

Multi-Tier Network Architecture for Real-Time Cocoa Pathology Estimation through Lightweight Deep Learning

Ali Ahmad¹[0000-0001-5530-7374], Vinie Lee Silva-Alvarado¹[0009-0000-5857-3248], Sandra Sendra¹[0000-0001-9556-9088] (*), and Jaime Lloret¹ [0000-0002-0862-0533]

¹ Instituto de Investigación para la Gestión Integrada de Zonas Costeras, Universitat Politècnica de València, Carrer del Paranimf 1, 46730 Grao de Gandia (Valencia), Spain
aahmad1@upv.es, vlsilalv@doctor.upv.es, jllloret@dcom.upv.es,
sansenco@upv.es

Abstract. This study presents a hierarchical hybrid framework for real-time detection of major cocoa diseases and pests in precision agriculture. The proposed system integrates ESP32-S3 edge nodes with a Raspberry Pi 5 gateway through a dual Wi-Fi and LoRaWAN communication framework, enabling reliable operation across large plantation environments. Visual inference is executed locally using an enhanced YOLOv8n model incorporating Squeeze-and-Excitation attention and C3Ghost lightweight feature extraction modules, thereby limiting data transmission to compact detection metadata and reducing communication overhead. Experimental evaluation indicates that the proposed YOLOv8n+SE+C3Ghost architecture improves disease detection sensitivity, achieving higher recall and overall fitness than the baseline YOLOv8n, while reducing the total number of model parameters by 23.63%. Although the inclusion of attention mechanisms increases computational load and inference latency, the model maintains real-time performance suitable for edge deployment. Explainable artificial intelligence analysis further demonstrates that the modified architecture produces more coherent and spatially focused activation over infected pod regions, particularly for frosty pod rot, while suppressing background noise. The combination of improved detection robustness, reduced model complexity, and resilient long-range connectivity establishes the proposed framework as an effective and scalable solution for early phytosanitary monitoring in cocoa plantations and beyond.

Keywords: Precision agriculture, edge AI, fruits, deep learning, computer vision.

1 Introduction

Cocoa (*Theobroma cacao* L.) is a major plantation crop cultivated in tropical regions, underpinning the global chocolate industry and supplying raw materials increasingly demanded by the cosmetic and pharmaceutical sectors[1], [2]. Despite its economic

importance, cocoa production is severely constrained by biotic stresses, which significantly reduces yield stability and farmer income.

Approximately 40% of global cocoa production is lost annually to major diseases, including black pod caused by *Phytophthora* species, witches' broom (*Moniliophthora perniciosa*), cocoa swollen shoot virus disease, vascular streak dieback (*Oncobasidium theobromae*), and frosty pod rot (*Moniliophthora roreri*) [2]. Black pod disease is widespread, with multiple *Phytophthora* species reported across cocoa-growing regions [3], while frosty pod rot remains particularly destructive in Latin America, where losses range from moderate yield reductions to complete crop failure [4]. In parallel, insect pests further exacerbate production losses. Among more than fifty cocoa-associated insect species, mirid bugs, especially the tea mosquito bug complex (*Helopeltis theivora*, *H. bradyi*, and *H. antonii*), cause extensive damage to cherelles, shoots, and pods, leading to yield losses of up to 40% in affected regions [5], [6].

Early and accurate detection of disease symptoms and pest damage on cocoa fruits is therefore essential for effective crop management. Traditional approaches based on visual inspection and laboratory diagnostics are labor-intensive, subjective, or poorly suited for rapid in-field decision-making. Recent advances in computer vision and deep learning have enabled real-time object detection using lightweight models deployable on edge devices; however, their practical adoption depends on achieving a balance between accuracy, inference speed, and computational efficiency. In this context, this study evaluates the YOLOv8n model for simultaneous detection of black pod, frosty pod rot, and mirid bug damage on cocoa fruits and proposes an enhanced YOLOv8n+SE+C3Ghost architecture that integrates channel attention and lightweight feature extraction. The proposed model demonstrates improved detection performance with competitive computational efficiency, offering a unified vision-based framework for real-time monitoring of major cocoa fruit diseases and insect pests.

