Coffee & Connections

Rethinking Security in the Age of AI

A practical, interactive conversation about what AI can do, what it cannot do, and how to integrate it into workflows responsibly.

Massimiliano Albanese
Professor & Associate Chair for Research, IST
Executive Director, IDIA
Director, ISEAI

An interactive conversation, not a lecture

1. AI reality check

What AI is, what it is not, and why it sometimes looks more capable than it is.

2. Workflow examples

Life sciences, discovery, business operations, and security use cases.

3. Hype vs. reality

Where AI helps, where it fails, and what should never be automated blindly.

4. Do / Don't

Practical rules for adoption.

ASK THE ROOM

Where are you today: curious, experimenting, piloting, or already using AI daily?

What AI is ... and what it is not

AI is a broad field. Machine learning is one family of methods. Generative AI is one very visible recent capability.

AI

Systems that perform tasks associated with intelligence: classification, prediction, generation, optimization, and decision support.

ML

A subset of AI where systems learn patterns from data rather than being explicitly programmed for every rule.

GenAI

Models that generate text, code, images, plans, summaries, and other outputs from learned patterns and prompts.

Key point: impressive output is not the same as truth, judgment, or accountability.

Impressive? Yes. Does it make sense?

Can AI reason? A practical answer

It can approximate reasoning patterns from language and data.

It does not inherently possess grounded understanding, goals, or causal models.

It becomes more reliable when grounded, constrained, and verified.

ASK THE ROOM

When would you trust AI output in your organization? When would you refuse to?

Hype vs. reality

Hype

AI will replace expertise.

AI knows the answer.

AI can safely automate anything.

AI removes the need for governance.

Reality

AI can augment expertise.

AI proposes outputs that need context.

Automation requires boundaries and monitoring.

Governance becomes more important, not less.

The adoption question is not “Can we use AI?”
It is “Where can AI create value safely?”

Concrete example: AI in scientific discovery

AI accelerates

Hypothesis generation
Search over large design spaces
Pattern detection

Science validates

Experiments
Causal interpretation
Expert judgment

Bottom line

AI can make discovery faster, but predictions become science only when tested and interpreted in the real world.

ASK THE ROOM

What is the equivalent of "experimental validation" in your business process? A pilot?

A controlled test? A compliance review? A security assessment?

The responsible workflow pattern

Security and governance should be designed into the workflow, not added after adoption.

ASK THE ROOM

Where is the human checkpoint in your current AI workflow?

When AI enters a workflow, the attack surface changes

Data leakage

Sensitive data, customer data, proprietary code, contracts, and internal deliberations may leave controlled environments.

Prompt and output manipulation

Inputs can steer systems toward unsafe or misleading outputs; polished answers can hide errors.

Automation risk

A small error becomes a large incident when AI output triggers downstream actions automatically.

Misuse and deception

Deepfakes, phishing, fraud, and social engineering become cheaper and more scalable.

Do this / Avoid this

Do this

Start with specific workflows.

Define acceptable data and prohibited data.

Keep a human accountable for consequential decisions.

Test outputs against known examples.

Document prompts, assumptions, and limits.

Monitor after deployment.

Avoid this

Do not paste sensitive data into unmanaged tools.

Do not automate approvals without controls.

Do not treat AI confidence as evidence.

Do not skip validation because the output sounds polished.

Do not ignore security, legal, and operational stakeholders.

A simple adoption maturity ladder

Move up only when value, risk, governance, and workforce readiness move together.

Start small, but do not stay informal.

Make security part of the operating model.

Measure impact: quality, speed, risk reduction, user adoption.

