

Enhancing Detection of *Leishmania* spp. Amastigotes in Canine Lymph Node Smear Images: Evaluating the Effectiveness of Synthetic Data in Augmenting Existing Datasets

Dimitrios Tsikos¹, Irene Chatzipanagiotidou², Dimitris Dimitriadis(✉)¹[0000-0002-9404-0331], Constantina N Tsokana¹[0000-0003-4180-6310], George Valiakos²[0000-0002-3869-5026], Labrini V Athanasiou, and Grigorios Tsoumakas¹[0000-0002-7879-669X]

¹ Aristotle University of Thessaloniki, Greece

dimtsikos@gmail.com, {dndimitri,greg}@csd.auth.gr, ctsokana@vet.auth.gr

² University of Thessaly, Greece

irenechatzi1@hotmail.com, georgevaliakos@uth.gr, lathan@vet.uth.gr

Abstract. Leishmaniosis is a parasitic mammalian disease that severely affects humans and dogs. Early diagnosis is crucial and associated with improved prognosis and treatment outcomes. A key diagnostic component is the detection of *Leishmania* amastigotes, the etiological agent of the disease, in cytologic preparations via microscopy. However, reliance on operator expertise limits its accessibility in veterinary clinics. Deep learning offers a promising approach for automating *Leishmania* amastigote detection, yet data limitations and the time-consuming, error-prone nature of real data annotation process remain significant challenges. This study explores the use of synthetic data to address these challenges and improve deep learning performance in detecting *Leishmania* amastigotes in microscopic images from canine lymph node aspirates. We propose an automated, two-stage synthetic data generation approach. First, structured representations of healthy and infected cells are created based on real microscopy data, incorporating randomized morphological features and material properties to mimic optical characteristics. Then, these elements are assembled into composite images with controlled variations in spatial arrangement, lighting, and perspective to enhance dataset diversity. The final output is annotated images designed for training object detection models. By supplementing real datasets with synthetic images, we address data scarcity and imbalance issues, improving model accuracy and generalization. Our results show that incorporating synthetic data significantly enhances deep learning models' ability to detect *Leishmania* amastigotes, offering a promising solution for veterinary diagnostics. Additionally, we introduce a new dataset that combines both original and synthetic data, contributing to further research into this important zoonotic disease.

Keywords: *Leishmania*, canine lymph node smears, synthetic data, object detection, deep learning, artificial intelligence

1 Introduction

Leishmaniosis is a vector-borne disease caused by protozoan parasites of the genus *Leishmania*. Among the 30 identified species [1], *Leishmania infantum* is the most widespread [2] and is responsible for Visceral Leishmaniasis (VL) in humans and Canine Leishmaniosis (CanL) in dogs, both of which can be fatal if left untreated³ [3]. The World Health Organization (WHO) reports that human Leishmaniasis occurs in over 90 countries, with an estimated 50,000 to 90,000 new VL cases annually, though only 25–45% are officially reported⁴.

Transmission occurs through the bite of infected female sandflies during blood feeding from their hosts [4]. Following the sandfly bite, metacyclic promastigotes are phagocytized by the host's immune cells, where they transform into amastigotes and rapidly multiply, leading to cell lysis. Once released, *Leishmania* amastigotes can invade new cells [5]. The outcome of the infection depends on both host factors and the virulence of the parasite [5,6,7].

Early and accurate diagnosis is crucial for improving prognosis and preventing further parasite dissemination within the host and its environment [8]. In dogs, treatment efficacy varies depending on symptom severity, parasitic load, and the host's immune response. While treatment aims to control clinical signs and improve the dog's quality of life, complete elimination of the parasite is rare, and the risk of relapse remains substantial [5,6,7,3]. Therefore, continuous monitoring of dogs affected by CanL is essential.

Diagnosing CanL is challenging and involves a combination of clinical examination, clinicopathological testing, cytology, serology, and molecular diagnostic techniques [9]. Parasitological diagnosis provides a conclusive identification of infection by directly observing amastigotes in cytologic preparations [10] from lymphoid organs (like bone marrow, lymph nodes, and spleen), liver, and skin, which typically harbor a high parasitic load [11,12].

Cytologic examination of stained lymph node smears is commonly used in routine laboratory settings due to the less invasive nature of lymph node aspiration compared to bone marrow and spleen biopsies [11]. However, the specificity of lymph node smear microscopy heavily relies on the operator's expertise, with high sensitivity achieved when a minimum of 100, ideally 1,000, oil immersion fields of high-quality smears are examined [12]. Moreover, the repeatability and reproducibility of this method remain uncertain [10]. Therefore, this method is infrequently performed in veterinary clinics due to a lack of experience among general practitioners in identifying *Leishmania* amastigotes via microscopy and the time constraints associated with examining the required number of fields.

To address these challenges, automating the detection of *Leishmania* amastigotes in microscopic images without the need for specialized equipment or extensive expertise is crucial. Diagnostic tools that enable accurate and efficient analysis of lymph node aspirate images would provide general practitioners with a useful tool for diagnosing *Leishmania* infection in dogs, ultimately saving time

³ <https://www.cdc.gov/parasites/leishmaniasis/epi.html>

⁴ <https://www.who.int/news-room/fact-sheets/detail/leishmaniasis>

and enhancing diagnostic capabilities for more effective monitoring of treatment outcomes and disease progression.

Deep learning models have the potential to enhance the diagnosis of *Leishmania*. Several studies have demonstrated promising results in both humans and animals using deep learning techniques [14,17]. However, a major challenge remains: the limited availability of data and the time-consuming process of acquiring it. In this paper, we address this issue by introducing a synthetic data generation technique that has not been previously applied to *Leishmania* diagnosis. Our results, obtained using a state-of-the-art deep learning model for object detection to identify *Leishmania* amastigotes, show that incorporating synthetic data improves performance compared to using only real images.

