www.editada.org

International Journal of Combinatorial Optimization Problems and Informatics, 17(2), 2026, 416-425. ISSN: 2007-1558. https://doi.org/10.61467/2007.1558.2026.v17i2.1139

Video-based detection of the parasite responsible for causing Chagas disease

Alan Villanueva-Paredes,[1] Carlos Brito-Loeza[1,], Ricardo Legarda-Sáenz[1,*], Francisco Escobedo-Ortegón[2,*], and Hugo Ruiz Piña[2,*]*
[1] Universidad Autónoma de Yucatán.
[2] Centro de Investigaciones Regionales Dr. Hideyo Noguchi.
[*] E-mails: carlos.brito@correo.uady.mx, rlegarda@correo.uady.mx, eortegon@correo.uady.mx, rpina@correo.uady.mx

Abstract. This study introduces a technique for identifying the presence of the parasite Trypanosoma cruzi in video recordings, the primary cause of Chagas disease. Early diagnosis—especially during the acute phase—is vital to prevent serious complications. The primary contribution of this work is a computational algorithm capable of detecting the presence of Chagas parasites in capillary tubes containing blood samples. While prior algorithms relied on stained blood samples, this study employs unstained blood samples, utilizing the motility of living parasites for detection. The proposed approach combines optical flow estimation using Farnebäck's algorithm with a classification stage where several machine learning models were evaluated. Among these, the best performance in terms of accuracy was achieved by a convolutional neural network based on the ResNet-18 architecture. The dataset consists of 24 videos totaling 32 minutes of recording. The results demonstrate optimal performance, achieving a F1-score of 0.9383.
Keywords: Chagas disease, video-based parasite detection, machine learning.

Article Info
Received May 20, 2025.
Accepted Dec 11, 2025.

1 Introduction

Chagas disease is a chronic condition caused by the Trypanosoma Cruzi (T. cruzi) parasite found mainly in Latin American regions. Just in Mexico, the estimation is that more than one million people are infected with this parasite (Sanchez-Patiño et al., 2021). Triatomine bugs, commonly known as kissing bugs, are a type of reduviid bug that can carry T. cruzi in their guts and feces. They feed on pets, wild animals, or human blood by night when they are sleeping. It is common for the bugs to defecate while feeding. If a person scratches around the wound, the kissing bug feces may get into the bite, and that person is infected. Another less common way to get infected is during childbirth, by ingesting contaminated food, blood transfusions, etcetera (Khan et al., 2011).

Once infected, Chagas disease evolves in the individual in two stages; the first stage called the acute phase, and the second stage called the chronic phase. The acute phase initiates when the parasite enters the body and lasts for a few weeks or months. During this phase, people may experience fever, tiredness, headache, loss of appetite, vomiting, rash, or may not present any symptoms at all. This can make it hard to diagnose Chagas disease. The chronic phase, on the other hand, may take years to develop. It is in this phase when people experience heart or other life-threatening problems. Early diagnosis of Chagas disease is relevant. Throughout the acute phase, medications are prescribed to kill the parasite. However, if Chagas disease reaches the chronic phase, medications will no longer be able to cure the disease. They may be prescribed to young people only to slow the progression of the disease and manage some health complications that may appear.

Detecting T. cruzi is most straightforward during the acute phase of the disease when the parasites are actively circulating in the bloodstream. Thus, the examination of blood samples serves as a means to verify the presence of parasitemia in a human individual. The most prevalent method for diagnosis involves creating a small layer of blood on a microscope slide, known as a stained smear, which is then examined. Nevertheless, this procedure needs a certain amount of time for the processing of samples. The blood samples are stained by saturating the slide containing the sample with dye and allowing it to sit undisturbed for approximately 7 minutes. Next, a fixative solution consisting of phosphate-buffered saline is applied onto the slide and left

undisturbed for an additional 7 minutes. Ultimately, the extraction of liquids is accomplished by employing distilled water, thereby rendering the blood smear prepared for scrutiny under a microscope. As previously mentioned, this procedure needs expertise in manipulating laboratory apparatus and around twenty minutes for each sample preparation.

In this paper, we present a completely different approach to detecting T. cruzi in blood samples. We no longer use blood smears but capillary tubes where the parasite is still alive and moving. We take advantage of this and design a pipeline with two main steps: optical flow to detect movement and classification to confirm the presence of the parasite.

The rest of the paper is organized as follows: Section II contains a review of different methods for T. cruzi detection; Section III sketches our proposal with a new type of data.

