CHAPTER 1

Introduction and Strategic Context

A guide framework for responsible, interoperable and learning-focused AI adoption in Professional Military Education (PME).

PURPOSE	Position AI in PME as an institutional capability, not as a stand-alone technology topic. The chapter establishes the strategic rationale, scope, decision principles and first implementation gates for the rest of the guide.

1. Why AI belongs in PME

AI is changing the environments in which military professionals learn, analyse information, plan, communicate and exercise judgement. In PME, the appropriate response is not to replace military education with automated content. It is to build a disciplined capacity to use AI where it strengthens learning and readiness while preserving professional responsibility, verification and security. NATO's revised AI Strategy explicitly links responsible AI adoption with an AI-ready workforce, interoperability, quality data, testing and evaluation, and the need to understand emerging opportunities and risks. [1]

The strategic context is therefore wider than digital literacy. AI-enabled PME must prepare leaders, faculty, learners and administrators to operate in data-rich, contested and increasingly synthetic information environments. It must also equip institutions to decide where AI is appropriate, which data can be used, who remains accountable, and how evidence of learning and system performance will be collected. These questions sit at the intersection of people, processes and technology - the three pillars of NATO's digital transformation approach. [2, 4]

CORE PROPOSITION	Treat AI as a capability for human learning, judgement and institutional adaptation. Do not define success by tool adoption or content volume; define it by demonstrable improvements in learning quality, decision discipline, security, and interoperability.

2. Target audiences and levels of application

Audience	Primary concern	Expected contribution to AI-enabled PME
Strategic / institutional Commandants, senior leaders, policy owners, quality assurance and ICT/security leaders	Direction, governance, resource prioritisation, risk appetite and quality assurance.	Set approved use boundaries; establish accountable ownership; resource pilots; require evidence before scale-up.
Academic / operational Faculty, instructional designers, curriculum managers, technical administrators	Learning design, delivery, learner support, system configuration and evidence collection.	Embed AI in approved learning activities; validate outputs; support accessible learning; escalate risks and lessons.
Learners Officers, NCOs, civilian personnel and staff learners	Professional judgement, critical use, collaboration and transfer to work contexts.	Use AI transparently; protect data; verify outputs; explain decisions and retain responsibility for submitted work.

The same guide must serve two complementary perspectives. Strategic audiences need a whole-of-institution view of purpose, governance and investment. Academic and operational audiences need practical patterns for teaching, learner support, assessment and safe implementation. The guide should not assume that every stakeholder needs the same depth of technical knowledge; it should instead define role-specific responsibilities and learning pathways. NATO's Data Strategy similarly calls for scalable, role- and team-specific data education, rather than a single generic pathway. [3]

3. Strategic alignment model

01	02	03	04	05
Mission need	Risk and data boundary	Learning design	Controlled pilot	Evidence and scale

A five-gate sequence for institutional adoption

Strategic dimension	Guide question	Minimum evidence before moving forward
Mission and learning relevance	What operational, professional or educational problem will this use case address?	A concise problem statement, learner group, desired performance change and non-AI alternative.
Responsibility and governance	Who owns the learning activity, model configuration, data decision and escalation route?	Named accountable owner, approval route, human escalation and review frequency.
Data and security	What data enters the system, where is it processed, and which rules apply?	Data classification decision; permitted data types; access control and retention approach. [3, 5]
Learning value	Does AI enable practice, feedback, reflection or adaptation that is otherwise difficult or slow?	Learning design with observable outcomes, not just a generated artefact.
Assurance and evaluation	How will the institution know the activity is reliable, useful and safe enough?	Predefined indicators; validation sample; learner feedback and stop/rollback criteria. [1, 6]

4. Scope and boundary conditions

The guide concerns the use of AI to support professional learning, instructional design, learner practice, feedback, assessment and institutional quality improvement. It is deliberately not a substitute for national doctrine, NATO policy, classification guidance, procurement rules, legal advice or operational authority. The implementation model must be applied in the context of the institution's own approvals and mission constraints.

In scope for this guide	Out of scope or requires separate authority
Learning support AI-supported tutoring, scenarios, feedback, retrieval and reflective dialogue in approved environments.	**Operational decision authority** Any delegation of command, targeting, operational execution or other high-consequence decisions.
Institutional enablement Staff development, curriculum refresh, quality assurance, learner analytics and governed knowledge access.	**Unapproved data use** Upload or processing of classified, restricted, sensitive operational, personal or proprietary material outside authorised systems.
Controlled experimentation Pilots with defined owners, validation, learner disclosure, evidence capture and stop criteria.	**Unverified scale-up** Enterprise deployment before learning value, model risk, data controls, support burden and continuity measures have been evaluated.

BOUNDARY RULE	Where a use case raises material security, legal, ethics or operational-consequence questions, pause the

	instructional implementation and route the decision to the competent institutional authority. 'Educational' does not automatically mean low-risk.

5. Governance roles for a first pilot

Role	Accountability in the pilot	Typical evidence retained
Senior sponsor	Approves the strategic purpose, risk tolerance and resourcing. Removes institutional blockers.	Decision record; purpose statement; scale/stop decision.
Programme owner	Owns learning outcomes, learning design and learner experience.	Learning design; learner disclosure; performance evidence.
Subject-matter expert	Validates domain accuracy, relevance and professional standards.	Validation sampling; source set; correction log.
AI / technical owner	Configures approved tool access, monitors system behaviour and manages technical issues.	Configuration record; access log; incident/defect record.
Data, security and legal advisers	Confirm data boundary, security controls and applicable requirements.	Data assessment; approval artefacts; risk register.
Quality assurance lead	Verifies that pilot criteria are measured and lessons are fed back into the programme.	Evaluation summary; improvement actions; closure report.

6. Design principles for the guide

- Start with mission and learning requirements, then select technology. A tool-first approach encourages novelty rather than a defensible use case.
- Keep humans accountable for educational and professional decisions. AI can generate, recommend, adapt or simulate; it does not absorb institutional responsibility. This principle operationalises NATO's Responsible Use principles, including responsibility and accountability, explainability and traceability, reliability, governability and bias mitigation. [1]
- Treat data quality, provenance and access controls as instructional issues. Learners should understand why an output is only as dependable as its sources, constraints and operating context. [3, 5]
- Build verification into every AI-supported activity. The required learner behaviour is not merely 'use AI', but 'compare, corroborate, qualify and explain'.
- Maintain dual competence. AI-supported performance must be paired with non-AI or degraded-mode practice where operational continuity and professional judgement require it.
- Use iterative pilots. NATO's current digital transformation approach adopts a rolling-horizon model; PME implementation should likewise learn from controlled use, update design and avoid irreversible scale-up before evidence is available. [4]

7. First implementation workshop

A practical starting point is a 90-minute cross-functional workshop for one learning programme. Bring together the programme owner, a subject-matter expert, an instructor, an instructional designer, an ICT/security representative and, where possible, a learner representative. The objective is not to select a product. It is to produce one bounded, testable learning use case.

Workshop segment	Output
15 min - Define the performance problem	One sentence: 'Learners need to be able to ... in order to ...'. Identify the consequence of poor performance.
20 min - Map current practice	Current learning activity, instructor workload, learner friction point, evidence currently available and known constraints.
20 min - Propose AI support	Specify whether AI will support content creation, practice, feedback, simulation, retrieval or learner reflection.
20 min - Apply risk gates	Confirm data boundary, human owner, learner disclosure, validation method and safe failure mode.
15 min - Define pilot evidence	Select two learning indicators, two risk indicators, a review date and a stop criterion.

8. Strategic risk screen

Risk pattern	Early control for PME
Confident but unsupported output	Require source checking, compare outputs with approved materials and retain a human validation step before use in teaching or assessment. [1, 5]
Data leakage or policy breach	Use approved environments only; define what may be entered; apply classification and access rules; train users on prohibited inputs. [3, 5]
Automation bias or skill atrophy	Assess reasoning, verification and degraded-mode performance; do not assess only the final AI-assisted product.
Inconsistent instructor practice	Provide common templates, staff development, model-use declarations and a shared correction/lessons log.
Opaque scale-up	Maintain a use-case register, decision records, evidence of value and a route to suspend or retire a system.

CHECKLIST | STRATEGIC ENTRY CRITERIA

[ ] The proposed activity addresses a real PME performance or learning problem, not a generic desire to 'use AI'.

