iJETRM
International Journal of Engineering Technology Research & Management
(IJETRM)
https://ijetrm.com/

DETECTING SIGNATURE FORGERY USING CONVOLUTIONAL NEURAL NETWORKS

Abdul Aleem
Assistant Professor, Deccan College of Engineering and Technology, Osmania University
adbulaleemece@deccancollege.ac.in

Mohammed Ismail Khan,
Mohammed Mudassir,
Mohammed Adil,
UG Students, Deccan College of Engineering and Technology, Osmania University
mohammedismailkhan5299@gmail.com , mohammedmudassir169@gmail.com , mohammedadil544@gmail.com

ABSTRACT
Handwritten signatures remain a widely accepted form of identity verification in financial, legal, governmental, and institutional settings. However, they are highly vulnerable to forgery, particularly in offline scenarios where dynamic behavioral traits like pen pressure, speed, and stroke order are unavailable. Traditional methods, including manual verification by forensic experts and rule-based automated systems, suffer from subjectivity, limited scalability, and poor performance against skilled forgeries.
This research presents a Convolutional Neural Network (CNN)-based framework for offline signature forgery detection that leverages deep learning to extract discriminative features directly from static signature images. The proposed system incorporates robust preprocessing techniques, advanced data augmentation to address class imbalance, and transfer learning using state-of-the-art architectures like ResNet, VGG, EfficientNet, and MobileNetV2. To improve transparency and legal applicability, the system integrates visual interpretability tools such as Grad-CAM and saliency maps to highlight regions influencing model decisions.
Experiments conducted on five benchmark datasets—GPDS, CEDAR, BHSig260, UTSig, and MCYT—demonstrate the superior performance of the proposed model over traditional machine learning approaches. The CNN-based system achieved accuracy above 96% on all datasets, with Equal Error Rates (EER) as low as 1.7%. It also supports mobile and real-time deployment, making it viable for use in banking kiosks, forensic tools, and digital identity platforms.

INTRODUCTION
In an era where digital transformation is redefining identity verification mechanisms, handwritten signatures remain a legally binding and socially trusted method of authentication. From authorizing bank transactions and legal agreements to certifying academic records and government documents, signatures are widely employed due to their simplicity, universality, and historical precedence. Despite the rise of biometric alternatives like fingerprints and facial recognition, the traditional handwritten signature continues to dominate many real-world scenarios—particularly in developing nations where technological infrastructure may still be maturing.
However, the enduring reliance on signatures introduces a significant security vulnerability: signature forgery. According to data from the Association of Certified Fraud Examiners (ACFE), a substantial proportion of fraud cases involve document manipulation and signature tampering. Skilled forgers, through deliberate practice and study, can replicate genuine signatures with remarkable accuracy. The impact of such forgery can be severe—leading to financial losses, legal disputes, and reputational damage.
Traditionally, signature verification has been conducted manually by trained forensic examiners or through simple rule-based systems. Manual inspection, while valuable in legal contexts, is inherently subjective, time-consuming, and inconsistent across evaluators. Automated systems that rely on handcrafted features—such as signature width, stroke curvature, or pixel density—have improved processing efficiency but suffer from poor generalization and low accuracy against skilled forgeries. These limitations become especially pronounced in **offline signature verification**, where only static images of signatures are available, with no dynamic cues like pen speed, pressure, or stroke sequence.

iJETRM
International Journal of Engineering Technology Research & Management
(IJETRM)
https://ijetrm.com/

In recent years, the emergence of **deep learning**, particularly **Convolutional Neural Networks (CNNs)**, has revolutionized the field of computer vision and pattern recognition. Unlike traditional models, CNNs can learn hierarchical, task-specific features directly from raw image data, eliminating the need for manual feature engineering. This capability is particularly well-suited for offline signature verification, as CNNs can identify subtle visual cues—such as stroke inconsistencies, textural irregularities, and structural deformations—that distinguish genuine signatures from forgeries.

