observed between patients but also within individual tumors. The complex nature of ovarian cancer histology often requires pathologists to review multiple slides from different regions of a tumor to make an accurate diagnosis. Given this inherent complexity, the classification of ovarian cancer subtypes from histopathological images has seen significant progress due to advancements in machine learning and deep learning techniques. Initial efforts utilized traditional image processing and classical machine learning methods for cancer classification. For instance, studies applying texture analysis and support vector

¹Archbishop Mitty High School, San Jose, CA, USA. ²Machine Learning (High School Lab Member) at OrsulicLab, David Geffen School of Medicine at the University of California, Los Angeles (UCLA), Los Angeles, USA. ✉email: samaira@coderbunnyz.com

machines (SVM) in various cancer types, such as breast cancer, achieved promising accuracies and demonstrated the potential of these methods⁸. These early works laid the groundwork for more sophisticated approaches in ovarian cancer classification.

Literature review

With the advent of deep learning, particularly convolutional neural networks (CNNs), the accuracy and robustness of classification systems have improved markedly. Hussain and Jabbar⁹ developed a CNN-based model specifically for predicting and diagnosing ovarian cancer in pre- and post-menopausal women. Their model achieved a remarkable accuracy of 94% in distinguishing between cancerous and healthy cells. Komura and Ishikawa¹⁰ also explored deep learning methodologies for digital pathology, emphasizing their effectiveness in ovarian cancer diagnosis and subclassification, with reported classification accuracies of up to 93%. Recent approaches have focused on integrating multi-scale features to better capture the complex tissue architecture in histopathological images. Yildirim et al.¹¹ proposed a U-Net model with attention module for gland segmentation in H&E histopathological images, which can be applied to ovarian cancer classification. Given the limited availability of annotated histopathological images, transfer learning and data augmentation techniques have become crucial. Janowczyk and Madabhushi¹² utilized transfer learning to enhance the detection of cancerous regions in histopathology slides, achieving an accuracy of 88%.

While these studies demonstrate significant progress in applying deep learning to ovarian cancer classification, they primarily rely on high-performance computing systems and cloud-based solutions. The proposed method in this study addresses a critical gap by adapting these advanced techniques to run on low-cost, portable hardware. This approach not only maintains high accuracy but also significantly improves accessibility, particularly in resource-constrained settings. The use of a Raspberry Pi-based system represents a novel and optimal solution to the problem of limited access to advanced diagnostic tools in many regions worldwide.

Problem statement and hypothesis

Current standard diagnostic techniques for ovarian cancer, including surgical biopsies followed by histopathological analysis, as well as blood tests like Cancer Antigen 125 (CA125) and other molecular tests, require advanced medical infrastructure and specialized expertise. This high level of resource requirement makes these diagnostic methods largely inaccessible in many lower-resource regions. As a result, these areas often experience higher mortality rates from ovarian cancer due to delayed or missed diagnoses. Hence, there is an urgent need for affordable and accessible diagnostic solutions, especially in low-income facilities where costly technologies are out of reach. Addressing this disparity is crucial for ensuring equitable access to quality cancer care worldwide. The development of low-cost, portable diagnostic devices, such as the framework proposed in the 'OVision' project, marks a significant step forward in addressing these disparities. By leveraging the computing power of Raspberry Pi, these devices aim to dramatically reduce costs while maintaining high diagnostic accuracy. This innovative approach promises to bridge the gap in access to advanced cancer diagnostics, especially in resource-constrained settings, thus laying the foundation for a more equitable and inclusive future in precision medicine¹³.

The overarching question driving this work was how to leverage cutting-edge technologies to develop a framework for accurate, low-cost cancer diagnostic tools that can be deployed in a decentralized manner worldwide¹⁴. Specifically, the study investigated whether recent advances in deep learning and portable computing could be combined to create an affordable, self-contained device capable of reliably detecting cancer subtypes from histopathology data.

The central hypothesis of the study was that by deploying sophisticated deep neural networks on low-power devices like the Raspberry Pi, we could eliminate reliance on cloud computing, internet connectivity, and extensive infrastructure. This would enable the development of precise, standalone diagnostic tools at a fraction of current costs, helping bridge the divide in access to precision oncology.

To test this hypothesis, a compact, self-powered device driven by deep learning algorithms was designed and built. Utilizing histopathological biopsy images as inputs, the device demonstrated 95% accuracy in detecting ovarian cancer subtypes with just a few button clicks. This proof-of-concept validated overarching premise of the study, showing the remarkable potential for affordable, decentralized cancer diagnostics powered by portable artificial intelligence.

Materials and methods

Conceptual framework

This study presents a novel approach to address the economic and infrastructure barriers in advanced oncology care, with the goal of improving accessibility worldwide. The framework leverages recent advancements in machine learning and low-cost computing to develop an affordable and reliable solution for ovarian cancer classification. Figure 1 illustrates the conceptual workflow of this approach starting with data pre-processing (Fig. 1a), deep learning and training the model (Fig. 1b) and deploying the model (Fig. 1c). This flow aims to enhance care delivery, particularly in communities with limited access to precision diagnostics.

Image dataset

Eighty Hematoxylin and Eosin (H&E) -stained¹⁵ ovarian cancer tissue specimens were obtained from Simon Fraser University, comprising 29 high-grade serous (SC), 9 low-grade serous (LC), 21 clear cell (CC), 11 endometrioid (EC), and 10 mucinous (MC) cases. These specimens originated from a Trans-Canadian study conducted in 2006^{16,17} and the histologic subtype was determined by trained pathologists involved in that study.