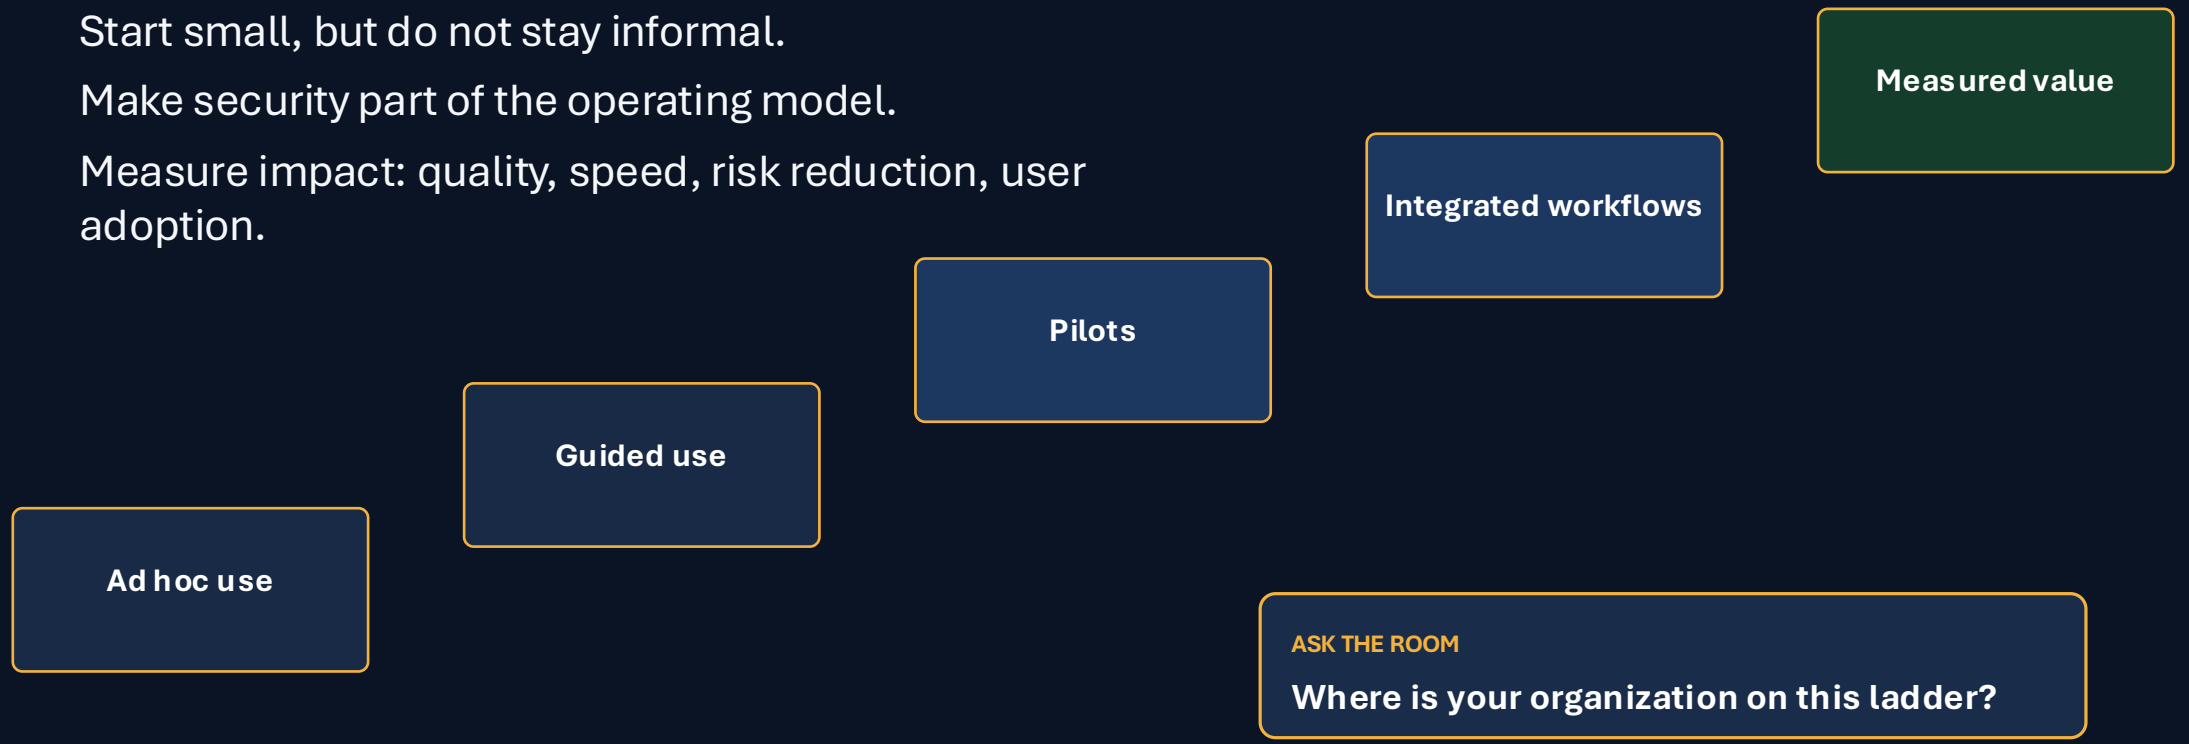

ASK THE ROOM

Where is your organization on this ladder?

