DEEP LEARNING WITH OPTICAL FLOW FOR AUTOMATED LUNG SLIDING DETECTION

Maroš HLIBOKÝ*, Katarína IŠTOŇOVÁ, Laura PITUKOVÁ, Marek BUNDZEL
Department of Cybernetics and Artificial Intelligence, Faculty of Electrical Engineering and Informatics,
Technical University of Košice, Letná 9, 042 00 Košice, Slovak Republic
E-mail: *maros.hliboky@tuke.sk

ABSTRACT

This study explores the application of optical flow techniques for the automated classification of lung sliding in lung ultrasound—a key indicator for assessing lung function and diagnosing pneumothorax. Optical flow is used to capture subtle pleural motion between frames. We evaluate both traditional machine learning models (Random Forest, Gradient Boosting) and deep learning architectures (CNN, ResNet-18), using various preprocessing methods to optimize flow representation. Among deep models, the CNN achieved the highest overall accuracy (80.8%) and F1-score (73.7%), while the more complex ResNet-18 reached a high recall (94.5%) but suffered from lower precision (62.6%), indicating a tendency to over-predict the positive class. These findings highlight the trade-offs between model complexity and generalization in limited-data scenarios. The results underline the importance of model selection and tailored preprocessing in improving diagnostic reliability.

Keywords: *optical flow, lung ultrasound, lung sliding, deep learning, machine learning, pneumothorax*

1. INTRODUCTION

Lung sliding, the subtle rhythmic movement of the pleural layers observed during respiration, is a vital diagnostic feature in thoracic ultrasonography [1]. Accurate detection and classification of lung sliding can assist in the timely diagnosis of critical conditions such as pneumothorax. However, traditional diagnostic methods rely heavily on clinicians' experience and visual assessment, which may introduce subjectivity and variability in interpretation.

Recent advancements in medical imaging and machine learning offer promising opportunities to automate such diagnostic tasks, thereby enhancing accuracy, consistency, and efficiency [2]. In particular, optical flow—originally developed in the field of computer vision—has shown potential for capturing dynamic motion in sequential imaging data. Optical flow quantifies the apparent motion of structures between consecutive frames [3], making it especially suitable for identifying subtle respiratory movements, such as lung sliding, in ultrasound sequences.

This study leverages optical flow techniques to extract temporal motion patterns from thoracic ultrasound videos, with the goal of automatically classifying the presence or absence of lung sliding using machine learning models. By translating motion information into a structured representation, we aim to develop a pipeline that supports robust classification without the need for handcrafted features.

In recent years, deep learning models, particularly Convolutional Neural Networks (CNNs) and Residual Networks (ResNets), have demonstrated remarkable performance in various image classification tasks [4]. These models are capable of learning hierarchical features directly from data and are especially well-suited to complex, high-dimensional inputs such as medical images [5].

To investigate the effectiveness of different modeling approaches, this study compares a simple convolutional neural network (CNN) with a deeper Residual Network, specifically ResNet-18. While ResNet-18 is known for its ability to train deep architectures efficiently due to residual connections [6], the custom CNN serves as a lightweight, computationally less demanding baseline [7]. By contrasting these two architectures, we aim to evaluate the trade-offs between model complexity, training efficiency, and classification performance in the context of lung sliding detection using optical flow data.

2. RELATED WORKS

The classification and detection of lung sliding, especially through lung ultrasound, has seen significant advancements with the application of deep learning techniques. Two notable studies in this field highlight the effectiveness of these approaches.

The classification and detection of lung sliding, especially through lung ultrasound, has seen significant advancements with the application of deep learning techniques. Two notable studies in this field highlight the effectiveness of these approaches. The first, Automated Detection of Pneumonia in Lung Ultrasound Using Deep Video Classification for COVID-19 [8], addressed the urgent need for rapid and accurate diagnostic tools during the COVID-19 pandemic. Lung ultrasound, being cost-effective and widely accessible, was used for diagnosing acute respiratory distress syndrome (ARDS) in COVID-19 patients. This study proposed a method leveraging the Kinetics-I3D network to classify lung ultrasound videos without requiring preprocessing or frame-by-frame analysis. Their approach achieved an accuracy of 90% and an average precision score of 95% with 5-fold cross-validation, demonstrating its potential to assist clinicians in quickly and reliably diagnosing pneumonia in COVID-19 patients. The second study, Detecting the Absence of Lung Sliding in Ultrasound Videos Using 3D Convolutional Neural Networks [9], focused on the crucial task of identifying the absence of lung sliding, which is important

for diagnosing conditions like pneumothorax. Unlike traditional methods that process static images, this study explored the use of deep learning models to analyze short video sequences. Various ResNet model architectures were evaluated on a real-world dataset, with the video-based approach outperforming traditional static image processing methods. This demonstrated that deep learning models, particularly those capable of handling temporal data, can more accurately detect the absence of lung sliding, providing valuable diagnostic information.