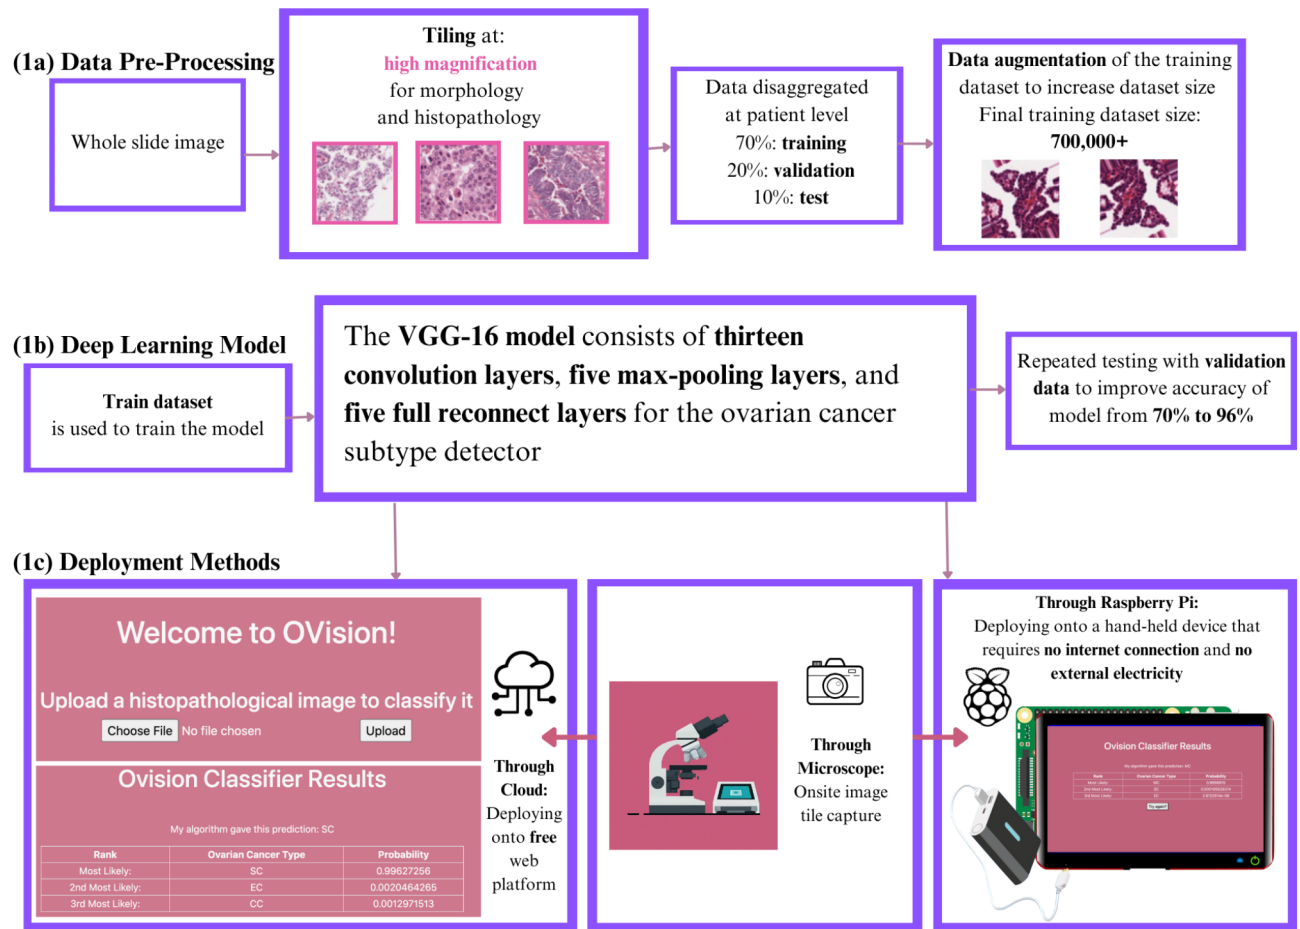


Fig. 1. Overview of the proposed comprehensive and innovative solution. (a) Data Pre-Processing to prepare the dataset - collection of samples, patch extraction for region of interest by using 20x magnification, tiling, dropping tile images with (> 50%) white space, data augmentation and dividing the data sample to train, validation and test samples. The division of train, validation and test set was per patient based to avoid any data leaks. The model will only see train and validate data and will never see the test sample patient to allow independent testing of the final model. (b) Training the deep learning model – Iterative process to create the final deployment model. (c) Deploying the saved model from the previous step to cloud and handheld RaspberryPi platforms to enable the ease of access to pathologists worldwide. The microscope with a camera represented in the flow can help local onsite capture images of 20x magnification objective (aligned with the trained model) that can be processed directly to achieve the classification results. (Code, Installation and Usage guide available on GitHub, OC data used in this work is available from Simon Fraser University^{16,17}). GitHub link to download and deploy OVision framework - <https://github.com/samairamehta/OVision-Framework/>.

Our study recognizes the complex role of stroma in tumor progression and its potential to reflect characteristics specific to each cancer subtype. This approach is supported by recent research highlighting the importance of the tumor microenvironment in ovarian cancer progression and classification. For instance, a study¹⁸ on LC carcinoma revealed significant differences in the tumor microenvironment compared to serous borderline ovarian tumors (SBTs), with fibroblast activation protein (FAP) identified as a key differentiator. Additionally, research¹⁹ has shown that machine learning models can achieve high accuracy in classifying ovarian carcinoma histotypes, capturing subtle features in both tumor cells and surrounding stroma.

Dataset structure

To transform the 80 whole slide images into the final dataset, a patient-level split was implemented to prevent data leakage, with 70% allocated for training, 20% for validation, and 10% for testing. This approach ensures data integrity throughout the dataset preparation process. The process, illustrated in Fig. 2, followed these steps:

1. Patch extraction:

- 20 non-overlapping patches extracted from each whole slide image at 20X magnification.
- The slide includes representative areas of both tumor and stromal tissue. This approach allows our deep learning model to learn from the full spectrum of tissue characteristics.

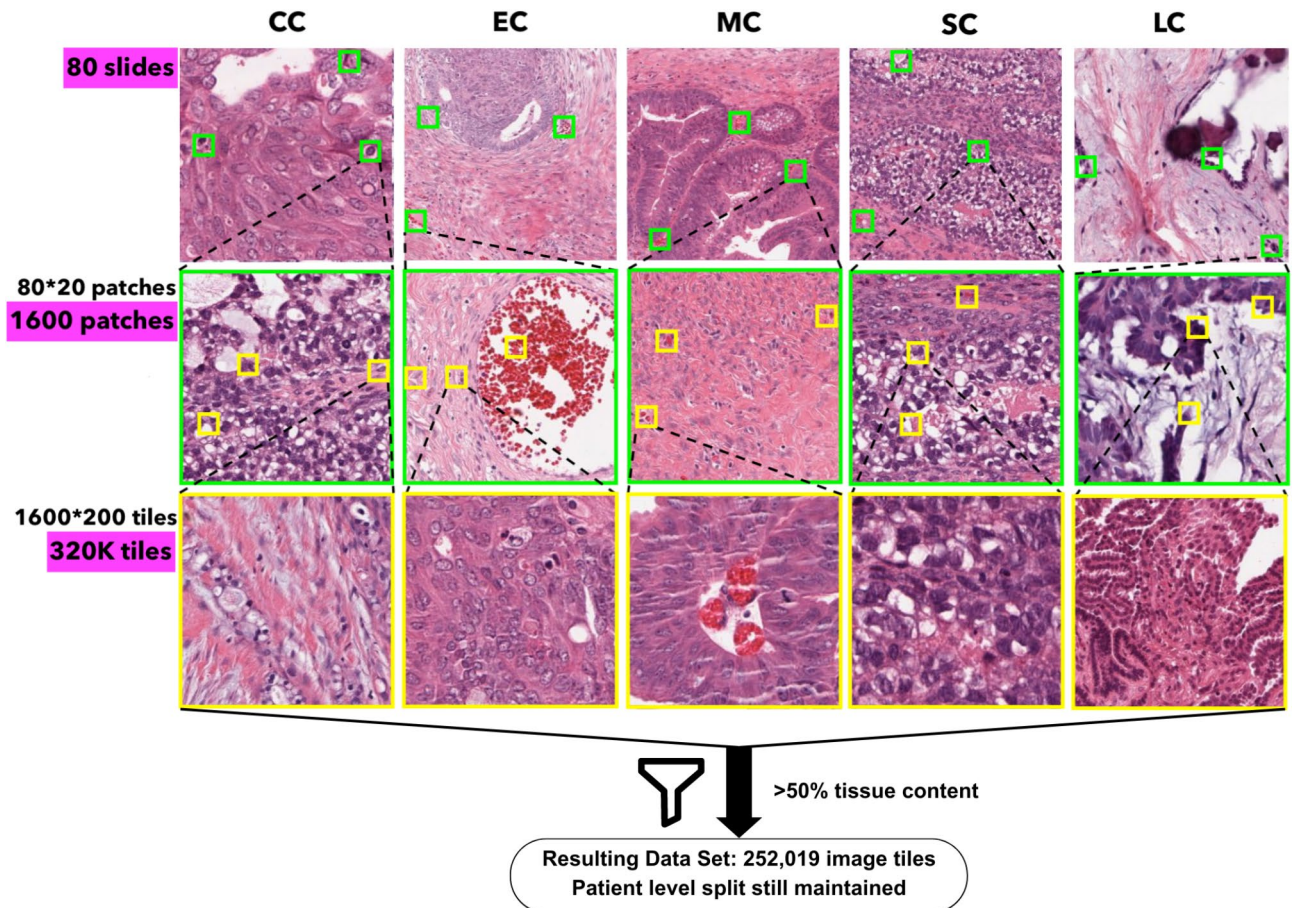


Fig. 2. Dataset Preparation Workflow: From whole slide images to filtered image tiles. Maintaining the patient-level splitting, the process includes patch extraction, tiles generation, and tile filtering to ensure > 50% tissue in each time, resulting in 252,019 histopathology image tiles for model training and evaluation.