Concrete examples across Nexus234 priority industries

Life sciences & forensics

Discovery support, evidence triage, literature synthesis, lab workflow documentation, and quality checks.

Aerospace & defense

Threat intelligence, mission planning support, maintenance analytics, and secure human-machine decision support.

Semiconductors & electronics

Design review assistance, supply-chain risk analysis, defect classification, and IP protection workflows.

Data centers, AI & cyber

Energy-aware operations, incident response, anomaly detection, and resilience planning.

Same adoption pattern, different domain constraints.

ASK THE ROOM

Which example is closest to your work?

ISEAI: why George Mason is building this center

Mission

Advance secure and resilient infrastructure in the era of AI through interdisciplinary research and external partnership.

Research

AI-enabled cyber security, cyber-physical resilience, threat detection, situational awareness, and human-machine collaboration.

Workforce

Practical AI fluency, applied training, student engagement, and pathways into high-demand roles.

Partnership

A convening platform for industry, government, and academia to work on shared problems without getting lost in the funding mechanics.

Why now?

AI amplifies attacks
Resilience is the goal
Infrastructure raises stakes
No one can solve this alone

College of Engineering and Computing
CENTER FOR INFRASTRUCTURE SECURITY IN THE ERA OF AI
George Mason University®

ISEAI: why George Mason is building this center

Towards AI-Driven Human-Machine Co-Teaming for Adaptive and Agile Cyber Security Operation Centers

