

FUTURE-FIT PROJECT LEADERSHIP



Project
Management
Institute.
Finland

Tam Abaku



How Project Managers Lead Ethical AI in Government



M.O.R.E. Than Compliance

How Project Managers Lead Ethical AI in Government

Someone must own the ethical outcome of government AI. Today we find out who that is. The profession has long avoided this simple question.



80km/h



What You Will Walk Away With

Our Session Objectives

1

Recognize Structural Risk

Identify bias entry points beyond training data across the full system lifecycle.

2

Master Compliance Thresholds

Understand the mandates of the new EU AI Act and what they require of you.

3

Implement M.O.R.E.

Operationalize ethical governance using a clear, actionable framework.

4

Deploy 5 Governance Gates

Embed ethical checkpoints directly into your next Agile sprint.

Past arrests sent patrols back to the same minority neighbourhoods.

PREDPOL / GEOLITICA | Lum & Isaac, 2016; The Markup, 2021

Black defendants falsely flagged high-risk at nearly twice the rate.

COMPAS | ProPublica, Machine Bias | May 2016

CONTEXTUAL BACKGROUND AND DEFINITION

Wrongly arrested on an unverified facial recognition match.

WILLIAMS v. CITY OF DETROIT | ACLU | Arrest 2020, settled 2024

Thousands of families falsely accused of benefit fraud. A government resigned.

TOESLAGENAFFAIRE | Netherlands | Cabinet resigned Jan 2021



A Government That Trusted the Algorithm

The Quiet Shift to Automated Decisions

77%

Governments Using AI

Now deploy AI in public services — a quiet
revolution in how states operate.

60M+

Citizens Affected
Annually

Receive algorithmically shaped decisions on
benefits, bail, and borders.



Critical question: what happens when the machine is wrong?

The World These Projects Actually Operate In

Why the Public Sector is Fundamentally Different

No Opt-Out

Citizens cannot opt out of automated government systems. There is no alternative service to turn to.

No Competition

No competitors or alternative services exist for the public — the state holds a monopoly on these decisions.

Legal Reality

"The algorithm decided" is not a legal defense. Human accountability cannot be delegated to a model.

Encoded Discrimination

Legacy data encodes historical discrimination as absolute fact — and systems learn from it.

Catastrophic Consequences

Public AI failures topple governments, not just corporate reputations.

3906 6468



The High Cost of Unowned Ethical Outcomes

Four Stories We Should Never Repeat

1

Netherlands 2005–2019

35,000 families wrongly flagged for fraud. The cabinet resigned.

2

Netherlands 2014–2020

Postcode profiling banned for human rights violations.

3

UK 2020

280,000 student grades downgraded by algorithm; reversed under heavy protest.

4

UK Ongoing

Risk tool missed an offender; four lives lost in Killamarsh.



Every project had a manager. None owned the ethical outcome.

Four Voices When Things Go Wrong

The Accountability Gap in Siloed Teams



Developer

"The model performed exactly as specified."

Legal Team

"The specification was reviewed and approved."

Data Scientist

"The data was provided by the client."

Project Manager

"The project was delivered on time and budget."



Everyone is technically correct, yet nobody owned the outcome.



What We Mean When We Say AI Ethics

Moving Beyond Feelings to Frameworks

AI Ethics is not a philosophical debate — it is an **operational framework** applied across every lifecycle phase of a system.

Systems missing these components are **not ready for public service**.

Fairness

Mandatory test

Transparency

Mandatory test

Legal Rights

Core baseline

Accountability

Core baseline

Trauma labelled for under \$2 an hour. [OPENAI / SAMA](#) | [TIME investigation](#) | [Jan 2023](#)

Bookbots cleared for
“romantic or sensual” chats
with children.

[META](#) | [GenAI: Content Risk Standards](#) | [Reuters, Aug 2025](#)

History rewritten. Image
generation paused.

[GOOGLE GEMINI](#) | [Image generation pause](#) | [Feb 2024](#)

Six court cases cited.
None existed.

[MATA v. AVIANCA](#) | [No. 22-cv-1461 \(PKC\)](#) | [S.D.N.Y. 2023](#)

Trauma labelled for
under \$2 an hour.

[OPENAI / SAMA](#) | [TIME investigation](#) | [Jan 2023](#)

Books allegedly used to train AI without consent. [AUTHORS GUILD v. OPENAI](#) | [S.D.N.Y.](#) | [Filed Sept 2023](#) Six court cases cited. None existed. [MATA v. AVIANCA](#) | [No. 22-cv-1461 \(PKC\)](#) | [S.D.N.Y. 2023](#)

Wrongly arrested on a facial recognition match. [WILLIAMS v. CITY OF DETROIT](#) | [ACLU](#) | [Arrest 2020, settled 2024](#)

STRUCTURAL RISKS: BIAS ENTRY POINTS

Higher false-match rates for Black and Asian faces. [NIST](#) | [NISTIR 8280](#) | [Dec 2019](#)

Higher false-match rates
for Black and Asian
faces.

