

Comparative Deep Learning Models for Cow and Buffalo Classification Using CNN and Attention Mechanisms

Ankit Kumar Gupta*, Gulshan Kumar, Ayush Ranjan, Anchal Kaundal

School of Computer Science and Engineering

Lovely Professional University

Phagwara-144411, Punjab, India

Email: ankit963134@gmail.com, gushan76542@gmail.com, ayushranjan9821@gmail.com, anchalkaundal786@gmail.com

Abstract—Telling a cow apart from a buffalo sounds like something any farmer can do at a glance, but building a machine that does it reliably from a photograph is a harder problem than it first appears. Both species share similar body sizes, live in the same farm environments, and are photographed under the same inconsistent lighting and angles. In regions like South Asia where both animals are kept on the same farms, automated species identification would directly support livestock registration, veterinary management, and supply chain verification — yet no published study has tackled this binary classification problem head-on with a controlled multi-model comparison. This paper does exactly that. We evaluated five deep learning architectures — a Baseline CNN trained from scratch, MobileNetV2, ResNet50, a CNN with Morphological Feature Fusion, and an Attention-Based CNN using the Convolutional Block Attention Module (CBAM) — on a cleaned and balanced dataset of 1,713 outdoor cow and buffalo images curated from two Kaggle sources covering eleven Indian breeds. Every model was trained under the same conditions: same dataset split, same optimizer, same batch size, same image normalization. ResNet50 achieved the best overall performance with 99.22 % accuracy, 99.26% precision, 99.26% recall, 99.26% F1-score, and 99.96% AUC-ROC, outperforming both the MobileNetV2 baseline and the CBAM attention model. An ablation study confirmed that CBAM did not improve results over ResNet50 alone, indicating that the deep pretrained features extracted by ResNet50 are already sufficient for this classification task without additional attention complexity. Grad-CAM maps revealed that the model focuses on horn shape, the dewlap, and body silhouette — the exact anatomical features that separate the two species — rather than background elements like fencing or ground texture.

Index Terms—Cow and Buffalo Classification, Convolutional Neural Networks, Transfer Learning, CBAM Attention, ResNet50, Precision Livestock Farming

I. INTRODUCTION

Anyone who has spent time on a mixed farm in South Asia knows that cows and buffalo are kept together routinely — same pasture, same shed, sometimes even the same herd. Globally, buffalo numbers have crossed 200 million while cattle are well past a billion, and a large share of both populations occupy overlapping farming regions [1]. Getting the two species confused is not just an embarrassing mistake — it leads to wrong medication being administered, livestock records that cannot be trusted, mislabelled dairy products reaching

consumers, and compensation being paid out for the wrong animal [2]. Each of these failures happens today because no reliable automated tool exists to make the distinction quickly and accurately in real field conditions.

Visually, a trained eye can tell them apart without much effort. Buffalo are broader through the neck, carry distinctively curved horns that sweep outward and back, and have a darker, more wrinkled skin and a heavier dewlap. Bos indicus cattle, the dominant breed across the region, tend to carry a dorsal hump, lighter colouring, and a more upright horn profile where horns are present at all. The trouble is that neither horn shape nor body colour is a reliable signal in a photograph taken at a distance, at an angle, in inconsistent light, with the animal partially hidden behind a fence. Building a model that handles all of that while still picking up on the right structural differences is the core challenge this work addresses [3].

Over the past decade, convolutional neural networks have become the go-to solution for exactly this kind of recognition task [24]. Rather than training a network from zero on a small domain-specific collection — which rarely goes well — researchers have found that starting from weights learned on large general datasets and then adapting them to the target problem gives far better results for the effort involved [10]. Two architectures that appear repeatedly in this kind of work are MobileNetV2 [9], valued for how little compute it needs, and ResNet50 [5], valued for how well it generalises to new domains. Beyond transfer learning, attention modules have more recently emerged as a practical way to squeeze more accuracy out of a pretrained backbone by directing the network's focus toward the regions of an image that actually matter for the decision [8]. Whether that combination of ideas works for livestock species separation specifically is something the literature has not yet settled.

Searching through what has been published, the cow-vs-buffalo binary problem has not been studied as a task in its own right. Work exists on distinguishing cattle breeds from one another [3], [4], on picking buffalo out of broader livestock populations [1], and on monitoring and tracking farm animals generally [6], but a side-by-side comparison of multiple deep learning approaches on this specific two-class problem has not

appeared. The work by Pan et al. [1] — the closest thing to a predecessor for this study — dealt with telling buffalo breeds apart from each other, not with the prior question of whether the animal is a buffalo at all. Naik et al. [2] brought ResNet50 and EfficientNetB3 into a mixed-species detection scenario but left out morphological feature extraction and any form of attention mechanism. Together these two studies define the edge of what is currently known and point toward where more work is needed.

This paper closes that gap. We make three specific contributions:

Contribution 1: We assembled, cleaned, and balanced a dataset of 1,713 outdoor cow and buffalo images curated from two publicly available Kaggle sources, covering five cow breeds and six buffalo breeds, providing a documented benchmark collection for this binary classification task.

