AI Hub:

Accelerating AI Adoption Across the Enterprise

technical brochure

Table of Contents

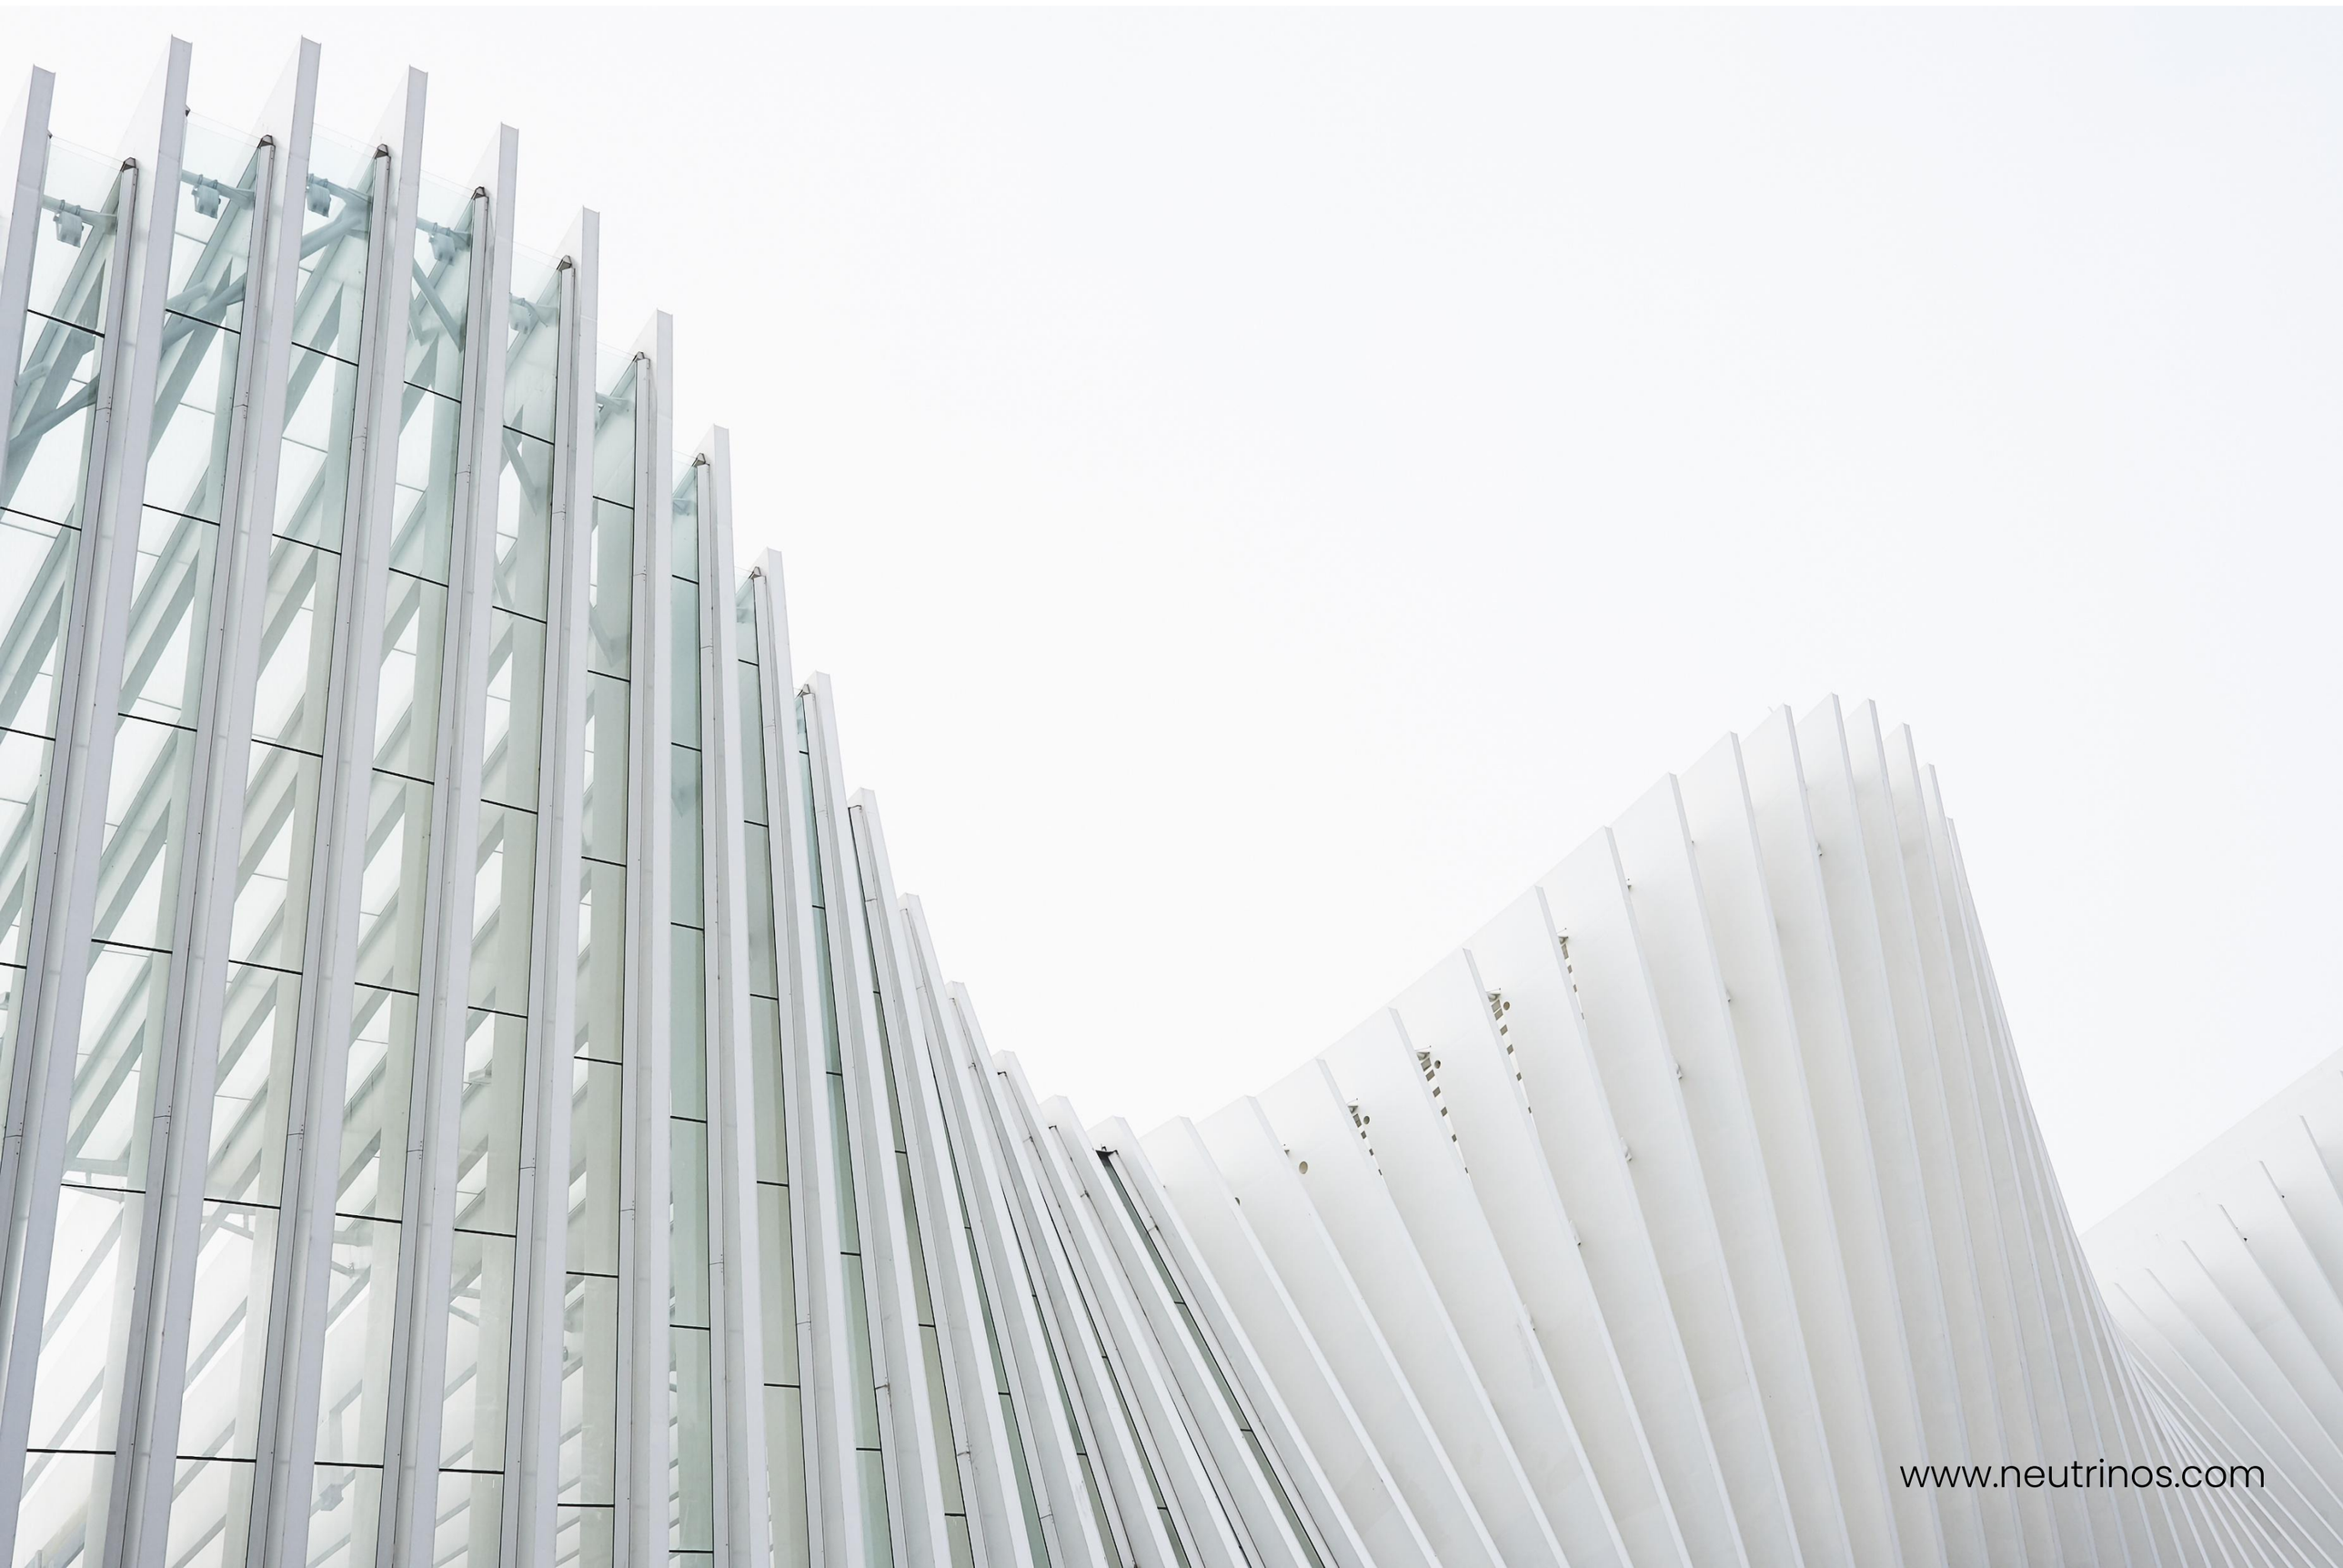

AI Hub: Overview

AI Hub on the Neutrinos platform is a comprehensive framework that facilitates the seamless integration of Artificial Intelligence (AI) and Machine Learning (ML) capabilities into enterprise applications. It empowers businesses to embed AI-driven functionalities— such as Natural Language Processing (NLP), Generative AI (GenAI), predictive analytics, and automation—without requiring deep technical expertise.

The AI Hub offers a wide range of utilities, including **ready-to-use AI models** that can analyse and process both text and images. These models can be easily integrated into existing applications, enabling rapid adoption of AI features.

In addition, technical teams have the flexibility to build and **train custom AI/ML models** directly on the platform, leveraging RESTful APIs to integrate AI as a Service into their applications.