The remainder of this manuscript is structured as follows. Section 2 presents a detailed review of related work, followed by the methodological framework and model architecture in Section 3. The experimental setup and test bench are described in Section 4, while the results and performance analysis are discussed in Section 5. Section 6 concludes the study and outlines future research directions

2 Related Work

Research on automated crop disease and pest identification has progressed from rule-based image processing toward data-driven machine learning and deep learning frameworks [7]. Early studies relied on handcrafted features and deterministic rules, exemplified by the fuzzy rule-based FRADD system proposed by Kour and Arora [8], which achieved 91.66% accuracy for apple scab detection but remained sensitive to illumination changes, background complexity, and occlusion. Such limitations motivated the adoption of convolutional neural networks, which enable automatic feature learning and improved robustness under field conditions. Within cocoa production, lightweight CNN architectures have demonstrated strong performance for disease recognition. For instance, Godmalin, et al. [9] evaluated EfficientNetB0, MobileNetV3Small, and

NASNetMobile for cacao pod classification, reporting superior performance for EfficientNetB0 with accuracy, recall, and F1-score of 94%. Further work by the same authors extended CNN-based approaches to disease severity classification, where MobileNetV3Small achieved 91% accuracy on a dataset of 993 images [10]. In parallel, hybrid pipelines combining classical feature extraction with shallow classifiers were explored. Achmad and Areni [11] applied Local Binary Pattern features with Artificial Neural Networks for cocoa pest detection, attaining limited accuracy of 66%, while subsequent comparisons of LBP, GLCM, HSV, and GLCH features with Support Vector Machines indicated that HSV-based representations improved performance to 80.95%, yet remained inferior to deep CNN-based methods [12].

Recent studies have increasingly focused on end-to-end object detection models optimized for mobile and edge deployment. For example, Kumi, et al. [13] demonstrated the feasibility of smartphone-based cocoa disease detection using SSD MobileNet V2, achieving confidence scores above 80% for swollen shoot and black pod diseases. In a related industrial application, Achmad and Areni [14] demonstrated that carefully engineered color and texture features combined with optimized SVM classifiers could achieve near-perfect accuracy for cocoa bean quality assessment. Advances in one-stage detectors have further improved performance, with YOLOv9 achieving precision and recall above 99% for binary cocoa pod rot detection [15], and YOLOv5 attaining an F1-score of 95.65% on a benchmark cocoa dataset [16]. Nevertheless, these approaches primarily address binary or single-disease scenarios. Multi-class, fruit-level detection of black pod, frosty pod rot, and mirid bug damage remains largely unexplored, particularly under real-time constraints. In contrast, the proposed YOLOv8n + SE + C3Ghost architecture integrates channel-wise attention and lightweight feature generation to enhance discrimination of subtle disease and pest symptoms while maintaining computational efficiency, thereby addressing key limitations identified in existing literature [17].

3 Methodology

This section describes the synergistic integration of the hardware and software frameworks. First, the hierarchical network topology and node operation logic are detailed. Subsequently, the structural design of the optimized YOLOv8n model is presented. Finally, the specific training parameters and dataset distribution are defined.

3.1 Node Architecture and Wireless Sensor Network (WSN)

The proposed hierarchical network topology is illustrated in Fig. 1 [18]. The architecture integrates Wi-Fi and LoRaWAN communication protocols. Edge clusters consist of multiple ESP32-S3 sensing nodes. These nodes perform high-speed data exchange via Wi-Fi. A Raspberry Pi 5 unit serves as the local aggregator. The aggregator manages data within a 50-meter radius. Local frames are converted into long-range LoRaWAN packets. Information travels several kilometers to the master station. This hybrid

approach ensures data backhaul efficiency. The power consumption remains optimal at the sensing level.




Fig. 1. Illustration of WSN architecture for cocoa disease detection.

The proposed logical workflow adopted in this study is presented in Fig. 2. The process initiates by waking up the sensing node. Deep sleep states minimize the total energy consumption. The ESP32-S3 module captures a high-resolution crop image. The diseases model detection executes visual inference at the edge. Detected pathologies trigger the packaging of specific metadata. The node establishes a secure local Wi-Fi connection. Data packets are forwarded to the nearby aggregator. Absence of anomalies initiates an immediate sleep return. This sequence optimizes the sensor's operational duty cycle.




