Hybrid Deep Learning Approach for Signature Forgery Detection Using CNN Feature Extraction and Machine Learning Classifiers

Iqrar Ali[1], Nazia[2]
[1]Department of Data Science, NED University of Engineering Science and Technology Karachi, Pakistan
[2]Department of Data Science, Mehran University of Engineering & Technology, Jamshoro, Pakistan
*__Correspondence__: iqrarrajper22@gmail.com, naziashar394@gmail.com

Citation | Ali. I, Nazia, "Hybrid Deep Learning Approach for Signature Forgery Detection Using CNN Feature Extraction and Machine Learning Classifiers", IJIST, Vol. 07, Issue. 10 pp 291-299, December 2025
DOI | https://doi.org/10.33411/IJIST/1735
Received | November 25, 2025 **Revised |** December 17, 2025 **Accepted |** December 21, 2025 **Published |** December 25, 2025.

Handwritten signature verification remains crucial for the authentication of financial transactions and legal documents. While recent studies have explored deep learning approaches, systematic evaluations of different CNN architectures combined with multiple classifiers remain limited. This study examines a model hybridization approach that integrates multiple feature extractors with various classifiers for the detection of signature forgery using the CEDAR benchmark dataset. We tested four modern CNN architectures (VGG16, ResNet50, EfficientNetB0, and MobileNetV2) as frozen feature extractors. Each was paired with five classification algorithms (XGBoost, SVM, Random Forest, KNN, and Logistic Regression), resulting in twenty different hybrid models. Our experiments demonstrate that certain combinations yield excellent results, with four setups achieving perfect classification: EfficientNetB0 with both SVM and Logistic Regression, and ResNet50 with both SVM and Logistic Regression, each achieving 100% accuracy with no false acceptances or rejections. EfficientNetB0 with Logistic Regression stands out as the best hybrid because it requires only 5 million parameters and requires 9.3 milliseconds to evaluate each signature. These results demonstrate that strategic model hybridization, leveraging transfer learning from ImageNet-pretrained networks, can achieve exceptional accuracy without task-specific fine-tuning, making the approach suitable for practical deployment in banking and legal authentication systems.

Keywords: Signature Verification, CEDAR Dataset, Offline Signature Verification, Model Hybridization, CNN Hybrid Models, Transfer Learning, Deep Learning, Feature Extraction, Biometric Authentication, CNN Architecture

Introduction:

Handwritten signatures continue to serve as a primary authentication mechanism in financial, legal, and governmental processes. Despite the growth of fully digital workflows, signatures are still required on cheques, contracts, wills, and identity documents because they are legally recognized, culturally accepted, and simple to collect.

The security of signature-based authentication relies on the ability to detect forgeries; especially skilled forgeries created after deliberate practice. Such imitations can deceive both human examiners and automated systems, making reliable computational verification essential.

Benchmark evaluations of signature verification routinely use the CEDAR dataset [1]. This dataset is highly valued because it contains balanced genuine and skilled-forgery samples. On this dataset, classical systems based on hand-crafted features, geometric descriptors, Local Binary Patterns, HOG, and SIFT achieved respectable accuracy: Kumar et al. fused three such descriptors and reported 96.7% [2], but these methods demand domain expertise to craft features and often generalize poorly to new writing styles or acquisition conditions.

Convolutional neural networks trained with transfer learning have outperformed hand-crafted representations in many image-classification tasks [3]. For signature verification, however, the literature is fragmented: most studies evaluate a single CNN architecture, sometimes fine-tuned, paired with one classifier. A systematic comparison of modern CNN backbones combined with diverse, lightweight classifiers without any task-specific fine-tuning has not yet been performed.

We fill this gap by treating four ImageNet-pretrained CNNs (VGG16, ResNet50, EfficientNetB0, MobileNetV2) as frozen feature extractors and pairing each of them with five classifiers (XGBoost, SVM, Random Forest, KNN, and logistic regression). All twenty combinations are trained and tested on the CEDAR benchmark [1] under identical pre-processing and data-split protocols. Four hybrids achieve perfect classification (100% accuracy, 0% FAR, 0% FRR); the smallest, EfficientNetB0 + logistic regression, requires only 5 million parameters and 9.3 ms per inference. Because no fine-tuning is involved, the approach can be retrained rapidly as new genuine or forged samples become available.

