Curriculum as Infrastructure: Integrating Open-Source AI Ecosystems and Experiential Learning into Computer Science and Engineering Pedagogy

Dr. Zahmeeth Sayed Sakkaff, West Virginia University Tech and West Virginia University

Dr. Zahmeeth Sayed Sakkaff is an Assistant Professor at WVU Institute of Technology and West Virginia University, whose nearly two decades of work in AI, machine learning, and data science spans both rigorous computational research and university-level teaching across AI, ML, databases, and big data. At Argonne National Laboratory's Data Science and Learning Division and the University of Chicago's CASE, she designed AI-powered experimental pipelines, built large-scale biological data models, and automated high-throughput computational workflows — experiences that directly shape her approach to bridging production-scale research practices and engineering pedagogy. Recognized with multiple Best Paper and Best Poster awards, she brings a sustained focus on open-source ecosystems, reproducibility, and systems-level thinking into her courses and mentorship. Her current work integrates national cyberinfrastructure into CS and engineering curricula, grounding student projects in the same deployment, validation, and reproducibility demands they will face in research and industry.

Dr. Venkat Sai Suman Lamba Karanam, Bowling Green State University

Dr. Karanam is an Assistant Professor in the Department of Computer Science at the Bowling Green State University. Dr. Karanam's research and teaching spans large scale data analysis, networking and distributed computing, computational biology and data science broadly.

Dr. Kenan Hatipoglu, West Virginia University Institute of Technology

Kenan Hatipoglu is the chair of Leonard C. Nelson School of Engineering and professor of Electrical Engineering at West Virginia University Institute of Technology. He completed his Master of Science degree in Electrical Engineering at University of Louisville, Kentucky in 2008 and joined Tennessee Tech University in 2009 to pursue his Ph.D. in Electrical (Power) Engineering. He completed his graduate study in August 2013. He is an active member of IEEE and ASEE.

Dr. Yan Wu, Bowling Green State University

Dr. Yan Wu is an Associate Professor of Computer Science at Bowling Green State University and serves as the department's Internship Coordinator. Her research focuses on software assurance and software engineering, specifically the development of secure and reliable systems. As a dedicated educator, she teaches a variety of cybersecurity courses, where she helps bridge the gap between theoretical engineering and practical defense.

Work in Progress (WIP) | Paper ID #50061 | Board # 107
West Virginia University Institute of Technology

Curriculum as Infrastructure: Integrating Open-Source AI Ecosystems and Experiential Learning into Computer Science and Engineering Pedagogy

Dr. Zahmeeth Sayed Sakkaff[1,2] · Dr. Venkat Sai Suman Lamba Karanam[3]. Dr. Kenan Hatipoglu[1] · Dr. Yan Wu[3]
[1]WVU Institute of Technology · [2]West Virginia University · [3]Bowling Green State University

ASEE 2026 Annual Conference | June 21-24, 2026, | Charlotte, NC

Work in Progress (WIP) | Paper ID #50061 | Board # 107
Bowling Green State University

BACKGROUND & MOTIVATION

The Problem

Engineering and computer science graduates learn computational methods in coursework but often lack preparation for real production systems—especially reproducibility, deployment, and working with diverse hardware[1].

Students learn tools — but not how to deploy them in shared, production-like environments.

Our Approach: Curriculum as Infrastructure

We use national-scale cyberinfrastructure (NSF/DoE) so students can deploy and validate their work under real-world constraints such as shared resources, diverse hardware, and reproducibility requirements

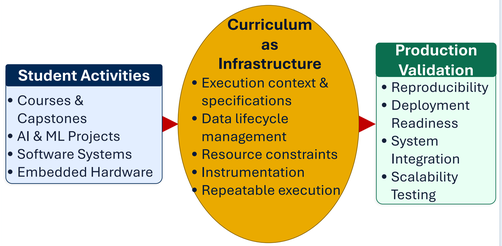

Figure 1: Curriculum as Infrastructure Model: *Student work is embedded in a curriculum that uses production-scale cyberinfrastructure to turn prototypes into reproducible, deployable, and validated artifacts while supporting Bloom's Taxonomy and ABET-aligned outcomes.*

Pilot Scope

- Multiple course offerings + senior capstone programs
- Two state universities: WVU Tech & BGSU
- CS, Information Systems & Electrical Engineering
- Platforms: NRP, Chameleon Cloud, COSMOS, FABRIC, POWDER

Course design incorporates insights from our ATPESC collaboration with Argonne National Laboratory, enriching the curriculum with exascale computing and data-science methods.