Fig. 2. Proposed node operation algorithm

The SE module recalibrates channel-wise features through global average pooling. This mechanism prioritizes relevant textures of cocoa pod diseases [19]. It suppresses background noise from complex tropical agricultural environments. The reduction ratio is set to sixteen for efficiency. This ensures high performance without increasing the model size. The neck uses C3Ghost modules to replace standard convolutions C2f. Ghost convolutions generate more features using fewer linear operations. This design significantly reduces the total number of parameters. It maintains a rich representation for multi-scale disease detection. The architecture remains lightweight for real-time edge device deployment. Layers 13 to 28 focus on hierarchical feature extraction. The model integrates spatial information and semantic context effectively. Four detection heads process different scales of cocoa fruits. This approach improves the detection of small mirid damage. High-resolution features are preserved for accurate disease localization.

3.4 Model Training and Explainable Artificial Intelligence Analysis

The model was trained for 100 epochs. A batch size of 32 and 640×640 resolution were used. This training configuration is directly informed by the dataset characteristics visualized in Fig. 4. The dataset [22] was partitioned into three distinct functional subsets. Specifically, 70% of the samples were for training. The remaining 30% was split equally for validation and test. Whereas Explainable Artificial Intelligence (XAI) techniques were applied to interpret and compare the decision-making behavior of the baseline YOLOv8n and the proposed YOLOv8n+SE+C3Ghost models. Gradient-based class activation mapping was employed to generate heatmaps highlighting image regions contributing most strongly to model predictions. A representative Frosty Pod Rot sample was selected to assess how architectural modifications influenced spatial attention under complex field conditions.




Fig. 4. Visualization of the training dataset labels: (a) Instance count per category where A is Black Pod Rot, B is Frosty Pod Rot, C is Healthy Pod, and D is Mirid Bug Damage; (b) Bounding box distribution across all frames; (c) Spatial distribution of object centers (x, y); (d) Correlation and distribution of object dimensions (width, height).

4 Results

In this section, the experimental results are presented. The performance of the modified YOLOv8n model is evaluated through quantitative metrics. First, results of the detection accuracy measured using mAP are provided. Then, the model evolution is

presented across the training epochs to ensure convergence stability. Finally, the results for XAI are presented followed by limitations and future prospective.

4.1 Performance of the Rust Severity Index Calculation

The numerical stability and learning efficiency of the proposed architectural modifications are illustrated in Fig. 5. In Fig. 5a both models show rapid precision growth during early epochs. Both models show rapid precision growth during early epochs. The baseline (YOLOv8n) performs better during the first 60 epochs. Then, the modified architecture (YOLOv8n + SE+ C3Ghost) surpasses the baseline until the end. The SE module and C3Ghost helps the network focus on discriminative features. The proposed model achieves superior refinement in the final stage.

In Fig. 5b the baseline YOLOv8n model exhibits a lower total loss. This lower value often indicates a simpler feature mapping. However, it might also suggest a risk of overfitting. The modified architecture model shows a continuous improvement trend. The attention mechanism requires more time to stabilize weights. The proposed model prioritizes complex features over simple errors. A higher loss does not prevent achieving superior precision.




Fig. 5. Training dynamics across 100 epochs. (a) Evolution of mAP@50; (b) Convergence of Total Loss.

4.2 Overall Performance

In Table 1 summarizes the data obtained during the experimental phase for each architectural variant. The optimized model achieves a reduction of 23.63% in total parameters. This decrease is primarily due to the C3Ghost modules. These modules generate feature maps using lightweight linear operations. Despite having fewer parameters, the computational cost (GFLOPs) increased by 11.36%. This rise is caused by the additional processing load of the SE attention mechanism. Consequently, the inference time increased by 32.98%. However, this latency remains sufficiently low to guarantee real-

time performance on edge computing devices. The models were evaluated using the validation set.

