Lightweight Embedded Framework for Edge Document Authenticity Detection

Emma Bigaret Salma Samiei Stephane Cauchie
AI Research and Development Team, Idakto, Angers, France
{Emma.Bigaret, salma.samiei, stephane.cauchie}@idakto.com

Abstract

The rise of generative models has intensified the threat of highly realistic identity-document forgeries. We propose a compact, system-level framework that unifies visual, textual, and security cues within a calibrated decision pipeline. Unlike isolated detectors, our approach integrates heterogeneous modules into a resource-constrained architecture, ensuring both robustness and interpretability. Evaluated on public benchmarks (MIDV, DLC, FMIDV, MIDV-Holo) and a large in-house dataset spanning 128 countries, the framework achieves competitive performance in both detection and localization. Its efficiency and accuracy make it a practical candidate for mobile and embedded ID verification.

1. Introduction

Identity document verification is a cornerstone of security in border control, eID infrastructures, and financial services, where authenticity checks are essential to prevent fraud and identity theft. The emergence of powerful generative models has intensified the threat of sophisticated forgeries, ranging from synthetic portraits to text inpainting and the manipulation of embedded security features. Recent surveys and research highlight that deepfake-driven identity fraud is no longer a theoretical risk but a pressing challenge for verification systems [5, 8]. Manual inspection has become insufficient, and traditional single-modality solutions often fail under real-world conditions, motivating the need for scalable, interpretable, and multimodal machine learning frameworks [10].

Prior studies have explored isolated detection strategies, such as the analysis of guilloché patterns using CNNs [1], graph-based OCR consistency checks [7], or handcrafted symmetry descriptors [3]. Others have proposed datasets and benchmarks for document forensics, such as DLC-2021 for liveness [9], or end-to-end architectures for forgery detection and localization [6]. While valuable, these efforts remain limited to specific subproblems or dataset conditions. What is still missing is a holistic, system-level approach that unifies complementary signals—visual, textual, and structural—into a calibrated pipeline deployable on constrained platforms.

In this paper, we address this gap by introducing a compact verification framework that treats document authenticity as probabilistic inference over heterogeneous cues. Unlike monolithic detectors, our approach orchestrates multiple interpretable modules within a resource-efficient architecture, enabling both robustness to unseen attack types and practical deployment on mobile or embedded devices.

The main contributions of this work are threefold:

- A system-level pipeline that integrates heterogeneous modules into a unified, interpretable decision process,
- A deployment-ready architecture optimized for constrained environments,
- A comprehensive evaluation across public datasets (MIDV, DLC, FMIDV, MIDV-Holo) and a diverse in-house dataset spanning 128 countries.

2. Related Work

Research on document forgery detection has traditionally addressed individual cues rather than full verification pipelines. One well-studied direction is the analysis of embedded security patterns such as guilloché designs. CNN-based approaches have been developed to distinguish authentic from forged textures by exploiting fine-grained geometric regularities [1], and the FMIDV dataset has supported methods based on contrastive or adversarial learning [3]. These works demonstrate the discriminative value of structural features but remain confined to a narrow set of artifacts, limiting their applicability in broader verification contexts.

Textual manipulations have been studied through Optical Character Recognition (OCR). In particular, graph-based representations of OCR outputs, where bounding boxes are modeled as nodes with spatial relations, have proven effective in detecting tampering through misalignments and scaling inconsistencies [7]. Such methods are precise in capturing local textual edits but overlook complementary dimensions of document integrity, such as holograms, portrait consistency, or liveness cues.