- The slides used represented heterogeneous histologies, reflecting the complexity represented in clinical settings.
 - Total patches: 80 slides \times 20 patches = 1,600 patches.
2. Tile generation:
- 200 tiles of 224×224 pixels generated from each patch.
 - Total tiles : 1,600 patches \times 200 tiles = 320,000 tiles.
3. Tissue content filtering:
- Tiles with > 50% tissue content (file size > 15kB) retained.
 - Resulting dataset: 252,019 histopathology image tiles. Figure 3 represents 3 samples for each sub-type from final dataset.

The proportion of tissue content in each tile was determined by analyzing file size—a method validated through visual correlation of 50 sample tiles. Image tiles with more tissue content typically have larger file sizes due to increased detail. Tiles exceeding 15kB in size, indicating more than 50% tissue content, were retained for further processing.

To ensure balanced class distributions and mitigate bias, dataset normalization was performed, and data augmentation techniques were applied. Categories with fewer instances were weighted higher during training via oversampling. For training data only, additional augmentation was applied, including rotations (90° , 180° , 270°) and other transformations. This strategy expanded the training dataset to over 700,000 images, enhancing the model's ability to generalize across various image orientations and transformations.

To address potential overfitting concerns, the 80 patients were stratified by cancer subtype before splitting, maintaining a similar distribution across sets. The test set was completely held out from all stages of model development and hyperparameter tuning. The process was repeated 5 times with different random seeds, with the reported 95% accuracy being the average performance across these independent runs.




Fig. 3. Sample Slide Exploration - Example patches from the training set, along with their tumor histology labels (Pathologists determined these histologic subtype classes), were classified into ovarian cancer subtypes: MC: Mucinous Carcinoma, LC: Low Grade Serous Carcinoma, CC: Clear Cell Carcinoma, SC: High Grade Serous Carcinoma, and EC: Endometrioid Carcinoma.

Deep learning model

The study compared two deep learning convolutional neural network (DLCNN) architectures to classify ovarian cancer subtypes: VGG-16^{20,21} and the recent EfficientNetV2B0²². Both models were selected for evaluation based on their distinct strengths: VGG-16 for its proven performance in medical image classification, and EfficientNetV2B0 for its computational efficiency.

The VGG-16²³ model, pretrained on over 14 million ImageNet images, is known for its effectiveness in transfer learning, ability to learn robust visual representations from large-scale datasets, and compatibility with resource-constrained devices. Its straightforward structure facilitates easier interpretation of learned features, crucial for medical applications. In contrast, EfficientNetV2B0 is designed to optimize both accuracy and efficiency, potentially offering faster training and inference times.

The model selection process involved the following steps:

1. Training both VGG-16 and EfficientNetV2B0 on the ovarian cancer dataset.
2. Evaluating performance using metrics including accuracy, precision, recall, and F1-score for each cancer subtype and overall.
3. Assessing model generalization across different subtypes.
4. Considering model interpretability for medical applications.

The evaluation results showed that VGG-16 significantly outperformed EfficientNetV2B0 in overall accuracy (86% vs. 49%) and across all subtypes. VGG-16 achieved higher F1-scores for each cancer subtype: Clear Cell (CC: 0.89 vs. 0.85), Endometrioid (EC: 0.86 vs. 0.26), Low-grade serous (LC: 0.89 vs. 0.66), Mucinous (MC: 0.87 vs. 0.36), and High-grade serous (SC: 0.82 vs. 0.54). The performance evaluation as detailed in Table 1 utilized the following metrics:

- Accuracy: Overall correctness of predictions (VGG-16: 86%, EfficientNetV2B0: 49%).
- Precision: Proportion of correct positive predictions (VGG-16 macro avg: 0.85, EfficientNetV2B0: 0.65).
- Recall: Proportion of actual positive cases correctly identified (VGG-16 macro avg: 0.89, EfficientNetV2B0: 0.61).
- F1-score: Harmonic mean of precision and recall (VGG-16 macro avg: 0.87, EfficientNetV2B0: 0.53).
- Class-specific metrics: Detailed performance for each ovarian cancer subtype.

These metrics provided a comprehensive evaluation across all ovarian cancer subtypes, assessing both overall accuracy and class-specific performance. This approach is crucial for medical applications where misclassification costs may vary between subtypes.

Based on this rigorous evaluation process and the superior performance metrics, VGG-16 was selected as the final model for ovarian cancer subtype classification. Its higher accuracy, consistency, and reliability across all subtypes make it well-suited for the specific challenges of ovarian cancer classification in clinical settings.

To adapt the VGG-16 model for the ovarian cancer subtyping task, 5 dense layers were added to enable multi-class prediction of the 5 major histologic subtypes (SC, LC, CC, EC, MC). Figure 4 provides a visual representation of this VGG-16 architecture, while Table 2 lists the layers, parameters, and dimensions of the adapted model, highlighting the fully connected dense layers and their role in facilitating the classification of ovarian cancer subtypes.

The training process was optimized through careful hyperparameter tuning²⁴ to achieve faster convergence and better early performance. This tuning process for the VGG-16 model involved the following steps, focusing on key parameters known to affect model performance:

- Initial grid search: A series of experiments was conducted varying batch size (16, 32, 64), optimizers (Adam, SGD, RMSprop), learning rates (0.001, 0.0005, 0.0001), and momentum values (0.8, 0.9).
- Performance evaluation: Each configuration was evaluated based on initial accuracy after one epoch of training.
- Fine-tuning: Based on initial results, the most promising configurations were further refined.
- Final selection: The best-performing configuration was chosen for full model training.

Particular attention was paid to tuning the learning rate and momentum for the SGD optimizer, as these parameters significantly impact the model's ability to converge on an optimal solution for fine-tuned pre-trained models.

Table 2 presents the results of six hyperparameter tuning experiments for the VGG16 model, each with different combinations of batch size, optimizer, learning rate, and momentum. The initial accuracy after one epoch of training was used as a quick indicator of each configuration's potential. The configuration from Run 4 (highlighted in bold in Table 3) was selected for full model training due to its superior initial accuracy of 75%. This configuration (batch size: 32, SGD optimizer, learning rate: 0.0001, momentum: 0.9) provided the best balance between learning speed and stability, proving particularly effective for fine-tuning the pre-trained VGG16 model on the ovarian cancer dataset.

Based on the hyperparameter tuning results, the final model was trained with the following configuration:

Class	EfficientNetV2B0			VGG16		
	Precision	Recall	F1-score	Precision	Recall	F1-score
CC	0.99	0.74	0.85	1.00	0.81	0.89
EC	0.73	0.16	0.26	0.82	0.90	0.86
LC	0.58	0.76	0.66	0.82	0.97	0.89
MC	0.22	0.97	0.36	0.86	0.88	0.87
SC	0.72	0.44	0.54	0.75	0.90	0.82
Macro avg	0.65	0.61	0.53	0.85	0.89	0.87
Weighted avg	0.71	0.49	0.50	0.88	0.86	0.87
Overall accuracy	49%			86%		

Table 1. Initial comparison of DLCNN model performance: EfficientNetV2B0 vs. VGG16. VGG16 demonstrated superior performance for the current classification over EfficientNetV2B0, newer DLCNN model and hence VGG16 was chosen for our study.




Fig. 4. VGG-16 Network Architecture adapted for classification of OC sub-types.

