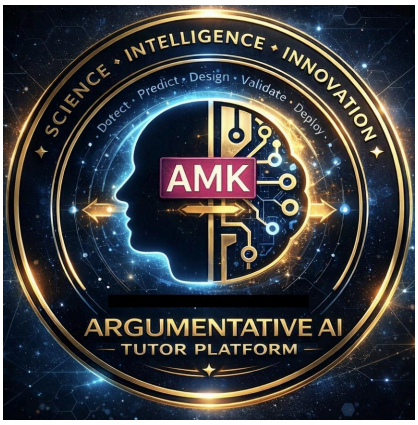

AI Lecture Notes

Sample Capstone Project Aligned to the AI Certification Program and the AMK Research Lab Program

Anchor sample: Fiduciary Monte Carlo–Driven AI Advisory Architecture for Real-Time Healthcare and Financial Decision Support

Prepared as a teaching resource with diagrams, charts, governance controls, and research design guidance.

Who this is for	What students learn	Expected outputs
AI certification learners, final-project students, lab researchers	Problem framing, lifecycle design, governance, evaluation, ethics, and capstone packaging	Proposal, architecture diagram, prototype plan, evaluation matrix, compliance map, and presentation slides

Learning outcomes

- Explain how an applied AI capstone links AI foundations, data, models, security, ethics, and deployment controls.
- Translate a research idea into a lifecycle with datasets, methods, evaluation metrics, and governance checkpoints.
- Use a sample capstone to discuss real-world risks such as bias, prompt injection, data leakage, instability, and overreliance.
- Map a project to practical governance frameworks including NIST AI RMF, the NIST Generative AI Profile, ISO/IEC 42001, UNESCO, OECD principles, HIPAA, and GDPR.

1. Why capstones matter in an AI certification program

A capstone turns theory into a defensible artifact. Instead of stopping at model accuracy, students must define the problem, identify stakeholders, justify datasets, specify controls, and show how the system will be measured and governed.

For the AMK Research Lab model, the capstone should also demonstrate research maturity: a clear gap, a tractable architecture, auditable experimentation, and a pathway from prototype to deployment.

This sample uses a fiduciary AI advisory concept because it forces students to combine machine learning, probabilistic reasoning, cybersecurity, sentiment analysis, governance, and sector-specific regulation in one coherent project.

2. Anchor sample capstone concept

Proposed sample title:

A Fiduciary Monte Carlo–Driven AI Advisory Architecture for Real-Time Healthcare and Financial Decision Support with Sentiment Governance and Cybersecurity Stability Guarantees

Teaching value of this sample

- It is broad enough to demonstrate modern AI practice, but structured enough to teach system boundaries.
- It combines predictive AI, conversational AI, uncertainty quantification, monitoring, and controls.
- It naturally supports interdisciplinary assessment across technical, ethical, security, and compliance domains.