Parallel efforts have investigated liveness and presen-

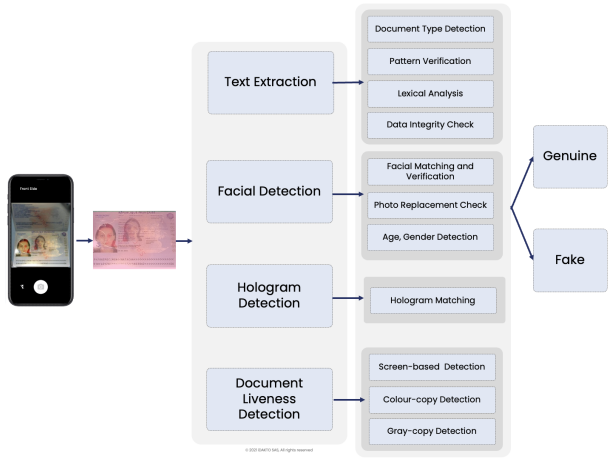

Figure 1. Overview of the proposed document forgery detection pipeline. The system integrates localization, text analysis, security feature verification, liveness assessment, and portrait consistency into a unified decision process.

tation attack detection, originally in biometric authentication [5, 8] and later adapted to documents. The DLC-2021 dataset [9] introduced benchmarks for distinguishing genuine physical IDs from recaptures and printouts, enabling progress in this area. Yet, these systems typically operate as standalone modules and are not designed for integration with other forgery checks.

More recently, attempts have been made to design holistic approaches that go beyond isolated cues. End-to-end architectures using CNNs or transformers aim to detect forgeries directly from document images, combining localization and classification in a single model [6]. Datasets such as MIDV-2020 and DLC-2021 [9] have played an important role in standardizing evaluation. However, these monolithic models tend to prioritize accuracy over interpretability and often require significant computational resources, which restricts their suitability for deployment in mobile or embedded scenarios.

3. Methodology

Our framework approaches document forgery detection as the integration of multiple complementary signals into a unified verification process. The overall pipeline is illustrated in Figure 1. Each stage has been designed to contribute independent evidence of authenticity while remaining computationally efficient, enabling deployment on mobile and embedded platforms. The modular structure also allows individual components to be updated or replaced without redesigning the entire system, which is crucial for adapting to evolving forgery techniques.

The pipeline begins with document localization in unconstrained scenes. To achieve this, we employ YOLOv8-based instance segmentation, which offers a favorable balance between inference efficiency and accuracy under varying illumination and background conditions. Once localized, the document undergoes textual analysis through OCR, relying on Tesseract and MLKit engines for robustness across languages and fonts. Key fields such as the Machine Readable Zone (MRZ), Card Access Number (CAN), and document serial identifiers are extracted and cross-validated against canonical templates. This step not only recovers textual content but also serves as an early defense against tampering by flagging inconsistencies at the character or layout level.

Beyond textual validation, the system incorporates verification of embedded security features that are inherently difficult to replicate digitally. Guilloché patterns, which consist of fine-line repetitive structures, are analyzed using contour regularity and repetition statistics, as counterfeit reproductions typically introduce measurable geometric deviations. An example is shown in Figure 2. In parallel, holograms are examined using a ResNet50 network trained on the MIDV-Holo dataset. Since holographic responses change with illumination and viewpoint, they provide a discriminative cue that static printouts or screenshots cannot replicate. Together, these checks secure the structural and material integrity of the document.

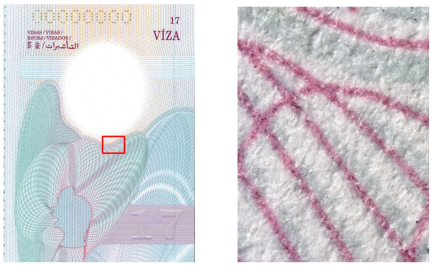

(a) Motif guilloché sur un document russe (b) Zoom sur le motif guilloché

Figure 2. Example of guilloché pattern verification from an identity document. The fine-line geometry is highly sensitive to digital forgeries.

A further layer of protection comes from liveness assessment, which ensures that the document being inspected is a physical object rather than a reproduction. We adopt a ResNet50 classifier trained on the Document Liveness Challenge (DLC-2021) dataset, exploiting reflection cues, noise characteristics, and device-specific artifacts to distinguish genuine captures from photocopies or screen replays. Unlike handcrafted reflection descriptors, the learned features generalize more effectively across acquisition conditions, thereby mitigating the risk of injection attacks where forged documents are introduced via digital display or reprinting.

In addition to verifying structural and physical properties, the system enforces global layout and biometric consis-

tency. Template matching is performed using the Structural Similarity Index (SSIM) against a reference library covering 128 issuing countries. Deviations in field placement or spacing often reveal unauthorized modifications. Complementing this, portrait consistency is checked by comparing the primary and ghost images present in many identity documents. For this task, we rely on the SFace model [2], which verifies whether both portraits correspond to the same individual. As illustrated in Figure 3, this mechanism provides a safeguard against face-swap forgeries that alter the main image while leaving the ghost image intact.

Figure 3. Comparison of primary and ghost portraits for intra-document consistency verification. Inconsistencies reveal potential face-swap forgeries.

Outputs from all modules are normalized into comparable confidence scores and fused into a global authenticity decision. This decision-level integration ensures redundancy: if one signal becomes unreliable under specific capture conditions, others can compensate. The resulting architecture thus combines efficiency, interpretability, and robustness, making it well suited to real-world verification scenarios where diverse attack vectors and constrained computational budgets coexist.

4. Experiments and Discussion

The pipeline was evaluated on a set of public benchmarks together with a diverse in-house dataset covering 128 issuing countries. Each module was tested on the dataset most relevant to its function: document detection on MIDV-2020, liveness verification on DLC-2021, hologram analysis on MIDV-Holo, and portrait consistency on MIDV/DLC. Template verification was carried out on the in-house corpus, where a wide variety of layouts and fonts are represented. The quantitative results are summarized in Table 1.

Figure 4 presents representative cases across modalities: the original document, a screen recapture, a high-quality copy, and a grayscale print. These examples highlight the strengths of liveness detection in rejecting recaptures, but also illustrate the difficulty of distinguishing certain copies from genuine originals. Figure 5 shows examples of document detection across different capture conditions.

Table 1. Performance summary across main pipeline components.

Task	Dataset	Accuracy (%)
Doc detection	MIDV-2020	99.7
Liveness	DLC-2021	96.0
Hologram	MIDV-Holo	94.0
Portrait cmp.	MDV/DLC	84.0
Template match	In-house	94.0

These visualizations illustrate the robustness of YOLOv8 in handling cluttered backgrounds and varying illumination, which is critical for reliable preprocessing before downstream analysis. In addition to such quantitative measures, qualitative inspection is essential for understanding error modes.

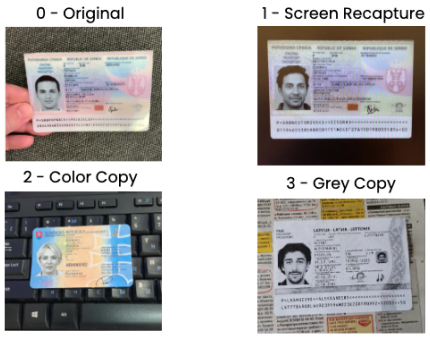

Figure 4. Qualitative examples of document analysis across modalities: (0) original identity document, (1) screen recapture, (2) high-quality color copy, and (3) grey-scale copy. These cases illustrate typical challenges for liveness detection and authenticity verification.

The main advantage of the system lies in the combination of heterogeneous signals. OCR and template checks validate the textual integrity of the document, guilloché and hologram analysis secure structural elements, liveness verification defends against replay attacks, and portrait consistency prevents identity substitution. When these cues are fused, the system is less dependent on any single component: errors in one channel can be compensated by evidence from others. This redundancy is particularly valuable in real-world settings, where acquisition conditions vary and adversaries exploit multiple attack surfaces.

Practical deployment requires that such verification be carried out under constrained computational budgets. The choice of YOLOv8 for segmentation and ResNet50 backbones for classification tasks reflects a compromise between accuracy and efficiency, allowing near real-time operation on mobile devices. The modular design further makes it possible to distribute tasks between edge devices and the

Figure 5. Examples of identity document detections on different images.

cloud, offering flexibility for scenarios ranging from airport kiosks with stable connectivity to smartphone-based verification where privacy concerns may preclude cloud processing.