Layer (type)	Output shape	Number of learnable parameters
input_1 (Input Layer)	[(224, 224, 3)]	0
block1_conv1 (Conv2D)	(224, 224, 64)	1,792
block1_conv2 (Conv2D)	(224, 224, 64)	36,928
block1_pool (MaxPooling2D)	(112, 112, 64)	0
block2_conv1 (Conv2D)	(112, 112, 128)	73,856
block2_conv2 (Conv2D)	(112, 112, 128)	147,584
block2_pool (MaxPooling2D)	(56, 56, 128)	0
block3_conv1 (Conv2D)	(56, 56, 256)	295,168
block3_conv2 (Conv2D)	(56, 56, 256)	590,080
block3_conv3 (Conv2D)	(56, 56, 256)	590,080
block3_pool (MaxPooling2D)	(28, 28, 256)	0
block4_conv1 (Conv2D)	(28, 28, 512)	1,180,160
block4_conv2 (Conv2D)	(28, 28, 512)	2,359,808
block4_conv3 (Conv2D)	(28, 28, 512)	2,359,808
block4_pool (MaxPooling2D)	(14, 14, 512)	0
block5_conv1 (Conv2D)	(14, 14, 512)	2,359,808
block5_conv2 (Conv2D)	(14, 14, 512)	2,359,808
block5_conv3 (Conv2D)	(14, 14, 512)	2,359,808
block5_pool (MaxPooling2D)	(7, 7, 512)	0
flatten (Flatten)	(25088)	0
fc7 (Dense)	(1024)	25,691,136
prob (Dense)	(5)	5,125

Table 2. VGG-16 network architecture details. This table presents the layer-wise structure of the VGG-16 model, including each layer’s type, output shape, and number of parameters. The network consists of 13 convolutional layers (Conv2D), 5 max-pooling layers (MaxPooling2D), and 3 fully connected layers (including the final output layer). The input shape is 224 × 224 × 3, and the model outputs 5 classes.

- Batch size: 32.
- Optimizer: Stochastic Gradient Descent (SGD).
- Learning rate: 0.0001.
- Momentum: 0.9.

Run	Batch size	Optimizer	Learning rate	Momentum	Initial accuracy (one epoch run)
1	16	Adam	0.001	0.9	67%
2	32	SGD	0.0001	0.8	69%
3	64	RMSprop	0.0005	0.9	71%
4	32	SGD	0.0001	0.9	75%
5	32	Adam	0.0001	0.9	68%
6	16	SGD	0.0001	0.9	70%

Table 3. DLCNN model hyperparameter tuning experiments for VGG16 model to select best training parameters. The bold row was chosen as starting hyperparameters for the full training.

This configuration allowed for effective fine-tuning of the pre-trained VGG16 model without drastically altering its learned features. The DLCNN was developed using the TensorFlow framework and Keras API and trained on Kaggle cloud GPUs (NVidia T4 × 2 with 16GB GDDR6 memory) for approximately 3 h and 35 min.

This approach resulted in a more efficient learning process, allowing the model to achieve optimal performance with fewer epochs compared to traditional training methods²⁵. The systematic hyperparameter tuning and optimization process significantly contributed to the model's efficiency and effectiveness in classifying ovarian cancer subtypes.

Cloud deployment

To enable global accessibility, the trained DLCNN model was deployed onto a cloud web platform hosted on DigitalOcean servers which are virtual private servers (VPS) provided by DigitalOcean, a cloud infrastructure provider. These servers, called “Droplets,” offer scalable computing resources for hosting websites, applications, and development environments. Users can quickly deploy, manage, and scale applications on these reliable, cost-effective cloud servers. The final Keras models were exported in HDF5 format and uploaded to cloud virtual machines provisioned through DigitalOcean.

A user-friendly web interface (Fig. 5) with file upload functionality was built using Python Flask, allowing medical professionals worldwide to securely submit histopathology images for analysis. Upon upload, the cloud platform returns the most likely ovarian cancer subtype, ranked predictions for other subtypes and corresponding confidence scores.

Edge deployment

Edge computing refers to processing data near the source of data generation, rather than relying on a centralized location that may be far away. In the context of AI and machine learning, edge deployment means running models on local devices instead of cloud servers, enabling faster processing, reduced latency, and operation without constant internet connectivity.

For settings lacking reliable internet or power infrastructure, OVision was implemented on an embedded Raspberry Pi 4 (2GB RAM) device (Figs. 6 and 7) for offline, portable usage. This edge deployment strategy ensures that advanced cancer diagnostics can be performed in remote or resource-constrained environments. The Raspberry Pi was chosen for its optimal balance of computational power, energy efficiency, and cost-effectiveness, making it an ideal choice for edge deployment in healthcare applications. To manage data without internet connectivity, the device employs multiple strategies: saving results to local patient directories, enabling USB-based data transfer, and offering a print option for immediate physical records.

The trained models were exported in TensorFlow Lite format, optimized specifically for the Pi's ARM CPU. This optimization allows for efficient model execution on the limited computational resources of the embedded device. The Raspberry Pi's ARM architecture is particularly well-suited for running TensorFlow Lite models, offering a good trade-off between performance and power consumption.

A custom graphical user interface (GUI) was developed using Python Flask, enabling intuitive loading of image files and executing predictions via simple button controls. This GUI was specifically designed to enhance user interaction with the system, making it more accessible and user-friendly. The interface integrates features for local data management, including options to save results to patient-specific directories, initiate USB transfers, and print diagnostic reports directly from the device.

The edge device is entirely self-powered by a rechargeable lithium battery pack, eliminating the need for external electricity. This feature ensures true portability and independence from power outlets, making it suitable for use in areas with unreliable or non-existent power infrastructure. The Raspberry Pi's low power consumption characteristics make it an ideal choice for battery-powered operation, enabling extended use in the field.

With both cloud and edge deployments, OVision allows seamless access to advanced cancer diagnostics worldwide, overcoming barriers of internet connectivity, computing resources, and infrastructure constraints. This dual deployment strategy ensures that the system can be utilized effectively in a wide range of settings, from well-equipped hospitals to remote clinics in developing regions. In offline scenarios, diagnostic results are stored in local patient directories on the Raspberry Pi. These results can be transferred via USB to other devices when needed or printed immediately for physical record-keeping. When internet connectivity becomes available, the locally stored data can be easily uploaded to a central database, ensuring data consistency across different deployment environments.

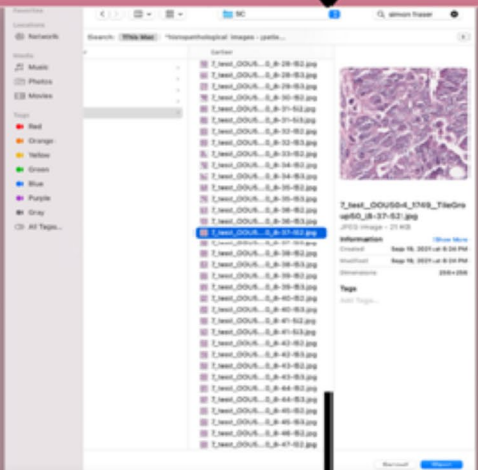
1. welcome page

Welcome to OVision!

Upload a histopathological image to classify

Choose File No file chosen Upload

2. upload image



3. final results

Ovision Classifier Results

My algorithm gave this prediction: SC

Rank	Ovarian Cancer Type	Probability
Most Likely:	SC	0.99627256
2nd Most Likely:	EC	0.0020464265
3rd Most Likely:	CC	0.0012971513

Try again?

Fig. 5. OVision Web Platform - three step process of using the subtype detector – Visit platform, Choose File, Upload gives final results.

Results


To validate the model's performance, a strategy of 5 independent runs was employed. This approach, while not a full 5-fold cross-validation, provides a robust assessment of the model's consistency and generalizability. A comprehensive statistical evaluation was then conducted on the test data set for each run. Standard metrics, including accuracy, sensitivity (recall), F1-score, and area under the receiver operating characteristic curve (AUROC), were calculated. (Table 4). The results from all 5 runs were averaged to provide a more reliable estimate of the model's performance and to account for potential variability in the data splits. This methodology was also applied in simulated science fair demonstrations and university research collaborations, where the model's performance was showcased across multiple runs to illustrate its consistency to attendees and potential research partners.