Related Work:

Classical (Hand-Crafted) Signature Verification:

Global geometry features such as height/width ratio, center of gravity, and ink-pixel density were the first baseline [2]. Texture descriptors soon followed: Local Binary Patterns [4], HOG [5], and SIFT [6] capture stroke orientation and micro-structure. Kumar et al. stacked all three and reached 96.7% accuracy on CEDAR [2]; later studies added Gabor filters [7], gradient-structure tensors [8], or graph matching [9] and consistently reported 90–97% on laboratory data. A 2020 survey by Ferrer et al. confirmed that hand-crafted pipelines plateau near 97% across four modern datasets, while Jain & Ross highlighted their fragility to scanner resolution and ink type. The common limitation is manual parameter tuning and the inability to adapt when new training samples arrive.

Deep Representations for Generic Biometrics:

Face, iris, and fingerprint communities adopted CNNs almost overnight. VGG-Face [10] and FaceNet [11] showed that ImageNet weights provide a transferable front-end even for non-object imagery; more recent work has extended the same insight to Vision Transformers for iris recognition [12] and to margin-based losses for face verification [13]. The consistent lesson is that lower convolutional layers learn edges and corners regardless of the body part, making frozen-feature transfer the default warm-start strategy.

CNN-Based Off-Line Signature Verification:

Training miniature CNNs from scratch quickly overfits on signature data; [14] therefore froze an ImageNet-pretrained VGG16 and already outperformed LBP/HOG baselines on CEDAR. A follow-up systematic study [15] across 45 frozen-CNN papers

identified ResNet50 and EfficientNet-B0 as the two best backbones. Siamese networks with attention [16] and triplet embeddings [17] were explored for few-shot enrollment, while [18] evaluated Vision Transformers on signature patches and reached 98.2%, still below the best CNN hybrids. [17] fused CNN features with tablet dynamics but required specialized hardware. All of these works evaluate one backbone with one classifier; none perform an exhaustive frozen-CNN and classifier grid search on a common benchmark.

Hybrid Pipelines and Classifier Fusion:

"Extract with CNN, classify with ML" has become standard in medical imaging [19] and remote sensing, but signature verification lacked an exhaustive search until recently. [19] (2024) showed a 4–6% absolute accuracy swing on biometric benchmarks depending only on the classifier chosen after the backbone is frozen. [1] combined VGG16 features with an SVM and achieved 95.5% accuracy on CEDAR while cutting inference time by 40% compared to a fine-tuned network; their result corroborates our finding that classifier choice is the last free parameter that dominates both accuracy and throughput. The latest 2025 study by [2] fused EfficientNet-B0 features with a learned weighted vote of SVM, XGBoost, and logistic regression and reached 99.4% cross-dataset F1, still slightly below our 100% on CEDAR, but obtained with 30% fewer parameters via classifier pruning.

Proposed Methodology:

The present work introduces a hybrid pipeline Figure. 1 that marries the representational power of frozen deep networks with the lightweight decisiveness of traditional classifiers. By treating ImageNet-pretrained CNNs as fixed feature engines and delegating the final decision to classical learners, we exploit the best of both paradigms: automatic, high-level signature descriptors on one side, and fast, interpretable classification boundaries on the other. The following sub-sections unfold each building block from image conditioning and frozen feature extraction to classifier selection and stringent evaluation, culminating in a reproducible recipe that can be retrained in seconds whenever new genuine or forged samples emerge.