RESEARCH QUESTIONS

RQ1 — Production Readiness & Reproducibility

How does integrating national cyberinfrastructure improve students' production-readiness—measured by rerun success, artifact completeness, and stress-test robustness?

Metrics: Rerun success rate | Artifact completeness | Stress-test pass rate

RQ2 — Bloom's[2] Taxonomy & ABET Outcomes

How do infrastructure-based milestones shape student performance across Bloom's Taxonomy and ABET outcomes using rubric and milestone evaluations?

Metrics: Rubric scores per Bloom level | ABET outcome mapping

RQ3 — AI-Assisted Learning in LMS

How does LMS-based AI support shape engagement, reflection, and readiness for infrastructure-centered projects without displacing higher-order cognition?

Metrics: LMS interaction logs | Pre/post self-efficacy | Reflective depth

Table 1: Planned Deliverables, Data Sources & Expected Outcomes (RQ1–RQ3)

Research Focus	Data / Deliverables Collected	Evaluation Methodology	Expected Results / Evidence
RQ1: Production-readiness & reproducibility	Validation bundles; automation scripts; logs; stress-test reports; infra traces	Rerun success; artifact completeness; milestone comparison	Higher reproducibility; redeployment success; stronger systems thinking
RQ2: Learning outcomes (Bloom + ABET)	Rubric scores mapped to Bloom + ABET; peer/instructor evaluations	Rubric analysis; longitudinal comparison; outcome aggregation	Improved higher-level Bloom performance; stronger ABET skills
RQ3: Engagement & reflective practice	Reflections; LMS logs; surveys; AI-assisted study interactions	Reflections; LMS logs; surveys; AI-assisted study interactions	Deeper reflection; better study habits; improved project readiness
Cross-cutting: feasibility & scalability	Instructor effort logs; infra usage; support requests; deployment overhead	Descriptive + feasibility analysis; scalability assessment	Evidence that national cyberinfrastructure is sustainable in curriculum

STUDY DESIGN

- Multi-university pilot (WVU Tech, Bowling Green State);
- CS/IS and EE programs across semesters and capstones.
- Data: validation bundles, rubric scores, reflections, LMS logs.
- Quantitative: reproducibility metrics, rubric trends.
- Qualitative: thematic coding of reflections and LLM prompt logs

National-scale cyberinfrastructure is a core curriculum element, enabling student planning, execution, and validation on shared, production-like heterogeneous testbeds.

METHODS

National-Scale Cyberinfrastructure

National-scale cyberinfrastructure is a core curriculum element, enabling student planning, execution, and validation on shared, production-like heterogeneous testbeds[3].

Research-Grade Platforms Used: NRP[4] · Chameleon · COSMOS[5] · POWDER[6] · FABRIC · FABRIC[7]

Students gain hands-on experience with Jupyter[8], Coder, AI & LLMs, Raspberry Pi[9], and FPGA[10]—enabling projects that expose real-world constraints.

4-Stage Integration Framework

- **i Provisioning & Access** — Accounts, compute/storage allocations, least-privilege access with predefined quotas and approval workflows
- **ii Execution Environments** — Versioned containers/software stacks eliminate it works on my machine' failures across heterogeneous hardware
- **iii Workflow Orchestration** — Scripted pipelines from data ingestion to training/inference; machine-readable config files, logs, and outputs
- **iv Artifact Management** — Standardized metadata (inputs, params, versions, timestamps) + version control integrated with the infrastructure

Figure. 2: Bloom's Taxonomy aligned with Curriculum as Infrastructure

Blooms Taxonomy ↔ **Curriculum as Infrastructure**

5. **Create** — *End-to-end systems* · *Production validation*
4. **Evaluate** — *Reproducibility checks* · *Trade-off analysis*
3. **Analyze** — *Performance profiling* · *Debugging & logging*
2. **Apply** — *Scripted workflows* · *Small-scale experiments*
1. **Understand** — *Tool fundamentals* · *System concepts*

Validation Bundle Requirements

(a) Re-execution guide: how to rerun and verify outputs
(b) Automation scripts for evaluation and execution
(c) Structured logging and metrics capture
(d) Stress-tests: resource limits, invalid inputs, partial avail

NSF REPETO-style reproducibility practices adopted for packaging code, docs, and rerun scripts[11,12].

