COURSE SYLLABUS

Human–AI Integration in Government Workflows

Participants will work through real-world scenarios and practical frameworks to evaluate when and how to keep “humans in the loop,” foster effective collaboration, and ensure accountability. The course emphasizes actionable strategies for building trustworthy hybrid systems and meeting evolving policy and operational requirements.

Group classes (Live Online) and onsite training are available for this course. For more information, email onsite@graduateschool.edu or visit: https://sdfm.graduateschool.edu/courses/human-ai-integration-in-government-workflows



CustomerRelations@graduateschool.edu • (888) 744-4723

Course Outline

Module 1: Understanding AI Agents

- What AI agents are and how they differ from automation, chatbots, and decision-support tools
- Core characteristics of AI agents: goal-driven behavior, multi-step execution, autonomy, and context-awareness
- The agentic AI landscape today: desktop agent tools, enterprise agent platforms, developer tools, and AI assistants with tool access
- How agentic tools work under the hood: the AI model, planning loop, tool access, and containment boundary
- Live demonstration of an agentic tool performing a multi-step government workflow task
- Human oversight roles: human-in-the-loop, human-on-the-loop, and human-in-command
- Common misconceptions and key risks, including prompt injection
- Current federal AI policy direction and its implications for government agencies

Module 2: Evaluating Agentic Tools for Government Work

- The evaluation mindset: understanding the workflow before evaluating the tool
- Identifying agent roles in a workflow: intake, analysis, recommendation, and escalation
- The Agent Evaluation Blueprint: goal, trigger, inputs, actions, boundaries, and oversight
- Permission models and data access: folder-level vs. organization-wide, network access, and least privilege
- Prompt injection: what it is, how it works, and why agentic tools are uniquely vulnerable
- Questions to ask before saying yes: a practical pre-approval checklist
- Hands-on activity: evaluate a realistic agentic tool proposal for a government workflow

Module 3: Where Agentic Tools Fit — and Where They Don't

- Appropriate use cases: case triage and routing, document analysis and extraction, internal coordination, and monitoring and alerts
- High-risk or inappropriate uses of agentic tools in government
- Operational risks: over-automation, over-reliance, and rubber-stamping

- Drift: data drift, concept drift, and objective drift — and how to detect them
- Early operational warning signs that an agentic tool may be failing
- Hands-on activity: red team a deployed agentic tool scenario to identify risks and recommend action

Module 4: Governing Agentic Tools in Your Organization

- Why governance is essential — and why no centralized federal AI regulator is coming
- Governance vs. technical controls: policies, oversight bodies, and accountability structures
- Legal, ethical, and procurement considerations: FedRAMP, ATO, vendor data handling, and records retention
- Security for agentic tools: access controls, prompt injection defenses, anomaly monitoring, and incident response
- Lifecycle management: design, pilot, deploy, monitor, update or retire — including regulatory sandbox alignment
- Performance monitoring and metrics: accuracy, override rates, equity indicators, and user satisfaction
- Workforce readiness and change management: role clarity, training on tool limitations, and avoiding fear and over-trust
- Deployment readiness checklist: a practical gate review before any agentic tool goes live
- Hands-on activity: conduct a readiness gate review for a proposed agentic tool deployment