

Automated ICAO/ISO Compliance Validator for ID Portraits

Mohanned Naji Albibas¹, AL-Mokhtar Abdullah Faraj², Musa Faneer³

¹Administrative Control Authority, Tripoli, Libya

²University of Ghryan, Ghryan-Libya

³The Libyan Academy, Tripoli, Libya

mohn5150@gmail.com

Received:12/06/2026

Accepted: 19/06/2026

Published: 31/07/2026

الملخص

لا يزال التحقق اليدوي من صور وثائق الهوية ومطابقتها للمعايير الدولية عملية كثيفة العمالة وعرضة للأخطاء، مما يحد من قابلية التوسع في سير عمل الهوية الرقمية عالية الحجم. تقدم هذه الورقة البحثية أداة تحقق آلي مفتوحة المصدر تفرض برمجياً متطلبات وثيقة منظمة الطيران المدني الدولي (ICAO) رقم 9303 والمعيار ISO/IEC 19794-5 الخاصة بصور الهوية الشخصية. يدمج النظام المقترح واجهة برمجة تطبيقات المهام الحديثة من MediaPipe لاستخراج علامات الوجه بدقة، مع فحوصات قائمة على قواعد هندسية لنسبة الرأس إلى الإطار (70-80%)، وموقع خط العينين (50-60% من الأعلى)، وتفاوت وضعية الرأس ($\pm 10^\circ$)، إلى جانب تقييمات بيئية لتجانس الخلفية وتناسق الإضاءة. تتيح البنية المعيارية الأمانة للخيوط المتعددة-Thread (safe) مزدوجاً: واجهة برمجة تطبيقات REST عديمة الحالة باستخدام Flask لتكامل الأنظمة، وواجهة رسومية متعددة المنصات باستخدام Tkinter للتحقق التفاعلي، وكلاهما يُنتج تقارير JSON منظمة تتضمن حالة كل فحص، ومقاييس كمية، وتوصيات قابلة للتنفيذ. يُظهر التقييم على مجموعة بيانات موسعة ومنسقة مكونة من 980 صورة شخصية عبر ثلاث مستويات صعوبة دقة عالية (0.97) واستدعاء (0.95) لاكتشاف الامتثال العام، مع متوسط زمن انتقال إجمالي يبلغ 225 ± 22 مللي ثانية على أجهزة المعالج المركزي. تقلل أداة التحقق من عبء المراجعة اليدوية بنسبة تزيد عن 85% مع الحفاظ على قابلية التدقيق الكاملة من خلال معالجة مؤقتة تحافظ على الخصوصية. من خلال الجمع بين التنفيذ المتوافق مع المعايير والنشر الجاهز للإنتاج، يوفر هذا العمل أساساً قابلاً للتكرار وقابلاً للتشغيل دون اتصال لسير عمل تصوير الهوية القابل للتوسع في البيئات محدودة الموارد.

الكلمات المفتاحية: الامتثال لمنظمة الطيران المدني الدولي (ICAO)، المعيار ISO 19794-5، التحقق من الصور الشخصية، علامات الوجه، مراقبة الجودة الآلية، واجهة برمجة تطبيقات مهام MediaPipe، واجهة برمجة تطبيقات Flask REST

Abstract

The manual verification of identity document portraits against international standards remains a labor-intensive and error-prone process, limiting scalability in high-volume digital identity workflows. This paper presents an automated, open-source compliance validator that programmatically enforces ICAO Doc 9303 and ISO/IEC 19794-5 requirements for ID portrait photographs. The proposed system integrates MediaPipe's modern Tasks API for robust face landmark extraction, geometric rule-based checks for head-to-frame ratio (70–80%), eye-line positioning (50–60% from top), and head pose tolerance ($\pm 10^\circ$), alongside environmental assessments for background uniformity and lighting consistency. A modular, thread-safe architecture enables dual deployment: a stateless Flask REST API for system integration and a cross-platform Tkinter graphical interface for interactive validation, both producing structured JSON reports with per-check status, quantitative metrics, and actionable recommendations. Evaluation on an expanded, curated dataset of 980 portrait images across three difficulty tiers demonstrates high precision (0.97) and recall (0.95) for overall compliance detection, with average end-to-end latency of 225 ± 22 ms on CPU hardware. The validator reduces manual review overhead by $>85\%$ while maintaining full auditability through ephemeral, privacy-preserving processing. By bridging standards-aligned rule enforcement with production-ready deployment, this work provides a reproducible, offline-capable foundation for scalable identity photography workflows in resource-constrained environments.

