

CDAC: Content-Driven Access Control Architecture for Smart Farms

Ghadeer I. Yassin¹ ^a and Lakshmish M. Ramaswamy² ^b

¹*School of Computing, Binghamton University, Binghamton, NY, U.S.A.*

²*School of Computing, University of Georgia, Athens, GA, U.S.A.*

Keywords: Smart Farm, Internet of Things, Edge Machine Learning, Image Classification, MobileNet, Access Control.

Abstract: One of the smart farms' pivotal components involves leveraging vast quantities of imagery data to inform decision-making and improve farm outcomes. With the increasing integration of image data in smart farms, ensuring secure and efficient access to these data sets is crucial. This paper proposes a novel Content-Driven Access Control (CDAC) architecture designed specifically for smart farming environments, where access requests to image data are evaluated based on the visual content of the images. The CDAC architecture employs a novel technique to assess the relevance of access requests to specific image contents by enriching access control requests with useful image content information with the help of an edge machine learning classification model that provides a fast and small-weight solution to classify images near their source in the smart farm. This approach goes beyond traditional access control methods by considering the information within images, allowing for more granular and content-aware permissions. To validate the effectiveness of the CDAC architecture, a series of experiments were conducted using a dataset of agricultural images. Results demonstrate that the proposed architecture is a valuable solution for regulating access to smart farm images based on the visual content of the images. Additionally, the architecture is proven suitable for deployment on smart farm edge devices.


1 INTRODUCTION


Smart farms represent a transformative paradigm in modern agriculture, where the integration of cutting-edge technologies has given rise to a wealth of collected images that redefine how we understand and manage farm and agricultural operations. Cameras, drones, and sensors deployed across smart farms capture diverse visual data, ranging from real-time snapshots of crops and livestock to high-resolution aerial views of entire fields. These images serve as invaluable datasets, providing insights into crop health, environmental conditions, and the overall status of the farm. The collected imagery forms the foundation for informed decision-making, enabling farmers and stakeholders to monitor, analyze, and respond promptly to dynamic agricultural scenarios. However, the sheer volume of image data generated necessitates sophisticated approaches to image management, including restriction to image retrieval based on their contents.

The current landscape of access controls in smart

communities, and specifically in smart farming, highlights significant shortcomings, particularly in the domain of regulating access to image data. Traditional access control mechanisms predominantly rely on broad resource attributes such as image source or given title, overlooking the nuanced nature of image content.

In implementing robust security measures for governing access to images based on their contents within the realm of smart farming, the integration of image classification models is paramount. However, image classification for access control purposes faces multiple challenges. Those challenges include the security concerns raised by transmitting sensitive farm images to the cloud-based systems where image classification models can be deployed and the impracticality of deploying these models on smart farm edge limited capability devices as a counter solution. Conversely, there's a pressing demand for instantaneous classification of farm-collected images to expedite access control decisions. Moreover, banking solely on image classification results for access control decisions proves inefficient due to the inability to ensure a 100% accuracy rate from the classification models.

^a  <https://orcid.org/0000-0001-5135-7622>

^b  <https://orcid.org/0000-0002-4567-4186>

To tackle these limitations, our paper introduces a novel smart farm content-driven access control architecture composed of an edge machine learning classifier that provides a fast and small-weight solution to classify images near their source in the smart farm and employs a semantic attribute-based access control mechanism that enriches access control requests with useful image content data to reduce undesired access to farm images. Our research contribution includes:

- Identifying the limitations of traditional access control methods in regulating accessibility to image resources based on their visual contents, paving the way for more granular and content-aware access control solutions.
- Creating a novel content-driven access control architecture for the smart farm domain.
- Introducing an instant classification technique of newly captured farm images enabling swift access control decisions while considering smart farm images' different contents security levels.
- Assessing the effectiveness of the suggested architecture on edge devices to better align with the unique requirements of smart farming.

The paper is structured as follows. In Section 2, we draw the motivation behind this research and in Section 3 we lay out important background information. While, in Section 4, we provide a detailed overview of the proposed architecture. In Section 5 we describe the prototype implementation of the architecture as well as the different evaluation experiments. In Section 6 we discuss our evaluation results and in section 7, we discuss important related works. Finally, in Section 8, we draw our conclusions.

2 MOTIVATION

Our research is motivated by the need to address the limitations of existing access control solutions in managing image data access within smart farming systems. In these systems, a wide range of stakeholders with different roles or specialties might require them to access and examine different imagery data related to their work. Conversely, it is imperative to enforce restrictions to prevent these stakeholders from accessing or reviewing imagery data outside the scope of their expertise. Take, for instance, an agronomist dedicated to fruit cultivation within the framework of a smart farming system. The agronomist's proficiency centers on the meticulous evaluation of fruit crops, requiring the analysis of visual data to assess factors such as health, growth, and potential challenges specific to various fruits. In this context, the

agronomist necessitates exclusive access to images featuring fruits. However, their role does not extend to the examination of all-encompassing smart farm imagery, which may include images of other stakeholders working inside the smart farm, agricultural equipment visuals, or livestock monitoring images. These broader datasets might encompass security-sensitive content not directly relevant to the agronomist's responsibilities.

To address the agronomist's access needs and safeguard other farm images, Access control policies should be meticulously formulated to grant the agronomist exclusive access to images featuring fruits while simultaneously restricting their access to images containing other contents. The access control solution not only enhances the agronomist's ability to focus on critical aspects of fruit health but also safeguards images unrelated to their specialization, contributing to the effectiveness of image management practices within the smart farming system.