Table 1. Ablation study of the proposed architectural modifications for coffee leaf rust detection

Model	mAP@50	mAP@50:95	Parameters	GFLOPs	Inf (ms)
Yolov8n	0.8215	0.6147	3006428	8.088	4.73
Yolov8n + SE + C3Ghost	0.8466	0.6458	2295858	9.007	6.29

The metric improvements are visually compared in the radar chart as shown in Fig. 6. The proposed model excels in all categories. The performance metrics reveal a distinct trade-off between the model variants. The baseline YOLOv8n architecture achieves a higher Precision of 68.46%. This represents a superior value compared to the 64.64% obtained by the YOLOv8n + SE + C3Ghost model. This indicates that the baseline is more accurate in its positive detections. However, the modified model architecture demonstrates a higher Recall of 78.17%. This suggests that the structural modifications improve the ability to detect a larger total number of disease instances. The Fitness score reflects the overall effectiveness of the models. The baseline YOLOv8n reaches a fitness value of 50.43%. In contrast, the modified architecture shows a fitness of 54.51%. While the modified model version captures more lesions (higher Recall), it also introduces more false positives (lower Precision). The models were evaluated using the test set.




Fig. 6. Radar chart comparison of model performance metrics for test set.

The model performance is further analyzed using a normalized confusion matrix. Fig. 7 illustrates the classification accuracy across all disease categories. The diagonal values represent high true positive rates for each class. Some minor confusion exists between diseased pods and the background. This is expected due to complex agricultural environment textures.

Compared with prior cocoa-focused studies, most existing works emphasize either image-level classification or binary disease detection, whereas the present framework performs simultaneous multi-class, fruit-level object detection under real-time constraints. While higher-capacity detectors such as YOLOv9 report near-saturated

accuracy for binary tasks [15], the proposed architecture demonstrates a more favorable balance between accuracy, computational efficiency, and deployment feasibility for resource-constrained agricultural environments.




Fig. 7. Confusion matrix of test set for (a) YOLOv8n and (b) YOLOv8 + SE + C3Ghost.

Table 2. Comparative analysis of proposed YOLOv8n+SE+C3Ghost model against related cocoa disease and pest detection studies.

Target task	Model / Method	Detection / Classification Level	Performance metrics	Model efficiency	Reference
Cocoa pest attack	LBP + ANN (ELM)	Image-level classification	Accuracy = 66%	Fast inference (0.59 s)	[11]
Cocoa disease detection	SSD MobileNet V2	Object detection	Confidence > 80%, mAP $\approx$ 30	Mobile-friendly	[13]
Cocoa disease detection	YOLOv5m	Object detection	F1 = 95.65%	Moderate complexity	[16]
Cocoa pod rot	YOLOv9	Object detection	mAP@50 = 0.99	High-capacity detector	[15]
Frosty pod rot stages	MobileNetV2	Image-level classification	Accuracy = 95%, F1 = 0.94	Lightweight CNN	[17]
Black pod + frosty pod rot + mirid bug	YOLOv8n + SE + C3Ghost	Fruit-level object detection	mAP@50 = 0.8466; mAP@50:95 = 0.6458	2.30 M params; 9.01 GFLOPs; 6.29 ms	This study

4.3 XAI Analysis Results

The baseline YOLOv8n exhibited dispersed and localized activations, primarily focusing on limited pod regions, indicating a weak association between global disease texture and prediction confidence (Fig. 8b). In contrast, the YOLOv8n+SE+C3Ghost model produced broader and more coherent activation patterns that covered most infected pod surfaces with high-confidence responses (Fig. 8c). These results indicate

that the Squeeze-and-Excitation mechanism enhanced channel-wise feature recalibration, enabling the model to emphasize disease-specific visual cues while suppressing background interference from foliage and branches. Overall, the XAI analysis confirms that the proposed architecture improves interpretability and localization reliability for frosty pod rot detection.




Fig. 8. Compared the heatmaps. (a) Original image; (b) YOLOv8n + heatmap; (c) YOLOv8n + SE + C3Ghost + heatmap.

4.4 Limitations and Future Prospective

The current visual detection system is constrained by environmental and technical factors. Recognition accuracy is sensitive to ambient lighting, and dense cocoa canopies may impede Wi-Fi signal propagation. LoRaWAN's limited bandwidth also restricts the transmission of high-resolution imagery.

Future work will focus on expanding the training dataset to improve model generalization and enhance precision and recall. Integrating adaptive energy harvesting could increase node autonomy, while multi-spectral sensors may enable earlier disease detection. Mesh networking is another avenue to extend coverage of local sensing clusters, and field validation across diverse regions will assess system robustness under varying environmental conditions.

