ALGEBRA BERNAYS UNIVERSITY

WORKSHOP SERIES | GEANT RESEARCH COMMUNITY

Cloud-Enabled Research Workshops

Microsoft Fabric & Azure AI for Academic Research Workflows

MODULE 03
Developing AI-Enhanced Research Workflows
Wednesday 14 May 2026

MODULE 04
Practical Application of Azure AI Services
Thursday 29 May 2026

MODULE 05
Experimental Design in Machine Learning
Monday 29 June 2026

Prepared by

Mario Pilija

Microsoft Certified Trainer (MCT)

Algebra Bernays University

Delivered for GEANT - pan-European Research & Education Network

Leading Learning Partners Association

ALGEBRA BERNAYS UNIVERSITY

Background & Purpose

Research institutions across Europe face a critical shift in how data is collected, managed, analysed, and shared. Funders including NWO and the European Commission expect open, FAIR-compliant data practices. The Netherlands Reproducibility Network reflects a growing recognition that reproducibility and data governance are foundational to credible science. Yet the practical infrastructure remains fragmented for most research groups: data scattered across local drives, pipelines that exist only in someone's head, collaboration via emailed spreadsheets, and analysis that cannot be re-run six months later when a journal reviewer asks.

These workshops address the infrastructure layer of that problem - not by claiming technology alone solves the challenge, but by showing how cloud platforms can make good research data practices significantly easier to implement and sustain. Attendees leave with a concrete understanding of how Microsoft Fabric and Azure AI support the full research data lifecycle: from ingestion and governance through to AI-assisted analysis, reproducible pipelines, and systematic AI-enhanced research workflows, illustrated through live demonstrations using real research datasets.

WHO SHOULD ATTEND

- **Academic researchers** at any career stage working with substantial datasets or multi-institution collaborations
- **Research data managers,** data stewards, and research support staff
- **Research IT professionals,** infrastructure teams supporting research groups
- **Graduate students & postdocs** involved in data-intensive research

Technologies Covered

Microsoft Fabric | OneLake & Lakehouses | Dataflow Gen2 | Data Factory Pipelines | Apache Spark

MLflow Tracking | Fabric Data Agent | Microsoft Purview | Content Understanding

Azure AI Foundry | Model Catalog | RAG & Grounding | Prompt Engineering | Code Interpreter

Azure AI Language | Azure AI Search | Knowledge Mining | Azure AI Vision

Azure Machine Learning | AutoML | Confidential Computing | Federated Learning

Foundry Agent Service | Connected Agents | MCP Protocol | Agent Framework

Evaluation Metrics | Content Filters | Responsible AI | EU AI Act

MODULE 03

Developing AI-Enhanced Research Workflows with Azure AI Foundry

Wednesday 14 May 2026 | 120 min | Online | 12:00-14:00 CET

50-100 participants | English-language delivery

Leading Learning Partners Association

ALGEBRA BERNAYS UNIVERSITY

Date & time	Wednesday 14 May 2026, 12:00-14:00 CET
Duration	120 minutes
Format	Online, instructor-led with live demonstrations
Audience	50-100 participants

Session Description

This module introduces Azure AI Foundry as a platform for structuring AI-driven research initiatives. Participants examine how experiments are defined, monitored, and compared through demonstrations, while ensuring methodological consistency and traceability throughout several study iterations. The programme emphasises the shift from informal model utilisation to systematic, verifiable research processes.

A particular focus addresses EU data residency: which AI models are hosted within European data zones, how to verify data processing guarantees, and practical guidance for researchers operating under GDPR and institutional data governance requirements. Azure AI Foundry is presented as a research workbench where models are selected like instruments, prompts designed like protocols, and AI agents built to autonomously analyse datasets with full traceability.

The session culminates in a live demonstration of an AI agent using Code Interpreter to perform statistical analysis on a research dataset, showing how natural language instructions can drive reproducible, auditable computational research workflows.

Learning Outcomes

- Navigate Azure AI Foundry to select, benchmark, and deploy models appropriate for a specific research task
- Design system messages and prompt strategies that ensure domain-specific, reproducible AI outputs
- Explain the RAG pattern and describe when to use grounding, fine-tuning, or prompt engineering for research applications
- Describe how AI agents with Code Interpreter can autonomously analyse datasets and generate reproducible research outputs
- Apply evaluation metrics to validate AI-generated outputs with the same rigour as any research method
- Identify responsible AI practices that align with research ethics and EU regulatory requirements

