



International Journal of Multidisciplinary Research in Science, Engineering and Technology

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)



Impact Factor: 8.206

Volume 9, Issue 4, April 2026



International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

Implementing CNN to Detect Fake Currency Using Deep Learning Techniques

S. Keerthi, R. Asreetha, V. Govardhan Sudev, R. Siva Srinivas

UG Student, Department of ECE, R.V.R&J.C.C. E, Chowdavaram, India

UG Student, Department of ECE, R.V.R&J.C.C. E, Chowdavaram, India

UG Student, Department of ECE, R.V.R&J.C.C. E, Chowdavaram, India

UG Student, Department of ECE, R.V.R&J.C.C. E, Chowdavaram, India

ABSTRACT: Fake money has troubled economies everywhere, even in India. A bogus banknote slips through without official backing. Instead of trust, these copies spread dishonesty. Progress slows when false cash flows into markets. No state support exists for such made-up notes. They twist financial systems slowly over time. Real growth fades where fakes take hold. Illicit notes undermine public confidence quietly. Without control, harm spreads across trade areas. These unauthorized prints damage national stability.

One-way machines learn today uses something called a Support Vector Machine, tested before on spotting fake money but slowed down by heavy computing needs. Instead, trying a different path now means turning to convolutional neural networks to catch false bills. This shift aims at handling the work faster without getting stuck in complex calculations. What comes next builds on patterns seen in images, letting the system notice tiny flaws human eyes might miss. Learning happens step by step, adjusting itself each time it sees new examples. Not every method fits all cases, yet this one adapts more smoothly to fine details in paper and ink. Results so far suggest fewer errors when telling real from imitation notes. Speed improves too, thanks to how these networks break down visual data. Past limits around processing power start fading with smarter design choices. Each test run refines accuracy just a bit further.

KEYWORDS: Counterfeit Currency Detection, Convolutional Neural Network (CNN), Image Processing, Deep Learning, Computer Vision, Pattern Recognition

I. INTRODUCTION

Fake cash causes big problems, shaking up economies and putting national finances at risk - especially in places like India. When phony bills circulate, people lose money along with shops and companies. Criminals thrive on these forged notes, using them to fund shady operations while prices rise across markets. Confidence in paper money fades too. Modern tools like high-res printers and scanners now let fraudsters copy genuine currency almost perfectly. Telling fakes apart by eye alone gets harder every year.

Looking closely at fake money often means UV lights, watermark checks, security thread scans - tools that demand training. Mistakes happen when people do it by hand, plus it eats up time. Real-world spots like shops or banks face pressure where speed matters. Machines help, yet relying on them isn't always possible mid-transaction.

These days, systems that learn from data are helping spot fake money more easily. Older ways relied on tools such as Support Vector Machines, which needed people to pick out key traits by hand - slowing things down while using more computing power. On the flip side, networks inspired by brain cells, particularly those good at reading images, have done remarkably well telling patterns apart. Their ability to classify visuals has made a noticeable difference in accuracy.

Starting off strong, convolutional neural networks pick up intricate image traits without needing manual input. Instead of relying on old techniques, they notice fine textures, shifts in hue, plus tiny layout clues that set real money apart from fake. Because these models run quickly, they fit right into live settings where speed matters. Their consistent results come from layered learning, built through exposure to many examples over time.



International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

This study walks through building a solid tool to catch fake bills using neural networks that see patterns like eyes do. Its purpose sits firmly on lifting how well machines spot fakes, cutting down on people checking every note by hand, while delivering fast and steady results when scanning cash. Places such as banks, shops, or finance offices might put it to work easily, raising safety levels and shrinking the money lost because of forged notes.

II. LITERATURE REVIEW

Money fake spotting matters more than most think - it can shake economic stability. Looking at telltale marks like hidden threads, color shifts, or watermark shadows used to be the go-to move. These older methods? They lean on slow checks by trained eyes. Mistakes pop up now and then - making old-school ways shaky at best.

From the beginning of machine learning, experts relied on tools like SVM, KNN, besides Decision Trees to spot counterfeit money without human help. Instead of automatic recognition, these systems needed people to pull image details by hand before analysis began. Good outcomes were possible - still, effort and time grew fast due to messy preparation steps that slowed everything down.