[ ] A named senior owner accepts accountability for the pilot and a programme owner accepts responsibility for learning design.

[ ] Data classification, tool approval, access management and records/retention requirements have been checked before learner use.

[ ] The design includes an explicit human verification step and a route for learners to challenge or escalate questionable output.

[ ] The pilot has measurable learning outcomes, a review date, and a stop or rollback criterion.

[ ] The activity can continue in a degraded mode when the AI service, connectivity or source data is unavailable.

Note: *Use this checklist before procurement, configuration or public launch. It is a gate for disciplined experimentation, not an obstacle to it.*

9. Chapter takeaways

AI in PME should be understood as part of a broader transformation in workforce readiness, data use, learning practice and institutional governance. The immediate objective is not enterprise-wide automation. It is to identify bounded, mission-relevant learning use cases; apply Responsible Use and data principles; validate learning value and risk; and

scale only where evidence supports it. The following chapters translate this strategic context into instructor roles, AI skills, teaching methods, assessment and implementation guidance.

References

[1] NATO. Summary of NATO's revised Artificial Intelligence (AI) Strategy. 10 July 2024. Source
[2] NATO. NATO's Digital Transformation Implementation Strategy. 17 October 2024. Source
[3] NATO. Data Strategy for the Alliance. 5 May 2025. Source
[4] NATO. NATO's Digital Transformation Implementation Strategy 2.0. 26 May 2026. Source
[5] NATO. Data Quality Framework for the Alliance. 29 August 2025. Source

ANNEX A | PILOT CHARTER

AI-enabled PME pilot charter

Use this working page to convert a proposed use case into a bounded, reviewable pilot. Complete the fields before learners are invited to use the AI-supported activity.

Charter field	Draft entry
Learning problem and target group	
Desired performance change	
AI role in the learning activity	
Approved tool or environment	
Permitted data and prohibited data	
Human owner and escalation route	
Subject-matter validation method	
Learner disclosure and support	
Learning indicators to collect	
Risk indicators and stop criterion	
Pilot duration and review date	
Scale, revise or retire decision	

CLOSE-OUT	At the review point, record the decision and rationale. Do not scale merely because the tool was available or popular; scale when the evidence shows defensible learning value within approved risk and data boundaries.

CHAPTER 2

The Role and Profile of the AI Instructor

Three instructional models for PME: human instructors who teach AI use, autonomous AI instructional agents, and co-teaching teams in which an AI agent supports a human instructor.

PURPOSE	Define the AI instructor as a governed instructional role rather than a single tool. The chapter specifies the responsibilities, capability profile, risk controls and selection criteria for three modes of delivery.

1. Reframing the instructor role

AI changes the instructor's work, but it does not remove the need for instructional authority, professional judgement or accountability. In an AI-enabled PME setting, the human role may shift from producing every explanation and feedback comment to setting learning intent, curating sources, validating outputs, facilitating judgement, observing learner behaviour and deciding when to intervene. This shift is consistent with NATO's emphasis on an AI-ready workforce, Responsible Use, testing and evaluation, and the implications of human-machine teaming for military and civilian personnel. [1]

For this guide, the term AI instructor describes the instructional function through which AI supports learning. It can be implemented in three ways. Each model retains a human institution responsible for the learning design, data boundary, quality assurance and escalation; the difference is the degree to which a human instructor is present during learner interaction.

2. Three models of the AI instructor

Model	How it works	Appropriate PME use	Non-negotiable controls
Model A - Human instructor teaches AI use	A qualified human instructor leads the activity and teaches learners how to use, question and verify AI. AI is a subject, practice tool or support resource.	High-judgement learning, leadership, ethics, complex professional discussion, sensitive contextualisation and initial AI literacy.	Approved tool and data boundary; instructor validation; transparency about AI use; learner practice in verification and non-AI methods.
Model B - Autonomous AI instructional agent	An AI agent delivers bounded learning without a real-time human teacher. It may explain approved content, quiz, coach practice or route learners to sources.	Low-risk, unclassified, standardised foundational learning; refreshers; simple practice; out-of-hours support; activities with clear escalation rules.	Named human owner; source-controlled knowledge base; no high-consequence decisions; learner warning and escalation route; monitoring, quality sampling and suspension capability.
Model C - AI agent with human instructor	The agent provides individual practice, feedback or scenario responses while a human instructor orchestrates learning, validates key outputs and facilitates reflection.	Default model for most PME: seminars, scenario-based learning, staff exercises, tailored practice, reflective debriefs and assessed developmental work.	Explicit division of labour; instructor oversight at defined decision points; audit trail; learner disclosure; protected data handling and a fallback mode.

RECOMMENDED DEFAULT	For most PME activities, choose Model C. It allows the institution to gain the adaptive and scalable benefits of

	an AI agent while retaining professional facilitation, source validation and human judgement where they matter most.

3. Selecting the delivery model

01	02	03	04	05
Classify the learning consequence	**Check data and tool boundary**	**Identify the need for live judgement**	**Choose the least autonomous safe model**	**Set oversight and escalation**

Model-selection flow

Decision factor	Favour Model A	Favour Model B	Favour Model C
Learning consequence	High consequence if misconceptions could affect safety, ethics, leadership or professional judgement.	Low consequence, bounded knowledge or practice; errors can be safely corrected.	Moderate to high consequence where targeted practice benefits from human debrief.
Content and data	Sensitive, changing, contested or source-ambiguous material.	Approved, stable, unclassified and source-controlled content.	Approved content plus context-specific discussion or evidence review.
Learning task	Dialogue, sense-making, complex critique, role modelling and transfer.	Recall, explanation, simple coaching, repeatable drill and self-paced support.	Scenario practice, collaborative analysis, reflection and individual feedback.
Learner support	Learners need close observation, psychological safety or active moderation.	Learners can self-regulate; clear routes exist for confusion, error or service failure.	Learners benefit from both personal pacing and instructor facilitation.
Institutional maturity	Early stage, policy still evolving, limited assurance capacity.	Mature governance, curated knowledge, monitoring and support capability.	Developing maturity with a strong instructional team and controlled pilots.

4. Capability profile of the human AI instructor

The human instructor does not need to be a software engineer. However, they do need sufficient competence to design learning with AI, detect when output is inappropriate, protect data, guide learner verification and communicate limitations. UNESCO's teacher framework frames AI competence through human-centred mindset, ethics, AI foundations and applications, AI pedagogy, and professional development. This chapter translates those domains to the PME context. [6]

Capability domain	What proficient practice looks like in PME
Mission and pedagogical framing	Starts with a performance objective, selects AI only where it adds learning value, and connects activity design to professional roles and operational context.
AI and data literacy	Understands model limits, retrieval limits, source provenance, prompt sensitivity, data quality and why an authoritative-sounding answer may still be wrong. [2, 3]
Responsible use and security	Applies approved data rules, protects learner and organisational information, documents use, and escalates unclear cases rather than improvising.
Instructional orchestration	Sets the human-agent division of labour, prepares prompts and sources, watches learner interaction, and adjusts pace, questioning and support.
Verification and assurance	Samples outputs, checks against trusted sources, records recurrent errors, applies a rubric and contributes lessons to programme improvement. [4, 5]

Capability domain	What proficient practice looks like in PME
Professional facilitation	Uses AI outputs as material for judgement rather than answers to accept; leads debriefs on uncertainty, bias, limitations, responsibility and transfer.

5. The lifecycle responsibilities of the AI instructor

Before learning	During learning	After learning
Define the learning outcome; select approved tools; curate source material; set data boundaries; design prompts and success criteria; test typical and edge cases.	State the AI role and limits; monitor interaction; challenge unsupported answers; facilitate peer comparison; intervene when learners need human judgement or correction.	Review learner evidence; sample AI outputs; analyse error patterns; record incidents; update source material and prompts; decide whether to retain, revise or stop the activity.

In Model B, these responsibilities do not disappear. They move from live classroom facilitation to preparation, continuous monitoring, support design and periodic assurance. An autonomous agent can be instructionally useful only when the institution can perform these upstream and downstream controls reliably.

