



Chair of Mobile Business & Multilateral Security

What You Can't See Can Hurt You!

Governance of the Invisible: Mobile Applications, AI, and Trust

January 13th, 2025 – Frankfurt

Derya S. Esen
External PhD Candidate

Goethe University Frankfurt
www.m-chair.de



Agenda

- 1 Mobile Apps: The Hidden Iceberg of Business Risk
- 2 The "Black Box" Paradox
- 3 Building Trust in the Black Box
- 4 The AI Regulation Tsunami

Mobile Apps: The Hidden Iceberg of Business Risk

Understanding the real architecture behind mobile business applications and where the critical risks actually live.



The top corners of the page feature decorative geometric shapes. In the top-left corner, there are two overlapping squares: a larger light blue one and a smaller, darker blue one. In the top-right corner, there are two overlapping squares: a larger light blue one and a smaller, darker blue one. The main content area is white with a thin horizontal line separating it from the top decorative elements.

The Iceberg of Mobile Business

When evaluating mobile applications, most stakeholders focus on what's visible: the user interface, features, and performance. But the mobile app is merely the tip of the iceberg.

The real business model—and the overwhelming majority of security, compliance, and operational risk—lies hidden beneath the surface in the cloud infrastructure that powers it.

The bottom-right corner of the page features decorative geometric shapes, consisting of two overlapping squares: a larger light blue one and a smaller, darker blue one, mirroring the design in the top corners.

What Users See vs. What Drives the Business

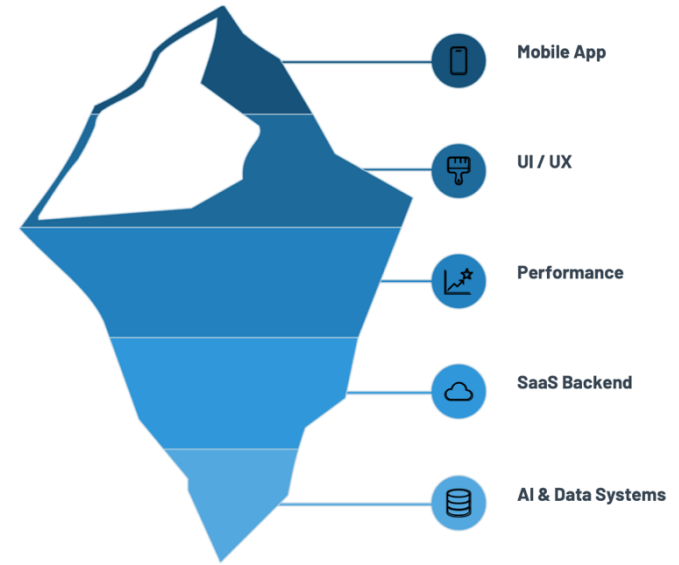
Above the Surface: The Mobile App (10%)

The visible layer where users interact

- Polished UI/UX design and interface elements
- App performance and response times
- Feature functionality and user workflows
- Client-side validation and basic security

Just like an iceberg, 90% of your mobile business infrastructure exists below the visible surface—in cloud services, databases, and third-party integrations.

This is where the majority of business risk, security vulnerabilities, and compliance obligations actually reside. Yet it's often the least scrutinized during product evaluations.



Below the Surface: The SaaS Backend (90%)

The invisible infrastructure powering everything

- AI models and machine learning pipelines
- Training datasets and data warehouses
- Third-party API integrations and dependencies
- Compliance frameworks and audit trails
- PII storage, encryption, and access controls

The Visible Layer: Mobile App Components

1

User Interface & Experience

Visual design, navigation flows, accessibility features, and responsive layouts across device types.

2

Performance Metrics

Load times, responsiveness, battery consumption, network efficiency, and crash-free sessions.

3

Feature Set

Core functionality, integrations, offline capabilities, and differentiating product features.

4

Client-Side Security

Basic authentication, local data encryption, certificate pinning, and app sandboxing controls.



Where Mobile Apps Store Your Business Risk

Third-Party API Dependencies

Every external service integration multiplies your attack surface. One compromised vendor can expose your entire user base.

PII Storage Architecture

Personal health information, financial data, biometrics—all subject to strict regulations. Where does it live? How long? Who has access?

AI Model Training Data

Your models learned from somewhere. Can you prove that training data was legally obtained and properly anonymized?