2 Review

To the best of our understanding, all existing computational techniques for detecting T. cruzi use still images of stained blood samples. For instance, (Uc-Cetina et al., 2015) is based on the use of an AdaBoost algorithm, (Morais et al., 2022) uses standard machine learning algorithms on images captured with a cell phone camera, (Pereira et al., 2022) and (Sanchez-Patiño et al., 2021) more recently developed deep learning algorithms based on convolutional networks for blood sample images and histopathological images respectively. Finally, (Ojeda-Pat et al., 2022) is the only work known to propose an algorithm for parasite segmentation.

Despite the achievements of the aforementioned works, there remain numerous challenges that must be addressed in order to prepare them for practical application in diagnosing real-life situations. For example, the standard dimensions of a microscope slide are 75 by 26 millimeters, whereas Trypanosomes have a length ranging from 12 to 30 micrometers. Hence, locating a parasite within a vast area under the microscope can be challenging.

In the upcoming section, our approach includes optical flow methods in addition to classification algorithms to identify regions of intense motion. A quite common method for achieving this is through the use of optical flow algorithms. Current state-of-the-art techniques for computing optical flow rely on neural networks. Several networks used for optical flow computation include FlowNet (Dosovitskiy et al., 2015), DeepFlow (Weinzaepfel et al., 2013), LiteFlow (Hui et al., 2020), and others. Models for all of these networks can be obtained with parameters that have been trained using either one or both of the most often used databases connected with optical flow, such as the KITTI Vision Benchmark suite (Geiger et al., 2012) or MPI Sintel optical flow dataset (Butler et al., 2012). Regrettably, these models exhibit poor performance when used to compute the optical flow of films depicting blood samples. For this work, we utilize a traditional method for computing optical flow. We chose the dense optical approach described in (Farnebäck, 2003) because of its flexibility and simplicity of implementation.

3 Proposal

As opposed to the stained smear method, we suggest a more robust technique involving placing a small amount of blood in a 1mm diameter capillary tube, followed by centrifugation to separate the blood components. Subsequently, the sample is placed under a microscope, with the focus solely on the area where the red blood cells terminate, as this is where the parasites will have clustered. There exists a significant distinction between the two approaches. The stained smear approach yields static images, whereas the capillary tube method produces videos. In the latter, the parasite remains alive and in constant motion. There is also no need to stain the blood sample in the tube. In Figure 1, we illustrate the disparity between the two approaches. The left image in Fig. 1 displays a stained blood sample that contains a T. cruzi. The right image in Fig. 1 displays a single frame extracted from a video clip captured from a capillary tube, showcasing a compact cluster of T. cruzi parasites. The parasites in the video frame are hardly discernible based on their shape or color. Nevertheless, by analyzing a sequence of successive frames, skilled technicians can effortlessly distinguish the parasites based on their distinct motion, which sets them apart from other objects present in the video. It is noteworthy that the T. cruzi parasite exhibits a remarkably rapid and erratic motion in comparison to other things captured in the videos, such as erythrocytes, leukocytes, and other type of parasites.

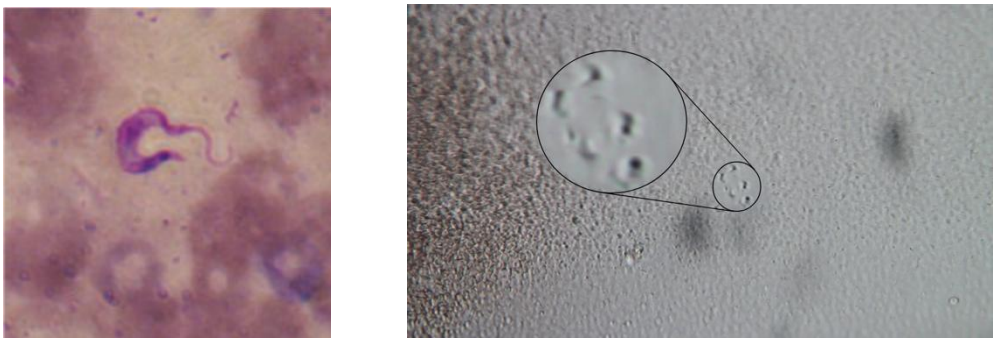

Fig. 1. To the left, a stained smear of blood sample showing a T. cruzi parasite. To the right, a frame extracted from a video footage of a capillary tube showing a group of parasites.

This study presents a computational technique that aims to replicate the inherent human capability to detect and track the fast and unpredictable movement of an object, specifically the T. cruzi. The technique comprises of two distinct stages: object detection and image classification.

