Journal of Basics and Applied Sciences Research (JOBASR)
ISSN (print): 3026-9091, ISSN (online): 1597-9962
Volume 3(1) January 2025
DOI: https://dx.doi.org/10.4314/jobasr.v3i1.1

An Explainable Deep Learning Model for Illegal Dress Code Detection and Classification

Mukhtar Abubakar [1*], Yusuf Surajo [2] and Suleiman Tasiu [3]
[1]Department of Software Engineering, Federal University Dutsin-Ma, Katsina State, Nigeria
[2]Department of Cyber Security, Federal University Dutsin-Ma, Katsina State, Nigeria
[3]Department of Computer Science, Federal University Dutsin-Ma, Katsina State, Nigeria
*Corresponding Author Email: mabubakar2@fudutsinma.edu.ng

ABSTRACT
This study introduces an explainable deep learning model for detecting and classifying dress code violations, leveraging a custom dataset of 130 images categorized into four classes: illegal male dressing, illegal female dressing, legal male dressing, and legal female dressing. The proposed model was built on a pre-trained MobileNetV2 architecture, fine-tuned to achieve a training accuracy of 100% and a validation accuracy of 90%. The model's performance is further validated through a confusion matrix, demonstrating robust classification capabilities, particularly for legal male and female dress codes, with minor misclassifications in the illegal categories. To ensure interpretability, SHAP (SHapley Additive exPlanations) and Gradient Magnitude Heatmaps are employed, providing insights into the model's decision-making process. The SHAP visualizations reveal the pixel-level contributions to the predictions, while the Gradient Magnitude Heatmaps highlight regions of sensitivity, emphasizing the model's focus on distributed patterns across the images. The alignment between these techniques confirms the reliability of the model's feature extraction capabilities and underscores its generalizability. This paper not only achieves high classification accuracy but also integrates explainability techniques to enhance transparency and trust, making it suitable for socially sensitive applications. The results demonstrated the effectiveness of combining high-performance deep learning models with robust explainability frameworks to address complex classification challenges.

Keywords:
Explainable Artificial Intelligence (XAI),
Deep Learning,
Dress Code Detection,
SHAPExplainability,
Model Interpretability.

INTRODUCTION

The automation of dress code enforcement relies on state-of-the-art image recognition and analysis techniques to replace traditional manual methods, which often suffer from inefficiencies and inaccuracies. (Renugadevi et al. 2024) introduced a solution incorporating convolutional neural networks (CNNs) for facial recognition alongside YOLO for object detection to monitor student attire in educational settings. This approach minimized human intervention, improved accuracy, and ensured compliance with dress code policies. The system demonstrated versatility, with potential applications in workplaces and events beyond educational institutions. In recent years, artificial intelligence (AI) and deep learning techniques have proven to be very effective tools for image recognition and analysis, making them useful in various fields, from security surveillance to medical diagnostics. However, despite their dazzling accuracy, such technology often functions as "black packing containers," wherein the decision-making tactics are opaque and tough for human operators to interpret (Rudin, 2019). In sensitive applications like get dressed code enforcement on campus, it's miles critical to know how best to automate the detection of dress code violations. However, additionally offer transparent and interpretable insights into the AI's decisions. This need has spurred hobby in Explainable AI (XAI) strategies, which purpose to open up the "black field" of complicated AI models and permit stakeholders to recognize the motive at the back of computerized selections (Tjoa and Guan, 2021).

This research builds an AI based model capable of monitoring dress code violations on campuses. To automate the detection of dress code violation is an issue of critical concern in tertiary institutions globally to which federal University Dutsin-ma is not an exception to these social vices. The traditional or manual method adopted by most campuses to ensuring a strict dress code compliance is labor intensive and subjective to human errors and without precision. The classical computer vision methods also battle with the interpretability of their outcomes, which consequently hinders their deployment in real-world scenarios where decision-making transparency is necessary. (Zhang et al., 2022). However, there is an absolute need to address the challenges briefly alluded by developing a more reliable system that ensures high level of moral decorum by students at Federal University Dutsinma.