Keywords: ICAO Compliance, ISO 19794-5, Portrait Validation, Face Landmarks, Automated Quality Control, MediaPipe Tasks API, Flask REST API

Introduction

The global expansion of digital identity ecosystems has intensified demand for automated, standards-compliant portrait photography workflows. International specifications such as ICAO Doc 9303 [1] and ISO/IEC 19794-5 [2] define stringent requirements for identity document photographs, including precise head-to-frame ratios (70–80% of image height), eye-line positioning (50–60% from the top edge), head pose tolerance ($\pm 10^\circ$ yaw/pitch/roll), background uniformity, and minimum resolution (600×600 pixels at 300 DPI equivalent). Manual verification of these criteria remains labor-intensive, subjective, and prone to inconsistency, particularly in high-volume enrollment scenarios or resource-constrained environments where specialized photographic expertise may be unavailable [3].

Recent advances in computer vision and deep learning have enabled automated background removal, face detection, and geometric alignment [4]–[6]. However, most existing solutions address isolated subtasks—such as semantic segmentation via U²-Net [7] or landmark-based face centering [8]—without integrating rule-based compliance validation into a cohesive, reproducible pipeline. Commercial APIs and proprietary desktop tools often operate as black boxes, limiting transparency, customization, and offline deployment [9]. Furthermore, few systems provide structured, machine-readable reports that support auditability, batch processing, or seamless integration into larger identity verification workflows.

To address these gaps, this paper presents an automated, open-source compliance validator that programmatically enforces ICAO/ISO portrait requirements using modern, accessible tooling. The proposed system leverages MediaPipe's Tasks API [10] for robust, thread-safe face landmark extraction, followed by geometric rule-based checks for head ratio, eye position, and pose estimation, alongside environmental assessments for background uniformity and lighting consistency. A modular architecture enables dual deployment: a stateless Flask REST API for system integration and a cross-platform Tkinter graphical interface for interactive validation, both producing structured JSON reports with per-check status, quantitative metrics, and actionable recommendations.

The primary contributions of this work are as follows:

1. A rule-based compliance engine with configurable ICAO/ISO thresholds, implemented as a reusable, thread-safe Python class.
2. Integration of MediaPipe's modern Tasks API for robust landmark extraction, eliminating dependency on legacy inference wrappers.
3. Dual-deployment framework supporting both interactive desktop validation and programmatic API access, with shared core logic.
4. Structured metadata reporting with actionable recommendations, enabling auditability and reducing manual review overhead.
5. Comprehensive evaluation on an expanded, diverse dataset of 980 portrait images, demonstrating high precision (0.97) and recall (0.95) for overall compliance detection.

The remainder of this paper is organized as follows. Section II reviews related work in portrait segmentation, face alignment, and automated compliance checking. Section III details the system architecture and core methodology, including landmark extraction, geometric checks, and environmental assessments. Section IV outlines implementation specifics, software dependencies, and optimization strategies. Section V presents experimental results, quantitative evaluations, and compliance assessments. Section VI discusses limitations and potential mitigations, while Section VII concludes the study and outlines directions for future research.

Related Work

The automation of identity document photography has evolved from heuristic image processing to deep learning-driven pipelines. However, significant gaps remain in end-to-end standards compliance, rule-based validation transparency, and production-ready deployment. This section reviews prior approaches across four dimensions: traditional compliance checking, deep learning for portrait analysis, commercial validation services, and the positioning of the proposed work.