7. Final Raspberry-Pi powered device

Fig. 6. Final Medical Device Framework - RaspberryPi powered compact (2.5 inch by 3 inch) final medical device with LCD, battery pack and optional mini keyboard. The final trained and validated model has been uploaded to the device and device is configured ready to accept image from USB drive and classify providing fast and accurate results in matter of seconds.

Confusion matrix (Fig. 8) visualized the true and false positive rates across different subtypes. The high diagonal values indicate an excellent classification model under this framework, demonstrating the classifier's accuracy in correctly identifying each subtype. Off-diagonal values are minimal, reflecting a low rate of misclassification and further validating the robustness and reliability of the model. The normalization process ensures that the performance metrics are comparable across subtypes, providing a clear visual representation of the classifier's effectiveness. The evaluation revealed promising results, with an overall accuracy of 95%. The model demonstrated high precision, recall, and F1-scores across all subtypes, ranging from 0.90 to 1.00. The weighted average precision, recall, and F1-score were 0.95, 0.95, and 0.95, respectively, indicating robust and well-balanced performance. The macro-averaged precision, recall, and F1-score were 0.95, 0.95, and 0.95, respectively. These comprehensive metrics, coupled with visual representations of the confusion matrices and AUC-ROC curve (Figs. 8 and 9), provide a thorough evaluation of the model's solid predictive capabilities. In educational settings and medical training simulations, these visualizations served as powerful tools to explain the model's performance to students, science fair attendees, and medical professionals, enhancing their understanding of AI in medical diagnostics. The system's performance can also be used in training scenario



8. Process of using hand-held device for classification

Fig. 7. The three-step process of using the hand-held device – Once the data is available on the device using any external storage (e.g. USB), the data can be uploaded selecting file and the classification results will be available with a click.

Class	Metric	Mean	Std-Deviation
CC	Precision	0.978	1.4%
	Recall	0.986	1.6%
	F1-score	0.980	0.8%
EC	Precision	0.952	2.7%
	Recall	0.926	3.7%
	F1-score	0.938	0.9%
LC	Precision	0.950	4.8%
	Recall	0.968	2.0%
	F1-score	0.946	2.6%
MC	Precision	0.940	3.8%
	Recall	0.898	6.5%
	F1-score	0.918	2.9%
SC	Precision	0.928	3.3%
	Recall	0.962	1.3%
	F1-score	0.946	1.5%
Macro avg.	Precision	0.952	1.5%
	Recall	0.947	1.2%
	F1-score	0.948	0.8%
Weighted avg.	Precision	0.950	0.8%
	Recall	0.948	0.8%
	F1-score	0.948	0.8%
Overall accuracy	Accuracy	94.8%	0.8%

Table 4. Performance metrics of DLCNN model for ovarian Cancer subtype classification (*n* = 5 runs).

where medical students could compare their manual diagnoses to the AI’s predictions, showcasing its potential as a learning tool in pathology education.

A central goal of this study was to create a cancer diagnostic solution that could function independently, without requiring internet connectivity or extensive infrastructure. To accomplish this, the deep learning model for ovarian cancer subtyping was deployed across two complementary platforms - a cloud-based web interface and a standalone device powered by a Raspberry Pi microcomputer. This dual-deployment strategy not only serves potential clinical needs but also provides versatility for educational demonstrations at science fairs, technology exhibitions, and in classroom settings. It was also presented as a case study in a simulated healthcare startup incubator, demonstrating its potential for commercialization and widespread impact.

For the cloud deployment, Digital Ocean’s hosting services were utilized, to make the classifier freely available worldwide through a simple graphical web interface. After completing the training and validation of the DLCNN, the final neural network weights were exported and deployed them on Digital Ocean’s virtual servers. This allows medical professionals globally to visit the website, upload histopathology image files, and instantly receive detailed subtyping results and confidence scores, without any local software installation required. In a potential remote learning scenario, this cloud-based system can be integrated into an online course for medical professionals, demonstrating how it could be used to enhance diagnostic skills training in a digital environment.

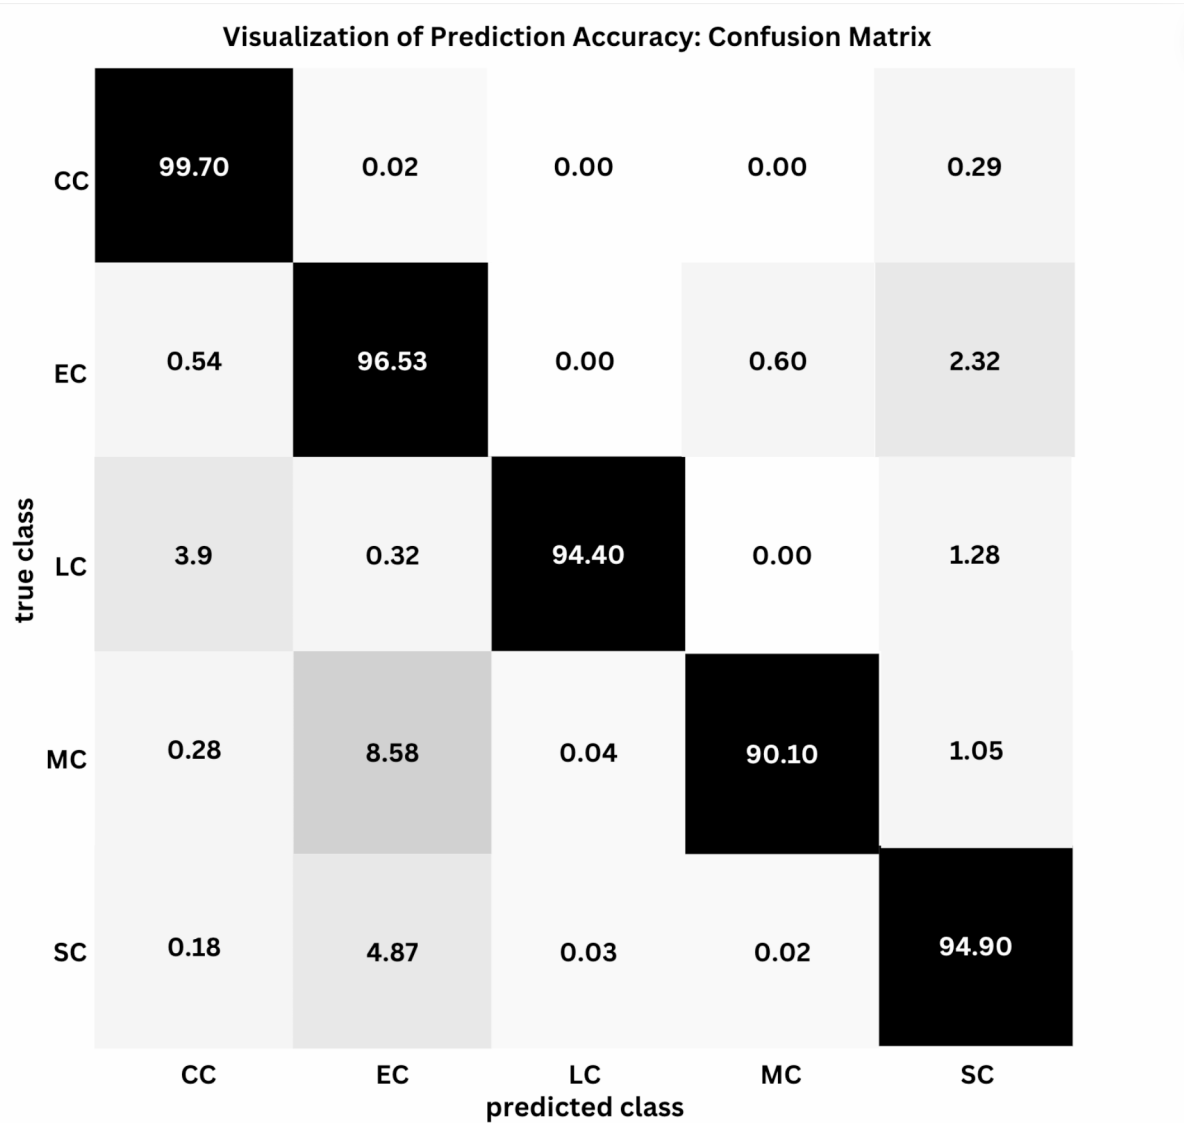


Fig. 8. Normalized (%) confusion matrix for the DLCNN subtype classifier. This matrix illustrates the performance of the classifier across different subtypes, with each cell representing the percentage of tiles correctly and incorrectly classified.