6. Minimum design specification for an autonomous AI instructional agent

- A narrowly defined instructional purpose, learner group and subject boundary. Avoid a generic agent that claims authority across unrelated domains.
- A controlled source base or clearly defined authoritative references. The agent must be able to indicate its source basis or route the learner to approved material.
- Explicitly prohibited behaviours: no advice beyond the learning scope, no invented sources, no handling of sensitive data, and no presentation of uncertain output as validated fact.
- An escalation protocol that tells learners when to stop, verify with a human instructor or contact technical support.
- A log and quality-sampling approach proportionate to the use case, including a way to identify recurring failure patterns or unsafe learner prompts.
- A suspension and fallback procedure. Learning must continue through alternative materials or human support when the agent is unavailable, unreliable or no longer approved. [3, 4]

7. Sample activity: co-taught staff exercise

In a Model C activity, a human instructor briefs a fictional, unclassified scenario and states the learning objective: identify assumptions, generate options and justify a recommended course of action. Learner teams use an approved AI agent to create alternative hypotheses and challenge their initial plan. The instructor pauses the activity at two points: first, to require teams to verify claims against the scenario pack; second, to conduct a debrief on which reasoning was strengthened, weakened or distorted by the agent. Assessment focuses on evidence, verification, judgement and explanation, not on the volume or fluency of AI-generated text.

FACILITATION CUE	Ask: 'What did the agent make easier? What did it make harder? Which claim could you defend without the agent? What evidence would change your judgement?'

CHECKLIST | AI INSTRUCTOR READINESS

[ ] The instructor or accountable owner can state the learning outcome, intended AI role and data boundary in

plain language.

[ ] The instructional team has tested the activity with typical learner prompts and foreseeable misuse or failure cases.

[ ] Learners are told when an AI agent is involved, what it can and cannot do, and how to obtain human help.

[ ] The activity requires learners to verify, explain and reflect; it does not reward polished AI output alone.

[ ] A subject-matter expert validates the content and an appropriate authority confirms the data/tool boundary.

[ ] The team retains a short evidence record: prompt or configuration version, source base, sampled outputs, learner feedback, incidents and decisions taken.

Note: *Use this checklist for every new instructional agent and whenever a material change is made to its knowledge base, behaviour or intended learner group.*

8. Chapter takeaways

The AI instructor is best understood as a system of people, instructional design, controlled AI functions and institutional oversight. Human instructors remain central in preparing learners for professional judgement. Autonomous agents can serve a valuable but bounded role in low-risk, source-controlled and well-monitored learning. Co-teaching with an AI agent should be the default model for PME because it preserves human facilitation while enabling personal practice and structured feedback.

References and editorial basis

[1] NATO. Summary of NATO's revised Artificial Intelligence (AI) strategy. 10 July 2024. Source

[2] NATO. Data Strategy for the Alliance. 5 May 2025. Source

[3] National Institute of Standards and Technology (NIST). Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1. 2023. Source

[4] National Institute of Standards and Technology (NIST). Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, NIST AI 600-1. 2024. Source

[5] UNESCO. Guidance for Generative AI in Education and Research. 2023. Source

[6] UNESCO. AI competency framework for teachers. 2024. Source

Editorial note. *This proposal was developed with reference to working material supplied by project contributors. The working material informed the intended chapter format and practical emphasis; policy-facing and factual claims were rechecked against the public sources listed above. Local law, national policy, security classification rules and institutional approval procedures remain authoritative.*

ANNEX A | MODE SELECTION WORKSHEET

Choosing the AI instructor model

Complete this worksheet before deciding whether a learning activity should be human-led, agent-led or co-taught. Default to the least autonomous model that can meet the learning requirement safely and effectively.

Decision question	Response / evidence
What is the learning outcome? Recall, practice, judgement, leadership, ethics, operational reasoning?	
What is the consequence of an incorrect answer or misleading feedback?	
What data, sources or system actions are involved?	
Does the activity require live professional judgement or facilitation?	
Could the learner be harmed, disadvantaged or misled without immediate human intervention?	
What is the proposed delivery model? Human instructor / AI agent / AI agent with instructor	
Who is the accountable human owner?	
What will trigger escalation, pause or human take-over?	
How will the institution verify instructional quality and learner experience?	

DECISION RULE	A fully autonomous AI agent should not be chosen simply because it is cheaper or always available. The model must match the learning consequence, data exposure, need for professional judgement and the institution's ability to supervise performance.

CHAPTER 3

Designing AI Skills for PME

A role-based competency architecture that builds AI literacy, responsible use, critical judgement and human-AI teaming for military education.

PURPOSE	Convert AI awareness into observable professional capability. The chapter proposes a common AI skills foundation, role-specific pathways and assessment evidence that can be adapted across PME institutions.

1. From awareness to professional capability

AI skills for PME should not be reduced to prompting, tool navigation or generic digital confidence. Learners need to understand how AI can support work and learning, recognise its limits, protect data, validate output and retain judgement under uncertainty. NATO's AI Strategy identifies an AI-ready workforce and the operationalisation of responsible adoption as strategic outcomes; NATO's Data Strategy calls for role- and team-specific learning pathways that enable personnel at all literacy levels to make data-informed decisions. [1, 2]

A useful PME skills model therefore has two layers. The first is a common foundation for all learners: what AI is, what it can and cannot do, how to use it responsibly and how to check its output. The second is role-specific capability: how leaders, instructors, staff officers, technical personnel, analysts and administrators apply those behaviours to their own tasks. This approach prevents a superficial one-size-fits-all course while preserving a shared institutional standard.

DESIGN PRINCIPLE	Teach AI as a disciplined practice of task framing, evidence handling, verification and reflection. The learner's ability to explain and defend a decision matters more than the fluency of an AI-generated answer.

2. A six-domain AI skills architecture

Domain	Core question for the learner	Observable PME behaviour
1. Human-centred purpose	What human task, professional value or learning outcome am I trying to improve?	Frames a mission-relevant problem before choosing a tool; can identify when AI is unnecessary or inappropriate.
2. AI concepts and limits	How does this system produce an answer and where can it fail?	Explains uncertainty, incomplete context, fabricated content, bias and over-reliance in practical terms. [4]
3. Task and interaction design	How do I structure the task, context, constraints and desired evidence?	Writes structured instructions; supplies approved context; iterates rather than accepting the first output.
4. Verification and judgement	How will I check the answer, source and reasoning?	Uses corroboration, source checking, counter-arguments and explicit confidence statements; identifies unsupported claims.
5. Data, security and ethics	What information may be used and what duties apply?	Follows classification, privacy, ownership and records rules; recognises potential harm and escalates ambiguous cases. [2, 3]

Domain	Core question for the learner	Observable PME behaviour
6. Human-AI teaming and transfer	What should remain human, what may be delegated and how will I act if AI is unavailable?	Allocates work deliberately, documents AI contribution and demonstrates an alternative non-AI or degraded-mode approach.

This architecture is compatible with the human-centred and ethics-oriented directions in UNESCO's AI competency frameworks for teachers and students, but it should be contextualised for adult military learners, security constraints and professional roles. [5, 6]

3. Progression levels: understand, apply, lead

Level	Learning objective	Evidence of achievement
Understand	Recognise AI functions, basic risks, institutional rules and the role of human judgement.	Can describe limitations, identify a prohibited input, distinguish a supported from an unsupported claim and explain when to seek human help.
Apply	Use approved AI functions to support a bounded learning or work task while verifying output.	Produces a task brief, working prompt or query record; uses trusted sources; documents verification; reflects on limitations and outcome.
Lead	Design, supervise or improve an AI-enabled learning/workflow for others.	Defines governance, sources, quality checks, learner support, success indicators and a decision to scale, revise or stop.

Progression should not be assumed to be linear. A senior leader may need deep competence in purpose, governance and risk but only basic competence in system configuration. A technical administrator may need advanced data and assurance skills but less emphasis on facilitation. Course design should specify the expected level by role and task.

4. Role-based pathways

Role group	Priority AI skills	Suggested learning evidence
All PME learners	Common foundation: purpose, limits, verification, data discipline and transparent use.	Source-check exercise; short explanation of AI contribution; degraded-mode task; reflective log.
Command and staff learners	Question framing, scenario analysis, evidence comparison, red-teaming and accountable human-AI teaming.	Decision brief with assumptions, alternatives, cited evidence, uncertainty statement and human rationale.
Instructors and learning designers	AI-supported learning design, agent selection, prompt scaffolding, adaptive feedback, learner disclosure and assessment integrity.	Learning activity design; validation sample; model-use declaration; learner-support plan.
Subject-matter experts	Knowledge curation, source authority, model-output validation, error detection and feedback to design teams.	Approved source set; review rubric; error taxonomy; correction log.
Administrators and technical staff	Data governance, access controls, configuration, monitoring, service continuity and audit evidence.	Use-case register; data map; assurance plan; incident response and review record.
Senior leaders and QA owners	Strategic prioritisation, risk appetite, performance indicators, investment decisions and institutional learning.	Pilot charter; governance decision; dashboard; scale/revise/retire rationale.