Compliance Framework Gaps

ISO 42001, SOC 2, HIPAA, GDPR—each framework has different requirements. Missing even one can block market access.



The Hidden Layer: Cloud Infrastructure Reality

Beneath every mobile app lies a complex ecosystem of cloud services, data processing pipelines, and interconnected systems. This is where your actual business logic, sensitive data, and critical operations live.

Understanding this infrastructure is essential for evaluating true business risk, ensuring regulatory compliance, and protecting customer data at scale.



Critical Backend Components

AI Models & ML Pipelines

Training infrastructure, model versioning, inference engines, and bias detection systems requiring constant monitoring.

Data Storage & Processing

Customer PII, transaction records, audit logs, and training datasets with complex retention policies.

Third-Party Dependencies

Payment processors, authentication providers, analytics services, and external APIs creating risk vectors.

Where the Real Risk Lives

90%

Infrastructure Risk

Of total security vulnerabilities exist in backend systems and cloud services, not mobile clients.

73%

Data Breaches

Of major data breaches originate from compromised cloud infrastructure or API vulnerabilities.

\$4.5M

Average Cost

Average cost of a data breach involving cloud-based systems and third-party integrations.





Why This Matters for Product Evaluation

01

Shift Security Focus Downward

Evaluate infrastructure security, data governance, and compliance frameworks—not just app features.

03

Assess Compliance Architecture

Verify GDPR, CCPA, SOC 2, and industry-specific compliance is baked into infrastructure, not bolted on.

02

Question Third-Party Dependencies

Understand every external service, their security posture, data access, and potential failure modes.

04

Demand Transparency

Request architecture diagrams, data flow maps, incident response plans, and disaster recovery procedures.

Four Pillars of Mobile-Cloud Security Governance



Visibility

You cannot secure what you cannot see. Implement full-stack observability from mobile edge to cloud core, with real-time monitoring of data flows and model behavior.



Auditability

Every algorithmic decision must have a traceable lineage. Regulators demand explanations—your architecture should generate them automatically.



Accountability

Establish clear ownership for AI outcomes. Define who approves model updates, monitors performance, and responds to failures.



Adaptability

Regulatory requirements will evolve. Build compliance systems that can absorb new frameworks without complete re-architecture.

What's Below the Waterline?

The mobile app your users see is just the beginning. The real questions are:



Can you see every component of your cloud infrastructure?



Can you explain how your AI models make decisions?



Can you prove compliance to regulators in multiple jurisdictions?



Can your security architecture scale with your business?

The organizations that thrive in the next decade won't be those with the best apps—they'll be those with the most transparent, auditable, and compliant infrastructure beneath them.



The Cost of Getting It Wrong

Regulatory Penalties

- GDPR fines up to **€20M or 4% of global revenue**
- EU AI Act violations: **€35M or 7% of turnover**
- HIPAA breaches: **\$1.5M per violation category**

Operational Impact

- Market access blocked in regulated regions
- Forced suspension of AI-powered features
- Emergency security audits consuming 6-12 months

Reputational Damage

- Loss of customer trust (irreversible)
- Competitive disadvantage in security-conscious markets
- Difficulty raising capital post-incident



The question isn't whether you can afford to invest in security—it's whether you can afford not to.

Key Takeaway: Look Below the Surface

The mobile app is what users see. The cloud infrastructure is what determines your risk.

When evaluating mobile business applications, dedicate the majority of your due diligence to understanding what lies beneath: the SaaS backend, data architecture, compliance frameworks, and third-party dependencies.

This is where your real business exposure lives—and where informed decisions make the difference between security and catastrophe.






The "Black Box" Paradox

Understanding the challenge of auditing Deep Learning systems and why opacity creates critical governance risks.

The Transparency Challenge in AI

Modern AI systems, particularly deep learning models, operate as opaque decision-makers. While they deliver impressive results, their internal reasoning remains hidden—creating a fundamental paradox for organizations that need accountability.

This opacity poses serious challenges for auditors, regulators, and stakeholders who must validate fairness, compliance, and reliability in automated systems.



Why "Black Box" Matters



Regulatory Risk

Inability to explain decisions creates compliance vulnerabilities and legal exposure across industries.



Fairness Concerns

Hidden biases can perpetuate discrimination without detection or accountability mechanisms.