[NIST](#) | [NISTIR 8280](#) | [Dec 2019](#)

Wrongly arrested on a facial
recognition match.

[WILLIAMS v. CITY OF DETROIT](#) | [ACLU](#) | [Arrest 2020, settled 2024](#)

Books allegedly used to
train AI without
consent.

[AUTHORS GUILD v. OPENAI](#) | [S.D.N.Y.](#) | [Filed Sept 2023](#)

Public photos
turned into explicit
deepfakes.

[DEFIANCE ACT](#) | [U.S. Senate](#) | [2024](#)

Chatbots cleared for “romantic or sensual” chats with children. [META](#) | [GenAI: Content Risk Standards](#) | [Reuters, Aug 2025](#)


Books allegedly used to train AI without consent. [AUTHORS GUILD v. OPENAI](#) | [S.D.N.Y.](#) | [Filed Sept 2023](#) Six court cases cited. None existed. [MATA v. AVIANCA](#) | [No. 22-cv-1461 \(PKC\)](#) | [S.D.N.Y. 2023](#)

The Problem Is Bigger Than Your Data

Bias Enters Far Beyond the Training Set

Cleaning data only solves 25% of the problem. Bias enters systems at 7 distinct points across three levels:

🔗 The project manager owns all 7 entry points.



Data-Level
Historical, Representation, and Measurement Bias

Algorithm-Level
Learning Bias and Evaluation Bias

System-Level
Deployment Bias and Feedback Loop Bias

Past arrests sent patrols back to the same minority neighbourhoods.

PREDPOL / GEOLITICA | Lum & Isaac, 2016; The Markup, 2021

Black defendants falsely flagged high-risk at nearly twice the rate.

COMPAS | ProPublica, Machine Bias | May 2016

COMPLIANCE THRESHOLDS

Wrongly arrested on an unverified facial recognition match.

WILLIAMS v. CITY OF DETROIT | ACLU | Arrest 2020, settled 2024

Thousands of families falsely accused of benefit fraud. A government resigned.

TOESLAGENAFFAIRE | Netherlands | Cabinet resigned Jan 2021

Our Profession Has Given the Answer

Mandates from Global Governing Bodies

PMI AI Standard 2024

Fairness is an explicit practitioner obligation — not optional guidance.

EU AI Act

Legal duties for risk management; fines up to **€30M** for non-compliance.

PMBOK 8th Edition

Stewardship explicitly extends to societal impacts — beyond scope, time, and cost.




Delivering a biased AI on time is an absolute project failure.



The Law Has Already Decided

Regulatory Realities Under the EU AI Act



Trauma labelled for under \$2 an hour. [OPENAI / SAMA](#) | [TIME investigation](#) | Jan 2023

Bookbots cleared for
“romantic or sensual” chats
with children.

[META](#) | [GenAI: Content Risk Standards](#) | [Reuters, Aug 2025](#)

History rewritten. Image
generation paused.

[GOOGLE GEMINI](#) | [Image generation pause](#) | Feb 2024

Six court cases cited.
None existed.

[MATA v. AVIANCA](#) | No. 22-cv-1461 (PKC) | S.D.N.Y. 2023

Trauma labelled for
under \$2 an hour.

[OPENAI / SAMA](#) | [TIME investigation](#) | Jan 2023

Books allegedly used to train AI without consent. [AUTHORS GUILD v. OPENAI](#) | S.D.N.Y. | Filed Sept 2023

Wrongly arrested on a facial recognition match. [WILLIAMS v. CITY OF DETROIT](#) | [ACLU](#) | Arrest 2020, settled 2024

SELECTED FRAMEWORK: M.O.R.E.

History rewritten. Image generation paused. [GOOGLE GEMINI](#) | [Image generation pause](#) | Feb 2024

Higher false-match rates for Black and Asian faces. [NIST](#) | [NISTIR 8280](#) | Dec 2019

Higher false-match rates
for Black and Asian
faces.

[NIST](#) | [NISTIR 8280](#) | Dec 2019

Wrongly arrested on a facial
recognition match.

[WILLIAMS v. CITY OF DETROIT](#) | [ACLU](#) | Arrest 2020, settled 2024

Books allegedly used to
train AI without
consent.

[AUTHORS GUILD v. OPENAI](#) | S.D.N.Y. | Filed Sept 2023

Public photos
turned into explicit
deepfakes.