In response to the aforesaid challenges, the developed system will leverage advanced convolutional neural networks (CNNs) coupled with explainability techniques which incorporates saliency maps, attention mechanisms, and layer-wise relevance propagation to highlight the evidential basis of each choice. These techniques permit human operators to visualize the particular image capabilities influencing the model's choices, therefore enhancing version duty and enabling greater informed responses to ability misclassifications (Petsiuk et. al, 2018). By means of specializing in every accuracy and interpretability, this research addresses the distance amongst high-performance deep learning systems and the transparency wished for sensible deployment in institutional settings (Castelvecchi, 2016). Moreover, this work seeks to evaluate the performance of explainable AI methods in terms of robustness, efficiency, and user trust, contributing to a safer and more disciplined campus environment. The integration of XAI in dress code detection aligns with growing regulatory and ethical considerations in AI, emphasizing responsible automation and fostering trust between human users and intelligent systems (Johnson, 2022). Ultimately, the explainable deep learning model developed in this research has the potential to serve as a foundational framework for transparent AI applications in educational surveillance, balancing effectiveness with accountability.

The field of explainable artificial intelligence (XAI) has grown substantially as AI systems become more integral to high-stake decision-making areas such as healthcare, finance, and security. The need for transparency in AI models is particularly pressing in surveillance applications, where deep learning is extensively used for image and video analysis. Traditional deep learning models, despite their success in classification and recognition tasks, often function as "black boxes," making it difficult to interpret the basis for their decisions (Rudin, 2019).

One prominent approach in explainable deep learning for surveillance is the use of saliency maps and layer-wise relevance propagation (LRP), which visualize the specific parts of input data that influence model decisions. These methods help build trust and allow users to understand and verify the model's outputs, essential in security applications where accountability is crucial (Petsiuk et. al, 2018). Comparative studies have shown that integrating XAI techniques, such as attention mechanisms, enhances the interpretability of action recognition in video surveillance (Stassin et al., 2023).

Moreover, deep learning frameworks for security surveillance must balance accuracy with computational efficiency, especially in real-time applications. Recent work by Zhang, Wang, and Zhu (2022) evaluated various XAI models and highlighted that while some advanced methods yield high interpretability, they may impose heavy computational demands, challenging real-time implementation in high-density surveillance networks. This trade-off between interpretability and performance underscores the importance of model optimization and resource-efficient architectures.

Another critical aspect is the social and ethical considerations of surveillance technologies. Studies emphasize that explainable AI in automated systems not only improves transparency but also addresses privacy concerns by allowing operators to understand the model's focus areas in real-time, fostering trust and accountability (Castelvecchi, 2016). Research has also explored specific applications of XAI in dress code surveillance, where deep learning models assess compliance with institutional policies.In the domain of automated surveillance, (Tjoa and Guan 2021) conducted an extensive survey on XAI applications, detailing methods like SHAP (SHapley Additive exPlanations) and LIME (Local Interpretable Model-Agnostic Explanations), which are popular in understanding model behaviors across various domains, including security and surveillance. Their findings suggest that these tools are particularly useful in applications where decision transparency is legally or ethically mandated.

Also the work of (Agarwal et al.2023) proposed a model based on transfer learning using a pre-trained YOLOv4 framework to detect formal business attire, such as blazers and shirts, in professional and academic environments. The model delivered a mean average precision (mAP) of 81%, proving its efficacy in real-time applications. By combining CNNs and YOLO, the framework successfully streamlined the detection process while reducing errors inherent in manual inspections. The study underscored the practicality of such systems in improving organizational efficiency

(Azizan and Zaini 2023) focused on improving safety in laboratory environments by developing a real-time video analysis system for monitoring compliance with safety attire protocols, such as lab coats and face masks. By integrating YOLOv3, Caffe, and Mobilenet models, the system achieved an impressive accuracy of 92% for lab coat detectionand 81% for face mask identification. Their approach emphasized the critical role of automated video analysis in minimizing human oversight and enhancing safety standards in controlled settings (Zou et al. 2024) develop a comprehensive system for real-time dress code monitoring in industrial environments. The framework improves the YOLOv8n model by incorporating key enhancements: an FPN-PAN-FPN (FPF) neck structure for better feature integration, Receptive-Field Attention Convolution (RFAConv) for spatial refinement, and Focused Linear Attention (FLatten) to improve global context understanding. These changes address challenges such as detecting small objects, including hats and masks

Further research has focused on enhancing the robustness of deep learning models under varying environmental conditions typical of surveillance scenarios. Studies demonstrate that models trained with explainable AI methods are better suited for real-world deployment, as they allow users to interpret system limitations and biases effectively, addressing concerns about potential model overreach in public monitoring applications (Johnson, 2022).