The remainder of this paper is structured as follows. Section 2, provides an overview of related work. Section 3 describes our dataset, detailing the image acquisition process and preprocessing steps. Section 4 presents the methodology for generating synthetic data. In Section 5, we outline the experimental setup, while Section 6 discusses the results. Finally, Section 7 concludes the paper.

2 Related Work

Early methods relied on classical image processing techniques, such as Difference of Gaussians (DoG) filtering for segmentation [18], which provided automation but lacked robustness due to variations in staining and imaging conditions [19]. As research progressed, ML emerged as a viable alternative.

In 2018, a U-Net model was introduced to segment *Leishmania* parasites in microscopy images [20]. The dataset, containing 45 images of *Leishmania* infected macrophages, faced class imbalance, which was mitigated through augmentation strategies. However, model generalization was constrained by the dataset’s limited size. By 2022, the Viola-Jones algorithm was implemented for binary classification of infected and non-infected cells [21]. The dataset comprised 300 images from 50 slides prepared from skin scrapings obtained from human patients suspected for Leishmaniasis. Although, their approach offered computational efficiency, precision and recall remained suboptimal compared to state-of-the-art deep learning models. Similarly, ML classifiers was employed, including K-Nearest Neighbors, Naïve Bayes, Support Vector Machines, and Logistic Regression, to diagnose CanL [22]. Using tabular data from 340 canine cases, which included physical examination records and serological test results, their best model—logistic regression—achieved an accuracy of 75%.

Deep learning has emerged as a more effective approach for *Leishmania* parasites detection in comparison with traditional ML methods. In bibliography the two main tasks applied in *Leishmanias* parasites detection problem is classification, which determines whether an image contains infected cells, and object detection, which identifies and localizes parasites within an image. While classification is computationally simpler and widely used, object detection presents a more challenging problem, requiring accurate spatial localization alongside classification.

Classification-based approaches have demonstrated strong results. CNN-based feature extraction and color space transformations were applied to detect visceral *Leishmania* amastigotes, using a dataset of 150 images obtained from bone marrow slides from human patients [23]. The model improved detection accuracy but struggled with false positives. Similarly, pre-trained CNNs such as ResNet and InceptionV3 were used to classify microscopic images from bone marrow smears of human patients as positive or negative for the presence of *Leishmania* amastigotes using a dataset consisting of 150 microscopic images [24]. In order to overcome the obstacle of limited data, data augmentation techniques-like contrast adjustments and flipping- were applied. The techniques led to a significant improvement in classification accuracy. In 2024, the LeishFuNet was introduced, a deep learning model that achieved the impressive 98.95% accuracy and 98,92% F1-score in detecting *Leishmania* amastigotes using a dataset of 239 Giemsa-stained microscopic images [25]. In this case transfer learning was used to address data scarcity issues, making it a scalable solution for leishmaniosis diagnostics.

Object detection methods, while more complex, provide additional information for parasite localization. [26] focused on automating parasite detection and counting by applying segmentation techniques such as the Otsu method and morphological operations, using a dataset of manually labeled Giemsa-stained microscopy images. Their method excelled in recall, effectively identifying intracellular parasites, although precision remained a challenge due to false positives in heavily stained regions. YOLOv5 and Faster R-CNN were fine-tuned on a dataset of 1858 manually labeled images, including 244 Cutaneous Leishmaniasis cases, 68 Visceral Leishmaniasis cases, and 1420 monocyte samples [14]. The model achieved a mean average precision (mAP) of 73%, with a precision of 68% and recall of 69%, marking a significant improvement over traditional segmentation methods. Augmentation strategies were crucial in enhancing model robustness and performance, especially for detecting parasites in low-contrast images.

The limited availability of labeled datasets is a critical challenge in Leishmaniosis research. The main causes behind that are the difficulties in obtaining well-annotated microscopy images, interlaboratory staining variations, and ethical concerns surrounding data sharing. In order to overcome this data scarcity, many data augmentation techniques are employed.

Basic augmentation pipelines, which contain rotation flips, zoom transformations and contrast modifications, have proven effective in increasing dataset diversity. [14] used a combination of rotation, flipping, scaling, translation, and shearing to create diverse representations of *Leishmania* parasites in Giemsa-stained images, ensuring that their YOLOv5-based detection model could perform well under different staining conditions. [26] applied contrast modifications and color space transformations by converting images to the hue, saturation, and intensity (HSI) model, improving the distinction between parasites and background noise. These augmentation strategies played a crucial role in mitigating

the risks of overfitting and ensuring that models could effectively handle real-world variations in microscopic imaging.

Lack of data is a regular challenge across all sectors that tend to adapt machine learning solutions. However, in recent years, there has been a growing shift toward synthetic data to further bridge the gap. Techniques like generative adversarial networks and real-world simulations are now being leveraged to create diverse, high-quality training data, improving model performance in scenarios where acquiring real data is costly or impractical. [27] explored the use of synthetic datasets in deep learning for computer-vision-assisted manufacturing tasks, demonstrating how artificially generated data can improve model accuracy and robustness. Similarly, [28] conducted a comprehensive survey on the creation and use of synthetic data in computer vision and medical imaging. They highlighted that synthetic data could address challenges related to limited patient populations, inconsistent data quality, and imbalanced disease stage distributions. By incorporating synthetic data, the models achieved improved accuracy and generalizability, effectively mitigating biases introduced during data collection and improving the objectivity and consistency of medical imaging applications.