5 Conclusion

This study confirms the potential of edge AI for real-time phytosanitary monitoring in cocoa plantations. By performing visual inference directly on ESP32-S3 edge nodes, the proposed framework reduces latency and communication overhead while avoiding reliance on cloud-based processing. The enhanced YOLOv8n+SE+C3Ghost model supports this edge-centric design by improving detection performance within the computational constraints of field-deployable devices. The hybrid network topology, combining short-range communication with LoRaWAN backhaul, enables scalable deployment across large and remote plantations. The integration of SE and C3Ghost modules enhances feature discrimination while limiting model complexity, ensuring reliable

detection of disease and pest symptoms under real field conditions. Future work will focus on dataset expansion and further architectural optimization to improve generalization and robustness.

Funding. Part of this work was funded by the Grant PRE2021-100809 funded by Agencia Estatal de Investigación (AEI) MICIU/AEI/10.13039/501100011033 and by European Social Fund Plus (ESF+). The author V.L.S.-A. is supported by a contract from the Universitat Politècnica de València through the “Programa de Ayudas de Investigación y Desarrollo (PAID-01-24)”. This work has also received funding from the PRIMA Programme supported by the European Union, project PAVER, and it is supported by MCIN/AEI/10.13039/501100011033. Funders had no role in the design and execution of the study.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article

References

1. G. S. Madhu, A. T. Rani, B. M. Muralidhara, S. Rajendiran, V. Venkataravanappa, and S. Sriram, “First report of *Phytophthora tropicalis* causing black pod of Cacao (*Theobroma cacao*) in India,” *Australasian Plant Pathology*, vol. 52, no. 6, pp. 591–593, 2023.
2. J. S. Minimol, B. Suma, A. Mary, T. K. Shija, S. Jose, and S. K. Plappallil, “Evaluating black pod resistance in cocoa (*Theobroma cacao* L.) and its relationship with phenotypical and biochemical characteristics,” *Beverage Plant Research*, vol. 4, no. 1, pp. 0–0, 2024, doi: 10.48130/bpr-0024-0028.
3. G. S. Madhu, A. T. Rani, B. M. Muralidhara, S. Rajendiran, V. Venkataravanappa, and S. Sriram, “First report of *Phytophthora tropicalis* causing black pod of Cacao (*Theobroma cacao*) in India,” *Australasian Plant Pathology*, vol. 52, no. 6, pp. 591–593, Nov. 2023, doi: 10.1007/s13313-023-00947-3.
4. D. L. Jiménez, J. C. Alvarez, and S. Mosquera, “Frosty pod rot: a major threat to cacao plantations on the move,” *Trop Plant Pathol*, vol. 47, no. 2, pp. 187–200, Oct. 2021, doi: 10.1007/s40858-021-00472-y.
5. C. Mohan, A. Josephraj Kumar, S. H. Thube, E. K. Saneera, and M. Mani, “Pests and Their Management in Cocoa,” in *Trends in Horticultural Entomology*, Singapore: Springer Nature Singapore, 2022, pp. 1441–1455. doi: 10.1007/978-981-19-0343-4_61.
6. S. H. Thube et al., “*Xylosandrus crassiusculus* (Motschulsky) on Cocoa Pods (*Theobroma cacao* L.): Matter of Bugs and Fungi,” *Insects*, vol. 13, no. 9, p. 809, Sep. 2022, doi: 10.3390/insects13090809.
7. G. Abbas, A. Mehmood, M. Carsten, G. Epiphaniou, and J. Lloret, “Safety, Security and Privacy in Machine Learning Based Internet of Things,” *Journal of Sensor and Actuator Networks* 2022, Vol. 11, Page 38, vol. 11, no. 3, p. 38, Jul. 2022, doi: 10.3390/JSAN11030038.
8. V. P. Kour and S. Arora, “Fruit Disease Detection Using Rule-Based Classification,” in *Smart Innovations in Communication and Computational Sciences: Proceedings of ICSICCS-2018*, Springer, 2019, pp. 295–312. doi: 10.1007/978-981-13-2414-7_28.
9. R. A. Godmalin, C. J. Aliac, and L. Feliscuzo, “Classification of Cacao Pod if Healthy or Attack by Pest or Black Pod Disease Using Deep Learning Algorithm,” in *2022 IEEE*

