Noctua AI — Technical Product Sheets

Welcome to the Noctua AI Technical Product Sheets. This document outlines our modular, compliant, and custom AI solutions engineered for high-stakes sectors including Defense, Healthcare, Finance, Smart Cities, and Industrial Manufacturing.

NoctuaEdge™ — Real-Time Edge AI Platform

A lightweight, real-time AI deployment framework built for disconnected, low-latency, and high-risk environments. Supports embedded AI on drones, wearables, and IoT devices. Deployable across ARM, WebAssembly, and Linux OS.

Key Features:

- - Inference latency <10ms on edge devices
- - Supports ONNX, TensorFlow Lite, and PyTorch Mobile
- - Zero-connectivity inference with secure fallback modes
- - Integrated model distillation and compression pipeline

ForgeAlign™ — Responsible AI Alignment Toolkit

An advanced model alignment and behavioral tuning suite leveraging human-in-the-loop feedback, ethical guardrails, and reinforcement learning from human preferences (RLHF).

Key Features:

- - Ethical alignment dashboard and behavioral tuning
- - Support for transparency protocols (model explainability, saliency maps)
- - Bias detection & mitigation engine
- - Continuous RLHF training support with ethical override logic

NeuroSafety™ — AI Risk & Compliance Platform

A full-stack governance layer for monitoring, logging, auditing, and validating AI behavior in regulated industries. Supports real-time explainability, regulatory reporting, and data traceability.

Key Features:

- - Audit trails for decision history and model outputs
- - Compliance alignment with HIPAA, GDPR, EU AI Act, ISO/IEC 42001
- - Role-based access and review workflows
- - Real-time monitoring of model confidence and safety thresholds

NoctuaOps™ — MLOps & Custom AI Engineering

Custom-built MLOps solutions and fully tailored AI applications for enterprise deployment. Scalable from R&D to production-grade systems with lifecycle management.

Key Features:

- - CI/CD pipelines for model updates and drift monitoring
- - Multi-cloud and on-premise deployment support
- - Collaborative model governance & stakeholder approval flows
- - Custom UIs and dashboards for operational AI use