Although the overall performance is strong, several sources of error persist. False positives—authentic documents incorrectly flagged as forged—are often linked to reflections on holographic elements, unusual color casts, or print quality variations that mimic tampering. Genuine security foils, for instance, can trigger the hologram detector under certain lighting conditions. Recent work on edge-aware representations, such as the edge attention–edge concatenation (EA–EC) module [4], suggests that explicitly modeling boundary discontinuities can mitigate such misclassifications. Incorporating these ideas into our pipeline would likely reduce both false positives and negatives.

The evaluation also highlights broader limitations. Most available datasets are biased toward Latin-script documents, limiting generalization to documents using Arabic, Cyrillic, or Chinese scripts. Hologram verification is currently performed on static images, while temporal responses under varying illumination would provide a stronger signal. High-quality copies remain difficult to distinguish from genuine originals, and although our in-house dataset is geographically diverse, it cannot fully represent the spectrum of real-world issuing authorities.

Addressing these issues opens several directions for future work. Developing multilingual OCR models would extend coverage beyond Latin scripts, while dynamic hologram verification could exploit temporal cues for greater resilience. Federated learning across institutions offers a way to enlarge training sets without centralizing sensitive data. Finally, integrating edge-focused modules such as EA–EC promises to improve robustness against benign artifacts and adversarial manipulations. Pursuing these directions would enhance both the accuracy and the adaptability of document authentication systems in practice.

5. Conclusion

We introduced a compact and modular pipeline for identity document forgery detection that integrates document localization, OCR, security feature analysis, liveness verification, and portrait consistency into a unified decision process. Evaluated on multiple public benchmarks and a diverse in-house dataset, the system achieves competitive performance while remaining lightweight enough for deployment on mobile and embedded platforms. Unlike prior approaches limited to a single modality, our framework leverages complementary cues to provide robustness and interpretability, raising the barrier for attackers and offering a practical foundation for real-world identity verification.

References

[1] Musab Al-Ghadi, Zuheng Ming, Petra Gomez-Krämer, and Jean-Christophe Burie. Identity documents authentication based on forgery detection of guilloche pattern. In *arXiv preprint arXiv:2206.10989*, 2022. 1

[2] Brandon Amos, Bartosz Ludwiczuk, and Mahadev Satyanarayanan. Openface: A general-purpose face recognition library with mobile applications. In *CMU School of Computer Science*, 2016. 3

[3] Y. Y. Bae. Enhancing document forgery detection with edge ... *Symmetry (MDPI)*, 17(8):1208, 2025. 1

[4] Yong-Yeol Bae, Dae-Jea Cho, and Ki-Hyun Jung. Enhancing document forgery detection with edge-focused deep learning. *Symmetry*, 17:1208, 2025. 4

[5] Sayan Banerjee, Sumit Kumar Yadav, Ankit Dhara, and Md Ajij. A survey: Deepfake and current technologies for solutions. *CEUR Workshop Proceedings*, 2025. 1, 2

[6] Anjith George and Sébastien Marcel. Edgedoc: Hybrid cnn-transformer model for accurate forgery detection and localization in id documents. *arXiv preprint arXiv:2508.16284*, 2025. 1, 2

[7] Hailey Joren, Otkrist Gupta, and Dan Raviv. Ocr graph features for manipulation detection in documents. *arXiv preprint arXiv:2009.05158*, 2020. 1

[8] Yisroel Mirsky and Wenke Lee. The creation and detection of deepfakes: A survey. *ACM Computing Surveys*, 2020. arXiv preprint arXiv:2004.11138. 1, 2

[9] D. V. Polevoy. Document liveness challenge dataset (dlc-2021). *Journal of Imaging*, 8(7):181, 2022. 1, 2

[10] A. Vaidya. A conceptual model for ai-powered identity verification. *arXiv preprint arXiv:2503.08734*, 2025. 1