

PURDUE
UNIVERSITY®

Executive Integration Packet

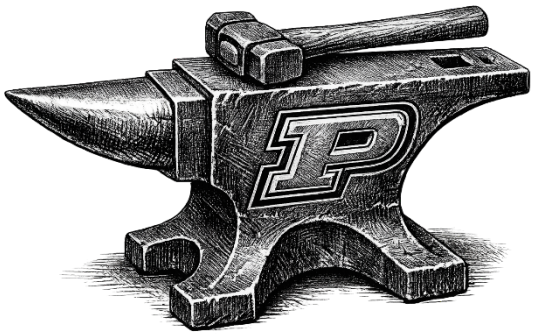
Operations Executive Program 2.0

September 2026

(this page intentionally left blank)

The Learner Experience

Throughout each day, learners will progress through four primary modules spanning more than a dozen sessions. The anvil graphic symbolizes areas of high importance to the executives completing the program. A key to the program is that each day learners examine ways to apply what they've learned and to self-assess their strengths and weaknesses in each topic area.



End-to-End Leadership	This module will empower leaders to influence strategically, align and mobilize process owners, and consider enterprise-wide approaches for people, processes, and projects.
Financial Insight	This module will equip leaders with the skills to analyze market position, make informed P&L decisions, and build compelling business cases that drive strategic and operational success
The AI-Ready Executive	This module will introduce foundational AI concepts and examine how operations leaders can prepare themselves for AI implementation for individuals, teams, and across the enterprise.
Data Interrogation & Fluency	This high-impact module teaches leaders to ask better questions, assess quality of analysis, and build clear, aligned processes.
Fundamental Future of Ops	This module equips participants with modern process leadership skills through sessions on emerging process theory, leading process owners horizontally, and managing end-to-end client journey processes. Learners will gain practical tools to drive cross-functional alignment and elevate process ownership across the organization.

PURDUE UNIVERSITY

Contents

The Learner Experience....3
Learning Modules, Outcomes, & Session Descriptions....5
Day 1....6
Leadership: Executive Leadership: Putting the "E" in OEP....7
AI 1: Beyond the Buzzword: Math not Magic WORKSHEETS & RESOURCES....8
AI 1: Glossary of AI Terms....10
AI 1: Copilot: Executive Summary....16
AI 1: Executive Reference Guide: AI Architecture & Model Comparison....18
AI 1: Beyond the Buzzword: Math not Magic....20
AI 2: AI for the Bank: The Internal AI Landscape....22
Day 2....23
Finance: Market Analysis: Understanding Market Position and the Cost of Inaction....24
Data 1: Questions and Elements: Interrogation Checklist....25
Data 1: Questions and Elements: Interrogating Data....26
Data 2: Process Intelligence: Streamlining with Structure & Smart Technology....28
Day 3....29
Process 1: Client Focused Process Design: Human-Centered Design Thinking....30
Fair Oaks Farms: Farm History, Innovation, and Data....31
Day 4....32
Process 2: Process Thinking: Emerging Process Theory....33
Process 3: Process Thinking: Executive Leadership of Processes & Process Owners....34
Data 3: The Next Chapter: Advanced Storytelling with Data....37
Day 5....38
Process 4: Tactical Management of Processes and Process Owners....39

Learning Modules, Outcomes, & Session Descriptions

Leadership	**Driving E2E Continuous Improvement with Enterprise & Ownership Mindset** Learning Outcome Drive end-to-end continuous improvement by fostering enterprise thinking, strengthening ownership, and breaking down silos to enable system-wide accountability and collaboration.
Finance	**Driving Performance: Financial Insight, Market Analysis, & Intelligent Operations** Learning Outcome Enhance performance by applying financial acumen, market and enterprise metrics, and predictive operational tools to optimize efficiency, resource allocation, and competitive positioning
Artificial Intelligence	**The AI Ready Executive: AI Fundamentals, Applications, & Strategy** Learning Outcome Develop AI readiness by mastering foundational concepts, evaluating use cases, aligning with enterprise strategy, and leading responsible adoption to drive business value.
Data	**From Ambiguity to Action: Data Fluency for Executive Decision-Making** Learning Outcome Build data fluency by defining problems, asking the right questions, analyzing data-driven insights, and translating findings into clear actions and measurable outcomes.
Process	**The Fundamental Future of Ops: Trends and models for E2E Process Design & Leadership** Learning Outcome Lead the future of operations by applying emerging technologies, process innovation frameworks, and end-to-end design principles to optimize performance and drive enterprise transformation.

