ENTRUST IDENTITY VERIFICATION SERVICES

AI TRANSPARENCY NOTICE

Entrust Identity Verification Services Artificial Intelligence ("AI") Transparency Notice
This Notice provides transparent information about the AI system you are using, outlining the system's capabilities, limitations, biases, data usage practices, and safety measures. We want to empower you with the knowledge needed to understand how this AI works, what it can and cannot do, and how your data is handled. By providing this transparency, we enable you to make informed decisions about your interactions with our AI technology.

As you read this notice, keep in mind to always double check important information provided by our AI. Our technology is designed to assist, not replace expert judgment. For decisions that may have health, financial, or legal implications, we strongly recommend consulting appropriately qualified professionals or obtain independent legal counsel before taking action.

ENTRUST'S IDENTITY VERIFICATION (IDV) TECHNOLOGY

Entrust's IDV (Identity Verification) product enables clients to remotely process, verify and onboard users. The inputs to this system are a photo/video of an identity document and a photo/video of the applicant. Entrust's decision engine uses these inputs to return an accept/consider signal, together with appropriate context and detailed breakdowns on how that result was produced.

The IDV product is driven by four core technology segments, each of which relies on AI systems. The four segments are (1) **facial similarity**; (2) **biometric anti-spoofing and liveness detection** (facial images); (3) **automated data extraction** from documents; and (4) **anti-fraud** technology on identity documents. These segments fulfil the tasks of data extraction, biometric matching, document validation and biometric validation, all requiring a positive result for an accept signal to be returned. While the collection of multiple AI systems serving each of these technology segments together form the IDV application, it should be noted that the IDV application contains a substantial amount of non-AI based algorithms as well. It is the combination of AI/ML based algorithms with non-AI based algorithms, i.e., "traditional" computer code, that together creates Entrust's product.

INTENDED USE AND AI FEATURES

Entrust's responsible use of AI enables faster and more consistent data processing, resulting in a smoother and more secure experience for end users. Our systems are designed to minimize human review of personal data wherever possible, helping to reduce unnecessary exposure of end user information.

AI is used across our Identity Verification (IDV) product portfolio to:

- Indicate whether a photo or video of a government ID appears genuine or may show signs of manipulation and/or fraud.
- Classify the document type, for instance a passport or driver's license.
- Extract relevant information from the document image or video, such as name or date of birth
- Indicate whether the facial biometrics provided by an end user appear genuine and corresponds to the individual shown on the identity document.

Our solutions are designed to identify patterns that may indicate authenticity or potential fraud, helping organizations make informed trust and safety decisions. Results should be interpreted as risk indicators rather than definitive determinations and should be used in conjunction with other controls and verification procedures.

DATA TYPES

- Photo/video of a government ID
- Photo/video of facial biometrics

AI USED / MODEL TYPE

All applications use deep learning models, with specific paradigms depending on the application. These models are typically pre-trained then fine-tuned over our own datasets to optimize for performance (both accuracy and fairness).

- Binary Classification Models: Document fraud detection and biometrics liveness detection use binary classification models to assess the probability of the input image/video if being genuine (as opposed to being altered in any fraudulent manner, e.g. being a deepfake, or tampered with some kind of deliberate obstruction).
- Multi-Class Classification Models: Document classification uses a multi-class classification model to return the most probable document type given an input document image.
- Vision-Language Models (VLM): Document extraction uses a vision-language model (VLM) for document intelligence
- Face Recognition Models: Our application employs supervised deep learning models to perform facial comparisons between live-captured selfie or video and an image sourced from a government-issued identity document, or other previously captured selfies or videos. The model produces a numerical representation ("embedding") of each face image and measures the similarity between embeddings to determine the likelihood that both images belong to the same individual, otherwise known as an embedding-based, 1:1 or 1:N biometric verification.

HOW WE TRAIN OUR AI

Our AI models are optimized towards performance (both accuracy and fairness) using a systematic benchmarking protocol. They are trained and fine-tuned, then thoroughly benchmarked against datasets collected from end user data. We use **Quality Check (QC)** data uniformly sampled from production to obtain unbiased estimates of the true false acceptance/rejection rates of a given model. Additionally, we employ a feedback process called **Human in the Loop (HITL)**, whereby human analysts evaluate rate outputs to continuously improve model performance in accordance with Entrust's responsible AI principles.

RESPONSIBLE DATA USE PRINCIPLES

To help you understand how we handle your data when you use Entrust's AI-powered system, please consider the following key points:

- Our use of AI involves the processing of end user personal data. We process such personal data only where we have a valid legal basis for doing so. To learn more about our use of personal data, visit the IDV Product Privacy Notice.
- We restrict access to our AI models and related data to only those individuals who require it for their specific job responsibilities.
- We encrypt the data in transit and in storage.