In line with these directions and current trends, we integrate deep learning and synthetic data generation to detect *Leishmania* amastigotes, an approach that has not been previously explored. The promising results of the newly generated images, combined with the high accuracy of deep learning techniques, offer fresh insights to address this challenge and similar issues within the research community. Furthermore, the annotated dataset accompanying this work, along with the open-source synthetic data generation code, serves as a valuable resource for further study.

3 The Dataset

In this section, we present the dataset used in our study. Specifically, we detail the acquisition phase, during which we captured microscopic images of lymph node samples from dogs with canine leishmaniosis suspected of leishmaniosis in laboratory settings. We then describe the pre-processing phase, where the dataset was prepared for the object detection task. In addition, we introduce a utility dataset that we explored to enhance the performance of our models.

3.1 Acquisition

The lymph node smears used in this study were archived materials from previous studies conducted in the Diagnostic Laboratory, Clinic of Medicine, Faculty of Veterinary Medicine, School of Health Sciences, University of Thessaly. The preparation of the lymph node smears followed a specific protocol. A non-aspiration fine-needle biopsy technique was performed using a 21-gauge needle attached to a 10-mL syringe, targeting either the prescapular or popliteal lymph nodes of dogs. After collection, the material was placed on 76 x 26 mm glass

slides for microscopy and the overlapping content was smeared using the squash method.

The lymph node smears were air-dried, fixed in methyl alcohol, and stained with Giemsa. We conducted a light microscopy examination of the lymph node smears using an optical microscope (OLYMPUS model BX4, Olympus, Germany) with a $1000\times$ magnification lens. Microscopic examination included 10 to 1,000 oil immersion fields (OIFs, $\times 1000$), depending on the detection and density of amastigotes. A lymph node smear was considered positive when *Leishmania* amastigotes were identified as round to oval organisms, measuring 2–5 μm in diameter, with an eccentric nucleus, a kinetoplast exhibiting more intense basophilic staining compared to the nucleus, and a visible cellular membrane. Smears that tested positive were stored for later image acquisition using an OLYMPUS model BX4 microscope (Olympus, Germany) coupled with a ZEISS Axiocam ERc 5s 5-megapixel all-in-one microscope camera (ZEISS, Germany) and a BioBlue.Lab microscope (Euromex, Holland) coupled with a CMEX 5 digital camera (Euromex, Holland).

3.2 Preprocessing

The 201 microscopic images⁵ collected during the acquisition phase underwent further processing to prepare them for the object detection task. First, we annotated the dataset using the Computer Vision Annotation Tool (CVAT)⁶, a recommended tool for such tasks [29]. Each image was labeled to identify instances of *Leishmania* amastigotes, ensuring high-quality ground truth data. We reviewed and annotated the images, marking regions containing parasites. The annotated dataset was then exported in the YOLO (You Only Look Once) [30] object detection format. To enhance diversity and improve model generalization, we supplemented the dataset with an additional 128 publicly available images from infected human patients [17].

The dataset consists of images with varying resolutions, which posed a challenge for uniform processing. Since the model selected for training requires input images of a fixed rectangular shape, we applied padding to standardize all images before resizing them to 1280×1280 pixels. This preprocessing step ensures consistency across the dataset and optimizes compatibility with the training framework.

4 Synthetic Data Generation

The process of generating synthetic data involved two primary pipelines⁷. Both pipelines were developed using Blender’s python API [31]. The first pipeline

⁵ The dataset can be accessed at <https://doi.org/10.5281/zenodo.15700017>

⁶ <https://www.cvat.ai/>

⁷ The source code for generating the synthetic data can be accessed at <https://github.com/tsikinio/Synthetic-data-for-Leishmania-spp.-amastigotes-detection.git>

was responsible for generating individual images of cells, either infected with parasites or not.

This step aimed to create a diverse set of cell images (Figure 1) that could later be used to construct synthetic microscopic samples. To achieve this, various procedural generation techniques were implemented to introduce natural variability in the shape, size, and texture of the cells. These characteristics were randomized within predefined limits to ensure that the dataset encompassed a broad spectrum of biological diversity. To characterize the cells and infected cells, we examined the key features of *Leishmania* amastigotes as detailed in previous studies [12,13]. These features include the size and shape of the amastigotes, their relative dimensions, internal staining patterns, properties of the kinetoplast and nucleus, and the characteristics of the host cell. In addition to these biological traits, it is also important to consider parameters related to synthetic data generation — such as the number of images, cells per image, and parasites per image—which are not directly tied to the biological characterization but are essential for the image synthesis process.⁸

Additionally, custom materials were created and applied to each cell and parasite to replicate the visual characteristics of real biological samples, such as texture, staining patterns, and structural details. This approach enhances the perceived authenticity of the synthetic dataset, ensuring it visually aligns with what practitioners typically observe in real microscopic images.

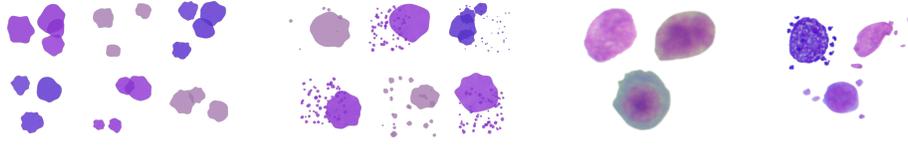


Fig. 1. (a) Synthetic images of healthy cells, (b) synthetic images of infected cells, (c) healthy cells extracted from real images, and (d) infected cells extracted from real images.