In another relevant work, Roy et al. proposed a deep learning model based on U-Net for segmenting lung regions in ultrasound images [10]. Their architecture was trained on annotated LUS data and achieved high Dice similarity scores, proving its utility for preprocessing or assisting downstream diagnostic models.

Another relevant work by Huang et al. [11], proposed an Attention U-Net for segmenting pleural effusions in ultrasound images. Their model significantly outperformed a standard U-Net, achieving a Dice coefficient between 0.83 and 0.90 versus 0.79–0.86 for U-Net, demonstrating how attention mechanisms can enhance the segmentation of pleural structures.

Another promising direction was explored by Ardon et al. [12], who developed a spatiotemporal deep learning framework that combines convolutional and recurrent layers (CNN-LSTM) to identify artifacts such as lung sliding and comet-tail artifacts in LUS videos. Their work demonstrated the added value of combining spatial and temporal features for robust LUS interpretation.

These studies illustrate the significant advancements in lung ultrasound analysis through deep learning. Automating the interpretation of ultrasound videos and improving diagnostic accuracy can greatly enhance clinical workflows and support healthcare professionals in delivering timely and reliable diagnoses.

These studies illustrate the significant advancements in lung ultrasound analysis through deep learning. Automating the interpretation of ultrasound videos and improving diagnostic accuracy can greatly enhance clinical workflows and support healthcare professionals in delivering timely and reliable diagnoses.

3. OPTICAL FLOW BACKGROUND

Optical flow refers to the pattern of apparent motion of objects, surfaces, or edges in a visual scene caused by the relative motion between an observer (either a camera or the eye) and the scene. The concept was first introduced by psychologist James Gibson in the 1940s to describe the visual stimulus that allows people to perceive motion [13]. In computer vision, optical flow is a crucial tool for estimating the motion of objects in consecutive video frames, based on the changes in pixels' intensities.

3.1. Principles of Optical Flow

Optical flow works under the assumption that the pixel intensities of an object remain consistent from one frame to the next, provided the motion is small. This principle allows the calculation of the displacement vectors (or flow vectors), which indicate the direction and speed of motion of pixels between frames. These vectors are computed based on the derivatives of pixel intensities with respect to time and space, typically using differential methods such as the Lucas-Kanade method or the Horn-Schunck algorithm [14].

3.2. RAFT Optical-Flow Network

RAFT (Recurrent All-Pairs Field Transforms for Optical Flow) is a deep network architecture for optical flow [15]. RAFT, illustrated in Figure 3, extracts per-pixel features, builds multi-scale 4D correlation volumes for all pairs of pixels, and iteratively updates a flow field through a recurrent unit that performs lookups on the correlation volumes. RAFT achieves state-of-the-art performance. RAFT has strong cross-dataset generalization as well as high efficiency in inference time, training speed, and parameter count [16]. Figure 3 illustrates the architecture of the RAFT model. This architecture is designed for optical flow estimation, showcasing the key components and their interactions. The figure depicts the process flow from input images through feature extraction, correlation volumes, and iterative updates to the final optical flow output.

3.3. Farneback Optical Flow

The Farneback optical flow method is the algorithm used to estimate the motion between two images based on polynomial expansion. Developed by Gunnar Farneback [17], this method offers a computationally efficient approach for calculating dense optical flow, which means it estimates the motion for every pixel in the image sequence. Fig. 2 illustrates the Farneback optical flow algorithm. This figure demonstrates the process of estimating motion between two images using the Farneback method, highlighting the steps involved in computing dense optical flow fields.

3.3.1. Key Principles of Farneback Optical Flow

The Farneback algorithm approximates each neighborhood of both the current and previous frames as quadratic polynomials. These approximations are derived using Taylor series expansion, which provides a way to represent the pixel intensities in a local neighborhood. By matching these polynomial expansions between two consecutive frames, the algorithm computes the displacement fields that indicate where each pixel in the first frame moves in the second frame [17].