Contribution 2: We ran a head-to-head comparison of five architectures — Baseline CNN, MobileNetV2, ResNet50, a CNN with Morphological Feature Fusion, and a CBAM Attention CNN — under identical training conditions, so that performance differences reflect the models and not the experimental setup.

Contribution 3: We identify ResNet50 as the strongest model for this task, achieving 99.22% accuracy, backed by an ablation study confirming that CBAM attention did not yield additional gains on this dataset, and Grad-CAM visualizations showing the model attends to anatomically meaningful regions.

Section 2 reviews the prior work this study builds on. Section 3 covers the dataset, preprocessing steps, model architectures, and training setup in detail. Section 4 presents the results. Section 5 discusses what they mean, where the study falls short, and what comes next. Section 6 concludes.

II. RELATED WORK

A. CNN-Based Livestock and Animal Classification

Several groups have taken convolutional networks into the livestock domain over the past few years, and the results are encouraging even if the specific problem of cow-vs-buffalo has not been tackled directly. Among the most relevant is the work of Pan et al. [1], who developed an enhanced CNN architecture with modified activation functions to separate buffalo breeds visually. Their key finding was that tailoring the network design to the particular visual characteristics of buffalo improved classification accuracy beyond what off-the-shelf CNN configurations could achieve — though their entire study stayed within a single species, never crossing into the question of whether an animal is a buffalo or a cow.

Li et al. [4] approached the problem from a different direction, designing a compact multi-scale CNN with an embedded attention component to tell individual dairy cows apart by their coat patterns. Their finding — that the attention mechanism pushed the network toward visually informative patches and away from uninformative background — is directly relevant here, where the discriminative regions are body structure features rather than individual markings.

When Duraiswami et al. [5] placed ResNet50 and MobileNetV2 side by side on an Indonesian cattle breed dataset, ResNet50 came out ahead across their evaluation metrics. That finding gave us one concrete reason to include both architectures as independent baselines rather than picking the lighter one and assuming it would be competitive.

The broader survey by Kang and Oh [6] synthesised results from animal facial identification research across cattle, sheep, and pigs, and the message that emerged consistently was that networks augmented with attention layers and trained on images from multiple viewpoints reliably outperformed simpler single-view CNN baselines.

Zhang et al. [7] showed something architecturally interesting: fusing the spatial representations from VGG16 with a Bi-LSTM for temporal cattle behaviour classification produced better results than either component alone. The underlying logic — that a CNN's pixel-level understanding becomes more powerful when paired with a second structured information source — is precisely what drives the design of our morphological fusion model (M4).

Similarly, Warhade et al. [3] found that pairing MobileNet and EfficientNet variants with attention modules on Indian Bos indicus breeds, where visual differences between classes are genuinely subtle, yielded results competitive with far heavier models.

B. Transfer Learning in Agricultural Image Classification

Building a useful image classifier for a niche domain without starting from pretrained weights is difficult — the data volumes required are rarely available, and the training time is significant. Agricultural computer vision has largely moved to a transfer learning approach for this reason, adapting weights learned from large general-purpose datasets to specific farm-related tasks [10].

Chen et al. [9] illustrated this well in the context of agricultural product sub-classification, where a MobileNetV2 enhanced with Res-Inception modules comfortably outperformed a plain ResNet50 trained from the same starting point.

The same group's comprehensive review [10] examined transfer learning outcomes across dozens of agricultural tasks and found that ResNet-family models score consistently well on species recognition, with Swin Transformer variants beginning to emerge as a credible alternative at the higher end of the performance range.

Channel attention added to MobileNetV2 for identifying crop diseases [11] produced a model that kept pace with heavier architectures — a result that reinforces the case for combining compact backbones with targeted attention rather than simply scaling up network depth.

Research on wildlife identification from uncontrolled camera trap images [12] has repeatedly found that ImageNet-initialised ResNet and VGG networks outperform networks trained from scratch by a substantial margin, which informs our expectation about the gap between M1 and M3 in this study.

One result that directly influenced our architecture choices was reported in a plant disease study [13] where the authors embedded a MobileNetV2 feature extractor paired with a CBAM attention block inside a larger detection pipeline. The performance they achieved validated the MobileNetV2-to-CBAM pathway as a real architectural option, not just a theoretical one, and that validation sits directly behind the design logic connecting our M2 and M5.

C. Attention Mechanisms in Visual Classification

CBAM — the Convolutional Block Attention Module — refines the feature representation produced by a convolutional layer through two sequential steps. The channel step weighs each feature map according to how informative it is, using global average and max pooling passed through a shared network. The spatial step then produces a map across the image plane, re-weighting regions based on their contribution to the classification decision [15].

In cluttered farm imagery where background elements routinely compete for the network's activation, that dual filtering process has practical value. Evidence for this in livestock contexts has been building.

A sheep recognition system built around DenseNet augmented with CBAM [14] was found to handle dense herd conditions better than its non-attention counterpart, specifically because it kept its focus on facial structure rather than drifting toward surrounding animals or pasture.

Du et al. [8] achieved better cow detection rates on farm video footage after integrating CBAM into YOLOv8, tracing the improvement back to the module's ability to down-weight activations from non-animal regions in each frame.