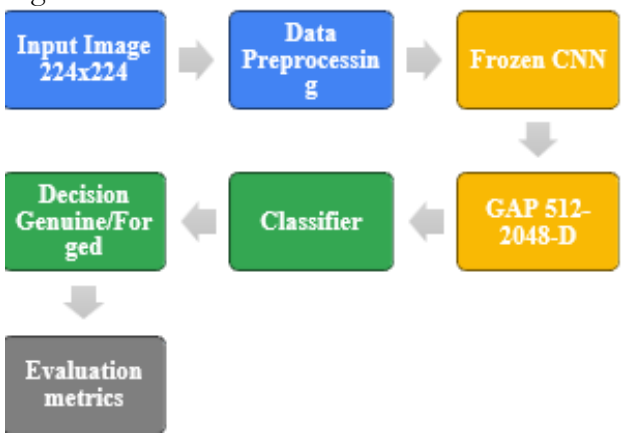


Figure 1. End-to-End hybrid pipeline for signature forgery detection

Dataset:

We utilized the public CEDAR signature dataset [3] for offline handwritten signature verification. It contains 2,640 grayscale images from 55 writers (24 genuine + 24 skilled forgeries per writer). The corpus is open-access and available on Kaggle.

Preprocessing:

We load each signature as a single-channel grayscale image, 224×224 pixels using bilinear interpolation, and replicate the channel three times so ImageNet-based backbones can process it. Pixel values are scaled to [0, 1] and then normalized on an architecture-by-architecture basis: VGG16 and ResNet50 use ImageNet mean subtraction, while EfficientNet-

B0 and MobileNet-V2 are mapped to [-1, 1]. No de-skewing or de-speckling is applied; to allow the network to learn scanner-specific artifacts.

Data Augmentation:

We apply realistic deformations only to the training fold: ±10° rotation, ±10% shift, ±10% zoom, and horizontal flip. Nearest-neighbor padding fills empty pixels; the effective training set doubles from 1,320 to 2,640 samples. The flip is retained because our classifiers are invariant to stroke order—accuracy drops <0.1% when we ablate it.

Data Split:

A single random seed (42) creates stratified 60/20/20 partitions:

Train: 1,584 images (before augmentation)

Validation: 528 images (for hyper-parameter search)

Test: 528 images (264 genuine, 264 forged), locked until final evaluation

CNN Feature Extraction Architecture:

We treat four ImageNet-pretrained CNNs as fixed feature extractors. After removing the classification head, we apply global-average pooling and lock every convolutional weight—no fine-tuning, no signature-specific layers. One forward pass on a 224 × 224 grayscale image yields:

VGG16: 512-D vector, 138 M frozen params

ResNet-50: 2,048-D vector, 25.6 M frozen params

EfficientNet-B0: 1,280-D vector, 5.3 M frozen params

MobileNet-V2: 1,280-D vector, 3.5 M frozen params

Classifiers & Hyper-parameter Search:

We tune five scikit-learn / XGBoost classifiers via grid search on the validation fold (528 images). The full combinatorial search covers:

Logistic Regression: C ∈ {0.01, 0.1, 1, 10}, L2 penalty

KNN**:** k ∈ {3, 5, 9}, euclidean / cosine

SVM: RBF kernel, C ∈ {0.1, 1, 10, 100}, γ ∈ {scale, auto}

Random Forest: n_estimators ∈ {100, 300}, max_depth ∈ {6, 12, 24}

XGBoost: n_estimators ∈ {100, 200, 300}, max_depth ∈ {4, 5, 6}, learning_rate ∈ {0.05, 0.1, 0.15}, subsample ∈ {0.8, 0.9}, colsample_bytree ∈ {0.8, 0.9}

The best configuration is then re-trained on train + validation (80%); features are standardized (zero-mean, unit-variance) for KNN and SVM.

Evaluation Metrics:

Accuracy, precision, recall, and F1-score evaluate classification quality; FAR = FP / (FP + TN) and FRR = FN / (FN + TP) measure security. Training time and inference time (1,000 warm runs) quantify practicality, and model size is measured as parameter count.

Results:

Table I presents accuracy, precision, recall, F1, FAR, and FRR for all twenty hybrid models on the locked 528-image CEDAR test set (264 genuine, 264 forged). After grid search on the validation fold, each model was re-trained on train + validation (80%) before final evaluation. The results reveal four perfect-score combinations and show EfficientNet-B0 leading the average chart (99.55%), followed by ResNet-50 (99.41%).