Day 1

Driving E2E Continuous Improvement with Enterprise & Ownership Mindset

Leadership 1

Description:
Empower teams to own their work, strengthen cross-functional collaboration, and align processes to deliver consistent business results.

Session Objective(s): Define your leadership role to empower team ownership, recognize and overcome barriers to enterprise thinking, and implement strategies and action plans that foster system-wide accountability, cross-functional collaboration, and value-added executive impact.

Faculty Lead: Dr. Brad Alge

BofA Video: Tom Scrivener

The AI Ready Executive: AI Fundamentals, Applications, & Strategy

Artificial Intelligence 1

Description:
Understand core AI concepts, explore real-world use cases and governance, and apply design thinking to integrate AI into processes for smarter, more effective operations.

Session Objective(s): Build foundational AI literacy by understanding core concepts, terminology, and use cases, critically evaluating outputs for quality and risk, and establishing a practical framework for responsible and effective AI adoption.

Faculty Lead: Mr. Mat Trampski

The AI Ready Executive: AI Fundamentals, Applications, & Strategy

Artificial Intelligence 2

Description:
Understand the needs, applications, and implementation goals for AI that fits the Bank's strategic goals.

Session Objective(s): Understand how to responsibly leverage AI to amplify strengths without over-reliance, align with the bank's vision for AI and Copilot adoption, and foster a culture that embraces and sustains AI-driven ways of working.

Bank of America Lead: Mr. Hank Weaver

Leadership: Executive Leadership: Putting the "E" in OEP

Individual-Reflection (5-minutes):

Table Discussion (15-minutes):

Large-Group Open Discussion (10-minutes):

Notes:

PROMPT #1



Identify one or two opportunities or challenges in your everyday role where an AI solution might be leveraged for some improvement?

Challenge/Opportunity #1 –

Challenge/Opportunity #2 –

PROMPT #2



Label your challenges or opportunities as probabilistic, deterministic, or both. (circle one)

Challenge/Opportunity #1 – probabilistic deterministic both
Challenge/Opportunity #2 – probabilistic deterministic both

PROMPT #3



Answer the following questions:

Regarding your challenges and opportunities, where can AI create the most value?

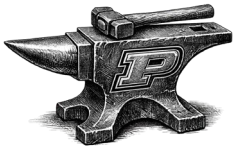
At what level of the organization is your challenge and opportunity focused (individual, team, enterprise)?

PROMPT #4



Create the first draft of an idea, project, or concept to share with your leadership that addresses your challenge or opportunity, in which AI (or maybe not AI) could be leveraged for some functional improvement. Write out the first stages of a problem statement and an informal request for support.

Additional Notes :



probabilistic vs deterministic

Activity

Compare & Contrast Probabilistic vs Deterministic Systems

In this exercise, you will work with peers at your table group.

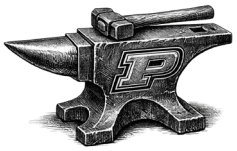
- Analyze and answer questions based on prompts and scenarios.
- Be prepared to discuss your answers.

Scenario Analysis

Analyze the following two scenarios. Determine which system type is being described and justify your choice based on how the system arrives at its conclusion.

- **Scenario A:** A loan processing tool that automatically rejects any applicant with a credit score below 650, regardless of other factors.
- **Scenario B:** A fraud detection tool that flags a transaction because it "shares 85% of the characteristics" of known money-laundering patterns in the current economic climate.

QUESTION: Which scenario represents a deterministic system, and which is probabilistic? Justify your answer.	**Answer:**
QUESTION: How does the shift from “rules” to “inference” change how you would explain rejection to a customer?	**Answer:**



Activity

- *Continued*

probabilistic vs deterministic

Trade Offs: Rule-based Determinism vs. Probabilistic Prediction

Traditional business processes, such as automated payroll systems or inventory management workflows, typically rely on deterministic logic: if the employee worked X hours at Y rate, the paycheck is Z; if stock falls below a specific threshold, a reorder is triggered. These systems operate on strict, pre-defined rules where the same input always produces the same output. Conversely, AI systems are probabilistic, meaning they use statistical inference and pattern recognition to generate the most likely "best" outcome based on high-dimensional data patterns.