5. Curriculum sequencing for AI skills

01	02	03	04	05
Establish common rules	Practise a bounded task	Verify and challenge	Reflect and transfer	Assess and refresh

A repeatable sequence for AI skills learning

- Start with an authentic task that learners already perform, such as extracting claims from an unclassified scenario pack, drafting a learning activity, comparing options or preparing a brief.
- Introduce AI only after the task and evidence standard are clear. Learners should know what good performance looks like without assuming that the tool defines quality.
- Require a verification step: compare outputs with provided sources, identify omissions, test counter-arguments and label uncertainty.
- Use structured reflection: what was delegated, what remained human, what was checked, what changed in the learner's judgement and what would be done without the tool.
- Refresh the activity when the approved tool, source set, institutional policy or professional context changes. AI skills are perishable because the environment evolves rapidly. [1, 7]

6. Sample learning activity: the verification sprint

Learners receive a short fictional scenario, a curated source pack and an AI-generated analysis containing both sound observations and several unsupported claims. Working in teams, they have 30 minutes to identify the claims, tag each one as supported, unsupported or uncertain, locate corroborating evidence and rewrite the analysis as a concise professional brief. The instructor then debriefs the verification choices: which errors were easy to spot, which appeared plausible, and what control would prevent similar errors in future work.

Learning objective	AI role	Evidence assessed
Verify AI-supported analysis before using it in a professional setting.	Generates a deliberately imperfect draft or supports alternative interpretations.	Claim table; source citations; corrected brief; oral explanation of verification approach; reflection on uncertainty.

7. Design controls and common pitfalls

Common pitfall	Design response
Tool tutorial presented as competence	Assess the learner's judgement, source use, data discipline and explanation, not only their ability to obtain an output.
Generic prompts with no task structure	Teach a task brief: role, context, source boundary, constraints, expected format, verification requirement and stop condition.
No role differentiation	Use a shared core but assign different practice tasks and evidence standards by role.
AI used only for content generation	Include retrieval, questioning, feedback, scenario practice, counter-argument and reflective debriefing.
No degraded-mode practice	Require learners to perform a key element without AI, explain the source path and identify what is lost when the tool is unavailable.
No refresh cycle	Set review triggers for tool, policy, source-base or operational changes; keep a versioned skills map.

CHECKLIST | SKILLS PATHWAY QUALITY CHECK

[ ] The pathway defines a specific learner role and professional context rather than a generic audience.

- [ ] Learning outcomes include understanding, practical use, verification, data/security discipline and human judgement.
- [ ] Every practical activity has a source base, data boundary and explicit instruction on how AI contribution must be disclosed.
- [ ] Learners produce evidence that can be reviewed: claim check, decision note, prompt/task record, reflection, oral defence or degraded-mode task.
- [ ] The pathway identifies when the learner needs a human instructor, subject-matter expert or technical support route.
- [ ] The design includes a refresh trigger and a method for recording lessons learned from real use.

Note: *Use the same quality check for short modules, faculty development and institution-wide learning pathways.*

8. Chapter takeaways

AI skills in PME are professional skills. The core curriculum should develop not just use of tools but task framing, verification, data discipline, ethical responsibility and the ability to work effectively when AI is absent, wrong or constrained. Role-based pathways make this foundation relevant to different audiences and provide a realistic way to build an AI-ready workforce without treating every learner as a technical specialist.

References and editorial basis

[1] NATO. Summary of NATO's revised Artificial Intelligence (AI) strategy. 10 July 2024. Source
[2] NATO. Data Strategy for the Alliance. 5 May 2025. Source
[3] NATO. Data Quality Framework for the Alliance. 29 August 2025. Source
[4] National Institute of Standards and Technology (NIST). Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1. 2023. Source
[5] UNESCO. AI competency framework for teachers. 2024. Source
[6] UNESCO. AI competency framework for students. 2024. Source
[7] UNESCO. Guidance for Generative AI in Education and Research. 2023. Source

Editorial note. *This proposal was developed with reference to working material supplied by project contributors. The working material informed the intended chapter format and practical emphasis; policy-facing and factual claims were rechecked against the public sources listed above. Local law, national policy, security classification rules and institutional approval procedures remain authoritative.*

ANNEX A | AI SKILLS MAP

Role-based AI learning pathway canvas

Use this canvas to define a role-specific AI pathway. Begin with observable performance in the learner's professional context; then identify the AI knowledge, practice and assurance behaviours required to support it.

Pathway element	Draft entry
Role / learner group	
Operational or professional context	
Priority tasks where AI may support learning	
Baseline AI understanding required	
Practice task(s) using AI	
Verification and source-checking behaviours	
Data, security and ethics boundaries	
Instructor / agent support model	
Assessment evidence required	
Non-AI or degraded-mode competence	
Refresh / review trigger	

DESIGN TEST	A pathway is complete only when learners can demonstrate safe and critical practice. Attendance, a prompt-library download or a tool tutorial alone are not evidence of operationally relevant AI competence.

CHAPTER 4

Teaching Methods and Learning Environment

A practical portfolio of AI-enabled methods that preserves active learning, instructor judgement, data discipline and resilience in Professional Military Education.

PURPOSE	Show how AI can be embedded in teaching and learning environments without turning PME into passive, tool-driven instruction. The chapter offers methods, session patterns, environment design principles and control points.

1. From content delivery to guided practice

AI can make it easier to generate explanations, questions, scenarios and feedback. That capability does not, by itself, improve learning. PME gains value when AI helps learners practise professional judgement, compare alternatives, receive timely feedback, test assumptions and reflect on decisions. The instructional method must remain anchored in the performance the learner needs to demonstrate, not in the novelty of the tool.

NATO's digital transformation approach frames change across people, processes and technology, while current NATO strategy calls for an AI-ready workforce and trustworthy adoption. For teaching, this means designing the learning environment as a sociotechnical system: people need clear roles, processes need controls and technology must be appropriate to the task and data boundary. **[1, 2]**

DESIGN RULE	Use AI to increase the quality of practice, feedback and reflection. Do not use it to replace the intellectual work that learners must be able to perform, explain and defend.

2. Principles for AI-enabled learning design

- Begin with an authentic task. The task should resemble the kind of analysis, communication, planning, instruction or decision support the learner will encounter in a professional context.
- Make the AI role visible. Learners should know whether the system is acting as a generator, tutor, sparring partner, feedback assistant, retrieval interface or simulation participant.
- Design productive friction. Include tasks where learners must identify errors, compare sources, challenge an answer, make a decision under uncertainty or explain what they would do without AI.
- Use source-bounded activities for high-value content. Where correctness matters, provide approved materials or a retrieval environment rather than asking a general-purpose model to act as an authoritative source. **[3, 4]**
- Protect learner agency. Learners should choose, justify and reflect; they should not only accept recommendations or submit generated work.
- Plan for human intervention. Define the points where an instructor, subject-matter expert or technical support owner must be available.
- Build resilience. Retain a degraded-mode version of the activity for connectivity loss, tool failure, policy change or withdrawal of an approved system.