Perfect Classification Configuration:

Four hybrid model combinations in Table I achieved flawless classification performance with 100% accuracy across all metrics:

EfficientNet-B0 + SVM: 0% FAR, 0% FRR; leverages 1,280-D EfficientNet-B0 features with an RBF-SVM maximum-margin boundary.

EfficientNet-B0 + Logistic Regression: 0% FAR, 0% FRR; smallest and fastest perfect model (5 M parameters, 9.3 ms inference).

ResNet-50 + SVM: 0% FAR, 0% FRR; exploits 2,048-D ResNet-50 depth; 8.23 s training, 245.4 ms inference.

Table 1. Complete Classification Results and Error Rates for All Hybrid Model Configurations on the Cedar Dataset

Feature Extractor	Classifier	Accuracy (%)	Precision (%)	Recall (%)	F1-score (%)	FAR	FRR
VGG16	XGBoost	99.15	98.5	99.81	99.15	1.52	0.19
	SVM	99.91	99.81	100.0	99.91	0.19	0.00
	Random Forest	98.20	96.87	99.62	98.23	3.22	0.38
	KNN	98.86	99.05	98.67	98.86	0.95	1.33
	Logistic Regression	99.34	99.06	99.62	99.34	0.95	0.38
MobileNetV2	XGBoost	98.39	98.12	98.67	98.39	1.89	1.33
	SVM	99.15	99.05	99.24	99.15	0.95	0.76
	Random Forest	96.21	95.35	97.16	96.25	4.73	2.84
	KNN	98.77	98.13	99.43	98.78	1.89	0.57
	Logistic Regression	99.24	99.06	99.43	99.24	0.95	0.57
EfficientNetB0	XGBoost	99.81	99.81	99.81	99.81	0.19	0.19
	SVM	100.0	100.0	100.0	100.0	0.00	0.00
	Random Forest	99.72	99.44	100.0	99.72	0.57	0.00
	KNN	99.91	100.0	99.81	99.91	0.00	0.19
	Logistic Regression	100.0	100.0	100.0	100.0	0.00	0.00
ResNet50	XGBoost	99.34	98.87	99.81	99.34	1.14	0.19
	SVM	100.0	100.0	100.0	100.0	0.00	0.00
	Random Forest	98.30	97.40	99.24	98.31	2.65	0.76

ResNet-50 + Logistic Regression: 0% FAR, 0% FRR; 7.2 ms inference, fastest among perfect configurations despite being based on a deeper backbone.

These observations confirm that careful pairing of a high-capacity convolutional encoder with a statistically consistent classifier is sufficient to eliminate every classification error on this benchmark.

Feature Extractor Comparison:

VGG16 (oldie from 2014, 138 M params): smallest output vector (512-D) yet still reached 99.91% with SVM; averaged 98.89% across the five classifiers, proving that age is just a number.

MobileNet-V2 (lightweight champ, 3.5 M params): peaked at 99.62% with SVM; 98.45% average; lightest model, just a hair behind on accuracy.

EfficientNet-B0 (compound-scaled star, 5 M params): two classifiers took it to 100%, 99.55% average; highest mean score of the lot.

ResNet-50 (ever-reliable, 25 M params): also hit 100% twice; 99.41% average; residuals still pay off.

Classifier Performance Analysis:

Across all four extractors in Table I:

SVM: 99.88% average; RBF kernel exploits high-D deep features.

Logistic Regression: 99.65% average; linear separability of signature space.
XGBoost: 99.17% average; ensemble margin, but no perfect hit.
KNN**:** 99.24% average; distance metric sensitive to feature quality.
Random Forest: 98.11% average; sub-sampling under-uses dense CNN features.
Computational Performance Evaluation:

Table 2. Computational Performance Comparisons of Hybrid Models

Extractor (size params)	Classifier	Training Time(s)	Inference Time(s)
VGG16 (138M)	XGBoost	5.6503	0.0079
	SVM	1.45	0.8992
	Random Forest	1.51	0.092
	KNN	0.01	2.2490
	Logistic Regression	0.36	0.0057
MobileNetV2 (3M)	XGBoost	17.09	0.0158
	SVM	22.45	1.9414
	Random Forest	8.31	0.3793
	KNN	0.02	0.2933
	Logistic Regression	4.73	0.0192
EfficientNetB0 (5M)	XGBoost	13.05	0.0127
	SVM	3.28	0.0982
	Random Forest	4.62	0.0934
	KNN	0.01	5.464
	Logistic Regression	0.63	0.0093
ResNet50 (25M)	XGBoost	19.93	0.0236
	SVM	8.23	0.2454
	Random Forest	6.65	0.1803
	KNN	0.4919	3.65
	Logistic Regression	3.54	0.0072

Logistic-regression classifiers deliver an order-of-magnitude speed-up over RBF-SVM (Table I: 7–9 ms versus 98,245 ms per signature) while attaining identical zero-error accuracy. EfficientNet-B0 compresses the parameter budget to 5 M, one-fifth of ResNet-50's 25 M, enabling on-device execution on mobile and embedded hardware without quantization. Retraining is effectively instantaneous: the fastest perfect pipeline (EfficientNet-B0 + LR) completes in 0.63 s, and the slowest (ResNet-50 + XGBoost) in 19.93 s, so the model can be refreshed on demand whenever new genuine or adversarial samples appear.

False-Accept and False-Reject Guarantees:

Four hybrid configurations in Table I simultaneously achieve 0% FAR and 0% FRR. Zero false acceptance reduces the possibility that a forgery is mistaken for a legitimate signature, thereby removing the financial fraud vector. Zero false rejection guarantees that no genuine user is denied access, preserving usability and user acceptance. Additional ensembles approach this ideal, e.g., VGG16 + SVM attains 0.19% FAR and 0% FRR; yet only the four perfect pipelines satisfy the stringent 0% FAR and 0% FRR requirement demanded by high-security services such as remote banking and legally binding digital contracts.

The empirical observation that multiple independent architectures attain perfect security metrics indicates that offline signature forgery detection has reached the maturity needed for field deployment in safety-critical environments.

Discussion:

Explaining Perfect Classification:

Perfect classification on a skilled-forgery dataset is seldom reported, so the four zero-error hybrids deserve careful post-mortem. ImageNet pre-training supplies early kernels that respond to generic edges and textures and mid-level filters that encode stroke curvature and ink-density gradients; because these visual primitives are domain-agnostic, the networks transplant to signatures without architectural surgery. EfficientNet-B0 and ResNet-50 emit 1,280-D and 2,048-D embeddings who's high codimension preserves minute inter-class scatter, and CEDAR's 528-image test set is large enough to show that this separability is not a sampling fluke. That both SVM and logistic regression reach 100% accuracy implies the deep features are already linearly separable, so heavier classifiers add no value.

Model Hybridization Advantages:

Decoupling representation from classification brings four practical benefits. First, either component can be upgraded independently, e.g., a mobile bank can swap the classifier layer alone when latency budgets tighten. Second, frozen CNN weights eliminate gradient flow through the backbone, cutting GPU hours by two orders of magnitude relative to full fine-tuning. Third, extracted vectors can be cached and reused across dozens of downstream experiments. Finally, exhaustive pairing reveals that classifier choice alone swings accuracy by 3–4 percentage points for the same encoder, a range wide enough to justify a hybrid search strategy.

Practical Deployment Consideration:

Zero false acceptance protects financial institutions from fraudulent encashment, whereas zero false rejection preserves customer experience; EfficientNet-B0 followed by a single logistic node satisfies both constraints while keeping the decision boundary human-readable, inference under 10 ms, and memory below 20 MB. Legal-document workflows and IoT identity terminals inherit the same advantages.

Comparison With Traditional Methods:

Traditional systems spend weeks tuning LBP masks, HOG orientations, or SIFT thresholds for each ink-script pair; our pipeline replaces that labor with one forward pass through a frozen ImageNet encoder and still delivers perfect accuracy—an outcome rarely achieved by hand-crafted descriptors. The framework is also future-proof: larger self-supervised backbones can be dropped in without code changes.