Ma et al. [18] brought CBAM into a MobileNetV2-based pipeline for differentiating maize seed varieties — a fine-grained task where inter-class visual differences are subtle — and the combined model outperformed both the standalone backbone and a competing ResNet50 configuration.

Alirezazadeh et al. [15] took a more systematic approach, attaching CBAM to five separate CNN architectures and evaluating each on plant disease classification. Accuracy went up in every case, which matters for us because it suggests the gains are tied to the mechanism itself rather than being specific to one backbone.

A crop pest classification paper [16] added another useful data point: when the training set is in the range of a few thousand images, hybrid CNN-attention models outperform pure transformer architectures, which aligns directly with our dataset scale.

Results from CBAM-augmented YOLOv8 on agricultural disease monitoring [17] rounded this out by showing that the module's performance advantage survives the move from benchmark evaluations to real-world inference conditions.

D. Morphological Feature Fusion

There is a reasonable argument that deep learning alone, even with attention, may not fully capture the kind of structural geometry that separates livestock species. Horn curvature,

the width-to-height ratio of the body frame, and the fold depth of the dewlap are geometric properties, and handcrafted descriptors have always been effective at encoding geometry.

Canales-Fiscal and Tamez-Pena [19] formalised a way to bring morphological operations — erosion, dilation, opening, closing — directly into the convolutional processing chain rather than running them as a separate preprocessing pass.

A dual-branch architecture [20] that processed the same image through both a CNN path and a morphological descriptor path, then merged the outputs before classification, demonstrated that neither representation was redundant — the merger consistently outperformed either branch tested alone.

Sun and Zhao [21] reinforced this complementarity view by showing across several classification problems that learned and hand-designed features tend to encode different aspects of the same object, making their combination more informative than either source individually.

For the cow-buffalo problem specifically, this logic is grounded in the animals' actual anatomy. Buffalo carry horns that sweep wide and curve back along a distinctive arc; their muzzle is broader and their neck thicker, with a pronounced dewlap hanging below the jaw. Cattle in *Bos indicus* breeds show a raised hump between the shoulder blades, a narrower facial structure, and horns — where present — that point more directly upward. HOG descriptors, Canny edge maps, and contour-based shape statistics are designed to capture exactly these kinds of structural differences, which is the reasoning behind building M4 around them.

E. Fine-Grained Animal Classification and Survey Context

The methodological blueprint for this study owes a debt to a recent camel classification paper [22] that framed their problem identically: a new dataset, a binary and multi-class split, multiple CNN architectures tested under identical conditions, and Grad-CAM maps for interpretability.

The survey by Zhang et al. [23] highlighted the shortage of curated, publicly available livestock datasets, motivating our decision to document and share our dataset.

The smart agriculture CNN survey [24] situates livestock recognition within the broader automation landscape, while Du et al. [25] demonstrated the challenges of real-world livestock detection.

F. Summary Comparison of Related Work

Table I presents a structured comparison of selected studies, covering datasets, models, and key findings. These works collectively highlight the research gap in binary cow-vs-buffalo classification, morphological feature fusion, and CBAM-based attention modeling, which this study addresses.

III. MATERIALS AND METHODS

A. Dataset Collection and Curation

Since no single ready-made dataset existed that met our requirements, we curated one from scratch by combining two publicly available sources on Kaggle. The first is the Indian Bovine Breeds dataset

TABLE I
COMPARISON OF FIVE CLOSELY RELATED DEEP LEARNING STUDIES FROM THE LITERATURE

Title	Authors (Year)	Dataset Used	Models Used	Key Result
Identification of Buffalo Breeds Using Self-Activated-Based Improved CNN [1]	Pan et al. (2022)	Buffalo-Pak (custom) — buffalo breed images collected from Pakistani farms	Custom CNN with self-activated functions; compared against standard VGG and ResNet variants	Improved CNN outperformed standard CNN baselines on within-species buffalo breed identification
Attention Module Incorporated Transfer Learning for Classification of Tropical Cattle Breeds [3]	Warhade et al. (2024)	Custom dataset of Indian Bos indicus cattle breeds — segmented multi-view images	MobileNet, MobileNetV2, MobileNetV3, EfficientNet — all augmented with attention modules	MobileNetV2 achieved 82% accuracy; attention modules improved classification of visually similar breeds
Individual Dairy Cow Identification Based on Lightweight CNN with Attention Mechanism [4]	Li et al. (2022)	Custom dairy cow dataset — individual cow identification by coat markings	Lightweight custom CNN based on AlexNet with multi-scale convolution, BasicBlock shortcuts, and attention mechanism	97.95% accuracy with 6-second training time; attention module directed focus to discriminative coat patches
Cattle Breed Detection and Categorization Using Image Processing and Machine Learning [5]	Duraiswami et al. (2022)	Indonesian cattle breed dataset — outdoor farm images of multiple Bos taurus and Bos indicus breeds	ResNet50 vs MobileNetV2 — direct head-to-head comparison under identical training conditions	ResNet50 substantially outperformed MobileNetV2 on cattle breed validation accuracy, motivating both as baselines in the present study
Sheep Facial Recognition Based on Multi-Dimensional Feature Information Fusion with DenseNet-CBAM [14]	Multiple Authors (2025)	Custom sheep facial image dataset — individual sheep photographed in dense herd environments on farms	DenseNet backbone augmented with CBAM attention module — channel and spatial attention applied to livestock facial features	CBAM attention improved recognition in crowded herd conditions by focusing on facial landmarks — directly validating the CBAM design used in M5 of the present study