Traditional and Heuristic Compliance Checking

Early automated ID photo systems relied on rule-based techniques such as chroma-key thresholding, edge detection, and manual region-of-interest (ROI) cropping [1], [2]. While computationally lightweight, these heuristic methods struggle with complex lighting, non-uniform backgrounds, and fine hair boundaries, often requiring manual post-processing to meet International Civil Aviation Organization (ICAO) specifications [3]. Commercial desktop tools improved usability through graphical interfaces but remained closed-source, cost-prohibitive, and incapable of dynamic compliance verification without user intervention [4]. Furthermore, traditional pipelines lack automated rule enforcement for head-to-frame ratios, eye-line positioning, and pose tolerances, leaving compliance validation entirely to human inspectors.

Deep Learning for Portrait Analysis

The introduction of convolutional neural networks significantly advanced foreground-background separation and facial landmark localization. U-Net architectures established a strong baseline for biomedical segmentation, later adapted for portrait matting via nested structures [5]. Open-source ecosystems have wrapped U-Net variants into accessible inference pipelines, enabling plug-and-play background removal without custom training [6]. While transformer-based models offer zero-shot generalization, their computational overhead and lack of portrait-specific optimization make them less suitable for lightweight compliance workflows [7]. Modern face detection libraries (e.g., dlib, MTCNN, MediaPipe) improve landmark precision and centering accuracy [8], [9], yet most academic implementations stop at detection without integrating downstream rule-based compliance logic. Consequently, advanced segmentation and alignment models are rarely coupled with automated standards validation.

Commercial and Cloud-Based Validation Services

Recent commercial APIs and cloud platforms offer automated portrait cropping, background replacement, and basic quality scoring. These services achieve high accuracy but operate as opaque black boxes, limiting transparency, customization, and offline deployment [10]. Subscription-based pricing, network dependency, and rate limiting further restrict scalability in institutional or resource-constrained environments. Crucially, most commercial tools do not expose structured compliance metadata, audit trails, or configurable thresholds aligned with ICAO Doc 9303 or ISO/IEC 19794-5, forcing organizations to rely on manual verification for regulatory adherence.

Positioning of This Work

Unlike prior approaches that isolate segmentation, alignment, or quality assessment, this paper introduces a cohesive, open-source compliance validator that integrates modern landmark extraction, geometric rule enforcement, and environmental checks within a single framework. By providing both a cross-platform desktop interface and a stateless REST API, the system bridges research-grade computer vision with real-world deployment requirements. The proposed validator addresses the fragmentation, opacity, and inflexibility of existing solutions by delivering transparent, configurable, and offline-capable compliance checking.

Table 1 Comparison of Compliance Validation and Standardization Approaches

Method	Open Source	Offline Capable	Rule Transparency	Compliance Coverage	Deployment Type
Thresholding/Chroma-Key	Yes	Yes	High	Low	Desktop CLI
Commercial ID Software	No	Yes	Low	Medium (Manual)	Closed GUI
Cloud APIs (e.g., remove.bg)	No	No	None	High (Black-box)	Cloud REST
U-Net / DeepLab Scripts	Yes	Yes	Medium	Low (Detection only)	Research Code
Proposed Validator	Yes	Yes	High	High (Configurable)	GUI + REST API

System Architecture & Methodology

The proposed system implements a deterministic, rule-based validation pipeline that transforms raw portrait photographs into structured compliance assessments. Unlike generative or corrective pipelines that modify input images, this architecture focuses on precise measurement, threshold enforcement, and audit-ready reporting. The

methodology integrates modern neural landmark extraction, classical geometric analysis, and environmental image statistics within a modular, thread-safe framework. This section details the standards mapping, processing sequence, core validation algorithms, and deployment architecture.

Standards Reference: ICAO Doc 9303 & ISO/IEC 19794-5

International identity document photography is governed by stringent biometric and compositional requirements. ICAO Doc 9303 [1] and ISO/IEC 19794-5 [2] specify that portrait photographs must maintain a head-to-frame ratio between 70% and 80% of the image height, with the eye line positioned at 50% to 60% from the top edge. Head pose deviations must not exceed $\pm 10^\circ$ in yaw, pitch, or roll to ensure frontal visibility of biometric features. Backgrounds must be uniform, light-colored, and free of shadows or patterns, while illumination must be even, avoiding harsh contrasts or backlighting that obscure facial landmarks. The proposed validator translates these qualitative guidelines into quantifiable computational metrics, enabling automated, reproducible compliance checking without manual intervention.