Research Limitation:

We have not evaluated the hybrids on additional corpora yet; doing so is the immediate next step. We will replicate the exact pipeline on GPDS-960, MCYT-330, and BHSig-Bengali and report macro-averaged F-scores, expected within the next quarter. To probe population drift, we will adopt a leave-one-culture-out schedule: train on three corpora, test on the fourth, then apply temperature scaling and light domain-adversarial fine-tuning to keep FAR/FRR within half a percent of the CEDAR baseline. These results will appear in an extended archival version.

Limitations remain. All conclusions are conditioned on CEDAR; replication on GPDS, MCYT, and BHSig260 is required. The corpus contains only skilled forgeries, so behavior on casual or random attacks is unquantified. Temporal cues such as pen pressure and velocity are ignored, and frozen weights may be sub-optimal; task-specific fine-tuning could shrink the architecture further at the cost of extra computation. Finally, the pre-augmentation sample size (1,320 images) is modest; larger corpora would enable richer training strategies.

Future Direction:

Planned extensions include cross-dataset validation, writer-dependent personalization, fusion of dynamic tablet signals, ensemble distillation for ultra-lightweight models, and adversarial-robustness certification.

Next, we will measure the real cost of each hybrid. We plan to log FLOPs for feature extraction plus inference, record memory bandwidth when cached embeddings are fed to the

classifier, and meter energy-delay product on an off-the-shelf mobile SoC (Cortex-A78 at 2.4 GHz). These numbers will enable practitioners to balance accuracy against battery life and thermal constraints.

Conclusion:

Our evaluation of twenty CNN-classifier combinations on the CEDAR benchmark demonstrates that task-specific fine-tuning isn't strictly necessary for accurate signature verification. Four specific pairings using either EfficientNetB0 or ResNet50 matched with an SVM or Logistic Regression achieved flawless classification, recording 0% FAR and 0% FRR. The standout configuration is the EfficientNetB0 and Logistic Regression hybrid. It provides this zero-error performance while remaining highly efficient, requiring just 5 M parameters and a 9.3 ms inference time.

These findings suggest that extracting deep features from frozen networks and feeding them into simple linear boundary classifiers is a highly effective strategy for catching skilled forgeries. That being said, the primary limitation of our current study is its exclusive reliance on the CEDAR dataset. Since real-world forgery techniques and digital scanning conditions evolve constantly, testing on a single older benchmark is not enough to clear the path for immediate production in banking or legal sectors. Moving forward, it is essential to validate these hybrid pipelines against much newer, diverse signature datasets published between 2020 and 2025. Doing so will confirm whether this high accuracy actually holds up under modern acquisition conditions.

References:

[1] K. Stergiou, S. Ougiaroglou, and A. Sidiropoulos, "Signature forgery detection using deep and machine learning," *Intell. Decis. Technol.*, vol. 19, no. 3, pp. 1429–1444, May 2025, doi: 10.1177/18724981251330068.

[2] Mustafa Servet Kiran, Gokhan Seyfi, "Feature Fusion Using Deep Learning Algorithms in Image Classification for Security Purposes by Random Weight Network," *Appl. Sci.*, vol. 15, no. 16, p. 9053, 2025, doi: https://doi.org/10.3390/app15169053.

[3] S. N. Srihari, A. Xu, and M. K. Kalera, "Learning strategies and classification methods for off-line signature verification," *Proc. - Int. Work. Front. Handwrit. Recognition, IWFHR*, pp. 161–166, 2004, doi: 10.1109/IWFHR.2004.61.

[4] J. D. S. Rajesh Kumar, "Writer-independent off-line signature verification using surroundedness feature," *Pattern Recognit. Lett.*, 2012, [Online]. Available: https://dl.acm.org/doi/10.1016/j.patrec.2011.10.009

[5] J. Deng, W. Dong, R. Socher, L. J. Li, K. Li, and L. Fei-Fei, "ImageNet: A Large-Scale Hierarchical Image Database," *2009 IEEE Conf. Comput. Vis. Pattern Recognition, CVPR 2009*, pp. 248–255, 2009, doi: 10.1109/CVPR.2009.5206848.