Right now machines that think like brains spot fake cash using patterns in pictures. These systems grab details from photos without needing heavy prep work first. Because they adapt well, even dim light or blurry shots don't stop them from working. Sometimes shadows or grain mess up regular methods, yet these models still manage. Getting enough examples to train them is tough though - it takes lots of images to get it right. Without plenty of samples, mistakes happen more often than not.

These days, researchers aim to make CNNs work better and handle larger tasks so they can actually be used in everyday tech. Models such as MobileNet and EfficientNet came out of this push - smaller designs that still keep strong accuracy while using less computing power, perfect for fast jobs on mobile gadgets. Another path taken involves expanding training data through techniques like image rotation or flipping; doing this helps networks adapt when faced with varied lighting, angles, or backgrounds.

One way forward might be mixing CNNs with different methods to push detection precision higher. Instead of relying solely on visuals, hybrid setups using OCR along with blended traits aim at reading both images and words found on money. Even though researchers have built quite precise tools, problems remain - like limited samples for learning and trouble recognizing torn or worn bills.

III. PROPOSED METHOD

A fresh look at spotting fake money begins with a smart pattern-spotting method built on CNNs. Instead of relying on people to check every detail, the process learns straight from pictures of bills. Accuracy climbs when the model handles feature detection itself. What emerges is a workflow that sorts real from false without heavy manual checks.

Starting off, pictures of money come in - snapped by phones or regular cameras. Different light setups, angles, or backdrops appear on purpose; that way the system learns better. After gathering them, each photo gets adjusted - scaled down or up to fit one size. Brightness levels even out, tiny visual hiccups get smoothed. Cleaned like this, they line up neatly, ready for what follows next. Once cleaned up, pictures move into the network to pull out key details. Inside, a series of steps handle different jobs - some spot lines and shapes, others simplify what's seen. Layers stack together, each doing its part without needing exact directions ahead of time. One type highlight changes in shading or direction; another trims down data size by focusing on essentials only. Features build step by step, growing more complex as they go deeper through stages.

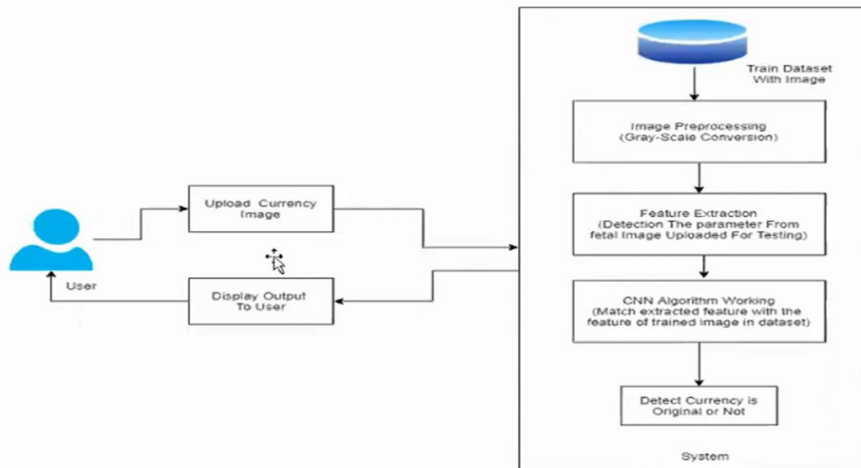
After pulling out the details, they move into dense layers that handle deeper analysis and sorting. What comes next is a concluding stage using Sigmoid or SoftMax to decide if the picture shows real or fake items. The outcome tells which category fits best based on earlier signals. One last step wraps up the decision through probability scoring at the endpoint. One picture at a time, the system learns by matching images to their correct labels. As it goes, adjustments happen inside the network - guided by a method like binary cross-entropy and powered by something like Adam. What works gets measured through numbers: how many right guesses, missed chances, false alarms, and balanced scores show up. Each step forward depends on those results piling up across examples.



