Felix Schaller

AI Strategy, Safety and Technical Due Diligence

Advisory for regulated, high-stakes and safety-critical AI systems

Service Sheet Overview

This document outlines how I support organizations, investors and engineering leadership across four decision layers: strategic AI readiness, technical validation, execution governance and autonomy safety research.

Executive Summary

Most organizations do not fail with AI because they lack ambition. They fail because they invest in systems that are not technically feasible, not governable, not certifiable, or not aligned with operational constraints.

I advise companies, investors and engineering leadership on how to evaluate, de-risk and structure AI initiatives before capital, engineering resources or reputation are committed.

I help decision-makers distinguish between AI that is impressive in a demo and AI that can survive engineering reality.

Core Advisory Services

AI Strategy and Feasibility Advisory

Assess what AI can realistically do, where hidden risks exist, and whether an organization should build, buy, partner or wait. Deliverables include feasibility assessment, use-case prioritization, architecture constraints and an executive decision memo.

Technical Due Diligence

Independent assessment of AI, autonomy or deep-tech companies for investors, consultancies and corporate strategy teams. Outputs include red flags, maturity review, roadmap credibility, dependency analysis and technical risk summary.

SOTIF / Autonomy Safety Risk Review

A 3-5 day sprint for ADAS, AV, perception, localization and sensor-fusion teams under ISO 21448 / SOTIF and ISO 26262 constraints. Outputs include risk register, validation strategy and evidence roadmap.

Architecture Review for Regulated AI

Review whether AI platforms, autonomy stacks and safety-critical systems support reliability, traceability, verification and governance. Outputs include subsystem dependency maps, validation recommendations and roadmap adjustments.

Interim / Fractional AI Safety Leadership

Temporary or fractional leadership for AI safety, autonomy governance, certification readiness, safety case structuring, audit preparation and investor-facing technical communication.

Engagement Formats

Advisory Sprints	2-10 days for executive risk briefings, feasibility assessments and targeted technical reviews.
Interim Mandates	4-12 weeks for program stabilization, technical handover, safety case structuring and stakeholder alignment.
Fractional Support	1-3 days per week for longer-term architecture governance, AI safety leadership and certification strategy.

Relevant Expertise

- ISO 26262 and ISO 21448 / SOTIF
- AI safety and certification feasibility
- Safety-critical software and toolchains
- Autonomy, ADAS and perception systems
- SLAM, localization and sensor fusion risk
- LiDAR / camera perception pipelines
- Model-based engineering and formal methods concepts
- Technical due diligence and architecture validation
- Regulated AI architectures and governance
- Non-LLM, declarative and verifiable AI systems
- Audit readiness and evidence generation
- Build-vs-buy and vendor risk evaluation

Selected Background

Felix Schaller is a senior AI and software safety expert with more than 15 years of experience across automotive, aerospace, defense, simulation and regulated engineering environments.

He has contributed to ISO 26262 toolchain qualification, SOTIF-related autonomy safety concepts, perception-system risk analysis, evidence generation strategies and model-based approaches to constraining AI behavior.

As founder of XIXUM, he develops non-LLM, declarative cognitive AI systems designed for verifiable reasoning and traceable decision-making in high-stakes environments.

His published research includes work on semantic models for constraining pattern recognition in modern AI systems.

Ideal Clients

- AI and autonomy startups preparing for fundraising or enterprise deployment
- Investors evaluating AI, autonomy or deep-tech companies
- Corporates building AI strategies in regulated or high-risk sectors
- ADAS / AV teams preparing SOTIF or ISO 26262 activities
- Consulting firms needing independent technical depth
- Public-sector or GCC organizations developing AI governance and safety programs

Engagement Philosophy

AI strategy should not begin with hype, vendors or model selection. It should begin with constraints: what must be explainable, deterministic, verifiable and governable - and which failure modes are unacceptable.

The goal is not merely to adopt AI. The goal is to build AI systems that can be trusted, governed and deployed.

AI Strategy Maturity Matrix

Model-based assessment of AI readiness, dependencies and strategic execution capability

Many organizations approach AI strategy by asking which tools, vendors or models they should adopt. The more important question is whether the organization is mature enough to deploy AI reliably, responsibly and economically.

FelixSchallerCOM provides a model-based AI Strategy Maturity Matrix to assess how prepared an organization is for AI adoption, scaling, governance and risk-controlled deployment.