The Farneback method has several important features. It produces dense optical flow, calculating motion for every pixel rather than only for specific points or features within an image. It is computationally efficient, making it suitable for real-time applications where processing power may be limited. Furthermore, the method is relatively robust against noise due to its polynomial expansion approach, which smooths out local variations and noise [17].

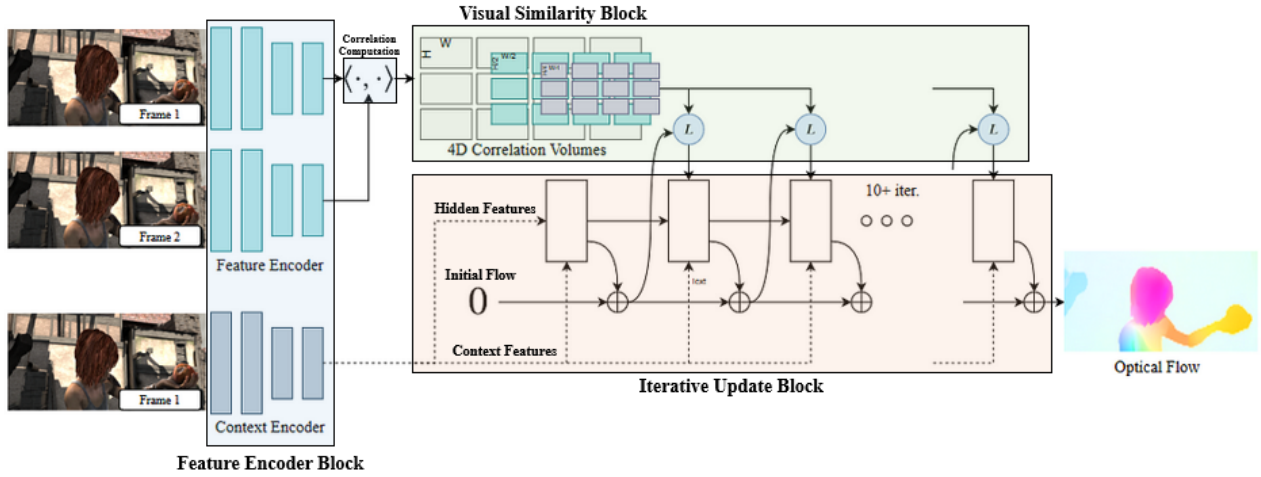

Fig. 1 RAFT architecture [18]

4. DATASET

The dataset used in this study consists of 197 lung ultrasound (LUS) videos acquired using a single ultrasound device (LOGIQ E10, GE Healthcare) at the Clinic of Thoracic Surgery, Jessenius Faculty of Medicine in Martin. All data were obtained in a controlled clinical environment, ensuring consistent imaging parameters across the samples. While this guarantees homogeneity in image quality, it also introduces the potential risk of model overfitting to device-specific features, which should be addressed in future multi-source datasets.

Each video clip was captured at a frame rate of 30 frames per second with a resolution of 640×480 pixels, and an average duration of 4 to 6 seconds. Videos were annotated by clinical experts to indicate whether lung sliding was present or absent in each case.

To avoid data leakage and ensure generalizability across patients, a patient-wise train-test split was implemented, such that no patient's data appeared in both sets.

Figure 2 shows a representative frame from a lung ultrasound sequence, illustrating the pleural line and surrounding tissue structures that are key to assessing lung sliding.

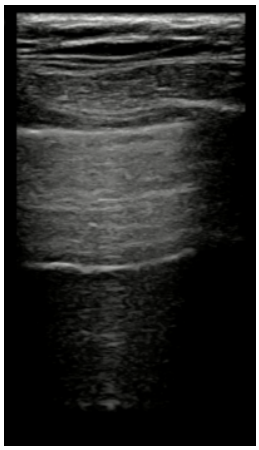

Fig. 2 Example frame of lung ultrasound image

Preprocessing and Flow Generation

Optical flow fields were computed using the Farneback algorithm from OpenCV, applied between every fifth frame to balance motion sensitivity with computational efficiency. Importantly, optical flow was extracted only from annotated segments in which either lung sliding or its absence was clearly visible, as identified by clinical reviewers. This targeted approach ensured that training data corresponded to diagnostically relevant motion.

Class Imbalance Handling

The original dataset consisted of 152 videos labeled as "lung sliding present" and 45 labeled as "absent." To address this imbalance, synthetic oversampling of the minority class was performed by replicating and augmenting the "absent" class segments using spatial transformations. These included random cropping, ±10% affine tilt, and vertical flipping. This strategy equalized the class distribution and helped prevent model bias toward the majority class.