In this work, we use two essential elements of computer vision: object detection and image classification. The first, involves detecting and locating objects in images or videos while the second is the process of assigning a specific label to an image or a specific region of interest (ROI) within an image.

In this study, our focus lies in the identification of objects (parasites) exhibiting swift motion. To do this, we employ optical flow approaches, which entail accurately calculating the direction and magnitude of movement for either all or a subset of pixels within an object between two consecutive frames and then making a classification through machine or deep learning techniques.

4 Database

The video database we possess comprises of 24 videos taken from different blood samples from different infected mice. Overall, we have 32 minutes of footage, with a frame rate of 20 frames per second. This footage was captured at the Dr. Hideyo Noguchi Regional Research Center, which is affiliated with the Autonomous University of Yucatán. Some sections of the recordings contain depictions of T. cruzi, and some others do not. Furthermore, videos exhibit a wide range of approaches, varying illuminations, and different levels of motion when capturing the parasites.

Additionally, we possess a collection of images derived from the videos. To construct this collection, we used a generic optical flow algorithm to identify ROIs in the video where there is significant motion and to extract corresponding images from those ROIs. This database has 9547 images, each sized 70 by 70 pixels, in which some level of movement was detected. Then, skilled technicians manually classified each image as containing or not a parasite obtaining a total of 4,776 positive images and 4,771 negative ones. Several illustrations from this database are displayed in Fig. 2.

Fig. 2. Top row: ROIs containing a parasite. Bottom row: ROIs without a parasite

5 Methodology

The Fig. 3 provides an extensive illustration of the pipeline for detecting T. cruzi parasites in a video footage, showcasing the detailed steps of our proposed methodology.

The first step in the algorithm is to find the optical flow, that is, the vectors of motion, we use Farnebäck's algorithm (Farnebäck, 2003). This algorithm involves partitioning the image into smaller sections and estimating the movement inside each section using polynomial expansion. The technique computes the flow vectors by comparing the polynomial representations of consecutive frames, which indicate the movement of pixels from one frame to another. Throughout some experiments, we found the following optimal parameters for our video frames: image scale to build pyramids for each image equal to 0.5, the number of pyramid layers equal to 3, the average window size selected is 15, the number of iterations to be performed at each pyramid level equal to 3, the size of the pixel neighborhood equal to 5, and the standard deviation of the Gaussian kernel used to compute smooth derivatives equal to 1.2.

The second stage involves using the pixel's movement vector to identify ROIs that exhibit significant movement, considering the specific video being analyzed. We employ a thresholding approach to carry out this task. It is preferable for this threshold to be adaptable to accurately reflect the situations depicted in the video. To get the most effective threshold value, a Gaussian filter is used on the movement vector of two consecutive frames. This involves sliding a window with dimensions of 70 × 70 over the optical flow array to calculate the average magnitude of the motion. The threshold for motion magnitude is originally determined by selecting the highest value among all the averages calculated for each pair of successive frames in the video. The preceding approach involves analyzing the whole video to decide the most suitable thresholding value. Long videos may not be advisable for this approach. Alternatively, one can employ a dynamic threshold that is regularly updated after a certain number of frames. This allows for the thresholding value to be adjusted whenever a new maximum value is discovered.

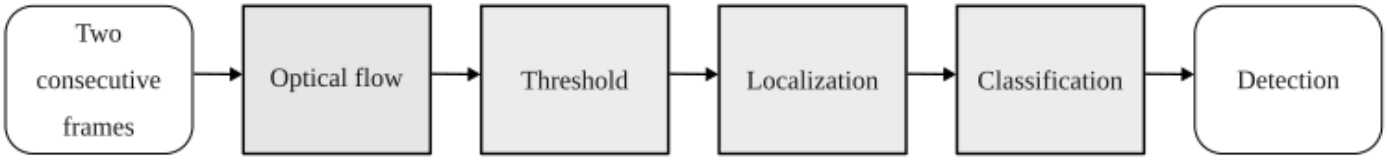

Fig. 3. Video footage analysis pipeline to detect T. cruzi parasites.

The third phase involves isolating the regions of rapid motion by employing a connected components method.

The fourth step involves eliminating ROIs that exhibit rapid motion and do not correlate to parasites. This is achieved by employing a classification algorithm. The incorrect detections mostly result from the motions of the organelles present in the blood, the sedimentation movement that occurs when the capillary tube is allowed to settle after centrifugation, and the unintentional displacement of the camera. We conducted experiments with various classifiers, including support vector machines (SVM) (Vapnik, 1997), logistic regression and decision trees (James et al., 2023), random forest (RF) (Breiman, 2001), multi-layer perceptron (MLP) (Rosenblatt, 1958), and residual convolutional networks (ResNet) (He et al., 2016).

The practical realization of the aforementioned algorithms requires the choice of several parameters. Here, we provide a full list of the essential variables for each algorithm that was evaluated.

SVM. The support vector machines approach allows for the variation of three parameters: the trade-off parameter C, the kernel parameter α, and the kernel type.

Logistic regression. The parameters for logistic regression (LG) are the regularization parameter C and the solver type.

Decision Tree. The decision tree (DT) adjusts splitting criterion, the maximum depth, the maximum features, and the α value.

Random Forest. The random forest (RF) model is characterized by three parameters: the maximum depth, which determines the maximum number of levels in each decision tree; the maximum number of characteristics to examine at each split; and the number of estimators, which refers to the number of decision trees in the forest.

Multi-Layer Perceptron. In the multi-layer perceptron (MLP), the hidden layers, activation function, solver, α, and learning rate can change.

Residual Network. Residual networks (RN) in conventional implementations often employ rectified linear units (ReLu) as activation functions and essentially modify the number of hidden layers. The typical values consist of 18, 34, 50, 101, and 154 layers. Pre-trained models require input images that have been normalized.

To select the optimal parameters for each classifier, we used a grid search technique with a 5-fold cross-validation where at each fold, the dataset is divided into two parts: 80% for the training set and 20% for the test set. In Table I, we display the optimal parameters and scores for each classifier.

6 Results

In this section, we evaluate our machine and deep learning models both quantitatively and qualitatively.

Quantitative assessment

In the context of medical image-based diagnosis, the performance of machine learning models must be assessed using metrics that reflect clinical priorities. Four commonly used metrics are **accuracy**, **precision**, **recall (sensitivity)**, and **F1-score**, each providing distinct insights into diagnostic reliability.

- **Accuracy** represents the proportion of correctly classified cases among all cases. While useful as an overall indicator, accuracy can be misleading in imbalanced datasets typical of medical imaging, where healthy cases often outnumber diseased cases. A model may achieve high accuracy by predominantly predicting the absence of disease, yet fail to detect true positives.
- **Precision** measures the proportion of true positive predictions among all positive predictions. High precision is clinically relevant because it minimizes false positives, reducing unnecessary treatments, patient anxiety, and additional diagnostic procedures.
- **Recall (Sensitivity)** quantifies the proportion of actual positive cases correctly identified. In clinical practice, recall is critical since false negatives—missed disease cases—can lead to delayed treatment and adverse outcomes. Screening applications often prioritize recall to ensure that as many true cases as possible are detected.

- **F1-score** is the harmonic mean of precision and recall, providing a balanced metric when both false positives and false negatives carry significant consequences. This is particularly important in healthcare scenarios where datasets are imbalanced and both types of errors have serious implications.

Table 1. Performance metrics for each mode and the optimal values for the parameters obtained using a grid search with 5-fold cross validation

Classifier	Parameter	Optimal Value	Accuracy	Recall	Precision	F1-score
SVM	C	10	0.8800 ± 0.0052	0.8591 ± 0.0158	0.8968 ± 0.0072	0.8774 ± 0.0065
	alpha	0.1				
	Kernel	Radial basis				
Logistic Regression	C	100	0.7998 ± 0.0096	0.7353 ± 0.0162	0.8445 ± 0.0112	0.7861 ± 0.0113
	Solver	SAGA				
Decision Tree	Criterion	Gini	0.8000 ± 0.0104	0.8092 ± 0.0303	0.7949 ± 0.0065	0.8017 ± 0.0139
	Maximum depth	20				
	Maximum number of features	Auto				
	alpha	0.001				
Random Forest	Maximum depth	50	**0.9004 ± 0.0037**	**0.9049 ± 0.0109**	**0.8970 ± 0.0093**	**0.9009 ± 0.0038**
	Maximum number of features	log2				
	Number of estimators	512				
Multi-Layer Perceptron	Number of neurons per hidden layer	100	0.8043 ± 0.0071	0.7410 ± 0.0111	0.8489 ± 0.0126	0.7912 ± 0.0074
	Activation function	ReLU				
	Solver	SGD				
	alpha	0.0001				
	Learning rate	Adaptive				

In Table 1, we display the performance metrics values computed for each classification algorithm tested together with their standard deviations using 5-fold cross-validation: accuracy, recall, precision, and F1-score. The first observation is that Logistic Regression (LR) underperforms compared to the rest of the models. This is because LR relies on strong assumptions of linearity between predictors and the log-odds of the outcome. This is an indication that the problem of Chagas parasite classification is complex and non-linear. Tree-based models perform better due to the implicit variable selection handle the problem more robustly. On the other hand, Support Vector Machines (SVM) outperform the single decision trees model most likely because the classes are separable by a clear margin, as it optimizes the decision boundary by maximizing this margin. Additionally, SVM usually demonstrates strong performance on small to medium-sized datasets which is the case of our dataset. SVM does this by utilizing support vectors rather than modeling the entire data distribution. Finally, Random Forest models surpass the Multilayer Perceptron (MLP) because is relatively robust and less sensitive to parameter choices. RF models are less prone to overfitting on small datasets like ours due to their ensemble nature.

With all the above considerations, we observe that effectively the best result was obtained with the Random Forest algorithm that achieved **accuracy = 0.9004**, **recall = 0.9049**, **precision = 0.8970**, and **F1-score = 0.9009**, indicating strong overall performance. The high recall obteined by RF suggests that the model effectively identifies most positive cases, which is essential for early

diagnosis, while the high precision reduces the likelihood of unnecessary follow-up procedures. The F1-score confirms a balanced trade-off between sensitivity and specificity, making the approach suitable for clinical applications where both error types must be minimized.

The results obtained with Random Forest bring some insigth about the problem faced. Random Forest is an ensemble learning method based on decision trees that can deliver superior results compared to other traditional machine learning algorithms due to its ability to model complex, nonlinear relationships between features without requiring extensive preprocessing or feature engineering. This makes it particularly effective for heterogeneous datasets, such as our dataset fot the clinical diagnostics of the Chagas disease, where interactions among variables are rarely linear. In our algorithm pipeline, the preprocessing of the recordings was kept minimal just applying a soft smoothig with the Gaussian filter. Unlike single decision trees, Random Forest mitigates overfitting by aggregating predictions from multiple trees trained on bootstrapped samples, improving generalization performance. It is also robust to noise and outliers, which are common in real-world clinical data. For instance, in our dataset, due to the dynamic nature of the parasite and other living organism in the sample there is always the presence of noise and outliers. Compared to algorithms like Support Vector Machines or Multi-Layer Perceptron neural networks, Random Forest often requires less parameter tuning and scales well to medium-sized datasets, offering a practical balance between accuracy, robustness, and interpretability.

Going further with Deep Learning classifiers.

Although the Random Forest algorithm obtained a reasonably high accuracy value, we chose to evaluate the effectiveness of a few deep learning architectures. Unsurprisingly, deep learning convolutional networks far outperformed all other machine learning models. We trained the residual networks using mini batches of RGB images and resized the images to dimensions of 224 pixels in height and width. The image pixel values were mapped into a range of [0, 1] and subsequently normalized using mean values of [0.485, 0.456, 0.406] and standard deviation values of [0.229, 0.224, 0.225]. Only 10 epochs were run for each model.

We evaluated Residual Networks (ResNet) with depths of 18, 34, and 50 layers, as we did not find any enhancements by increasing the depth beyond these values. The ResNet model with 18 layers (ResNet-18) achieved an impressive accuracy of 0.9379 on the test set, while the models with 34 and 50 layers achieved accuracies of 0.9447 and 0.9451 respectively. In practice, the ResNet-18 network is probably the best option due to its reduced number of parameters. The recall, precision and F1-score for ResNet-18 were also excellent reaching 0.9587, 0.9179, and 0.9383. These values, together with the standard deviations are displayed in Table 2.

The positive results of these initial experiments indicate that further research is required with cutting-edge neural networks.

Table 2. Performance metrics for the Residual Network model

Model	Accuracy	Recall	Precision	F1-score
ResNet-18	0.9379 ± 0.0077	0.9587 ± 0.114	0.9179 ± 0.0127	0.9383 ± 0.0089

Table 3. Performance comparison between the Random Forest and the Residual Network mean performance metrics.

Metric	Random Forest	ResNet-18	Absolute Δ
Accuracy	0.9004	0.9379	+0.0375
Recall	0.9049	0.9587	+0.0538
Precision	0.8970	0.9179	+0.0209
F1-score	0.9009	0.9383	+0.0374

In Table 3, we present the comparison between the best two learning models. These are the clinical implications of

- **Sensitivity (Recall) and Missed Diagnoses:**
  ResNet-18's higher recall (0.9587 vs 0.9049) reduces the miss rate from 9.51% to 4.13%, a ~56% relative reduction in missed positive cases. In clinical screening—especially in the acute phase of diseases where early detection is critical—this directly lowers the risk of false negatives and delayed treatment.

- **Precision and Unnecessary Follow-ups:**

Precision improves from 0.8970 to 0.9179, decreasing the false discovery rate (1 − precision) from 10.30% to 8.21% (~20% relative reduction). This helps limit unnecessary confirmatory tests, patient anxiety, and avoidable interventions.

- **Overall Balance (F1-score):**
  The F1-score increase (0.9383 vs 0.9009; +0.0374) indicates a better balance between catching true cases (recall) and maintaining reliable positive calls (precision). This balance is essential in typical imbalanced clinical datasets, where accuracy alone can be misleading.

- **Accuracy and Operational Reliability:**
  The 3.75-point accuracy gain provides stronger overall correctness, but should be interpreted alongside recall/precision given class imbalance and the asymmetric costs of errors in medicine.

Qualitative assessment

Finally, to evaluate qualitatively our algorithm, in Fig. 4, we demonstrate the result of applying Farneback's optical flow method. The intensity of the color in the top row and left image corresponds to the magnitude of the motion vector, while the color itself indicates the direction of motion. In this study, we do not make use of the direction of movement. Therefore, Fig. 4 top row and right image display the result of the threshold in relation to the intensity of motion. In the last image of Fig. 4, in the bottom row, the final forecast is displayed. In this prediction, red boxes represent ROIs with intense motion but not caused by the presence of a parasite. On the other hand, green boxes indicate ROIs where the likelihood of finding a parasite is quite high.

7 Conclusions

We have developed an innovative proof-of-concept technique for identifying the presence of the parasite responsible for Chagas disease in blood samples. Unlike previous approaches, our method does not require stained images, removing one of the critical variables that earlier staining-based techniques often suffered from. Instead, it relies on the distinct movement patterns of live parasites, detected through an algorithm that first identifies regions of interest (ROIs) with significant motion and then applies a classification method to confirm the presence of parasites within those ROIs.

To validate the classification stage, we tested several machine learning models, with the best performance achieved by a convolutional neural network based on the ResNet-18 architecture, reaching an impressive F1 score of 0.9383. While these results are highly promising, this work is based on a limited dataset consisting of 24 videos totaling 32 minutes of recording, and therefore should be considered a proof of concept. Future work should focus on constructing a robust dataset that includes different *Trypanosoma cruzi* strains, multiple mouse strains, and even human blood samples to ensure generalizability and clinical applicability. Additionally, exploring modern deep learning architectures—such as those incorporating attention mechanisms and transformer-based models—may further improve detection accuracy and robustness.

Although we considered variations in illumination during video recording, the sensitivity of the method to this variable must be extensively tested, and complementary techniques may be required to mitigate its impact. It is important to note that, in practice, obtaining a single positive detection is sufficient to confirm positive parasitemia. Although the model may produce a small number of false positives among positive detections, a key advantage of the methodology is that the position of each detection in the video track is recorded for subsequent expert review, ensuring reliability and traceability in clinical applications.

Moreover, this technology has the potential to significantly automate the diagnostic process by speeding up and simplifying it. Unlike traditional methods that require the intervention of trained clinicians and complex staining procedures, our approach only needs a small blood sample and centrifugation of the capillary tube before placing the sample under a microscope for algorithmic analysis. This streamlined workflow could reduce human error, lower costs, and make early detection more accessible in resource-limited settings.

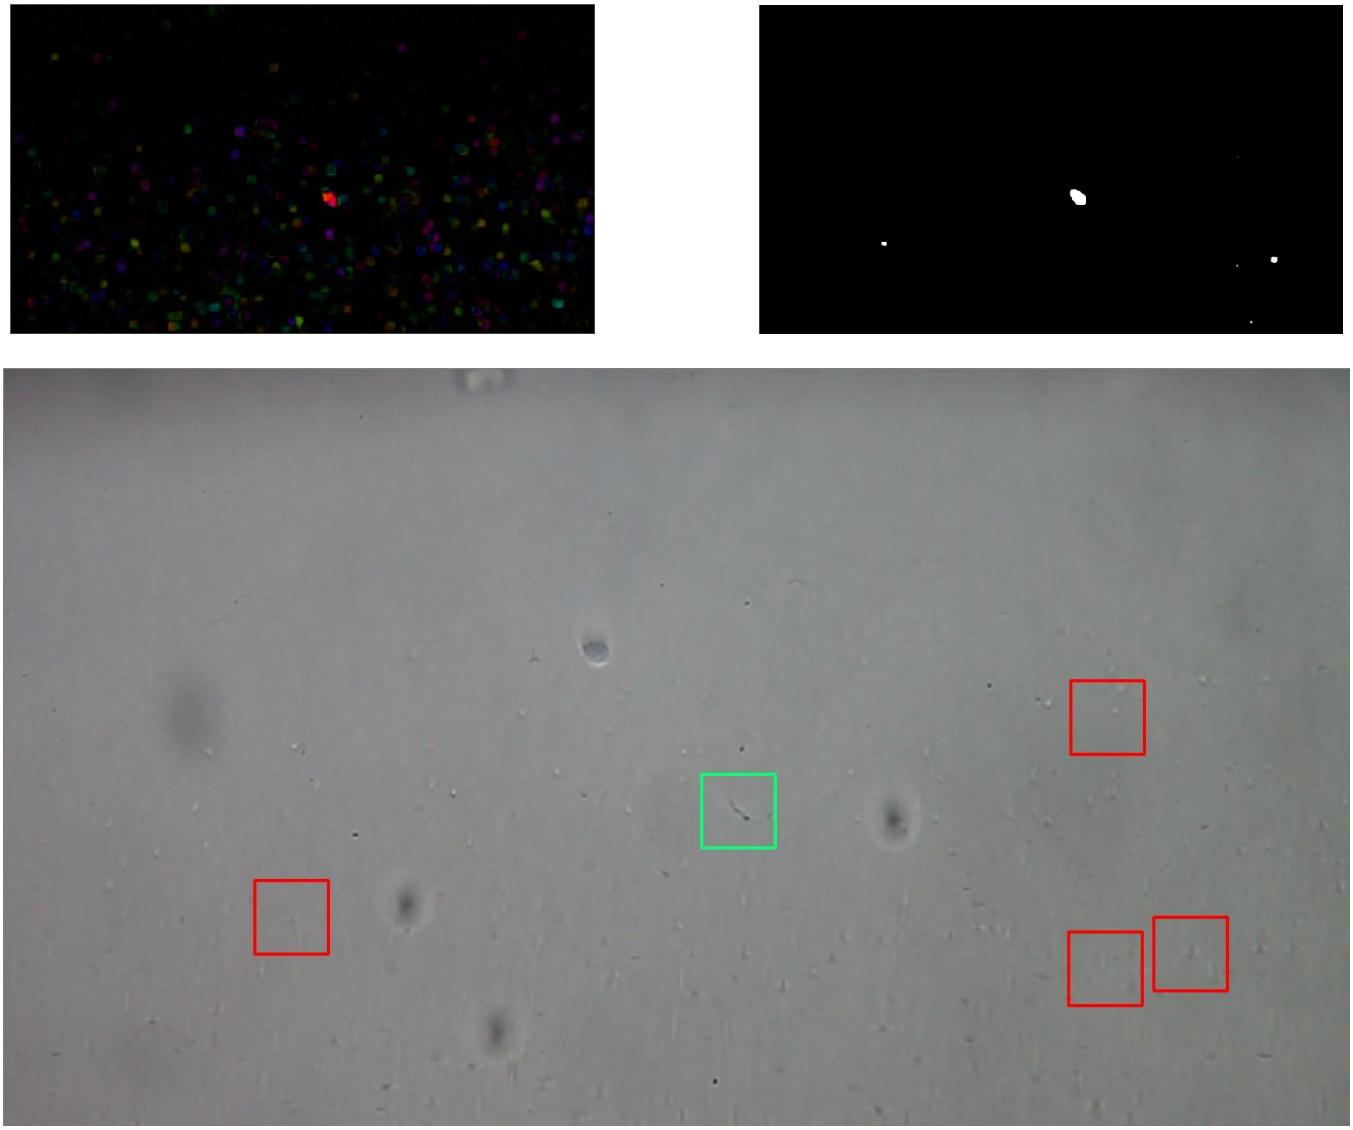

Fig. 4. In the top row to the left, the dense optical flow computed using Farneback's method is shown. The colors indicate both the direction and intensity of the pixel's motion. In the top row to the right, a binary image representing the magnitude of the motion vectors is presented. White color means large values of magnitude above some threshold. In the bottom row bounding boxes around ROIs are presented. Red boxes indicate ROIs with high intensity motion but categorized as non-parasitic ones. Green boxes indicate ROIs with high intensity motion having a parasite.

References

Breiman, L (2001). Random Forests. *Machine Learning* 45, 5–32.

Butler, D. J. and Wulff, J. and Stanley, G. B. and Black, M. J. A naturalistic open source movie for optical flow evaluation. In *Computer Vision – ECCV 2012: 12th European Conference on Computer Vision, Florence, Italy, October 7-13, 2012, Proceedings, Part VI 12*, (pp. 611-625). Springer Berlin Heidelberg.

Dosovitskiy A, Fischer P, Ilg E, Hausser P, Hazirbas C, Golkov V, Van Der Smagt P, Cremers D, Brox T. (2015). Flownet: Learning optical flow with convolutional networks. In *Proceedings of the IEEE international conference on computer vision* (pp. 2758-2766). IEEE.

Farnebäck, G. (2003). 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 Berlin Heidelberg.

Geiger A., Lenz, P., & Urtasun, R. (2012). Are we ready for Autonomous Driving? The KITTI Vision Benchmark Suite. In *2012 Conference on Computer Vision and Pattern Recognition* (pp. 3354-3361). IEEE.

He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 770-778). IEEE.

Hui, T. W., Tang, X., & Loy, C. C. (2020). A lightweight optical flow cnn—revisiting data fidelity and regularization. *IEEE transactions on pattern analysis and machine intelligence*, 43(8), 2555-2569.

James G., Witten, D., Hastie, T., Tibshirani, R., & Taylor J. (2023). An Introduction to Statistical Learning: With Applications in Python. Springer.

Jung, T., Anzaku, E. T., Ozbulak, U., Magez, S., Van Messem, A., & De Neve, W. (2020). Automatic detection of Trypanosomosis in thick blood smears using image pre-processing and deep learning. In *International Conference on Intelligent Human Computer Interaction* (pp. 254-266). Springer International Publishing.

Khan, M. I., Acharya, B., Singh, B. K., & Soni, J. (2011). Content based image retrieval approaches for detection of malarial parasite in blood images. *International Journal of Biometrics and Bioinformatics (IJBB)*, 5(2), 97.

Morais, M. C. C., Silva, D., Milagre, M. M., de Oliveira, M. T., Pereira, T., Silva, J. S., … & Nakaya, H. I. (2022). Automatic detection of the parasite Trypanosoma cruzi in blood smears using a machine learning approach applied to mobile phone images. *PeerJ*, 10, e13470.

Ojeda-Pat, A., Martin-Gonzalez, A., Brito-Loeza, C., Ruiz-Piña, H., & Ruz-Suarez, D. (2022). Effective residual convolutional neural network for Chagas disease parasite segmentation. *Medical & Biological Engineering & Computing*, 60(4), 1099-1110.

Pereira, A. S., Mazza, L. O., Pinto, P. C., Gomes, J. G. R., Nedjah, N., Vanzan, D. F., ...& Soares, J. G. (2022). Deep convolutional neural network applied to Trypanosoma cruzi detection in blood samples. International Journal of Bio-Inspired Computation, 19(1), 1-17.

Poostchi, M., Silamut, K., Maude, R. J., Jaeger, S., & Thoma, G. (2018). Image analysis and machine learning for detecting malaria. *Translational Research*, 194, 36-55.

Purwar, Y., Shah, S. L., Clarke, G., Almugairi, A., & Muehlenbachs, A. (2011). Automated and unsupervised detection of malarial parasites in microscopic images. *Malaria journal*, 10, 1-11.

Rosenblatt, F. (1958). The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6), 386–408.

Sanchez-Patiño, N., Toriz-Vazquez, A., Hevia-Montiel, N., & Perez-Gonzalez, J. (2021, November). Convolutional neural networks for chagas' parasite detection in histopathological images. In *2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)* (pp. 2732-2735). IEEE.

Tek, F. B., Dempster, A. G., & Kale, I. (2006). "Malaria parasite detection in peripheral blood images". *BMVA*, 347-356

Uc-Cetina, V., Brito-Loeza, C., & Ruiz-Piña, H. (2015). Chagas parasite detection in blood images using AdaBoost. *Computational and Mathematical Methods in Medicine*, 2015(1), 139681.

Vapnik, V.N. (1997). The Support Vector method. In *International conference on artificial networks* (pp. 261-21). Springer, Berlin, Heidelberg.

Weinzaepfel, P., Revaud, J., Harchaoui, Z., & Schmid, C. (2013). DeepFlow: Large displacement optical flow with deep matching. In *Proceedings of the IEEE international conference on computer vision* (pp. 1385-1392). IEEE.