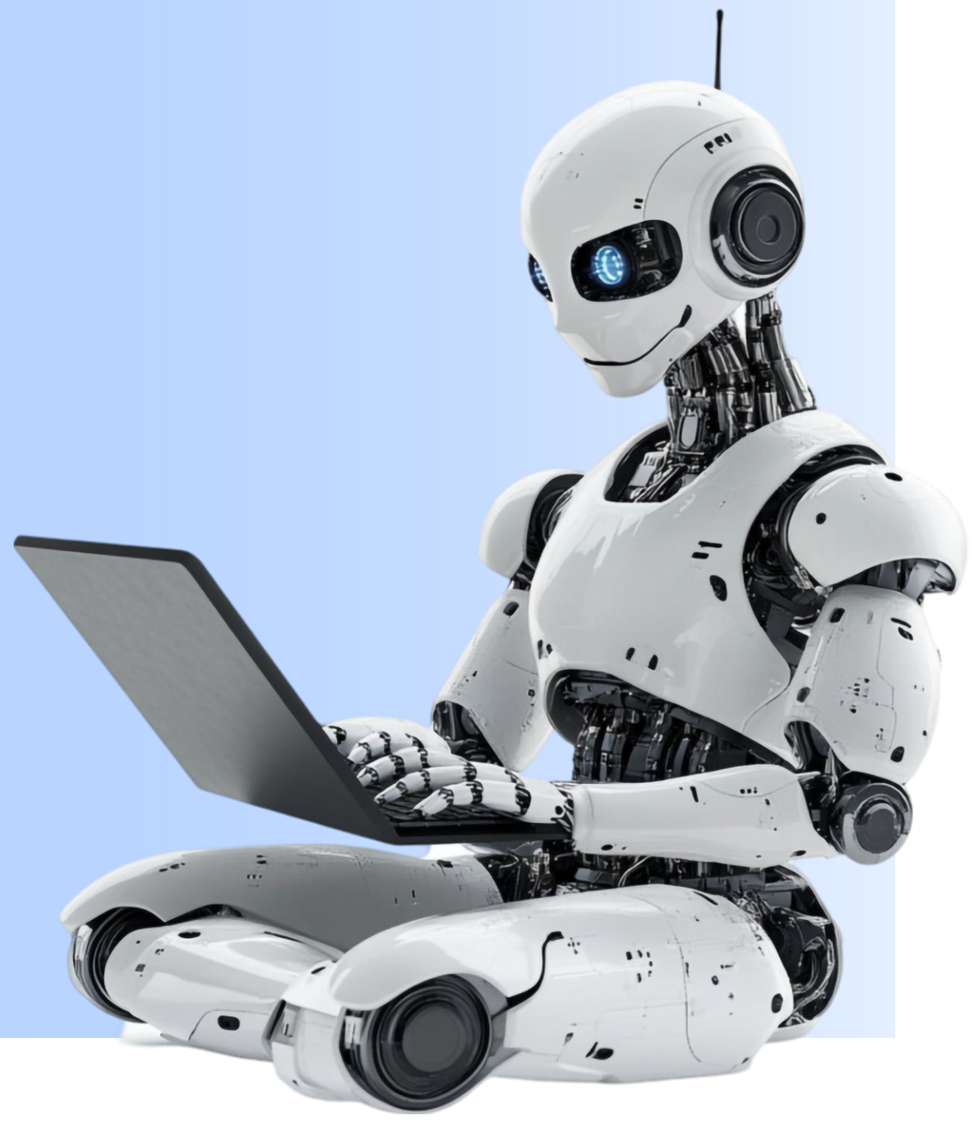

By integrating AI Hub into existing business workflows, organizations can automate decision-making, detect fraudulent activities, and minimize manual interventions.

The platform is designed to scale with business needs, making it suitable for organizations of all sizes.

Models

AI Hub offers the capability to create and support a wide range of models for various AI tasks, including text classification, text extraction, document classification, and document extraction. Each model is organized into quality tiers—Medium, High, and Best—allowing users to select or build the most appropriate option based on performance requirements and available resources.

