https://doi.org/10.31891/csit-2025-3-11
UDC 004.8

Nataliia MELNYKOVA, Anna VECHIRSKA
Lviv Polytechnic National University

ANALYSIS OF OPTIMIZER AND HYPERPARAMETER INFLUENCE ON YOLO IN THERMAL LANDMINE DETECTION

This paper investigates the impact of optimizer choice and hyperparameter tuning on the performance of the YOLO deep learning model for landmine detection in thermal images. The aim of this work is to study the effect of different optimizers and parameter configurations on model accuracy and training stability. The object of the study is the process of detecting landmines in thermal imagery using deep neural networks.

A dataset of thermal landmine images annotated in YOLO format was used for training. The experiments were conducted with the YOLOv11n architecture initialized with pre-trained weights. The varied parameters included the optimizer (SGD or Adam), learning rate, and batch size. Each model was trained for 50 epochs, and performance was evaluated using mAP, precision, and recall metrics.

The study provides a comparative analysis of the influence of Adam and SGD optimizers on the accuracy and stability of YOLO when trained on a limited dataset of thermal landmine images. The results suggest that, given appropriate configuration, SGD is capable of achieving performance competitive with adaptive methods, despite their popularity. The experiments also confirm the feasibility of achieving high detection accuracy even with a relatively small dataset.

All configurations achieved high mAP values. The Adam optimizer enabled a faster initial reduction in loss functions, whereas SGD provided smoother and more stable training dynamics. The highest precision and recall were obtained in the experiment with SGD at a learning rate of 0.01 and batch size of 64, making this configuration the most promising for further research.

The findings on optimizer and hyperparameter selection can be applied to improve the efficiency of automated thermal image analysis systems based on unmanned aerial vehicles, contributing to safer and faster detection of explosive hazards.

Keywords: object detection, deep learning, landmine detection, YOLO, thermal imaging, optimizers

Наталія МЕЛЬНИКОВА, Анна ВЕЧІРСЬКА
Національний університет «Львівська політехніка»

АНАЛІЗ ВПЛИВУ ОПТИМІЗАТОРІВ І ГІПЕРПАРАМЕТРІВ НА ПРОДУКТИВНІСТЬ YOLO У ЗАДАЧІ ВИЯВЛЕННЯ МІН НА ТЕПЛОВИХ ЗОБРАЖЕННЯХ

У цій роботі досліджується вплив вибору оптимізатора та налаштування гіперпараметрів на продуктивність моделі глибокого навчання YOLO для виявлення мін на теплових зображеннях. Метою роботи є вивчення впливу різних оптимізаторів і конфігурацій параметрів на точність моделі та стабільність навчання. Об'єктом дослідження є процес виявлення мін на теплових зображеннях за допомогою глибоких нейронних мереж.

Для навчання використано набір даних тепловізійних зображень мін, анотований у форматі YOLO. Експерименти проводилися на архітектурі YOLOv11n із попередньо навченими вагами. Варіювалися оптимізатори SGD або Adam, швидкість навчання та розмір пакету. Моделі навчалися протягом 50 епох, а якість оцінювалася за метриками mAP, precision та recall.

Робота надає порівняльний аналіз впливу оптимізаторів Adam та SGD на точність і стабільність роботи YOLO в умовах обмеженого набору тепловізійних даних мін. Показано, що при належному налаштуванні SGD може забезпечувати перспективні результати, незважаючи на популярність адаптивних методів. Також в роботі підтверджуються можливість високоточного виявлення мін навіть за обмеженого обсягу даних.

Усі конфігурації досягли високих значень mAP. Оптимізатор Adam забезпечив швидше початкове зниження функції втрат, тоді як SGD продемонстрував більш плавну та стабільну динаміку навчання. Найвищий precision і recall показав експеримент із використанням SGD, швидкістю навчання 0.01 та розміром пакету 64, що робить цю конфігурацію перспективною для подальших досліджень.

Отримані результати щодо вибору оптимізатора та гіперпараметрів можуть бути використані для підвищення ефективності автоматизованих систем аналізу тепловізійних зображень з безпілотних літальних апаратів, сприяючи безпечнішому та швидшому виявленню вибухонебезпечних предметів.