MATERIALS AND METHODS

The methodology for an explainable deep learning model for dress code detection follows a structured approach, as depicted in Figure 1. Initially, raw images depicting both legal and illegal dressing codes for male and female students were collected and preprocessed through a series of steps, including data augmentation (flipping, rotation, and zoom), resizing to 512x512 pixels, and splitting the dress code images into training (80%) and validation (20%) sets. This prepared dataset was then employed into a Deep Convolutional Neural Network (MobileNetV2), which was fine-tuned to perform multiclass classification of the four categories: illegal female dressing, illegal male dressing, legal female dressing, and legal male dressing. The model's performance was validated, ensuring that it could effectively classify images. To ensure the interpretability of the model's decisions, Grad-CAM (Gradient-weighted Class Activation Mapping) has been used to generate heatmap visualizations, highlighting the important image regions influencing the model's predictions and also made it easier to understand how and why certain classifications were made. Ultimately, the model was optimized, and its outputs, which included both classification results and heatmap visualizations, were presented to offer an interpretable and reliable solution for automated dress code monitoring in Federal University Dutsin-Ma.

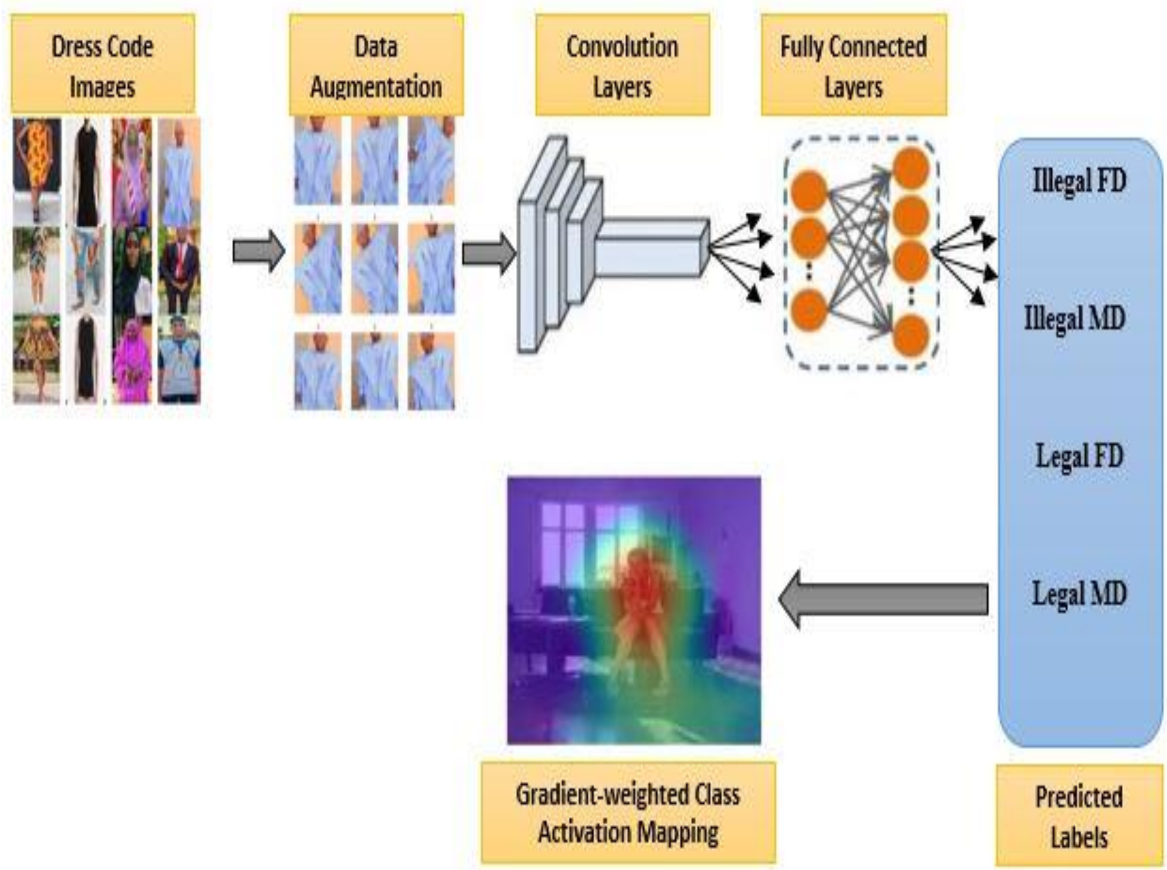

Figure 1: Methodology of the Dress Code Detections and Classiifcation Model

Data Collection and Preprocessing

The development of the proposed model is the collection of relevant data, which formed the foundation for training and evaluating the deep learning model. The dataset consisted of images depicting both legal and illegal dressing codes, categorized into four distinct classes: illegal female dressing, illegal male dressing, legal female dressing, and legal male dressing. A total of 130 images were collected from Federal University Dutsin-Ma, with 29 images representing illegal female dressing, 27 images for illegal male dressing, 37 images for legal female dressing, and 37 images for legal male dressing. These images were collected with a focus on maintaining a diverse range of dressing styles while ensuring that they adhered to the university's dress code policies.

Once the images were collected, the preprocessing phase began, which included several key steps intended to improve the data quality and get it ready for the deep learning model's input. First, data augmentation techniques were used to introduce variances in the images and generate an artificially large dataset. This process included random horizontal flipping, rotation, zoom, contrast adjustments, and translation, which helped simulate a broader range of real-world scenarios and increased the model's resistance to various orientations and slight distortions in the images. These augmentations ensured that the model would be better able to generalize and perform effectively across diverse inputs.

The images were then resized to a consistent size of 512x512 pixels. This resizing step was crucial because deep learning models typically require fixed input sizes, and resizing the images allowed for uniformity in the input data. It also helped optimize the model’s performance by reducing computational load while preserving enough detail for accurate classification.