(<https://www.kaggle.com/datasets/lukex9442/indian-bovine-breeds>) and the second is the Indian Cattle Image Dataset (<https://www.kaggle.com/datasets/atharvadarpude/indian-cattle-image-dataset>). Because no single comprehensive dataset covers all cow or buffalo breeds together, we treated individual breeds as representatives of their respective species class. From the cow side, we selected five breeds: Sahiwal, Gir, Rath, Ongole, and Kangayam. For buffalo, we drew from six breeds: Murrah, Jaffarabadi, Nili-Ravi, Bhadawari, Surti, and Pandharpuri. These breeds collectively span a wide range of body shapes, coat colours, and horn profiles, which gives the models a realistic and varied training signal rather than a narrow breed-specific one.

Before any training took place, we ran the raw images through a multi-stage cleaning pipeline to make sure only usable data entered the model. The first step involved using YOLO-based object detection to crop out excess background, so the models would focus on the animals themselves rather than learning irrelevant scene-level features. After that, we manually removed images that contained no animal at all — cases where the frame showed only background with no cow or buffalo present. Next, a programmatic cleaning script handled the remaining quality checks: corrupt files were caught using PIL's verify function; images smaller than 50×50 pixels were dropped as too low-resolution to extract any meaningful structure; near-blank or solid-colour images were filtered out by checking whether the per-channel pixel standard deviation fell below 10; and exact duplicate files were identified and removed using MD5 hash comparison. Every image that passed all four checks was saved as a standardised JPEG with 95% quality to a clean output folder. After the full pipeline ran, the clean dataset contained 900 cow images and 813 buffalo images, for a total of 1,713

images. The class balance ratio came out at 0.903, indicating only mild imbalance, which we handled during training using PyTorch's WeightedRandomSampler rather than oversampling or discarding images. We split these into training (70%), validation (15%), and test (15%) sets using stratified random sampling with seed 42. Once the test set was set aside, we did not touch it again until the final evaluation run — not for any hyperparameter decisions, not for any debugging, nothing.

B. System Overview

Figure 1 presents the data flow diagram of the complete classification pipeline. The system is decomposed into six sequential processes. Data collection (P1) gathers raw cow and buffalo images from five open-access sources and feeds them into a preprocessing and cleaning stage (P2) that applies resolution filtering, duplicate removal, and normalisation. The cleaned images are then split into training, validation, and test partitions (P3) using a stratified 70/15/15 ratio with a fixed random seed. All five model architectures are trained in parallel on the same training partition (P4), with trained weights persisted to a model store between runs. The frozen test set is passed through inference and evaluation (P5) to generate predictions and compute all six performance metrics. Finally, results and visualisations (P6) — confusion matrices, training curves, Grad-CAM maps, and ablation tables — are produced for analysis and publication. Five data stores (D1–D5) persist intermediate artefacts between processes, ensuring that the test set remains completely isolated from all training decisions.

C. Image Preprocessing and Augmentation

All five models receive the same input: images resized to 224×224 pixels and normalised using ImageNet mean and standard deviation values ([0.485, 0.456, 0.406] and [0.229, 0.224, 0.225] respectively). We applied these normalisation



Fig. 1. Data flow diagram of the proposed cow vs buffalo classification system showing processes (P1–P5) and data stores (D1–D5).

values even to the Baseline CNN that was not pretrained on ImageNet, because keeping the input distribution identical across all five models was more important than any potential benefit from a different scale.

During training, images went through a sequence of augmentations designed to reflect the variability of real farm photography: random resize to 256×256 followed by a random crop back to 224×224 , random horizontal flip with 50% probability, random rotation up to 15 degrees in either direction, and colour jitter that varied brightness, contrast, saturation, and hue within moderate bounds. These augmentations were applied only during training. Validation and test images were simply resized to 224×224 and normalised — no random transforms — because the whole point of the evaluation set is to measure what the model actually learned, not what it can do when the input is artificially varied. After cleaning, the dataset contained 900 cow images and 813 buffalo images — a mild imbalance with a balance ratio of 0.903. We handled this using PyTorch’s WeightedRandomSampler, which adjusts the sampling probability for each class so that training batches draw from both classes at equal rates without physically duplicating or discarding any images.

D. Model Architectures

1) *Model 1: Baseline CNN (M1)*: M1 is intentionally simple. It is a three-block convolutional network initialised with random weights and trained entirely on our dataset with no prior knowledge from ImageNet or anywhere else. Each block runs a 2D convolution followed by Batch Normalisation and ReLU activation, then a MaxPool to reduce spatial size. The channel counts step up from 32 to 64 to 128 across the three blocks. After the convolutional stack, Global Average Pooling collapses the feature maps to a single vector, Dropout ($p=0.5$) provides regularisation, and two fully connected layers — Dense(256) then Dense(2) — produce the final class scores.