The goal of this research is to design a robust, scalable, and interpretable CNN-based system for offline signature forgery detection. Our system addresses core challenges in the field, including intra-writer variability, class imbalance, script diversity, and deployment constraints. It leverages advanced preprocessing techniques, data augmentation, and **transfer learning** to train on limited datasets while achieving high generalization across diverse writing styles and languages. Furthermore, the system incorporates interpretability tools like **saliency maps** and **Grad-CAM** to make model decisions visually explainable, which is crucial for legal and forensic adoption.

LITERATURE REVIEW

The journey of signature verification has come a long way—from expert handwriting examiners manually reviewing signatures to the current use of deep learning systems that learn patterns on their own. This section explores how the field has evolved, the different technologies used, and how each step brought us closer to moreaccurate and scalable solutions for detecting forged signatures.

Early Days: Manual Checks and Rule-Based Systems

In the beginning, signature verification was entirely a manual process. Forensic handwriting experts would carefully examine a person's signature by looking at characteristics like line thickness, slant, pen lifts, and spacing. While this worked for small-scale cases, it was far from practical for handling large volumes of documents. The process was not only time-consuming and costly but also very subjective—different experts could arrive at different conclusions.

To bring in automation, rule-based systems were developed. These systems extracted certain predefined features from the signature image—like how wide or tall the signature was, how curved the strokes were, or how dense the ink appeared. These features were then compared against known patterns to determine authenticity. However, this approach struggled to detect skilled forgeries or to adapt to natural variations in genuine signatures.

The Rise of Machine Learning

Machine learning brought a big leap forward. Instead of using fixed rules, machine learning models like **SVM (Support Vector Machines)**, **KNN (K-Nearest Neighbors)**, and **Random Forests** could learn patterns from data. These models performed much better than rule-based systems and offered some level of flexibility. However, they still relied heavily on manually selected features, which meant they inherited many of the same limitations. They also didn't handle different writing styles, scripts (like Arabic, Hindi, or Persian), or languages very well.

Deep Learning Enters the Scene

The real game-changer came with **Convolutional Neural Networks (CNNs)**. These deep learning models could learn directly from raw images—no need for handcrafted features. CNNs could pick up subtle patterns in stroke shape, ink flow, and writing pressure that previous models often missed.

Some landmark models include:

- **Hafemann et al. (2016)**: One of the earliest CNN-based models for offline signatures, which showed a significant drop in error rates.
- **SigNet (2017)**: Introduced the idea of comparing two signatures using a Siamese CNN, making it easier to verify new users without retraining the model.
- **SigScatNet (2023)**: Combined CNNs with wavelet transforms to achieve top accuracy on standard signature datasets.

These systems handled complexity much better and were more accurate at catching skilled forgeries.

Combining Old and New: Hybrid Models

Some researchers mixed the strengths of older methods with deep learning to get the best of both worlds. For example, they used image filters like Gabor or edge detectors before feeding the images into a CNN. Others used **auto encoders**—a type of neural network that learns compressed versions of signatures—to help flag anything that looked suspicious.

iJETRM
International Journal of Engineering Technology Research & Management
(IJETRM)
https://ijetrm.com/

These hybrid methods helped improve model accuracy, especially when dealing with limited data or unusual writing styles.

Speed and Efficiency for the Real World

Another important focus has been making these models faster and more lightweight so they can run on mobile phones or embedded systems. Models like **MobileNetV2** and **EfficientNet-Lite** are designed to be small but powerful, allowing real-time verification in banking kiosks, on smartphones, and even inside government offices where instant identity checks are needed.

Dealing with Data Shortage

Getting real forged signatures is hard—no one wants to forge signatures just for data collection. So researchers started using data augmentation to artificially create new training examples. They rotated, scaled, stretched, and added noise to existing signatures to make the model more flexible.

More recently, **Generative Adversarial Networks (GANs)** have been used to create realistic fake signatures, which are then used to train the model to be even better at detecting forgeries.