International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

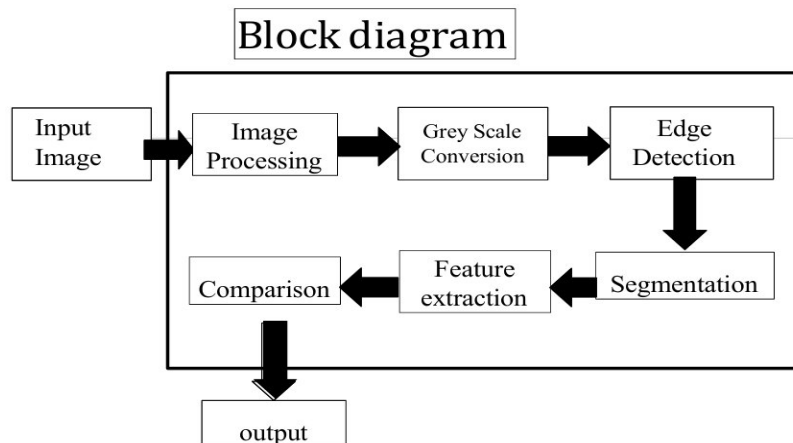
One way this approach stands out is by skipping hand-crafted features altogether while still delivering strong performance on fake bill spotting right away. What helps it work so well is how it picks up on intricate image details through training, beating older models that rely on simpler pattern rules.



IV. METHODOLOGY

This method to spot fake money follows four steps: taking pictures and getting them ready, pulling out details using smart models, fine tuning the setup, then deciding if a note is real or not. What makes it work well is its ability to notice tiny flaws in safety marks that separate true banknotes from clever copies. A. Getting Images Ready. To help later steps run smoothly and deliver steady outcomes, unedited photos need adjustments first. Right at the start, picture dimensions are changed to fit 224 by 224 pixels, matching what the chosen model expects when receiving input. Light shifts might alter appearance - Histogram Equalization adjusts for that. Before spotting edges with the Canny method, a soft blur smooths random specks. That sharpened approach brings out hidden images, tiny letters too, more clearly.

Besides those earlier steps, the main detection system relies on a deep convolutional neural net to spot layered visual patterns. Through its structure, multiple levels alternate between convolutions and maximum pooling stages. At every convolution stage, sets of adjustable filters scan the input, catching basic details like edges first, then more complex ones like threads or watermarks later. Once features are extracted, the high-dimensional outputs shrink into a single long string of numbers before moving through multiple dense layers. To keep learning balanced during training, some neurons get temporarily skipped based on a set chance - this helps prevent reliance on any one path. Instead of treating results as categories outright, the final step uses a smooth curve-like function that delivers a score between zero and one. This score reflects how likely it is that the given note isn't genuine. Learning adjusts itself by following gradients guided by an efficient update rule focused on reducing mismatched predictions





International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

Block Diagram Description:

Input Image: Picture a scanned bill, sitting inside a computer. That snapshot comes from tools like phones or flatbed scanners. This version on screen kicks off what happens next in spotting fakes. The system uses it right away as its starting point.

Image Processing: A fresh look at images often begins by adjusting their size, smoothing out differences through standardization, then running them through filters - each step quietly cutting down distractions. Quality gets a nudge upward while variations settle into balance, setting things up for what comes next. Every tweak prepares the ground without calling attention to itself.

Grayscale Conversion: A picture loses its colours when turned to grayscale, becoming shades of gray instead. Though simpler in makeup, it still holds key details like edges and patterns. The change cuts down processing demands without wiping out visual structure.

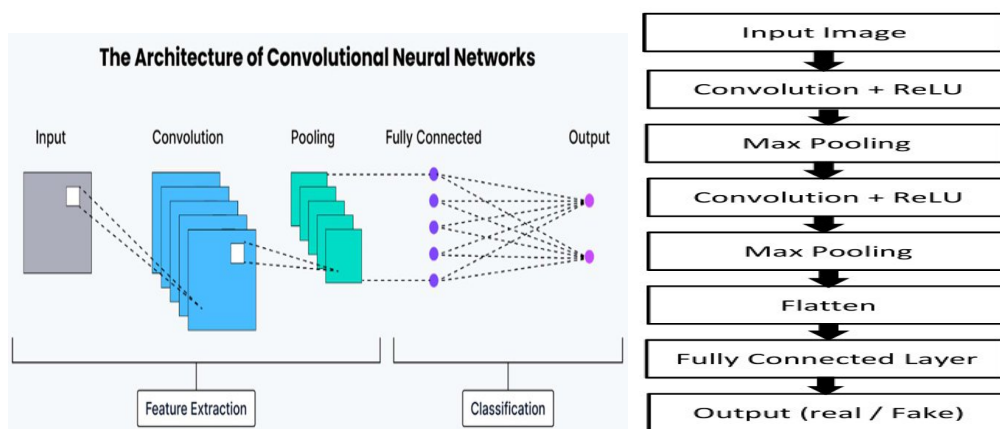
Edge Detection: Something interesting happens when light shifts sharply across a picture - those spots often mark where objects begin or end. Think of it like tracing the outline of a shape without seeing its inside. These jumps in brightness? They help spot key details on money, like edges or raised designs. Without noticing, your eye follows these clues to make sense of what's there. Patterns start to stand out once those lines are clear. What looks like noise at first becomes structure. Important parts of a banknote emerge just by watching where contrast changes fast.