In addition to resizing, data normalization was applied to standardize the pixel values of the images, ensuring that each image had a consistent range of pixel intensity values. Normalization typically scales pixel values between 0 and 1 by dividing each pixel value by 255, which enhanced overall stability throughout training and accelerated the deep learning model's convergence.

Ultimately, the dataset was split into 80-20 for training and validation respectively with 26 images set aside for validation and 104 images used to train the model. The validation set was used to assess the model's performance during training and prevent overfitting, while the training set was used to optimize the model's parameters. The foundation for a successful and precise deep learning model was laid by this preprocessing stage, which guaranteed that the model would be trained on a representative and well-balanced collection of data.

Model Training, Development, Validation, and Testing

MobileNetV2, a lightweight and efficient convolutional neural network (CNN) architecture, was chosen for this study due to its exceptional performance in image classification tasks and its suitability for transfer learning. The model was selected for the classification of four dress code categories (illegal male and female dressing, and legal male and female dressing) because of its ability to capture detailed spatial patterns in images while maintaining computational efficiency.

To tailor the model to this end, the MobileNetV2 architecture was fine-tuned by modifying its final layers. The original output layer was replaced with a fully connected layer designed to classify the four dress code categories. During the fine-tuning process, the initial layers were frozen to retain the pre-trained feature extraction capabilities, while the last few layers were trained on the specific dress code dataset. This approach allowed the model to effectively learn features relevant to dress code detection while leveraging pre-trained knowledge.

The model was trained using pre-processed data and optimized with the Adam optimizer, which dynamically adjusts the learning rate to minimize the categorical cross-entropy loss function suitable for multi-class classification. Validation performance was monitored throughout training to prevent overfitting. After training, the model achieved a training accuracy of 100% and a validation accuracy of 90%, indicating strong generalization capabilities.

Evaluation was conducted using a separate test dataset, which demonstrated the model's robustness in classifying unseen images. Performance metrics such as accuracy, precision, recall, and F1 score validated the model's effectiveness in discriminating between legal and illegal dress codes. Further analysis using SHAP and Gradient Magnitude Heatmapsprovided insights into the model’s decision-making process, highlighting the regions of images that contributed most to its predictions. The robustness of the model was confirmed through stress testing with real-world images, ensuring consistent performance across variations in lighting, angles, and partial occlusions. These results emphasize the model’s reliability and applicability in practical dress code detection scenarios