Ключові слова: ідентифікація об'єктів, глибинне навчання, розпізнавання мін, YOLO, теплові зображення, оптимізатори

Received / Стаття надійшла до редакції 23.08.2025
Accepted / Прийнята до друку 19.09.2025

Introduction

Unexploded ordnance (UXO) and landmines are among the most hazardous remnants of armed conflicts, posing a persistent and serious threat to civilian populations. These explosive hazards may remain either on the surface or buried underground in diverse terrain conditions, making humanitarian demining operations complex and resource intensive. Such missions require substantial time, specialized equipment, and trained personnel, who are often exposed to significant risk. Automating the detection of explosive hazards is an important research area because it can make humanitarian demining faster and safer for the people involved.

In recent years, advances in computer vision and deep learning have enabled the development of highly accurate object detection systems. Convolutional neural networks have demonstrated exceptional capability in automatically analysing large volumes of imagery and detecting objects based on learned feature representations. Among most popular object detection frameworks, the YOLO (You Only Look Once) family of models stands out for its ability to achieve high detection accuracy while maintaining real-time processing speeds. These characteristics make YOLO a suitable candidate for aerial and unmanned aerial vehicle-based detection tasks, including thermal imagery analysis [1-4].

Despite the progress in computer vision and deep learning, the detection of UXO and landmines in thermal imagery remains challenging due to variations in environmental conditions and limited annotated datasets.

The performance of YOLO is heavily influenced by its hyperparameters and the choice of optimization algorithm during training. Hyperparameters such as the learning rate and batch size can significantly impact both convergence speed and final detection accuracy. Similarly, the optimizer governs how the network's weights are updated, affecting the stability and generalization ability of the trained model.

This study aims to analyze how different optimizers and hyperparameter configurations affect YOLO's performance on thermal images of landmines, including those located on the surface or buried at varying depths.

Related works

Several studies [1–4] have explored various approaches to detecting landmines, including traditional computer vision techniques, machine learning models, and deep neural networks. R-CNN and Faster R-CNN have been widely applied in such tasks [1, 3], while in recent years YOLO has been gaining increasing popularity for object detection in similar contexts [2, 4].

While works [1–4] address the general landmine detection problem, the detection of buried mines or mines in areas with dense vegetation remains a significant challenge. In object detection tasks, the choice of optimizer plays a key role in determining the convergence speed and final accuracy of the model. Therefore, conducting an analysis of different optimizers can offer valuable insights into improving detection accuracy and robustness in the context of landmine identification.

Analysis of optimizer and hyperparameter influence on YOLO in thermal landmine detection

In this study, the YOLO (You Only Look Once) deep learning model was employed. YOLO belongs to the category of single-stage object detectors and has demonstrated high effectiveness in image-based detection tasks, including thermal imagery. Unlike two-stage methods such as Faster R-CNN, YOLO processes the input image in a single step, predicting both bounding box coordinates and object classes simultaneously. This approach ensures high throughput and minimizes delays in dynamic environments. Furthermore, YOLO performs a global analysis of the entire image, allowing it to account for contextual information and reducing the likelihood of false detections caused by local artifacts [2, 4, 5-7].

During the training of neural networks, the optimizer plays a crucial role in adjusting the model's weights to minimize the loss function. It defines the strategy for updating parameters to progressively approach a global or local minimum of the loss surface. Two of the most widely used optimizers in YOLO-based models are Stochastic Gradient Descent (SGD) and Adam.

SGD is a classical optimization method that updates model weights using only a subset (mini-batch) of the training data at each step, making it more computationally efficient than full-batch gradient descent. To improve convergence, SGD is often combined with momentum, which helps smooth parameter updates and maintain a consistent search direction [8].

Adam, on the other hand, is an adaptive optimization algorithm. It automatically adjusts the learning rate for each model parameter, considering both the current gradient values and their historical changes. Adam computes updates based on the exponentially weighted moving average of past gradients (capturing direction) and the exponentially weighted moving average of squared gradients (capturing scale), enabling more stable and often faster convergence in the early training stages [8, 9].

Due to its adaptive learning rate adjustment, Adam often achieves faster convergence, particularly in the early stages of training, and tends to be less sensitive to hyperparameter selection. However, it can underperform compared to SGD in terms of generalization ability [8, 9].

The performance of the models in this study was assessed using the following key metrics: intersection over union (IoU), mean average precision (mAP), precision, and recall.

Experiments

For training, the Landmines Detection Dataset was used, obtained via the Roboflow platform [10]. The dataset consists of thermal images, annotated to indicate the presence of landmines. It is organized into three subsets – train, validation, and test – each containing both images and corresponding annotations in YOLO format (text files specifying bounding box coordinates). The training set comprises 836 images, of which 77 are background (containing no objects). The validation set contains 240 images, 30 of which are background, and the test set includes 118 images, 7 of which are background. All images have a resolution of 740x480 pixels.