Text Classification & Extraction

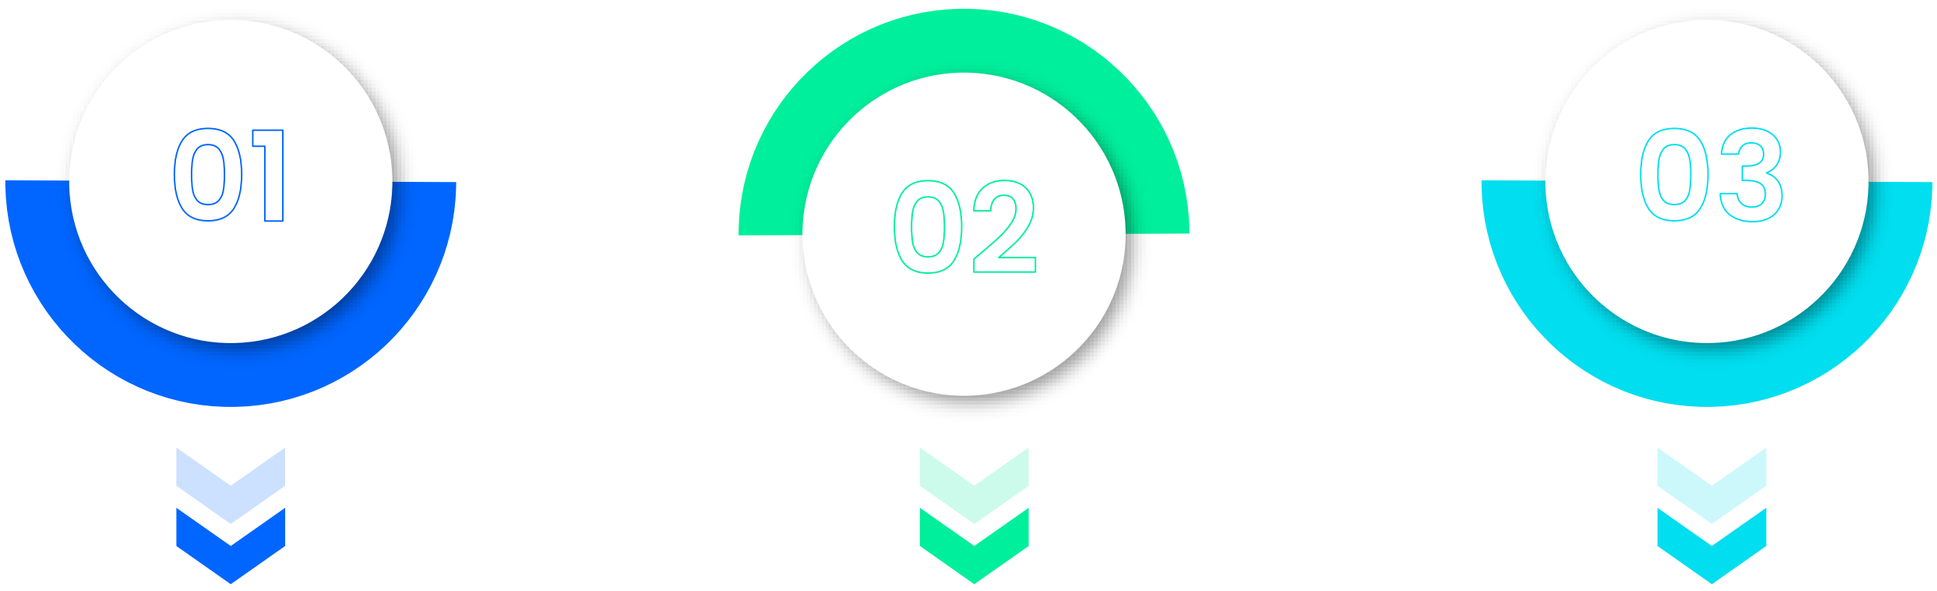

Best

google/electra-base-discriminator

High

lightGBM, catBoost, XGBoost, random forest

Medium