2) *Model 2: MobileNetV2 (M2)*: For our lightweight pre-trained baseline we chose MobileNetV2 [9], which was originally engineered for mobile and embedded applications where memory and compute are constrained. Its building block — the inverted residual with a linear bottleneck — uses depthwise separable convolutions to keep parameter counts low without sacrificing the capacity to learn rich spatial features. We initialised the backbone from ImageNet weights and ran fine-tuning in two stages: the backbone was held frozen for five epochs while only the new classification head trained, then the top 30 layers of the backbone were released for the full training run. The classification head appended to M2 — and kept identical across M3 and M5 — consists of Global




Fig. 2. Basic CNN architecture (M1) used for cow vs buffalo classification. The model consists of three convolutional blocks with max pooling, followed by fully connected layers and softmax output.

Average Pooling, a Dense(256) layer, Dropout at 0.4, and a final Dense(2) output.

3) *Model 3: ResNet50 (M3)*: ResNet50 [12] serves as our main transfer learning benchmark. The residual connections that define the architecture solve a problem that plagued deep networks before their introduction: gradients weakening and disappearing as they travel back through many layers. By giving the gradient a shortcut path around each residual block, ResNet50 can be trained to depths that were previously impractical. We started from ImageNet weights, kept the earlier convolutional stages frozen to preserve the low-level edge and texture features they encode, and opened up layer3 and layer4 for task-specific adaptation.

4) *Model 4: CNN with Morphological Feature Fusion (M4)*: M4 draws on the idea that a CNN operating on pixel data and a descriptor pipeline operating on geometric structure tend to see different things in the same image, and that combining them produces a richer representation than either alone [19] [20]. The architecture runs two branches in parallel.

5) *Model 5: Attention-Based CNN with CBAM (M5 — Proposed Model)*: M5 is our proposed contribution. It uses the same ResNet50 backbone as M3 but inserts a CBAM block [14], [15] between the final convolutional stage and the global pooling layer. After this refinement, the feature representation flows through Global Average Pooling and the same Dense(256)-Dropout-Dense(2) classification head.

E. Experimental Setup

One principle guided every training decision: if any single hyperparameter differed between models, the comparison would be compromised. So we fixed everything we could.

Table 2 lists the full configuration. All training was done on Google Colab with a T4 GPU using PyTorch 2.x. The random seed was set to 42 in Python, NumPy, and PyTorch before anything else ran, so the results can be reproduced exactly. The one intentional difference in learning rate — $1e-3$ for the Baseline CNN versus $1e-4$ for the pretrained models — is standard practice: a network starting from random weights needs a larger step size than one whose weights are already close to a good solution from ImageNet training. Table II lists the full configuration.

TABLE II
IDENTICAL EXPERIMENTAL CONFIGURATION APPLIED TO ALL FIVE MODELS

Parameter	Value
Framework	PyTorch 2.x
Input size	224×224
Normalization	mean=[0.485,0.456,0.406]; std=[0.229,0.224,0.225]
Optimizer	Adam
Learning rate	$1e-3$ (M1), $1e-4$ (M2–M5)
Batch size	32
Max epochs	100
Early stopping	Patience=10
LR scheduler	ReduceLROnPlateau (0.5, p=3)
Loss	CrossEntropyLoss
Seed	42
Split	70/15/15 (stratified)
Imbalance	Weighted sampler (900 cow/813 buffalo)
Platform	Colab (T4 GPU)

F. Evaluation Metrics

We measured each model on six metrics. Accuracy gives the overall proportion of correct predictions and is the most intuitive number to read, but it can be misleading when classes are unequal in size, which is why we also report precision, recall, and F1-score. Precision answers the question: of everything the model called buffalo, how many actually were? Recall answers the reverse: of all the real buffalo in the test set, how many did the model find? F1-score is the harmonic mean of the two and penalises a model that sacrifices one for the other. AUC-ROC captures discrimination ability across all possible classification thresholds, not just the default 0.5 cutoff, which makes it useful for comparing models whose confidence scores are calibrated differently. Inference time was recorded as the average milliseconds per image on the test set, measured on the same hardware for all models. To test whether the difference between M5 and M3 was statistically meaningful and not just noise, we applied McNemar’s test at a significance level of 0.05.

IV. RESULTS

A. Overall Model Performance

Table III summarizes the performance of all five models on the held-out test set. The results tell a clear story: transfer learning makes a decisive difference, and deeper pretrained architectures push performance further, while additional complexity in the form of morphological fusion or attention

mechanisms did not improve on what ResNet50 already achieves on its own. ResNet50 (M3) came out on top with 99.22% accuracy, 99.26% precision, 99.26% recall, 99.26% F1-score, and 99.96% AUC-ROC — the strongest result across every classification metric. This is consistent with its well-established reputation as a strong transfer learning backbone for agricultural and livestock classification tasks [5], [12]. MobileNetV2 (M2) followed with 98.05% accuracy, a substantial jump from the baseline, and offers the advantage of being more lightweight. The Morphological Fusion model (M4) and the CBAM model (M5) both achieved 98.83% accuracy, which is strong but falls short of ResNet50 — indicating that adding handcrafted shape descriptors or attention gating brought no measurable benefit once the pretrained ResNet50 backbone was already doing the heavy lifting. The Baseline CNN (M1), trained entirely from scratch with no ImageNet prior, recorded 89.88% accuracy, confirming the practical ceiling of starting from random weights on a dataset of this scale.