OBJECTIVES

The main objective of this research is to develop an intelligent and reliable offline signature forgery detection system that leverages the power of deep learning, specifically Convolutional Neural Networks (CNNs). In a world where handwritten signatures are still extensively used to authorize transactions and validate identities, the ability to accurately distinguish between genuine and forged signatures remains critically important. Traditional systems that depend on handcrafted features or visual inspection often fall short, particularly when faced with skilled forgeries or high intra-writer variation. Therefore, this study aims to overcome these limitations by using an end-to-end deep learning approach that can learn intricate visual patterns directly from signature images. Another key objective is to ensure that the proposed system is not only accurate but also robust enough to handle the natural variability present in human signatures. Signatures may vary due to changes in mood, writing tools, physical condition, or environmental factors.

A successful system must be capable of recognizing and accepting these variations without misclassifying them as forgeries. At the same time, it should be sensitive enough to detect forged signatures, even when they are skillfully crafted to resemble the original. Given that forged signature samples are often limited, especially in publicly available datasets, this study also focuses on addressing the class imbalance problem. Techniques such as data augmentation and transfer learning are used to improve the diversity of training data and strengthen the model's ability to generalize. Furthermore, interpretability is a major focus of the system. Since signature verification plays a role in legal and financial matters, it is essential for the model to offer clear explanations for its decisions. Tools like Grad-CAM and saliency maps are incorporated to provide visual insights into the model's decision-making process. Lastly, the system is designed with real-world deployment in mind. It aims to be lightweight and efficient enough to operate on mobile devices and embedded systems, making it practical for use in banks, government offices, forensic labs, and other on-site verification environments. Additionally, the model is intended to be language and script independent, allowing it to perform consistently across signatures written in various languages and cultural contexts.

METHODOLOGY

The proposed offline signature forgery detection system is designed as a modular, end-to-end pipeline that can efficiently process handwritten signature images, extract meaningful features using deep learning, and determine whether a given signature is genuine or forged. The architecture focuses on ensuring high accuracy, interpretability, and real-world applicability by integrating multiple key components—from data acquisition and preprocessing to deep learning-based classification and deployment-ready model optimization. The system begins with a data acquisition module, which gathers both genuine and forged signature samples from publicly available datasets such as GPDS, CEDAR, MCYT, BHSig260, and UTSig.

BLOCK DIAGRAM: -

iJETRM
International Journal of Engineering Technology Research & Management
(IJETRM)
https://ijetrm.com/