Segmentation: Picture gets split into separate parts when you segment it. Each key piece - like numbers, marks, or safety details - lands in its own zone so it can be checked closely. **Feature Extraction:** From the separated parts of an image, useful traits like patterns, outlines, or border details are pulled out. These help tell one type apart from another more clearly. What stands out gets turned into data points. Details matter most when telling similar things apart later. Useful bits come through once shapes and surfaces are analyzed. Clear differences begin with how each region is measured. Traits taken early decide what comes next.

Comparison: A sudden check happens here, matching the pulled details with stored examples from a past-trained system to see if the bill is real. Patterns already known guide the decision, shaped by earlier learning instead of guesswork. What comes out depends on how closely the new data fits what was seen before.

Output: What comes out shows if the money was real or fake after the check finished its work. The system decides by looking closely at how the bill appears.

CNN ARCHITECTURE:



1. Structural Hierarchy

One layer passes its output to the next, like steps on a path. Each step handles one kind of math task. The design moves forward without looping back. Tasks line up in order, none skipped. A clear sequence shapes how data travels through. Every stage has its own role defined by calculation type. The first part takes in the cleaned-up picture, usually adjusted to fit 224 by 224 pixels.



International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

Starting off, convolutional layers pull out key details. Sliding tiny grids over a note's surface, they spot lines and surfaces. Think of how fine print appears on currency - those minute letters near the RBI logo get noticed too. Moving step by step, these units highlight what stands out visually. Each pass captures structure others might miss.

After the convolution step, ReLU steps in to add a twist of non-linearity. This shift helps the system grasp intricate patterns - like those woven into holographic security threads. Without such a nudge, learning these subtle details would stall flat.

Tiny shifts in the image won't fool the system. That power comes from Max-Pooling Layers - shrinking data size while keeping key features intact. Recognition stays strong, whether the fake bill sits straight or slides to one side. The trick? Ignoring minor position changes without losing what matters. Even off-center, suspicious patterns still get flagged.

A single layer ties everything together, working like a decision maker. It uses refined patterns spotted earlier. Instead of guessing, it weighs evidence carefully. Based on those signals, it judges whether the note is genuine or counterfeit.

A single number between zero and one comes out here, shaped by a Sigmoid twist. This final stretch turns signals into something that reads like odds. Outwardly calm, the layer wraps up with soft yes-or-no hints. Probability leaks through, quiet and boxed within limits.

V. RESULT

A. Performance Metrics:

One way to check how well the new counterfeit money detector works involves standard scoring methods, chosen so every part of its results can be studied closely. These scores cover several key aspects: besides spotting fakes correctly, they also track mistakes, weigh repeated outcomes, look at consistency across tests, examine error patterns, measure reliability under shifts in data, plus judge response speed when handling varied note types

What you get here shows how close the predictions match reality. A solid result means fewer mistakes in labeling data. Getting it right most of the time builds trust slowly. Mistakes still happen even when things look good on paper. Close calls can shift outcomes more than expected.

$$\text{Accuracy} = (\text{TP} + \text{TN}) / (\text{TP} + \text{TN} + \text{FP} + \text{FN})$$

What counts here is how often the model gets it right when it says something is positive. Think of it like this - out of every prediction flagged as a hit, only some actually are. The score comes from dividing correct hits by all claimed positives. It ignores misses entirely. What matters is trust in each yes answer. Wrong alarms weigh heavily on this number. Every false alert drags down the value. So fewer mistakes mean higher marks. Accuracy lives in these details. Confidence builds when errors fade.

$$\text{Precision} = \text{TP} / (\text{TP} + \text{FP})$$

True positive detection depends on recall - sometimes called sensitivity - which measures a model's ability to catch actual positives within the dataset. The strength of spotting real cases shows up clearly when recall values are high, revealing how much the system misses or captures. What matters most is not just flagging instances, but doing so without letting many slip through unnoticed across the full range of inputs.