To further enrich the collection of the infected and healthy cells generated by the initial pipeline, we extracted samples of both infected and healthy cells from some real images (Figure 1). These manually cropped samples were used along with the synthetic cells images for synthetic image generation.

Once a large and diverse repository of individual cell images was created, the second pipeline was used to generate full synthetic images.

This stage involved placing the previously generated cell images onto background textures. Background selection was performed to ensure visual consistency with real microscopy slides. The infected and non infected cells were po-

⁸ The Appendix A.1 provides additional details on the parameters involved in synthetic data generation, specifying which are related to biological traits and which pertain solely to the image synthesis process.

sitioned randomly but under controlled constraints to prevent significant overlaps, ensuring a natural spatial distribution within each image. The cells were also subjected to transformations such as random scaling, rotation, and lighting variations to enhance realism. A key aspect of this pipeline was the structured randomization, which ensured that each synthetic image remained unique while adhering to biological plausibility. The parameters for this process can be found in Appendix A.2.

The output of this pipeline consisted of synthetic images accompanied by their corresponding annotations. The annotation process was automated, using the known placements of infected cells within each image. Using camera projection methods, precise bounding boxes were generated and saved in the YOLO object detection format, ensuring compatibility with the real dataset.

The entire process involved extensive parametrization, allowing for controlled variation in cell shapes, parasite distributions, lighting conditions, and other critical factors. This high degree of parameterization was essential in creating a dataset capable of improving the generalization performance of the machine learning model, and it is one of the greatest advantages of synthetic data.

To further enhance the realism of our synthetic dataset, we incorporated advanced image processing techniques aimed at bridging the gap between synthetic and real microscopic images. Specifically, we applied Gaussian noise and blur, both of which are commonly observed in real microscopy because of sensor imperfections and other environmental factors. By simulating these characteristics, we aim to make our synthetic images more photorealistic, thus improving the ability of the model to generalize to real-world data (Figure 2).

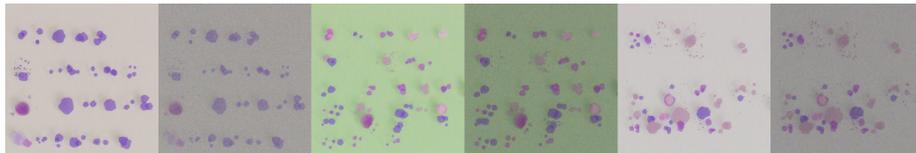


Fig. 2. Synthetic images before noise augmentation (left) and after noise augmentation (right).

The generation of one annotated synthetic images in our approach took in average 1.25 ± 0.083 seconds, whereas the image processing step, which involved applying Gaussian noise and blur to 130 images, required negligible time.

5 Experimental Setup

This section outlines the chosen configurations and environment for our experiments. First, we describe the datasets used for the experiments. Then, we detail the deep learning model selected for the object detection task and justify

our choice. We then describe the training process and conclude with the post-processing phase, which played a key role in enhancing models' performance. All experiments were conducted on a machine equipped with an AMD Ryzen 9 5900X 12-core CPU, 16 GB of RAM, and an NVIDIA GeForce RTX 4090 GPU.

5.1 Datasets

As mentioned earlier, in Section 3, our dataset consists of 201 microscopic images of canine lymph node smears supplemented with 128 publicly available images -as detailed in Section 3 -resulting in a total of 329 instances. This dataset was used to perform 5-fold cross-validation, where each fold involved a different split of the full dataset into training, validation, and test sets with a 70-15-15 % ratio. Using this approach, we trained five models exclusively on real microscopic images (RMI).

Next, we expanded these datasets by generating 330 synthetic images (SI) through our proposed method (Section 4), increasing the total number of instances to 639. The number of synthetic images was chosen to maintain a 1:1 ratio with the real data, as introducing a bias toward synthetic data could lead the trained model to generalize better on synthetic images rather than real-world samples. The synthetic data was split using a 70-15-15% distribution and integrated into the RMI datasets, resulting in five augmented sets. Using these expanded datasets, we trained five models incorporating real microscopic images and synthetic images (RMI+SI).

5.2 Model selection

Leishmania amastigotes detection is a quite complex task that requires high accuracy. At the same time inference speed is a critical factor considering that the proposed approach can be used in real-time applications in laboratories or veterinary clinics. To identify the most suitable model for our case, we conducted multiple training runs using different architectures. Specifically, we evaluated three state-of-the-art models [14]: YOLO v11, Faster R-CNN ResNet152 V1 (1024×1024) [15], and SSD ResNet101 V1 FPN (1024×1024) [16]. All models were trained on one of the folds of the combined RMI+SI dataset, using the same hyperparameters. Specifically learning rate was set to 0.005, batch size to 8 while all other parameters were kept at their default values as defined by the Ultralytics library ⁹. Their performance was evaluated on the validation set containing both real and synthetic images. Based on the results of this experiment (as presented in Table 1), YOLOv11 emerged as the most effective model among those evaluated.

The architecture of YOLOv11 builds upon previous versions with key improvements, such as a more efficient backbone network for feature extraction, advanced neck components that refine spatial relationships, and an optimized

⁹ <https://docs.ultralytics.com/>

Models	mAP50	Recall
YOLO v11	0.692	0.652
Faster R-CNN Resnet152 V1 1024x1024	0.564	0.467
SSD ResNet101 V1 FPN 1024x1024	0.569	0.523

Table 1. Performance comparison of the three models

detection head for accurate bounding box regression and classification. The inclusion of transformer-based attention mechanisms enhances the model’s capability to focus on relevant regions, making it particularly effective for detecting small objects, like *Leishmania* amastigotes. Additionally, the redesigned anchor-free detection mechanism reduces computational overhead while maintaining high detection precision. These architectural advancements make YOLOv11 the optimal choice for our use case. In our experiments, we used also the YOLOv11 pre-trained weights since our dataset is too small for training from scratch.