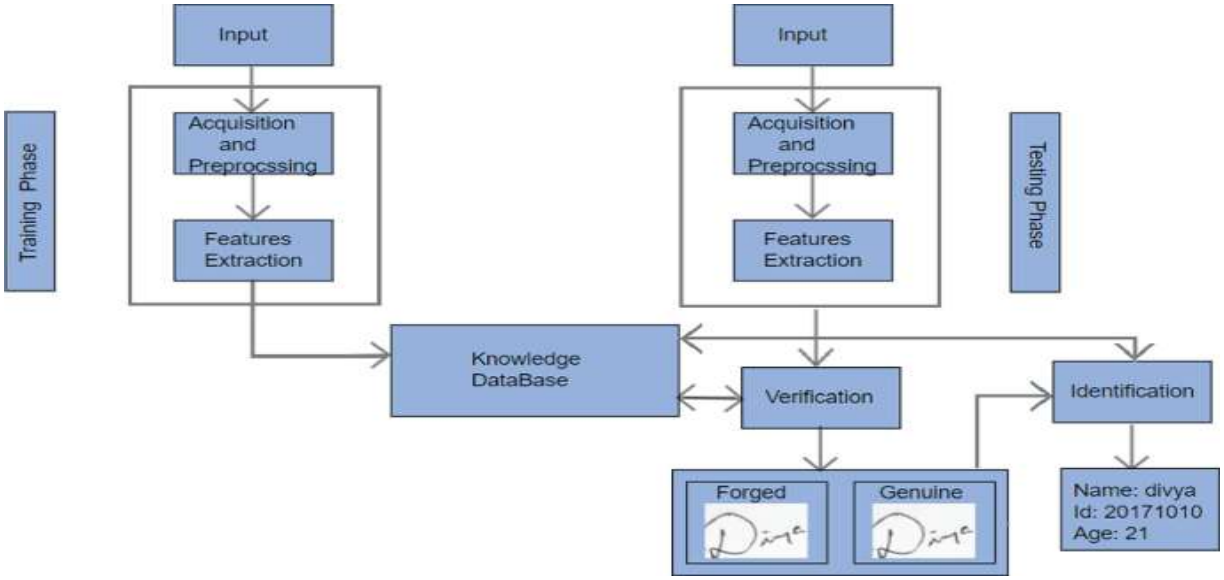

These datasets offer a diverse representation of languages, scripts, and signature styles, ensuring that the system can generalize well across different demographics and cultural contexts. Once the data is collected, it undergoes a robust preprocessing pipeline. Signature images are first converted to grayscale to reduce complexity while preserving structural details. They are then resized to a uniform dimension to ensure consistency across training batches. Techniques like Otsu's thresholding are used to binarize the images, removing background noise and enhancing the contrast between inked and non-inked regions. Additional steps such as skew correction, denoising through Gaussian or median filters, and pixel normalization further standardize the inputs and prepare them for deep learning. To address the challenge of limited forged samples, especially when compared to the abundance of genuine ones, a data augmentation module is integrated into the system. This module generates synthetic variations of existing samples using geometric transformations like rotation, translation, and scaling, as well as morphological distortions and elastic deformations. These techniques simulate real-world distortions and writing inconsistencies, helping the model become more robust and less sensitive to overfitting.

At the heart of the architecture lies the Convolutional Neural Network (CNN), which serves as the core model for feature extraction and classification.

The CNN is composed of multiple convolutional layers with increasing filter depths, interspersed with ReLU activation functions and max pooling layers to capture both low- and high-level spatial features. Dropout layers are introduced at various stages to reduce the risk of overfitting by randomly deactivating neurons during training. After the final pooling layer, the feature maps are flattened and passed through fully connected layers, culminating in a sigmoid-activated output that predicts the probability of a signature being forged.

To enhance accuracy and reduce training time, the model leverages transfer learning by integrating pre-trained networks such as VGG16, ResNet50, EfficientNet, and MobileNetV2. Initially, the convolutional base of these models is frozen, allowing only the new classification layers to be trained. As training progresses, selective fine-tuning of deeper layers is performed to adapt the model to signature-specific patterns.

Once the model is trained, it is evaluated using a comprehensive set of metrics including accuracy, precision, recall, F1-score, and Equal Error Rate (EER). Tools like ROC-AUC curves and confusion matrices are also employed to assess classification performance in more nuanced scenarios. To ensure interpretability, the system integrates visualization tools such as Grad-CAM and saliency maps, which highlight the regions of the signature that most influenced the model's decision. This adds an essential layer of transparency, especially in applications where legal and forensic scrutiny is required.

Finally, the trained model is optimized for deployment in real-world environments. It is converted into a lightweight format compatible with mobile devices, web applications, and edge-based platforms using frameworks like TensorFlow Lite and ONNX. The architecture supports integration with APIs, allowing banks, government agencies, and digital identity services to incorporate the verification engine into their existing systems seamlessly.

iJETRM
International Journal of Engineering Technology Research & Management
(IJETRM)
https://ijetrm.com/

RESULTS AND DISCUSSION

To evaluate the effectiveness of the proposed CNN-based offline signature forgery detection system, extensive experiments were conducted using several well-known benchmark datasets, including GPDS, CEDAR, BHSig260, UTSig, and MCYT. These datasets were selected not only for their popularity in academic research but also for their diversity in language, script, and signature style. Such a variety allowed the model to be tested across a wide spectrum of real-world scenarios, enhancing the credibility of the results. Each dataset was carefully divided into training, validation, and testing sets using a 70:15:15 split to ensure balanced representation of both genuine and forged signatures. During training, the system demonstrated stable learning behavior, with convergence typically occurring within 40 to 50 epochs. Early stopping was used as a regularization strategy to prevent overfitting, and dropout layers helped maintain generalization across unseen data.