QUESTION: In high-stakes organizational operations, when is the "certainty" of a deterministic process a liability, and when is the "flexibility" of a probabilistic system a risk?	**Answer:**
QUESTION: Identify a standard operating procedure in your firm where "100% consistency" might actually be hindering innovation or responsiveness to market shifts.	**Answer:**

AI 1: Glossary of AI Terms

Agent: An autonomous AI entity or program designed to perceive its environment, make independent decisions, and execute multi-step actions using external tools to achieve a specific, overarching goal without requiring constant human intervention.

Agentic AI: A more autonomous class of AI systems ("agents") that don't just answer questions, but can plan multi-step actions, use external tools, and make independent decisions to achieve a specific goal.

Algorithm: A step-by-step set of instructions or rules given to a computer to complete a task or solve a problem.

AP (Application Programming Interface): A software intermediary that allows different applications to talk to each other. In this context, it is the bridge used to connect your existing enterprise software to external or internal AI models.

Artificial Intelligence (AI): The broad field of computer science dedicated to creating systems capable of performing tasks that typically require human intelligence, such as reasoning, learning, and problem-solving.

Auditability: The ability to reconstruct and review the decision-making process of an AI system. This includes tracking model versions, training logs, prompt histories, and system outputs to satisfy regulatory compliance and internal security audits.

Backpropagation: The process of sending error signals backward through a neural network during training. It uses the mathematical chain rule to calculate how much each individual weight contributed to the final error, allowing the model to adjust itself accurately.

Bias: A systematic error where an AI model favors certain outcomes, demographics, or viewpoints over others, usually because the training data it learned from was skewed or unrepresentative.

Chain-of-Thought (CoT): A prompting technique that explicitly instructs the AI to break down complex problems and show its step-by-step reasoning process before giving the final answer, which significantly improves logic and math accuracy.

Context Window: The maximum amount of text (measured in tokens) that an AI model can read, remember, and process at one single time during a conversation or task. If a conversation exceeds this limit, the model starts "forgetting" the earliest parts of the chat.

Controls / Guardrails: System-level restrictions and filters placed around an AI model to restrict its inputs and outputs. They prevent AI from generating inappropriate content, leaking confidential data, or straying from its intended business function**.**

Data Lineage: The end-to-end lifecycle of data, tracking its origins, movements, and transformations over time, critical for proving compliance, debugging model biases, and verifying the quality of training data.

Deep Learning: A specialized type of Machine Learning that uses deep neural networks (networks with many layers) to analyze highly complex data like video, audio, and raw text.

Deterministic: The opposite of probabilistic. A deterministic system follows rigid rules where a specific input will always produce the exact same output (like a traditional calculator or standard code).

Embedding: The process of converting text, images, or data into a string of numbers (a vector) that captures its deeper conceptual meaning. This allows AI systems to understand relationships and similarities between different pieces of information, rather than just matching literal keywords.

Encoder-Decoder: A classic two-part structural framework in neural networks. The Encoder compresses the input data into a dense mathematical representation, and the Decoder reads that math to generate a brand-new output (like translating English text into French).

Explainability / Interpretability: The practice of designing artificial intelligence models so that humans can understand how and why they make specific decisions or predictions.

Few-Shot Prompting: A technique where you provide the AI with a few examples of the desired input and output format within the prompt itself before asking it to complete the actual task.

Fine-Tuning: The process of taking an existing, pre-trained AI model and training it further on a smaller, specialized dataset. This adapts the model to excel at specific corporate tasks, industry jargon, or unique business logic.

Generative AI: AI models designed to create entirely new content, such as text, images, code, audio, or video, by predicting the most likely sequence of data based on what they've learned.

Gradient Descent: The primary optimization algorithm used to train AI. It uses calculus to calculate the direction of error and takes small, iterative steps "downhill" to find the point of lowest possible error.