However, The existing access control solutions in smart communities, particularly in smart farming, reveal substantial limitations, particularly regarding the management of image data access. Traditional access controls mainly focus on broad resource attributes such as image IRI, source, or provided title, failing to consider the nuanced nature of image content. This lack of granularity hampers control over stakeholders' access to specific content within the vast array of image data generated in smart farming. To overcome these limitations, a more sophisticated access control approach is needed for smart farming systems to ensure that stakeholders have precise access to the image data relevant to their expertise while maintaining security by limiting accessibility to other contents.

Efficient access control solutions within smart farms should identify images based on their contents and meticulous access control policies should be formulated to exclusively grant the stakeholders access to images bearing the specified content while concurrently imposing restrictions on their access to unrelated images. This bespoke approach fortifies security by preventing access to irrelevant or security-sensitive imagery resulting in an advancement in both the efficiency and security of image management practices within the smart farm. On the other hand, establishing robust security measures to manage access to images based on their content within the context of smart farming, requires the incorporation of machine learning models for image classification. While conventional methods, involving the processing of image data in cloud-based systems for machine learning tasks, exhibit effectiveness, they raise significant security concerns due to the transmission of sizable

volumes of potentially sensitive visual information.

This underscores the vital transition towards deploying such models at the edge of the smart farming network, closer to the data sources. However, the computational demands inherent in machine learning models encompass intensive computations, resulting in models that are computationally intensive and reliant on high-performance computing resources. These challenges are particularly pronounced when extending deep networks to edge devices in smart farming, where constraints such as limited battery capacity and memory further compound the intricacies.

Hence, there is an imperative need to develop models tailored to the constrained computational resources of edge devices in smart farms, recognizing the unique limitations of these environments. This strategic approach ensures not only enhanced security but also optimized performance in real-world applications within the context of smart farming.

Additionally, it's imperative to recognize that achieving 100% classification accuracy is unrealistic. Consequently, an access control architecture may face the risk of making decisions based on inaccurate classification results. This risk becomes particularly severe when granting access to images containing highly sensitive visual content. Conversely, denying access due to incorrect classification, while less severe, can still disrupt essential tasks. Therefore, it's essential for the access control architecture to address the imperfections of image classification results and implement measures to mitigate these risks effectively.

Furthermore, smart farms generate vast amounts of farm image data daily, which serve as the foundation for swift decision-making in smart farming operations. Therefore, the access control architecture should not restrict access solely to pre-classified images. Instead, it should integrate instant image classification techniques to ensure that access control requests are promptly evaluated based on image contents.

3 BACKGROUND INFORMATION

Enabling an access control solution to assess access requests to images based on their content requires comprehensive information about the image content. These details can be generated with the help of edge-suitable image classification models. Subsequently, enriching them within access control requests to serve as specifics for the targeted resource, namely the farm image. Consequently, the access control model can accurately evaluate the request based on the informa-

tion about the image content.

While various solutions for semantically enriching access requests with metadata about targeted resources exist, tailored approaches are necessary for smart farms due to their distinctive characteristics (Yassin and Ramaswamy, 2022). Hence, we build upon the semantically enriched access control architecture proposed in (Yassin and Ramaswamy, 2023) and incorporate within it a quantized Mobilnet model to perform image classification tasks on smart farm edge devices.

In this section, we briefly describe the important components of our proposed solution.

3.1 Semantically-Enriched Access Control




Figure 1: RE Component.

In the semantically enriched access control architecture (Yassin and Ramaswamy, 2023), the traditional ABAC architecture was modified to include a new RE component with two modules namely attribute processor and semantic reasoner as described in Figure 1. The attribute processor module extracts information from the access requests and forwards them to the semantic reasoner. Utilizing a smart farm-specific ontology, the semantic reasoner deduces supplementary attributes about the request entities, which are then returned to the attribute processor for incorporation into the request. This process empowers the Policy Decision Point (PDP) to make well-informed access control decisions.

3.2 Mobilnet Models

Mobilenets are machine learning classification models that are proven suitable for edge devices. It is a family of neural network architectures specifically crafted for efficient computation on mobile devices with constrained computing power and memory resources (Howard, 2017). It strategically employs

depth-wise separable convolutions, breaking down the standard convolution into depth-wise and point-wise convolutions. This approach significantly reduces the number of parameters and computations required for image recognition and classification tasks, enhancing overall efficiency. Furthermore, the MobileNet V2 model incorporates a technique known as 'bottlenecking' to further decrease computational complexity by compressing the input feature map before processing it with convolutional layers.

MobileNet models' architecture design, with its emphasis on depth-wise separable convolutions, aims to address the challenges of deploying deep learning models on mobile devices with limited resources. This makes it well-suited for smart farm applications. In this paper, MobileNet-V1 and MobileNet-V2 were selected for our use case.

3.3 Model Quantization

Quantization is a commonly applied technique to networks, aiming to reduce the number of bits needed to represent weights without compromising accuracy. One prevalent method is the utilization of lower precision formats such as 16-bit floating-point (FP16) and 32-bit floating-point (FP32) for representing weights during the quantization process (Cheng et al., 2017). This allows for a reduction in the memory footprint and computational demands of the model, making it more suitable for deployment on resource-constrained edge devices and mobile applications.

In our proposed architecture, we have employed the quantization approach as a key strategy for model compression. The optimal representation for our specific use case has been determined through an empirical analysis conducted directly on edge devices.

4 PROPOSED ARCHITECTURE

The pipeline of the proposed architecture is illustrated in Figure 2. Our proposed architecture employs a multi-path approach to enhance access control decisions. When the system receives an access request for image resources, it initiates a dynamic process to determine whether the targeted resource is a pre-classified image or a new, unclassified image.