However, studies have showed that some clinics in underserved regions lack reliable internet access^{26,27} and having a no-internet device option can help improve access to healthcare diagnostics for these areas. Hence to compliment the connectivity barrier, a compact, self-powered physical device was developed, by directly installing the model onto a Raspberry Pi microcomputer. This portable unit does not require an internet connection or external power source to operate. Users simply need to connect it to a display and battery pack, then utilize an intuitive button interface to load images and generate subtype predictions.

Remarkably, deploying the model at the “edge” on this low-cost IoT device enabled rapid and highly accurate performance. In validation experiments using a fully held-out test set, the Raspberry Pi generated subtype predictions in under 10 s for all cases, with most classifications occurring within just 2–7 s. Despite its compact form factor, it achieved 95% accuracy in correctly identifying the five major ovarian cancer subtypes from histopathology data alone. This rapid response time proved particularly engaging in science fair demonstrations and public health awareness campaigns, allowing for real-time interactions and fostering discussions about AI applications in healthcare. In a simulated interdisciplinary research demonstration, computer science and bioengineering students collaborated to explore the system’s potential for further optimization and integration with other medical technologies.

This high level of performance is clearly visualized through the receiver operating characteristic (ROC) curves and confusion matrices. The area under the ROC curve exceeded 99%, indicating excellent discrimination between subtypes. The confusion matrix’s strong diagonal demonstrates the high true positive rates across classes, ranging from 90% for MC, up to 99% for CC.

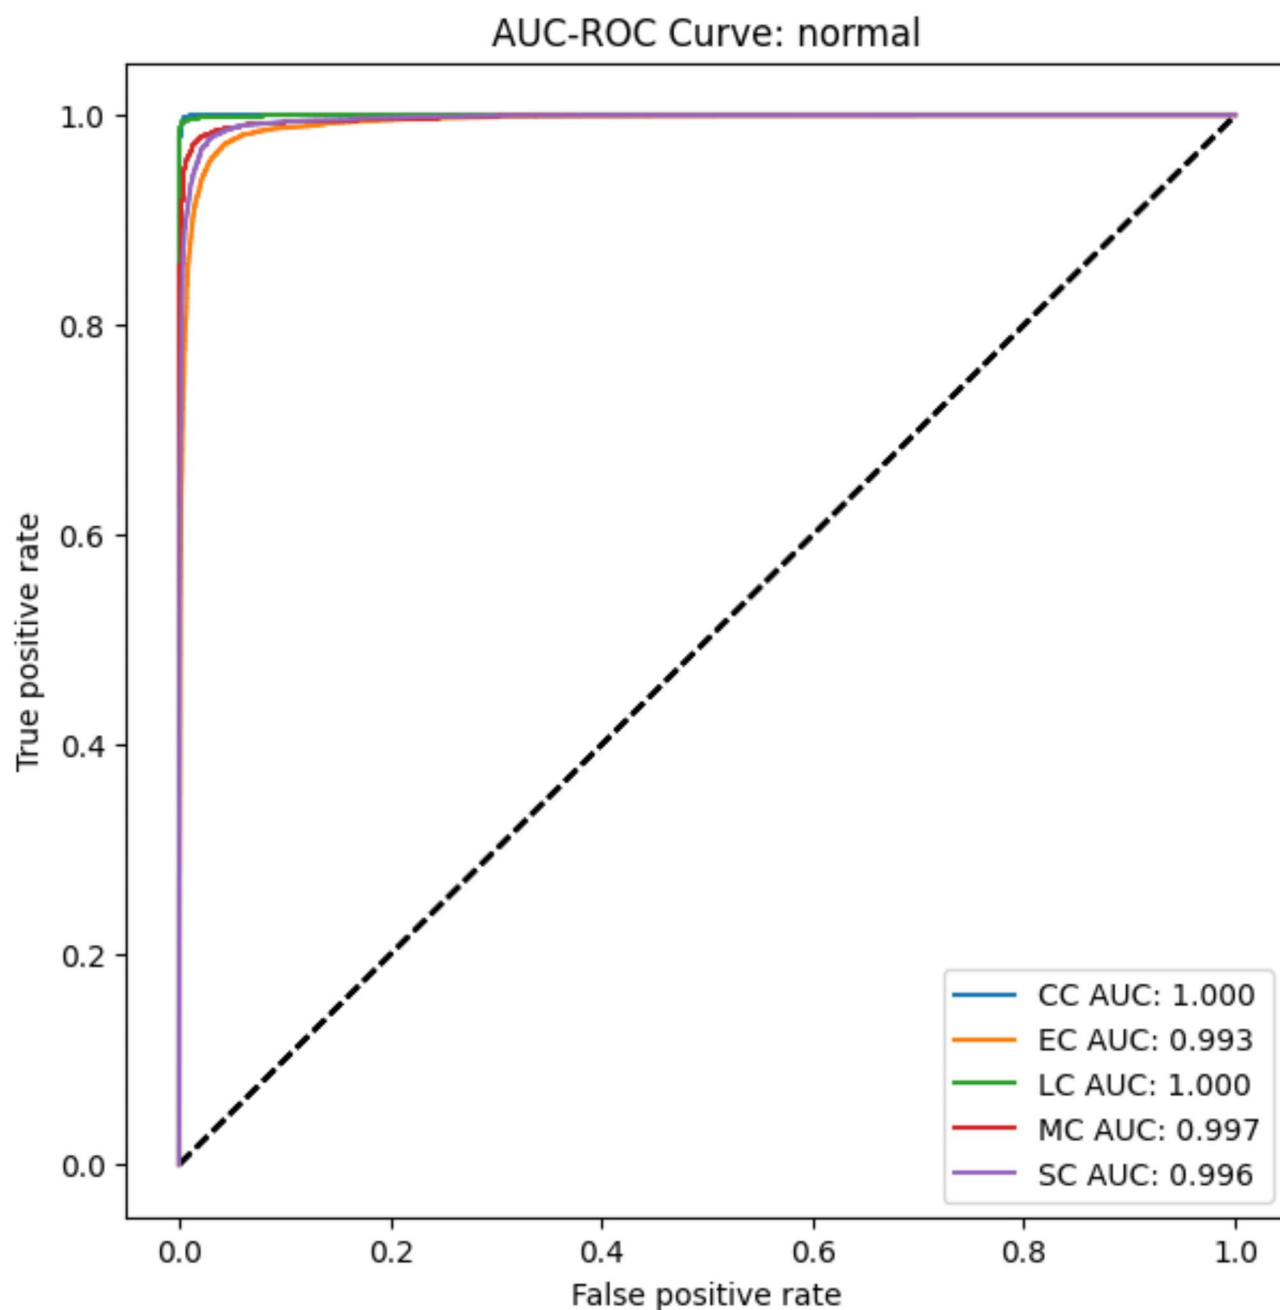


Fig. 9. The Area under the ROC curve (AUC) graph for the DLCNN subtype classifier - AUC results are considered excellent for AUC values between 0.9-1, good for AUC values between 0.8-0.9 and reasonable for lower values. The model achieved AUC value between 0.9-1 for all classes and in excellent results range.