Research-gap illustration

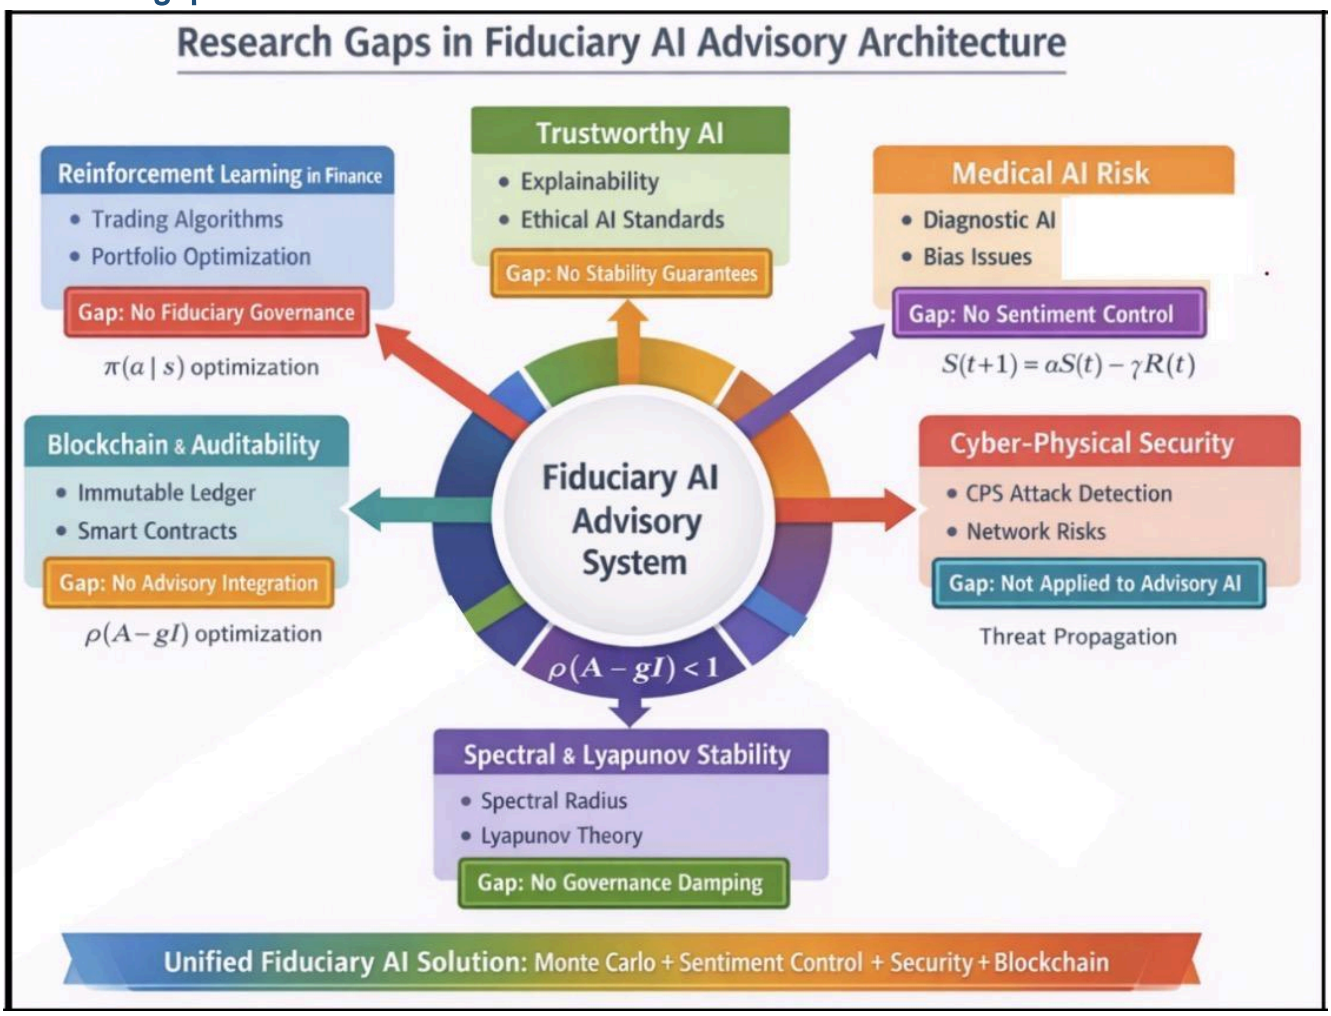

Figure 1. User-provided research-gap illustration for the fiduciary AI advisory concept.

3. Alignment with AI Certification competencies

The chart below shows how a single capstone can cover multiple certification outcomes at once.