Data Augmentation and Storage

Data augmentation was applied during preprocessing to improve generalization. The transformations included random cropping to emphasize different anatomical regions, an affine tilt of up to 10% to simulate varied probe angles, and vertical flipping to mimic real-world variability in probe positioning. Processed flow representations were saved in both TIFF and PNG formats. TIFF files were used to preserve high-fidelity data for archival and detailed analysis, while PNG files supported fast loading during model training.

Optical Flow Representations

As illustrated in Figures 3, 4, and 5, multiple visualization strategies were applied to interpret motion. Figure 3 shows a flow field generated using the pretrained

RAFT model, highlighting pleural movement. Figure 4 adds directional encoding via sine and cosine components, while Figure 5 emphasizes magnitude variation through color intensity.

5. EXPERIMENTS AND RESULTS

This section presents the experimental evaluation and results of various machine learning and deep learning approaches applied to the classification of lung sliding using optical flow data derived from thoracic ultrasound sequences. We detail the preprocessing steps used to extract meaningful motion representations and describe the implementation of both traditional machine learning classifiers and convolutional neural networks. The performance of each method is reported and analyzed to identify the most effective approach for this task.

In this experimental setup, the optical flow was represented using three channels: magnitude, angle, and gradient [19]. The magnitude channel was clipped and normalized to the range [0, 1] and stored as 32-bit floats. The angle channel was clipped between $-\pi$ and π and also stored as 32-bit floats. Gradients were computed using a Sobel operator (kernel size = 5) applied to the magnitude channel to emphasize edge information related to motion discontinuities.

5.1. Machine Learning Methods

We evaluated several traditional machine learning methods for classifying optical flow data extracted from lung ultrasound videos. These models served as baseline approaches for comparison with more advanced deep learning architectures.

The optical flow images were preprocessed by converting the flow magnitude maps into grayscale images. Each frame was then flattened into a one-dimensional feature vector representing pixel-wise motion intensity. While this method discards spatial structure, it allowed us to test whether general statistical patterns in motion could support reliable classification.

The best-performing classifiers were the Gradient Boosting Classifier (GBC) and the Random Forest (RF), both implemented using the scikit-learn library. These ensemble-based models were chosen for their robustness and ability to handle noisy or nonlinear data.

Hyperparameter Settings

For the RF classifier, we used 100 estimators, allowing each decision tree to grow without a maximum depth limit, so that trees expanded until all leaves were pure or contained fewer than two samples. Gini impurity was employed to evaluate the quality of splits, and a fixed random seed ensured reproducibility.

For the GBC, we configured the model with 100 boosting stages and a learning rate of 0.1 to control the contribution of each tree. Each tree was limited to a maximum depth of 3 to mitigate overfitting. The full dataset was used for each boosting iteration. A fixed random seed was also applied to maintain consistency across experiments.

Unless otherwise stated, other parameters were set to default values. No extensive grid search or cross-validation tuning was performed, as the focus was on evaluating general feasibility rather than optimal performance.

Performance and Observations

Table 1 summarizes the performance of both classifiers. The Gradient Boosting Classifier achieved an accuracy of 67.3%, with a precision of 66.2% and a recall of 74.2%. The Random Forest model achieved 64.5% accuracy, with a precision of 63.9% and a recall of 67.1%.

Table 1 Performance metrics of the two machine learning methods

Model	Accuracy	Precision	Recall
GBC	0.673	0.662	0.742
Random Forest	0.645	0.639	0.671

While these methods slightly outperformed other traditional classifiers (e.g., logistic regression, k-NN), their predictive power remained limited.

5.2. RAFT

Experiments were made by creating a dataset using images processed by trained RAFT models, known for their advanced optical flow estimation capabilities. The objective was to leverage RAFT's state-of-the-art motion detection to enhance our classifier's ability to analyze and predict accurately based on our ultrasound data.

Fig. 3 Extracted optical flow using RAFT pretrained models

However, contrary to our expectations, the results from this approach were underwhelming. The classifier trained on RAFT-generated optical flow data achieved an accuracy

of only 62%, with a precision of 59% and recall of 65%, which are notably lower compared to other optical flow methods employed in this study. These metrics indicate that despite RAFT's robustness in motion estimation, the derived features did not enable effective learning, possibly due to overfitting or insufficient domain adaptation.