PILOT PROJECTS & FINDINGS

Senior Capstone Projects — WVU Tech

1. **Hardware Monitoring Platform** — A Windows Presentation Foundation (WPF)-based administrative platform for operational monitoring and organizational efficiency visualization
2. **Financial Advisory App** — Blazor/C# full-stack app integrating Azure, GitHub, SSMS, MS Entra & MudBlazor with hardware acceleration
3. **AI Malware Detection System** — PyTorch/TensorFlow steganographic malware detection with web frontend and dedicated network testbed
4. **Student Government Portal** — Centralized site with club subpages, unified calendars, and subscription systems for future maintenance

Robotics & Gaming in Education — Investigating robotic/gaming technologies as effective tools for teaching and learning (edge-cloud integration)

Preliminary Evidence: LLM Prompt Logs (Fall 2025)

Anonymized prompt logs showed students extending beyond models to infrastructure, execution, and reproducibility.

Key finding: Prompts shifted from model-centric to infrastructure- and validation-aware reasoning, consistent with *Curriculum as Infrastructure* objectives; students explicitly addressed data handling, execution environments, and reproducibility.

Table 2: LLM Prompt Logs—Shift Toward Infrastructure-Aware Reasoning (Fall 2025)

Prompt Category	Representative Student Prompt	Milestone
ML Debugging & Preprocessing	"How should I handle class imbalance when accuracy is high but F1 is low?"	Model training
Evaluation & Metrics	"Which metric is most appropriate for this dataset, and why?"	Validation
Infrastructure / Execution	"How can I run this experiment on OSC without dependency conflicts?"	Deployment
Experiment Management	"How should I structure experiments for reproducibility across runs?"	Pipeline design
Reporting & Writing	"How should I present results and limitations in the final report?"	Final report

AI Learning Tool in LMS (Pilot)

StudyStash AI supports concept review and study reflection in Canvas/Blackboard Ultra—not graded work—preserving academic integrity and reserving higher-order Bloom's tasks for students.

CONCLUSIONS

- Infrastructure-first framing builds systems-level thinking.
- Students identify production constraints before independent projects.
- LLM prompt logs show a shift toward infrastructure-aware reasoning.
- CI-based milestones span all Bloom's Taxonomy levels.
- ABET-aligned rubrics are embedded in infrastructure-centered milestones.
- NSF REPETO practices improve artifact reproducibility and reusability.[12]

FUTURE WORK

- Quantitative reproducibility study across student cohorts
- Longitudinal rubric analysis aligned with Bloom's Taxonomy
- LMS AI tool effectiveness on engagement, self-efficacy, and reflection depth
- Responsible AI integration strategies for LMS platforms
- Scale the pilot to additional universities and disciplines

ACKNOWLEDGMENTS

Supported by the 2025–26 WVU Tech Faculty & Staff Mini Grant Program (Cycle II), this work draws on ATPESC insights and collaborations with Argonne National Laboratory to enrich the course with exascale computing and data science methodologies. The pilot also integrates NSF- and DoE-supported open cyberinfrastructure platforms.

SELECTED REFERENCES

[1] Sculley et al., Hidden Technical Debt in ML Systems, NeurIPS, 2015.
[2] Krathwohl, A Revision of Bloom's Taxonomy, Theory Into Practice, 2002.
[3] Lanubile et al., Teaching MLOps via Project-Based Learning, ICSE-SEET, 2023.
[4] Weitzel et al., The National Research Platform, PEARC '25, ACM, 2025.
[5] Raychaudhuri et al., COSMOS: A City-Scale Programmable Wireless Testbed, MobiCom, ACM, 2020.
[6] Breen et al., POWDER: Platform for Open Wireless Data-Driven Research, WiNTECH, ACM, 2020.
[7] Baldin et al., FABRIC: National-Scale Programmable Experimental Network, IEEE Internet Computing, 2019.
[8] Johnson, Benefits and Pitfalls of Jupyter Notebooks in Teaching, SIGITE, ACM, 2020.
[9] Upton & Halfacree, Raspberry Pi User Guide, Wiley, 2016.
[10] Boutros & Betz, FPGA Architecture: Principles and Progression, IEEE Circuits & Systems Magazine, 2021.
[11] Saeed & Fund, Scaffolding Reproducibility for ML Classroom, ACM REP, 2024.
[12] REPETO: Reproducibility in Experimental Evaluation, NSF Report, 2024.