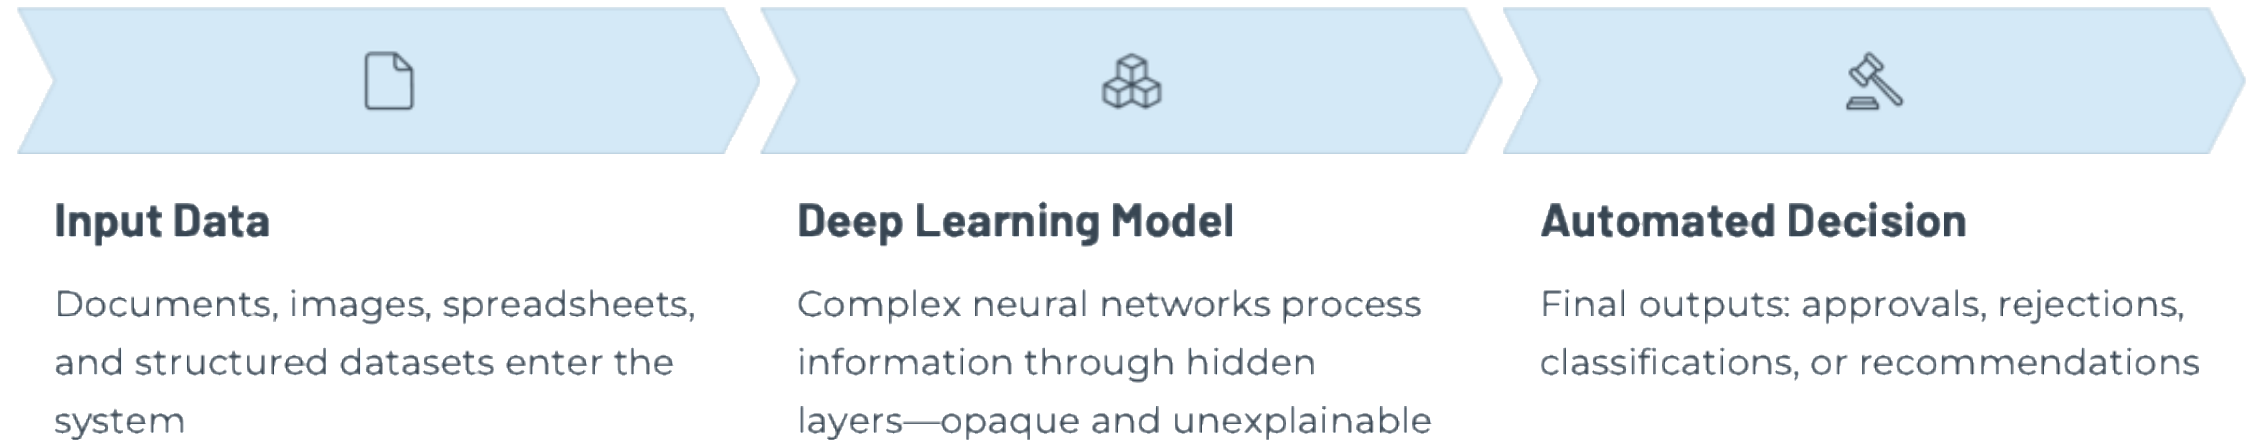


Audit Impossibility

Traditional verification methods fail when decision logic remains fundamentally inaccessible to inspection.

The Black Box Flow

Deep learning systems transform diverse inputs into automated decisions through opaque internal processes that resist traditional audit approaches.



❏ The critical question: **Why did the model make this specific decision?** Without visibility into the middle stage, accountability becomes impossible.

Inside the "Black Box"



What's Really Happening?

- **Millions of parameters** adjust during training, creating complex decision boundaries
- **Non-linear transformations** across multiple hidden layers obscure direct input-output relationships
- **Emergent behavior** that even designers cannot fully predict or explain
- **Mathematical complexity** beyond human comprehension at scale

Why Traditional Methods Fail

No Clear Logic Trail

Unlike rule-based systems, deep learning lacks explicit if-then decision paths that auditors can follow and verify.

Technical Expertise Gap

Understanding model architecture requires specialized data science knowledge rarely found in audit teams.

Scale and Complexity

Models with billions of parameters process data through hundreds of layers—far exceeding human review capacity.