Hallucination: When an AI model confidently generates an answer that sounds plausible but is completely factually incorrect or fabricated.

Heuristics: Practical, rule-of-thumb strategies or shortcuts used to solve problems quickly when finding an absolute perfect solution is impossible or too slow.

Inference (Phase): The operational phase where a fully trained AI model is put to work in the real world, taking in new, unseen data and outputting a prediction or response.

Large Language Model (LLM): A massive neural network trained on vast amounts of text data. It understands, summarizes, generates, and predicts human language (e.g., Gemini, GPT-4).

Latency: The time delay between a user submitting a request (a prompt) and the AI system delivering the response. Low latency is critical for real-time user experiences like customer service chatbots.

Loss Function: A mathematical formula that measures the "wrongness" of an AI's prediction. It compares the model's guess to the actual truth and outputs a penalty score; the higher the score, the worse the guess.

Machine Learning (ML): A subset of AI focused on building systems that learn from data and improve their accuracy over time without being explicitly programmed for every scenario.

Matrix: A two-dimensional grid of numbers (rows and columns). Inside an AI, matrices are used to store massive blocks of data and the "weights" (connection strengths) of a neural network.

Model Drift: The degradation of an AI model's performance or accuracy over time, often because real-world data patterns have changed since the model was last trained**.**

Neural Network: An AI architecture inspired by the human brain. It consists of layers of interconnected digital "neurons" that pass data to one another, processing complex patterns to make a final prediction.

PII (Personally Identifiable Information): Any data that can be used to distinguish or trace an individual's identity (e.g., SSNs, names, email addresses). Protecting PII is a primary legal requirement when handling enterprise AI datasets.

Predictive AI: A subset of artificial intelligence that uses historical data, statistical modeling, and machine learning to identify patterns and forecast future outcomes, trends, or behaviors (e.g., predicting equipment failure, customer churn, or financial market trends).

Probabilistic: Systems driven by probability and likelihood rather than absolute certainty. AI models are probabilistic; they don't look up a fixed answer but instead calculate that a certain word or label has a, say, 92% chance of being the correct output.

Prompt Engineering: The practice of carefully designing, structuring, and phrasing prompts to get the most accurate, relevant, and high-quality responses from an AI model.

Prompt: The natural language input, question, or instruction given to an AI model to guide its output.

Responsible AI: A strategic framework for developing and deploying AI systems that are ethical, transparent, unbiased, and accountable. It ensures AI alignment with corporate values, legal requirements, and societal norms.

Retrieval-Augmented Generation (RAG): An architecture that optimizes LLM outputs by querying an external, secure knowledge base (like internal company documents) to fetch up-to-date, information before generating a response. It significantly reduces "hallucinations" without the cost of retraining a model.

Robotic Process Automation (RPA): Software technology that automates repetitive, rules-based digital tasks by mimicking human interactions with computer systems (such as copying and pasting data, moving files, or filling out forms).

Sensitive Data Handling: The protocols, security controls, and governance practices required to protect confidential information (like financial records, intellectual property, or PII) from being exposed or inadvertently used to train public AI models.

Small Language Model (SLM): A highly optimized, compact alternative to an LLM. These models have fewer parameters but are highly capable, designed to run efficiently on local hardware like smartphones or laptops without needing massive cloud data centers.

Tensor: A multi-dimensional grid of numbers. If a vector is 1D (a line) and a matrix is 2D (a flat grid), a tensor can be 3D (a cube), 4D, or higher. It is the ultimate container for complex data like color videos or multi-layered network data.

Throughput: The volume of data or number of requests an AI system can successfully process within a specific timeframe (often measured in tokens per second). It dictates how well your AI infrastructure scales under heavy concurrent usage.

Token: The basic unit of data processed by an LLM. A token isn't always a full word; it is often a fragment of a word or a single syllable (e.g., the word "learning" might be split into "learn" and "ing").

Training (Phase): The initial process where an AI model is exposed to a dataset, calculates its errors, and adjusts its mathematical parameters to learn a task.

Vector Database: A specialized database designed to store and quickly search complex data represented as numerical values (vectors). It is the foundational storage engine used to power RAG and semantic search applications.