Time	Topic	Content
0:00	**Recap and Framing**	Quick recap of Sessions 1-2. Today: the AI layer. Problem: researchers using ChatGPT with no versioning, evaluation, or audit trail.
0:05	**Azure AI Foundry as Research Infrastructure**	Microsoft Foundry platform. Hubs and projects as research infrastructure. Model catalog: 11,000+ models. Benchmarks: quality, accuracy, cost, latency. Deployment options. EU data residency. Foundry SDK.
0:22	**Systematic Prompt Engineering for Research**	Prompt anatomy: system message, user prompt, completion. Prompt engineering as experimental design: controlling temperature, top-p, max tokens. Optimization 2x2: RAG vs fine-tuning.
0:35	**DEMO 6: Foundry Portal - Model Selection & Playground**	Navigate ai.azure.com. Model catalog, benchmarks across 3 models. Chat playground with research system message. Temperature effects. Compare two models side by side.
0:47	**Grounding AI in Research Data**	Ungrounded AI: hallucinations, fabricated citations. RAG pattern: retrieve, augment, generate. Vector indexes. Azure AI Search. Fine-tuning for domain adaptation. When to use each.
1:02	**Evaluating AI Outputs & Responsible AI**	Manual and automated evaluation: coherence, fluency, groundedness. Content filters and prompt shields. Responsible AI 4-layer model. EU AI Act alignment.
1:15	**AI Agents for Research Workflow Automation**	AI agents: model + knowledge + tools. Foundry Agent Service. Knowledge tools (File Search, AI Search, Fabric). Action tools (Code Interpreter, Functions). Connected agents. MCP integration.
1:27	**DEMO 7: Research Data Agent & Code Interpreter**	Create an agent that analyses a research dataset. Natural language instructions drive statistical analysis. Agent generates charts and insights autonomously. Every step logged, reproducible, shareable.
1:42	**End-to-End Architecture & Next Steps**	Visual recap of research AI architecture. Microsoft Learn paths. Azure for Research credits. Fabric trial options. GitHub lab exercises: mslearn-ai-studio, mslearn-ai-agents.
1:50	**Q&A**	Open discussion.

Practical Application of Azure AI Services in Research Settings

Thursday 29 May 2026 | 120 min | Online | 12:00-14:00 CET

50-100 participants | English-language delivery

Leading Learning Partners Association

ALGEBRA BERNAYS UNIVERSITY

Date & time	Thursday 29 May 2026, 12:00-14:00 CET
Duration	120 minutes
Format	Online, instructor-led with live demonstrations
Audience	50-100 participants

Session Description

This module explores integrating prebuilt Azure AI Services into research workflows for tasks such as text analysis and image comprehension. Illustrative use cases emphasise the appropriateness and constraints of these services, examining instances where they enhance research and situations that require bespoke methodologies or more rigorous validation to uphold scientific integrity.

Participants work through practical scenarios: extracting entities and sentiment from research text, building searchable knowledge bases from document collections, processing instrument logs and field data sheets at scale, and detecting PII in human-subjects data. Each capability is framed within the context of research methodology - when pre-built services are sufficient and when custom approaches are warranted.

This module serves as an application-focused extension to the workshop series, connecting core cloud literacy with everyday research practices while highlighting scientific rigour, reproducibility, and ethical use of AI.

Learning Outcomes

- Identify which Azure AI Services apply to common research data processing tasks
- Apply text analytics (sentiment, entities, key phrases, PII detection) to research text data
- Design a knowledge mining pipeline that makes research document collections searchable
- Use Content Understanding and Document Intelligence to extract structured data from unstructured research documents
- Describe how multimodal AI capabilities can process diverse research data types
- Integrate AI Services outputs with the Fabric data platform for downstream analysis

Time	Topic	Content
0:00	**Recap and Framing**	Recap M1-M3. Today: applying Azure AI Services directly to research tasks. From platform concepts to practical tooling.
0:05	**Azure AI Services Landscape**	AI Services portfolio overview. Language, Vision, Speech, Content Understanding, AI Search. Single-endpoint deployment. Pre-built vs custom models. Research use cases.
0:20	**Text Analytics for Research Data**	Sentiment analysis on surveys. Key phrase extraction from abstracts. Named entity recognition. PII detection for human-subjects data. Custom text classification.
0:35	**DEMO 8: Text Analysis on Research Abstracts**	Analyse research abstracts using Azure Language. Extract key phrases, entities, sentiment. Detect PII in participant responses.
0:47	**Knowledge Mining for Research Archives**	Azure AI Search indexing. Built-in AI skills. Custom skills. Knowledge stores. Semantic ranking. Vector search for semantic similarity.
1:00	**DEMO 9: Knowledge Mining Pipeline**	Build AI Search index over research documents. Configure skills. Query with filters and facets. Semantic ranker improving relevance.
1:12	**Document Intelligence & Content Understanding**	Reading text from scanned documents. Custom analysers for research forms. Batch processing instrument logs and field data sheets.
1:25	**DEMO 10: Document Processing Pipeline**	Submit research documents to Content Understanding. Extract structured fields. Integration path: data flows into Fabric lakehouse.
1:35	**Multimodal AI for Research**	Image analysis for field research. Combining vision + language. Biodiversity surveys, archaeological documentation, satellite imagery.
1:48	**Architecture & Next Steps**	AI Services architecture. Integration with Fabric and Foundry. Certification: AI-3003, AI-3004, AI-3022. Learn paths.
1:55	**Q&A**	Open discussion.