lightGBM, catBoost, random forest

Document Classification

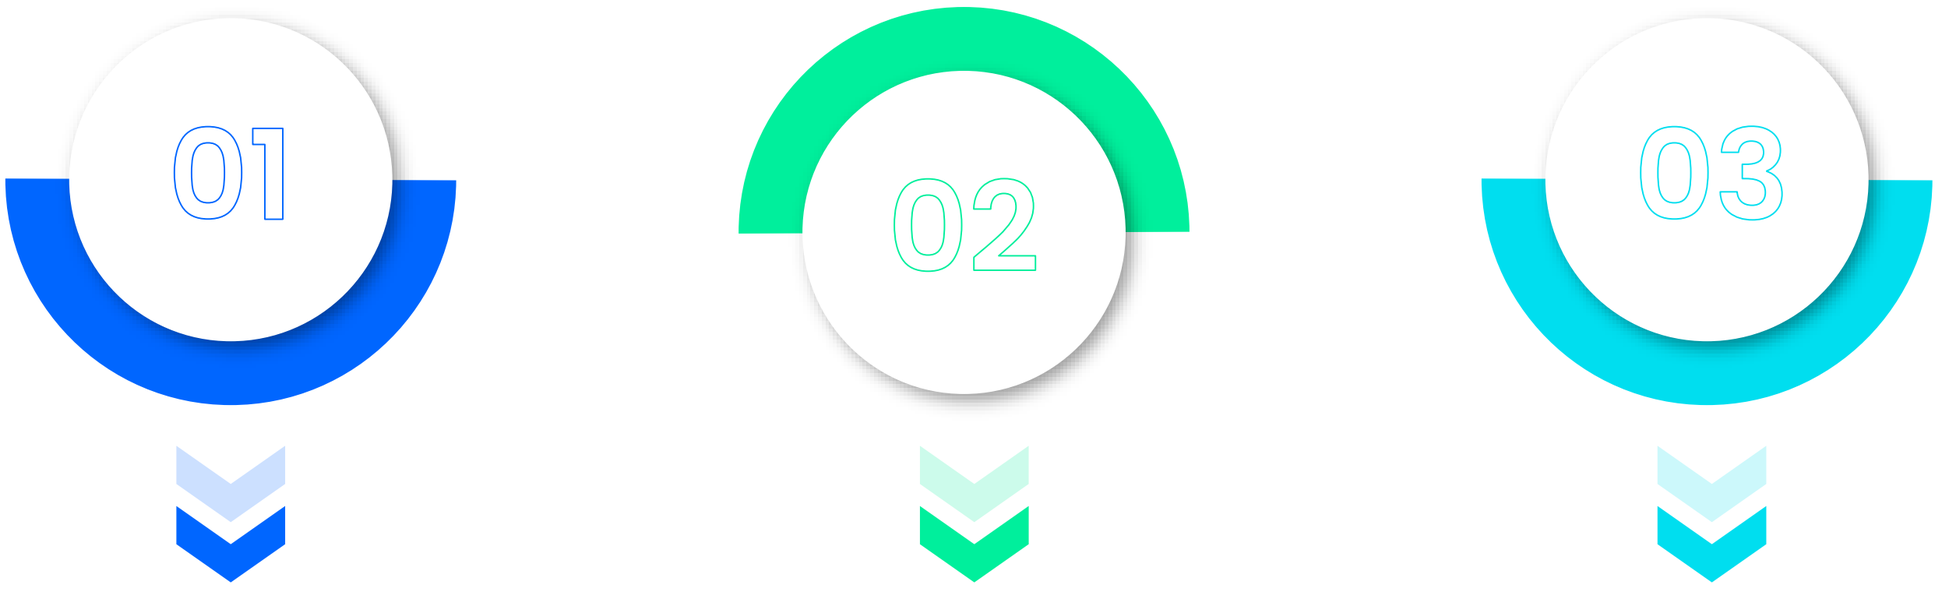

Best

timm_image with advanced backbones such as swin_large_patch4_window7_224 **or** vit_base_patch16_224