Vector: In AI, a vector is a structured list of numbers that represents a piece of data (like a word or an image pixel) in mathematical space. It gives data a geometric position so the AI can measure how "close" or related things are.

Weights and Biases: The adjustable parameters inside a neural network. Weights determine how much importance to give a specific input, while biases act like an offset to shift the model's calculations for a better fit.

Zero-Shot Prompting: Asking an AI to perform a task without providing any examples beforehand, relying entirely on the model's pre-existing training to figure out the correct response.

AI 1: Copilot: Executive Summary

What Copilot Is

Copilot is Microsoft's enterprise AI companion that enhances knowledge work by combining large language models with your organization's data and Microsoft's security framework. It is not a chatbot — it is an intelligence layer embedded across Microsoft 365, Windows, and business applications.

How Copilot Works (At a Glance)

1. Understands Your Intent

Copilot interprets natural language prompts and determines what information, tools, or context are needed.

2. Retrieves the Right Data

It pulls from Microsoft 365, Graph-connected systems, and approved web sources — always respecting existing permissions.

3. Generates Grounded Output

Copilot produces summaries, drafts, insights, and recommendations based on retrieved data and organizational policies, delivered directly in the workflow.

Retrieval-Augmented Generation (RAG): Why It Matters

RAG is the architecture that ensures Copilot "looks up the facts before it answers."

- **Retrieve:** Copilot searches trusted internal sources (documents, emails, chats, files, systems).
- **Generate:** It uses that information to create accurate, context-aware responses.

This reduces hallucinations and ensures outputs reflect your organization's actual knowledge.

How Copilot Uses RAG

- Retrieves only the data the user is permitted to access.
- Grounds responses in that data for accuracy and relevance.
- Provides traceability by showing or linking to source documents.

What Makes Copilot Different

Enterprise-Grade Security

Built on Microsoft's security, privacy, and compliance controls — including role-based access, DLP, conditional access, and audit frameworks.

Deep Workflow Integration

Embedded directly into Outlook, Teams, Word, Excel, PowerPoint, Dynamics 365, and Windows — enabling context-rich assistance.

Grounded, Traceable Responses

Copilot uses your data, not generic internet patterns, to produce relevant and defensible outputs.

Responsible AI by Design

Aligned with Microsoft's Responsible AI principles: transparency, fairness, reliability, and accountability.

Extensible Across the Enterprise

Supports plugins, connectors, and custom copilots built with Copilot Studio.

Executive Takeaways

- Copilot is an enterprise-safe AI layer, not a consumer chatbot.
- It combines LLMs with your business data and Microsoft's security architecture.
- RAG ensures accuracy, relevance, and traceability.
- Deep integration across Microsoft 365 drives real productivity gains.
- Copilot accelerates decision-making, reduces cognitive load, and modernizes operations.

Strategic Value for Leaders

Copilot enables organizations to:

- Scale productivity across teams
- Reduce administrative overhead
- Improve decision quality with grounded insights
- Shift employee time toward high-value work
- Drive AI-enabled operational transformation

Copilot turns information into intelligence — and employees into augmented decision-makers.

AI 1: Executive Reference Guide: AI Architecture & Model Comparison

To maximize ROI, leadership must recognize that "Generative AI" is no longer a single tool. Modern AI infrastructure is divided into **five distinct model classes**, each optimized for different cost structures, processing speeds, and cognitive workloads.

NOTE: This document is for awareness and may not be applicable in restricted corporate environments.

The Enterprise Model Matrix

Model Classification	Core Business Capability	Top Market Examples	Strategic Use Case
1. Frontier Generalists	High-nuance text generation, advanced communication, brand voice replication.	• OpenAI GPT-5.5 • Anthropic Claude Sonnet 4.6	Ghostwriting, marketing localization, legal drafting, complex customer engagement.
2. Reasoning & Logic Models	Multi-step problem solving, advanced mathematics, structural engineering.	• OpenAI o3 Pro • DeepSeek R1 • Gemini 3.1 Ultra	Supply chain optimization, financial audit validations, deep competitive strategy simulations.
3. Multimodal & Long-Context	Ingesting and cross-referencing massive, multi-format enterprise data pools.	• Google Gemini 3.1 Pro • Llama 4 Scout (Meta)	Analyzing hours of video, cross-referencing 1,000-page operational manuals with invoices.
4. Agentic Coders	Autonomous software development, codebase debugging, and system integration.	• Claude Opus (4.6/4.7) • GPT-5.5 Codex	Legacy code migration, rapid product prototyping, automated software patching.
5. High-Velocity / Utility	Low-cost, high-speed automated data processing and micro-routing.	• DeepSeek V4 Flash • Gemini 3.1 Flash-Lite • GPT-5.5 mini	High-volume email sorting, basic data extraction, first-tier customer support routing.