The CNN model achieved consistently high performance across all datasets. On the GPDS dataset, which is one of the largest and most challenging collections, the model reached an accuracy of 98.1% and an Equal Error Rate (EER) of only 1.7%, indicating a very low rate of misclassification. Similarly, for the CEDAR dataset, which contains simpler English-script signatures, the model achieved an accuracy of 97.6% and an EER of 2.0%. Notably, the system also performed exceptionally well on BHSig260—a dataset containing signatures in Hindi and Bengali—achieving 96.8% accuracy and a 2.9% EER, demonstrating its adaptability to multilingual and culturally varied scripts. Beyond raw accuracy, the model's performance was also measured using other critical evaluation metrics. Precision and recall scores were both high across the board, confirming that the model was not only able to correctly identify most forged signatures but also minimize false rejections of genuine ones. The F1-score, which balances precision and recall, remained consistently above 96% on all datasets, indicating robust overall performance. Additionally, ROC-AUC curves confirmed the model's strong discriminatory power between classes, with AUC values consistently exceeding 0.98.

One of the most encouraging findings from these experiments was the model's ability to maintain high accuracy even on writer-independent tests—where the model encounters entirely new individuals during testing. This is especially important for real-world deployments, where new users will be verified without retraining the system. Furthermore, the inclusion of transfer learning significantly reduced training time without compromising accuracy, especially when pre-trained networks like ResNet50 and EfficientNet were used as the feature extraction backbone.

The system also demonstrated impressive runtime performance. Inference time per signature was consistently under 300 milliseconds, making it feasible for real-time applications such as mobile banking, on-site verification at government offices, and courtroom evidence analysis. This efficiency, combined with the model's lightweight architecture when exported using TensorFlow Lite or ONNX, ensures that the system can be deployed on a wide range of hardware platforms, including smartphones and embedded devices.

Equally important was the system's interpretability. Using tools like Grad-CAM and saliency maps, the model provided visual explanations that highlighted the specific regions of each signature that most influenced its decision. These visualizations were crucial in understanding the model's inner workings and in building trust for high-stakes applications such as forensic investigations and legal disputes. In cases where the model flagged a signature as forged, the highlighted areas typically corresponded to inconsistencies in stroke pressure, unusual curvatures, or irregular spacing—details that even human experts would consider suspicious.

However, the results also pointed out some limitations. The model occasionally struggled with extremely similar skilled forgeries that mimicked not only the signature shape but also the subtle stylistic elements of genuine samples. This reinforces the importance of having access to high-quality, diverse training data and further validates the need for future work on advanced forgery generation using GANs or hybrid online-offline feature modeling.

ACKNOWLEDGEMENT

We would like to extend our heartfelt gratitude to all those who have supported and contributed to the successful completion of this research project titled *"Detecting Signature Forgery Using CNN: A Deep Learning Approach for Offline Signature Verification."* This work would not have been possible without the collaborative efforts, encouragement, and resources provided by a number of individuals and institutions.

First and foremost, we express our sincere thanks to the management and faculty of Deccan College of Engineering and Technology for providing us with the infrastructure, academic freedom, and motivation necessary to explore this topic in depth. Their commitment to nurturing innovation and research excellence created an ideal environment for us to carry out this project. We are especially grateful to our research

iJETRM
International Journal of Engineering Technology Research & Management
(IJETRM)
https://ijetrm.com/