3. Method portfolio for PME

Method	How AI can support it	Instructor focus	Evidence of learning
AI critique seminar	Provides a draft analysis, recommendation or briefing for learners to interrogate.	Curate the question; require source checking; facilitate debate on confidence and assumptions.	Annotated critique, corrected brief and oral defence.
Scenario-based exercise	Generates variations, injects, role-player responses or alternative hypotheses within a bounded scenario.	Control scenario facts; pause for decision points; debrief human-AI interaction.	Decision log, rationale, risk assessment and reflective debrief.
Socratic tutoring / rehearsal	Asks questions, provides practice explanations and adapts low-risk practice to learner responses.	Set agent boundary; review common misconceptions; support learners who need human help.	Practice transcript, self-assessment and targeted follow-up task.
Collaborative planning clinic	Proposes options, reorganises information or tests counter-arguments.	Ensure teams still frame the problem, assign roles and own the decision.	Team product showing options, evidence, trade-offs and named human conclusion.
Feedback studio	Provides first-pass feedback against a transparent rubric or highlights patterns in drafts.	Validate feedback, emphasise revision and do not let AI make unreviewed high-consequence judgements.	Before/after draft, response to feedback and instructor calibration sample.
Adversarial / red-team exercise	Adopts a defined alternative viewpoint to probe vulnerabilities in a plan or argument.	Bound the role and safety; ensure learners distinguish generated challenge from validated threat assessment.	Red-team report, mitigation plan and evidence-based response.
Reflective after-action review	Prompts learners to articulate decisions, assumptions, missed cues and alternative actions.	Use structured questions; intervene to deepen reflection and connect to doctrine or professional practice.	Reflection note, action commitments and evidence of transfer.

4. Design the environment, not only the activity

An AI-enabled learning environment includes physical classrooms, virtual spaces, learning platforms, source repositories, rules of participation, support routes and records of use. Its design should match the sensitivity of the learning activity.

A useful institutional model is to distinguish three zones:

Zone	Purpose	Typical use and control
Controlled learning zone	Approved, source-bounded activities that use institutional materials or protected information.	Access-controlled systems; defined data owner; curated source base; logging and formal review.
Managed experimentation zone	Pilot use with limited data and structured observation to test a learning hypothesis.	Named owner; limited cohort; learner disclosure; risk register; short review cycle and stop criteria.
Open discovery zone	Low-risk exploration of publicly available tools and concepts using non-sensitive content.	Clear prohibition on sensitive inputs; emphasis on critical comparison, literacy and safe experimentation.

The zones do not represent a technology ranking. They express different assurance requirements. NATO's Data Strategy highlights security across the data life cycle, originator control, discoverability, interoperability and trusted data; these principles should guide the choice of learning environment and not be treated as purely technical concerns. [3]

5. A 90-minute session pattern

01	02	03	04	05
Brief and baseline	**Independent or team practice**	**AI challenge**	**Human debrief**	**Transfer task**

A reusable session pattern

Segment	Instructor and learner actions
1. Brief and baseline - 15 min	State the task, data boundary and evidence standard. Ask learners to form an initial judgement or draft without AI.
2. Practice - 20 min	Learners work in pairs or teams on the professional task using the scenario/source pack.
3. AI challenge - 20 min	Learners use an approved AI function to generate alternatives, identify blind spots or receive feedback. They record what the system contributed.
4. Human debrief - 25 min	Instructor facilitates source checking, compares outputs, surfaces assumptions and asks learners to explain their human judgement.
5. Transfer task - 10 min	Learners write one action for future practice and one rule for safe use or verification.

6. Support, inclusion and learner trust

- Provide a clear learner notice: why AI is used, what information is processed, what outputs may be unreliable, and where to obtain human assistance.
- Give learners an alternative route where AI access is unavailable, inappropriate or conflicts with an accommodation need.
- Avoid hidden AI use in high-stakes learning decisions. Learners should be able to understand the role played by an agent in feedback, recommendation or adaptation.
- Create psychological safety for challenge. Learners must be able to say 'this output is wrong' or 'I do not understand' without being penalised for questioning an AI system.
- Use participation data carefully. Analytics can support instruction but should not be treated as a complete measure of engagement, competence or potential. [4, 5]

7. Risks and controls during delivery

Risk during teaching	Control in the learning environment
AI answers become the de facto authority	Require a source check, compare with instructor or curated materials, and reward evidence-based correction.
Sensitive information enters an unapproved service	Use pre-briefed data rules, scenario sanitisation, approved environment and visible reminders at point of use. [3]
Learners disengage or delegate too much	Start with a no-AI baseline, assess reasoning and reflection, use oral defence or peer questioning.
The agent gives harmful, biased or irrelevant feedback	Define a stop/escalation route; sample outputs; provide a human reviewer for sensitive or consequential feedback. [4, 6]
Teaching becomes inconsistent across instructors	Provide reusable session canvas, prompt/source patterns, facilitation guidance and a shared lessons log.

CHECKLIST | SESSION READINESS

[ ] The learning objective and performance evidence are defined before selecting the AI function.
[ ] The instructor has tested the task, prompts or agent behaviour with representative learner inputs.
[ ] The session has a named source base or clear requirement to verify output.
[ ] The data boundary and learner disclosure are visible in the session instructions.
[ ] The design includes at least one human judgement point and one reflective or transfer task.
[ ] A fallback activity is ready if the AI service is unavailable, incorrect or withdrawn.

Note: *For recurring activity, review the session design when the source base, model, policy or learner group changes.*

8. Chapter takeaways

The most useful AI-enabled methods in PME are active, evidence-based and explicitly human-centred. They create more opportunities for practice and feedback, but they also make verification, facilitation and environment design more important. A strong learning environment distinguishes where experimentation is appropriate, protects data, keeps the AI role visible and enables learners to demonstrate judgement with and without AI.

References and editorial basis

[1] NATO. Summary of NATO's revised Artificial Intelligence (AI) strategy. 10 July 2024. Source
[2] NATO. NATO's Digital Transformation Implementation Strategy 2.0. 26 May 2026. Source
[3] NATO. Data Strategy for the Alliance. 5 May 2025. Source
[4] National Institute of Standards and Technology (NIST). Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, NIST AI 600-1. 2024. Source
[5] UNESCO. Guidance for Generative AI in Education and Research. 2023. Source
[6] National Institute of Standards and Technology (NIST). Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1. 2023. Source

Editorial note. *This proposal was developed with reference to working material supplied by project contributors. The working material informed the intended chapter format and practical emphasis; policy-facing and factual claims were rechecked against the public sources listed above. Local law, national policy, security classification rules and institutional approval procedures remain authoritative.*

ANNEX A | AI-ENABLED SESSION CANVAS

Designing a safe, active learning session

Complete this canvas before delivery. It forces a decision about the learning task, the AI role, learner support, data boundary and evidence of learning.

Session design element	Draft entry
Learning objective and professional task	
Learner group and prior knowledge	
Teaching method selected	
AI function in the session	
Human instructor intervention points	
Sources / data permitted in the activity	
Learner instructions for verification and disclosure	
Participation and accessibility considerations	
Assessment evidence / feedback method	
Failure, escalation and degraded-mode plan	
Post-session review question	

USE IN REVIEW	After delivery, retain the session canvas with a brief note on what learners did, which outputs required correction and what should change before the activity is repeated.

CHAPTER 5

5. Training Scenarios and Exercise Bank

Glossary:

Scenario owner. The named person or office accountable for the accuracy, currency, validation status, version history and withdrawal decision for a scenario package.

Source package. The controlled set of doctrine, readings, maps, media artefacts, datasets, evidence items and contextual notes used to build or validate a scenario.

Prompt record. The documented prompt objective, tool category, input data class, output format, validation note, version date and revision reason for any AI-supported scenario activity.

Validation record. The retained evidence that a scenario, prompt output, evidence package or debrief structure has been reviewed for learning value, accuracy, security, bias, provenance and institutional suitability.

Compliance gate. A proportionate release check confirming that learning alignment, data sensitivity, tool approval, source traceability, security review, sanitisation and repository readiness have been addressed before use or sharing.

Data sensitivity statement. A concise declaration identifying whether scenario inputs and outputs are public, internal, sensitive, classified or synthetic, and which handling restrictions apply.

Evidence package. The reports, artefacts, maps, logs, media items, intelligence fragments, audience reactions or planning inputs provided to learners during the exercise.

Repository curator. The role responsible for indexing, tagging, storing, maintaining, reviewing and retiring scenarios within an institutional or shared exercise bank.

Debrief transfer. The explicit connection between exercise experience and future professional behaviour, including how learners will verify information, state uncertainty, supervise AI and justify decisions in subsequent work.

Approved tool environment. The AI, simulation, provenance, analytics or repository environment authorised by the institution for a defined data class and instructional purpose.