5.3 Model Training

We followed three different training processes for evaluating the impact of synthetic data for detecting *Leishmania* amastigotes. In the first approach, the model was trained exclusively on real data. In the second, training was performed using a combination of real and synthetic data. Finally, in the third approach, the model was initially trained on both real and synthetic data, followed by a fine-tuning phase using only real data.

The training processes were monitored using real-time logging, with key performance metrics computed after each iteration. The model’s performance was evaluated based on mAP, Precision, Recall, and F1-score. The training process was also tracked using box loss and focal loss metrics, along with evaluation metrics on the validation set, to ensure efficient learning and prevent overfitting. Once the training was over, the best model was saved and ready for further evaluation.

Training on Real Data We trained and validate five different models on the five folds arised from the RMI dataset. The objective was to evaluate the models’ performance when trained purely on real data. The hyperparameters were selected to maximize model’s performance. Specifically, all models were trained with an early stopping mechanism, using a patience threshold of 50 epochs - meaning training would terminate if no significant improvement was observed over 50 consecutive epochs. As a result, the five models were trained for varying durations, ranging from 101 to 131 epochs depending on convergence behavior. Leveraging the hyperparameter tuning feature provided by Ultralytics, the initial learning rate was set to 0.0063, while the final learning rate was adjusted to 0.00951. All other parameters were kept at their default values as defined by the Ultralytics library. Additionally, computational limitations necessitated the utilization of a batch size of 8. Finally, the data augmentation features offered

by Ultralytics were enabled using their default settings. These included random color transformations such as hue, saturation and brightness adjustments, along with random translations, scaling and horizontal flipping. These enhancements aimed to increase the variability of the datasets.

Utilizing Synthetic Data Two processes were followed for evaluating the impact of synthetic data. In the first process, the models were trained on the folds derived from the RMI+SI dataset. In the second, the best-performing models from this initial process were further fine-tuned on the RMI dataset (RMI+SI+RMI) for a few epochs to further enhance their performance.

During the first process, the five models were trained for varying durations, ranging from 133 to 150 epochs, due to the early stopping mechanism with a patience threshold of 50 epochs. Following hyperparameter tuning the initial learning rate was set to 0.00503 and the final learning rate to 0.0085. Additionally a batch size of 8 was selected balancing efficiency and process’ computational cost. All the other parameters, including data augmentation, were set to the default values provided by Ultralytics library.

Once the initial training process was completed, the best-performing model from each fold was saved for further refinement. To enhance its ability to generalize to real-world data, the models’ weights were fine-tuned using the RMI dataset consisting exclusively of real microscopic images. In order to avoid catastrophic forgetting- the models’ disposition to forget the previously learned knowledge- a selective fine-tuning approach was adopted. Specifically, the backbone layers were frozen, preserving the feature extraction capabilities developed during pretraining. This allowed the training process to focus solely on adjusting the external layers, which are responsible for higher-level decision-making, ensuring that the models would adapt effectively to the real data without compromising its foundational learned features.

To ensure stable convergence and effective adaptation to real data during fine-tuning, specific adjustments were made to the training configuration. A lower batch size of 2 was chosen to allow more precise weight updates, given the smaller dataset and the need for refinement without drastic parameter shifts. Additionally, after hyperparameter tuning the initial learning rate was set to 0.00918 and the final learning rate to 0.0098. Finally the patience value was set to 5 epochs, since the model had already been exposed to the real images and it was more prone to overfitting. All other parameters, including data augmentation, were retained as initially configured.

After 12 to 14 epochs, the best-performing models were saved and ready for further evaluation.

5.4 Post-processing and validation

Once the training process was over, we applied tiling inference to enhance detection accuracy using the open source framework Slicing Aided Hyper Inference (SAHI) [32]. This approach is particularly useful when detecting small objects,

such as *Leishmania* infected regions, where every detail matters [33]. By analyzing smaller regions of the image, the model can focus on fine-grained features, leading to improved detection performance compared to processing the entire image at once. In our approach, the tiles have a slight overlap to ensure that no cell is partially cropped. Inference is performed on each tile with a confidence threshold of 0.3 as an optimal threshold effectively balances precision and recall.

After processing image tiles, we used Non-Maximum Suppression (NMS) [34] to remove duplicate bounding box detections. Because overlapping tiles can lead to multiple detections of the same object, NMS selects the highest-confidence prediction and eliminates others based on their overlap. We used Intersection over Union (IoU) as the overlap metric, with a threshold of 0.2. If the IoU between two boxes exceeds 0.2, the box with the lower confidence score is discarded. This ensures only the most accurate detections are retained. We also managed the overlapping region between adjacent tiles to be 0.2 and tile size to be 640X640.

6 Results

To assess the impact of synthetic data on the training process, a comparative study was conducted using three distinct training strategies. The first approach involved training exclusively on the RMI dataset. The second used a combined dataset of RMI and SI (RMI+SI) for training. The third built upon the second by applying an additional fine-tuning phase using only the RMI dataset (RMI+SI+RMI). Evaluation was carried out on the test sets of each fold of the RMI dataset. Each fold included a different subset of images in the test set to ensure an objective evaluation and minimize the risk of biased performance results. No data augmentations were applied to any test set. Table 2 presents the evaluation metrics of the trained models. For each training approach, the table reports the mean values and standard deviations computed across the five cross-validation folds.