Overall Validation Workflow

The pipeline operates as a sequential, stateless processing chain. Given an input image, the system first validates resolution and aspect ratio against minimum thresholds (600×600 pixels). A MediaPipe-based face detector then isolates the primary facial region and extracts 478 3D facial landmarks. These coordinates serve as the foundation for geometric proportion calculations, while the original pixel data undergoes environmental analysis for background uniformity and lighting consistency. Each check produces a structured status flag (PASS/FAIL), quantitative metric, and threshold comparison. Results are aggregated into a compliance verdict, with failed triggers mapped to actionable remediation steps. The final output consists of a machine-readable JSON report and an annotated preview image highlighting compliance boundaries and alignment guides.




Figure 1 End-to-End Compliance Validation Workflow

Core Compliance Engine

The validation logic is encapsulated within a modular engine that decouples inference from presentation layers. All checks operate independently but share a unified configuration interface, enabling parameter tuning for regional or institutional standards.

1. Face Detection & Landmark Extraction

Robust facial localization is achieved using the MediaPipe Tasks API FaceLandmarker model [3], which provides 478 precise 3D facial coordinates normalized to the image frame. The inference engine runs in a thread-safe manner using a synchronization lock, allowing concurrent processing across GUI and API interfaces without race conditions. Unlike heuristic detectors, the neural landmarker maintains high accuracy across diverse skin tones, lighting conditions, and mild occlusions. If no face is detected or multiple faces are present, the pipeline immediately flags the input as non-compliant and logs a structured warning.

2. Geometric Checks

Geometric compliance is evaluated through three primary metrics derived from normalized landmark coordinates:

- Head-to-Frame Ratio: Calculated as the vertical distance between landmark 10 (top of head) and landmark 152 (chin), divided by total image height. Values outside the 0.70–0.80 range trigger a framing violation.

- Eye Line Position: Determined by averaging the Y-coordinates of the left eye (landmark 33) and right eye (landmark 263), then normalizing by image height. The compliant zone spans 0.50–0.60 from the top edge.
- Head Pose Estimation: To eliminate the inaccuracies of heuristic approximations, exact 3D head pose is now computed using OpenCV's `cv2.solvePnP` algorithm. The system maps a subset of MediaPipe's 3D facial landmarks (e.g., nose tip, chin, and bilateral eye corners) to a canonical 3D face model. By solving the Perspective-n-Point problem, the pipeline extracts precise Euler angles (yaw, pitch, and roll). This deterministic approach ensures strict adherence to the $\pm 10^\circ$ ICAO tolerance limits, significantly reducing false rejections caused by heuristic drift.

3. Environmental Checks

Portrait quality extends beyond facial proportions to encompass capture conditions. Background uniformity is assessed by sampling the outer 10% border region of the image, converting it to grayscale, and computing pixel intensity variance. A variance threshold of ≤ 150.0 ensures the absence of distracting patterns, shadows, or gradients. Lighting consistency is evaluated using global grayscale mean intensity (target range: 80–200) and contrast standard deviation (target range: 20–120), effectively flagging underexposed, overexposed, or high-contrast captures that violate ICAO illumination guidelines.

4. Recommendation Generator

To bridge the gap between technical validation and user action, the system maps each failed check to a human-readable remediation directive. For example, a head ratio violation triggers "Adjust subject distance or crop to achieve 70–80% head-to-frame ratio," while a background failure recommends "Use uniform, light-colored background without shadows or patterns." These recommendations are dynamically compiled and embedded directly into the JSON report and GUI display.

Dual-Deployment Framework

The methodology is implemented across two complementary interfaces to support both interactive verification and automated system integration:

1. *Desktop GUI*: Built with Tkinter and ttk, the application features an intuitive canvas for image preview, one-click validation, and collapsible report panels. Heavy inference tasks are delegated to daemon threads to maintain UI

responsiveness, while structured logging captures processing metadata for audit trails.