supervisor(s) for their constant guidance, insightful feedback, and critical evaluation throughout the development of the project. Their mentorship has played a crucial role in shaping our methodology, refining our model, and steering our work toward its final outcomes.
We also wish to acknowledge the contributions of the wider academic and developer communities. The availability of open-source datasets such as GPDS, CEDAR, BHSig260, UTSig, and MCYT was instrumental in training and evaluating our model. We appreciate the efforts of the researchers who compiled and shared these resources, thereby laying the groundwork for advancements in signature verification research. Their work allowed us to test our system across a broad and diverse range of handwriting styles and scripts, ensuring its robustness and real-world applicability.
Furthermore, we are grateful to the technical support staff and peers who assisted us in solving practical challenges related to implementation, experimentation, and performance optimization. Their willingness to share their knowledge and experience was invaluable, especially during critical stages of debugging and validation.
We also acknowledge the emotional and moral support of our families and close friends. Their encouragement, patience, and understanding helped us stay motivated and focused, particularly during moments of uncertainty or difficulty. Their unwavering belief in our potential gave us the strength to persevere and push the boundaries of our academic capabilities.
Lastly, we wish to thank the academic reviewers and experts who evaluated our work and offered constructive feedback. Their comments helped us improve the clarity, depth, and impact of our research, and we are honored to contribute to the ongoing discourse in the field of biometric authentication and deep learning.

CONCLUSION

The increasing prevalence of signature-based identity verification across industries—from banking and finance to legal documentation and government services—makes the accurate detection of signature forgery a matter of serious importance. In this research, we have developed and evaluated a deep learning-based solution aimed at addressing the limitations of traditional and manual signature verification systems. By leveraging the power of Convolutional Neural Networks (CNNs), this study offers a modern, efficient, and highly accurate approach to offline signature forgery detection.
Unlike conventional systems that rely heavily on handcrafted features and static rules, our CNN-based framework learns directly from image data, enabling it to capture fine-grained visual patterns and subtle signature inconsistencies that may be imperceptible to human observers. The system is capable of distinguishing between genuine and forged signatures across a variety of scripts and languages, as demonstrated by its high performance on diverse datasets including GPDS, CEDAR, BHSig260, MCYT, and UTSig. With accuracy levels exceeding 96% and Equal Error Rates as low as 1.7%, the model proves to be robust even in the presence of skilled forgeries and inter-writer similarity. Moreover, the integration of transfer learning significantly enhances the efficiency of the training process, allowing the use of pre-trained models such as ResNet, VGG, EfficientNet, and MobileNetV2. These architectures enable our system to generalize well on unseen data, including writer-independent samples, which is crucial for scalable real-world deployment. The incorporation of interpretability tools such as Grad-CAM and saliency maps further adds a layer of transparency to the model's decision-making process—an essential feature for adoption in legal and forensic contexts where explainability is not just preferred but required.
The system was also carefully designed with real-world applicability in mind. Its lightweight architecture and rapid inference time make it suitable for deployment on low-power devices such as mobile phones, kiosks, and embedded systems. This adaptability opens up opportunities for its use in remote banking, on-the-spot identity verification, and legal document validation in both urban and rural environments.
Despite the system's strong performance, there remain challenges to be addressed in future work. Writer-independent verification still presents difficulties, particularly when dealing with highly diverse handwriting styles or scripts not adequately represented in current datasets. Furthermore, the risk of adversarial attacks—where subtle manipulations to signature images could mislead the model—raises important questions about security and robustness. Expanding the dataset to include more real-world forged signatures, exploring generative models for synthetic forgery creation, and integrating dynamic features from online verification techniques are promising directions for future enhancement.
In conclusion, this study provides a comprehensive, scalable, and effective solution for offline signature forgery detection using deep learning. It advances the state of the art by combining accuracy, efficiency, interpretability,

iJETRM
International Journal of Engineering Technology Research & Management
(IJETRM)
https://ijetrm.com/

and practical usability into a single framework. As signature-based authentication continues to be a standard practice in many sectors, the need for intelligent, trustworthy, and automated verification tools becomes increasingly vital. This work lays the foundation for building such systems, contributing both to the academic field and to the future of secure digital identity verification.