MODULE 05

Experimental Design in Machine Learning for Research

Monday 29 June 2026 | 120 min | Online | 12:00-14:00 CET

50-100 participants | English-language delivery

Leading Learning Partners Association

ALGEBRA BERNAYS UNIVERSITY

Date & time	Monday 29 June 2026, 12:00-14:00 CET
Duration	120 minutes
Format	Online, instructor-led with live demonstrations
Audience	50-100 participants

Session Description

This module emphasises machine learning as a research methodology rather than as an implementation assignment. Participants examine the translation of research objectives into machine learning problem formulations, the selection and evaluation of models, and the interpretation and reporting of results in a scientifically rigorous manner. Scenarios address common research issues, including data leakage, overfitting, and misinterpretation of performance metrics.

A dedicated focus addresses distributed patterns for data confidentiality: how Azure ML can train models on data that cannot leave the customer's on-premises environment. Azure Confidential Computing, Azure Arc-enabled ML, and federated learning patterns are presented as practical solutions for institutions bound by GDPR, health data directives, or institutional data governance policies.

This module serves as an application-focused extension connecting core cloud literacy with everyday research practices, strengthening programme objectives around scientific rigour, reproducibility, and ethical use of AI while establishing Azure as an integrated research ecosystem.

Learning Outcomes

- Set up an Azure Machine Learning workspace and configure compute for research experiments
- Design and run ML experiments with full tracking using MLflow for parameters, metrics, and artefacts
- Apply hyperparameter tuning and AutoML to systematically optimise models for research tasks
- Identify which distributed/confidential computing pattern applies to specific data residency and privacy requirements
- Deploy trained models to real-time and batch endpoints with monitoring for production research use
- Integrate Azure ML with Fabric and Foundry for end-to-end research AI workflows

Time	Topic	Content
0:00	**Recap and Framing**	Recap M1-M4. Today: rigorous ML experimental design. Problem: ML on laptops with no tracking, reproducibility, or scalability.
0:05	**Azure ML as a Research Lab**	Azure ML workspace. Compute targets. Data assets. Environments for reproducibility. Comparison with Fabric notebooks - when to use which.
0:22	**ML Experiment Lifecycle**	MLflow tracking: parameters, metrics, artefacts. Command jobs. Hyperparameter tuning with sweep jobs. ML pipelines. AutoML for non-ML researchers.
0:37	**DEMO 11: ML Experiment with MLflow Tracking**	Azure ML workspace. Training script on EWQ dataset. Track with MLflow. Compare runs. Register best model. Every experiment versioned.
0:50	**Distributed & Confidential ML Patterns**	Data that cannot leave the institution. Azure Confidential Computing: TEE-based training. Arc-enabled ML: on-premises Kubernetes. Federated learning. GDPR and health data compliance.
1:05	**DEMO 12: Responsible AI Dashboard**	Create Responsible AI dashboard. Inspect fairness, error analysis, interpretability. Validate before publishing ML results.
1:15	**Model Deployment**	Real-time and batch endpoints. Blue-green deployment. Model monitoring and drift detection. MLOps with GitHub Actions automation.
1:30	**From Experiment to Production**	ML pipelines as reproducible workflows. Automated retraining. Integration: Azure ML + Fabric + Foundry. Certification: DP-100 retiring, replaced by AI-300.
1:42	**Architecture & Next Steps**	Complete ML architecture for research. Azure ML + Fabric + Foundry integration. Learn paths. Azure for Research credits.
1:55	**Q&A**	Open discussion.

Themes Across All Three Modules

Reproducibility	Every demo illustrates reproducibility in practice: auditable queries, logged experiments, version-controlled pipelines, lineage from raw data to published result.
FAIR Data	The medallion architecture, OneLake catalog, and Purview integration are mapped to Findable, Accessible, Interoperable, and Reusable data principles throughout.
Governance	Governance is framed for researchers: data provenance, sensitivity classification, audit logs for ethics boards, and funder-required data management plan support.
Collaboration	Cross-institutional data sharing, access control without data duplication, and shared workspaces for multi-partner projects are covered as practical scenarios.
Accessibility	Tools are demonstrated at multiple levels: visual no-code interfaces, notebook-based workflows for computational researchers, and SQL/KQL for data professionals.
Honest Framing	All three modules acknowledge that obstacles to research data management are primarily cultural, not technical. The workshops address the infrastructure layer without overclaiming.

DELIVERY NOTES

All three workshops are delivered online in English by a Microsoft Certified Trainer (MCT). Demonstrations are performed live on pre-provisioned Azure infrastructure. All demonstrations use publicly available or purpose-built research-domain datasets. Participants do not need an Azure subscription to attend; hands-on lab access can be arranged through Azure for Research or Fabric trial programmes referenced at session close.

EXCEL AT WHAT YOU LOVE DOING.

LIGHT THE SPARK.

Prepared by

Mario Pilija

Microsoft Certified Trainer (MCT)

Algebra Bernays University

Document prepared for

GEANT - pan-European Research & Education Network

Cloud-Enabled Research Workshop Series | 2026

Leading Learning Partners Association

ALGEBRA BERNAYS UNIVERSITY