Through these complementary cloud and edge computing deployments, powered by rigorously validated deep learning architecture, a flexible yet highly accurate diagnostic tool for ovarian cancer was developed. This AI solution can be readily utilized across the full spectrum of care settings globally, from urban medical centers with reliable internet access, to remote clinics operating fully offline. It provides a powerful aid for streamlining clinical decision-making and enabling precision treatment planning. Additionally, the system's versatility extends to educational contexts, where it serves as an interactive tool for teaching AI concepts and inspiring future innovations in medical technology at science fairs, in classroom settings, and during technology exhibitions. In a webinar, the tool was used to demonstrate the intersection of AI and healthcare to a global audience of medical and technology professionals, highlighting its potential for improving cancer diagnostics worldwide.

These diverse demonstrations and educational applications form a crucial foundation for future controlled deployment exercises in rural clinics and hospitals. By thoroughly testing and refining the system in these varied settings, the way has been paved for more extensive real-world trials, ensuring that the OVision platform is



Fig. 10. Deployment option for use case for OVision.

robust, user-friendly, and ready for practical implementation in healthcare settings where a significant impact on patient care can be made.

Discussion

The results demonstrate the remarkable potential of the proposed OVision platform, which leverages deep learning to detect ovarian cancer subtypes rapidly and accurately from histopathology images. With an overall accuracy of 95% for subtyping detection across a held-out test set, OVision demonstrates performance that is on par with existing deep learning approaches²⁸, while delivering critical diagnostic insights in under 10 s per case. The rigorous data splitting strategy and multiple independent runs provide strong evidence that this accuracy is a reliable measure of the model's generalization capability. The use of patient-level splitting ensures that the model's performance is not artificially inflated by data leakage or biased subtype distribution. The consistency of high performance across different runs further supports the robustness of the model. Medical professional can either choose to use the web platform deployed on cloud OR the compact handheld device to obtain the classification results. (Fig. 10) These results when associated with phenotype clinical data of the patient, will help professionals create targeted diagnosis plan for the patient, leading to a more promising response from patient.

While the current study focuses on ovarian cancer, the OVision framework has been designed with broader applicability in mind. The architecture and methodologies employed have the potential to be adapted for various cancer types, underscoring the significance of OVision's contribution to the field of AI-assisted cancer diagnostics.

For instance, the same deep learning architecture could theoretically be applied to lung adenocarcinoma histologies or prostate cancer grading, provided appropriate datasets are used for training and validation. The image preprocessing steps, the model's ability to handle heterogeneous tissue samples, and the rapid inference capabilities of the Raspberry Pi deployment could all be valuable in these contexts as well. The framework's adaptability is rooted in several key features:

1. Flexible image processing pipeline: The preprocessing steps can be adjusted to accommodate different staining techniques or tissue characteristics specific to other cancer types.
2. Modular deep learning architecture: The VGG-16 model used in this study can be fine-tuned or replaced with other architectures more suitable for specific cancer types, while maintaining the overall framework.
3. Scalable deployment: The Raspberry Pi platform allows for easy scaling and adaptation to different clinical settings, regardless of the cancer type being diagnosed.
4. Generalizable performance metrics: The evaluation methodology, including multiple independent runs and comprehensive performance metrics, can be applied to validate the model's effectiveness across various cancer types.

However, it's important to note that such applications would require rigorous testing, validation, and potentially, modifications to the model architecture to account for the specific characteristics of each cancer type. Future research could explore:

1. Adapting the framework for other common cancers such as breast, colorectal, or skin cancer.
2. Investigating the framework's effectiveness in rare cancer diagnostics, where the portability and low cost of the system could be particularly beneficial.
3. Exploring multi-cancer classification models within the OVision framework.

While studies on other cancer types have not been conducted in this work, the potential for broader application underscores the significance of OVision's contribution to the field of AI-assisted cancer diagnostics. The adaptation of this framework to other cancer types could be explored in future research, which would further validate its versatility and impact on cancer diagnosis more broadly.

In addition, by classifying each tile, the tool can provide percentages of each tumor histology detected within the slide. This capability enhances diagnostic precision, offering a detailed overview of the heterogeneity within each tissue sample. Such granularity is crucial for understanding the complexity of tumor histology and tailoring personalized treatment plans. The ability to quantify the presence of different histology within a single tissue sample allows for a more nuanced understanding of the tumor environment, which can impact treatment decisions. This detailed insight is particularly valuable for creating personalized treatment plans that are more likely to be effective, improving patient outcomes and advancing the field of precision oncology.

The true impact of OVision lies in its potential to streamline ovarian cancer diagnosis and enable more personalized, timely treatment planning. By reliably identifying subtypes like clear cell carcinoma that increase clotting risk, physicians can immediately prescribe preventative anticoagulants without waiting for additional

tests. Integrating OVision's outputs with other clinical data points allows clinicians to develop comprehensive, precision treatment regimens tailored to each patient's specific disease characteristics and risk factors.

The Raspberry Pi's utility as a diagnostic tool in remote areas is exemplified by its deployment in rural clinics with limited resources. In such settings, healthcare workers can use the device to perform on-site ovarian cancer subtyping without needing to send samples to distant laboratories, which can be time-consuming and costly. The conceptual workflow involves a healthcare worker collecting a histopathology sample, preparing it on a slide, and capturing an image using a digital microscope (Fig. 1c). Once the image is captured, it is saved onto a USB drive. The Raspberry Pi, equipped with USB ports, allows for easy transfer of these digital files. The healthcare worker can then insert the USB drive into the Raspberry Pi, which processes the data and provides a subtype prediction within seconds. This method leverages the portability and low power consumption of the Raspberry Pi, making it ideal for use in remote areas without reliable internet access. This rapid turnaround enables immediate clinical decision-making, such as initiating targeted treatments or referring patients to specialized care, thereby improving patient outcomes in resource-limited environments. The Raspberry Pi's ability to operate independently of internet connectivity, combined with its low cost and portability, makes it an ideal solution for bringing advanced diagnostic capabilities to underserved regions.

While these results are highly encouraging, several important scientific questions remain. Further study is needed to assess OVision's real-world clinical performance across diverse patient populations and healthcare settings. Extending the deep learning models to incorporate additional predictive biomarkers could enhance their diagnostic value. There is also an opportunity to apply similar AI-powered approaches to other cancer types beyond ovarian cancer like lung adenocarcinoma histologies, prostate cancer grading and others.

Crucially, it is clearly recognized that OVision is not a standalone diagnostic - it is intended to augment and complement existing methodologies within the overall cycle of care. Its subtyping and biomarker outputs are powerful decision support tools but should be synthesized with other data inputs by medical professionals. Overreliance on any single AI system for high-stakes medical decisions would be unwise.

In conclusion, this work validates a transformative new paradigm for leveraging artificial intelligence²⁹ and low-cost computing to deliver premium precision oncology decision support, even in resource-limited settings. The OVision platform provides affordable, decentralized access to accurate, rapid cancer diagnostics that have the potential to optimize treatment pathways, mitigate adverse effects, and ultimately improve patient outcomes globally. As AI capabilities continue advancing and real-world deployments expand, OVision represents a harbinger for a future of democratized, patient-centric cancer care for all.

Data availability

The datasets generated and/or analyzed during the current study are available in the <https://github.com/samaira-mehta/OVision-Framework/> repository and at <https://mia-www.cs.sfu.ca/download/>.