As illustrated in Figure 3, the extracted optical flow fields, generated using pretrained RAFT models, effectively capture the subtle pleural movements that are indicative of lung sliding. These flow patterns visually emphasize the motion dynamics between consecutive ultrasound frames, which are often challenging to detect through raw imaging alone.

5.3. Resnet18 with magnitude, sine and cosine of angle

Experiments were conducted by constructing a dataset comprising three channels—magnitude, sine of the angle, and cosine of the angle—derived from optical flow data. Various preprocessing strategies were tested to optimize the input for a ResNet-18 model. These included min-max normalization applied to the entire dataset as well as the default normalization typically used for ResNet architectures. Additionally, different random cropping approaches were explored, focusing both on the area beneath the pleura and on the entire ultrasound frame.

Despite these preprocessing efforts, the model's performance remained limited. The ResNet-18 architecture, due to its relative depth and complexity, exhibited signs of overfitting on the limited dataset. Specifically, the model achieved a high recall of 94.5% but suffered from lower precision at 62.6%, indicating a tendency to misclassify negative samples as positive. This imbalance in performance metrics suggests that the model learned patterns too specific to the training data, failing to generalize well to unseen samples. These results imply that the dataset size and diversity were insufficient for effective training of a deep convolutional network such as ResNet-18. Consequently, either a simpler architecture or a substantially larger and more varied dataset would likely be necessary to improve generalization and classification reliability.

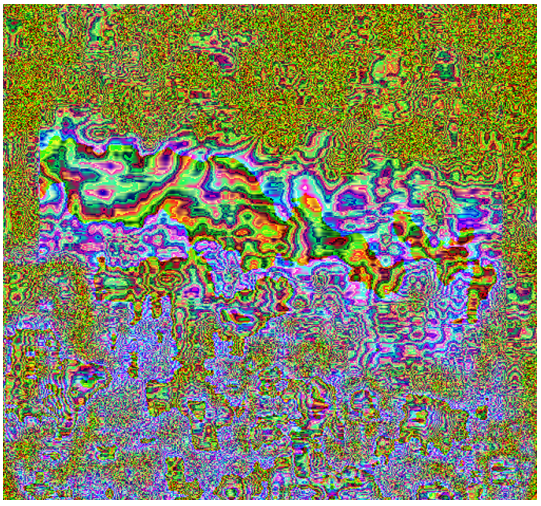

Fig. 4 Optical flow with magnitude, sine, and cosine of angles

As illustrated in Figure 4, the three-channel optical flow representation captures the magnitude and directional components (sine and cosine of the angle) of motion, providing a detailed visualization of pleural dynamics critical for lung sliding detection.

5.4. Simple CNN and Resnet18 with magnitude

In this experiment, we created a specialized dataset by extracting optical flow images using the Farneback method, focusing specifically on the magnitude component of motion. The images were processed into three channels, each representing different transformations or perspectives of the magnitude data to capture subtle variations in motion intensity. To ensure consistent scaling across the dataset, min-max normalization was applied, which scaled all values uniformly between the minimum and maximum observed values. This normalization helps improve model training stability and performance. The processed images were saved as PNG files, balancing image quality and efficient loading during training.

As illustrated in Figure 5, the three-channel magnitude representation effectively captures spatial distribution and intensity of pleural motion in ultrasound frames. Each channel encodes subtle variations, helping the model detect flow patterns linked to lung sliding. This richer representation aims to improve feature extraction, but classification results indicate that more data or improved model design may be needed for reliable diagnosis.

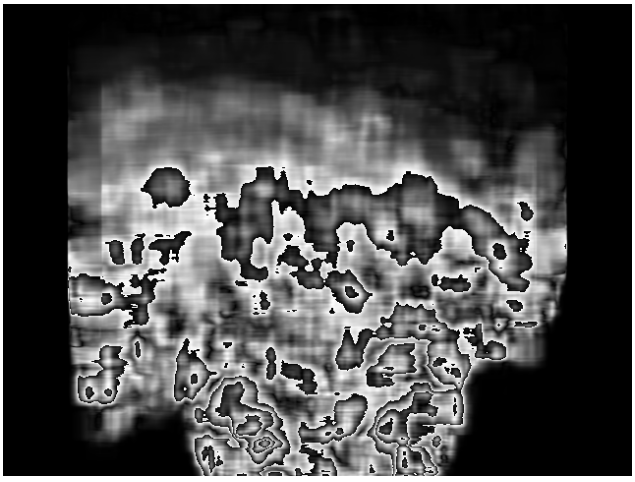

Fig. 5 Optical flow with 3 channels of magnitude

The scores presented in Table 2 confirm that the simple CNN achieves the highest overall accuracy of 80.8% and the best F1-score. In contrast, the ResNet-18 model, despite its significantly higher complexity, appears to overfit the limited dataset. This overfitting is reflected in its very high recall of 94.5%, paired with a substantially lower precision of 62.6%, suggesting that the model tends to classify many samples as positive, resulting in a high number of false positives.