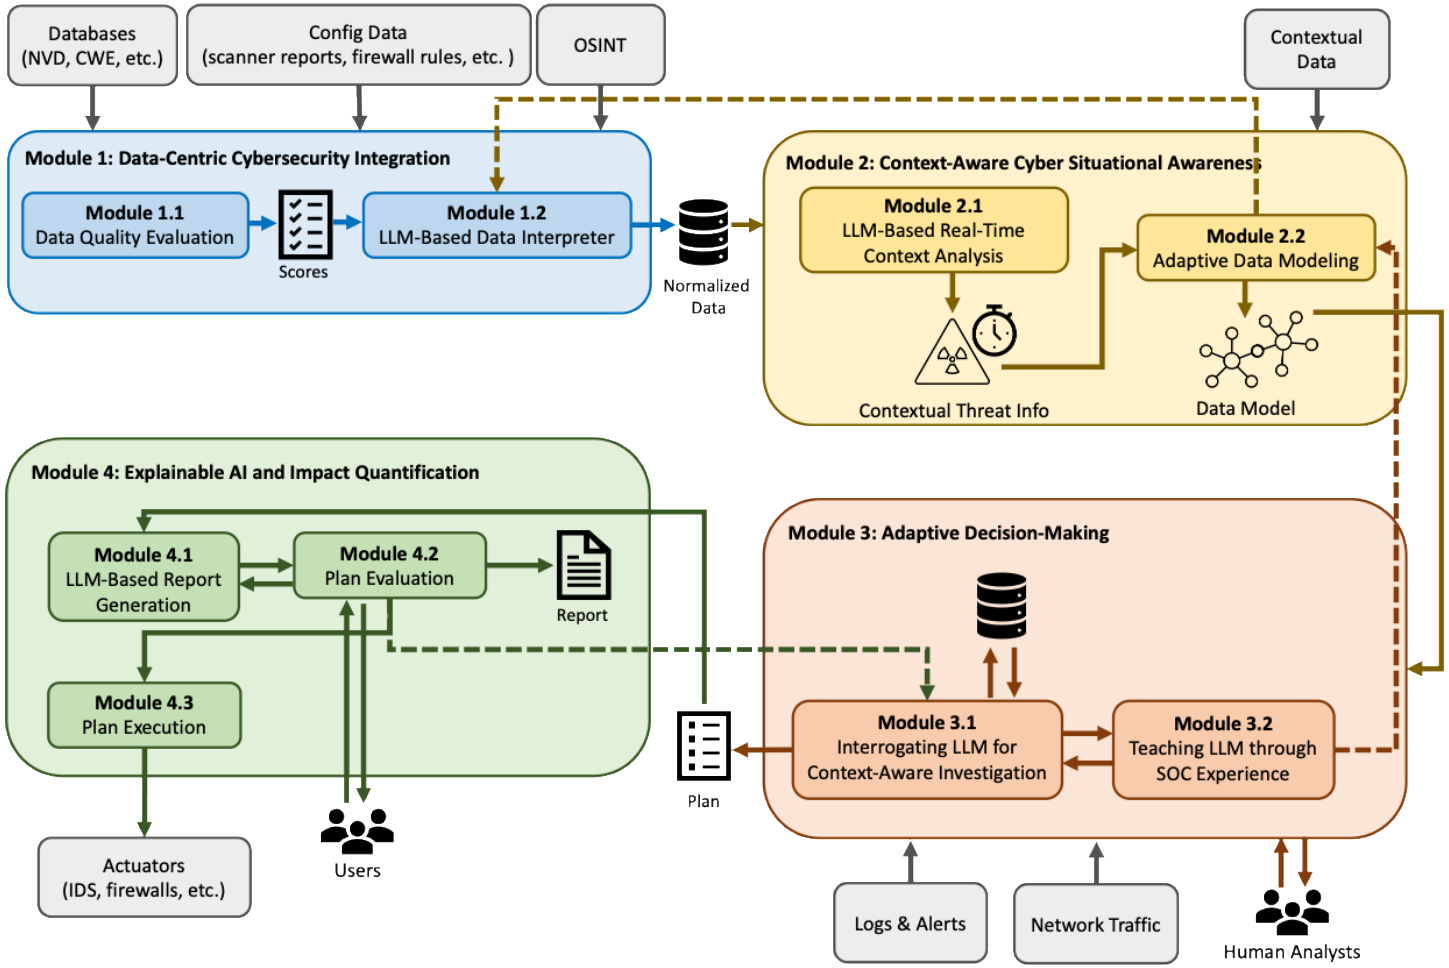

GEORGE MASON UNIVERSITY

UNIVERSITY of SOUTH FLORIDA

READ OUR PAPER

Towards AI-Driven Human-Machine Co-Teaming for Adaptive and Agile Cyber Security Operation Centers