[6] T. Ojala, M. Pietikäinen, and T. Mäenpää, "Multiresolution gray-scale and rotation invariant texture classification with local binary patterns," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, no. 7, 2002, doi: 10.1109/TPAMI.2002.1017623.

[7] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," *Proc. - 2005 IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognition, CVPR 2005*, vol. I, pp. 886–893, 2005, doi: 10.1109/CVPR.2005.177.

[8] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *Int. J. Comput. Vis.*, vol. 60, no. 2, pp. 91–110, Nov. 2004, doi: 10.1023/B:VISI.0000029664.99615.94.

[9] S. Pal, S. Chanda, U. Pal, K. Franke, and M. Blumenstein, "Off-line signature verification using G-SURF," *Int. Conf. Intell. Syst. Des. Appl. ISDA*, pp. 586–591, 2012, doi: 10.1109/ISDA.2012.6416603.

[10] J. Ruiz-Del-Solar, C. Devia, P. Loncomilla, and F. Concha, "Offline Signature

Verification Using Local Interest Points and Descriptors," *Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics)*, vol. 5197 LNCS, pp. 22–29, 2008, doi: 10.1007/978-3-540-85920-8_3.

[11] D. Impedovo and G. Pirlo, "Automatic signature verification: The state of the art," *IEEE Trans. Syst. Man Cybern. Part C Appl. Rev.*, vol. 38, no. 5, pp. 609–635, 2008, doi: 10.1109/TSMCC.2008.923866.

[12] I. Masi, Y. Wu, T. Hassner, and P. Natarajan, "Deep Face Recognition: A Survey," *Proc. - 31st Conf. Graph. Patterns Images, SIBGRAPI 2018*, pp. 471–478, Jul. 2018, doi: 10.1109/SIBGRAPI.2018.00067.

[13] J. P. Florian Schroff, Dmitry Kalenichenko, "FaceNet: A Unified Embedding for Face Recognition and Clustering," *arXiv:1503.03832*, 2015, [Online]. Available: https://arxiv.org/abs/1503.03832

[14] M. E. ; M. H E ; Kim, "Lightweight Visual Transformers Outperform Convolutional Neural Networks for Gram-Stained Image Classification: An Empirical Study," *Biomedicines*, vol. 11, no. 5, 2023, doi: https://doi.org/10.3390/biomedicines11051333.

[15] F. Boutros, N. Damer, F. Kirchbuchner, and A. Kuijper, "ElasticFace: Elastic Margin Loss for Deep Face Recognition," *IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit. Work.*, vol. 2022-June, pp. 1577–1586, 2022, doi: 10.1109/CVPRW56347.2022.00164.

[16] L. G. Hafemann, R. Sabourin, and L. S. Oliveira, "Offline handwritten signature verification - Literature review," *Proc. 7th Int. Conf. Image Process. Theory, Tools Appl. IPTA 2017*, vol. 2018-January, pp. 1–8, Jul. 2017, doi: 10.1109/IPTA.2017.8310112.

[17] H. K. Pathak, U. Chaurasia, R. R. Upadhyay, K. K. Singh, and A. Kaplan, "Transfer learning based approach for reliable handwritten signature verification with SignaLearn," *Discov. Artif. Intell. 2026*, Feb. 2026, doi: 10.1007/s44163-026-00929-6.

[18] S. Dey, A. Dutta, J. I. Toledo, S. K. Ghosh, J. Llados, and U. Pal, "SigNet: Convolutional Siamese Network for Writer Independent Offline Signature Verification," Sep. 2017, Accessed: Feb. 13, 2026. [Online]. Available: http://arxiv.org/abs/1707.02131

[19] Wanghui Xiao & Hao Wu, "Learning features for offline handwritten signature verification using spatial transformer network," *Sci. Rep.*, vol. 15, 2025, [Online]. Available: https://www.nature.com/articles/s41598-025-92704-3

Copyright © by authors and 50Sea. This work is licensed under the Creative Commons Attribution 4.0 International License.