The model was trained using the pre-trained YOLOv11n architecture [11], with training performed in the Google Colab environment on a GPU.

During the experiments, the following parameters were varied:

– optimizer (SGD or Adam);

– batch size – the number of samples processed simultaneously before updating the weights, influencing both the stability and the speed of training;

– learning rate – the step size used when updating the weights, affecting the convergence behaviour of the model.

Several experiments were conducted, and the parameters of the five configurations that achieved the best results are presented in Table 1.

Table 1

Experiment parameters

№	Optimizer	Learning rate	Batch size
1	SGD	0.01	32
2	SGD	0.005	32
3	SGD	0.01	64
4	Adam	0.001	64
5	Adam	0.001	32

Each model was trained for 50 epochs, with Box loss and Distribution focal loss monitored on both the training and validation datasets [12].

In experiments using the Adam optimizer, both loss functions started with a sharper decrease compared to SGD. Example is shown on Fig.1.

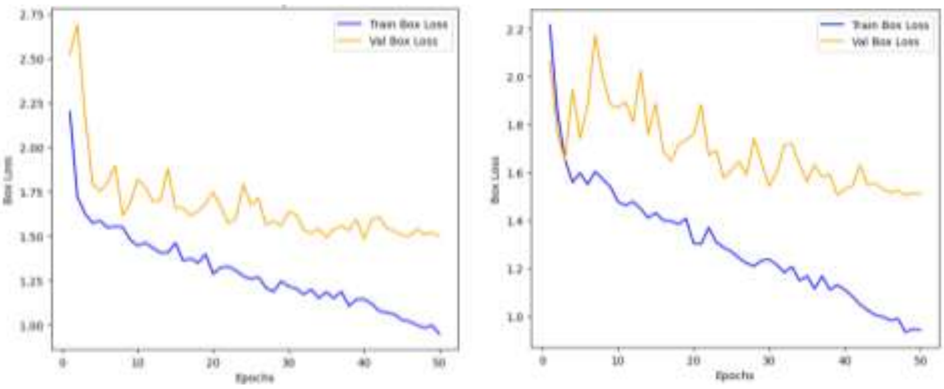

Fig.1. Box loss function for different experiments: a) experiment 4; b) experiment 3

Among the SGD runs, experiment 1 showed the most fluctuations in Box Loss, whereas experiment 2 exhibited the most stable and uniform decrease in loss. For the experiments with the Adam optimizer (experiments 4 and 5), the loss decreased rapidly during the initial epochs, followed by a more gradual and stable reduction after around the 10th epoch. Overall, these trends indicate that Adam enables faster initial convergence, while SGD tends to produce smoother and more consistent loss trajectories across epochs.

A certain degree of overfitting was observed across all loss curves, reflected in the gap between the training and validation loss curves. This effect is likely due to the relatively small size of the training dataset.

For each experiment, the performance metrics were recorded, including mAP@0.5 (mean average precision at an intersection over union threshold of 0.5), mAP@0.5:0.95 (mean average precision averaged across IoU thresholds from 0.5 to 0.95 in steps of 0.05), precision, and recall. The results are summarized in Table 2.

Table 2

Model performance evaluation results

№	mAP@0.5	mAP@0.5:0.95	Precision	Recall
1	0.978	0.631	0.950	0.938
2	0.980	0.614	0.953	0.926
3	0.981	0.641	0.970	0.950
4	0.984	0.647	0.932	0.950
5	0.976	0.638	0.949	0.942

All models demonstrated high mAP@0.5 scores (ranging from 0.976 to 0.984), indicating the overall effectiveness of the detector. The highest generalized precision across varying IoU thresholds was achieved in experiment 4, suggesting that this model most accurately localized objects. The highest precision value was observed in experiment 3, while the highest recall was achieved by experiments 3 and 4. The results of experiments

3 (SGD optimizer, learning rate 0.01, batch size 64) and 4 (Adam optimizer, learning rate 0.001, batch size 64) appear the most balanced. Considering the small difference in mAP but a more notable advantage in precision for experiment 3, further work could reasonably focus on this configuration.

Conclusions

This study investigated the impact of hyperparameters and optimizer choice on the performance of the YOLO deep learning model in the task of detecting landmines in thermal imagery.