High

timm_image with checkpoint caformer_b36.sail_in22k_ft_in1k (enhanced security)

Medium

timm_image with checkpoints such as resnet50 **or** caformer_b36.sail_in22k_ft_in1k (Optimized for speed & size)

Document Extraction

01	02	03
Best	**High**	**Medium**
yoloV3_mobilenetv2_8xb24-320-300e_coco (YOLO Models)	yolox-l	yolo-s

Assistant Model

TwelveLabs · Meta · LUMA AI · OpenAI

Azure · stability.ai · cohere · ANTHROP\C

Nova · MISTRAL AI_ · deepseek · Writer AI

In addition to the listed models, the platform offers advanced configuration options that allow users to fine-tune models by customizing hyperparameters. This enables greater control over model performance, ensuring alignment with specific business and data requirements.

Data Retention

In AI Hub, although organizational data is hosted on Neutrinos servers, it is securely containerized within dedicated tenant spaces. This architecture ensures data isolation, enhances security, and supports compliance with industry and regulatory requirements.

The data retention on the platform happens through 3 broad categories:

1. **Training data associated with specific prediction** (classification), extraction models, or even the Assistant (chatbot – generative model) is retained only as long as the corresponding models exist on the platform. Once a model is deleted, all related data and artifacts are purged. Neutrinos ensures that no residual or derivative data remains in any shared systems, maintaining strict data privacy and integrity standards.

2. **Inference data is governed by configurable retention rules** at individual model level. These rules can be defined under the Retention and Others section at the time of training. Based on these configurations, inference data can be automatically purged, offering organizations control over data lifecycle management.

3. Additionally, **audit log data**—which includes transactional details such as transaction IDs, timestamps, and other model-related metadata (excluding any organization-specific information)—is retained even after the associated models are deleted or the subscription is terminated.

Data Privacy

While the platform leverages Large Language Models (LLMs), Optical Character Recognition (OCR), and generative AI (such as the Assistant chatbot) for extraction, it also incorporates robust data privacy controls.

A key feature is PII Masking, which ensures that personal, organizational, and sensitive information is masked before being transmitted to LLMs. This prevents unintended data exposure during model inference, ensuring that any sensitive content—whether extracted from documents or entered as user input—is protected throughout the model execution lifecycle.