BIAS

POTENTIALLY INHERITED BIASES

Like most AI systems trained on real-world data, our models may reflect inherited biases from the datasets used to train and improve it. The diversity of training data varies across countries and demographic groups, which may lead to differences in performance for some populations or document types. We actively monitor these factors and apply measures such as data rebalancing, retraining, and fairness testing to reduce unintended differences in performance across user groups.

Potential sources of inherited bias include:

- **Training Data Composition**: Publicly available datasets used in early model development may overrepresent certain regions, facial characteristics, or imaging conditions.
- **Image Quality & Document Design**: Variations in photo quality, printing standards, holograms, or security features may affect performance across document types.
- **Capture Conditions**: Differences in lighting, camera devices, or how images are collected can influence model outcomes across environments.

BIASES MITIGATION EFFORTS

Entrust's AI systems are developed and continuously refined to minimize unintended bias and promote fair outcomes across demographic and geographic groups. To mitigate bias, we apply:

- **Balanced and adaptive sampling techniques** during model training to ensure diverse regional and demographic representation.
- **Regular performance evaluations** across countries, continents, and genders to identify and address disparities.
- **Dynamic sampling adjustments** that increase training emphasis on underrepresented groups to improve accuracy.
- **Fine tuning on real-world domain specific data** (such as Selfie to document comparisons) to strengthen model generalization across different image types and conditions.

These measures significantly reduce performance differentials; however, no AI model is entirely free from bias. We remain committed to transparency, continuous evaluation, and ongoing improvements as part of our responsible AI development process.

For more information on our commitment to developing innovative technology responsibly, you can refer to our ICO Sandbox Report.

LIMITATIONS

Our AI may be subject to certain limitations:

- **Scope**: Our AI is specifically designed for fraud prevention
- **Language**: Our AI supports multilingual inputs based on languages that use Latin and Arabic scripts
- **Context:** Our AI evaluates inputs as either fraudulent or genuine
- **Hallucinations**: In limited cases, our AI may generate inconsistent results if information is inaccurately extracted from a source document.

HUMAN OVERSIGHT / HUMAN IN THE LOOP

Entrust provides identity verification services for customers utilizing a Human in the Loop (HITL) model and is central to how our AI systems operate. Through this process, trained analysts regularly review, validate, and refine model outputs to improve its performance and alignment in accordance with Entrust's responsible AI principles.

Before models are released to production, experts evaluate its performance using analytic dashboards and controlled A/B testing to confirm that results meet our standards for quality and fairness. Human reviewers also oversee data labeling and verification, ensuring that the training data used to improve our systems is accurate and representative. In certain products, human specialists remain available in real time to assess results when the AI detects unusual or uncertain inputs.

FEEDBACK

We welcome your feedback for the purpose of improving our AI systems. To provide details on potential inaccuracies or to share your experience, email us at AI@Entrust.com.

CUSTOMER CONTROLS

We believe customers should have full visibility and control over their data, helping them stay transparent with their end users about how their data contributes to improving our services.

Our identity verification (IDV) systems use machine learning to enhance accuracy, detect anomalies, and improve overall performance. Participation in these processes is entirely at the customer's discretion. Our IDV operates on an **opt-out** basis. When a client opts out, their checks are excluded from any machine learning activities, including model training. Customers have full control over whether end user data may be used to help develop and refine our services and can customize at any time through account configuration.

When a customer opts out of machine learning, they are also automatically excluded from our **Quality Control (QC)** process. As part of our standard practice, approximately three (3%) percent of all checks are randomly reviewed each month to validate accuracy and support ongoing improvements. These QC insights help strengthen the reliability of our models, but opting out removes the customers' checks from this feedback cycle.

Even when customers choose not to participate in machine learning or QC, our systems *may process* **pseudonymized and aggregated data** for legitimate business purposes such as analytics, benchmarking, and statistical reporting. Aggregated data **does not** include any personally identifiable information (PII).

CONTINUOUS IMPROVEMENT

We treat bias monitoring and fairness evaluation as a continuous process. Our AI models are regularly reviewed and updated to enhance performance and address issues. When disparities are identified, we adjust our training data, sampling strategy, or calibration parameters to improve balance and accuracy.

These efforts are part of our broader commitment to responsible AI development, ensuring that our verification technology remains as accurate and transparent as possible across all user populations. We update this notice annually or when significant changes occur to ensure you have the most current information.

Last Updated: October 28, 2025