Table 2 Performance metrics of the two models

Model	Accuracy	Precision	Recall	F1 Score
Simple CNN	0.808	0.704	0.772	0.737
ResNet-18	0.784	0.626	0.945	0.753

We explored two different neural network architectures to analyze our dataset of images generated using the Farneback optical flow method, focusing solely on the magnitude channel. The architectures tested were ResNet-18 and a simple Convolutional Neural Network (CNN), each evaluated for their effectiveness in processing and classifying magnitude data from optical flow.

Among these, the simple CNN, which leveraged targeted preprocessing and normalization of the magnitude data, showed the most promising results in our experiments, highlighting the benefits of focused data preparation in improving model performance for motion analysis tasks.

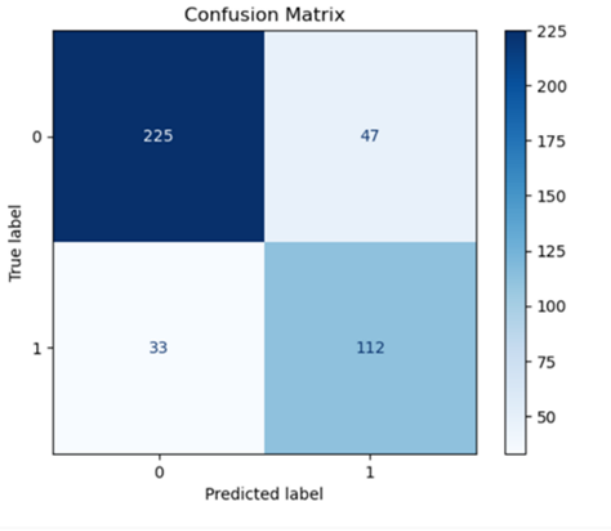

Fig. 6 Confusion matrix of simple CNN model

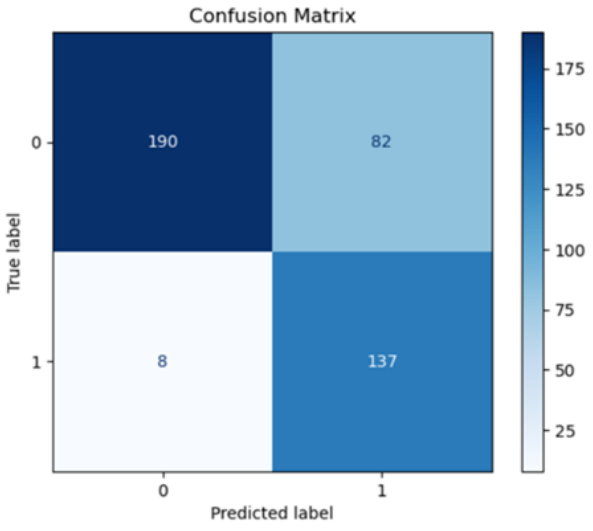

Fig. 7 Confusion matrix of ResNet-18 model

Figures 6 and 7 present the confusion matrices for the simple CNN and ResNet-18 models illustrating their performance in classifying the presence of lung sliding in ultrasound image data. Confusion matrices are presented only for the CNN and ResNet models, as these methods were selected for detailed performance analysis due to their superior results and relevance to the study objectives. For the other evaluated methods, summary performance metrics such as accuracy, precision, recall, and F1-score are reported to enable a comprehensive comparison.

6. CONCLUSION

This study addressed the challenges of classifying lung sliding in ultrasound images by applying various machine learning and deep learning techniques to optical flow representations. We evaluated both traditional classifiers—such as Gradient Boosting and Random Forest—and deep learning architectures, including a simple CNN and ResNet-18, using datasets derived primarily from the magnitude component of Farneback optical flow. Our results show that the simple CNN achieved the highest overall accuracy (80.8%), with balanced precision (70.4%) and recall (77.2%), demonstrating robust detection capabilities. Although ResNet-18 did not surpass the simple CNN in accuracy, its notably higher recall (94.5%) is a result of overfitting due to the model's complexity relative to the limited size of the dataset. This overfitting manifests as a lower precision and a higher false positive rate, reflecting the model's tendency to classify many samples as positive. Such a trade-off is critical in clinical diagnostics, where missing lung sliding (false negatives) must be minimized, but excessive false positives can also reduce reliability. Traditional machine learning methods performed reasonably but did not match the overall performance levels of the deep learning models.