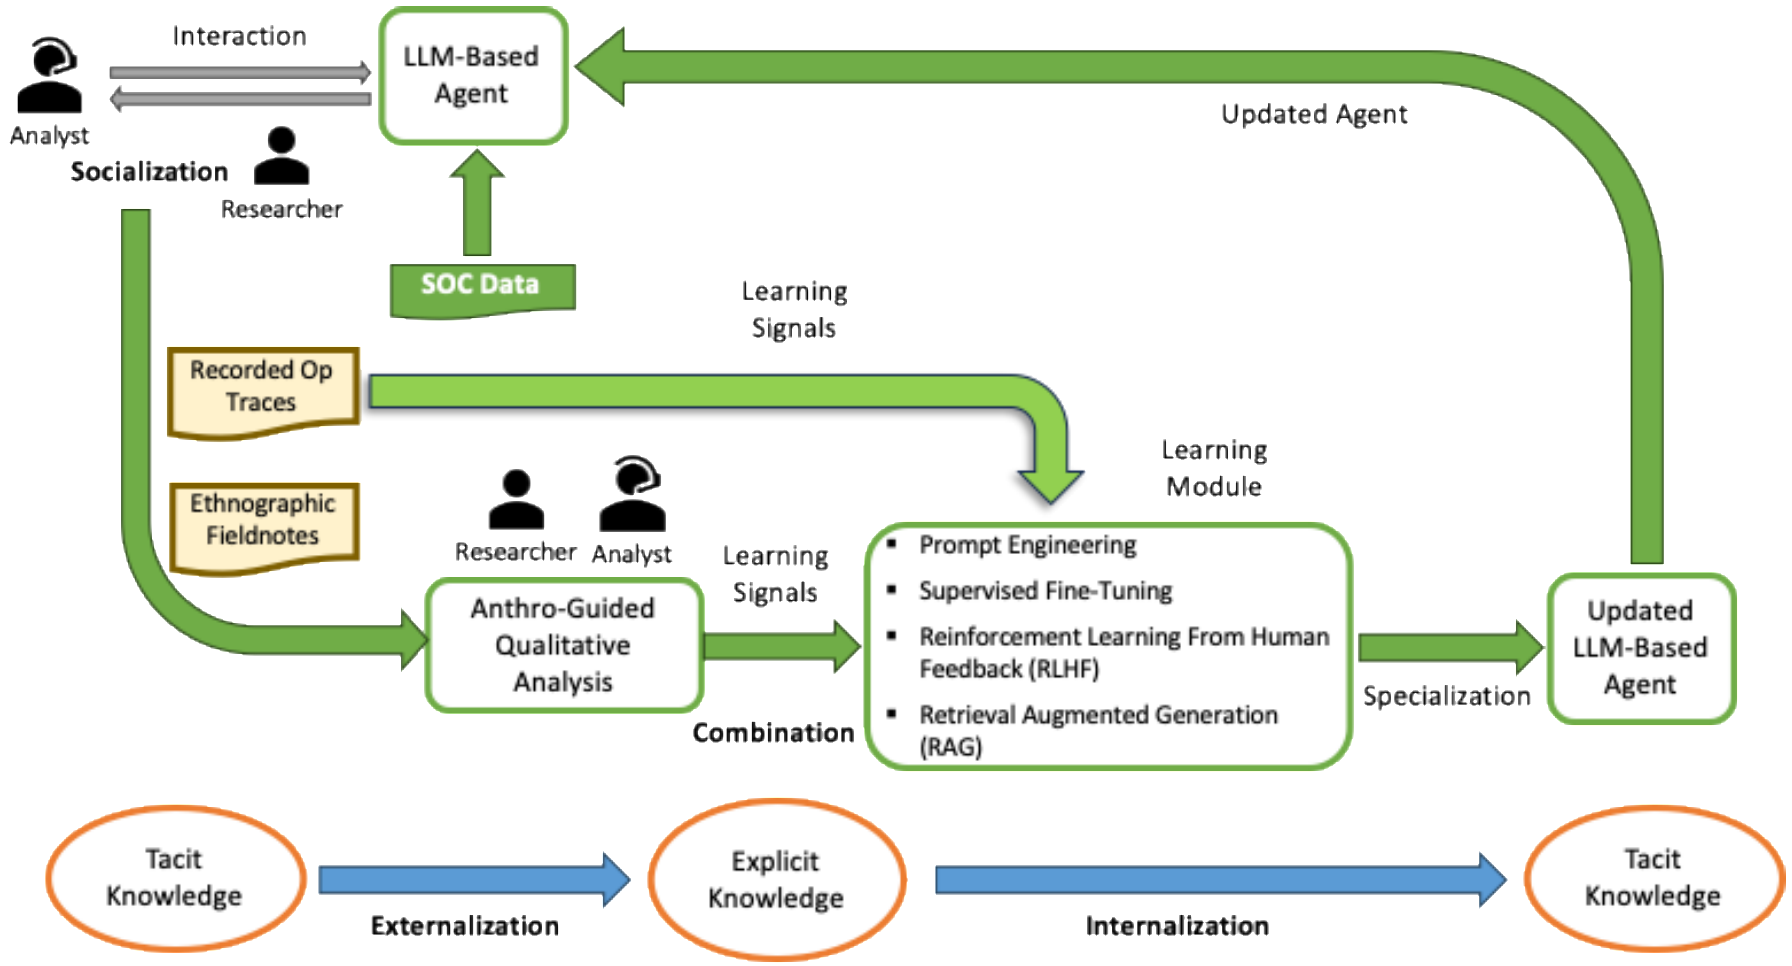

Project NoéMI: AI fluency as a security strategy

Project NoéMI™ by newpush®

Explorers

Everyone develops baseline AI fluency: what to delegate, what to verify, and what not to expose.

Practitioners

A smaller group builds and maintains AI-enabled workflows under supervision and governance.

Accelerators

Leaders define oversight, risk management, ethics, ROI, and organizational guardrails.

AI adoption is not only a technology problem; it is a workforce and governance problem.

Interactive mini-exercise: pick one workflow

Choose a workflow you know well. In small groups, answer four questions:

1. Value

Where could AI save time, improve quality, or reveal useful patterns?

2. Risk

What could go wrong if the output is wrong, biased, leaked, or misused?

3. Validation

What evidence would make the output trustworthy enough for use?

4. Control

Who approves, monitors, and remains accountable?

ASK THE ROOM

Share one workflow where AI is promising — and one guardrail you would require.

Five takeaways

1. Start with the workflow, not the tool.
2. Treat AI output as a proposal, not proof.
3. Keep sensitive data inside controlled environments.
4. Use humans where judgment, accountability, or safety matters.
5. Build fluency, governance, and security together.

Responsible AI is not slower AI.
It is AI that can scale without creating avoidable harm.

Questions?

What should your organization do next to adopt AI safely and effectively?

Massimiliano Albanese
malbanes@gmu.edu
https://iseai.gmu.edu/