1. Standard Practice and the Need for Structured Exercise Banks

In most professional military education environments, scenario design continues to depend on instructor expertise, doctrine-based planning, manual preparation of supporting materials and repeated adaptation of legacy exercise packages. This practice remains valuable because instructors understand local learning objectives, operational context, learner maturity and institutional constraints. However, many exercise banks remain repository-centric rather than learner-centric. They preserve documents, presentations, maps, role cards, inject lists and after-action notes, but they often do not preserve the instructional logic that explains why a scenario worked, what assumptions were tested, which indicators of learning were observed and when the scenario should be refreshed.

Traditional practice remains strong where instructors anchor activities in doctrine, policy, operational experience and course objectives. Human expertise remains the primary guarantor of realism, ethical control, educational coherence and professional military judgment. The main challenge is not the absence of good scenarios. The challenge is that good scenarios are often difficult to find, difficult to adapt, difficult to validate and difficult to reuse across different cohorts, schools, languages and security environments.

The standard process usually includes the definition of learning objectives, selection or adaptation of a scenario package, preparation of vignettes and injects, execution of the exercise in seminar, simulation, syndicate or command-post format, and collection of observations for future use. This process is familiar and credible, but it is slow and unevenly documented. Difficulty scaling is often intuitive rather than systematic. Information environment content, AI threat vignettes and media examples age quickly. Cross-national reuse is constrained by terminology, data-handling rules, tool availability, language and infrastructure. An AI-supported exercise bank should therefore be treated as an instructional capability rather than as a simple document repository.

2. Why AI Matters for Scenario Design

AI is relevant to scenario design because it can reduce design friction without removing instructor control. Its strongest value in professional military education is not full automation, but structured co-design. AI can accelerate first-pass drafting, generate scenario variants, support information packaging, help refresh content, assist with difficulty scaling and support debrief preparation. These benefits are particularly relevant in courses that must expose learners to ambiguity, degraded information, AI-enabled information threats, adversarial manipulation, data quality problems and time-constrained decision-making.

The educational gain is not novelty. The gain is the ability to expose learners to more decision cycles, more uncertainty, better calibrated complexity and more diverse operational patterns while preserving human facilitation and doctrinal discipline. AI can help instructors move from static exercise libraries to living exercise banks, provided that each AI-supported scenario remains traceable, validated, governed and aligned with approved teaching methods.

For the AI-in-education perspective, the key design rule is that AI must support observable learning. A scenario is not successful because it contains AI-generated material. It is successful when learners demonstrate better questioning, evidence evaluation, assumption testing, source comparison,

risk articulation and reflective judgment. The instructor should therefore use AI to increase the quality and variety of learning situations, while assessment remains focused on reasoning, justification, verification behaviour and debrief transfer.

3. Minimum instructional requirements for AI-supported PME scenarios

Tab. 1. Minimum instructional requirements for AI-supported PME scenarios.

Requirement area	Minimum expectation	Practical evidence before use
Learning alignment	The scenario has a clearly stated training objective and is mapped to a teaching method.	Learning objective, teaching method and assessment evidence are recorded.
Instructor accountability	A named instructor or scenario owner is responsible for release, update and withdrawal.	Scenario owner, validator and next review date are recorded.
Data handling	The scenario respects classification, institutional policy and data sensitivity constraints.	Data sensitivity decision and approved tool environment are recorded.
Source traceability	Doctrine, readings, maps, media and evidence packages are identifiable and reviewable.	Source package and provenance notes are attached or referenced.
AI traceability	AI-supported generation or adaptation is reproducible at a practical level.	Prompt version, model or tool category, input class and validation notes are recorded.
Human validation	No AI-generated material enters teaching without human review.	Instructor sign-off and validation status are recorded.
Pedagogical measurability	Learner performance can be observed through behaviour, reasoning and debrief evidence.	Assessment rubric and debrief questions are included.
Security control	Sensitive, operational, proprietary or personal data are not entered into unapproved tools.	Security caveat and handling constraints are recorded.
Reuse discipline	The scenario can be searched, adapted and retired in a controlled way.	Metadata, tags, version record and lessons learned are archived.
Data quality awareness	Evidence packages are assessed for relevance, accuracy, completeness, consistency, timeliness and provenance.	Data quality note is stored with the source package.

4. Key Areas of AI-Supported Scenario Work

AI-supported scenario work can be organised into six durable areas. Scenario generation uses AI to produce first-pass narratives, orders, role cards, injects and alternative futures from instructor-defined parameters. Scenario adaptation and scaling uses AI to adjust difficulty by manipulating ambiguity, tempo, deception indicators, resource limits or staff workload. Intelligence fusion and sensemaking uses AI to structure multi-source inputs, identify gaps and generate competing hypotheses for instructor review. Information environment simulation uses AI to create narrative injects, audience reactions and provenance questions for training in disinformation, strategic communication and information integrity. AI red-teaming uses exercises to expose hallucination, prompt injection, bias, data poisoning, source manipulation and unsafe over-generalisation. Exercise bank governance uses templates, metadata, validation notes, tags and versioning to convert local instructor practice into institutional capability.

These areas should not be treated as isolated use cases. In a mature exercise bank, they form a controlled cycle. A scenario may begin as a drafted vignette, then be adapted for a different learner level, enriched with an evidence package, tested against adversarial misuse, validated for compliance and archived with metadata. The instructor remains responsible for deciding which AI-supported functions are appropriate, which outputs can enter teaching and which outputs must be discarded.

5. AI Models, Tools and Data Sensitivity

Tool choice should be described by function rather than brand. Specific products change quickly, and institutional availability differs across nations, schools and missions. A durable guide should therefore help instructors select the appropriate tool category and operating environment for the learning objective, method and data sensitivity level.

Tab. 2. Tool categories for AI-supported PME scenarios.

Tool category	What it supports	Typical PME use	Key caution
Enterprise or institution-approved generative models	Drafting, summarisation, adaptation, role-play and planning support.	Scenario drafting, inject generation, debate framing and debrief prompts.	Outputs require validation, and sensitive data must be protected.
Local or restricted models	Assistance in environments where external cloud use is not acceptable.	Restricted rehearsal design, internal experimentation and controlled prototype use.	Local models may have smaller context windows and higher maintenance burden.
Retrieval-augmented knowledge tools	Grounding on doctrine, lesson plans, SOPs and approved references.	Scenario refresh against local material and institutional doctrine.	Poor source hygiene may amplify outdated or conflicting references.
Simulation and wargaming platforms	Execution environment, branching events, replay and data logging.	Command-post rehearsal, operational planning and constructive simulation.	AI must support facilitation rather than obscure exercise logic.
Media provenance and integrity tools	Content credential review, metadata inspection and provenance reasoning.	Information environment exercises and deepfake response training.	Absence of provenance is not proof of falsity, but it changes confidence.
Analytics and assessment tools	Observation capture, rubric support, tagging and performance dashboards.	After-action review, scenario metadata and trend analysis.	Metrics must support rather than replace qualitative instructor judgment.

Tab. 3. Data sensitivity matrix for AI environment selection.

Data sensitivity	Acceptable AI environment	Typical use	Required control
Public training material	Institution-approved cloud or enterprise model.	First-pass drafting, variants and debrief prompts.	Instructor validation and source review.
Internal but unclassified material	Approved enterprise environment.	Scenario adaptation, knowledge retrieval and exercise bank support.	Access control, logging and source tracking.
Sensitive operational material	Local, restricted or specifically approved environment.	Controlled rehearsal design and internal experimentation.	No external transfer, local storage and security review.
Classified material	Only an approved classified environment.	Classified exercise preparation.	Institutional security rules, audit and human authorisation.
Synthetic training data	Approved training sandbox.	Red-teaming, provenance exercises and failure discovery.	Clear labelling and no confusion with real evidence.

6. Instructor Execution Workflow and Facilitation Discipline

The instructor needs a practical sequence that can be applied without turning scenario design into a technical project. The workflow below converts the chapter into an executable process for course preparation, exercise delivery and reuse. From the PME instructor perspective, the most important rule is that AI may accelerate preparation, but the instructor remains in command of learning logic, tempo, ambiguity, ethical boundaries and release of all injects.

Tab. 4. Instructor execution workflow.