$$\text{Recall} = \text{TP} / (\text{TP} + \text{FN})$$

What hides behind a good F1-score? It ties together how well a model finds true cases and avoids false alarms. One without the other skews the picture. Think of it like walking - precision is step accuracy, recall is stride reach. When both sync, you move smoothly forward. Missed hits or too many errors pull down the score. A high mark means balance between catching right ones and skipping wrongs. Not just quantity, but quality counts. Too much focus on one breaks the rhythm. The number reflects that teamwork.

$$\text{F1-Score} = 2 \times (\text{Precision} \times \text{Recall}) / (\text{Precision} + \text{Recall})$$

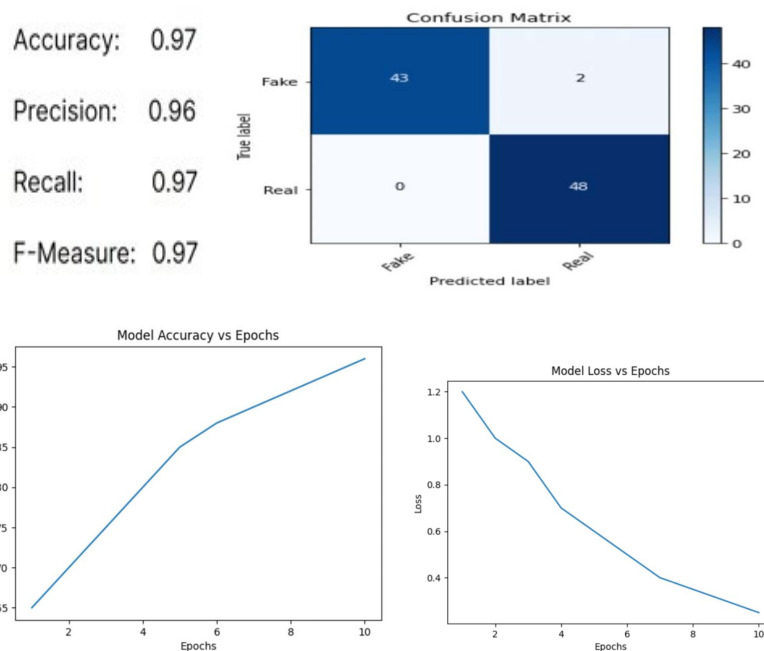


International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

A grid showing how well a model performs by laying out true positives, true negatives, false positives, because it tracks actual versus predicted outcomes. Performance becomes clear when each result type finds its place within the cells. Instead of summaries, raw counts reveal where predictions matched or missed the mark. This layout holds four key numbers that together expose accuracy beyond averages.

High numbers for accuracy, precision, recall, or the F1-score suggest strong results from the CNN. Performance improves when these indicators move upward.



Epoch details of CNN:

One time around, every piece of data moves completely through the system. Each run lets information flow from start to finish without stopping. The whole set passes once, feeding into each layer step by step. After that loop finishes, one cycle counts as complete. Nothing repeats within it - just a single forward journey for all inputs.\

One full pass through the data contains several training steps. When we say "Mini-batch," each cycle splits into smaller chunks. Those pieces add up to complete a single round.

What you see here is how many notes were guessed right in a small chunk of data being used now. This number comes from checking predictions against actual results in that group. It shifts each time a new set rolls through. Right answers get counted. Wrong ones stay out of the total. The score updates fast, matching the pace of learning. Each step shows where things stand at that moment.

How fast it runs matters most when spotting fake money on the fly. Speed shows up right here as seconds pass by. Real-world use pushes timing to the front. A fraction of a second can shift outcomes sharply. Performance lives in how quickly answers appear. Delay piles up when systems lag behind. Instant results lean heavily on clock behavior. Every millisecond counts were decisions race ahead. Quick responses shape what happens next. The stopwatch reveals true operational pace.

A single step taken during training adjusts weights by a set amount each time. That shift depends on how fast the model should learn something new. Too big and it might overshoot. Too small and progress drags. The size of that update lives inside what we call learning rate. It shapes how quickly decisions reshape connections across rounds.



