Face Recognition Product Suite

paravision

From Cloud to Edge, Paravision's comprehensive face recognition product suite offers all of the tools necessary to develop and deploy mission-critical biometric identification and authentication solutions across a wide range of applications, from identity verification to physical access control to air travel and government services.

Paravision face recognition has been optimized for leading chipsets, from Intel CPUs to NVIDIA GPUs to Ambarella SoCs, supports leading operating systems including Windows, Linux, and Android, and can be deployed at any level, from SDK to API-driven Docker containers.

Paravision has been repeatedly recognized by NIST FRVT as a top global provider and the most accurate U.S.-based face recognition technology provider across all use cases, including 1:1 verification and 1:N identification, including multi-million record databases, performance by age, performance with face masks, and the full range of image quality characteristics, from passport-quality to fully unconstrained images.

In short, Paravision powers its partners with world-class face recognition technology in a way that fits their technical expertise and deployment goals, supporting transformative solutions for the next generation identity, security, efficiency, and user experiences.

Face Recognition SDKs - From Cloud to Edge

For developers looking to deeply integrate Paravision face recognition into software applications, solutions, and services, Paravision offers a rich, multi-platform series of face recognition SDKs. For cloud, on-premises, and desktop applications, Paravision offers C++ and Python SDKs for Windows and Linux. For high-speed, high-precision Edge AI applications, Paraivison offers its SDK for Ambarella CVflow. And for mobile- or kiosk-based application development, Paravision offers its SDK for Android. All of these SDKs offer full face recognition pipeline functionality, including face detection, face landmark analysis, image quality analysis, template (embedding) generation, and 1:1 or 1:N matching.

To implement optimal systems architectures, developers can pick and choose aspects of Paravision's SDKs across platforms. For instance, face detection and image quality analysis could be used for mobile registration, passing cropped, quality-checked faces to backend systems for registration. Templates (embeddedings) can then be created on the server and distributed to networked Edge devices, which could then be used for full 1:N matching. Alternatively, the full face recognition pipeline could be deployed in each of these cases - whatever is appropriate for the application needs is achievable.

Face Recognition Engines - Rapidly Deployable, Massively Scalable

For developers looking to access Paravision face recognition functionality at a higher level, Paravision offers a series of Docker container-based face recognition engines. With all required face recognition capability fully integrated, these engines can be rapidly integrated into many standard cloud- or on-premises systems architectures, and are accessible by either REST or gRPC (and associated client libraries).

To enable face recognition at any scale, Paravision now offers Scaled Vector Search (SVS), a next-generation face recognition search engine capable of scanning databases of tens of millions of records, hundreds of times per second. Like Paravision's other engines, Paravision SVS is Docker container-based, fully modular, and cloud-ready while fully deployable on-premises.

Solutions for Video - Integration with Standard Video Security Systems

Paravision face recognition can integrate natively with standard IP video infrastructure using the Paravision Streaming Container. Accurate, efficient, and highly scalable, the Paravision Streaming Container is Docker-based and easily deployable with the Paravision Face Recognition SDK or any of Paravision's face recognition engines, including SVS. For more detail, please see the Paravision Streaming Container Data Sheet.

Presentation Attack Detection - Liveness and Anti-Spoofing

Ensuring the authenticity of biometric samples is critical to high confidence biometric identification and authentication in unattended or fully automated use cases. To protect against face recognition spoofing, Paravision offers an advanced Presentation Attack Detection (PAD) software toolset for both PC-based and embedded applications. Leveraging the latest advances in multispectral and 3D imaging, Paravision PAD is optimized for leading Ambarella- and Intel RealSense-based camera solutions. Paravision PAD is fully passive and operates in real-time, enabling advanced protection without compromising user experiences.