When an unclassified image is received, it undergoes classification using an edge-suitable image classification model. This model analyzes the visual content, extracting information that directly enhances the real-time access request. Simultaneously, the classification details, such as the Image IRI (Internationalized Resource Identifier) and the probability distri-

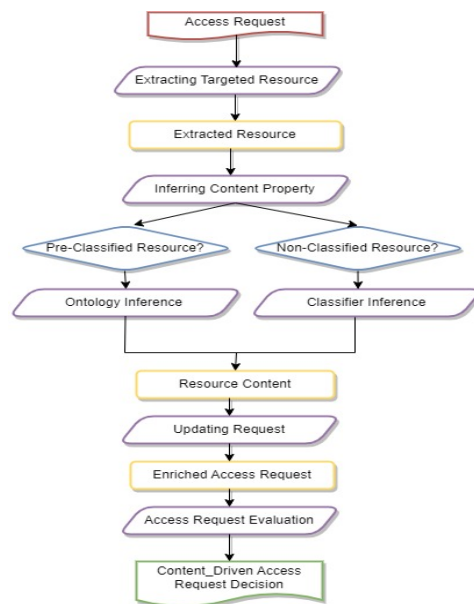


Figure 2: Proposed Pipeline.

bution of image content classes, seamlessly integrate into the smart farm ontology, creating a structured representation of the image content as described in Figure 3. This structured representation lays the foundation for subsequent access requests.

In cases where the image is pre-classified, the system taps into the semantic richness stored in the ontology to deduce image content information. This information is then utilized to enrich the access request.




Figure 3: Image Classification Pipeline.

In both scenarios, the access requests are enriched with relevant image content information. This augmentation empowers the system to make well-informed decisions about granting or revoking accessibility to images based on their contents rather than their general attributes (Source, Title, etc). This approach allows the system to gain a deeper understanding of the visual content, enabling it to make access decisions based on contextual insights derived from both pre-classified and dynamically classified images.

5 ARCHITECTURE SETUP

5.1 Dataset

The dataset selected for the model training and evaluation process is the Kaggle "Plants Type Dataset" which presents a robust collection of 30,000 high-resolution plant images. With a meticulous curation of 1,000 images per class, the dataset spans across 30 distinct plant classes and encompasses seven diverse plant types, including crops, fruits, vegetables, and herbs.

5.1.1 Dataset Split

To promote a comprehensive evaluation of model performance, The dataset was partitioned into training, validation, and testing subsets. The training set, which constitutes 70% of the data, is employed to train the model, enabling it to learn and generalize from diverse examples. Simultaneously, the validation set, comprising 15% of the dataset, is reserved for assessing the model's performance on unseen data, providing insights into its robustness. While the remaining 15% of the data were reserved for the phase of model testing.

5.1.2 Image Preprocessing

Before model training, a comprehensive preprocessing pipeline was implemented to optimize the dataset for enhanced neural network learning and generalization. Pixel values of images were normalized to a standardized range of [0, 1] to enhance the model's convergence during training. Simultaneously, all images were resized to uniform 224x224 dimensions, establishing a standardized input size to fit the input shape for the subsequent convolutional neural network input. Furthermore, images' randomness was introduced through shuffling images during training to enhance the model's adaptability by exposing it to varied patterns and features.

5.2 Experiments

For our experiments, we employed transfer learning techniques which involve retaining knowledge acquired while addressing one problem and applying it to solve a new, related problem enabling the model to adapt more swiftly to new data, a process that would be less efficient if started from scratch.

We utilized MobileNet models with pre-trained weights from ImageNet, a dataset containing millions of well-organized images, accessible globally to researchers (Deng et al., 2009). To enhance the models'

performance, fine-tuning was conducted in this study. Fine-tuning involves carefully adjusting the model's parameters to better respond to specific information. During this process, a trained model or a segment of it is unfrozen, and the training occurs anew on new data, utilizing a reduced learning rate. This slight modification of the already learned weights results in enhanced performance.

The primary method employed for fine-tuning in this investigation involves removing the last fully connected layer from selected pre-trained CNN models and replacing it with a new fully connected layer or layers, matching the number of classes in our dataset. In our case, we employed 30 classes due to the specific requirements of our dataset. It's important to note that we employed two strategies: one where the MobileNet models were used as feature-extractors with all layers frozen, and another with fine-tuning, where some layers were unfrozen. This dual approach allowed us to leverage the general features learned by MobileNet models while also adapting the model to the nuances of our specific classification task.

5.2.1 Hyperparameters Selection

The models were evaluated with the goal of classifying the 30 different classes found in the data set. A training duration of 25 epochs was utilized for calculating model metrics, and the choice of the number of epochs and learning rate parameters was made empirically. Past research has demonstrated promising outcomes with the use of Adam as an optimizer (Bera and Shrivastava, 2020), influencing our decision to adopt Adam as the optimizer for our models. Additionally, a batch size of 32 and a learning rate of 0.0001 were specified. The complete set of hyperparameters for all transfer learning models is outlined in Table 1.

Table 1: Hyperparamters For All Transfer Models.

Parameter	Value
Max Epochs	25
Min Patch size	32
Optimizer	Adam Optimizer
Learning Rate	0.0001
Loss Function	Categorical_Crossentropy

5.2.2 Model Quantization

To implement model quantization, the proposed framework utilizes the Tensorflow Lite Library, employing various quantization parameters to evaluate the impact on model size and performance. Specifically, we represented the selected model using both

16-bit floating-point precision (FP16) and 32-bit floating-point precision (FP32). Subsequently, a comprehensive evaluation was conducted on both representations to assess their effectiveness in balancing model size reduction and preserving accuracy.