The experiments showed that the Adam optimizer provides a faster initial decrease in loss functions, allowing the model to reach acceptable performance within the early epochs, although this can be accompanied by reduced stability. In contrast, the SGD optimizer yields a more gradual loss reduction and exhibits greater sensitivity to the choice of learning rate. The selection of hyperparameters, particularly the learning rate and batch size, was also found to have a substantial influence on model performance. Despite the popularity of Adam as an adaptive optimizer, the results indicate that with proper tuning, SGD can deliver competitive results.

Performance evaluation based on precision, recall, and mAP metrics demonstrated that high detection accuracy can be achieved through careful hyperparameter tuning and optimizer selection. The experiments achieved mAP@0.5 values of up to 0.984, mAP@0.5:0.95 up to 0.647, precision up to 0.970, and recall up to 0.950. These findings confirm that even with a limited dataset, it is possible to achieve an acceptable level of accuracy. However, further improvements could be obtained by expanding the dataset and incorporating regularization and data augmentation techniques.

References

1. Applying Deep Learning to Automate UAV-Based Detection of Scatterable Landmines / J. Baur et al. *Remote Sensing*. 2020. Vol. 12, no. 5. P. 859.
2. Bajić M., Potočnik B. UAV Thermal Imaging for Unexploded Ordnance Detection by Using Deep Learning. *Remote Sensing*. 2023. Т. 15, № 4. С. 967.
3. Vivoli E., Bertini M., Capineri L. Deep Learning-Based Real-Time Detection of Surface Landmines Using Optical Imaging. *Remote Sensing*. 2024. Т. 16, № 4. С. 677.
4. Soeleman M. A., Supriyanto C., Purwanto. Deep Learning Model for Unmanned Aerial Vehicle-based Object Detection on Thermal Images. Revue d'Intelligence Artificielle. 2023. Т. 37, № 6. С. 1441–1447.
5. You Only Look Once: Unified, Real-Time Object Detection / J. Redmon та ін. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, м. Las Vegas, NV, USA, 27–30 черв. 2016 р.
6. Khanam R., Hussain M. YOLOv11: An Overview of the Key Architectural Enhancements. arXiv.org. URL: https://arxiv.org/abs/2410.17725 (дата звернення: 10.08.2025).
7. Ultralytics. YOLO11 NEW. Home - Ultralytics YOLO Docs. URL: https://docs.ultralytics.com/models/yolo11 (дата звернення: 11.08.2025).
8. Ruder S. An overview of gradient descent optimization algorithms. arXiv.org. URL: https://arxiv.org/abs/1609.04747 (дата звернення: 12.08.2025).
9. Kingma D. P., Ba J. Adam: A Method for Stochastic Optimization. arXiv.org. URL: https://arxiv.org/abs/1412.6980 (дата звернення: 12.08.2025).
10. Landmines detection dataset Object Detection Model by Northumbria. Roboflow. URL: https://universe.roboflow.com/northumbria/landmines-detection-dataset (дата звернення: 10.08.2025).
11. Ultralytics. YOLO Performance Metrics. Home - Ultralytics YOLO Docs. URL: https://docs.ultralytics.com/guides/yolo-performance-metrics/#how-can-validation-metrics-from-yolo11-help-improve-model-performance (дата звернення: 13.08.2025).
12. What is DFL Loss in YOLOv8? | Understanding DFL Metrics. YOLOv8. URL: https://yolov8.org/what-is-dfl-loss-in-yolov8/ (дата звернення: 13.08.2025).

Natalia Melnykova **Наталія Мельникова**	DEng, Professor, Head of the Department of Artificial Intelligence Systems, Institute of Computer Science and Information Technologies, Lviv Polytechnic National University, Lviv, Ukraine. e-mail: melnykovanatalia@gmail.com https://orcid.org/0000-0002-2114-3436	доктор технічних наук, професор, зав. кафедри систем штучного інтелекту, Інститут комп'ютерних наук та інформаційних технологій, Національний університет «Львівська політехніка», Львів, Україна.
Anna Vechirska **Анна Вечірська**	student at the Department of Artificial Intelligence Systems, Institute of Computer Science and Information Technologies, Lviv Polytechnic National University, Lviv, Ukraine. e-mail: annvechirska@gmail.com https://orcid.org/0000-0002-3886-3026	студентка кафедри систем штучного інтелекту, Інститут комп'ютерних наук та інформаційних технологій, Національний університет «Львівська політехніка», Львів, Україна.