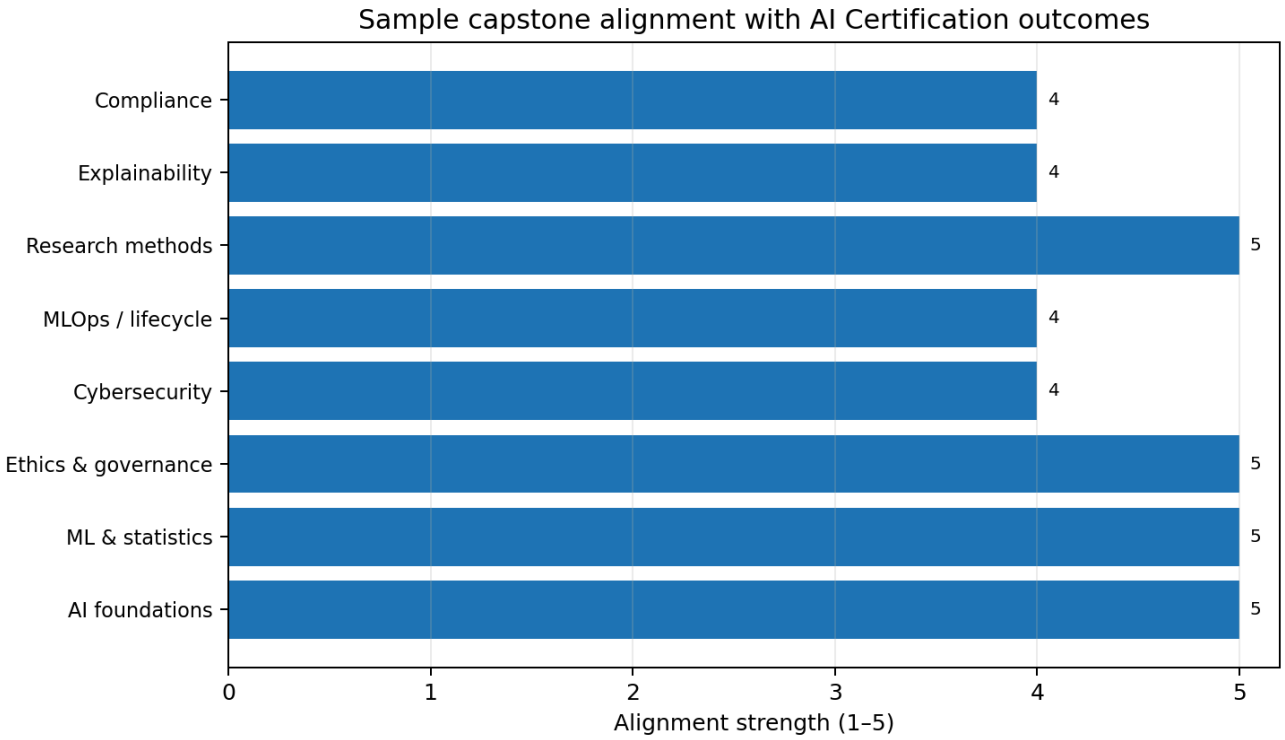

Figure 2. Illustrative alignment of the sample capstone with core AI certification outcomes.

Interpretation

High alignment scores in AI foundations, ML and statistics, research methods, and ethics indicate that the capstone can anchor an entire applied-learning module.

Slightly lower but still strong scores in compliance, explainability, cybersecurity, and lifecycle management reflect the need for supporting lectures, labs, and policy exercises.

4. Practical capstone lifecycle

Students should treat the project as a governed lifecycle rather than only a modeling exercise.

