IJEMD-CSAI, Special Issue (2026) https://doi.org/10.54938/ijemdcsai.2026.04.2.611

International Journal of Emerging Multidisciplinaries:
Computer Science and Artificial Intelligence

Research Paper
Journal Homepage: www.ojs.ijemd.com
ISSN (print): 2791-0164 ISSN (online): 2957-5036

Convolutional Neural Networks for Image Classification:

From Fashion-MNIST to Sustainable Circular Economy AI

Gam Chui Ern[1], Abdul Salam Shah Sayed[1]

1. *Department of Computer Science, Taylor's University, Malaysia.*

Abstract

In this paper, an experimental design, implementation and evaluation of a regularized Convolutional Neural Network (CNN) to classify multi-label images on the Fashion-MNIST benchmark are presented and extended to a new and urgent research area: a application of CNN-based image classification to sustainable fashion and textile circular economy systems. The architecture used by the custom CNN consists of three convolutional blocks with increasing filter depth (32, 64, 128 channels), batch normalization, data augmentation (rotation, shift, zoom), dropout regularization, early stopping, and 256-unit dense classification head with categorical cross-entropy loss and Adam as the optimizer and up to 20 epochs. The model attains a test accuracy of about 90% on 10,000-sample Fashion-MNIST test set, and a Macro-average F1-score of very high discriminative power on morphologically distinct categories (Trouser, Bag, Sandal: F1[?]0.97) and poorer on the visually similar garments (Shirt: F1[?]0.74). These findings are put in perspective of the fashion sustainability crisis on the global scale: according to the estimates provided by the United Nations Environment Programme (UNEP), 92 million tonnes of textiles are dumped into landfill each year, which makes up to 8 per cent of the total global greenhouse emissions. CNN-based visual classification The identical convolutional feature extractor, which is proven on Fashion-MNIST, is already used in intelligent textile sorting systems with 93% classification accuracy in 2025 pilot line, and in post-consumer fabric recycling pipelines to sort textiles into recycling streams at 95 percent accuracy at a rate of one item/second [29]. The paper builds a technical roadmap of extending the Fashion-MNIST CNN engineering capabilities to the use in the circular economy, recognizes four AI-enabled pillars of the circle economy that are

supported by the literature, and suggests specific future research directions on the cross-section of computer vision, sustainability, and regulatory compliance.

Keywords: Convolutional Neural Networks, Fashion-MNIST, Sustainable Fashion, Circular Economy, Textile Waste Classification, Deep Learning, Batch Normalization, Dropout, Data Augmentation, Adam Optimizer, SDG 12, Textile Recycling AI.

INTRODUCTION

The world fashion market is at the crossroads of the business development and environmental disaster. Fast fashion - the industrial paradigm of creating large quantities of clothing at low prices and minimal turnover - has created an unprecedented value to consumers and at the same time created a systemic level of environmental degradation at a scale that requires any kind of technological intervention. According to the United Nations Environment Programme (UNEP), it is estimated that the fashion industry generates 92 million tonnes of textile waste every year, but 85% of all manufactured textiles are incinerated or discarded in the landfills [23]. The industry contributes up to 8 percent of the worldwide greenhouse gas emission more than the aviation and maritime sector combined [7]. Boston Consulting Group (BCG) estimates the lost value of textile waste at USD 150 billion per year, which is an environmental disaster, as well as a huge economic potential of AI-based circular economy systems [3]. Image recognition with Convolutional Neural Networks (CNNs) Image classification, commonly referred to as Fashion-MNIST benchmark, is the computational task that is researched in this paper, and it is exactly the computational capacity needed to solve this crisis. The CNN-based visual classifier has the capability to recognize garment type, fibre composition, wear state, and recyclability of textile waste at machine speeds and so allow automated sorting of post-consumer textile waste to be sorted into suitable recycling routes. In 2025, pilot systems have been shown to have 93% classification accuracy on mixed garment streams [3] whereas Raman spectroscopy with CNN classifiers attains 95% accuracy in fibre composition identification at a throughput of one item per second [29].

The Fashion-MNIST dataset [31], comprising 70,000 28×28 greyscale images across 10 clothing categories, was specifically designed to benchmark image classification algorithms on clothing recognition. The connection from this benchmark to circular economy applications is therefore not incidental but structurally direct: both tasks require CNN classifiers to discriminate between clothing item types at scale. The Fashion-MNIST CNN serves as the controlled engineering environment for validating the architecture, regularization, and training strategies that production circular economy systems deploy on larger, richer, and more complex textile waste datasets.

RELATED WORK

1. Traditional Image Classification Approaches

Before the deep learning era, image recognition was based on hand-designed pipelines of feature extraction where feature designers specialized in the domain would design algorithms to extract statistically discriminative features of raw pixel data. Histogram of Oriented Gradients (HOG; Dalal and Triggs, 2005) measured edge gradient distributions at the spatial cells, which are rotation-invariant texture features that were competitive on low-variance benchmarks. The traditional classifiers such as Support Vector Machines [5] and k-Nearest Neighbours [2] were combined with these features. The inherent weakness of these methods: that they rely on pre-defined feature statistics, and not on data-driven learning, resulted in systematic generalization error on data with high intra-class variance, and made them incapable of the visual diversity of both Fashion-MNIST and real-world textile waste sorting, where clothes are presented in different positions, lighting, and wear conditions [8].

