

# PRAGMATIC AI GOVERNANCE

---

A Hands-on Course Syllabus Offered by [the RAI Fellow](#)

Instructor: Iman Saleh, PhD ([bio](#))

*Professional Training & Certification Series*  
© 2026 The RAI Fellow

## 1. Course Overview

This class is designed for AI Governance professionals to walk them through the journey of building an AI Governance program. The class uses the Governance Maturity Model described in the book “Pragmatic AI Governance” to provide a step-by-step methodology to building AI governance from the ground up, and scale it with the AI investment.

## 2. High-Level Course Structure

The table below details the course roadmap and core deliverables.

| Sessions and Duration | Topic                                             | Deliverable                                       |
|-----------------------|---------------------------------------------------|---------------------------------------------------|
| Session 1<br>2 Hrs    | * AI and The Probabilistic Revolution             | Questionnaire: Categorizing AI Paradigms          |
| Session 2<br>2 Hrs    | * AI Risks and Safeguards                         | Project Part I: Use Case Analysis                 |
| Session 3<br>1 Hr     | * The Anatomy of an AI Governance Program         | Project Part II: Use Case - Risk Heat Map         |
| Session 4<br>1 Hr     | * The AI Governance Maturity Model                | Questionnaire: Identify a Maturity Level          |
| Session 5<br>1 Hr     | * Principles and Policies                         | Project Part III: Drafting a Policy               |
| Session 6<br>1 Hr     | * Mastering Provenance: Data, Model, and Use Case | Project Part IV: Drafting a Governance Taxonomy   |
| Session 7<br>1 Hr     | * Enforcement                                     | Project Part V: Designing a Safeguard             |
| Session 8<br>1 Hr     | * Auditing and Reporting + Incident Management    | Project Part VI: A Governance Dashboard Blueprint |
| Session 9<br>2 Hr     | * Projects Presentation and Discussion            |                                                   |

### **3. Detailed Weekly Course Syllabus**

#### **Session 1: AI and the Probabilistic Revolution**

Learners will explore the fundamentals of AI, distinguishing between training and inference. We will trace the historical evolution of predictive technologies, from collaborative recommender systems to white-box supervised classifiers and opaque deep learning neural networks, concluding with a conceptual analysis of Generative and Agentic AI models.

#### **Session 2: AI Risks and Safeguards**

We will uncover the risks built into the AI lifecycle and the strategies used to address them, including. Learners will leave with a solid grasp of why responsible AI practices matters now more than ever, and how to build it into the AI pipelines from day one rather than retrofitting it later.

#### **Session 3: The Anatomy of an AI Governance Program**

Learners will learn why organizations must move past late-stage reactive firefighting to build proactive governance frameworks. We will analyze the key components of an AI governance program, the role of each component and how to connect them so the whole system operates seamlessly. Since the program is driven by risk, learners will draft a risk assessment heat map for a real-world use case that they will use in later sessions to apply the different stages of governance.

#### **Session 4: Principles and Policies**

Learners will learn how to write effective corporate AI policies. We'll cover how to translate high-level legal rules and principles into clear, measurable and enforceable policies and learners will get hands-on experience drafting a policy of their own.

#### **Session 5: Mastering Provenance: Data, Model, and Use Cases**

An in-depth study of lineage tracking across AI pipelines. Learners will design a machine-readable governance taxonomy and analyze how to construct continuous tracking systems across the three dimensions of provenance.

#### **Session 6: Enforcement**

Policy enforcement typically consists of governance review processes and technical controls. We'll weigh the pros and cons of each, showing learners how to enforce the rules without slowing down the business. Then, they will put their learnings into practice by designing a safeguard for their real-world use case.

#### **Session 7: Auditing and Reporting + Incident Management**

Learners will learn how to set up clear reporting requirements across the AI lifecycle, and discover how a strong incident management process elevates the entire governance program. Learners will design a blueprint for an AI Governance Dashboard.

#### **Session 8: Projects Presentation and Discussion**

We will conclude with a practical discussion on governing organizational AI, covering strategies for scaling, navigating cultural shifts, and addressing pushbacks. Learners will also present their projects to share insights and key takeaways.

## 4. Why Choose the RAI Fellow

There are plenty of "AI Ethics" bootcamps out there that will spend three days lecturing you on high-level moral principles. This is not one of them.

Here is why this program is the ultimate vehicle to transform you into a true AI governance guru:

### **Built on Actual Battle Scars, Not Textbooks.**

This curriculum is forged from years of real-world experience building, scaling, and occasionally firefighting AI and governance inside highly complex enterprise environments. We focus heavily on the "friction gap"; showing you how to replace slow, manual "DMV-style" review gates with automated, silent guardrails that protect the enterprise without slowing down your engineering pipelines.

### **Dual-Perspective Leadership**

Created and taught by Dr. Iman Saleh, an industry expert operating at the intersection of technical research and production engineering. Holding a PhD-level technical background, Dr. Saleh understands the mathematical foundations of probabilistic neural networks. Simultaneously, her background in corporate environments provides practical insight into deploying effective governance frameworks alongside aggressive AI deployment cycles.

### **Practical Deliverables Over Passive Listening.**

You won't just sit through slides. You will actively build. Throughout the program, you will construct a custom AI Risk Heat Map, draft a localized AI Policy Blueprint, and design a live Governance Dashboard architecture. You leave the program not just with a certificate, but with a portfolio of high-value governance assets.

## 5. Format

The course is delivered in one or more of the following formats:

### **Format 1: 3-Day Live Virtual/In-Person Intensive Workshop**

**What's Included:** 20 hours of live instruction including 12 hours of lectures and 8 hours of facilitated hands-on group work on the deliverables, real-time Q&A, and a Certificate of Completion

**Target Audience:** AI governance teams seeking immediate upskilling.

### **Format 2: 8-Week Guided Cohort (Self-Paced + Feedback)**

**What's Included :** Pre-recorded weekly lecture videos, access to a peer discussion forum, weekly assignment submissions with personalized written feedback from the instructor, and a Certificate of Completion.

**Target Audience:** Busy AI governance teams who need flexibility but still want rigorous accountability and professional feedback.

### **Format 3: On-Demand Access (Self-Study)**

**What's Included :** Instant access to the recorded sessions, slide decks, and project templates. No homework grading, no instructor Q&A, and no community access. Sessions are reviewed and updated annually, as needed.

**Target Audience:** AI governance teams looking for self-paced learning or a quick reference utility.

**For Pricing and Scheduling**

**Book a [Free Appointment](mailto:info@theraifellow.com) / [info@theraifellow.com](mailto:info@theraifellow.com)**