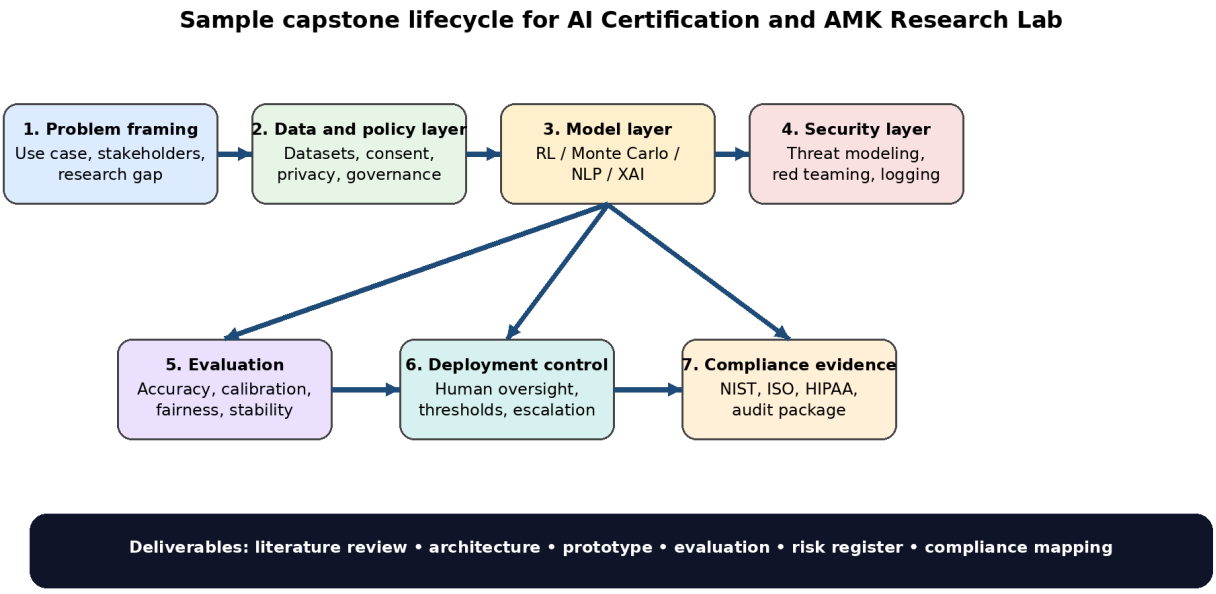

Figure 3. Sample capstone lifecycle showing how problem framing, modeling, security, evaluation, deployment, and compliance evidence fit together.

Phase-by-phase teaching notes

Phase	Main question	Typical artifacts	Certification focus
1	What is the real problem?	Problem statement, stakeholders, research gap	AI foundations
2	What data and rules govern the system?	Dataset inventory, consent logic, privacy map	Data governance
3	Which models and equations are needed?	Monte Carlo engine, RL policy, sentiment model, XAI plan	ML methods
4	How can the system fail?	Threat model, red-team scenarios, logging design	AI security
5	How will success be measured?	Metrics, baselines, error analysis, calibration	Evaluation
6	Who remains accountable?	Human override rules, approval thresholds	Ethics and operations
7	What proves compliance?	Audit trail, standards mapping, evidence pack	Compliance

5. Technical architecture summary

- Layer 1 – User interaction: advisory chatbot, dashboard, authentication, disclosures, and escalation routes.
- Layer 2 – Advisory core: Monte Carlo risk estimation, policy logic, recommendation engine, and confidence reporting.
- Layer 3 – Sentiment and behavioral intelligence: emotional-state analysis, tone control, and trust-signal management.

- Layer 4 – Cybersecurity and threat modeling: prompt injection defense, anomaly detection, model abuse monitoring, and propagation analysis.
- Layer 5 – Governance and dashboard control: explainability, decision logs, thresholds, and policy enforcement.
- Layer 6 – Secure data infrastructure: access control, encryption, backup, retention, and audit evidence.

Representative mathematical ideas

Monte Carlo advisory estimate: $E[D] = (1/N)\,\Sigma\, f(X_i, \theta_i)$

Sentiment regulation: $S(t+1) = \alpha S(t) + \beta E(t) - \gamma R(t)$

Threat propagation: $x(t+1) = (A - gI)x(t)$

Stability target: $\rho(A - gI) < 1$

6. Evaluation design for classroom or lab implementation

A strong capstone must define baselines, comparison groups, and measurable outcomes.

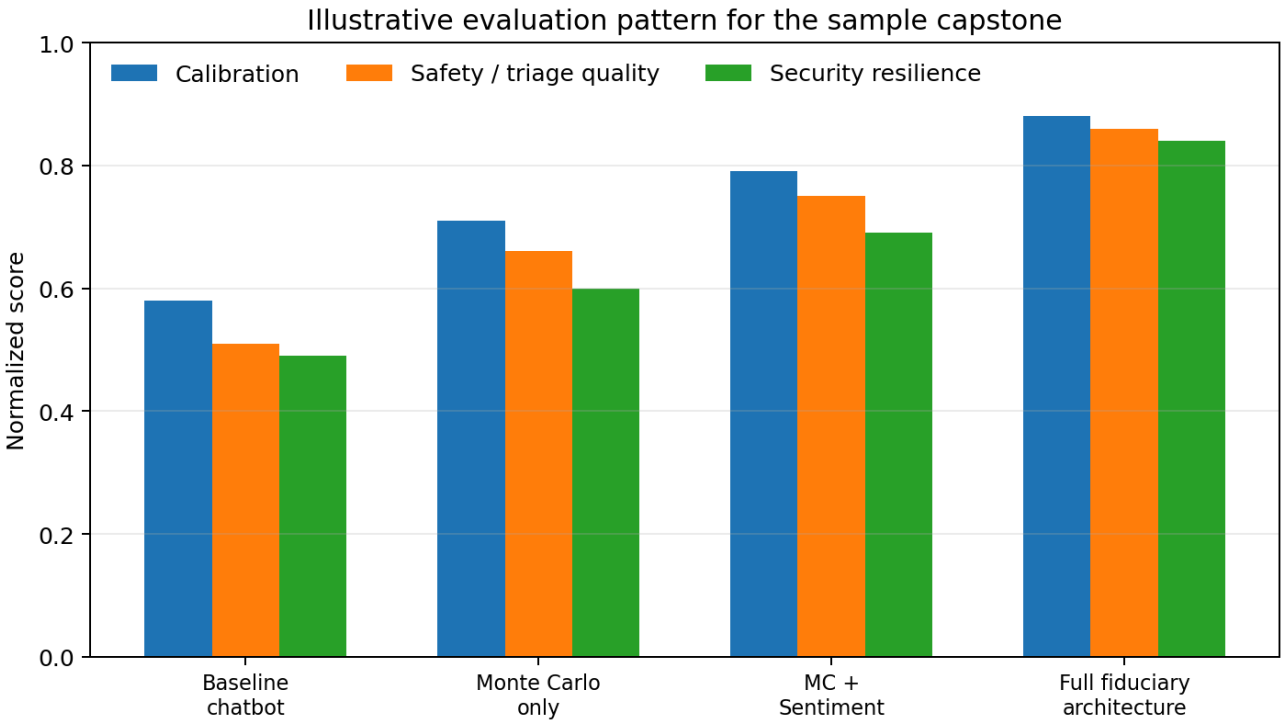

Figure 4. Example evaluation profile across four model configurations.

Recommended evaluation matrix

Dimension	Example metrics	Why it matters	Control response
Model quality	Precision, recall, F1, calibration	Checks whether recommendations are reliable	Retrain, recalibrate, rebalance data
Fiduciary safety	False-negative risk, harm bound, regret interval	Measures whether advice can harm users	Stricter thresholds and human review
Sentiment stability	Volatility index, convergence time	Prevents anxiety or panic amplification	Tone control, disclosure, escalation

Security resilience	Attack detection rate, MTTD, MTTM	Assesses prompt-injection and poisoning resilience	Hardening, sandboxing, incident response
Governance evidence	Audit completeness, explainability coverage	Shows accountability and traceability	Immutable logs, sign-off workflow

7. Risk, ethics, and control mapping

AI capstones become stronger when students explicitly show where the system can fail and which controls reduce each failure mode.

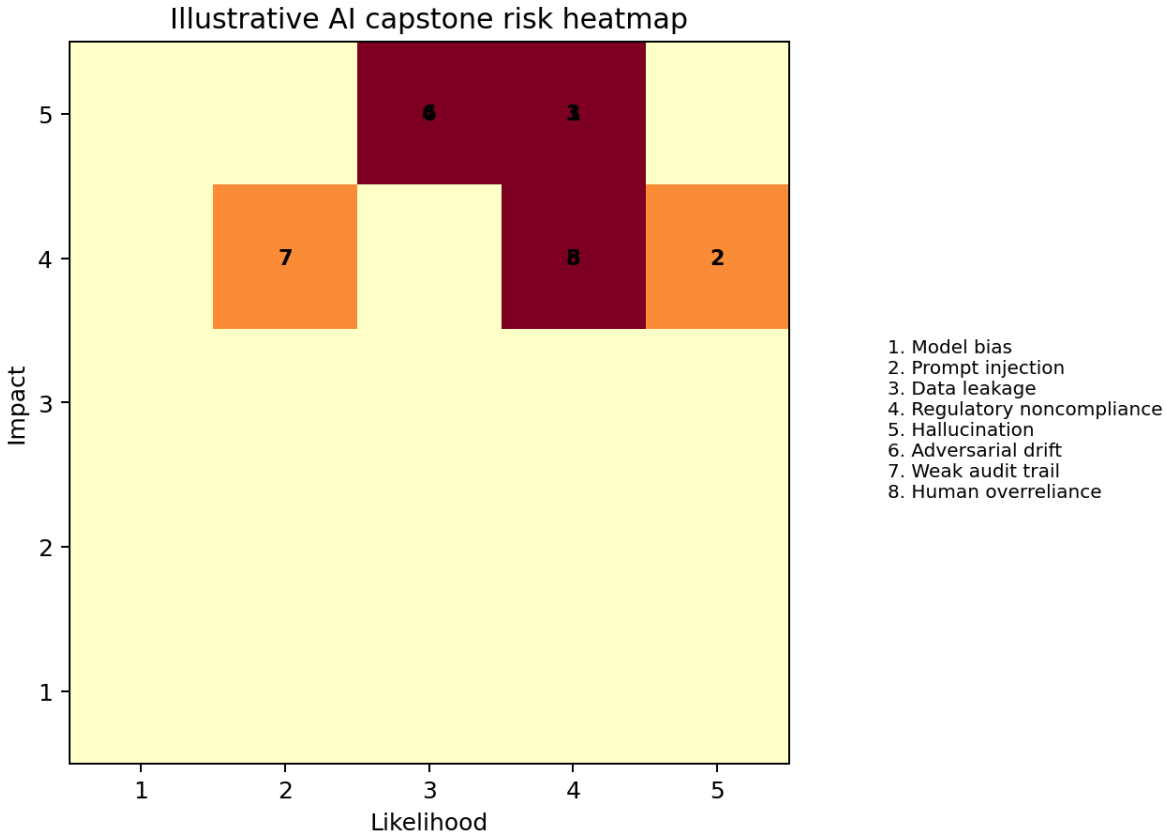

Figure 5. Illustrative heatmap of common capstone risks.

Risk-to-control examples

Risk	Manifestation	Primary control	Reference anchor
Bias	Unequal financial or health recommendations	Bias testing, subgroup evaluation, override policy	OECD; UNESCO
Prompt injection	Manipulated instructions or disclosure bypass	Input filtering, isolation, approval gates	OWASP; NIST GenAI
Data leakage	Exposure of medical or financial data	Access control, minimization, encryption	HIPAA; GDPR
Hallucination	Unsupported advice or fabricated rationale	Grounding, confidence disclosure, human review	NIST AI RMF

Model drift	Performance degradation over time	Monitoring, drift alarms, periodic validation	NIST AI RMF
Weak accountability	No evidence of who approved what	Immutable logs, audit trail, RBAC	ISO/IEC 42001

8. Compliance and governance anchors

NIST AI RMF 1.0 frames AI governance using Govern, Map, Measure, and Manage functions; this is especially useful for structuring capstone sections and assessment rubrics.

The NIST Generative AI Profile extends AI RMF thinking to generative and conversational systems, with attention to misuse, content integrity, security, privacy, and human-AI interaction.

OECD AI Principles and the UNESCO Recommendation reinforce human rights, transparency, fairness, accountability, and human oversight.

ISO/IEC 42001 provides an organizational management-system view of AI governance, useful for teaching policy, controls, and audit responsibilities.

For healthcare-flavored capstones, HIPAA highlights the need to protect the confidentiality, integrity, and availability of electronic protected health information.

For projects touching European personal data, GDPR principles remain central for lawful basis, minimization, transparency, and accountability.

9. Capstone roadmap and classroom pacing

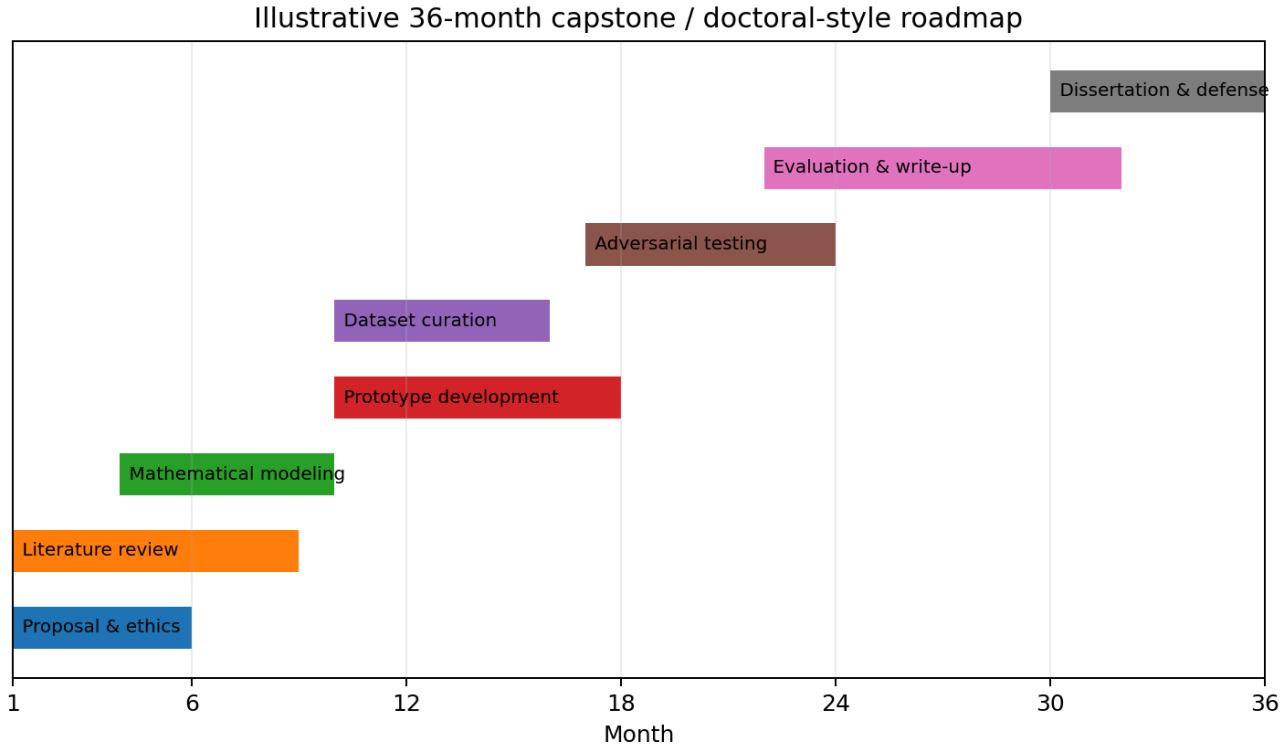

Figure 6. Example 36-month roadmap for a doctoral-style version of the capstone; the same structure can be compressed into a semester project.

Compressed 12-week certification delivery option

Week	Focus	Outputs
1–2	Problem framing and literature scan	Topic statement, gap matrix
3–4	Data governance and architecture	Dataset plan, architecture diagram
5–6	Model design and prototyping	Baseline model, Monte Carlo or NLP experiment
7–8	Security and ethics testing	Risk register, control matrix
9–10	Evaluation and explainability	Metrics dashboard, error analysis
11	Compliance and documentation	Standards mapping, audit evidence
12	Presentation and defense	Capstone report and slide deck

10. Practical discussion questions

- Why is a high-accuracy model still unsafe if uncertainty, human oversight, and auditability are missing?
- How does a fiduciary framing change the reward function or objective of an AI advisory system?
- Which risks belong at the model layer, and which belong at the governance layer?
- What evidence would a regulator, hospital, bank, or lab director ask for before approving deployment?

11. Suggested capstone grading rubric

Criterion	Weight	What strong work looks like
Problem definition and research gap	15%	Clear stakeholders, significance, novelty, and scope
Technical design and method choice	20%	Methods fit the use case and are well justified
Security, ethics, and governance	20%	Controls are explicit, realistic, and traceable
Evaluation design and metrics	20%	Baselines and metrics are appropriate and interpretable
Documentation and evidence quality	15%	Readable report, diagrams, tables, and citations
Presentation and defense	10%	Professional communication and responses to critique

12. Conclusion

This sample capstone is useful because it mirrors the complexity of real AI deployment. It does not treat AI as a single model, but as a governed system with data, models, people, risks, policies, and evidence.

That framing aligns well with the AI Certification program and with the AMK Research Lab emphasis on practical, research-grounded, and responsible AI.

Students can reuse this structure for adjacent projects in cybersecurity, health informatics, finance, education, or governed AI workforce systems.

References

1. European Union. (2024). Artificial Intelligence Act. Official Journal publication and implementation timeline resources.
2. International Organization for Standardization. (2023). ISO/IEC 42001: Artificial intelligence management systems.
3. National Institute of Standards and Technology. (2023). AI Risk Management Framework (AI RMF 1.0).
4. National Institute of Standards and Technology. (2024). Artificial Intelligence Risk Management Framework: Generative AI Profile.
5. OECD. (2024). OECD AI Principles.
6. OWASP Foundation. (2025). OWASP Top 10 for LLM Applications.
7. U.S. Department of Health and Human Services. (2024–2026). HIPAA Privacy and Security Rule guidance materials.
8. UNESCO. (2021/2024 update). Recommendation on the Ethics of Artificial Intelligence.
9. World Economic Forum. (2025). Future of Jobs Report 2025.
10. MITRE. (2025). ATLAS and SAFE-AI resources for adversarial threats to AI systems.