Intelligence Assured

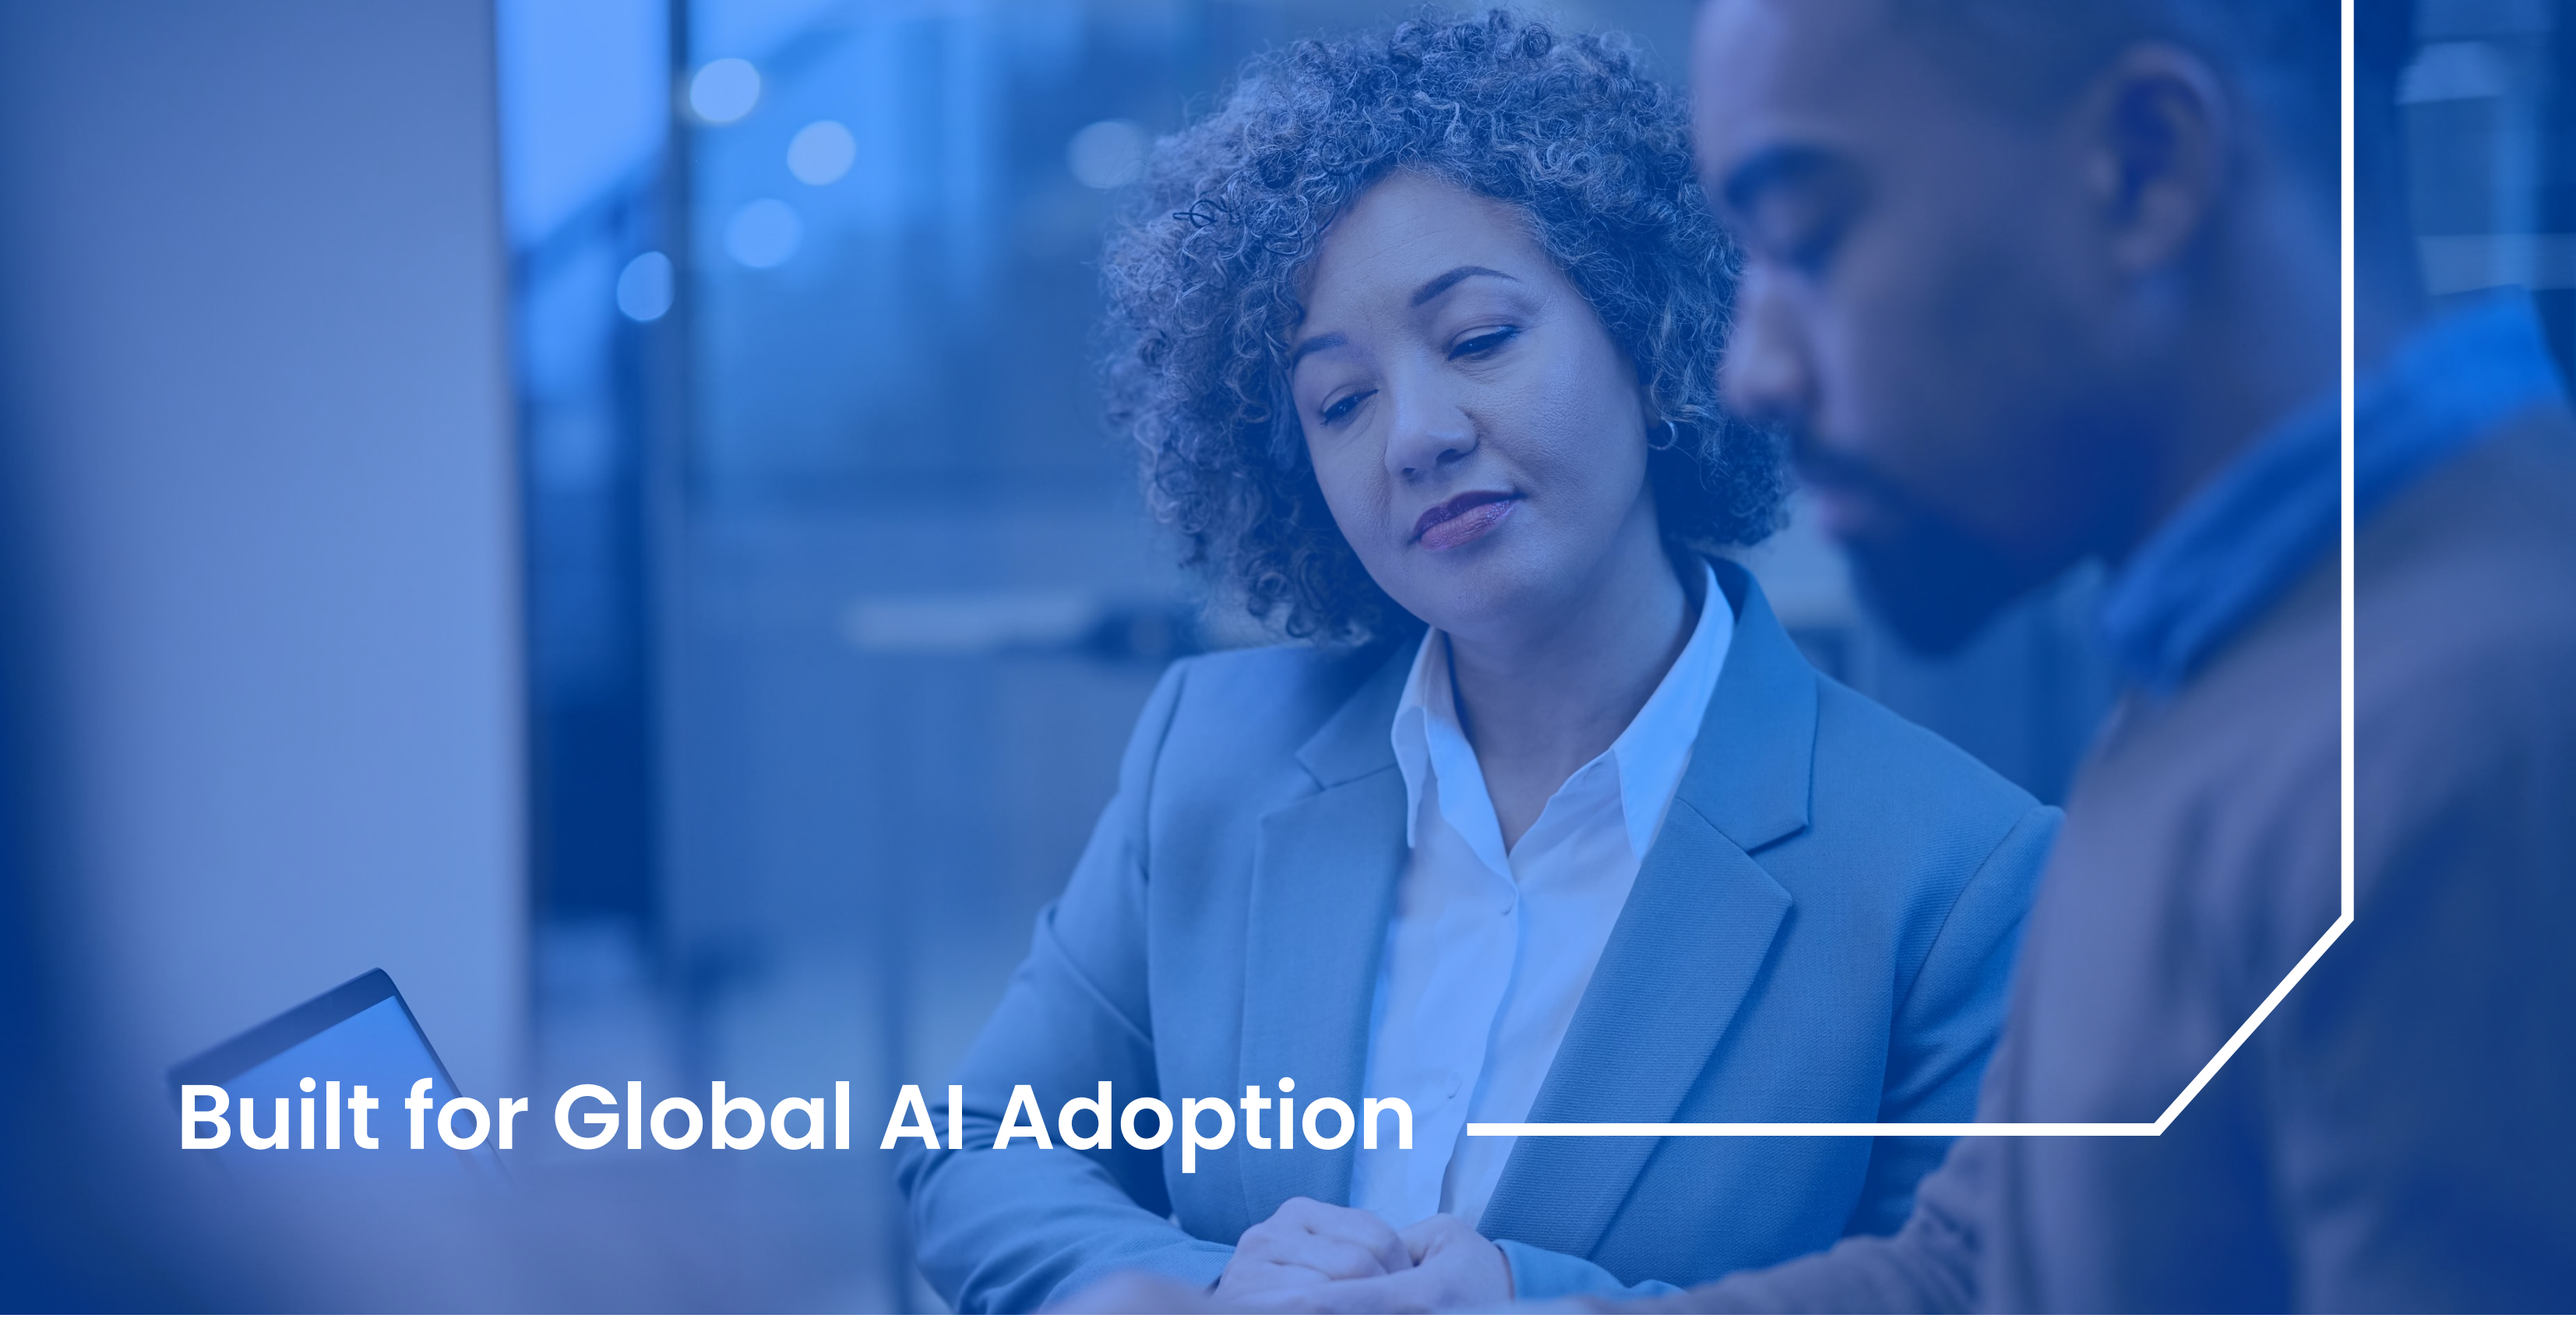

Built for Global AI Adoption

Summary

In conclusion, the AI Hub on the Neutrinos platform is a **robust, scalable framework** that enables the seamless integration of AI, ML, and Generative AI (GenAI) capabilities—including assistants and agent-based functionalities—into enterprise applications. It elevates AI from a standalone technology to a fully managed, service-oriented solution. The platform supports a comprehensive range of models for text and document classification and extraction, organized into quality tiers to meet varying performance and resource requirements.

Business users can **build, deploy, and manage models** with ease, without requiring deep technical expertise. Data is securely containerized within tenant-specific environments, governed by configurable retention policies and protected through PII masking to ensure compliance with **privacy and regulatory standards**. The platform also enables enterprise-grade automation, governance, and workflow orchestration.

Furthermore, AI Hub extends its utility to global use cases with advanced multilingual and multimedia support, including **built-in language translation** capabilities within models. This makes AI Hub as a future-ready solution for organizations looking to scale AI adoption across diverse domains and geographies.

comms@neutrinos.co

Neutrinos
Reinvent. Accelerated.