Real-World Implications

Healthcare Decisions

AI diagnoses without explanation erode physician trust and patient consent mechanisms.


Credit Approvals

Unexplained loan denials violate fair lending laws requiring adverse action explanations.

Hiring Systems

Opaque candidate screening perpetuates bias without accountability or recourse for applicants.

In each scenario, the inability to audit decisions creates legal, ethical, and operational risks that organizations cannot ignore.



The Central Question

Why

?

"The most critical question in AI governance isn't whether the system works—it's whether we can explain *why* it made each specific decision."

This question sits at the heart of responsible AI deployment. Without satisfactory answers, organizations face mounting pressure from regulators, customers, and internal stakeholders demanding transparency.

Approaches to Open the Box

01

Explainable AI (XAI) Tools

Techniques like SHAP, LIME, and attention visualization provide partial insights into model reasoning.

03

Hybrid Architectures

Combining deep learning with interpretable models creates decision systems that balance accuracy with explainability.

While none provide complete transparency, these approaches offer practical pathways toward more auditable AI systems.

02

Model Documentation

Comprehensive model cards and data sheets establish audit trails for training data and performance metrics.

04

Regulatory Frameworks

New standards like EU AI Act mandate transparency requirements for high-risk AI applications.

Moving Forward: Balancing Power and Transparency

The Path Ahead

Organizations must navigate the tension between AI's powerful capabilities and the governance imperative for explainability. Success requires:

- Investing in XAI capabilities early in development
- Building cross-functional teams bridging AI and audit expertise
- Establishing clear accountability frameworks
- Prioritizing transparency in high-stakes decisions



Key Takeaway: The black box paradox isn't insurmountable but solving it requires deliberate organizational commitment to transparency as a core design principle—not an afterthought.






Building Trust in the Black Box

Designing & Auditing AI-Driven Business Models

The Trust Paradox

Modern AI-driven platforms promise seamless experiences—a single tap to buy, sell, or transact. But beneath that polished interface lies a labyrinth of algorithms, data pipelines, and decision-making systems most users never see.

The challenge? Building products that inspire confidence while operating in inherent complexity.



△ CHALLENGE

Why "Black Box" Systems Fail Users

Opacity Breeds Distrust

When users can't understand how decisions are made, they hesitate. Studies show 67% of consumers avoid AI products they perceive as opaque.

Security Theater vs. Real Protection

Surface-level assurances don't address underlying vulnerabilities. Users need proof, not promises.

Compliance Isn't Enough

Meeting regulatory standards is baseline. True trust requires transparency that goes beyond checkboxes.



The Designer's Dilemma

You're Caught Between Two Worlds

Create intuitive, frictionless experiences that feel magical to users—while communicating the rigorous security and intelligence happening behind the scenes.

73%

Users Want Simplicity

Prefer minimal clicks and instant results

81%

Executives Demand Proof

Require auditable AI processes and security controls

Four Pillars of Trustworthy AI Design

A practical framework for product teams navigating the complexity.



Transparency

Show users how decisions are made without overwhelming them with technical detail.



Security by Design

Embed verification and audit trails from day one, not as afterthoughts.



Explainability

Provide context-appropriate explanations that match user sophistication levels.



Continuous Validation

Build feedback loops that adapt to emerging threats and user concerns.

The Audit-First Design Process



Map the Decision Tree

Document every AI decision point and data dependency before touching design tools.



Identify Stakeholder Needs

Different audiences require different transparency depths—map them explicitly.



Design with Layers

Create progressive disclosure systems that balance simplicity with depth.



Build Audit Mechanisms

Embed logging, monitoring, and explanation systems into the architecture.



Test and Iterate

Validate with both security auditors and end users continuously.

Red Flags in AI Product Design

Warning signs that your platform may be vulnerable to trust breakdowns.



No Explanation Layer

Users receive decisions without any context or reasoning provided.



Security Through Obscurity

Relying on complexity rather than proven cryptographic controls.



Single-Layer Interfaces

No way for technical stakeholders to access deeper system information.



Retroactive Compliance


Bolting on audit capabilities after launch instead of designing them in.



Measuring Trust: Beyond NPS Scores


Traditional Metrics Miss the Mark

Standard product metrics don't capture whether users trust your AI systems. You need specialized measurements that reveal confidence levels.