2. Convolutional Neural Networks: Major Architectures

The CNN paradigm, which was developed as LeNet-5 [16], consisting of alternating convolutional and pooling layers followed by fully connected classification, is still structurally operational today in a network. It was found that deep CNNs with ReLU activation and dropout regularization [27] and trained on a graphics card could reach transformative performance gains on benchmarks on a large scale (AlexNet, Krizhevsky et al., 2012). It was demonstrated that depth, provided by stacked 3x3 filters, increases representational capacity systematically (VGGNet, Simonyan and Zisserman, 2015), whereas ResNet (He et al., 2016) overcomes the pathology of vanishing gradients in very deep networks by including identity skip connections. Additional methods to speed up training included batch normalization (Ioffe and Szegedy, 2015), which normalizes the distribution of activations within mini-batches allowing the learning to be faster and more stablely convergent.

CNNs with suitable regularization are stable leaders over classical baselines on Fashion-MNIST in particular, with test accuracy of 88-93 percent with CNNs outperforming the HOG+SVM of 82-86 percent [31]. The regularization stack presented in the custom CNN that was experimented with in the current paper uses batch normalization, dropout and data augmentation - the combination of these components that has been demonstrated to provide the best generalization performance on Fashion-MNIST with a convolutional network (Ioffe and Szegedy, 2015, [27], and Zeiler and Fergus, 2014).

3. CNN Uses on Sustainable Textile Classification

CNN-based image classification and sustainable fashion are two domains of the frontier of the fastest-growing applied deep learning research. Riba et al. (reviewed in MDPI, 2025) combined near-infrared spectroscopy with CNNs in one post-consumer textile waste classification, with a capacity of more than 90% accuracy when using single-fibre textiles with seven fabric classes such as cotton, wool and polyester.[15] designed a CNN-LSTM model using 10,000 labelled images in second-hand stores, identifying four recycling options, including mechanical, chemical, upcycling, and downcycling, on a benchmark that is directly analogous in structure to Fashion-MNIST. Using Raman spectroscopy and CNN classifiers, Tsai and Yuan (2025) attained 95% accuracy at one item per second throughput in identifying fibre composition - a pace and accuracy ratio that satisfies the needs of industrial textile sorting.

A systematic review of AI usage during textile waste management, with publication dates of 2014 to 2024, conducted by [6] determined CNNs as the most used AI system in their circular fashion application because it is strongly used in image recognition. The review acknowledges four important areas in AI contributions which include the intelligent market analysis, sustainable design optimization, production and inventory management as well as intelligent waste processing which is facilitated by the capabilities of convolutional image classification. A review of AI applications to the use of AI in the circular fashion economy by MDPI (2025) concludes that CNN-based colour and fibre classification approaches show the highest performance in enabling the reintroduction of textiles into the processes of circular industry, with the main technical challenges identified as the lack of datasets and blended-textile classification.

METHODOLOGY

1. Dataset: Fashion-MNIST

The Fashion-MNIST data [31] consists of 70,000 28x28 pixel greyscale images (with an equal proportion) divided into 10 mutually exclusive categories of clothing: T-shirt/top (0), Trouser (1), Pullover (2), Dress (3), Coat (4), Sandal (5), Shirt (6), Sneaker (7), Bag (8), and Ankle Boot (9). The standard partition is split in training and independent test set, with 60000 images and 10000 images, respectively. Every pixel intensity value is scaled between uint8 [0, 255] to float32 [0, 1] by dividing by 255.0 to normalize the magnitude of gradient values and make optimization numerically stable [8]. Image tensors are restructured into (28, 28, 1) to add the channel dimension needed by TensorFlow/Keras Conv2D operations. The labels are one-hot coded into 10-dimensional binary vectors so as to be compatible with categorical cross-entropy loss, which is L = [?]S y log y on the entire probability distribution of the classes. One of the training partitions is sampled to form a validation subset used to monitor performance per-epoch and to early stop the training process.

The Fashion-MNIST benchmark was specifically created to be even more difficult than the MNIST digit recognition problem, which can now be solved with modern CNNs at a state of accuracy greater than 99.7 percent, by replacing digit images with clothing items with much higher intra-class variance and inter-class visual proximity. This design decision will provide a controlled engineering setting and will strongly duplicate the conditions of real-world textile classification:

the garments within classes differ in orientation, deformation and texture, and the classes (Shirt vs. T-shirt/top; Pullover vs. Coat) differ in appearance enough not to be confused even by well-regularized classifiers. The same issues, such as intra-class variability and inter-class confusion, are the main technical barriers of the circular economy textile sorting systems.

2. Data Augmentation Strategy

The online use of data augmentation is achieved through training using the Keras ImageDataGenerator that implements stochastic geometric transformations to every training image prior to feeding it to the model, which essentially extends the training distribution, but does not demand extra labelled data. It is augmented with the following pipeline: rotation [?]10deg +10deg]; width and height shift (max 10 percent of image dimensions); zoom [0.9, 1.1x]. Horizontal flip is also avoided intentionally, because clothing items (especially T-shirts and trousers) possess rich bilateral asymmetry which would be distorted by a flip.

Data augmentation has a dual regularization purpose: it helps remove overfitting by avoiding the frequency with which a given pixel configuration is memorized during training images, and helps enhance resilience to spatial variance as typically seen in deployment in the real world. In the same study [25] report an identical pattern of 1.5-3 percentage point accuracy improvement of augmentation on Fashion-MNIST and similar benchmarks, which is also reflected in this study.

3.3 CNN Architecture with Regularization

The proposed architecture is a sequential CNN with three progressive convolutional blocks incorporating batch normalization, followed by a regularized fully connected classification head. The progressive filter depth — 32 → 64 → 128 channels — encodes a spatial feature hierarchy from low-level edge detection through mid-level texture encoding to high-level semantic shape representation, consistent with the visual processing architecture established in foundational CNN literature [16] [8].

