

Pipeline Leak Detection Using Infrared Cameras: A Convolutional Neural Network Approach

João Vitor S. Mendes^{1*}, João Pedro Almeida¹, Rodrigo Dias Paolillo², Alexandre Adonai Silva¹, Rodrigo F. Bastos¹, Herman A. Lepikson¹

¹Robotics Department, SENAI CIMATEC University; ²Optics and Photonics Department, SENAI CIMATEC University; Salvador, Bahia, Brazil

Pipeline leak detection is crucial for maintaining pipeline safety, particularly in complex environments. This study proposes a novel approach that integrates infrared cameras with a convolutional neural network model, specifically VGG16, utilizing infrared cameras that do not inherently measure temperature. Our results indicate that this approach is highly effective, with the model achieving 100% accuracy on both the training and validation datasets, and a near-zero validation loss in a laboratory environment. The confusion matrix confirmed that there were no misclassifications, and the Receiver Operating Characteristic (ROC) curve demonstrated an Area Under the Curve (AUC) of 1.0. These findings underscore the model's potential for real-world pipeline monitoring applications.

Keywords: Infrared Cameras. Leak Detection. Computer Vision.

Pipeline leak detection is a critical concern in industries such as oil and gas, as well as water management, due to its significant economic and environmental implications. Effective detection of leaks is essential to minimize losses and prevent environmental damage [1,2].

Various methods have been developed for detecting pipeline leaks, including the use of ultrasonic sensors, thermal cameras, and infrared cameras. Ultrasonic sensors detect changes in acoustic signals within the pipeline but face limitations due to physical access constraints and environmental noise interference [3]. Thermal cameras identify variations in surface temperature, although their effectiveness can be compromised by environmental conditions and variations in ambient temperature [4]. Infrared cameras are particularly notable for their ability to detect temperature differences with high sensitivity without direct contact. These cameras capture emitted infrared radiation and convert it into visible images, enabling the detection of thermal anomalies that may indicate leaks [5].

Infrared cameras operate by detecting thermal radiation emitted by all objects above absolute zero. This radiation is captured by sensors within the camera and converted into a visible image that represents the thermal distribution of the observed objects. The main advantage of infrared cameras is their ability to perform inspections without physical contact and in low-visibility conditions. However, these cameras can be expensive and require precise calibration to ensure measurement accuracy [1,2].

To analyze and classify images captured by infrared cameras, advanced image processing techniques and machine learning algorithms are employed. Traditional image processing techniques may include filtering to enhance specific features and segmentation to isolate areas of interest. Machine learning algorithms, particularly Convolutional Neural Networks (CNNs), are widely used to recognize patterns and detect anomalies in thermal images [4]. Recent advancements in deep learning have shown significant improvements in the automatic and accurate detection of leaks. These methods involve training models with extensive datasets of infrared images, enabling efficient differentiation between leaks and other thermal characteristics [3,5].

The work presented by Xie and colleagues [6] introduces an innovative automated leakage detection method that combines infrared thermography (IRT) with the Faster R-CNN object

Received on 24 May 2025; revised 15 July 2025.

Address for correspondence: João Vitor S. Mendes. Av. Orlando Gomes, 1845 - Piatã, Salvador – BA – Brazil, Zipcode: 41650-010 E-mail: vitor.mendes@ieee.org.

J Bioeng. Tech. Health 2025;8(4):379-384
© 2025 by SENAI CIMATEC University. All rights reserved.

detection technique. This methodology utilizes a modified VGG16 network for feature extraction, allowing for the detection of finer leakage features and details in the original infrared images. The study reports high performance metrics, with mean Average Precisions (mAPs) of 1.00, 0.98, and 0.99 for detecting leaks at pipes, valves, and flanges, respectively, under various leakage scenarios. The results demonstrate that the proposed system is both practical and robust, maintaining high detection accuracy even in complex backgrounds and diverse operational conditions, such as varying ambient light, changes in camera angles, and pedestrian interference.

In this study, we propose a similar pipeline leak detection system of Xie and colleagues [6]. However, we use an infrared camera that does not natively measure temperature. Our objective is to assess whether this non-temperature-measuring infrared camera can achieve comparable performance to the system described by Xie and colleagues [6]. By evaluating this alternative approach, we aim to determine its effectiveness in real-world leak detection scenarios and contribute to the broader application of infrared imaging technologies in pipeline monitoring. This investigation will provide insights into the feasibility and performance of using non-temperature-sensitive infrared cameras for detecting pipeline leaks, further advancing the field of automated pipeline monitoring and safety.

Materials and Methods

This study presents the development and validation of a convolutional neural network (CNN) for detecting leaks in images. The method covers the experimental setup description and model training details, divided into data preparation, model architecture, training procedures, and performance evaluation in detail.