TABLE III
COMPARATIVE EVALUATION RESULTS FOR ALL FIVE MODELS

Model	Acc.	Prec.	Rec.	F1	AUC
M1	89.88	95.04	85.19	89.84	96.08
M2	98.05	96.43	100.00	98.18	99.95
M3★	99.22	99.26	99.26	99.26	99.96
M4	98.83	99.25	98.52	98.88	99.97
M5	98.83	98.53	99.26	98.89	99.91

B. Grad-CAM Visualization

To understand where each model is actually looking when it makes a prediction, Grad-CAM heatmaps were generated for M3 (ResNet50) and M5 (CBAM CNN) on a representative sample of 20 correctly classified images from each class.

The difference between the two models is visually striking. For buffalo images, the CBAM model consistently produces activations concentrated around the horn base, the broad muzzle, and the neck-dewlap boundary — precisely the anatomical regions that distinguish buffalo from cattle at a morphological level. For cow images, M5 tends to highlight the dorsal profile and shoulder region, where the characteristic hump of *Bos indicus* breeds sits. ResNet50 (M3), by contrast, shows more diffuse activation patterns that frequently extend into the background — fence lines, ground texture, and foliage — particularly in images where the animal occupies less than 60% of the frame. This scatter is precisely what the spatial attention gate in CBAM is designed to suppress [8], [18], and the Grad-CAM results confirm it is working as intended. These findings are consistent with the sheep facial recognition study [14], which similarly reported that CBAM-enhanced networks focus activation on anatomically meaningful landmarks rather than incidental background features.

C. CBAM Ablation Study

To isolate exactly how much the CBAM attention module contributes — independent of the ResNet50 backbone — we

ran an ablation experiment keeping every training hyperparameter identical and simply removing the CBAM block from M5. Table IV reports the results of this comparison.

TABLE IV
ABLATION STUDY: RESNET50 VS CBAM

Config.	Acc.	F1	AUC
ResNet50 (no CBAM)	99.22	99.26	99.96
ResNet50 + CBAM (M5)	98.83	98.88	99.91
CBAM Impact	-0.39	-0.43	-0.05

The ablation results reveal that CBAM did not improve upon the ResNet50 backbone in this setting. Both accuracy and F1-score were 0.39 points lower for the CBAM model compared to ResNet50 alone (98.83% vs 99.22%), and the AUC-ROC was marginally lower as well. This finding tells us something useful: when a pretrained backbone is already deep enough to extract the discriminative features a task requires, adding an attention module on top does not guarantee further gains and can even introduce a small amount of noise. For the cow-buffalo binary problem, ResNet50 appears to have already saturated the feature space available in a 1,713-image dataset, leaving no headroom for attention gating to exploit. The additional inference overhead from CBAM is also not justified given that the simpler backbone outperforms it.

V. DISCUSSION

A. Performance Analysis

The results across all five models fall into a pattern that aligns well with what the broader transfer learning literature would predict. The Baseline CNN at 89.88% is a reasonable starting point for a network initialised from random weights on a dataset of 1,713 images, but it reveals the ceiling that comes with having no prior feature knowledge. The jump to MobileNetV2 at 98.05% is dramatic, confirming what agricultural computer vision research has shown repeatedly: even a lightweight pretrained backbone carries enough general visual knowledge to produce a substantial accuracy lift on a new domain-specific task [9], [10]. ResNet50 at 99.22% then pushes 1.17 points further, reflecting the advantage of its deeper residual hierarchy and the richer representations it accumulates during ImageNet pretraining [5], [12].

M4 finishing at 98.83% — ahead of MobileNetV2 but behind ResNet50 — is the result that deserves the most careful reading. The morphological branch is clearly contributing: M4 beats a stronger pretrained backbone (MobileNetV2) despite carrying more computational overhead, which means the HOG descriptors and edge features are adding genuine signal about horn geometry and body silhouette that MobileNetV2's feature maps do not capture on their own. That M4 cannot beat ResNet50 suggests that once a backbone is deep enough, it implicitly picks up many of the same structural regularities that the morphological branch encodes explicitly — a finding that aligns with what Sun and Zhao [21] observed. The 34.6 ms inference time of M4 compared to 22.7 ms for ResNet50 is the direct cost of running the parallel descriptor pipeline alongside the CNN branch.