Strategic Blueprint: Matching the Model to the Problem

Class 1: Frontier Generalists (The Communicators)

- **What they do best:** These models focus on pattern prediction to generate highly polished, contextually aware text. They are highly empathetic and adaptive in tone.
- **Executive Takeaway:** Use these to elevate the baseline output of marketing, HR, and communications teams. They act as force multipliers for human writers.

Class 2: Reasoning Models (The Analysts)

- **What they do best:** Unlike generalists that guess the next word instantly, reasoning models use internal "Chain-of-Thought" processing to systematically think through an enterprise problem before answering.
- **Executive Takeaway:** Deploy these when accuracy and logic matter infinitely more than speed. Ideal for risk assessment, scientific R&D, and corporate finance.

Class 3: Multimodal & Long-Context (The Researchers)

- **What they do best:** They eliminate the need to pre-format data. They can read audio, video, charts, and text simultaneously across massive data windows (up to millions of words at once).
- **Executive Takeaway:** This is your solution for breaking down corporate information silos. Pass it massive, messy historical archives to find immediate operational efficiencies.

Class 4: Agentic Coders (The Engineers)

- **What they do best:** These models don't just write code snippets; they can act as autonomous agents that navigate entire software repositories to make multi-file updates.
- **Executive Takeaway:** Crucial for tackling technical debt. These models allow your existing software engineering teams to focus on architecture and strategy while automation handles maintenance.

Class 5: High-Velocity / Utility (The Workhorses)

- **What they do best:** Sub-second response times at a fraction of the cost of frontier models (often less than 5% of the API cost).
- **Executive Takeaway:** Essential for operational scalability. Do not waste budget using expensive reasoning models to sort basic text data; route high-volume, low-complexity automation tasks here.

Implementation Rule of Thumb:

Always deploy the leanest, fastest model that safely matches the cognitive complexity of the task. Over-specifying your AI architecture directly erodes project margins.

AI 1: Beyond the Buzzword: Math not Magic

Individual-Reflection (5-minutes):

Table Discussion (15-minutes):

Large-Group Open Discussion (10-minutes):

Notes:

TEMPLATE · DAY 1

Day 1 Capture: Process Design & AI Enablement

KEY CONSIDERATIONS:

- *Notice how the process is framed and bounded.*
- *Observe where work appears to slow, repeat, or be checked.*
- *Pay attention to how responsibility moves through the flow.*
- *Look at how effort, review, and escalation show up in the design.*

DAY 1 EXERCISE:

Take 20 minutes to review the Day 1 process pack artifacts. Then spend a further 20 minutes capturing the questions or observations you would raise with the Process Owner to surface process design and capacity improvement opportunities that would strengthen the effectiveness and management of the process.

No.	Question/Areas of Concern	Rationale

Reflection Question
Where do you see potential opportunities to apply AI to improve the efficiency and/or effectiveness of this process?

PURDUE UNIVERSITY®

AI 2: AI for the Bank: The Internal AI Landscape

Individual-Reflection (5-minutes):

Table Discussion (15-minutes):

Large-Group Open Discussion (10-minutes):

Notes:

Day 2

Driving Performance: Financial Insight, Market Analysis, & Intelligent Operations

Finance

Description:
Assess your institution's competitive standing and strengthen financial acumen by examining the economics of innovation, AI adoption, and operational performance. Interpret financial results, explore the risks and rewards of acting—or waiting—in a shifting market, and identify opportunities to improve efficiency, reduce costs, and optimize resource allocation.