2. *RESTful API*: The Flask-based backend exposes stateless endpoints (/api/v1/validate) supporting multipart file uploads and Base64-encoded payloads. CORS policies, payload size limits, and environment-driven configuration ensure production readiness. Each response includes a comprehensive JSON payload containing check statuses, quantitative metrics, compliance verdicts, and a Base64-encoded annotated image for visual verification.

Implementation Details

This section details the software architecture, computational optimizations, robustness mechanisms, and deployment configurations that enable the compliance validator to operate efficiently across both interactive desktop and scalable web environments. The implementation prioritizes modularity, hardware accessibility, and production-grade reliability, ensuring seamless integration into institutional identity workflows.

Software Stack & Core Dependencie

The system is implemented in Python 3.10+, leveraging a curated combination of computer vision, deep learning, and application frameworks. Facial landmark extraction is powered by the MediaPipe Tasks API FaceLandmarker model [1], which replaces legacy inference wrappers with a modern, thread-safe architecture. Matrix operations and geometric checks are handled by OpenCV and NumPy, while PIL/Pillow manages high-level image rendering and GUI canvas updates. The desktop interface is constructed using Tkinter with ttk theming, and the REST API is deployed via Flask and flask-cors for cross-origin compatibility [2], [3]. Optional dependencies are gracefully managed via conditional imports, ensuring the validator remains functional in constrained or air-gapped environments. All external models are auto-provisioned on first execution from Google's official CDN, eliminating manual setup requirements.

Computational & Memory Optimizations

To maintain responsive performance on consumer-grade hardware, several optimization strategies are embedded throughout the pipeline. Input images are dynamically pre-scaled to a maximum dimension of 1024 pixels prior to inference, significantly reducing memory consumption and landmark detection latency without compromising geometric measurement accuracy [4]. The MediaPipe landmarker

session is initialized once and cached via a singleton pattern, avoiding redundant weight loading across consecutive GUI or API requests. All annotation resizing operations exclusively utilize Lanczos-4 interpolation, which preserves high-frequency boundary details during preview generation. In the desktop interface, heavy inference and I/O tasks are delegated to daemon threads with real-time progress callbacks, preventing main-thread UI blocking. The API employs a stateless request-response model with configurable execution timeouts and strict payload limits (50 MB) to sustain throughput under concurrent load, while temporary files are explicitly deleted post-processing to prevent disk exhaustion.

Furthermore, the integration of `cv2.solvePnP` for pose estimation is optimized by caching the canonical 3D model points in memory, ensuring that the iterative pose refinement process adds negligible latency to the overall pipeline.

Robustness, Fallback Mechanisms, & Logging

Production reliability is enforced through multi-layered error handling and graceful degradation. A structured `ErrorLogger` class implements timestamped, UTF-8 encoded file logging with simultaneous console output, capturing full tracebacks, parameter states, and processing milestones for auditability [5]. If `MediaPipe` fails to detect a face or returns insufficient landmark confidence, the pipeline immediately flags the input as non-compliant, logs a structured warning, and skips downstream geometric checks to prevent cascading errors. The API validates all incoming payloads, enforces MIME-type restrictions, and returns structured JSON error responses with explicit categorization (e.g., `validation_failed`, `system_error`, `payload_too_large`). Thread-safe inference is guaranteed via a synchronization lock (`threading.Lock`) that serializes `.detect()` calls across concurrent Flask workers and GUI callbacks, eliminating race conditions without sacrificing throughput. Configuration mismatches and version-specific API changes are intercepted via dynamic `TypeError` and `AttributeError` handling, ensuring backward compatibility across `MediaPipe` updates.

Standards-Aligned Configuration Management

The validator is explicitly engineered to align with international identity document specifications, particularly ICAO Doc 9303 [6] and ISO/IEC 19794-5 [7] guidelines. All compliance thresholds are centralized within a `ComplianceConfig` dataclass, enabling parameter tuning for regional or institutional standards without code modification. Configuration is deployed identically across both interfaces: the GUI exposes real-time parameter controls via structured input fields, while the API leverages environment variables and a centralized `Config` class for deployment-agnostic settings (e.g., `HEAD_RATIO_MIN`, `BG_VAR_THRESHOLD`, `MAX_POSE_DEG`). All presets are validated against strict enumerations to prevent

malformed outputs, and metadata reporting ensures full traceability of applied thresholds, model versions, and processing latency. This configuration-driven architecture supports reproducibility, automated compliance auditing, and seamless integration into larger biometric preprocessing pipelines.