5.3 Smart Farm Ontology

Building upon the classification results from our model, we extended the ontology detailed in (Yassin and Ramaswamy, 2023) to encompass information regarding classified images. Using the RDFLib Library, our code dynamically loads the ontology and systematically incorporates image instances into the graph.

Each image instance from the dataset is assigned to the ontology class 'cropImageResource' and is linked to class probabilities, reflecting the likelihood of the image belonging to one of the 30 possible categories, along with associated confidence scores. Additionally, we group confidence scores for major classes based on the classification of 30 classes as follows:

- **Fruits:** waterapple, pineapple, pomelo, guava, mango, papaya, orange, banana, cantaloupe, watermelon, coconut, bilimbi.
- **Vegetables:** soybeans, pepper chili, spinach, shallot, sweet potatoes, ginger, kale, long beans, eggplant, cucumber.
- **Crops:** tobacco, melon, paddy, cassava, corn.
- **Herbs:** galangal, curcuma, aloe vera.

For instance, consider an image classification results as: (Waterapple, 25%), (Mango, 30%), (Pepper, 20%), (Orange, 15%), (Bilimbi, 10%). Therefore, the image classification for the broader categories would be: (Fruits,80%),(Vegetables:20%).

The resulting RDF triples, encapsulating these associations, are serialized and seamlessly integrated back into the ontology, enriching it with pertinent image-related data and enhancing the knowledge representation within the Smart Farm ontology framework. This augmentation facilitates a more comprehensive understanding of the images by encompassing valuable insights derived from image classification results.

5.4 Request Enricher- RE

The RE component, outlined in Section 3, plays a pivotal role in our content-driven access control architecture and was modified as shown in Figure 4. Within the RE component, the Attribute Processor module extracts crucial attributes from the access control request, notably the targeted resource, which in this

context is a Smart Farm Collected Image. These attributes are then forwarded to either the Semantic Reasoner module or the Image classification model directly and infer information about the image content, specifically its class and confidence score. The inferred details are subsequently sent back to the Attribute Processor.

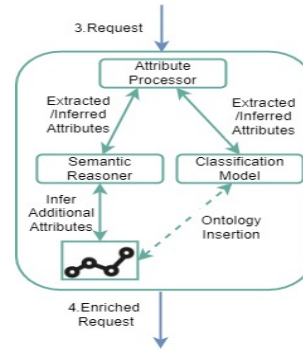


Figure 4: RE component changes.

5.4.1 Class Specific Threshold

Building upon the results of image classification, a single image may be assigned to multiple classes, each associated with distinct confidence scores. In this context, we consider a security-sensitivity level assigned to images corresponding to different classes to mitigate the imperfections of image classification results. For instance, images of herbs might carry a higher security sensitivity compared to images of fruits. Consequently, when the additional attributes are received, the Attribute Processor component undertakes a comparison between the image contents attributes and class-specific thresholds. If the confidence score surpasses the designated threshold, the image content attribute is deemed suitable for request enrichment; otherwise, it is disregarded.

5.4.2 Request Update

Upon receiving the additional attributes and comparing them to the specified sensitivity threshold, the Attribute Processor appends the selected attributes to the original request. The enriched request, now containing information about the targeted image content, is then sent to the Policy Decision Point (PDP) for evaluation. This content-driven approach enhances the access control decisions by incorporating detailed insights into the nature of the requested resource, In this context, the Image content.

5.5 Policy Administration

As detailed in Table 2, the access control policy set incorporates nuanced regulations governing access to resources (Smart Farm collected Images). This involves regulating access based on the content of the images and the subject's specialization, allowing them to access image contents in alignment with their expertise. The policies introduce the criteria considering the content of the images, and subject specializations. These policies are designed to facilitate a more granular and content-aware access control mechanism, ensuring that individuals can only access image content relevant to their specific area of expertise.

6 EVALUATION RESULTS

For evaluation purposes, The initial phases of evaluations targeted different models' evaluations, including training, validation, and testing, were performed on Google Colab. This cloud-based platform provided the computational resources necessary for the training phase, ensuring efficient model development.

The subsequent evaluation criteria were thoroughly tested on the Raspberry Pi 4 Model B to assess the performance and feasibility of our proposed system in a real-world, edge computing environment. The following evaluation criteria were tested:

- **Model Performance:** This assessment aims to identify the most effective transfer learning model configuration for the task by evaluating its performance across training, validation, and testing phases. The purpose is to discern which model excels in terms of accuracy, generalization, and effectiveness for the given application. This evaluation provides insights crucial for selecting the optimal model for deployment.
- **Model Evaluation on Edge:** Assessment of the performance of the selected transfer learning image classification model on a Raspberry Pi device. The evaluation includes inference time, quantization effects on accuracy, and model size considerations, aiming to identify the most suitable configuration for efficient deployment on edge devices.
- **Architecture Ability to Enrich Access Requests:** Assessment of the capability of the proposed architecture in handling access requests to images through enriching access requests with image content details.
- **Threshold Adjustment Effect:** Assessment of the effect of adjusting class-specific threshold in

mitigating the effect of making access request decisions based on inaccurate image classification results.

- **Average Request Evaluation Time:** Assessment of the time required on average for a single request to be processed on Edge device. The evaluation includes image content inference from both smart farm ontology and the image classification model.