System Architecture

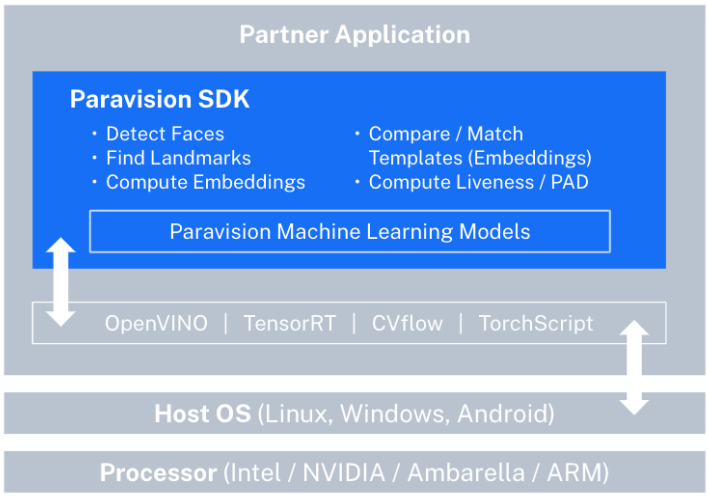

Face Recognition SDKs

System Architecture

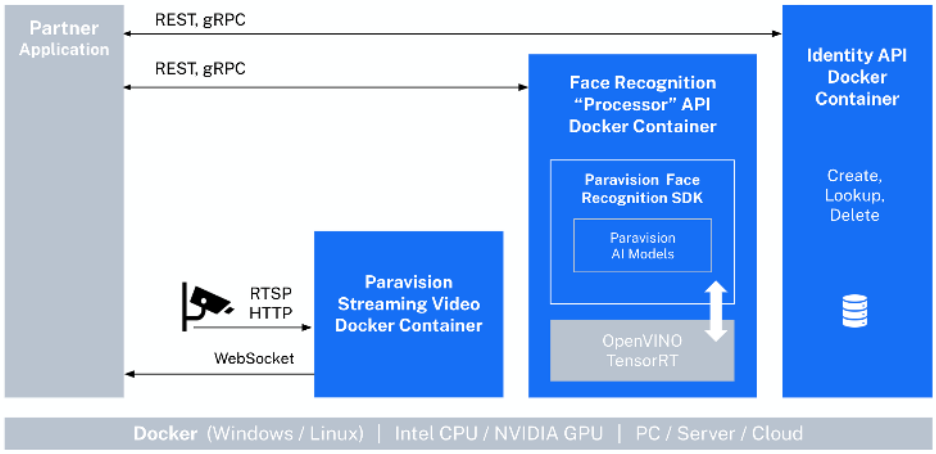

Face Recognition Docker-Based Engines

Technical Specifications

Face Recognition SDKs

Supported operating systems	Linux, Windows, Embedded Linux, Android
Programming languages	Linux / Window: C++, Python Embedded Linux: C++ Android: Kotlin
Supported compute platforms (Computer vision frameworks)	Intel CPU (OpenVINO) NVIDIA GPU (TensorRT) Ambarella (CVflow) ARM (TorchScript)
Supported functions	Face detection / bounding box detection Face landmarks detection Image quality analysis Template (embedding) generation 1:1 verification 1:N identification Presentation Attack Detection (Requires specific 3D / NIR cameras)

Technical Specifications

Face Recognition Engines

Deployment method	Docker container, supporting on-premises or cloud-based computing
Supported operating systems	Windows, Linux
Supported compute platforms (Computer vision frameworks)	Intel CPU (OpenVINO) NVIDIA GPU (TensorRT)
Face recognition engine APIs	REST gRPC, with Clients supporting Python, C++, C#, Node.JS, GoLang, Java, Ruby
Streaming engine protocols	RTSP, HTTP(s)
Supported functions	Face detection / bounding box detection Face landmarks detection Image quality analysis Template (embedding) generation 1:1 verification 1:N identification, including Scaled Vector Search for very large database sizes
Scaled Vector Search database size	Up to and greater than 10M+ records
Scaled Vector Search speed	10M record search in 50 msec, single Intel Xeon compute instance