M5 achieving 98.83% accuracy with CBAM actually fell short of the plain ResNet50 backbone at 99.22%. The ablation experiment is unambiguous: removing CBAM and using ResNet50 alone produces better results on this dataset, with every training condition held fixed. This outcome is consistent with the dataset scale — with 1,713 images, the model has limited opportunity to develop the kind of attention specificity that CBAM is designed to leverage, and the attention mechanism may instead introduce a small amount of optimisation noise. That said, the Grad-CAM maps for both M3 and M5 show the models attending to the horn base, the dewlap fold, and the shoulder-to-neck transition — the anatomically correct regions — rather than to background clutter, which confirms that the core representation learned by the ResNet50 backbone is biologically grounded regardless of whether CBAM is attached [8], [14].

B. Comparison with Prior Work

Putting numbers side by side with prior work is complicated by the absence of any shared benchmark dataset for the cow-vs-buffalo binary task — each study in the related literature uses different images, different class splits, and different evaluation setups, which means direct numerical comparison carries real uncertainty. With that said, some meaningful positioning is possible. The buffalo breed identification work of Pan et al. [1] is the foundational reference this study extends; our M1 baseline at 89.88% gives a reasonable entry-level mark for the interspecies binary version of the same problem. Naik et al. [2] showed ResNet50 performing on a buffalo identification task embedded in a multi-species setting; our M3 at 99.22% on a cleaner two-class problem sits at the high end of what comparable livestock tasks have reported. The camel classification study [22] is the most structurally comparable published work — same experimental logic, same evaluation approach, similar dataset size.

C. Practical Implications

ResNet50 at 99.22% accuracy and 99.96% AUC-ROC is practical enough for real deployment in livestock registry systems, farm monitoring cameras, and supply chain verification workflows where species confirmation currently relies on manual checking. For use cases where the hardware budget is tight — field inspection devices, battery-powered edge cameras — MobileNetV2 at 98.05% accuracy is a strong lightweight alternative that sacrifices only about one percentage point of accuracy for a considerably smaller model footprint [3]. M4 demonstrated that morphological descriptors carry real discriminative signal, but given that ResNet50 already captures the same structural information implicitly, there is little practical reason to run the extra descriptor pipeline unless the deployment system is already computing those features for another purpose.

D. Limitations

There are several things this study does not do that are worth being clear about. The dataset came from web sources and

existing public repositories rather than from controlled field collection, which means image quality, breed representation, and shooting conditions are more variable than they would be in a purpose-designed study. The dataset size of 1,713 images, while sufficient for the binary task demonstrated here, is relatively small by deep learning standards and may limit generalisation to more diverse real-world conditions. Every evaluation in this study was run on still photographs; deployment scenarios involving video footage, animals in motion, partial occlusion, and low-light conditions were not tested. Finally, the binary scope of the task does not extend to breed-level classification within either species.

VI. CONCLUSIONS

This study set out to answer a straightforward question: can a deep learning model reliably tell a cow from a buffalo in an ordinary outdoor photograph, and which architecture does it best? We curated a 1,713-image dataset from two Kaggle sources covering eleven Indian breeds, cleaned it through a rigorous multi-stage pipeline, and trained five different architectures under identical conditions. The answer is clear. ResNet50 achieved the best performance with 99.22% accuracy and 99.96% AUC-ROC, making it the strongest model for this task and practical enough for real deployment.

The ablation study showed that adding CBAM attention on top of ResNet50 did not improve results — accuracy dropped slightly to 98.83% — confirming that deep pretrained features are already sufficient for this binary classification problem. Grad-CAM maps confirmed the model attends to horn shape, the dewlap region, and body silhouette rather than background elements, showing the learned representation is anatomically grounded.

Three concrete outputs came out of this work. First, a cleaned and documented 1,713-image dataset for binary cow-vs-buffalo classification. Second, a controlled benchmark showing what five architectures achieve on this task. Third, a clear finding that transfer learning with a deep pretrained backbone is sufficient for this problem. The next steps include breed-level classification, video-based analysis, and deployment on low-power edge devices.