Transparency Index

Users who understand system reasoning



Audit Confidence

Compliance teams satisfied with trails



Perceived Control

Users feeling agency over AI decisions

Track these metrics quarterly alongside traditional product KPIs to catch trust erosion early.

The Future is Transparent by Design

"The companies that win in the AI era won't be those with the most sophisticated algorithms—they'll be the ones that make complexity comprehensible."

Your Next Steps

01

Audit Your Current State

Map where your product falls on the transparency spectrum

02

Design the Layers

Build progressive disclosure for different stakeholder needs

03

Embed Security

Make auditability a core architectural principle

04

Measure Trust

Implement metrics that capture confidence, not just satisfaction



The AI Regulation Tsunami

A comprehensive overview of the accelerating regulatory landscape transforming how businesses deploy artificial intelligence technologies.



EXECUTIVE BRIEF

Why This Matters Now

Business Impact

Non-compliance can result in fines up to €35M or 7% of global revenue under the EU AI Act.

Market Access

Regulatory compliance is now a prerequisite for operating in major markets worldwide.

Competitive Edge

Early adopters of compliance frameworks gain trust and market advantage.



The Accelerating Regulatory Timeline

Regulation isn't just coming—it's accelerating. What began with data privacy has evolved into comprehensive AI governance frameworks that fundamentally reshape business operations.

The Regulatory Tsunami: 2018–2025

2018: GDPR

EU establishes global data privacy standard, setting precedent for tech regulation

1

2

2023: Gen AI Boom

ChatGPT sparks explosive AI adoption, catching regulators' attention worldwide

3

2024: EU AI Act

Comprehensive AI legislation passes, creating first major regulatory framework

4

2025: ISO 42001

International standard for AI management systems becomes operational

📌 The pace is accelerating—regulations that took years to develop now emerge in months.

EUROPE

GDPR: The Foundation

What Changed in 2018

The General Data Protection Regulation established unprecedented requirements for data handling, consent, and individual rights. It fundamentally shifted how companies process personal information.

- Mandatory data protection by design
- Explicit consent requirements
- Right to explanation for automated decisions
- Hefty penalties for violations

€20M

Maximum Fine

Or 4% of global revenue



INNOVATION

2023: The Generative AI Revolution

ChatGPT's Impact

100 million users in 2 months—the fastest-growing consumer application in history, forcing immediate regulatory attention.

Enterprise Adoption

Organizations rushed to integrate generative AI, often outpacing their own governance capabilities and risk frameworks.

Regulatory Response

Governments worldwide accelerated AI legislation timelines, recognizing the urgent need for guardrails.

EU AI Act: The Game Changer



The EU AI Act introduces a risk-based classification system that determines compliance requirements. High-risk applications face rigorous assessment, documentation, and monitoring obligations.

ISO 42001: International Standard



Global AI Management

ISO 42001 provides a comprehensive framework for establishing, implementing, and maintaining an AI management system.

01

Risk Assessment

Systematic evaluation of AI systems

02

Governance Structure

Clear roles and accountability

03

Continuous Monitoring


Ongoing performance tracking

04

Documentation


Comprehensive audit trails

The Pressure Is Intensifying




Executives Concerned

About AI regulatory risk




Companies Unprepared

Lack compliance frameworks



Enforcement Growth

Expected in next 18 months

 The window for proactive compliance is closing rapidly as enforcement mechanisms come online.



Your Next Steps



Assess

Inventory your AI systems and classify risk levels



Mobilize

Form cross-functional compliance teams now



Document

Begin building audit trails and governance records

The regulatory tsunami is here. Organizations that act decisively today will navigate these waters successfully—those that wait risk being swept away.



Thank you! Any questions?

NEXT:

- Break for 10 min
- Case Study and Discussion
- Quick Skill Check
- Practical Handouts

Case Study: SafeRoute

You are designing SafeRoute, a mobile app for students and staff that provides real-time routing and an optional “Safety Score” for routes (lighting, incident reports, crowd levels, time-of-day patterns). Users can share an ETA with trusted contacts and submit incident/hazard reports (optionally anonymous).

Your tasks:

- Sketch a data flow including third parties.
- Identify top 3 risks
- Choose 5 controls you’d implement first
- Define 2 audit/assurance artifacts you want available in 6 months
- Name one business-model/design decision you’d change to reduce risk