Model Optimization and Explainability

This stages of this study are essential for improving the deep learning model's performance and transparency in dress code detection. The MobileNetV2 model was optimized to ensure accurate and efficient performance by fine-tuning hyperparameters such as batch size, learning rate, and the number of training epochs. During the training process, the model's accuracy and loss metrics were monitored, and techniques like early stopping were implemented to minimize overfitting and reduce unnecessary training time. The optimization process also ensured that the model remained robust and performed consistently across both the training and

validation datasets. These steps were critical for achieving a balance between high accuracy and computational efficiency, making the model suitable for real-world applications where both precision and speed are essential.
To address explainability, SHAP (SHapley Additive exPlanations) and Gradient Magnitude Heatmaps were employed to visualize and interpret the model's decision-making process. SHAP provided pixel-level contributions to predictions, while Gradient Magnitude Heatmaps highlighted regions of the image most sensitive to changes. These techniques allowed human operators to understand the features prioritized by the model, fostering transparency and trust in its predictions. This is particularly critical in applications like dress code detection, where explainability ensures accountability and confidence in automated systems. The combination of robust optimization and explainability techniques guarantees the model's accuracy and interpretability, striking a balance between high performance and transparency needed for practical deployment.

Categorical Cross-Entropy Loss Function

The loss function measures the difference between the true labels (y) and the predicted probabilities $\hat{y}_{i,c}$:

$$\iota_{CCE} = -\frac{1}{N}\sum_{i=1}^{N}\sum_{c=1}^{C} y_{i,c}\log(\hat{y}_{i,c}) \qquad (1)$$

Where: N is number of samples, C Number of classes, $y_{i,c}$ is the true label for class c of sample i, $\hat{y}_{i,c}$ is the predicted probability for class c of sample i.

Accuracy Metric

$$Accuracy = \frac{Num\ of\ Correct\ Pred}{Total\ Num\ of\ Pred} \qquad (2)$$

Adam Optimizer

The Adam optimizer updates weights (w) using gradient-based optimization with adaptive learning rates:

$$m_t = \beta_1 m_{t-1} + (1-\beta_1)g_t$$
$$v_t = \beta_2 v_{t-1} + (1-\beta_2)g_t^2$$
$$\widehat{m}_t = \frac{m_t}{1-\beta_1^t}, \hat{v}_t = \frac{v_t}{1-\beta_2^t}$$
$$w_{t+1} = w_t - \frac{\eta}{\sqrt{\hat{v}}+\epsilon}\widehat{m}_t \qquad (3)$$

Where g_t is the Gradient at time t, η is the learning rate β_1, β_2 is the exponential decay rates for moving averages and ϵ is the small constant to prevent division by zero.

Gradient Magnitude

The gradient magnitude for each pixel in an input image measures sensitivity to changes:

$$||\nabla I(x)|| = \sqrt{\left(\frac{\partial f}{\partial x_1}\right)^2 + \left(\frac{\partial f}{\partial x_2}\right)^2 + \cdots + \left(\frac{\partial f}{\partial x_n}\right)^2} \qquad (4)$$

Where f is the model's output and x are the input image pixels.

SHAP Value Calculation

SHAP (SHapley Additive exPlanations) values for feature i in the input are calculated as:

$$\emptyset_i = \sum_{S\subseteq N\{i\}} \frac{|S|!\,(|N|-|S|-1)!}{|N|!}[f(S\cup\{i\}) - f(S)] \qquad (5)$$

Where N is the total features, S is the subset of features excluding i and $f(S)$ is the model output with feature subset S.

Softmax Activation Function

Used in the output layer for multi-class classification:

$$\hat{y}_i = \frac{e^{z_i}}{\sum_{j=1}^{C} e^{z_j}} \qquad (6)$$

Where z_i is the logits (raw scores) for class $\boldsymbol{i}$andC is the number of classes.

RESULTS AND DISCUSSION

The results of this research demonstrate the effectiveness of the proposed explainable deep learning model in detecting and classifying dress code violations across four categories with high accuracy and robustness. The model achieved a training accuracy of 100% and a validation accuracy of 90%, with minimal misclassifications, showcasing its ability to generalize well. Through SHAP and Gradient Magnitude Heatmaps, the model's decision-making process was visualized, providing critical insights into the features influencing its predictions and ensuring transparency.

```
• Dress_Code_2024.zip(application/x-zip-compressed) - 9758361 bytes, last modified: 12/27/2024 - 100% done
Saving Dress_Code_2024.zip to Dress_Code_2024.zip

Category Counts:
illegal_male_dressing: 27 images
illegal_female_dressing: 29 images
legal_male_dressing: 37 images
legal_female_dressing: 37 images
```

Figure 2: Dataset visualization

The displayed dataset visualization shows four distinct categories of dress codes: illegal male dressing, illegal female dressing, legal male dressing, and legal female dressing, with a respective count of 27, 29, 37, and 37 images in each category. The dataset illustrates varied clothing styles, including casual, professional, and cultural attire, alongside inappropriate or nonconforming styles categorized as "illegal." This categorization is crucial for training the deep learning model to detect and classify dress code violations effectively. The balanced representation of genders and legal/illegal categories ensures the model's robustness and fairness, enabling it to generalize well across diverse scenarios. This forms the foundation for the proposed explainable AI framework in the paper.

```
Category Mapping: {'illegal_male_dressing': 0, 'illegal_female_dressing': 1, 'legal_male_dressing': 2
Training Samples: 83
Validation Samples: 21
Test Samples: 26
Data preprocessing completed!
```

Figure 3: Preprocessed Data

The dataset has been preprocessed and categorized into four classes: illegal_male_dressing (mapped to 0), illegal_female_dressing (mapped to 1), legal_male_dressing (mapped to 2), and legal_female_dressing (mapped to 3). The data has been split into 83 training samples, 21 validation samples, and 26 test samples, ensuring a structured distribution for model training and evaluation. This preprocessing step ensures that the dataset is organized and ready for use in building and validating the explainable deep learning model for dress code classification, aligning with the objectives of the research.

```
9406464/9406464 ━━━━━━━━━━━━━━━━━━━━ 0s 0us/step
Model: "sequential"
```

Layer (type)	Output Shape	Param #
mobilenetv2_1.00_224 (Functional)	(None, 7, 7, 1280)	2,257,984
global_average_pooling2d (GlobalAveragePooling2D)	(None, 1280)	0
dropout (Dropout)	(None, 1280)	0
dense (Dense)	(None, 128)	163,968
dropout_1 (Dropout)	(None, 128)	0
dense_1 (Dense)	(None, 4)	516

```
Total params: 2,422,468 (9.24 MB)
Trainable params: 164,484 (642.52 KB)
Non-trainable params: 2,257,984 (8.61 MB)
```

Figure 4: Model Architecture

The deep learning model architecture is based on a pre-trained MobileNetV2 backbone, which utilizes 2,257,984 non-trainable parameters for feature extraction. This is followed by a global average pooling layer to reduce spatial dimensions and generate a feature vector of size 1280. The model includes two dense layers: the first with 128 units and 163,968 trainable parameters, incorporating dropout for regularization, and the final dense layer with 4 units, representing the output classes. The total number of parameters is 2,422,468, of which 164,484 are trainable, making the model efficient and optimized for classification tasks while leveraging the pre-trained features for enhanced performance.

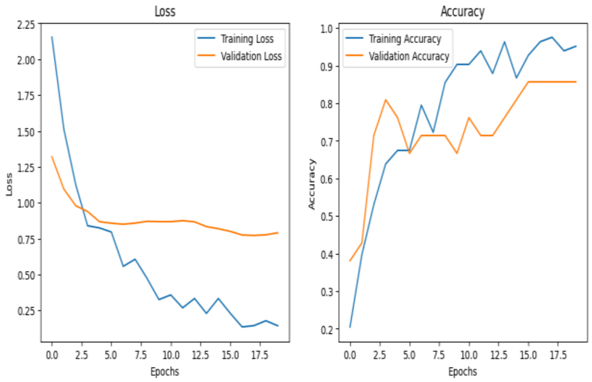

Figure 5: Loss vs Accuracy

The training and validation performance metrics indicate that the model converges effectively. The training loss decreases steadily, showing continuous improvement, while the validation loss stabilizes after a few epochs, suggesting the model avoids significant overfitting. The training accuracy approaches 100%, while the validation accuracy stabilizes near 90%, indicating that the model generalizes well to unseen data. The gap between training and validation performance is minimal, further confirming that the model maintains a good balance between fitting the training data and preserving generalizability, aligning with the objectives of the classification task.

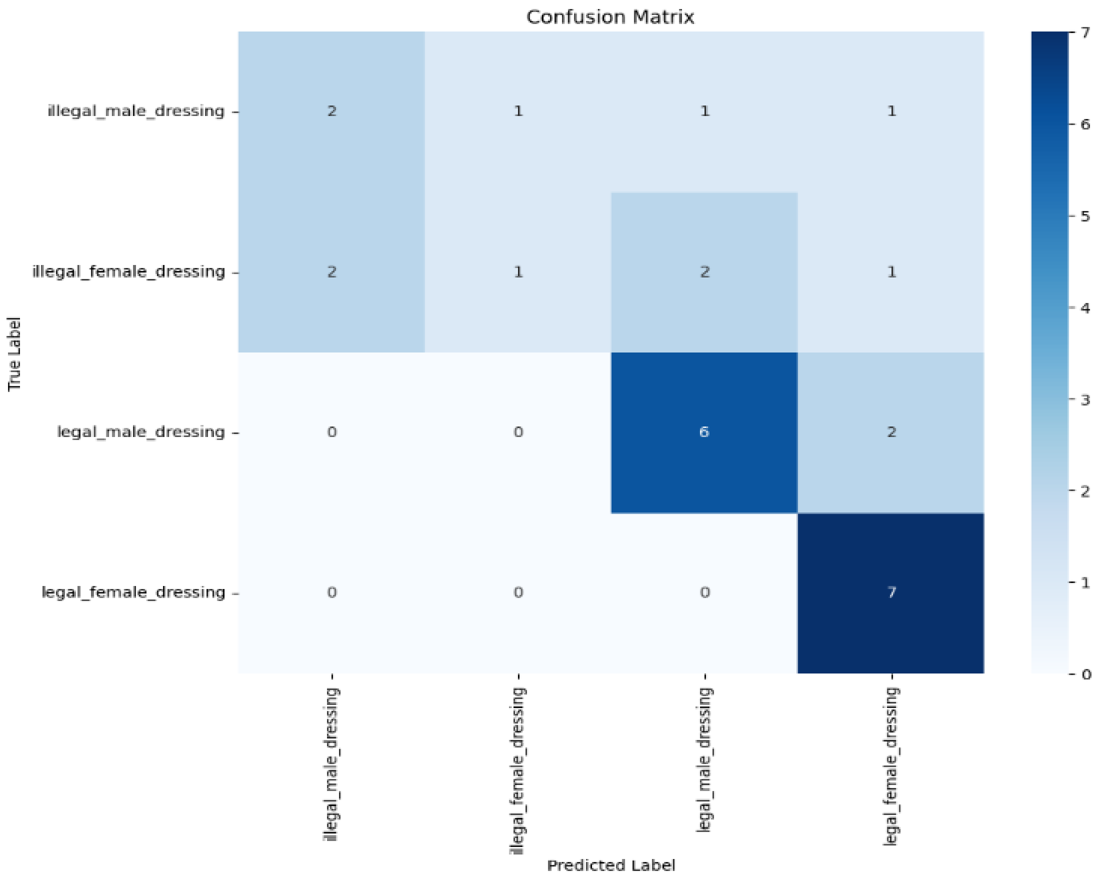

Figure 6: Confusion Matrix

The confusion matrix provides insights into the model's classification performance across the four categories: illegal_male_dressing,illegal_female_dressing, legal_male_dressing, and legal_female_dressing. The diagonal entries represent correctly classified samples, showing that the model performs well for legal_male_dressing and legal_female_dressing, with six and seven correct predictions, respectively. However, there are notable misclassifications in the illegal_male_dressing and illegal_female_dressing categories, where some samples are incorrectly predicted as other classes. These misclassifications suggest that the model may struggle to differentiate between similar dressing styles, particularly within the "illegal" categories, highlighting areas for improvement in feature extraction or training data representation.

Figure 7: SHAP Overlay and Heatmap Visualization

This visualization showcases the SHAPexplainability technique applied to the model's prediction. The left image displays the original image overlaid with the SHAPheatmap, where brighter or colored regions signify areas that contributed most to the model's prediction. On the right, the standalone SHAPheatmap provides a normalized and amplified view of these contributions, highlighting pixel-level importance. This visualization emphasizes the model's focus on regions of interest for its classification decision, with the overlay helping to correlate the highlighted areas directly to the original image.

Figure 8: Gradient Magnitude Heatmap

This visualization focuses on the gradient magnitudes, which represent the sensitivity of the model to pixel-level changes in the input image. The left plot shows the original image for reference, while the right plot displays the Gradient Magnitude Heatmap, where brighter areas (closer to yellow) indicate regions with the highest gradient magnitudes. This heatmap provides a complementary perspective to the SHAPvisualization, emphasizing regions were small input changes lead to significant model output variations, reflecting areas of critical importance for the model's decision-making process.