- International Conference on Artificial Intelligence in Engineering and Technology (IICAIET), IEEE, Sep. 2022, pp. 1–5. doi: 10.1109/IICAIET55139.2022.9936817.
10. R. A. Godmalin, C. J. Aliac, and L. Feliscuzo, “Cacao Pod Infection Level Classification Using Transfer Learning,” in 2023 IEEE Open Conference of Electrical, Electronic and Information Sciences (eStream), IEEE, Apr. 2023, pp. 1–6. doi: 10.1109/eStream59056.2023.10135062.
 11. Basri, Indrabayu, A. Achmad, and I. S. Areni, “Classification Model Evaluation with Feature Extraction for Pest Attacks Identification on Cocoa Pods,” in 2023 7th International Conference on New Media Studies (CONMEDIA), IEEE, Dec. 2023, pp. 150–155. doi: 10.1109/CONMEDIA60526.2023.10428262.
 12. Basri, Indrabayu, A. Achmad, and I. S. Areni, “Comparison of Image Extraction Model for Cocoa Disease Fruits Attack in Support Vector Machine Classification,” in 2022 International Conference on Electrical and Information Technology (IEIT), IEEE, Sep. 2022, pp. 46–51. doi: 10.1109/IEIT56384.2022.9967910.
 13. S. Kumi, D. Kelly, J. Woodstuff, R. K. Lomotey, R. Orji, and R. Deters, “Cocoa Companion: Deep Learning-Based Smartphone Application for Cocoa Disease Detection,” *Procedia Comput Sci*, vol. 203, pp. 87–94, 2022, doi: 10.1016/j.procs.2022.07.013.
 14. B. Basri, I. Indrabayu, A. Achmad, and I. S. Areni, “Cocoa bean quality identification using a computer vision-based color and texture feature extraction,” *International Journal of Advances in Intelligent Informatics*, vol. 11, no. 1, p. 86, Feb. 2025, doi: 10.26555/ijain.v11i1.1609.
 15. S. Sudjud, M. Jamil, and R. Melati, “Development of Cocoa Pod Rot Disease Identification Model Using You Only Look Once (YOLO-v9),” *Mathematical Modelling of Engineering Problems*, vol. 12, no. 4, Apr. 2025, doi: 10.18280/mmep.120418.
 16. G. Leon and J. Yauri, “A Benchmark Dataset and Baseline Model for Disease Detection in Peruvian Cocoa Fruits Using YOLOv5,” in *Proceedings of the 2025 8th International Conference on Systems Engineering - Cybersecurity & AI: Building a reliable digital future*, New York, NY, USA: ACM, Oct. 2025, pp. 23–30. doi: 10.1145/3771678.3771682.
 17. J. Alvarado, J. F. Restrepo-Arias, D. Velásquez, J. D. Martinez-Vargas, and M. Maiza, “Precision Agriculture Applied to Cocoa: ML Solutions for Classification of *Moniliophthora Roreri*,” in *International Conference on Intelligent and Fuzzy Systems*, 2025, pp. 79–87. doi: 10.1007/978-3-031-98565-2_10.
 18. D. Bri, M. Garcia, J. Lloret, and P. Dini, “Real deployments of wireless sensor networks,” *Proceedings - 2009 3rd International Conference on Sensor Technologies and Applications, SENSORCOMM 2009*, pp. 415–423, 2009, doi: 10.1109/SENSORCOMM.2009.69.
 19. J. Hu, L. Shen, S. Albanie, G. Sun, and E. Wu, “Squeeze-and-Excitation Networks,” *IEEE Trans Pattern Anal Mach Intell*, vol. 42, no. 8, pp. 2011–2023, 2020, doi: 10.1109/TPAMI.2019.2913372.
 20. K. Han, Y. Wang, Q. Tian, J. Guo, C. Xu, and C. Xu, “GhostNet: More Features From Cheap Operations,” in 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2020, pp. 1577–1586. doi: 10.1109/CVPR42600.2020.00165.
 21. G. Jocher, J. Qiu, and A. Chaurasia, “Ultralytics YOLO,” 2023, Ultralytics. [Online]. Available: <https://github.com/ultralytics/ultralytics>
 22. Roboflow, “CACAO Computer Vision Model.” Accessed: Dec. 10, 2025. [Online]. Available: <https://universe.roboflow.com/r3-r9fts/cacao-otprq>