Datasets	Recall	F1-score	mAP50
RMI	0.4954±0.0721	0.3853±0.0436	0.3854±0.0462
RMI+SI	0.6174±0.1226	0.5447±0.0645	0.5593±0.1254
RMI+SI+RMI	0.7220±0.0715	0.4678±0.0355	0.6019±0.0462

Table 2. Performance comparison of the three models

The models initially trained solely on real microscopic images (RMI) exhibited several critical limitations in their detection performance. They often failed in low-contrast regions, where infected cells blended into the background, and in densely populated areas, where overlapping structures led to frequent misdetections or omissions. Additionally, the models struggled to generalize across the morphological variability of infected cells, an essential aspect in the detection

of *Leishmania*, resulting in a high false-negative rate and occasional false positives, where background elements or healthy cells were incorrectly classified as infected.

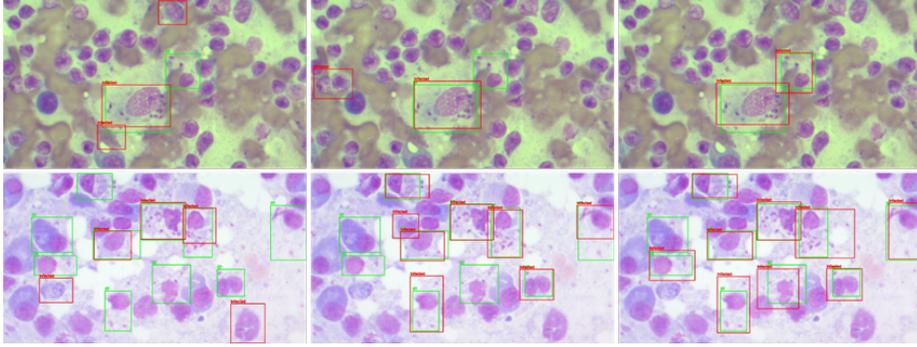


Fig. 3. Performance comparison results for two examples. The predicted bounding boxes are shown in red, while the ground truth annotations are displayed in green. The images on the left display the predictions from RMI, those in the center correspond to RMI+SI, and the ones on the right to RMI+SI+RMI.

Introducing synthetic data in the first fine-tuning stage (RMI+SI) led to a substantial improvement in detection capability, with mAP50 increasing by approximately 0.17 compared to training solely on real data. This stage also resulted in noticeable gains in recall and F1-score, indicating enhanced sensitivity and overall detection balance. However, these improvements were accompanied by increased variability across cross-validation folds, particularly in recall, suggesting a trade-off between performance and stability. The second fine-tuning phase on real data (RMI+SI+RMI) further improved mAP50 to its highest value and reduced standard deviations across metrics, indicating the models' more consistent generalization ability.

A performance comparison example of the three models is presented in Figure 3, using two images from one of the five test sets that were used for model evaluation.

7 Conclusions and Future Work

This paper presents a novel approach for detecting *Leishmania* amastigotes in canine lymph node smears through synthetic data generation. Specifically, we utilized a small dataset of microscopic images of canine lymph node smears to create synthetic data, aligning with current advancements in the field. The performance of this approach was evaluated using a state-of-the-art deep learning model for object detection. By adopting this technique, we address key challenges

in medical imaging for deep learning, such as limited dataset availability, labor-intensive manual annotation, and class imbalance.

Using five-fold cross-validation, our model achieved an average mAP50 of 60.19%, an average Recall of 72.20%, and an average F1-score of 46.78% demonstrating competitive performance in comparison to existing methods. Notably, [14] explored various object detection frameworks for *Leishmania* amastigotes detection in human samples, with YOLOv5 achieving a mAP of 73%, Precision of 68%, and Recall of 69%. The proposed solution approaches these results while utilizing only one-fifth of the real data used in the referenced study, demonstrating the effectiveness of the synthetic data. Additionally, the results suggest that synthetic data augmentation not only compensates for small dataset sizes but also enhances detection reliability by contributing to the creation of a more diverse and robust dataset.

Moreover, with minor modifications, particularly in annotation format, the proposed pipeline can be adapted for other tasks such as segmentation and classification. For instance, in the case of [20], it could effectively address class imbalance, providing a more balanced dataset for training without relying solely on traditional oversampling techniques.

Future research will focus on enhancing the realism of synthetic data to further improve model efficiency and generalization. A key objective is to explore whether an ML model trained exclusively on synthetic data can achieve competitive performance, potentially reducing the reliance on real-world annotated datasets. Additionally, further improvements in accuracy and confidence levels will be pursued by refining model architectures, optimizing training strategies, and incorporating more diverse synthetic variations. Furthermore, future efforts will focus on real-world deployment to ensure the model's robustness and practicality in veterinary diagnostics and field applications. Lastly, given that the synthetic data generation pipeline is highly task-specific, it is important to evaluate its performance on other tasks and datasets to enhance its generalization ability and applicability.