6.1 Model Performance

Figures 5, 6, 7, and 8 illustrate the Accuracy and Loss graphs for each model during both training and validation phases. The training graphs depict the evolution of network accuracy at 5-epoch intervals, showcasing the model's progress during training. Additionally, the figures display cross-entropy loss, providing insights into the models' overall classification accuracy. The discrepancy between training and validation accuracy is minimal, indicating robust generalization on the training dataset and effective performance on the validation dataset. Finetuned MobilNet_V1 and Finetuned MobilNet_V2 have the best comparable outcomes, with Finetuned MobilNet_V2 showing a slightly superior performance. Notably, all models demonstrate validation accuracy values surpassing their training accuracy, and the figures suggest an absence of overfitting, with the models converging within 25 epochs.

The analysis presented in Table 3 emphasizes that the models achieved their highest testing accuracy when subjected to fine-tuning, notably with Finetuned MobilNet_V1 reaching 85% accuracy and Finetuned MobilNet_V2 achieving 86%, showcasing a slight superiority in accuracy.

This superior performance of Finetuned MobilNet_V2 can be attributed to its underlying architecture, characterized by inverted residuals, linear bottlenecks, and shortcut connections. When fine-tuned, this architecture exhibited enhanced capabilities for the specific task at hand. The fine-tuning process allowed the model to adapt and specialize, resulting in improved accuracy compared to Finetuned MobilNet_V1 and other configurations.

6.2 Model Evaluation on Edge

The optimal model chosen for deployment on Raspberry Pi edge device was MobileNet-V2 FineTuned based on its relatively superior performance. Additionally, to assess the deployment's performance, two different quantization techniques, FP16 and FP32, were applied to the model. Table 4 describes the model performance on Edge.

Table 2: Sample Polices from the Policy Set.

#	Subject	Subject Property	Action	Resource	Resource Property
1	Regular Worker	Speciality: Paddy	Read	Image Resources	Content: Paddy
2	Agronomist	Speciality: Crops	Read	Image Resources	Content: Crops
4	Researcher	Speciality: Cross-Category-Analysis	Read	Image Resources	Content: Crops, Herbs

Table 3: Different Models Evaluation Metrics.

Model	Training				Validation				Testing			
	Acc.	Prec.	Rec.	F1	Acc.	Prec.	Rec.	F1	Acc.	Prec.	Rec.	F1
FeatureExtractor MobileNet-v1	0.86	0.92	0.82	0.8678	0.87	0.92	0.83	0.8741	0.77	0.84	0.74	0.7856
Finetuned MobileNet-v1	0.96	0.96	0.96	0.9583	0.97	0.97	0.97	0.9705	0.85	0.86	0.85	0.8556
FeatureExtractor MobileNet-v2	0.89	0.95	0.82	0.8797	0.89	0.95	0.83	0.8798	0.80	0.88	0.74	0.8020
Finetuned MobileNet-v2	0.97	0.96	0.96	0.9575	0.97	0.97	0.97	0.9696	0.86	0.87	0.85	0.8599

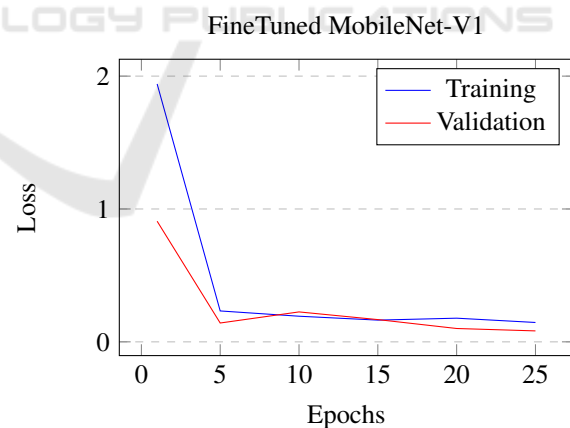

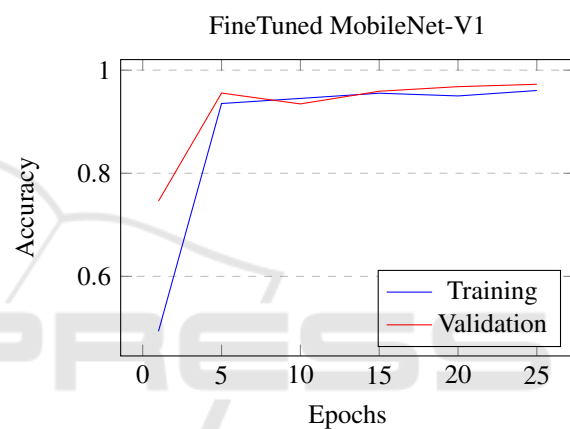
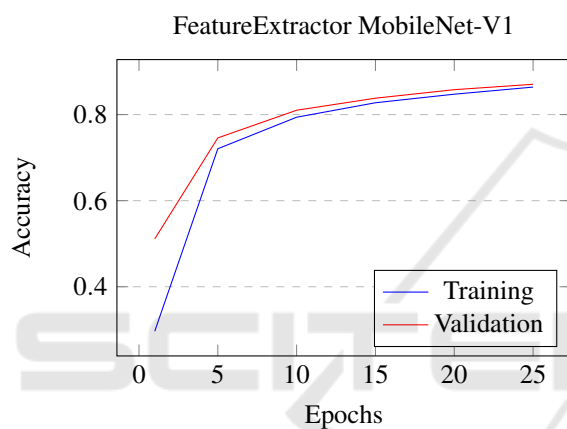


Figure 5: Accuracy and Loss for Training and Validation: FeatureExtractor MobileNet-V1.

Figure 6: Accuracy and Loss for Training and Validation: FineTuned MobileNet-V1.