The success of the deep learning models was largely dependent on careful preprocessing steps. Key factors included focused segmentation of the pleural region to isolate relevant motion, normalization of the data to standardize input values, and the targeted use of the magnitude channel from the optical flow, which effectively captured the motion characteristics critical for classification. In contrast, experiments involving pretrained RAFT models and datasets incorporating magnitude along with sine and cosine components of the flow angle produced unsatisfactory results. This highlights the importance of tailoring both the model architecture and the data representation specifically to the nature of ultrasound imaging data.

Overall, our findings demonstrate that deep learning architectures—especially simple CNNs and ResNet-18—offer a promising pathway for automated lung sliding classification. Such advancements have the potential to enhance diagnostic accuracy and reduce inter-observer variability in assessing respiratory conditions. Moreover, the diverse optical flow datasets generated in this study are made available for further research, encouraging continued development in this domain.

Finally, it is important to consider that the absence

of lung sliding may also be observed in non-pathological situations, such as during apnea. This should be taken into account when interpreting classification results in clinical settings to avoid potential misdiagnosis.

7. FUTURE WORK

In future research, we aim to compare the direction of optical flow in regions above and below the pleura. By analyzing these distinct regions, we hope to identify differences and similarities in optical flow patterns that could improve diagnostic accuracy.

Furthermore, in order to optimize the analysis for various kinds of ultrasound equipment, we intend to experiment with different subpleural region cropping percentages. We can identify the most efficient settings that improve the precision and dependability of optical flow measurements by methodically adjusting the cropped area. This strategy might result in enhanced diagnostic capacities across various ultrasonic technologies, making it possible to interpret ultrasonic images with more accuracy and dependability. In the end, these initiatives might improve ultrasound-based diagnoses in clinical settings.

ACKNOWLEDGEMENTS

In our research, the dataset was constructed using available ultrasound scans. The Jessenius Faculty of Medicine in Martin of Comenius University in Bratislava, the Clinic of Thoracic Surgery, and the Technical University of Košice are collaborating on an innovative research and development project funded by the Slovak Research and Development Agency (APVV-20-0232). The primary objective of this project is to explore and advance the application of artificial intelligence (AI) techniques in the processing and interpretation of lung ultrasonography data. This collaborative effort underscores the purpose of the available dataset, which is to support and drive forward this cutting-edge research in the field. We employed a patient-split strategy to ensure that images from the same patient were not present in both the training and testing subsets. This method helped to reduce bias and improve the generalizability of the model across different patients.

REFERENCES

[1] M. JAŠČUR, M. BUNDZEL, M. MALÍK, A. DZIAN, N. FERENČÍK, and F. BABIČ, "Detecting the absence of lung sliding in lung ultrasounds using deep learning," *Applied Sciences*, vol. 11, no. 15, p. 6976, 2021.

[2] M. L. GIGER, "Machine learning in medical imaging," *Journal of the American College of Radiology*, vol. 15, no. 3, pp. 512–520, 2018.

[3] D. SUN, S. ROTH, J. P. LEWIS, and M. J. BLACK, "Learning optical flow," in *European Conference on Computer Vision*, pp. 83–97, Springer, 2008.

[4] F. HE, T. LIU, and D. TAO, "Why resnet works? residuals generalize," *IEEE transactions on neural networks and learning systems*, vol. 31, no. 12, pp. 5349–5362, 2020.

[5] Z. LI, F. LIU, W. YANG, S. PENG, and J. ZHOU, "A survey of convolutional neural networks: analysis, applications, and prospects," *IEEE transactions on neural networks and learning systems*, vol. 33, no. 12, pp. 6999–7019, 2021.

[6] M. GAO, P. SONG, F. WANG, J. LIU, A. MANDELIS, and D. QI, "A novel deep convolutional neural network based on resnet-18 and transfer learning for detection of wood knot defects," *Journal of Sensors*, vol. 2021, no. 1, p. 4428964, 2021.

[7] J. WU, "Introduction to convolutional neural networks," *National Key Lab for Novel Software Technology. Nanjing University. China*, vol. 5, no. 23, p. 495, 2017.