References

1. Akhoundi, M., Downing, T., Votýpka, J., Kuhls, K., Lukeš, J., Cannet, A., Ravel, C., Marty, P., Delaunay, P., Kasbari, M. & Others Leishmania infections: Molecular targets and diagnosis. *Molecular Aspects Of Medicine*. **57** pp. 1-29 (2017)
2. Azami-Conesa, I., Gómez-Muñoz, M. & Martínez-Díaz, R. A systematic review (1990–2021) of wild animals infected with zoonotic *Leishmania*. *Microorganisms*, **9**, 1101 (2021)
3. Solano-Gallego, L., Miró, G., Koutinas, A., Cardoso, L., Pennisi, M., Ferrer, L., Bourdeau, P., Oliva, G. & Baneth, G. LeishVet guidelines for the practical management of canine leishmaniosis. *Parasites & Vectors*, **4**, 1–16 (2011)
4. Maroli, M., Feliciangeli, M., Bichaud, L., Charrel, R. & Gradoni, L. Phlebotomine sandflies and the spreading of leishmaniasis and other diseases of public health concern. *Medical And Veterinary Entomology*. **27**, 123-147 (2013)

5. Morales-Yuste, M., Martín-Sánchez, J. & Corpas-López, V. Canine leishmaniasis: Update on epidemiology, diagnosis, treatment, and prevention. *Veterinary Sciences*, **9**, 387 (2022)
6. Ribeiro, R., Michalick, M., Silva, M., Dos Santos, C., Frézard, F. & Silva, S. Canine leishmaniasis: an overview of the current status and strategies for control. *BioMed Research International*. **2018**, 3296893 (2018)
7. Saridomichelakis, M. Advances in the pathogenesis of canine leishmaniasis: epidemiologic and diagnostic implications. *Veterinary Dermatology*. **20**, 471-489 (2009)
8. Capistrano Costa, N., Souza Pereira, A., Silva, C., Souza, E., Oliveira, B., Ferreira, L., Hernandez, M. & Pereira, V. Exploring bioinformatics solutions for improved leishmaniasis diagnostic tools: a review. *Molecules*. **29**, 5259 (2024)
9. Paltrinieri, S., Solano-Gallego, L., Fondati, A., Lubas, G., Gradoni, L., Castagnaro, M., Crotti, A., Maroli, M., Oliva, G., Roura, X. & Others Guidelines for diagnosis and clinical classification of leishmaniasis in dogs. *Journal Of The American Veterinary Medical Association*. **236**, 1184-1191 (2010)
10. Guerra, J., Fernandes, N., Ressio, R., Magno, J., Kimura, L., Barbosa, J., Bertollo, D., Taniguchi, H., Hiramoto, R., Motoie, G. & Others Evaluation of cytopathological techniques for the diagnosis of canine visceral leishmaniasis with lymph node samples. *Journal Of Comparative Pathology*. **172** pp. 62-71 (2019)
11. Noli, C. & Saridomichelakis, M. An update on the diagnosis and treatment of canine leishmaniasis caused by *Leishmania infantum* (syn. *L. chagasi*). *The Veterinary Journal*. **202**, 425-435 (2014)
12. Saridomichelakis, M., Mylonakis, M., Leontides, L., Koutinas, A., Billinis, C. & Kontos, V. Evaluation of lymph node and bone marrow cytology in the diagnosis of canine leishmaniasis (*Leishmania infantum*) in symptomatic and asymptomatic dogs. *The American Journal Of Tropical Medicine And Hygiene*. **73**, 82-86 (2005)
13. Knottenbelt, C. (2016). RE Raskin, DJ Meyer, *Canine and Feline Cytology: A Color Atlas and Interpretation Guide*, Elsevier Saunders, 2016, ISBN 9781455740833, 544 pp. and 1200 colour illustrations; £97.99 (hard).
14. Tekle, E., Dese, K., Girma, S., Adissu, W., Krishnamoorthy, J. & Kwa, T. DeepLeish: a deep learning based support system for the detection of Leishmaniasis parasite from Giemsa-stained microscope images. *BMC Medical Imaging*. **24**, 152 (2024)
15. S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks," in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 28, 2015.
16. W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, "SSD: Single Shot Multibox Detector," in *Computer Vision – ECCV 2016*, B. Leibe, J. Matas, N. Sebe, and M. Welling, Eds., Lecture Notes in Computer Science, vol. 9905, Springer, 2016, pp. 21–37.
17. Sadeghi, A., Sadeghi, M., Fakhar, M., Zakariaei, Z., Sadeghi, M. & Bastani, R. A deep learning-based model for detecting *Leishmania* amastigotes in microscopic slides: a new approach to telemedicine. *BMC Infectious Diseases*. **24**, 551 (2024)
18. Leal, P., Ferro, L., Marques, M., Romão, S., Cruz, T., Tomás, A., Castro, H. & Quelhas, P. Automatic assessment of leishmania infection indexes on in vitro macrophage cell cultures. *Image Analysis And Recognition: 9th International Conference, ICIAR 2012, Aveiro, Portugal, June 25-27, 2012. Proceedings, Part II 9*. pp. 432-439 (2012)
19. Nogueira, P. Determining leishmania infection levels by automatic analysis of microscopy images. *ArXiv Preprint ArXiv:1311.2621*. (2013)