1. **Inference Time:** This refers to the inference time on the edge device, which is the duration taken by the model from receiving the input to delivering the evaluated output. It provides insights into the speed or efficiency of the model. As depicted in Table 4, the inference time of quantized models was significantly reduced for both FP32 and FP16

quantized models scoring 0.16 and 0.17 seconds respectively in comparison with the model without quantization which scored an average inference time of 0.62 seconds. This improvement enhances the suitability of these models for deployment on edge devices.

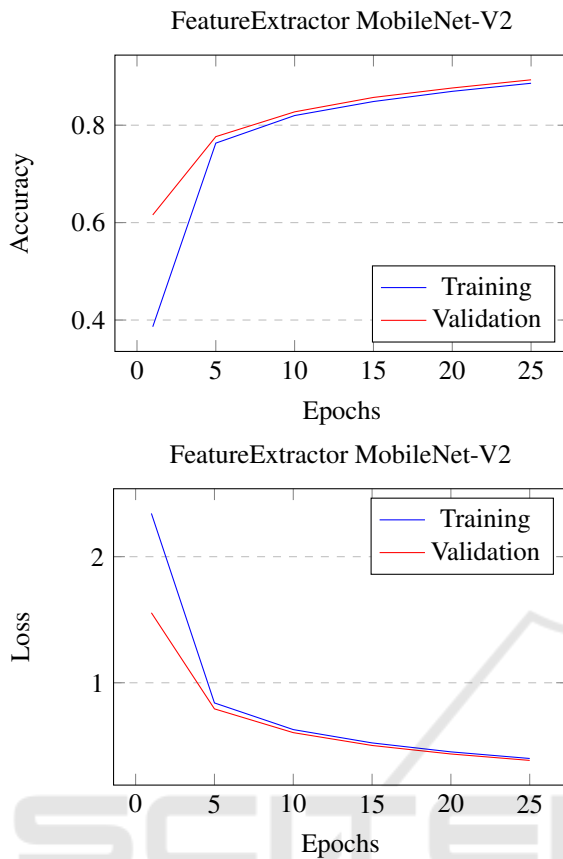


Figure 7: Accuracy and Loss for Training and Validation: FeatureExtractor MobileNet-V2.

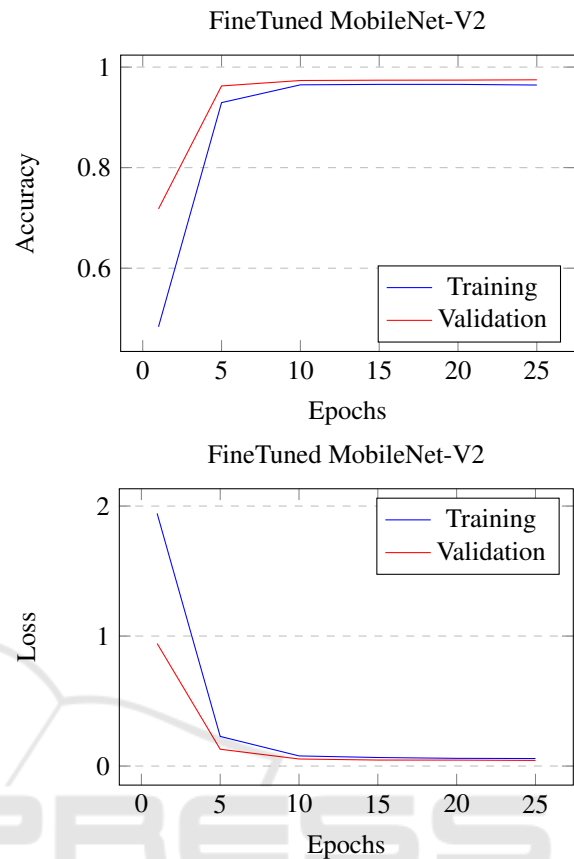


Figure 8: Accuracy and Loss for Training and Validation: FineTuned MobileNet-V2.

2. **Quantization Effect on Accuracy:** Table 4 illustrates that the testing accuracy of the MobileNet-V2 Finetuned model was 86%, while both quantized models with FP16 and FP32 achieved slightly reduced accuracy at 85% for both. However, the quantization effect on the chosen model did not significantly impact accuracy.
3. **Model Size:** This indicates the storage demand of the model on the edge device, influencing both device storage capacity and download bandwidth. As indicated in Table 4, the FP32 and FP16 models exhibited smaller sizes compared to the base model, which was 12.04 megabytes. FP16 had the smallest size of 3.02 megabytes. This reduction in size is indicative of the efficiency in model size achieved through quantization.

Taking into account all the evaluation criteria, it was observed that the test accuracy and inference time for both FP16 and FP32 models were nearly identical. However, the model size of the FP16 variant was approximately half the size of the FP32 model. Therefore, FP16 models were selected for deployment on

the Edge Raspberry Pi device, striking a balance between accuracy and model size efficiency.

6.3 Architecture Ability to Enrich Access Requests

To evaluate the efficacy of the proposed architecture in handling access requests to images based on their contents, we thoroughly evaluated it using different scenarios targeting various policies. These scenarios encompassed requests directed at both pre-classified images (where the image contents' classification information was previously saved in the ontology), and new, non-classified images (where the image content information has not been identified yet). The proposed Content-Driven access control architecture demonstrated its remarkable capability to precisely enrich requests with image content properties across the spectrum of testing scenarios, each aligned with different stored policies. As a result, our proposed architecture unequivocally proved its effectiveness in dynamically adapting to diverse access request scenarios and its agility in incorporating real-

Table 4: On Device (Raspberry Pi) Model Compression Effect on Fine-Tuned MobileNet V2.