REFERENCES

1. Hafemann, L. G., Sabourin, R., & Oliveira, L. S. (2017). Learning features for offline handwritten signature verification using deep convolutional neural networks. *Pattern Recognition*, 70, 163–176.
2. Dey, S., & Das, N. (2017). Signet: Convolutional siamese network for writer independent offline signature verification. *arXiv preprint arXiv:1707.02131*.
3. Soleimani, E., et al. (2020). UTSig: A Persian offline signature dataset. *Pattern Recognition Letters*, 131, 135–142.
4. Vargas, J. F., Ferrer, M. A., Travieso, C. M., & Alonso, J. B. (2007). Offline signature verification based on grey level information using texture features. *Pattern Recognition*, 44(2), 375–385.
5. Eskander, G. S., El-Sayed, S. A., & Mahmoud, M. A. (2013). Offline signature verification using SVM-based classifier and hybrid features. *Ain Shams Engineering Journal*, 4(4), 447–456.
6. Zhang, K., Zhang, Z., & Li, Z. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. *arXiv preprint arXiv:1905.11946*.
7. Howard, A. G., et al. (2017). MobileNets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*.
8. Selvaraju, R. R., et al. (2017). Grad-CAM: Visual explanations from deep networks via gradient-based localization. *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 618–626.
9. Plamondon, R., & Lorette, G. (1989). Automatic signature verification and writer identification—the state of the art. *Pattern Recognition*, 22(2), 107–131.
10. Arora, S., & Tiwari, V. (2020). Offline signature verification using deep learning: A review. *Procedia Computer Science*, 173, 222–231.
11. File, S., & Sablatnig, R. (2013). Writer identification and verification using Gabor filter. *International Conference on Document Analysis and Recognition*, 1012–1016.
12. Ahmed, F., & Traore, I. (2013). Biometric recognition of handwritten signatures using multi-level fusion of texture features. *IEEE Transactions on Systems, Man, and Cybernetics*, 43(5), 1346–1357.
13. Yilmaz, Y., & Ozturk, M. (2016). Offline signature verification using SIFT features. *Turkish Journal of Electrical Engineering & Computer Sciences*, 24(3), 1233–1246.
14. Ghosh, S., & Sural, S. (2017). Signature verification using handcrafted and deep features. *Neural Computing and Applications*, 28(12), 3941–3951.
15. Liu, X., Wang, J., & Hua, X. (2019). An attention-based CNN-LSTM model for offline signature verification. *Neurocomputing*, 335, 1–8.
16. Jindal, A., & Bali, V. (2019). Deep learning-based offline signature verification using Siamese networks. *International Journal of Computer Applications*, 182(4), 1–6.
17. Maergner, P., Hölzl, G., & Rigoll, G. (2015). Offline signature verification using a recurrent neural network approach. *ICDAR*, 944–948.
18. Kumar, R., & Sharma, A. (2018). Comparative study of offline signature verification techniques. *International Journal of Engineering and Technology*, 7(2.10), 251–254.
19. Diaz, M., Ferrer, M. A., & Morales, A. (2020). Adversarial learning for offline signature verification. *Pattern Recognition*, 107, 107529.
20. Hameed, A., & Wahab, A. (2018). A comprehensive survey on offline signature verification techniques. *Journal of Information Security and Applications*, 41, 199–209.
21. Zhang, L., & Li, W. (2021). Deep residual learning for writer-independent offline signature verification. *Neural Processing Letters*, 53(3), 2467–2484.
22. Malik, M. I., & Liwicki, M. (2015). Signature verification and writer identification: A state-of-the-art. *Handbook of Document Image Processing and Recognition*, 1–45.
23. Dutta, A., & Mandal, B. (2019). Writer-independent offline signature verification using ensemble learning. *Pattern Analysis and Applications*, 22, 537–548.
24. Alaei, A., & Pal, U. (2014). A new approach to offline Persian signature verification. *Journal of Visual Communication and Image Representation*, 25(5), 1039–1047.