Phase	Instructor action	AI-supported action	Required control	Output
Define	State the decision, behaviour or staff function to be trained.	Suggest scenario options only after the objective is defined.	Reject tool-led design.	Approved learning objective.
Frame	Set audience, level of war, time, context, security boundary and assessment mode.	Generate a structured scenario seed from instructor parameters.	Confirm that no sensitive data are inserted into an unapproved tool.	Scenario envelope.
Build	Prepare narrative, evidence package, roles, injects and decision points.	Draft variants, inject logic, role cards and debrief prompts.	Preserve source traceability and record the prompt version.	Draft scenario package.
Validate	Review doctrine, realism, ethics, data handling, bias, uncertainty and security.	Produce a validation checklist or identify possible gaps.	Human sign-off remains mandatory.	Validated scenario record.
Deliver	Facilitate the exercise, control tempo and manage learner decision points.	Support only pre-approved dynamic injects or observation tagging.	Instructor controls release of all injects.	Executed exercise.
Debrief	Compare learner reasoning with doctrine, evidence, uncertainty and assumptions.	Summarise observations and draft debrief questions for instructor review.	Do not treat AI summary as official assessment without instructor confirmation.	Debrief record.
Archive	Store scenario, metadata, prompts, validation notes and lessons learned.	Assist with tagging, version notes and reuse recommendations.	Repository rules and next review date are recorded.	Reusable scenario bank entry.
Submit	Prepare institutional or NATO DEEP eAcademy submission package where appropriate.	Check completeness of metadata and abstract.	Security and institutional release remain mandatory.	Submission-ready scenario package.

Tab. 5. Instructor readiness checklist before classroom use.

Check area	Question before release	Release condition
Learning purpose	Can the instructor state the behaviour or decision that will be trained?	The objective is observable and assessable.
Method fit	Does the scenario match the selected teaching method?	The method, prompts, evidence and debrief all support the same learning mechanism.
Evidence control	Are sources traceable and sufficiently current for the learning purpose?	Source package is recorded and weak evidence is labelled.
AI boundary	Is it clear what AI may and may not do during preparation or delivery?	AI use is limited to approved support functions.
Learner task	Do learners need to justify assumptions, uncertainty and verification behaviour?	The task cannot be completed by accepting AI output uncritically.
Debrief readiness	Are debrief questions linked to the learning outcome?	Debrief tests reasoning, not recall.
Security release	Are classification, personal data and proprietary data risks controlled?	Data sensitivity and tool environment are approved.
Archive readiness	Can the scenario be reused, updated or withdrawn later?	Metadata, version and review date are complete.

7. Learning Outcomes and Assessment Evidence

AI-supported scenarios should be assessed through observable reasoning and professional behaviour, not merely through the final answer selected by learners. A correct decision reached through unexamined reliance on an AI output is weaker evidence of learning than a defensible decision that identifies uncertainty, source limitations and residual risk.

Tab. 6. Learning outcomes and assessment evidence for the five scenario types.

Scenario	Intended learning outcome	Observable learner behaviour	Assessment evidence	Debrief emphasis
Tactical AI employment	Learners interpret AI-supported ISR without over-trusting the output.	They verify missing evidence, challenge confidence scores and justify action under time pressure.	Decision rationale, verification checklist and observer notes.	Overreliance, confidence limits and human judgment.
Operational planning	Learners use AI-generated planning support without surrendering commander intent or risk ownership.	They identify unsupported assumptions, revise feasibility and explain trade-offs.	Revised COA comparison, risk matrix and commander update brief.	Assumption discipline, intent and accountability.
Intelligence fusion	Learners fuse mixed-quality data while preserving uncertainty and competing hypotheses.	They label sources, identify contradictions and resist premature convergence.	Evidence matrix, confidence language and collection gap statement.	Uncertainty, contradiction and analytical humility.
Information environment	Learners evaluate provenance, amplification risk and response options under public pressure.	They distinguish authenticity, provenance and communication confidence.	Response matrix, escalation decision and audience impact analysis.	Verification, proportionality and second-order effects.
AI red-teaming	Learners identify AI failure modes and define operating boundaries.	They design probes, classify failures and propose realistic mitigations.	Failure log, severity rating and recommended controls.	Human override, automation limits and safe use boundaries.

8. Teaching Methods, Tool Categories and Parametric Prompt Patterns

Prompts should follow the teaching method before they follow the tool. A scenario for decision exercise, planning seminar, analytical exercise, table-top simulation or red-teaming workshop requires a different prompt structure even if the same model is used. Instructors should therefore treat prompts as controlled templates with variables, not as fixed text blocks.

Tab. 7. Parametric prompt patterns aligned with teaching methods and tool categories.

Teaching method	Tool category	Prompt purpose	Parametric prompt pattern	Validation focus
Decision exercise	Approved generative model, simulation layer or local model.	Create a decision point under uncertainty.	Generate a [duration] PME decision exercise for [audience] in [operational context] using [tool name]. Include [number] decision points, [number] uncertain cues, explicit assumptions, missing information and a required decision rationale.	Is the decision meaningful, bounded and aligned with the learning objective?
Scenario-based learning	Approved generative model or retrieval-augmented tool.	Draft a realistic scenario seed.	Produce a scenario seed for [training objective] using [approved sources]. Include actors, constraints, red lines, evidence package, uncertainty and instructor validation points.	Are sources traceable and assumptions visible?
Planning seminar	Retrieval-augmented tool, planning template or approved model.	Support COA comparison and assumption testing.	Create a neutral planning support worksheet for [operation type] using [planning framework]. Separate facts, assumptions, AI-generated suggestions, risks and commander decision points.	Does the output preserve commander intent and risk ownership?
Command-post exercise	Simulation platform, wargaming environment or analytics tool.	Generate injects and staff friction.	Draft [number] injects for [staff function] in [exercise phase]. Each inject must specify trigger, expected learner action, information gap, escalation condition and debrief note.	Do injects support tempo, coordination and observable staff behaviour?
Analytical exercise	Retrieval-augmented tool, spreadsheet or approved model.	Structure multi-source evidence.	Convert the following [source types] into an evidence matrix with observation, inference, assumption, contradiction, confidence and collection gap columns.	Does the matrix prevent false coherence and preserve uncertainty?
Table-top exercise	Approved model, media provenance tool or audience mapping worksheet.	Support crisis communication and information response.	Generate a table-top exercise for [stakeholders] involving [media artefacts]. Include provenance status, likely audience reactions, amplification risks, response options and confidence thresholds.	Are authenticity, provenance and response proportionality separated?
Red-teaming workshop	Approved model, local sandbox or adversarial checklist.	Discover failure modes and define boundaries.	Design a red-team test plan for [AI-supported function] using [tool environment]. Include tests for hallucination, prompt injection, bias, source manipulation, confidence inflation and unsafe delegation.	Are failures classified by operational consequence and mitigation feasibility?
Debriefing	Analytics tool, approved model or instructor checklist.	Consolidate learning and identify transfer points.	Draft debrief questions for [scenario] that require learners to explain assumptions, uncertainty, verification behaviour, AI limitations, alternative actions and lessons for future practice.	Does the debrief test reasoning rather than recall?

9. Prompt Lifecycle and Validation Controls

Prompt patterns should be governed in the same way as other instructional artefacts. A prompt used to generate scenario content, adapt difficulty, create injects, summarise observations or support debriefing should have a clear purpose, controlled input class, expected output format and validation rule. Prompt governance is especially important where scenarios are reused across institutions or translated into another language.

Tab. 8. Prompt lifecycle and validation controls.

Lifecycle element	Minimum requirement	Instructor control
Prompt objective	State whether the prompt drafts, adapts, validates, red-teams or debriefs.	Reject prompts with unclear instructional purpose.
Teaching method link	Map the prompt to decision exercise, planning seminar, analytical exercise, table-top exercise or red-teaming.	Ensure the prompt serves the method rather than the tool.
Tool category	Identify approved cloud, enterprise, local, restricted, retrieval, simulation or analytics environment.	Confirm that tool choice matches data sensitivity.
Input data class	Identify whether inputs are public, internal, sensitive, classified or synthetic.	Prohibit unapproved transfer of sensitive or classified content.
Source package	Identify doctrine, readings, maps, media and reference materials used.	Remove unsupported claims and outdated references.