Session Objective(s): Analyze the bank's competitive position using operating leverage, efficiency ratios, and key enterprise metrics such as ROTCE; interpret P&L statements and manage budgets to evaluate OpEx and capital allocation decisions; and apply forecasting, capacity planning, and predictive tools—including AI—to identify efficiency opportunities, optimize resource deployment, and support strategic decision-making.

Faculty Lead: Dr. Kevin Koharki

From Ambiguity to Action: Data Fluency for Executive Decision-Making

Data 1

Description:
Learn to define clear problem statements, craft the right questions, and evaluate data outputs to drive informed, high-impact decisions.

Session Objective(s): Identify and define problems by asking targeted questions to gather the right data, develop data-driven solutions, and evaluate outcomes using clear success metrics and performance measures.

Faculty Lead: Dr. Jim Stratton

From Ambiguity to Action: Data Fluency for Executive Decision-Making

Data 2

Description:
Explore how to simplify and standardize complex processes, uncover inconsistencies, and leverage AI for smarter, faster workflows.

Session Objective(s): Assess processes for ambiguity and consistency while identifying opportunities to apply AI and critically analyzing outputs to drive improvement.

Faculty Lead: Dr. Elcin Gunay

Finance: Market Analysis: Understanding Market Position and the Cost of Inaction

Individual-Reflection (5-minutes):

Table Discussion (15-minutes):

Large-Group Open Discussion (10-minutes):

Notes:

Data 1: **Questions and Elements: Interrogation Checklist**

☐ _____

☐ _____

☐ _____

☐ _____

☐ _____

☐ _____

☐ _____

☐ _____

☐ _____

☐ _____

☐ _____

☐ _____

Data 1: Questions and Elements: Interrogating Data

Individual-Reflection (5-minutes):

Table Discussion (15-minutes):

Large-Group Open Discussion (10-minutes):

Notes:

TEMPLATE · DAY 2

Day 2 Capture: Control, Risk, Metrics & Performance

KEY CONSIDERATIONS:

- *Observe where controls appear to intervene, and at what point in the process they do so.*
- *Pay attention to what the metrics emphasize — and what they remain silent on.*
- *Look for alignment (or misalignment) between risks, controls, tests, and reported results.*
- *Consider how quickly emerging issues would become visible using the current measures.*

DAY 2 EXERCISE:

Take 20 minutes to review the Day 2 process pack artifacts. Then spend a further 20 minutes capturing the questions or observations you would raise with the Process Owner to strengthen control framework, process performance and overall effectiveness.

No.	Question/Areas of Concern	Rationale

Reflection Questions

If you were the executive overseeing this process, what are the key questions you want answered on an ongoing basis about this process? Do the metrics provide answers?

Is this process a continuous or batch flow process? What potential improvements might increase the throughput?

2

PURDUE UNIVERSITY

Data 2: Process Intelligence: Streamlining with Structure & Smart Technology

Individual-Reflection (5-minutes):

Table Discussion (15-minutes):

Large-Group Open Discussion (10-minutes):

Notes:

Day 3

The Fundamental Future of Ops: Trends and models for E2E Process Design & Leadership

Process 3

Description:
Use design thinking to reframe problems, uncover customer insights, and co-create innovative practical solutions that work at scale.

Session Objective(s): Apply design thinking to enterprise challenges by using client-centered approaches and journey mapping to uncover pain points, reframe problems, and drive innovative solutions.

Faculty Lead: Mrs. Jennifer Christie

Fair Oaks Farms

Farm History, Innovation, and Data

Description:
Experience firsthand how a modern, high-tech dairy farm combines innovation, data, and technology to improve decision-making, operational efficiency, and long-term performance. Through an onsite farm experience, explore the evolution of the operation and see how real-time data, automation, and technology are used to monitor systems, optimize resources, and drive continuous innovation.

Session Objective(s): Examine how a modern dairy operation has evolved through the adoption of technology, automation, and data-driven practices. Observe firsthand how data is collected and used to support operational decision-making, improve efficiency, and optimize resources. Explore real-world applications of innovation and technology while considering how similar approaches to data, experimentation, and continuous improvement can be applied within the enterprise. Identify opportunities to translate lessons from the farm experience into new ways of thinking about innovation, operations, and data-informed decision-making.