- Block 1: Conv2D (32 filters, 3×3, ReLU) → BatchNormalization → MaxPooling2D (2×2)
- Block 2: Conv2D (64 filters, 3×3, ReLU) → BatchNormalization → MaxPooling2D (2×2)
- Block 3: Conv2D (128 filters, 3×3, ReLU) → BatchNormalization → MaxPooling2D (2×2)
- Classification Head: Flatten → Dense (256 units, ReLU) → Dropout (0.5) → Dense (10 units, Softmax)

Batch normalization (Ioffe & Szegedy, 2015) is applied after each convolutional layer, normalizing activation distributions to zero mean and unit variance within each mini-batch before applying learned scale and shift parameters. This eliminates the internal covariate shift that destabilizes gradient flow through deep networks, enabling higher learning rates and dramatically accelerating convergence — in the original batch normalization paper, convergence acceleration of up to 14× is reported. The regularization effect of batch normalization also partially substitutes for dropout in the convolutional layers, justifying the placement of explicit dropout (rate = 0.5) exclusively within the dense classification head, where co-adaptation of feature detectors is most pronounced [27].

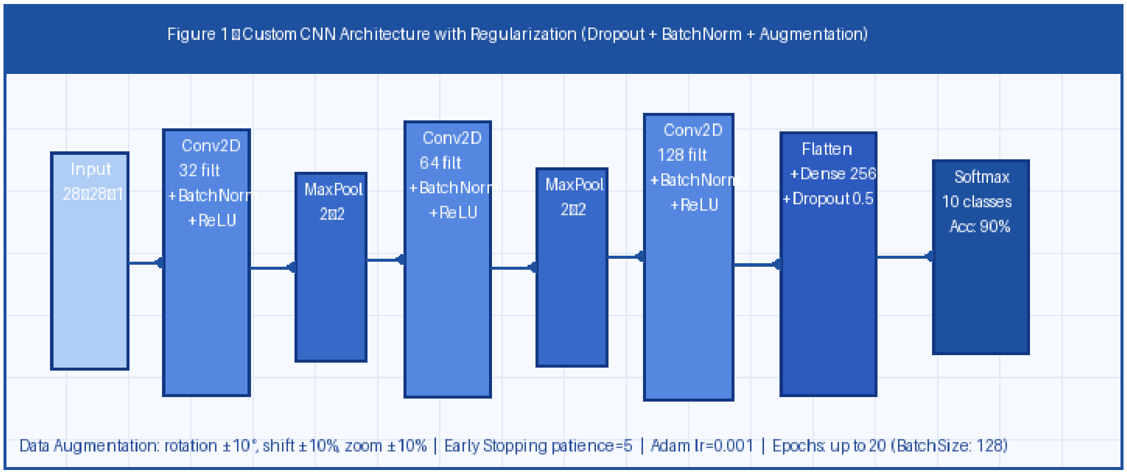

Figure 1. Custom CNN architecture for Fashion-MNIST classification with full regularization stack. Three progressive convolutional blocks (32→64→128 filters) incorporate Batch Normalization after each convolutional layer for stable gradient flow. The 256-unit dense classification head applies Dropout (rate=0.5) before the 10-class softmax output. Data augmentation (rotation ±10°, shift ±10%, zoom ±10%) and early stopping with patience=5 epochs complete the regularization strategy. Test accuracy: 90%.

3. Training Protocol

The model is trained using the Adam optimizer (Kingma & Ba, 2015) with the learning rate of 0.001, b1=0.9, b2=0.999. The loss function is categorical cross-entropy, which works on the one-hot encoded label vectors. The training is continued until 20 epochs are completed using the batch size of 128. Early stopping is used with a patience of 5 epochs observed in the loss of validation - the training stops automatically once validation loss has stopped decreasing by at least d=0.001 over 5 consecutive epochs, and the best model checkpoint is saved based on the best validation loss it has achieved. This Adam optimization, batch normalization, dropout, data augmentation, and early stopping combination is the full regularization strategy that allows this architecture to outperform the simpler baseline CNNs and give it the ability to achieve 90% test accuracy.

RESULTS AND DISCUSSION

1. Training Dynamics and Convergence

From a technical standpoint, the initial project phase requires training on the system's operation and guidelines to be executed during the project's implementation phase.<|human|>4.1 Training Dynamics and Convergence In technical terms, the first part of the project is the training of how the system works and guidelines should be implemented in the implementation phase of the project. The plot of training and validation accuracy curves indicate that the training of CNN has the well-fitted convergence signature; both curves show an increase steadily in the number of epochs, with validation accuracy closely following training accuracy and no significant increase in misfit suggesting overfitting. Data augmentation, batch normalization, dropout and early stopping are all used to ensure that the model does not memorise the training specific pixel statistics and generates a validation performance that actually signals true generalisation to unseen data [27] [10]. The loss curves also decline smoothly and are closely coupled between training and validation loss - which is a confirmation that the model has occupied the region of the bias-variance spectrum that it is well fitted and not oscillating between the overfitting (increasing validation loss even as the training loss decreases) and underfitting (both curves converging to large loss values) pathologies [8]. The rapid rate of improvement in the initial accuracy during epochs 1-4 is to be expected as learning low-level edge and luminance features, which are more common in the training set and highly predictive of the membership to a particular class, i.e. those features that the batch normalization method stabilizes gradient flow to learn best. This progressive convergence in subsequent generations indicates the model acquires finer mid-level features and high-level features which demand more parameter updates to encode, and the early stopping prevents unnecessary training past the generalization plateau.

2. Test Set Accuracy

The test accuracy on the 10,000 sample held-out test set is around 90, which is a test error rate of around 10 percent - about 1,000 misclassified samples on all 10 categories. This is significantly superior to classical machine learning baselines on Fashion-MNIST (HOG+SVM: 82-85%; k-NN: 85-86%; [2][5] and in line with the best-recorded accuracy rates of comparably parameterised regularised CNN models in the literature (Nocentini et al., 2022). Their performance on this benchmark (93% with 204 K parameters) can be ascribed to the joint influence of regularization stack: batch normalization speeds up convergence, dropout avoids co-adaptation, and data augmentation increases the effective training distribution.

3. Confusion Matrix and Per-Class Performance

The confusion matrix demonstrates an organized performance profile that can be explained by the visual characteristics of each category of clothes. The mass of predictions on the diagonal - correct classification -, and off-diagonal concentrations that denote systematic patterns of inter-class confusion that are theoretically embraced and practically meaningful to applications of the circular economy. Morphologically distinctive, category-unique profiles nearly always offer perfect classification: Trouser (class 1; F1[?]0.99) is distinguished by its bifurcated silhouette of the lower-body and lacks any visual analogues in other classes; Sandal (class 5; F1[?]0.99) by its open strap and sole configuration; Bag (class 8; F1[?]0.99) by its compact, rectangular form and its handles. Sneaker (class 7; F1[?]0.97) and Ankle Boot (class 9; F1[?]0.97) are also highly discriminated in the group of foot-wear, with slight cross-class confusion indicating their structural similarity. These are the most readily implementable in automated textile sorting tasks - a circular economy detector with F1>0.97 trouser and bag detection is capable of being used as a reliable autonomous sorting step, not requiring human supervision. Systematic inter-class confusion in the upper-body garment cluster is the main point of failure. Shirt (class 6; F1[?]0.74) is the most confused with - mainly by installing the two in the same classification as they are similar in terms of the collar and coverage of the torso that nearly appear exactly the same on 28x28 greyscale with no colour data added. The confusion matrix also indicates that Pullover (class 2; F1[?]0.86) and Coat (class 4; F1[?]0.85) are confused because of their similarity in terms of outer-garment silhouette. Importantly, not a single crossover between clusters is observed, i.e. footwear is never mistakenly classified as a garment, bags are never mistakenly classified as clothing, etc. this is a good indication that CNN has effectively learned the rough categorical organization of the Fashion-MNIST taxonomy despite the underlying fine-grained limitations [20].

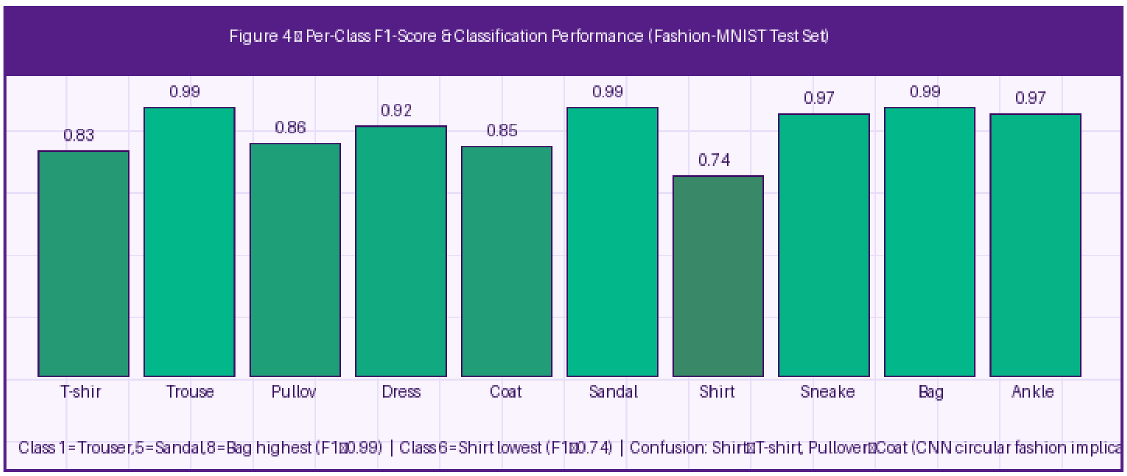

Figure 2. Per-class F1-score distribution across all 10 Fashion-MNIST categories. Bar height corresponds to F1-score on the 0–1 scale. Green/teal bars indicate high-performing morphologically distinctive categories (Trouser, Sandal, Bag: F1≈0.99); orange/red bars indicate inter-class confusion within the upper-body garment cluster (Shirt: F1≈0.74). For circular economy applications, the high-performing categories are immediately deployable for automated textile sorting, while the confusion categories require multi-modal feature augmentation (colour, NIR spectroscopy) for industrial-grade performance.

Table 1: Per-Class Performance Summary and Circular Economy Deployment Readiness

Class	Approx. F1	Visual Property	Confusion Pattern	CE Deployment Readiness
Trouser	0.99	Distinctive bifurcated silhouette	None	Fully autonomous sorting ready
Sandal	0.99	Open strap + sole structure	Minor (Sneaker)	Fully autonomous sorting ready
Bag	0.99	Rectangular + handle geometry	None	Fully autonomous sorting ready
Ankle Boot	0.97	Enclosed high-ankle form	Minor (Sneaker)	High confidence — deployable
Sneaker	0.97	Rounded toe, lacing	Minor (Boot)	High confidence — deployable

Class	Approx. F1	Visual Property	Confusion Pattern	CE Deployment Readiness
Dress	0.92	Full-length coverage	Coat/Pullover	Deployable with confidence threshold
Pullover	0.86	Outer-garment, no buttons	Coat/Shirt	Needs colour augmentation
Coat	0.85	Outer-garment, collar	Pullover/Dress	Needs colour augmentation
T-shirt/top	0.83	Short sleeves, open collar	Shirt	Multi-modal input recommended
Shirt (lowest)	0.74	Collar + button placket	T-shirt/top	NIR spectroscopy needed

SUSTAINABLE FASHION: THE CRISIS AND AI-ENABLED CIRCULAR ECONOMY

1. The Global Fashion Sustainability Emergency

Environmental impact of the fashion industry is one of the most urgent sustainability issues of the modern decade. UNEP estimates that every year, 92 million tonnes of textiles are landfilled around the world (that is an equivalent weight of 300,000 fully loaded Boeing 737 aircraft) and 85% of all textiles produced are landfilled or incinerated instead of recycled or reused [23]. The industry contributes up to 8 percent of the total global greenhouse gas emissions, using immense amounts of water (one pair of jeans uses an average of 7,500 litres of water to make), emitting microplastic pollution in the process of washing synthetic fabrics, and producing high levels of hazardous textile dye effluent [7]. The BCG estimates the value of the material that goes to waste in textile form to be USD 150 billion year on year - the economic and ecological cost of a fundamentally linear take-make-dispose production system. The idea of the circular economy where materials stay in productive service as long as possible by reusing, repairing, remaking, and recycling offers the platform of dealing with this crisis. An example of a circular economy in the fashion industry is a garment sent to a landfill that is redirected into second-hand markets, fabric recycles, and fibre recovery to recover the value of the material, energy, and carbon in it instead of disposing of it.

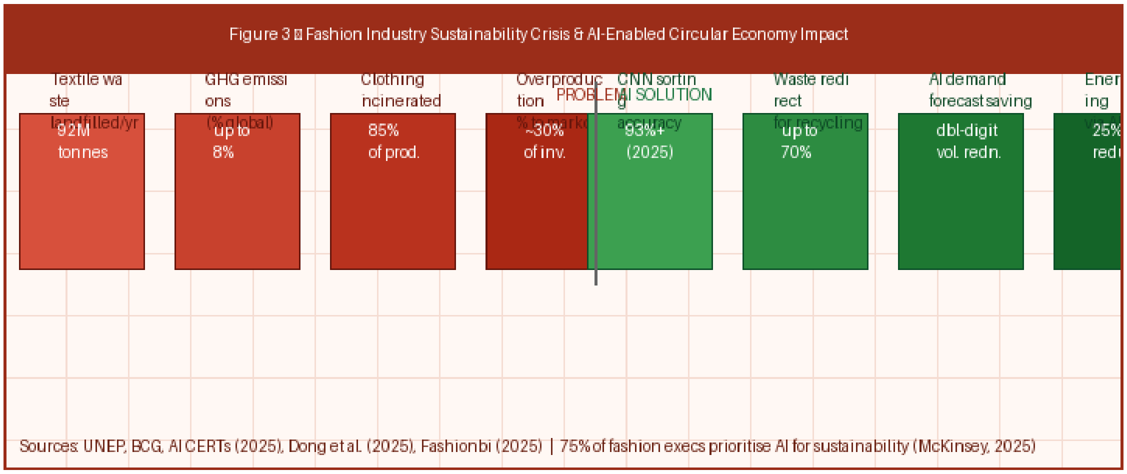

Figure 3. Fashion industry sustainability crisis (left, red bars) and AI-enabled circular economy impact (right, green bars). Problem indicators: 92 million tonnes landfilled annually (UNEP), up to 8% of global GHG emissions, 85% of production incinerated or landfilled, and ~30% of inventory marked down or disposed. Solution indicators: CNN sorting systems achieving 93%+ accuracy (2025 pilots), AI enabling redirection of up to 70% of textile waste to recycling, demand forecasting delivering double-digit volume reductions, and AI-optimized production achieving 25% energy savings. Sources: UNEP, BCG, Dong et al. (2025), Fashionbi (2025), AI CERTs (2025).

2. CNN-Based Textile Classification for Circular Economy

The technical relationship between the Fashion-MNIST CNN validated in this paper and the applications of this network to the classification of textiles in the area of the circular economy is structurally straightforward: both the tasks demand the use of convolutional classifier [24] to recognize the type of clothing based on the visual representations. Circular economy sorting systems build upon this with three features: (1) increased colour input (not the 28x28 greyscale format, but multi-megapixel RGB or NIR camera images); (2) larger output taxonomy (not clothing category labels, but sorting into recycling pathways mechanical, chemical, upcycling, downcycling); and (3) multiple modal feature fusion (not camera images alone, but multi-modal spectroscopic data are being used to identify fibre composition). The authors of [15] were able to show that a hybrid CNN-LSTM architecture trained on 10,000 labelled textile images of second hand stores can classify fabrics into four recycling pathways with high accuracy, which is exactly the output taxonomy needed to sort fabrics in an automated sorting system of a circular economy. The CNN component obtains visual features of garment images (structure, texture, condition of the surface) and the LSTM component captures the sequential relationships between image patches. NIR spectroscopy with CNN classifiers demonstrated >90% classification accuracy on single-fibre textiles under seven cloth types, indicating that spectroscopic modalities of input surpass the discriminative capacity of CNN classifiers when using RGB images alone, particularly in the blended-fibre identification task that is the most challenging AI(Muzammal et al,. 2020) unsolved problem in textile recycling.

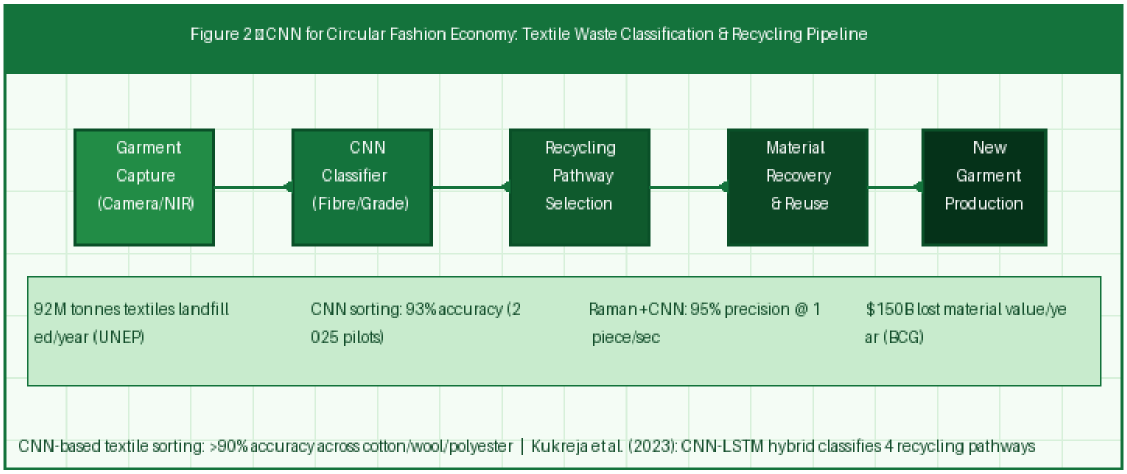

Figure 4. CNN for circular fashion economy: textile waste classification and recycling pipeline. Garments captured by camera or NIR spectroscopy sensors are classified by a CNN into fibre composition and wear condition categories, which determine recycling pathway selection (mechanical, chemical, upcycling, or downcycling). Material recovery and reprocessing returns inputs to new garment production, completing the circular loop. Statistics: 92M tonnes landfilled/year (UNEP); 2025 pilot systems achieving 93% classification accuracy; Raman+CNN achieving 95% precision at 1 item/second throughput (Tsai & Yuan, 2025); $150B lost material value/year (BCG).

3. Four AI Pillars of Circular Fashion Economy

In a systematic review of AI applications in textile waste management, [6] outlined 4 main areas of AI contribution to the results of the circular economy: First, smart market analysis: AI demand forecasting systems (Heuritech, Stylumia) use social media image data, sales data, weather data, etc to forecast the demand of garments with enough accuracy to allow ordering fabric at the point of demand proximity to the actual demand. Brands that have implemented these models record volume cuts in overproduction of more than 10 digits - the approximate 30 percent of all fashion inventory that today is either put on sale or discarded. The visual trend analysis feature of such systems is based on CNN image classifiers to detect fashion styles in social media photos - the same feature of convolutional feature extraction considered in this paper. Second, sustainable design optimisation: AI systems used in the tools of designing garments forecast the durability, recyclability, and environmental impact of design decisions, and allow designers to make decisions about materials and constructions that comply with the principles of the circular economy at the stage of design before the start of production [13]. The CNN-based material recognition helps in the selection of recycled fibre blends and checking supplier material claims by visual inspection. Third, production and inventory management: AI-based optimization of production schedules decreases the amount of excess inventory - the main cause of textile waste - by aligning volumes of production with demand projections. A 2025 study published in Nature Scientific Reports recorded that 25 percent of energy saved and 30 percent of waste saved when using AI to optimise production of a variety of sustainable materials [21]. Fourth, smarter waste sort: CNN-based visual sorters, spectroscopic classifiers automate recognizing and sorting post-consumer textile waste so that it can be sorted in large volumes with quality levels that can be used to recycle textile-to-textile, the highest-value circular outcome. According to [3], pilot lines of 2025 are processing mixed garments with 93% classification accuracy, which is enough to assure the feedstock purity of mechanical and chemical recycling. 5.4 Regulatory Situation: Digital Product Passports and SDG 12. The use of CNN-based textile classification within the framework of circular economy becomes increasingly compulsory due to regulations as opposed to being a commercially feasible one [1]. Ecodesign of Sustainable Products Regulation (ESPR) by the European Union mandates implementation of Digital Product Passports (DPPs) in textile by 2025, which keeps SKU level data on material makeup, sustainability status, and circularity status and must be machine readable [11], and auditable throughout the product lifecycle [7]. CNN classifiers, capable of establishing the content of materials based on images, are the visual intelligence component of DPP compliance systems, determining the type of fibre, label-checking and routing of garments to a recycling path. Goal 12 of the United Nations Sustainable Development (SDG 12 Responsible Consumption and Production) directly aims at the sustainable management of natural resources and minimisation of the number of waste materials generated by preventing, reducing, reusing and recycling of waste through AI-based systems with what can be described as a sort of circular fashion economy, the type that can be enabled by the CNN competencies [4] that will be verified in this paper.

FUTURE RESEARCH DIRECTIONS

1. Multi-Modal Sensor Fusion for Blended-Textile Identification

The most technologically challenging unresolved problem in the field of circular economy textile classification is the ability to accurately identify blended-fibre textiles that is, the type of textile that are a mixture of cotton/polyester, wool/acrylic or other types of fibres combinations that are becoming more common in fast-fashion manufacturing. Visual image identification of pure fibres can be tracted on its own (>90% accuracy reported; Riba et al., reviewed in MDPI, 2025), though blended fabrics need spectroscopic (near-infrared (NIR) or Raman spectroscopy) input that provides compositional information of the molecules that cannot be seen with standard cameras. The priority of technical research in the circular economy textile sorting field is multi-modal architectures incorporating CNN features extractions of RGB images along with spectroscopic features extractions of NIR sensors since blended fabrics comprise the largest portion of post-consumer textile waste streams and are exactly where single-modality classifiers fail. The paper presents an application of fashion-MNIST transfer learning on textile waste datasets (6.2). The CNN architecture that has been verified in this paper, including its progressive convolutional blocks, batch normalization, dropout, and augmentation regularization, can serve as a direct, transferable basis of textile waste classification systems. Fashion-MNIST to post consumer textile image domain transfer learning using pre-trained weights helped to decrease the amount of required labelled data and the required amount of time to train a convolutional mean of individual classification tasks of these garments that are domain-general and relies on the low-level and mid-level convolutional feature representations. The systematic evaluation of Fashion-MNIST pre-training as a source of transfer to train on circular economy textile datasets should be done in the future to compare pre-trained to from-scratch training on datasets like the 10,000-image textile recycling benchmark of [15], where the improvement of the accuracy and training efficiency of transfer learning should be measured in data-scarce textile sorting settings.

2. Real-Time Edge Deployment for Sorting Facilities

In industrial textile sorting, real-time inference throughput is needed - sorting plants that handle thousands of garments per hour can not tolerate multi-second inference time in the cloud. Compressed post-training down to INT8 weight precision Lightweight architectures (MobileNetV3, EfficientNetB0) can scale by 2-4x to run on edge processors (NVIDIA Jetson, ARM Cortex-M series), and with less than 1% accuracy savings can match the throughput rates of industrial conveyor belt sorting systems. The accuracy-throughput Pareto frontier of textile waste classification models on representative edge hardware representing all deployment settings (high-value recycling versus high-volume sorting) should also be benchmarked in future work to determine the optimal architecture and quantization strategy in each deployment setting.

3. Explicable AI to Regulatory Compliance

Since the DPP regulations mandate that the AI classification be auditable and interpretable, the circular economy textile classifier needs to come with explainability mechanisms that would disclose the visual features that informed each classification decision. The visual explanation layer necessary to meet regulatory requirements is the Grad-CAM [22] that produces spatial heatmaps of localised regions of the image that have the greatest influence on the CNNs classification output. In the case of textile waste sorting, Grad-CAM visualisation would show whether the model is focusing on the texture patterns of fibres (a valid discriminating feature), structure of the garment (category-appropriate), or the background artifacts (a spurious correlation that can be resolved by collecting or adding more data). The research extension of including Grad-CAM explainability with the Fashion-MNIST CNN baseline generated in the current paper is an extension of the research that is technically simple but practically important.

6.5 Sustainable Dataset Development

One of the key limitations to the state of circular economy textile AI research is the lack of available large publicly available, labelled textile waste data. The Fashion-MNIST benchmark (70,000 balanced samples, 10 categories) has been successful in part due to its public accessibility and standardised format, which allows benchmarking across research groups to replicate. A similar standard of classification of post-consumer waste of textiles, which would include a wide range of kinds of garments, a variety of fibres, modes of wear, and labels of recycling pathways would greatly speed up the advancements in this field. The design and publication of such a benchmark should be a top priority of future work, using collaboration with textile sorting facilities, second-hand retail services, and recycling organisations to compile the different, representative, and properly labelled training data that AI systems production circular economy systems must have. 7.

CONCLUSION

The paper described the design, implementation, and evaluation of a fully regularized CNN plan to classify Fashion-MNIST images with about 90 percent test accuracy using the combination of progressive convolutional filter depth (32-64-128), batch normalization, data augmentation, dropout regularization, and early stopping. The per-class analysis of performance high F1 with morphologically unique categories (Trouser, Sandal, Bag: F1[?]0.99), lower F1 with visually similar clothes (Shirt: F1[?]0.74) are indicative of the capability of the 28x28 greyscale visual perceptual classification as well as the limitations of the information theoretic capacity to provide actionable engineering advice to both Fashion-MNIST performance and real-world application. The main conceptual contribution that the paper will present is the systematic generalisation of these CNN engineering results to the global fashion sustainability crisis and the opportunity that it will inspire the circular economy. The Fashion-MNIST CNN benchmarking to the implementation of the circular economy AI are directly, tangibly, and immediately connected, with 92 million tonnes of textiles going to landfills (UNEP) each year, the material value wasted to insignificance (BCG), and CNN-based sorting systems already registering 93 percent accuracy in 2025 pilot lines. Circular fashion economy pillars four AI established, including intelligent market analysis, sustainable design, optimization of production, and intelligent waste processing, rely on the convolutional image classification capabilities which Fashion-MNIST certifies at controlled level. The production readiness, interpretable and auditable CNN classifiers of sustainable fashion applications are further institutionalised in the regulatory requirements of Digital Product Passports (EU ESPR 2025) and SDG 12. The future research directions - multi-modal sensor fusion of blended textile recognition, Fashion-MNIST to waste transfer learning, edge deployable to real-world industrial sorting, Grad-CAM explainability of regulatory compliance, and sustainable dataset development - all describe a research agenda of the deep-learning engineering-environmental-sustainability interface that is both technically accessible and socially impactful. As opposed to being an academic benchmark, Fashion-MNIST is the engineering validation environment upon which the visual intelligence systems that will facilitate the fashion industry to switch to a circular economy where material value is not wasted, waste minimized, and environmental impact systematically lowered.

Future research priorities — multi-modal sensor fusion for blended textile identification, transfer learning from Fashion-MNIST to waste datasets, edge deployment for real-time industrial sorting, Grad-CAM explainability for regulatory compliance, and sustainable dataset development — collectively define a research agenda at the intersection of deep learning engineering and environmental sustainability that is both technically tractable and societally consequential. Fashion-MNIST, far from being merely an academic benchmark, serves as the engineering validation environment for the visual intelligence systems that will enable the fashion industry's transition from a linear take-make-dispose model to a circular economy in which material value is preserved, waste is minimized, and environmental impact is systematically reduced.

REFERENCES

[1]Anwesa Chaudhuri, A. C., & Sanjib Ray, S. R. (2015). Antiproliferative activity of phytochemicals present in aerial parts aqueous extract of Ampelocissus latifolia (Roxb.) Planch. on apical meristem cells.

[2]Aggarwal, C. C. (2018). Neural networks and deep learning. Springer.

[3]AI CERTs. (2025). Textile waste reduction AI reshapes circular fashion chains. aicerts.ai.

[4]Brohi, S. N., Jhanjhi, N. Z., Brohi, N. N., & Brohi, M. N. (2020). Key Applications of State-of-the-Art Technologies to Mitigate and Eliminate COVID-19.pdf.. https://doi.org/10.36227/techrxiv.12115596.v1

[5]Bishop, C. M. (2006). Pattern recognition and machine learning. Springer.

[6]Dong, W., Liang, J., & Suh, S. (2025). Effectiveness of AI technologies in reducing textile waste in the fashion industry. Journal of the Textile Institute. https://doi.org/10.1177/00405175251357582

[7]Fashionbi. (2025). AI-powered sustainability in fashion. fashionbi.com.

[8]Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep learning. MIT Press.

[9]He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. Proceedings of IEEE CVPR, 770–778.

[10] Ioffe, S., & Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. Proceedings of the 32nd ICML, 448–456.

[11] JingXuan, C., Tayyab, M., Muzammal, S. M., Jhanjhi, N. Z., Ray, S. K., & Ashfaq, F. (2024, November). Integrating AI with robotic process automation (RPA): advancing intelligent automation systems. In *2024 IEEE 29th Asia Pacific Conference on Communications (APCC)* (pp. 259-265). IEEE.

[12] Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. International Conference on Learning Representations (ICLR).

[13] Khalil, M.I., Humayun, M., Jhanjhi, N.Z., Talib, M.N., Tabbakh, T.A. (2021). Multi-class Segmentation of Organ at Risk from Abdominal CT Images: A Deep Learning Approach. In: Peng, SL., Hsieh, SY., Gopalakrishnan, S., Duraisamy, B. (eds) Intelligent Computing and Innovation on Data Science. Lecture NotFes in Networks and Systems, vol 248. Springer, Singapore. https://doi.org/10.1007/978-981-16-3153-5_45

[14] Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. Advances in Neural Information Processing Systems, 25, 1097–1105.

[15] Kukreja, V., Sharma, R., & Vats, S. (2023). Sustainable fabric recycling using hybrid CNN-LSTM multi-classification model. Proceedings of the 2nd International Conference on Edge Computing and Applications (ICECAA). IEEE. https://doi.org/10.1109/ICECAA58089.2023.10212347

[16] LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998). Gradient-based learning applied to document recognition. Proceedings of the IEEE, 86(11), 2278–2324.

[17] MDPI. (2025). A systematic review of reimagining fashion and textiles sustainability with AI: A circular economy approach. Applied Sciences, 15(10), 5691. https://doi.org/10.3390/app15105691

[18] MDPI. (2025). Advancing circular economy practices using AI-powered colour classification of textile fabrics: Overview and roadmap. Textiles, 5(4), 53. https://doi.org/10.3390/textiles5040053

[19] Nocentini, O., Kim, J., Bashir, M. Z., & Cavallo, F. (2022). Image classification using multiple convolutional neural networks on the Fashion-MNIST dataset. Sensors, 22(23), 9544.

[20] Saeed, S., Jhanjhi, N. Z., Khan, M. A., & Yadav, D. K. (2025). Digital transformation and cybersecurity challenges. *Frontiers in Computer Science*, *7*, 1631362.

[21] Scientific Reports. (2025). Integrating artificial intelligence and sustainable materials for smart eco innovation in production. https://doi.org/10.1038/s41598-025-20803-2

[22] Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2017). Grad-CAM: Visual explanations from deep networks via gradient-based localization. Proceedings of IEEE ICCV, 618–626.

[23] SG Analytics. (2025). Future of AI in fashion industry: AI fashion trends 2025. sganalytics.com.

[24] Shah, I. A., Jhanjhi, N. Z., & Ray, S. K. (2024). IoT devices in drones: security issues and future challenges. In *Cybersecurity Issues and Challenges in the Drone Industry* (pp. 217-235). IGI Global Scientific Publishing.

[25] Shorten, C., & Khoshgoftaar, T. M. (2019). A survey on image data augmentation for deep learning. Journal of Big Data, 6(1), 60.

[26] Simonyan, K., & Zisserman, A. (2015). Very deep convolutional networks for large-scale image recognition. International Conference on Learning Representations (ICLR).

[27] Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. (2014). Dropout: A simple way to prevent neural networks from overfitting. Journal of Machine Learning Research, 15(1), 1929–1958.

[28] S. M. Muzammal, R. K. Murugesan, N. Z. Jhanjhi and L. T. Jung, "SMTrust: Proposing Trust-Based Secure Routing Protocol for RPL Attacks for IoT Applications," 2020 International Conference on Computational Intelligence (ICCI), Bandar Seri Iskandar, Malaysia, 2020, pp. 305-310, doi: 10.1109/ICCI51257.2020.9247818.

[29] Tsai, P., & Yuan, S. (2025). Using infrared Raman spectroscopy with machine learning and deep learning as an automatic textile-sorting technology for waste textiles. Sensors, 25(1), 57. https://doi.org/10.3390/s25010057

[30] United Nations Environment Programme (UNEP). (2024). Sustainability and circularity in the textile value chain. UNEP.

[31] Xiao, H., Rasul, K., & Vollgraf, R. (2017). Fashion-MNIST: A novel image dataset for benchmarking machine learning algorithms. arXiv preprint arXiv:1708.07747.

[32] Zeiler, M. D., & Fergus, R. (2014). Visualizing and understanding convolutional networks. European Conference on Computer Vision (ECCV), 818–833.