Model	Model Size (MB)	Avg Inference Time (Sec)	Test Accuracy
Original Model	12.04	0.62	86%
FP32	5.98	0.16	85%
FP16	3.02	0.17	85%

time classification outcomes into access decisions, thereby showcasing its robust performance in enhancing smart farm system security.

6.4 Class Specific Threshold Adjustment Effect

We evaluated the effect of adjusting different levels of class-specific thresholds for the Crops major class when the architecture receives access requests to images containing crops and targeting policy-2 described in Table 2. The test dataset contains a total of 750 images of different crop types out of 4500 total test images. Based on the model classification results with different confidence scores and using a customizable threshold specific to Crops, we evaluated the number of images that were classified as a type of crop and are considered suitable for access authorization as follows:

- **True Positives (TP):** Number of images that were classified as at least one of the crop classes with confidence score \geq crop class specific threshold and are indeed crop images.
- **False Positives (FP):** Number of images that were classified as one of the crop classes with confidence score \geq crop class specific threshold and are not crop images.

Table 5 illustrates how different thresholds influence the access authorization to images based on their content classification confidence scores. As for the 100% threshold, fewer images are deemed suitable for access authorization due to stricter criteria, leading to a decrease in access to 14.7% of crop images and eliminating all access to misclassified images. On the other hand, using a smaller threshold of 20% increases the number of crop images eligible for access to 84.5% but includes more misclassified images to 31.9%. Notably, Thresholds ranging from 40% to 80% showed a balanced decrease in both the true positive and false positive access percentage. This underscores the critical role of threshold selection in determining the trade-off between granting access to relevant images and mitigating the risk of granting access to irrelevant or misclassified images.

Given this observation, the class-specific threshold emerges as a crucial element in mitigating in-

correct image classification outcomes, particularly concerning the security sensitivity levels associated with the content. Consequently, integrating the class-specific threshold directly into the policies would enhance efficiency further. This approach will allow smart farm administrators to directly formulate policies that describe the desired confidence levels for image classification, aligning with the specific security requirements and content sensitivities of the farm images.

Table 5: Access Authorization Percentage based on Class Specific Thresholds.

Threshold	True Positives	False Positives
%	%	%
20	84.5	31.9
40	78	23.8
60	73.6	17.5
80	71.1	11.2
100	14.7	0

6.5 Average Request Evaluation Time

The higher Request Evaluation time for pre-classified images, as indicated in Table 6, can be attributed to the time required for Inference from Ontology. This inference process involves extracting information from a structured representation of image content, and its duration can be influenced by the size and complexity of the ontology. The average time for Inference from Ontology ranges from 1.713 to 1.797, reflecting the intricate nature of ontology-based reasoning. The ontology's size and the depth of semantic relationships within it can contribute to a slightly longer processing time. On the other hand, the Request Evaluation time for non-classified images is notably lower, ranging from 0.56 to 0.67. In this case, the system relies on the image classifier, which tends to operate more swiftly. The classifier's efficiency in inferring image content contributes to the shorter Request Evaluation time compared to the ontology-based approach. This nuanced difference in processing times underscores the trade-off between ontology-based reasoning and the efficiency of image classifiers, highlighting the system's adaptability to dif-

Table 6: Average Total Content-Driven Access Request Evaluation Time of Pre-Classified & non-Classified Images.

	Classified Image		Non-Classified Image	
	Inference(Ontology)	Request Evaluation	Inference(Classifier)	Request Evaluation
Low	1.713	2.113	0.17	0.56
Avg	1.752	2.192	0.18	0.63
High	1.797	2.267	0.20	0.67

ferent scenarios and requirements. Despite the variations, the overall system maintains practical time manageability, ensuring effective access control in diverse smart farm image scenarios.

6.6 Limitation

One limitation of our proposed approach to plant image classification lies in the lack of images containing occluded or underdeveloped plants in our dataset. It is important to train the model on diverse conditions of plants which is common in smart farm images. For instance, occluded plant images might include instances where leaves or other objects partially obscure the plant, challenging the model to correctly classify the image content. Similarly, images of underdeveloped plants could feature younger crops with fewer leaves or less visible growth patterns, necessitating the model's ability to differentiate between varying stages of plant development. The absence of this type of plant image could adversely affect the classification accuracy of real smart farm images. Consequently, misclassifications may occur, potentially increasing undesired access request decisions.

7 RELATED WORK

Over the years, researchers in access control have proactively harnessed the capabilities of machine learning to derive more efficient solutions. Numerous research endeavors have advocated for ML-based approaches to formulate access control policies that exhibit greater robustness compared to conventional methods (Bui and Stoller, 2020), (Cotrine et al., 2018), (Abu Jabal et al., 2020), (Karimi et al., 2021), (Zhou et al., 2019). Additionally, rather than employing ML to address the entire problem, it can be leveraged to aid in resolving specific aspects within the access control domain. Therefore, Researchers have utilized ML advancements to automate cumbersome tasks in access control, including the extraction of attributes from information or text presented in plain natural language (Alohaly et al., 2018), (Alohaly et al., 2019b), (Abdi et al., 2022), (Alohaly et al., 2019a),

(Sandhu, 2021), the mapping between roles and permissions (Zhou et al., 2019), the extraction of security rules from access logs (Cotrine et al., 2018), and (Karimi et al., 2021), and even the extraction of access control policies from user stories (Sandhu, 2021).

Furthermore, researchers advocated for the integration of machine learning into the decision-making process of access control, as demonstrated in recent studies such as (Liu et al., 2021) and (Nobi et al., 2022). These studies underscore the benefits of leveraging an ML model for heightened accuracy in access control decision-making. In this paradigm, access decisions pivot on a trained ML model rather than a pre-defined access control policy. Typically, these models render access control decisions—granting or denying access—by analyzing user and resource metadata along with associated attributes. These metadata and attributes encapsulate user and resource features that the ML model learns for subsequent access decisions.