The approach is methodically inspired by focus area maturity model research, including work associated with Prof. Sjaak Brinkkemper and the Privacy-by-Design Maturity Model research at Utrecht University. The aim is not to copy a privacy model, but to adapt the maturity-matrix logic - focus areas, dependencies and improvement paths - to AI strategy.

What We Assess

Data Readiness

Quality, ownership, structure, semantic consistency and operational accessibility of data.

Process Readiness

Decision points, automation potential, process clarity and operational dependencies.

Governance and Compliance

Policies, accountability, auditability, regulatory constraints and risk ownership.

Technical Architecture

System boundaries, integration capability, data pipelines, model lifecycle and scalability.

Use Case Maturity

Business relevance, feasibility, economic value, measurable outcomes and operational fit.

AI Risk and Safety

Failure modes, controllability, explainability, validation strategy and human oversight.

Organizational Capability

Expertise, stakeholder alignment, vendor dependency, change readiness and execution capacity.

Model-Based Evaluation

Unlike simple AI-readiness checklists, the maturity matrix treats AI strategy as a system of dependent conditions. A high score in one area does not compensate for missing prerequisites in another.

Example dependency logic	AI use cases depend on data quality. Automation depends on process clarity. Governance depends on accountability and auditability. Certification feasibility depends on architecture and evidence generation. Scaling depends on infrastructure and organizational capability.

Deliverables

- AI maturity assessment across core focus areas
- Dependency map of blockers and enabling conditions
- Maturity matrix with current and target states
- Prioritized improvement roadmap
- AI use-case feasibility ranking
- Build-vs-buy and vendor risk recommendations
- Executive summary for leadership, investors or board-level decision-making

Outcome

The goal is to show where an organization currently stands, which AI initiatives are realistic now, which initiatives are blocked by missing prerequisites, where investment creates value and where AI adoption introduces unacceptable risk.

AI strategy should not begin with model selection. It should begin with maturity.

Research note: This adaptation references the logic of focus area maturity models and Utrecht University's Privacy-by-Design Maturity research as a methodological inspiration, not as a one-to-one reuse of privacy assessment metrics.

SLAM, SafeWahr and SOTIF Research

Autonomy safety architecture, perception validation and featureless motion-estimation research

SAFEWAHR RESEARCH CONSORTIUM CONTRIBUTION

SafeWahr was a multi-institutional research project for SOTIF-aligned runtime safety monitoring of AI-based perception systems. My contribution focused on safety argumentation, SOTIF compliance process design, runtime monitor architecture and deterministic perception research.

Safety Argumentation and SOTIF Compliance Strategy

Developed the safety argumentation framework and SOTIF compliance process for Validas AG, including the transition from ISO 26262 functional safety into the SOTIF domain. Covered evidence generation, residual-risk argumentation and audit readiness.

Perception Architecture and Deterministic Algorithm Research

Contributed to runtime monitor architecture and conducted independent research into deterministic perception algorithms, including featureless motion estimation as an alternative to feature-based Kalman tracking in repetitive-pattern environments.

RUNTIME MONITOR ARCHITECTURE

Function Monitor

Validates perception output against expected operating parameters. Detects when model outputs violate performance bounds or consistency constraints.

Situation Monitor

Assesses whether the current scenario remains within the trained distribution. Flags Unknown-Unsafe conditions before they reach the decision layer.

Validity Monitor

Checks sensor inputs against the valid operational domain. Detects distribution shift, sensor degradation and environmental edge cases in real time.

IMPLICIT KALMAN AND FEATURELESS MOTION ESTIMATION

Independent research into featureless motion estimation as an alternative to classical feature-based Kalman tracking. Classical methods extract features and track them with Kalman filters, which can create systematic errors in repetitive or homogeneous environments such as road surfaces or structured SLAM scenes.

Approach

Hierarchical phase-shift analysis using Gabor filters, wavelet decomposition and FFT frequency decomposition across a resolution pyramid. Coarse motion is estimated first; higher resolutions register relative deviations.

Why it matters

Motion is registered via phase shift rather than feature displacement. This reduces dependency on contrast, feature extraction and repeated visual patterns at fine scale.

Current status

Proof of concept established. Steerable-filter implementation is in progress; IFFT aliasing at high frequencies is under investigation. Not yet production-ready; publication planned after completion.

> *Autonomy safety is not a checklist. It is a system of dependent evidence conditions. Sound safety argumentation, runtime monitoring architecture and deterministic perception alternatives together form a defensible foundation for regulated autonomous deployment.*