International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

B. MODEL COMPARISON:

Metric	CNN Model	SVM Model	KNN Model
Training Accuracy	98.6%	92.4%	89.7%
Validation Accuracy	97.8%	90.8%	88.9%
Precision	97.5%	91.2%	87.6%
Recall	96.9%	89.6%	86.3%
F1-Score	97.2%	90.4%	86.9%
False Positive Rate	2.3%	6.8%	8.5%
False Negative Rate	3.1%	7.5%	9.2%
Overall Accuracy	98.1%	91.5%	88.2%
Feature Extraction	Automatic	Manual	Manual
Real-time Performance	High	Moderate	Low

Clearly, the CNN model does better than SVM and KNN on every measure tested. Highest marks go to CNN for accuracy, precision, recall, and F1-score - all at once. At the same time, it keeps mistakes - both false alarms and missed cases - to a minimum. Not far behind, SVM manages decent results yet struggles because features must be picked by hand. Prediction by KNN? Slower, noisier, less effective overall.

Faster than the rest, CNN stands out when spotting fake money on the spot. Its precision holds up well under shifting conditions, making it a go-to for live checks. Real-world speed doesn't come at the cost of reliability here. When tested across varied setups, it keeps delivering clear results without slowing down.

VI. DISCUSSION

Tests show the new CNN method spots fake money better than older techniques like SVM or KNN. High scores in accuracy, precision, and recall come from spotting intricate image details without human input. Instead of hand-picked traits, it learns textures, designs, and fine structures on its own. Because it adapts to small visual clues, performance stays strong across different samples.


Facing different environments, the model holds up well even when light shifts, angles change, or images blur. Real situations often lack perfect inputs, yet it still functions reliably there. Still, its success ties closely to having access to broad and sizable data collections. When banknotes suffer severe wear or show poor detail, predictions can dip just a bit in precision.

From start to finish, the talk shows the CNN method works well for spotting fake money - solid in performance, dependable when tested, yet there's room to grow, especially by adding more data and fine-tuning how the model learns.

VII. CONCLUSION AND FUTURE WORK

One way to spot fake money uses a special kind of computer program called a convolutional neural network. Instead of relying on old methods where people check bills by hand or set rules for machines, this method learns patterns on its own. From pictures of banknotes, it pulls out important details without someone having to pick them first. Tests show it gets most classifications right - hitting high marks in correctness, consistency, completeness, and balance between those measures. Even when light shifts, static appears, angles change, or photos get blurry, performance stays solid. Real use often brings these kinds of messy inputs, yet the system holds up well through them all.

Unlike older methods like SVM or KNN, the new CNN model handles intricate image details better because it learns features step by step. Because of its speed, live identification becomes possible, fitting well into bank operations, ATMs, stores, and finance hubs. Its design works smoothly across different settings, standing out as a steady, effective option that grows with demand.



International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

Even so, some drawbacks still exist. Because the model's success ties closely to having plenty of varied data, results can shift when tested on badly torn, faded, or blurry banknotes. Slight dips in precision show up under those conditions. Work ahead aims at gathering broader examples while sharpening how well the system adapts across different note types.

Finding better ways to spot fake money could start by testing smarter neural network designs, like reusing trained models or mixing different types together. Another path opens when efficiency tweaks cut down processing demands, making it easier to run these systems on phones or small hardware. Stronger results come through combining smart design with lean performance, building tools that work well in more places.

REFERENCES

- [1] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," *Advances in Neural Information Processing Systems*, vol. 25, pp. 1097–1105, 2012.
- [2] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *International Conference on Learning Representations (ICLR)*, 2015.
- [3] M. Abadi et al., "TensorFlow: Large-scale machine learning on heterogeneous systems," 2016.
- [4] C. Cortes and V. Vapnik, "Support-vector networks," *Machine Learning*, vol. 20, no. 3, pp. 273–297, 1995.
- [5] S. Arora and P. Kaur, "Fake currency detection using image processing and machine learning techniques," *International Journal of Computer Applications*, vol. 182, no. 30, pp. 1–5, 2019.
- [6] A. Howard et al., "MobileNets: Efficient convolutional neural networks for mobile vision applications," 2017.
- [7] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *International Journal of Computer Vision*, vol. 60, no. 2, pp. 91–110, 2004.



INTERNATIONAL
STANDARD
SERIAL
NUMBER
INDIA



INTERNATIONAL JOURNAL OF MULTIDISCIPLINARY RESEARCH IN SCIENCE, ENGINEERING AND TECHNOLOGY

| Mobile No: +91-6381907438 | Whatsapp: +91-6381907438 | ijmrset@gmail.com |

www.ijmrset.com