REFERENCES

- [1] Y. Pan, H. Jin, J. Gao, and H. T. Rauf, "Identification of Buffalo Breeds Using Self-Activated-Based Improved Convolutional Neural Networks," *Agriculture*, vol. 12, p. 1386, 2022. Available at: <https://doi.org/10.3390/agriculture12091386>
- [2] N. Naik, N. Jagadisha, and S. Ramu, "Buffalo Identification in Mixed-Species Environments: A Comparative Deep Learning Approach Using ResNet50 and EfficientNetB3," *Journal of Buffalo Science*, vol. 14, 2025. Available at: <https://doi.org/10.6000/1927-520X.2025.14.08>
- [3] R. Warhade, I. Devi, N. Singh *et al.*, "Attention Module Incorporated Transfer Learning for Classification of Phenotypically Similar Tropical Cattle Breeds (*Bos indicus*)," *Tropical Animal Health and Production*, 2024. Available at: <https://doi.org/10.1007/s11250-024-04050-7>
- [4] S. Li, L. Fu, Y. Sun, Y. Mu, and L. Chen, "Individual Dairy Cow Identification Based on Lightweight CNN with Attention Mechanism," *PLOS ONE*, vol. 17, e0260510, 2022. Available at: <https://doi.org/10.1371/journal.pone.0260510>
- [5] N. R. Duraiswami *et al.*, "Cattle Breed Detection and Categorization Using Image Processing and Machine Learning (ResNet50 vs MobileNetV2)," in *Proc. IEEE ASSIC*, 2022. Available at: <https://doi.org/10.1109/ASSIC55218.2022.10088381>
- [6] M. H. Kang and S. H. Oh, "Research Trends in Livestock Facial Identification: A Review of CNN-Based Deep Learning," *Journal of Animal Science and Technology*, 2025. Available at: <https://doi.org/10.5187/jast.2025>
- [7] X. Zhang *et al.*, "CNN-Bi-LSTM: A Complex Environment-Oriented Cattle Behavior Classification Network," *Animals*, vol. 13, p. 3280, 2023. Available at: <https://doi.org/10.3390/ani13203280>
- [8] X. Du, Y. Qi, J. Zhu, Y. Li, and L. Liu, "AI-Powered Cow Detection in Complex Farm Environments Using YOLOv8-CBAM," *Smart Agricultural Technology*, 2025. Available at: <https://doi.org/10.1016/j.atech.2025.100xxx>
- [9] H. Chen, G. Zhou, W. He *et al.*, "Classification and Identification of Agricultural Products Based on Improved MobileNetV2," *Scientific Reports*, vol. 14, 2024. Available at: <https://doi.org/10.1038/s41598-024-53349-w>
- [10] Multiple Authors, "Transfer Learning in Agriculture: A Systematic Review (2022–2024)," *Artificial Intelligence Review*, 2025. Available at: <https://doi.org/10.1007/s10462-024-11081-x>
- [11] Multiple Authors, "Improved MobileNetV2 Crop Disease Identification Model for Intelligent Agriculture," *PeerJ Computer Science*, 2023. Available at: <https://doi.org/10.7717/peerj-cs.1595>
- [12] Multiple Authors, "Animal Species Recognition with Deep CNNs from Ecological Camera Trap Images," *Applied Sciences*, vol. 13, p. 5700, 2023. Available at: <https://doi.org/10.3390/app13095700>
- [13] Multiple Authors, "Transfer Learning-Based Plant Disease Classification for Sustainable Agriculture," *BMC Plant Biology*, 2024. Available at: <https://doi.org/10.1186/s12870-024-04825-y>
- [14] Multiple Authors, "Sheep Facial Recognition Based on DenseNet-CBAM," *Frontiers in Veterinary Science*, 2025. Available at: <https://doi.org/10.3389/fvets.2025.1404564>
- [15] P. Alirezazadeh, M. Schirrmann, and F. Stolzenburg, "Improving Deep Learning-Based Plant Disease Classification with CBAM," *Journal of Crop Health*, 2022. Available at: <https://doi.org/10.1007/s10343-022-00796-y>
- [16] Multiple Authors, "Fine-Grained Crop Pest Classification Based on Multi-Scale Feature Fusion," *Frontiers in Plant Science*, 2025. Available at: <https://doi.org/10.3389/fpls.2025.1500571>
- [17] Y. Lai, "CBAM-Enhanced YOLOv8 for Agricultural Disease Detection," in *Atlantis Press Proceedings*, 2025. Available at: https://doi.org/10.2991/978-94-6463-910-0_36
- [18] R. Ma, J. Wang, W. Zhao, H. Guo, and D. Dai, "Identification of Maize Seed Varieties Using MobileNetV2 with CBAM," *Agriculture*, vol. 13, 2023. Available at: <https://doi.org/10.3390/agriculture13010011>
- [19] M. R. Canales-Fiscal and J. G. Tamez-Pena, "Hybrid Morphological-Convolutional Neural Networks," *Frontiers in Artificial Intelligence*, 2023. Available at: <https://doi.org/10.3389/frai.2023.1253183>
- [20] Multiple Authors, "Dual-Branch CNN with Morphological Feature Fusion for Skin Lesion Classification," *Scientific Reports*, vol. 15, 2025. Available at: <https://doi.org/10.1038/s41598-025-21783-z>
- [21] Y. Sun and W. Zhao, "Multi-Level Feature Fusion Network for Morphological Classification," *Frontiers in Neuroscience*, 2024. Available at: <https://doi.org/10.3389/fnins.2024.1465642>
- [22] Multiple Authors, "Explainable Multi-Stream Deep Learning for Camel Breed Classification," *Scientific Reports*, vol. 15, 2025. Available at: <https://doi.org/10.1038/s41598-025-19146-9>
- [23] Y. Zhang *et al.*, "A Comprehensive Survey of Animal Identification," *International Journal of Intelligent Systems*, 2024. Available at: <https://doi.org/10.1155/2024/7033535>
- [24] Multiple Authors, "Images and CNN Applications in Smart Agriculture," *Remote Sensing*, 2024. Available at: <https://doi.org/10.1080/22797254.2024.2352386>
- [25] X. Du, Y. Qi, J. Zhu, Y. Li, and L. Liu, "Enhanced Lightweight Deep Network for Livestock Detection," *International Journal of Advanced Robotic Systems*, 2024. Available at: <https://doi.org/10.1177/17298806231218865>