Nevertheless, as far as our knowledge extends, there has been a notable absence of the application of machine learning techniques tailored to facilitate informed decision-making in access control scenarios.

8 CONCLUSION

In this paper, we introduced a content-driven access control architecture that addresses the critical issue of regulating access to the extensive collection of smart farm images. Unlike traditional approaches, this novel architecture prioritizes regulating access based on the visual content of the images. The key components of this architecture include an edge machine-learning classification model and a semantic attribute-based access control mechanism. The edge classification model operates efficiently by providing a quick and lightweight solution for classifying images near their source in the smart farm. Its speed and accuracy make it a valuable tool for enhancing the overall performance of the access control system. Moreover, the semantic attribute-based access control mechanism adds a layer of sophistication by enriching access control requests with detailed image content information. This enhancement reduces the likelihood

of undesired access to smart farm image data, ensuring that only authorized personnel can retrieve specific image content.

The effectiveness of this architecture is demonstrated by its ability to enrich access requests with necessary image information, facilitating a more informed decision-making process based on the visual content of the images. Additionally, the edge classification model's lightweight, fast, and accurate nature contributes to the overall efficiency of the system. Furthermore, the average total time taken to evaluate access requests for both pre-classified and new images is remarkably small, with only slight additional time required for pre-classified images due to the time taken for ontology inference. This efficiency ensures a seamless and swift access control process, contributing to the overall success and practicality of the proposed architecture in a smart farm environment.

REFERENCES

- Abdi, A. I., Eassa, F. E., Jambi, K., Almarhabi, K., Khe-makhem, M., Basuhail, A., and Yamin, M. (2022). Hierarchical blockchain-based multi-chaincode access control for securing iot systems. *Electronics*, 11(5):711.
- Abu Jabal, A., Bertino, E., Lobo, J., Law, M., Russo, A., Calo, S., and Verma, D. (2020). Polisma-a framework for learning attribute-based access control policies. In *Computer Security—ESORICS 2020: 25th European Symposium on Research in Computer Security, ESORICS 2020, Guildford, UK, September 14–18, 2020, Proceedings, Part I 25*, pages 523–544. Springer.
- Alohaly, M., Takabi, H., and Blanco, E. (2018). A deep learning approach for extracting attributes of abac policies. In *Proceedings of the 23rd ACM on Symposium on Access Control Models and Technologies*, pages 137–148.
- Alohaly, M., Takabi, H., and Blanco, E. (2019a). Automated extraction of attributes from natural language attribute-based access control (abac) policies. *Cyber-security*, 2(1):2.
- Alohaly, M., Takabi, H., and Blanco, E. (2019b). Towards an automated extraction of abac constraints from natural language policies. In *ICT Systems Security and Privacy Protection: 34th IFIP TC 11 International Conference, SEC 2019, Lisbon, Portugal, June 25-27, 2019, Proceedings 34*, pages 105–119. Springer.
- Bera, S. and Shrivastava, V. K. (2020). Analysis of various optimizers on deep convolutional neural network model in the application of hyperspectral remote sensing image classification. *International Journal of Remote Sensing*, 41(7):2664–2683.
- Bui, T. and Stoller, S. D. (2020). Learning attribute-based and relationship-based access control policies with unknown values. In *International Conference on Information Systems Security*, pages 23–44. Springer.
- Cheng, J., Wu, J., Leng, C., Wang, Y., and Hu, Q. (2017). Quantized cnn: A unified approach to accelerate and compress convolutional networks. *IEEE transactions on neural networks and learning systems*, 29(10):4730–4743.
- Cottrini, C., Weghorn, T., and Basin, D. (2018). Mining abac rules from sparse logs. In *2018 IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 31–46. IEEE.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. (2009). Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee.
- Howard, A. G. (2017). Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*.
- Karimi, L., Aldairi, M., Joshi, J., and Abdelhakim, M. (2021). An automatic attribute-based access control policy extraction from access logs. *IEEE Transactions on Dependable and Secure Computing*, 19(4):2304–2317.
- Liu, A., Du, X., and Wang, N. (2021). Efficient access control permission decision engine based on machine learning. *Security and Communication Networks*, 2021(1):3970485.
- Nobi, M. N., Krishnan, R., Huang, Y., Shakarami, M., and Sandhu, R. (2022). Toward deep learning based access control. In *Proceedings of the Twelfth ACM Conference on Data and Application Security and Privacy*, pages 143–154.
- Sandhu, R. (2021). Access control policy generation from user stories using machine learning. In *Data and Applications Security and Privacy XXXV: 35th Annual IFIP WG 11.3 Conference, DBSec 2021, Calgary, Canada, July 19–20, 2021, Proceedings*, volume 12840, page 171. Springer Nature.
- Yassin, G. I. and Ramaswamy, L. (2022). Effective & efficient access control in smart farms: Opportunities, challenges & potential approaches. *ICISSP*, pages 445–452.
- Yassin, G. I. and Ramaswamy, L. (2023). Harvesting security: A semantically enriched access control architecture for smart farms. In *2023 5th IEEE International Conference on Trust, Privacy and Security in Intelligent Systems and Applications (TPS-ISA)*, pages 335–343. IEEE.
- Zhou, L., Su, C., Li, Z., Liu, Z., and Hancke, G. P. (2019). Automatic fine-grained access control in scada by machine learning. *Future Generation Computer Systems*, 93:548–559.