Experimental Results & Evaluation

This section presents the quantitative and qualitative evaluation of the proposed compliance validator across rule-based accuracy, processing latency, system scalability, and alignment with ICAO/ISO identity document standards. All experiments were conducted to validate the pipeline's suitability for production-grade identity verification workflows and to benchmark its performance against manual inspection baselines.

Experimental Setup & Dataset

The system was evaluated on a workstation equipped with an Intel Core i7-12700K CPU, 32 GB RAM, and an NVIDIA RTX 3060 12 GB GPU, running Ubuntu 22.04 LTS with Python 3.10, OpenCV 4.8.0, and MediaPipe 0.10.14. To ensure statistical significance and robust representation of edge cases, the dataset was expanded to 980 portrait images. This expanded corpus was assembled from public repositories and controlled captures. To ensure statistical significance and robust representation of edge cases, the dataset was expanded to 980 portrait images. This expanded corpus was assembled from public repositories and controlled captures, specifically augmenting the "Hard" tier with diverse demographic variations, extreme lighting conditions, and complex occlusions (e.g., thick-rimmed eyewear, fine hair boundaries), spanning diverse lighting conditions (studio, natural, low-light), background complexities (uniform, cluttered, patterned), and subject variations (frontal, slight tilt, eyewear, facial hair). Images were partitioned into three difficulty tiers: Easy (plain backgrounds, frontal poses), Medium (textured backgrounds, mild occlusions), and Hard (complex lighting, fine hair boundaries, non-uniform illumination). Evaluation metrics included per-check precision/recall, end-to-end validation latency, peak memory consumption, and compliance accuracy against a manually annotated ground-truth set verified by certified ID photography auditors.

Quantitative Validation Performance

Processing accuracy was measured across each compliance rule, with MediaPipe landmark extraction serving as the foundation for geometric checks. As shown in Table 2, the validator achieves high precision and recall across all dimensional and environmental checks. Geometric rules (head ratio, eye position) demonstrate robust

Mohanned Naji Albibas.

performance due to the stability of normalized landmark coordinates, while environmental checks (background variance, lighting mean/contrast) exhibit slightly lower recall in Hard-tier images with extreme backlighting or low-contrast backgrounds. Pre-scaling inputs to a maximum of 1024 pixels reduced peak RAM usage by 34% and improved CPU throughput by 1.9× without degrading measurement accuracy.

Table 2 Per-Check Compliance Validation Metric

Compliance Check	Precision	Recall	F1-Score	Avg. Latency (ms)
Resolution	1.00	1.00	1.00	<1
Head Ratio	0.97	0.95	0.96	46 ± 7
Eye Position	0.96	0.94	0.95	43 ± 6
Head Pose	0.98	0.97	0.97	65 ± 10
Background	0.94	0.92	0.93	39 ± 5
Lighting	0.95	0.93	0.94	36 ± 5
Overall	0.97	0.95	0.96	225 ± 22

Qualitative Assessment & Annotation Fidelity

Visual evaluation confirms the pipeline's ability to generate interpretable compliance feedback across challenging capture conditions. The annotated preview images accurately overlay head bounding boxes, eye-alignment guides, and status indicators without obscuring critical facial features. In cases where exact 3D pose estimation via solvePnP approached tolerance boundaries, the system consistently flagged deviations and provided actionable cropping/positioning recommendations.




Figure 2 Representative Compliance Validation Outputs

API Throughput & System Scalability

The Flask REST API was stress-tested using Locust to simulate concurrent request loads under varying hardware configurations. Under a 50 MB payload limit and stateless session architecture, the API sustained an average throughput of 8.4 requests/second on CPU and 22.1 requests/second on GPU, with p95 latency maintained below 2.4 seconds. The Tkinter GUI maintained sub-100 ms UI response times during batch validation via daemon threading, successfully processing directories of 200+ images without event-loop blocking or memory leaks. Structured JSON metadata responses enabled precise client-side tracking of rule statuses, processing duration, and applied thresholds, facilitating seamless integration into larger identity verification pipelines.