[DEFIANCE ACT](#) | [U.S. Senate](#) | 2024

Chatbots cleared for “romantic or sensual” chats with children. [META](#) | [GenAI: Content Risk Standards](#) | [Reuters, Aug 2025](#)

Books allegedly used to train AI without consent. [AUTHORS GUILD v. OPENAI](#) | S.D.N.Y. | Filed Sept 2023

Six court cases cited. None existed. [MATA v. AVIANCA](#) | No. 22-cv-14 61 (PKC) | S.D.N.Y. 2023



The Framework You Already Have

Deploying M.O.R.E. for Responsible AI

M — Manage Trust

Publish system logic. Allow citizens to contest decisions made about them.



O — Own the Outcome

Name the accountable human explicitly. Ethical ownership cannot be anonymous.



R — Reassess Continuously

Monitor model degradation and population shifts throughout the system's life.

E — Expand Your View

Identify marginalized stakeholders missing from design rooms before launch.

27

Without M.O.R.E.

Average project success score without ethical governance framework.

94

With M.O.R.E.

Project success score with consistent M.O.R.E. application.

Past arrests sent patrols back to the same minority neighbourhoods.

PREDPOL / GEOLITICA | Lum & Isaac, 2016; The Markup, 2021

Black defendants falsely flagged high-risk at nearly twice the rate.

COMPAS | ProPublica, Machine Bias | May 2016

THE GOVERNANCE GATES

Wrongly arrested on an unverified facial recognition match.

WILLIAMS v. CITY OF DETROIT | ACLU | Arrest 2020, settled 2024

Thousands of families falsely accused of benefit fraud. A government resigned.

TOESLAGENAFFAIRE | Netherlands | Cabinet resigned Jan 2021

Governance Is an Operational Cadence

The Five Gates of Ethical AI Delivery

Failed fairness checks completely halt project progression.

01

G0 (Initiation)

EU AI Act classification and documented fairness definition.

02

G1 (Training)

Audit all 7 bias points and verify data consent.

03

G2 (Development)

Embed fairness parameters directly into the Definition of Done.

04

G3 (Go-Live)

Pass bias testing, validate HITL, and approve DPIA.

05

G4 (Post-Launch)

Deploy drift alerts, monthly dashboards, and quarterly validation.

Seven entry points, three levels, one accountable owner

Where bias enters the AI lifecycle, the gate that catches it, and who leads the Intervention



O1 of O7, data level

Historical bias

Where it enters


Training labels inherit past discrimination. The Toeslagenaffaire model learned nationality as a fraud signal from earlier audits.

Caught at

G1 Data approval

Who leads, and how

Data steward. Audit label provenance before the data is approved for use.



6536 6820

AI system to review public housing applications. We can either choose: Option A, or B. What do you vote for?

The Decision Nobody Is Making

Defining Fairness Before the First Sprint

What does "fair" mean for this specific population and context?

Four mathematically distinct, conflicting definitions exist.

Equal Opportunity $P(\hat{Y} = 1 \mid Y = 1, A = a) = P(\hat{Y} = 1 \mid Y = 1, A = b)$

Equal chances for equally qualified applicants (Benefits).

Equalized Odds $P(\hat{Y} = 1 \mid Y = 1, A = a) = P(\hat{Y} = 1 \mid Y = 1, A = b)$ $P(\hat{Y} = 1 \mid Y = 0, A = a) = P(\hat{Y} = 1 \mid Y = 0, A = b)$

Equal True Positive and False Positive Rates (Justice/Health).

Demographic Parity $P(\hat{Y} = 1 \mid A = a) = P(\hat{Y} = 1 \mid A = b)$

Equal outcomes to correct severe historical exclusion.

Predictive Parity $P(Y = 1 \mid \hat{Y} = 1, A = a) = P(Y = 1 \mid \hat{Y} = 1, A = b)$

Equal precision across groups (Risk scoring).

You cannot satisfy all four simultaneously; the choice is a values decision.

Trauma labelled for under \$2 an hour. [OPENAI / SAMA](#) | [TIME investigation](#) | [Jan 2023](#)

Bookbots cleared for
“romantic or sensual” chats
with children.

[META](#) | [GenAI: Content Risk Standards](#) | [Reuters, Aug 2025](#)

History rewritten. Image
generation paused.

[GOOGLE GEMINI](#) | [Image generation pause](#) | [Feb 2024](#)

Six court cases cited.
None existed.

[MATA v. AVIANCA](#) | [No. 22-cv-1461 \(PKC\)](#) | [S.D.N.Y. 2023](#)

Trauma labelled for
under \$2 an hour.