Received: 7 September 2024; Accepted: 24 February 2025

Published online: 28 February 2025

References

1. Ferlay, J. et al. *Global Cancer Observatory: Cancer Today* (International Agency for Research on Cancer, 2023). <https://gco.iarc.fr/today>
2. Sung, H. et al. Global cancer statistics 2020: GLOBOCAN estimates of incidence and mortality worldwide for 36 cancers in 185 countries. *Cancer J. Clin.* **71** (3), 209–249. <https://doi.org/10.3322/caac.21660> (2021).
3. Pramesh, C. S. et al. Priorities for cancer research in low- and middle-income countries: A global perspective. *Nat. Med.* **28** (4), 649–657. <https://doi.org/10.1038/s41591-022-01738-x> (2022).
4. Brinkhuis, F., Goettsch, W. G., Mantel-Teeuwisse, A. K. & Bloem, L. T. High cost oncology drugs without proof of added benefit are burdening health systems. *BMJ* **384**, q511. <https://doi.org/10.1136/bmj.q511> (2024).
5. Mitchell, A. Pathology job market struggling to meet increasing demand. Lighthouse Lab Services (2022). <https://www.lighthouselabservices.com/pathology-job-market-struggling-to-meet-increasing-demand/>
6. Schroeder, M. Deciphex scores \$11.5 M in bid to address pathologist shortage (2022). <https://medcitynews.com/2022/05/deciphex-scores-11-5m-in-bid-to-address-pathologist-shortage/>
7. Centers for Disease Control and Prevention (CDC). U.S. Cancer Statistics Data Visualizations Tool. (2024). <https://www.cdc.gov/cancer/ovarian/statistics/index.htm>
8. Kumaraswamy, E., Sharma, S. & Kumar, S. An invasive ductal carcinomas breast cancer grade classification using an ensemble of convolutional neural networks. *Diagnostics*, **13**(11), 1977. (2023). <https://doi.org/10.3390/diagnostics13111977>
9. Hussain, A. & Jabbar, S. A deep learning framework for the prediction and diagnosis of ovarian cancer in pre- and post-menopausal women. *Diagnostics* **13** (10), 1703. <https://doi.org/10.3390/diagnostics13101703> (2023).
10. Komura, D. & Ishikawa, S. Machine learning methods for histopathological image analysis. *Comput. Struct. Biotechnol. J.* **16**, 34–42. <https://doi.org/10.1016/j.csbj.2018.01.001> (2018).
11. Yildirim, Z., Samet, R., Hancer, E., Nemati, N. & Traore Mali, M. Gland segmentation in H&E histopathological images using U-Net with attention module. In *Proceedings of the 2023 Twelfth International Conference on Image Processing Theory, Tools and Applications (IPTA)* 1–6 (2023). <https://ipta-conference.com/ipta23/index.php/program>
12. Janowczyk, A. & Madabhushi, A. Deep learning for digital pathology image analysis: A comprehensive tutorial with selected use cases. *J. Pathol. Inf.* **7** (1), 29. <https://doi.org/10.4103/2153-3539.186902> (2016).
13. Raines, T. L. & Nkengasong, J. Bringing the benefits of precision medicine to lower-resourced settings. (Brookings, 2023). <https://www.brookings.edu/articles/bringing-the-benefits-of-precision-medicine-to-lower-resourced-settings/>
14. Yadav, K., Cree, I., Field, A., Vielh, P. & Mehrotra, R. Importance of cytopathologic diagnosis in early cancer diagnosis in resource-constrained countries. *JCO Global Oncol.* **8**, e2100337. <https://doi.org/10.1200/GO.21.00337> (2022).
15. Li, Y. et al. Hematoxylin and Eosin staining of intact tissues via delipidation and ultrasound. *Sci. Rep.* **8** (1), 12259. <https://doi.org/10.1038/s41598-018-30755-5> (2018).
16. Köbel, M. et al. Ovarian carcinoma subtypes are different diseases: Implications for biomarker studies. *PLoS Med.* **5** (12), e232. <https://doi.org/10.1371/journal.pmed.0050232> (2008).
17. Köbel, M. et al. Diagnosis of ovarian carcinoma cell type is highly reproducible: A transcanadian study. *Am. J. Surg. Pathol.* **34** (7), 984–993. <https://doi.org/10.1097/PAS.0b013e3181e1a3bb> (2010).

18. Vallejos, R. et al. Changes in the tumour microenvironment mark the transition from serous borderline tumour to low-grade serous carcinoma. *J. Pathol.* <https://doi.org/10.1002/path.6338> (2024).
19. Farahani, H. et al. Deep learning-based histotype diagnosis of ovarian carcinoma whole-slide pathology images. *Mod. Pathol.* **35** (12), 1983–1990. <https://doi.org/10.1038/s41379-022-01146-z> (2022).
20. Xu, Y. et al. Large scale tissue histopathology image classification, segmentation, and visualization via deep convolutional activation features. *BMC Bioinform.* **18**, 281. <https://doi.org/10.1186/s12859-017-1685-x> (2017).
21. Wu, M., Yan, C., Liu, H. & Liu, Q. Automatic classification of ovarian cancer types from cytological images using deep convolutional neural networks. *Biosci. Rep.* **38** (3), BSR20180289. <https://doi.org/10.1042/BSR20180289> (2018).
22. Tan, M. & Le, Q. EfficientNetV2: Smaller models and faster [preprint]raining [Preprint]. arXiv:2104.00298. (2021). <https://doi.org/10.48550/arXiv.2104.00298>
23. Simonyan, K. & Zisserman, A. Very Deep Convolutional Networks for Large-Scale Image Recognition [Preprint]. arXiv:1409.1556. (2014). <https://arxiv.org/abs/1409.1556>
24. Wojciuk, M., Swiderska-Chadaj, Z., Siwek, K. & Gertych, A. Improving classification accuracy of fine-tuned CNN models: impact of hyperparameter optimization. *Heliyon* **10** (5), e26586. <https://doi.org/10.1016/j.heliyon.2024.e26586> (2024).
25. Doğan, R. S. & Yılmaz, B. Histopathology image classification: Highlighting the gap between manual analysis and AI automation. *Front. Oncol.* **13**, 1325271. <https://doi.org/10.3389/fonc.2023.1325271> (2023).
26. Saeed, S. A. & Masters, R. M. Disparities in health care and the digital divide. *Curr. Psychiatry Rep.* **23** (9), 61. <https://doi.org/10.1007/s11920-021-01274-4> (2021).
27. Yu, J. & Meng, S. Impacts of the internet on health inequality and healthcare access: A cross-country study. *Front. Public. Health.* **10**, 935608. <https://doi.org/10.3389/fpubh.2022.935608> (2022).
28. Sadeghi, M. H., Sina, S., Omid, H., Farshchitabrizi, A. H. & Alavi, M. Deep learning in ovarian cancer diagnosis: A comprehensive review of various imaging modalities. *Pol. J. Radiol.* **89**, e30–e48. <https://doi.org/10.5114/pjr.2024.134817> (2024).
29. Orsulic, S., John, J., Walts, A. E. & Gertych, A. Computational pathology in ovarian cancer. *Front. Oncol.* **12**, 924945. <https://doi.org/10.3389/fonc.2022.924945> (2022).

Acknowledgements

I would like to extend my sincere gratitude to my mentors Dr. Sandra Orsulic, UCLA David Geffen School of Medicine, Stephanie Owens, Advanced Specialist Product Development - 3 M Interconnect products portfolio, and Dr. Arkadiusz Gertych, Cedars-Sinai Medical Center, for their invaluable guidance, support, and advise throughout the study and development of the portable device and framework. They also educated me with the process of submitting the results to a wider research community as a means to pursue the novel idea and platform for global outreach.

Author contributions

S.M. wrote the full manuscript.

Declarations

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1038/s41598-025-91914-z>.

Correspondence and requests for materials should be addressed to S.M.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2025