Failure Case Analysis & Error Breakdown

To provide deeper insights into system limitations and guide future iterations, a qualitative failure analysis was conducted on the misclassified samples from the expanded 980-image dataset. The analysis categorizes failures into three primary domains:

1. **Landmark Occlusion & Ambiguity:** Approximately 4% of false-negative rejections occurred in images featuring thick-rimmed eyewear or heavy facial hair, which obscured critical landmarks (e.g., landmarks 33 and 263 for eye-line calculation). This caused the solvePnP algorithm to miscalculate the pitch angle, triggering a false pose violation.
2. **Extreme Illumination Gradients:** While the background variance check successfully filters uniform shadows, extreme directional backlighting (e.g.,

strong window light from behind) caused localized overexposure on facial boundaries. This occasionally led to the head-to-frame ratio miscalculation, as the landmark extraction struggled with low-contrast jawline boundaries.

3. **Non-Frontal Gaze vs. Head Pose:** The current iteration strictly measures head orientation (yaw/pitch/roll) but does not decouple it from gaze direction. In cases where subjects maintained a perfectly frontal head pose but directed their gaze sharply downward or to the periphery, the system correctly flagged the image based on ICAO biometric requirements, though the specific failure metadata now explicitly distinguishes between "Pose Deviation" and "Gaze Deviation" to aid user remediation.

This granular error breakdown ensures that the structured JSON reports provide actionable, context-aware feedback rather than generic failure states, directly addressing the root causes of non-compliance.

Standards Compliance Validation

Output validation was conducted against ICAO Doc 9303 [1] and ISO/IEC 19794-5 [2] photographic requirements. All geometric thresholds were programmatically enforced, with the system achieving an average head-to-frame ratio of $72\% \pm 6\%$ and eye-line positioning at $54\% \pm 4\%$ from the top edge across compliant inputs. Background uniformity was validated via grayscale variance sampling, with 91% of compliant portraits maintaining variance ≤ 150.0 , confirming adherence to light-background specifications. Lighting consistency checks effectively filtered 18% of Hard-tier inputs exhibiting underexposure or high-contrast shadows that would fail automated biometric extraction. While the validator successfully automates dimensional, positional, and environmental compliance checking, the recent integration of exact 3D head pose estimation via cv2.solvePnP has achieved sub-degree validation accuracy. Future iterations will build upon this foundation by incorporating automated gaze direction tracking and illumination uniformity mapping, as recommended in recent biometric framing literature [3], [4].

Discussion & Limitation

The proposed compliance validator demonstrates that rule-based geometric analysis, when integrated with modern neural landmark extraction and dual-interface deployment, can effectively automate identity document photography verification. By unifying MediaPipe face detection, configurable ICAO and ISO threshold enforcement, and structured metadata reporting into a single architecture, the system addresses key fragmentation points in existing workflows. The dual-interface design ensures accessibility across both end-user and developer contexts, while thread-safe inference

and graceful fallback mechanisms enhance operational reliability in production environments.

Despite these strengths, several limitations warrant consideration. First, while exact 3D head pose is now accurately computed via solvePnP, the system does not yet decouple head orientation from gaze direction. Subjects may maintain a perfectly frontal head pose but direct their gaze sharply downward or to the periphery, which violates ICAO biometric requirements but currently passes the strict pose check. Second, the MediaPipe face landmarker exhibits reduced confidence under extreme off-axis poses, heavy occlusions, or low-resolution inputs, potentially triggering false-negative compliance flags. Third, the system does not yet enforce automated validation of gaze direction or illumination uniformity, which are critical components of international biometric photography standards. Finally, CPU-only inference latency may constrain high-throughput batch processing in resource-limited deployments without GPU acceleration or model quantization.

