Sustainable AI for Nuclear

Operationalizing AI Governance for Regulatory Trust

Pravallika (Pravi) Devineni, *Ph.D.*
Lead AI Scientist

DUKE ENERGY
BUILDING A SMARTER ENERGY FUTURE®

Who am I?

- Call me **Pravi**.
- Lead AI Scientist at Duke Energy
- AI Product Manager for nuclear business unit
- Driven by using AI to make life easier
- Ph.D. in CS from U. California Riverside
- Previously, Research Scientist at Oak Ridge National Lab

- **Goal**: Co-create safe and reliable AI systems

What is this about?

"AI isn't just software. It's a system that learns, adapts, and depends on context."

Why This Matters

- AI is moving from pilots → **functional workflows**
- Traditional software assurance does not apply to learning systems
- **Shared goal:** safe, secure, trustworthy systems over years, not months

Agenda

- The AI Workflow
- Sustainable AI
- The AI Iceberg
- DE’s AI Lifecycle controls
- Let’s collaborate!
- Takeaways

The AI Workflow

Software

data + rules = model

- Linear development
- Deterministic
- Rule based
- Explicitly programmed
- Fixed for same input
- Versioned changes

AI

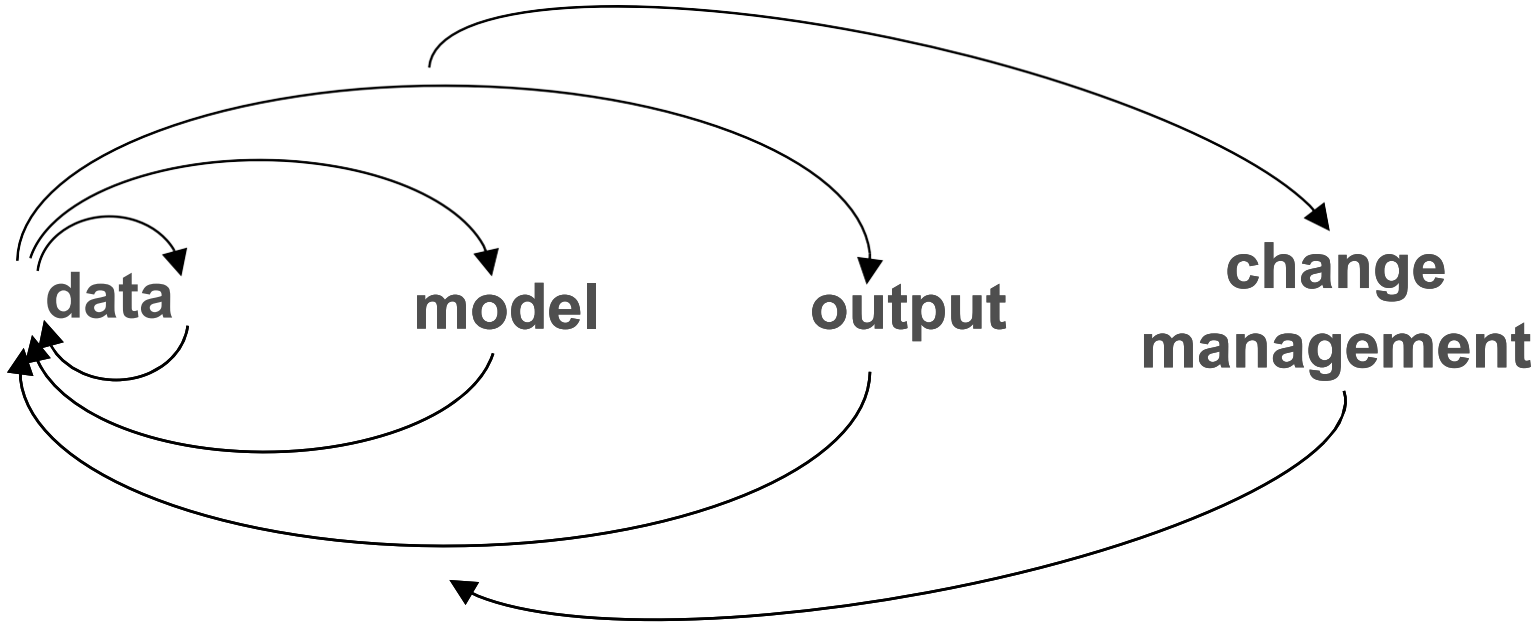

- Iterative development
- Probabilistic
- Learns from data
- Might vary across runs
- Often robust
- Model drift
- Requires monitoring and retraining

Sustainable AI

Supportable

What: AI must be owned, monitored, and have safe fallbacks

Why: Without clear oversight, systems can drift quietly into risk

Reproducible

What: Every AI decision can be replayed exactly as it happened.

Why: In nuclear, traceability is essential for trust and accountability

Auditable

What: AI decisions must be explainable, traceable, and tied to a moment in time

Why: Regulators & operators must be able to verify decisions

Aligned

What: AI must act consistently with human intent and nuclear safety principles

Why: Alignment ensures AI reinforces, rather than undermines, safety culture

The reality of successful AI

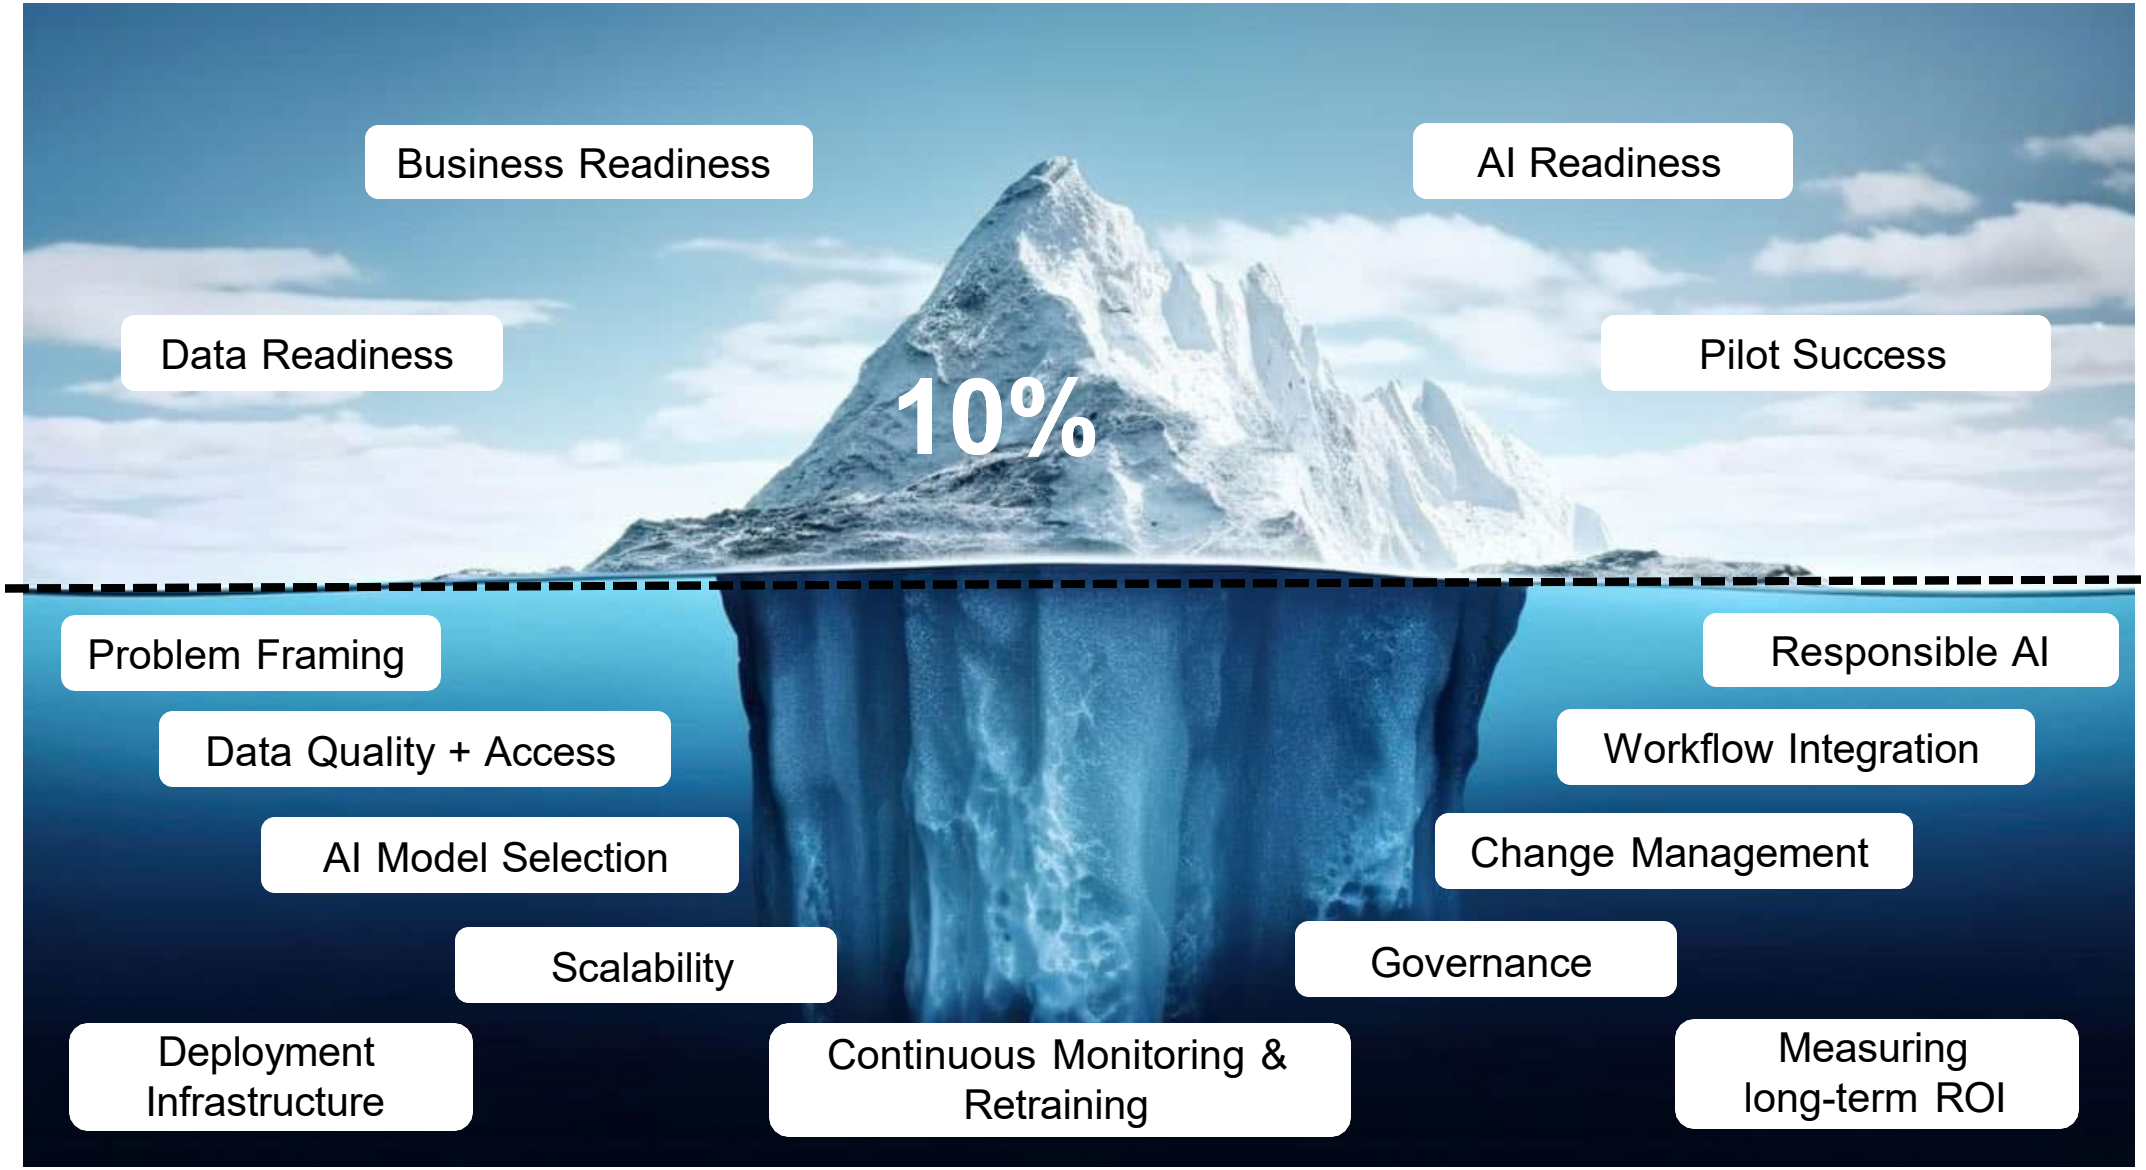

Duke Energy's AI Lifecycle Controls

1 Intake

Choosing wisely

- Should this be AI?
- Clear ownership
- Decide accountability

2 Build

Designing Responsibly

- Reliable data
- Track what's built
- Plan traceability

3 Validate

Testing before trust

- Stress-test edge cases
- Check safety limits
- Confirm with experts

4 Deploy

Rolling out safely

- Start small
- Phase in gradually
- Human review

5 Monitor

Watching over time

- Track performance
- Flag unusual behavior
- Roll back if needed

6 Retrain/ Decommission

Rolling out safely

- Approve updates
- Record changes
- Retire cleanly

The Ask: Co-defining Auditability with NRC

Draft audit packet (proposal)

- **Design intent:** scope/limits, risk tier, misuse risks
- **Reproducibility:** data snapshots, model/feature hashes, env pinning
- **Explainability:** stored decision traces + summaries
- **Lifecycle:** monitoring performance, drift alerts, retraining approvals, rollback records

Takeaways

- AI isn’t software; it changes over time.
- Sustainable AI is our north star: supportable, reproducible, auditable, aligned.
- At Duke, governance is built in, not bolted on.
- Trust comes from evidence — AI must be auditable.
- The future isn’t defined yet; we need to shape it together.

.

Questions

Pravi Devineni, *Ph.D.*

Lead AI Scientist

Duke Energy

Pravallika.Devineni@duke-energy.com

https://www.linkedin.com/in/pdevineni/

BUILDING A **SMARTER** ENERGY FUTURE®

Example: AI Artifact Audit Matrix

Area	Artifact
Design Intent	Use-case intake, scope/limits, risk card
Data/Model	Lineage graph, dataset snapshots (hash), model registry
Validation	Bias/safety/adversarial test results, acceptance memo
Deploy	Signed artifacts, approvals, phased rollout results
Monitor	Drift dashboards, alert/incident logs
Change	Retraining change requests, rollback records, Segregation of Duties evidence
Explainability	Decision replays, reason codes, Explainable AI faithfulness