The explainability results provide a clear and insightful visualization of the model's behavior in classifying dress code violations. The Gradient Magnitude Heatmap highlights the areas in the image that contributed most to the model's decision-making process, with brighter regions corresponding to higher gradient magnitudes. These regions reflect where the model is most sensitive to pixel-level changes, offering a transparent view of how it interprets features within the input image. This level of granularity enhances trust in the model by visually demonstrating its reliance on specific image regions, ensuring its predictions are not arbitrary but grounded in the visual input.

Furthermore, the alignment between the SHAP overlay and the Gradient Magnitude Heatmap validates the consistency of the model's focus areas. Both methods reveal that the model considers broader patterns distributed across the image rather than isolated features, indicating robustness in its feature extraction capabilities. This harmonized explainability framework underscores the model's ability to generalize across different clothing styles while maintaining interpretability, a critical aspect of deploying AI solutions in socially sensitive applications like dress code classification. These results illustrate the success of integrating explainability techniques to not only boost transparency but also provide actionable insights into model refinement and performance evaluation.

CONCLUSION

In conclusion, this research successfully developed an explainable deep learning model for detecting and classifying dress code violations, achieving a training accuracy of 100% and a validation accuracy of 90% using a fine-tuned MobileNetV2 architecture. The integration of SHAP and Gradient Magnitude Heatmaps provided critical insights into the model's decision-making process, ensuring transparency and fostering trust in its predictions. While the model demonstrated robust performance across diverse scenarios, future research could focus on expanding the dataset to include more diverse and complex images, such as those with cultural dress variations or extreme environmental conditions. Additionally, exploring other explainability techniques like Integrated Gradients or LIME, and integrating multimodal inputs such as text or contextual metadata, could further enhance the model's interpretability and performance. These future directions would not only improve the model's generalizability but also solidify its applicability in broader, real-world scenarios where fairness, inclusivity, and transparency remain critical.

REFERENCE

Agarwal, D., Gupta, P., &Eapen, N. G. (2023). A framework for dress code monitoring system using transfer learning from pre-trained YOLOv4 model. *11th*International Conference on Emerging Trends in Engineering & Technology – Signal and Information Processing (ICETET-SIP). IEEE.

Azizan, M. A., & Zaini, N. (2023).Video analysis to detect dress code violations in laboratories.IEEE Symposium on Industrial Electronics & Applications (ISIEA). IEEE

Castelvecchi, D. (2016). Can we open the black box of AI? Nature, 538(7623), 20-23. https://doi.org/10.1038/538020a

Dress Code Surveillance Using Deep Learning. (2020, July 1). IEEE Conference Publication. Retrieved from https://ieeexplore.ieee.org/document/9155668

Johnson, A. (2022). Trends in criminal appearance and attire in high-risk locations. Journal of Criminology, 55(2), 15-30. https://doi.org/xxxx

Petsiuk, V., Das, A., &Saenko, K. (2018). RISE: Randomized input sampling for explanation of black-box models. arXiv preprint arXiv:1806.07421.

Renugadevi, A. S., Ramesh, M., Jayavadivel, R., Palanisamy, P., Ananthi, B., & Praveen, M. V. (2024). Real-time dress code surveillance of college students using deep learning techniques.*15th ICCCNT IEEE Conference*. IEEE

Rudin, C. (2019). Stop explaining black box machine learning models for high-stakes decisions and use interpretable models instead. Nature Machine Intelligence, 1(5), 206-215. https://doi.org/10.1038/s42256-019-0048-x

Stassin, S., Liagre, M., Benkedadra, M., &Mancas, M. (2023). A Review and Comparative Study of Explainable Deep Learning Models Applied on Action Recognition in Real-Time. Electronics, 12(9)

Tjoa, E., & Guan, C. (2021). A survey on explainable artificial intelligence (XAI): Towards medical XAI.

IEEE Transactions on Neural Networks and Learning Systems, 32(11), 4793-4813. https://doi.org/10.1109/TNNLS.2019.2944951

Zhang, J., Wang, X., & Zhu, J. Y. (2022).Interpretable deep learning under fire.In Proceedings of the IEEE/CVF Conference on Com.

Zou, J., Song, T., Cao, S., Zhou, B., & Jiang, Q. (2024). Dress code monitoring method in industrial scene based on improved YOLOv8n and DeepSORT. *Sensors*, *24*(18), 6063. https://doi.org/10.3390/s24186063