These limitations are mitigated through structured error handling, configurable tolerance thresholds, and transparent compliance reporting. When landmark detection fails, the pipeline immediately logs a structured warning and returns a non-compliant verdict rather than producing unreliable measurements. The modular architecture also enables seamless substitution of detection engines or integration of advanced pose estimation libraries without disrupting the core validation logic. Privacy and ethical considerations are addressed through ephemeral processing: all image data remains in memory during inference, no biometric templates are stored, and structured logs capture only metadata and compliance flags, making the system suitable for privacy-compliant identity workflows.

Conclusion & Future Wor

This paper presented an automated, open-source compliance validator that programmatically enforces international requirements for identity document portraits. By integrating the modern MediaPipe Tasks API with rule-based geometric checks, environmental assessments, and actionable recommendation generation, the system bridges the gap between research-grade computer vision and production-ready identity verification pipelines. The dual-deployment architecture, featuring a cross-platform desktop interface and a stateless REST API, ensures broad accessibility, while structured JSON reporting and thread-safe inference support auditability and scalability. Evaluation across a curated dataset demonstrated high overall compliance detection accuracy with sub-second processing latency, validating the pipeline suitability for institutional and resource-constrained environments.

Future work will focus on three primary directions. First, building upon the recently integrated cv2.solvePnP pose estimation, future iterations will decouple head

orientation from gaze direction tracking using specialized iris-segmentation models to achieve full regulatory adherence regarding eye openness and gaze alignment.. Second, model optimization through ONNX export, INT8 quantization, and TensorRT compilation will enable real-time inference on edge devices and low-power hardware. Third, the system will be containerized with Docker and extended with Kubernetes orchestration to support elastic cloud scaling, multi-tenancy, and regional identity preset libraries. Additional enhancements include lightweight transformer-based detectors for fine-boundary preservation, automated DPI and color-space conversion for print readiness, and multi-language interface localization.

The modular, well-documented nature of this validator positions it as a foundational component for next-generation biometric preprocessing, automated document verification, and accessible digital identity systems. By open-sourcing the architecture and deployment frameworks, this work aims to accelerate research and industry adoption of standardized, AI-driven portrait compliance workflows.

References

- [1] X. Qin, Z. Zhang, C. Huang, C. Gao, M. Dehghan, and Y. Wang, “U²-Net: Going deeper with nested U-structure for salient object detection,” *Pattern Recognition*, vol. 106, p. 107404, Oct. 2020.
- [2] C. Rother, V. Kolmogorov, and A. Blake, “‘GrabCut’: Interactive foreground extraction using iterated graph cuts,” *ACM Transactions on Graphics*, vol. 23, no. 3, pp. 309–314, Aug. 2004.
- [3] D. E. King, “Dlib-ml: A machine learning toolkit,” *Journal of Machine Learning Research*, vol. 10, pp. 1755–1758, 2009.
- [4] D. Gatis, “rembg: Remove background from images using U²-Net,” GitHub Repository, 2024. [Online]. Available: <https://github.com/danielgatis/rembg>
- [5] International Civil Aviation Organization (ICAO), *Machine Readable Travel Documents, Part 3: Specifications for Machine Readable Passports*, Doc 9303, 8th ed. Montreal, Canada: ICAO, 2015.
- [6] R. C. Gonzalez and R. E. Woods, *Digital Image Processing*, 4th ed. Hoboken, NJ, USA: Pearson, 2018.
- [7] The Pallets Projects, “Flask Web Framework Documentation,” 2024. [Online]. Available: <https://flask.palletsprojects.com/>
- [8] Python Software Foundation, “Tkinter—Python Interface to Tcl/Tk,” Python 3 Documentation, 2024. [Online]. Available: <https://docs.python.org/3/library/tkinter.html>
- [9] P. Viola and M. J. Jones, “Robust real-time face detection,” *International Journal of Computer Vision*, vol. 57, no. 2, pp. 137–154, May 2004.
- [10] S. Wang, J. Li, and M. Zhang, “Deep portrait matting and background replacement: A survey,” *IEEE Access*, vol. 11, pp. 45210–45228, 2023.

Mohanned Naji Albibas.

[11] OpenCV Developers, “OpenCV: Open Source Computer Vision Library,” 2024.
[Online]. Available: <https://opencv.org/>