Faculty Lead: Mrs. Jacqueline McCloskey

Process 1: Client Focused Process Design: Human-Centered Design Thinking

Individual-Reflection (5-minutes):

Table Discussion (15-minutes):

Large-Group Open Discussion (10-minutes):

Notes:

Fair Oaks Farms: Farm History, Innovation, and Data

Individual-Reflection (5-minutes):

Table Discussion (15-minutes):

Large-Group Open Discussion (10-minutes):

Notes:

Day 4

The Fundamental Future of Ops: Trends and models for E2E Process Design & Leadership

Process 2

Description:
Explore the latest advances in process thinking and how they reshape modern organizational workflows.

Session Objective(s): Recognize how emerging and cutting-edge technology can influence process theory.

Faculty Lead: Dr. Pat Brunese

The Fundamental Future of Ops: Trends and models for E2E Process Design & Leadership

Process 3

Description:
Learn practical strategies for influencing and guiding process owners across functional boundaries.

Session Objective(s): Think beyond siloed ownership to evaluate and implement process improvements across upstream and downstream impacts, while understanding horizontal effects and confidently interrogating processes to drive system-wide optimization.

Faculty Lead: Dr. Pat Brunese

From Ambiguity to Action: Data Fluency for Executive Decision-Making

Data 3

Description:
Expand on prior storytelling with data techniques to better craft and communicate complex ideas up- and downstream within the organization with interactive exercises.

Session Objective(s): Develop advanced storytelling skills with data by transforming complex analysis into clear, compelling narratives that drive insight, influence decisions, and align stakeholders through effective visualization and communication.

Faculty Lead: Dr. Doug Pruim

Process 2: Process Thinking: Emerging Process Theory

Individual-Reflection (5-minutes):

Table Discussion (15-minutes):

Large-Group Open Discussion (10-minutes):

Notes:

Process 3: Process Thinking: Executive Leadership of Processes & Process Owners

Individual-Reflection (5-minutes):

Table Discussion (15-minutes):

Large-Group Open Discussion (10-minutes):

Notes:

TEMPLATE · DAY 3

Day 3 Capture: Process Efficiency & Storytelling w/ Data

KEY CONSIDERATIONS:

- *Notice where work appears to change hands, pause, or come back before completion.*
- *Observe which steps seem to consume the most effort relative to the value they create.*
- *Look for signs of exception handling, rework, or special routing becoming the norm.*
- *Reflect on where simplification, automation, or design changes could drive benefit.*

DAY 3 EXERCISE:

Take 15 minutes to review the Day 3 process pack artifacts. Then spend a further 15 minutes capturing the questions or observations you would raise with the Process Owner to meaningfully improve cost, speed, and overall experience.

No.	Question/Areas of Concern	Rationale

Reflection Questions
Where could potential bottlenecks exist in this process? Where might you find a potential Herbie?
Does the dashboard effectively tell the story about the performance of the process? How can it be improved?

3

PURDUE UNIVERSITY®

TEMPLATE · DAY 3

Reflection: Lessons Learned and Actions to Apply

Directive – Take 20 minutes reflect on key learnings from the entire week's program content. Be prepared to speed-share with another teammate – and potentially present to the broader team (1 per table to be selected on Friday).

Day 3 – Part B
Reflect on your most significant learnings and insights gained from the week. What will you will take back into your own organization to apply?

PURDUE UNIVERSITY®

Data 3: The Next Chapter: Advanced Storytelling with Data

Individual-Reflection (5-minutes):

Table Discussion (15-minutes):

Large-Group Open Discussion (10-minutes):

Notes:

Day 5

The Fundamental Future of Ops: Trends and models for E2E Process Design & Leadership

Process 4

Description:
Understand how to tactically lead processes and process owners to create a seamless, end-to-end client journey experience.

Module Objective(s): Develop the ability to tactically manage processes and process owners by driving accountability, optimizing performance, and ensuring alignment with operational goals and enterprise priorities.

Faculty Lead: Mr. Mike Hyzy

Process 4: Tactical Management of Processes and Process Owners

Individual-Reflection (5-minutes):

Table Discussion (15-minutes):

Large-Group Open Discussion (10-minutes):

Notes: