

DRAFT — UNOFFICIAL — NOT FOR OPERATIONAL USE

COURSE SYLLABUS

SL 5M



COURSE SYLLABUS — SL 5M: ADVANCED MACHINE LEARNING ENGINEERING

Maven Smart System (MSS) — USAREUR-AF

HEADQUARTERS
UNITED STATES ARMY EUROPE AND AFRICA
(USAREUR-AF)
Wiesbaden, Germany

DRAFT — NOT FOR OFFICIAL USE. FOR TRAINING PLANNING PURPOSES ONLY.

26 MARCH 2026

DRAFT — UNOFFICIAL — NOT FOR OPERATIONAL USE

COURSE SYLLABUS — SL 5M: ADVANCED MACHINE LEARNING ENGINEERING

MAVEN SMART SYSTEM (MSS) — USAREUR-AF

Field	Detail
Level	SL 5M — Advanced ML Engineer Specialist Track
Duration	5 days (40 hours)
Prerequisites	SL 4M complete (Go evaluation on file); 12+ months active ML engineering experience; demonstrated proficiency deploying and monitoring models in production; familiarity with MLOps concepts and distributed data frameworks
Audience	Senior ML engineers, ML platform engineers, data scientists leading production ML systems on MSS
Format	Advanced lab + architecture review + evaluated MLOps pipeline build
Location	MSS Training Environment (Python Transforms, Model Integration, Code Workspace with GPU allocation required)

PREREQUISITE WARNING: SL 5M is not required for most ML engineer billets. It is intended for engineers managing production ML systems, building the platform other data scientists use, or implementing enterprise MLOps governance.

BLUF: SL 5M addresses the full production lifecycle of ML systems at enterprise scale — not model building (covered in SL 4M), but the infrastructure, governance, and organizational patterns required to keep models working and trustworthy over time. Drift detection, automated retraining, model governance documentation, and responsible deployment are the core of this course.

LEARNING OBJECTIVES

#	Objective
1	Design and implement model drift detection pipelines: data drift (feature distribution shift), concept drift (label distribution shift), and prediction drift monitoring
2	Build automated retraining triggers and validation gates: define retraining criteria, automate the retraining pipeline, implement human approval gates before production promotion
3	Apply advanced feature engineering patterns: feature stores, time-aware feature construction, feature leakage detection at scale
4	Implement ensemble architectures and model stacking for operational reliability; document performance vs. complexity tradeoffs
5	Build a model governance package meeting USAREUR-AF standards: model card, data lineage, training data documentation, known limitations register, deprecation criteria
6	Design a responsible AI review process for operational ML systems: fairness evaluation, failure mode analysis, human review gate placement
7	Conduct a production model audit: evaluate a deployed model against its original model card; identify drift, gap, or scope creep; recommend action

PRE-COURSE CHECKLIST

Complete **7+ duty days before Day 1:**

- Confirm Code Workspace GPU allocation (C2DAO request — 7–10 duty days)
- Confirm Model Integration (model registry) access for training account
- Read TM-50M, Chapter 1 (Introduction and Scope) and Chapter 8 (ML Platform Architecture and Leadership) before Day 1
- Prepare a model card for a model you currently maintain (or maintained previously) — you will audit it on Day 4

DAILY SCHEDULE

Day 1 — Production Drift Detection

Time	Block	Method	Content
0800–0900	1	Seminar	SL 5M scope; the production model lifecycle; why models fail silently; drift taxonomy
0900–1100	2	Lab	Data drift detection: population stability index (PSI), Kullback-Leibler divergence, feature distribution monitoring pipelines
1100–1115	—	Break	
1115–1200	3	Lab	Concept drift: label distribution monitoring, prediction confidence trends, sliding window detection
1200–1300	—	Lunch	
1300–1500	4	Lab	Drift monitoring pipeline build: Foundry Transform that computes drift metrics on a weekly cadence; alert configuration
1500–1515	—	Break	
1515–1700	5	Lab	Drift response playbook: severity tiers, escalation path, retraining vs. rollback decision criteria

Evening reading: TM-50M, Chapter 2 (Automated Retraining) — trigger design and validation gate sections.

Day 2 — Automated Retraining and Feature Stores

Time	Block	Method	Content
0800–0830	—	Review	Drift detection questions; monitoring pipeline check
0830–1030	6	Lab	Automated retraining pipeline: trigger design, data version pinning, training environment reproducibility
1030–1045	—	Break	

Time	Block	Method	Content
1045– 1200	7	Lab	Validation gates: shadow mode deployment, A/B evaluation, human approval gate before production promotion
1200– 1300	—	Lunch	
1300– 1500	8	Lab	Feature stores: design principles, time-aware feature construction (point-in-time correctness), feature leakage detection
1500– 1515	—	Break	
1515– 1700	9	Lab	Feature store integration with Foundry Ontology: connecting feature pipelines to Object properties

Evening reading: TM-50M, Chapter 3 (Advanced Neural Architectures for Operational Data) — ensemble stacking and blending sections.

Day 3 — Ensemble Architecture and Responsible AI

Time	Block	Method	Content
0800– 0830	—	Review	Feature store design review
0830– 1030	10	Lab	Ensemble architectures: stacking, blending, dynamic ensemble selection; performance vs. complexity tradeoff documentation
1030– 1045	—	Break	
1045– 1200	11	Lab	Failure mode analysis: systematic failure mode identification, operational consequence mapping
1200– 1300	—	Lunch	
1300– 1500	12	Lab	Responsible AI for operational systems: fairness evaluation across subgroups, disparate impact analysis, bias documentation
1500– 1515	—	Break	
1515– 1700	13	Lab	Human review gate design: where to place human checkpoints in an automated ML pipeline; gate bypass audit trails

Evening reading: TM-50M, Chapter 8 (ML Platform Architecture and Leadership) — model card and limitations register templates.

Day 4 — Model Governance and Production Audit

Time	Block	Method	Content
0800–0900	14	Seminar	USAREUR-AF model governance package requirements; model card standard; deprecation criteria
0900–1100	15	Workshop	Participant model audit: audit the model card you prepared against the SL 5M governance checklist
1100–1115	—	Break	
1115–1200	16	Workshop	Peer model card exchange: critique a peer's model card for gaps, optimistic limitations, missing lineage
1200–1300	—	Lunch	
1300–1600	17	Lab	MLOps pipeline integration: connecting drift monitoring → retraining trigger → validation gate → production promotion in a single pipeline architecture
1600–1700	18	Prep	Evaluation scenario brief and preparation

Day 5 — Evaluated MLOps Pipeline

Time	Block	Method	Content
0800–1700	19	Eval	Evaluated exercise: Build a drift monitoring + retraining pipeline for a provided production model; produce a governance package (model card, failure mode register, deprecation criteria); defend to evaluator

Go standard: Drift monitoring pipeline functional; retraining trigger defined; governance package complete; responsible AI section present with at least one identified failure mode and its operational consequence.

PEER ADVANCED TRACKS

Track	Relevance to SL 5M
SL 5H (Advanced AI Eng)	AI inference integration with deployed ML models; shared governance requirements for AI/ML production systems
SL 5G (Advanced ORSA)	Ensemble and Bayesian methods overlap; ORSA analysts consuming ML predictions for COA analysis
SL 5L (Advanced SWE)	OSDK integration for model inference endpoints; CI/CD pipeline patterns applicable to ML pipeline deployment

USAREUR-AF Operational Data Team Syllabus SL 5M | Version 1.0 | March 2026

DRAFT