Model Training

The dataset consists of images organized into 'leak' and 'no-leak' categories. Each image is

resized to 224x224 pixels and normalized to a [0, 1] range to standardize input features. Labels are encoded as binary values, with '0' for 'no-leak' and '1' for 'leak'. To ensure balanced representation, the data is split into training and validation sets using stratified sampling, which preserves the proportion of each class.

The model utilizes the VGG16 architecture, a well-known CNN pre-trained on the ImageNet dataset. The base model, which includes multiple convolutional layers followed by max-pooling layers, is used as a feature extractor by excluding its top classification layer. The extracted feature maps are processed by a custom classification head designed as follows:

- **Flattening Layer:** Converts the 2D feature maps into a 1D vector.
- **Dense Layer:** A fully connected layer with 512 units and ReLU activation, described by

$$ReLU(x) = \max(0, x) \quad (1)$$

where x is the input to the layer.

- **Dropout Layer:** Applied with a rate of 0.5 to reduce overfitting by randomly setting half of the neurons to zero during training.
- **Output Dense Layer:** A single unit with sigmoid activation, which outputs a probability p for the binary classification:
where x is the input to the sigmoid function.

$$\sigma(x) = \frac{1}{1 + e^{-x}} \quad (2)$$

The model is compiled using the Adam optimizer, which adjusts the learning rate adaptively based on estimates of first and second moments of the gradients. The learning rate is set to (1×10^{-4}) . Adam is defined by the under rule:

$$\theta_{t+1} = \theta_t - \frac{\alpha}{\sqrt{v_t} + \epsilon} m_t \quad (3)$$

where (θ_t) is the parameter, (α) is the learning rate, (m_t) and (v_t) are estimates of the first and second moments of the gradients, respectively, and (ϵ) is a small constant for numerical stability. Binary cross-entropy is used as the loss function,

which measures the difference between the accurate labels y and the predicted probabilities p :


$$\text{Loss} = -(y \log(p) + (1 - y) \log(1 - p)) \quad (2)$$

The model is trained for ten epochs with a batch size of 32. During training, the model's performance is monitored using accuracy and loss metrics, and training progress is recorded with a CSV logger. The average inference time per image is calculated to assess computational efficiency.

Experiment Setup (Figure 1)

To experiment, it was essential to establish a system capable of simulating a fluid leak in a pipe. The experimental setup comprised a cyclical water system utilizing an aquarium pump connected to a PVC pipe. The water transported by the pump traverses the pipe and is subsequently returned to the same receptacle as the pump, thereby completing the cycle. The water was heated using a portable electric heater/blower, thus enabling the camera to identify thermal variations in the fluid in relation to the pipe. Two distinct PVC pipes were used: one with an aperture to simulate the leak and one without, serving as a control. Data was collected using an FLIR ADK camera, which was connected to a computer running a Python script. This script was programmed to take pictures at 100-millisecond intervals, capturing thermal images of the pipe and allowing for a detailed analysis of the fluid conditions and identification of the simulated leak.

Figure 1. Technical setup drawing.



Results and Discussion



The model's performance was comprehensively evaluated using several metrics, including accuracy, loss, confusion matrix, and the Receiver Operating Characteristic (ROC) curve. The following analysis provides a detailed examination of each metric, offering insights into the model's performance.

Figure 2 presents the training and validation accuracy over the course of the epochs. The training accuracy increased from 52.5% in the first epoch to 100% by the final epoch, indicating that the model effectively learned from the training data. Similarly, the validation accuracy also reached 100% towards the end of the training, suggesting that the model generalized well to the unseen validation data. This high accuracy is indicative of the model's strong performance; however, achieving 100% accuracy may also raise concerns about potential overfitting. It is crucial to ensure that the training and validation datasets are sufficiently diverse to mitigate this risk.

The model's performance was comprehensively evaluated using several metrics, including accuracy, loss, confusion matrix, and the Receiver Operating Characteristic (ROC) curve. The following analysis provides a detailed examination of each metric and offers insights into the model's performance.

Figure 2 presents the training and validation accuracy over the course of the epochs. The training accuracy increased from 52.5% in the first epoch to 100% by the final epoch, indicating that the model effectively learned from the training data. Similarly, the validation accuracy also reached 100% towards the end of the training, suggesting that the model generalized well to the unseen validation data. This high accuracy is indicative of the model's strong performance; however, achieving 100% accuracy may also raise concerns about potential overfitting. It is crucial to ensure that the training and validation datasets are sufficiently diverse to mitigate this risk.

Figure 3 illustrates the training and validation loss over the epochs. The training loss decreased

Figure 2. Training and validation accuracy.**Figure 3.** Training and validation loss.



significantly from 0.82 to 0.00188, while the validation loss dropped from 0.6775 to 0.0001258. This steady reduction in loss values indicates that the model improved consistently throughout the training process and maintained its performance on the validation set. The near-zero loss values for both training and validation phases underscore the model's effectiveness in minimizing error, although it is important to remain cautious about potential overfitting.

The confusion matrix shown in Figure 4 reveals a perfect classification result, with no false positives or false negatives. This matrix indicates that the model accurately identified all positive and negative cases, highlighting its reliability in distinguishing between classes. The absence of misclassifications reflects the model's high precision and effectiveness in the classification task.

Figure 4 shows the matrix with no misclassifications, reflecting the model's high accuracy in predicting the correct class for each sample.

Figure 5 presents the Receiver Operating Characteristic (ROC) curve, which demonstrates an Area Under the Curve (AUC) of 1.0. This perfect AUC indicates that the model has an exceptional ability to discriminate between positive and negative cases. The ROC curve reinforces the findings from the confusion matrix, showcasing the model's excellent performance in detecting the target class with an actual positive rate of 1.0 and no false positives. The results from the accuracy, loss, confusion matrix, and ROC curve collectively suggest that the model performs exceptionally well in detecting pipeline leaks. The high accuracy and low loss values indicate practical training and strong generalization capabilities. The perfect confusion matrix and ROC AUC further validate the model's robustness and reliability in classification tasks.

The results from the accuracy, loss, confusion matrix, and ROC curve collectively suggest that the model performs exceptionally well in detecting pipeline leaks. The high accuracy and low loss values indicate practical training and strong generalization capabilities. The perfect confusion matrix and ROC AUC further validate the model's robustness and reliability in classification tasks. Despite these positive results, it is important to consider the potential for overfitting, as indicated by the perfect performance across all metrics. Future work should involve evaluating the model on more diverse and realistic datasets to ensure its generalizability and effectiveness in real-world scenarios.

Figure 4. Confusion matrix.**Figure 5.** ROC curve.

Conclusion

This study evaluated a pipeline leak detection system that utilizes infrared thermography in conjunction with advanced machine learning techniques. Adapting a method based on a VGG16-based Faster R-CNN approach, we tested an infrared camera that does not natively measure temperature to see if it could achieve comparable performance. The model performed exceptionally well, reaching 100% accuracy and showing no misclassifications, as confirmed by the ROC curve's AUC of 1.0. However, perfect accuracy suggests potential overfitting, indicating a need for further validation on more diverse datasets. The performance of this model has achieved better results than those shown in Xie and colleagues [6]. However, it is necessary to consider that this work has been tested only in a laboratory in one single pipeline; the work

developed in Xie and colleagues [6] was tested in an environment more similar to an operational system and has more complex structures.

While our findings are promising, future research should focus on testing the model in more complex environments, improving model robustness against overfitting, and exploring alternative infrared cameras that might enhance detection accuracy. Additionally, implementing the model in real-time systems and conducting comparative analyses with other leak detection methods will help to assess its real-world applicability and identify areas for improvement.

References

1. Dogan E, Koc S, Koc M. Pipeline leak detection using infrared thermography: a review. *J Pet Sci Eng.* 2021;196:107741. Available from: <https://www.sciencedirect.com/science/article/pii/S0920410521000544>.
2. Liu Y, Xu W, Li Y. Advanced techniques for pipeline leak detection: a review and future directions. *Sensors* (Basel). 2022;22(15):5604. Available from: <https://www.mdpi.com/1424-8220/22/15/5604>.
3. Yao H, Huang J, Zhang C. A comprehensive survey on pipeline leak detection methods and technologies. *IEEE Trans Instrum Meas.* 2023;72:3000107. Available from: <https://ieeexplore.ieee.org/document/9428223>.
4. Kumar A, Sharma S, Gupta R. Enhanced pipeline leak detection using thermal imaging and machine learning techniques. *Sensors* (Basel). 2023;23(4):2123. Available from: <https://www.mdpi.com/1424-8220/23/4/2123>.
5. Cheng X, Li S, Wang C. Integration of infrared camera systems and AI for pipeline leak detection and monitoring. *Comput Environ Urban Syst.* 2023;90:101803. Available from: <https://www.sciencedirect.com/science/article/pii/S0197397523000834>.
6. Xie J, Zhang Y, He Z, Liu P, Qin Y, Wang Z, et al. Automated leakage detection method of pipeline networks under complicated backgrounds by combining infrared thermography and Faster R-CNN technique. *Process Saf Environ Prot.* 2023;174:39-52. Available from: <https://www.sciencedirect.com/science/article/pii/S0957582023002938>.
7. BSI Engineering. Fluid thermal expansion. BSI Engineering. 2020. Available from: <https://bsiengr.com/wp-content/uploads/2020/12/Fluid-Thermal-Expansion.pdf>. White Paper – Vol. 11.