[OPENAI / SAMA](#) | [TIME investigation](#) | [Jan 2023](#)

Books allegedly used to train AI without consent. [AUTHORS GUILD v. OPENAI](#) | [S.D.N.Y.](#) | [Filed Sept 2023](#) Six court cases cited. None existed. [MATA v. AVIANCA](#) | [No. 22-cv-1461 \(PKC\)](#) | [S.D.N.Y. 2023](#)

Wrongly arrested on a facial recognition match. [WILLIAMS v. CITY OF DETROIT](#) | [ACLU](#) | [Arrest 2020, settled 2024](#)

APPLIED CASE EXAMPLES

Higher false-match rates for Black and Asian faces. [NIST](#) | [NISTIR 8280](#) | [Dec 2019](#)

Higher false-match rates
for Black and Asian
faces.

[NIST](#) | [NISTIR 8280](#) | [Dec 2019](#)

Wrongly arrested on a facial
recognition match.

[WILLIAMS v. CITY OF DETROIT](#) | [ACLU](#) | [Arrest 2020, settled 2024](#)

Books allegedly used to
train AI without
consent.

[AUTHORS GUILD v. OPENAI](#) | [S.D.N.Y.](#) | [Filed Sept 2023](#)

Public photos
turned into explicit
deepfakes.

[DEFIANCE ACT](#) | [U.S. Senate](#) | [2024](#)

Chatbots cleared for “romantic or sensual” chats with children. [META](#) | [GenAI: Content Risk Standards](#) | [Reuters, Aug 2025](#)

Books allegedly used to train AI without consent. [AUTHORS GUILD v. OPENAI](#) | [S.D.N.Y.](#) | [Filed Sept 2023](#) Six court cases cited. None existed. [MATA v. AVIANCA](#) | [No. 22-cv-1461 \(PKC\)](#) | [S.D.N.Y. 2023](#)

A Sprint Planning Session with Teeth

Embedding Ethics into the Agile Cadence

Every user story carries a fairness acceptance criterion.

+5

Initial Story Points

Velocity completely recovered by sprint 3.

68%

Saved in Rework

Fairness blocker caught Day 7.

13

Days Earlier Detection

Bias issues detected than traditional paths.

Evolving User Stories

Legacy

"As an officer, I want to assess housing eligibility faster."

Ethical

"...ensuring approval rates do not vary by demographics by $\pm 5\%$."

Same Country. Same Law. Different Story. The Difference Governance Makes

Robodebt (Australia): A Cautionary Tale

- 500,000 wrongful notices issued.
- Zero oversight of automated systems.
- AU\$1.8 billion repaid to citizens.



ACT Government (Australia): A Model for Success

- Ethics framework published pre-deployment.
- Named Senior Responsible Officers (SROs) for accountability.
- Public transparency registers in place.
- Automated quarterly bias audits implemented.

The differentiator between these two outcomes is **five gates, four roles, and eight questions** that embed ethical governance throughout the AI lifecycle.

You Do

The Ultimate Ownership of AI Outcomes

"The public sector cannot afford to get this wrong."

"Governments cannot afford to opt out of innovation."

"Someone must ask the questions nobody else thought to ask."

"Owning the outcome nobody else claimed falls to you."

Move M.O.R.E. than compliance.

Public Sector AI Project RACI Matrix

Not just roles. **Named people**, for every decision.

	JC Joel Cinthio Project Manager	EC Dr El-naz Ceydo Agency Head / Sponsor	SU Sergio Uppals Data Scientist / AI Vendor	JN Jari Nurks Legal Counsel & Privacy Officer	ZT Zvitlana Tamo Chair, Citizen Advisory Board
01 Defining Algorithmic Use Case & Scope	R	A	C	C	I
02 Algorithmic Impact Assessment (AIA)	A	I	R	C	C
03 Bias Testing & Data Representative Audit	A	I	R	C	I
04 Human-in-the-Loop Override Protocol	R	A	C	C	I
05 Final System Deployment Authorisation	C	A	I	R	I
06 Continuous Audit Trail & Drift Monitoring	A	I	R	I	C

Defining Algorithmic Use Case & Scope

R Responsible: does the work **A** Accountable: the buck stops here **C** Consulted: provides input **I** Informed: kept updated

Built by **Tam Abaku**

The Five Things You Do By Monday

Immediate Action Items for Project Leaders

01

Classify your highest-stakes AI system under the EU AI Act.

02

Name your SRO and AI Ethics Lead for active projects today.

03

Add a fairness acceptance criterion to your next high-stakes user story.

04

Schedule a training data bias audit before the next model run.

05

Verify if a citizen can contest decisions in plain language.



Let's Connect & Open Q&A

Did We Meet Our Objectives Today?



Open Floor

What challenges do you foresee implementing the 5 gates?



Framework Checklist

[Download the full M.O.R.E implementation guide.](#)



Contact

Email:
tamabaku.et@proton.me



Network