20. Górriz, M., Aparicio, A., Raventós, B., Vilaplana, V., Sayrol, E. & López-Codina, D. Leishmaniasis parasite segmentation and classification using deep learning. *Articulated Motion And Deformable Objects: 10th International Conference, AMDO 2018, Palma De Mallorca, Spain, July 12-13, 2018, Proceedings 10*. pp. 53-62 (2018)
21. Zare, M., Akbarialiabad, H., Parsaei, H., Asgari, Q., Alinejad, A., Bahreini, M., Hosseini, S., Ghofrani-Jahromi, M., Shahriarirad, R., Amirmoezzi, Y. & Others A machine learning-based system for detecting leishmaniasis in microscopic images. *BMC Infectious Diseases*. **22**, 48 (2022)
22. Ferreira, T., Santana, E., Jacob Junior, A., Silva Junior, P., Bastos, L., Silva, A., Melo, S., Cruz, C., Aquino, V., Castro, L. & Others Diagnostic classification of cases of canine leishmaniasis using machine learning. *Sensors*. **22**, 3128 (2022)
23. Borges, A., Gonçalves, C., Dias, V., Sousa, E., Costa, C. & Silva, R. Visceral Leishmaniasis Detection Using Deep Learning Techniques and Multiple Color Space Bands. *International Conference On Intelligent Systems Design And Applications*. pp. 492-502 (2022)
24. Gonçalves, C., Andrade, N., Borges, A., Rodrigues, A., Veras, R., Aguiar, B. & Silva, R. Automatic detection of Visceral Leishmaniasis in humans using Deep Learning. *Signal, Image And Video Processing*. **17**, 3595-3601 (2023)
25. Sadeghi, A., Sadeghi, M., Sharifpour, A., Fakhar, M., Zakariaei, Z., Sadeghi, M., Rokni, M., Zakariaei, A., Banimostafavi, E. & Hajati, F. Potential diagnostic application of a novel deep learning-based approach for COVID-19. *Scientific Reports*. **14**, 280 (2024)
26. Portuondo-Mallet, L., Mollineda-Diogo, N., Orozco-Morales, R. & Lorenzo-Ginori, J. Detection and counting of Leishmania intracellular parasites in microscopy images. *Frontiers In Medical Technology*. **6** pp. 1360280 (2024)
27. Manettas, C., Nikolakis, N. & Alexopoulos, K. Synthetic datasets for deep learning in computer-vision assisted tasks in manufacturing. *Procedia CIRP*. **103**, 9th (2021)
28. Paproki, A., Salvado, O. & Fookes, C. Synthetic data for deep learning in computer vision & medical imaging: A means to reduce data bias. *ACM Computing Surveys*. **56**, 1-37 (2024)
29. Pangal, D., Kugener, G., Shahrestani, S., Attenello, F., Zada, G. & Donoho, D. A guide to annotation of neurosurgical intraoperative video for machine learning analysis and computer vision. *World Neurosurgery*. **150** pp. 26-30 (2021)
30. Redmon, J., Divvala, S., Girshick, R. & Farhadi, A. You only look once: Unified, real-time object detection. *Proceedings Of The IEEE Conference On Computer Vision And Pattern Recognition*. pp. 779-788 (2016)
31. Conlan, C. The Blender Python API. *Precision 3D Modeling And Add.* (2017)
32. Akyon, F., Altinuc, S. & Temizel, A. Slicing aided hyper inference and fine-tuning for small object detection. *2022 IEEE International Conference On Image Processing (ICIP)*. pp. 966-970 (2022)
33. Shewajo, F. & Fante, K. Tile-based microscopic image processing for malaria screening using a deep learning approach. *BMC Medical Imaging*. **23**, 39 (2023)
34. Neubeck, A. & Van Gool, L. Efficient non-maximum suppression. *18th International Conference On Pattern Recognition (ICPR'06)*. **3** pp. 850-855 (2006)

A Appendix

A.1 Parameter Details for Synthetic Elements Generation

We outline the parameters used in the generation of synthetic data, presenting the components of a synthetic image, and explain how these parameters correspond to the biological traits of *Leishmania* amastigotes.

1. **Number of Images:** Defines the number of synthetic cell images to generate.
2. **Number of Cells per Image:** Specifies the number of main cells appearing in each generating image.
3. **Number of Parasites per Image:** Specifies the number of parasites appearing per image.
4. **Parasite Radius Range:** Sets the minimum and maximum radius for the generated parasites. This parameter is related to the size of Amastigotes. The characteristic size %range is approx. 1-5 μm long by 1-2 μm wide which aids identification.
5. **Cell Radius Range:** Sets the minimum and maximum radius for the generated cells. This parameter is related to the size of the host cell. The characteristic size %range is approx. 15 - 30 μm .
6. **Cell shape:** The typical appearance of the host cell is round to oval in shape, with clearly defined cell borders and moderate cytoplasmic volume. The nucleus is round to slightly oval, basophilic, and centrally or slightly eccentrically located. The cytoplasm appears eosinophilic to lightly granular.
7. **Spread Area:** Determines the spread area for cell placement, controlling spacing. This parameter is related to size relativity. A key diagnostic element is detecting a small object (amastigote) located inside or immediately around a much larger object (host macrophage).
8. **Material Selection:** Controls the selection of materials applied to cells and parasites. This parameter is related to internal staining pattern. While overall stain intensity varies (e.g., old vs. newly stained slides), the relative intensity pattern (kinetoplast darkest, then nucleus, then amastigote cytoplasm, then host cell nucleus and cytoplasm) is typically preserved and diagnostically important.

Leishmania amastigote kinetoplast and nucleus properties are not included in our synthetic element generation process because these features were not consistently present in real images. As a result, it is more likely that generated images will omit these properties rather than include them.

A.2 Parameter Details for Synthetic Image Generation

Listed below are the parameters involved in generating the full synthetic image.

1. **Number of Images:** Defined as above
2. **Background Selection:** Randomly selects a background image from a list.

3. **Number of Non-Infected Cells per Image:** Specifies the number of non-infected cells appearing in each image.
4. **Number of Infected Cells per Image:** Specifies the number of infected cells appearing in each image.
5. **Cells Scaling Factor:** Applies random scaling to cells to introduce variation. This parameter is related to the size of host cells.
6. **Parasites Scaling Factor:** Applies random scaling to cells to introduce variation. This parameter is related to the size of Amastigotes.
7. **Rotation Factor:** Applies random rotation to cells to introduce variation.
8. **Lighting Variation:** Applies random lighting conditions.
9. **Camera Position:** Randomizes camera position to simulate different microscope perspectives.
10. **Image Resolution:** Sets the resolution of the generated images.