[8] S. E. EBADI, D. KRISHNASWAMY, S. E. S. BOLOURI, D. ZONOOBI, R. GREINER, N. MEUSER-HERR, J. L. JAREMKO, J. KAPUR, M. NOGA, and K. PUNITHAKUMAR, "Automated detection of pneumonia in lung ultrasound using deep video classification for covid-19," *Informatics in medicine unlocked*, vol. 25, p. 100687, 2021.

[9] M. KOLARIK, M. SARNOVSKÝ, and J. PARALIČ, "Detecting the absence of lung sliding in ultrasound videos using 3d convolutional neural networks," *Acta Polytechnica Hungarica*, vol. 20, no. 6, pp. 47–60, 2023.

[10] S. ROY, W. MENAPACE, S. OEI, B. LUIJTEN, E. FINI, C. SALTORI, I. HUIJBEN, N. CHENNAKESHAVA, F. MENTO, A. SENTELLI, *et al.*, "Deep learning for classification and localization of covid-19 markers in point-of-care lung ultrasound," *IEEE transactions on medical imaging*, vol. 39, no. 8, pp. 2676–2687, 2020.

[11] A. HUANG, L. JIANG, J. ZHANG, and Q. WANG, "Attention-vgg16-unet: a novel deep learning approach for automatic segmentation of the median nerve in ultrasound images," *Quantitative imaging in medicine and surgery*, vol. 12, no. 6, p. 3138, 2022.

[12] K. ARDON-DRYER and A. TAIRU, "Impacts of african dust storm particles on human lung epithelial cells," in *American Meteorological Society Meeting Abstracts*, vol. 101, p. 31, 2021.

[13] S. S. BEAUCHEMIN and J. L. BARRON, "The computation of optical flow," *ACM computing surveys (CSUR)*, vol. 27, no. 3, pp. 433–466, 1995.

[14] A. BRUHN, J. WEICKERT, and C. SCHNÖRR, "Lucas/kanade meets horn/schunck: Combining local and global optic flow methods," *International journal of computer vision*, vol. 61, pp. 211–231, 2005.

[15] Z. TEED and J. DENG, "Raft: Recurrent all-pairs field transforms for optical flow," in *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pp. 402–419, Springer, 2020.

[16] Z. TEED and J. DENG, "Raft: Recurrent all-pairs field transforms for optical flow," in *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pp. 402–419, Springer, 2020.

[17] G. FERNEBÄCK, "Two-frame motion estimation based on polynomial expansion," in *Image Analysis: 13th Scandinavian Conference, SCIA 2003 Halmstad, Sweden, June 29–July 2, 2003 Proceedings 13*, pp. 363–370, Springer, 2003.

[18] S. MALLICK, "Optical flow using deep learning: Raft," 2023. Accessed: 2024-05-14.

[19] G. MASON-WILLIAMS and F. DAHLQVIST, "What makes a good prune? maximal unstructured pruning for maximal cosine similarity," in *The Twelfth International Conference on Learning Representations*, 2024.

Received June 26, 2025, accepted July 17, 2025

BIOGRAPHIES

Maroš Hliboký is a PhD student at the Department of Cybernetics and Artificial Intelligence at the Technical University of Košice. He obtained his Bachelor's degree in 2019 and his Master's degree in 2021. His Bachelor's thesis focused on the development of the Rehapiano device, while his Master's thesis addressed the development of a convolutional neural network for semantic segmentation of aerial survey LiDAR images. His current research is dedicated to the application of deep neural networks in the medical domain, with a primary focus on image data—particularly lung ultrasound imaging.

Katarína Ištoňová completed her Master's degree at the Department of Cybernetics and Artificial Intelligence at the Technical University of Košice in 2025, with a focus on the implementation of deep neural networks for medical image data. Her Bachelor's thesis focused on the detection of pleural sliding in ultrasound data.

Laura Pituková is a Master's student at the Department of Cybernetics and Artificial Intelligence at the Technical University of Košice. Her research focuses on deep neural networks, capsule networks, and Kolmogorov–Arnold networks, primarily applied in the medical domain.

Marek Bundzel is an Associate Professor at the Department of Cybernetics and Artificial Intelligence at the Technical University of Košice, where he also earned both his Master's degree and PhD. His research interests lie in biologically inspired methods of artificial intelligence, with a primary focus on artificial neural networks and optimization techniques applied to computer vision and adaptive rehabilitation.