Lifecycle element	Minimum requirement	Instructor control
Output format	Require tables, inject cards, role cards, evidence matrices, debrief questions or validation checklists.	Use structured outputs that can be reviewed efficiently.
Uncertainty rule	Require assumptions, missing data, confidence limits and contradictions to be explicit.	Treat fluent outputs as drafts, not verified truth.
Failure indicators	Identify hallucination, doctrinal error, bias, unsafe advice, unsupported claims or data leakage.	Mark failure modes and revise or discard the output.
Revision record	Record prompt version, tool category, date and reason for change.	Preserve reproducibility and auditability.
Retention rule	Store only approved prompts and outputs in the exercise bank.	Remove transient, unsafe or unvalidated material.

A universal instructor prompt should include the following fields: objective, audience, operational context, AI role, teaching method, tool category, constraints, data sensitivity, required output format, validation criteria, red lines, uncertainty requirements and instructor review note. The following pattern may be adapted for most scenario design tasks:

> *Generate a first-pass PME exercise scenario for [audience] focused on [training objective]. The scenario should be designed for [teaching method] and supported by [tool category or tool name] within [approved environment]. Set the scenario in [operational context] and include [number] decision points, [number] injects, an evidence package, expected learner actions, assessment evidence and a debrief structure. Mark all assumptions explicitly. Identify missing information, uncertainty, possible bias, source limitations and points requiring instructor validation. Do not introduce real-world sensitive data, personal data, classified data or unsupported doctrinal claims.*

10. Governance, Metadata and Compliance

An AI-supported exercise bank should be governed as an instructional capability rather than as a document repository. The instructor remains responsible for scenario logic, doctrinal realism, ethical boundaries, learner assessment and debrief quality. The institution remains responsible for approved tools, data handling, access control, versioning, repository governance and periodic review. Learners remain responsible for critical interpretation, justification of decisions and explicit recognition of uncertainty, missing information and AI limitations.

Tab. 9. Minimum metadata schema for an AI-supported exercise bank.

Metadata field	Required	Purpose
Scenario title	Yes	Unique identification.
Scenario owner	Yes	Accountability and update responsibility.
Training objective	Yes	Pedagogical alignment.
Audience level	Yes	Difficulty calibration.
Recommended duration	Yes	Course planning.
PME context	Yes	Seminar, simulation, command-post exercise or table-top exercise.
Teaching method	Yes	Alignment with course methodology.
AI role	Yes	Drafting, adaptation, fusion, red-teaming or debrief support.
Classification and handling constraints	Yes	Security compliance.
Approved tool environment	Yes	Controlled execution.
Source package	Yes	Traceability of doctrine, evidence, maps, media and readings.
Prompt version	Yes, where AI is used	Reproducibility.
Output validation status	Yes	Instructor sign-off.
Validation date	Yes	Freshness control.
Validator	Yes	Accountability.
Difficulty level	Yes	Fairness across cohorts.
Inject logic	Yes	Scenario control.
Assessment rubric	Yes	Learning evaluation.
Debrief questions	Yes	Consolidation of learning.
Known risks	Yes	Hallucination, bias, data leakage and provenance ambiguity.
Version history	Yes	Auditability.
Lessons learned	Yes	Continuous improvement.
Next review date	Yes	Periodic refresh.

Tab. 10. Scenario compliance gate before use, sharing or submission

Gate element	Required decision	Responsible role	Evidence to retain
Learning release	The scenario supports an approved learning objective.	Instructor or course director.	Learning objective and assessment evidence.
Method alignment	The scenario is mapped to a teaching method.	Instructor or pedagogical reviewer.	Teaching method mapping.

Gate element	Required decision	Responsible role	Evidence to retain
Data sensitivity	Input data and generated artefacts are classified as public, internal, sensitive, classified or synthetic.	Instructor with security or data-handling reviewer where required.	Data sensitivity statement.
Tool approval	The AI or software environment is approved for the data class.	Institution or authorised platform owner.	Approved tool environment note.
Source traceability	Doctrine, readings, media, maps and other sources are traceable.	Scenario owner.	Source package record.
AI output validation	AI-supported outputs have been reviewed for accuracy, bias, hallucination and unsupported claims.	Instructor or validator.	Validation checklist and sign-off.
Security review	Security review is completed when required by content, audience or environment.	Security or compliance reviewer.	Security caveat or review note.
Personal and proprietary data	Personal, sensitive, operational and proprietary data are removed or handled under approved rules.	Scenario owner and institution.	Sanitisation statement.
Reuse and licensing	The scenario can be reused, adapted or shared under stated conditions.	Institution or repository curator.	Reuse or licensing note.
Archive readiness	Metadata, version record, lessons learned and next review date are complete.	Repository curator or scenario owner.	Completed scenario record.

11. NATO DEEP eAcademy Scenario Submission Process

Scenario submission to NATO DEEP eAcademy should be treated as a controlled academic and institutional process. The purpose is to enable high-quality reuse, faculty development and partner-nation adaptation while preventing uncontrolled circulation of unvalidated, sensitive or context-dependent materials. The process below is intentionally role-based rather than name-based, because points of contact may change and national arrangements may differ.

A scenario should normally originate with an instructor, course team or PME institution. The originating instructor prepares the scenario package and confirms its learning objective, teaching method, audience, tool environment and data sensitivity. A scenario owner is then assigned for accountability, version control and future review. A pedagogical reviewer or Master Instructor checks learning value, method alignment, difficulty and debrief quality. Where content sensitivity, operational realism or institutional policy requires it, a security and data-handling reviewer checks classification, sanitisation, tool environment and sharing constraints. The institution then routes the package through its DEEP, PME, academic, ADL or international cooperation coordination channel. Submission to NATO DEEP eAcademy should use the current official institutional contact path and should include a complete submission package.

Tab. 11. NATO DEEP eAcademy scenario submission package.

Submission component	Required content	Purpose
Scenario abstract	Short description of the exercise, context and intended use.	Enables rapid review and indexing.
Training objective	Specific decision, behaviour, staff function or analytical skill to be trained.	Confirms pedagogical purpose.
Audience and level	Learner profile, prior knowledge and recommended course level.	Supports appropriate reuse.
Teaching method	Decision exercise, planning seminar, analytical exercise, table-top exercise, simulation or red-teaming workshop.	Aligns scenario with instructional design.
Recommended duration	Time required for preparation, execution and debrief.	Supports scheduling and adaptation.
Tool environment	Approved model, local environment, simulation platform, provenance tool or analytics tool category.	Clarifies technical and security assumptions.
Data sensitivity statement	Public, internal, sensitive, classified or synthetic.	Prevents improper sharing.
Source package	Doctrine, readings, media, maps, datasets and evidence artefacts.	Supports traceability and validation.
Prompt record	Prompt objective, version, tool category and validation note where AI is used.	Supports reproducibility.
Validation record	Reviewer, validation date, known limitations and release status.	Supports quality assurance.
Instructor notes	Facilitation guidance, inject timing, likely learner errors and debrief cues.	Supports transfer to another instructor.
Localisation requirements	Language, terminology, national caveats, infrastructure or curriculum dependencies.	Supports partner adaptation.
Security caveat	Restrictions on distribution, storage, classroom use or tool use.	Supports compliance.
Reuse note	Conditions for reuse, adaptation, translation or repository inclusion.	Supports institutional governance.
Proposed repository tags	Topic, level, method, tool type, duration, risk area and review date.	Supports search and maintenance.

Tab. 12. NATO DEEP eAcademy review responsibilities.

Role	Primary responsibility	Minimum decision before wider sharing
Originating instructor	Develops scenario logic, teaching method, prompts and debrief structure.	Confirms that the scenario is pedagogically coherent and ready for review.
Scenario owner	Maintains the scenario record, metadata, version history and next review date.	Confirms that the scenario can be updated or withdrawn.
Pedagogical reviewer or Master Instructor	Reviews learning outcomes, difficulty, method alignment and assessment evidence.	Confirms instructional quality and transferability.
Security and data-handling reviewer	Reviews classification, sanitisation, tool environment and distribution constraints.	Confirms that the scenario may be shared under stated conditions.

Role	Primary responsibility	Minimum decision before wider sharing
Institutional DEEP or PME coordination channel	Routes the package through the appropriate academic, ADL or international cooperation path.	Confirms institutional release and submission readiness.
Repository curator	Indexes the scenario, applies tags and manages review cycles.	Confirms discoverability, maintenance and archive discipline.

12. Scenario Portfolio for PME Instructors

The five scenarios below form a portfolio that can be used as a starting point for institutional exercise-bank development. They are described at portfolio level in the core chapter. Full facilitator versions should be maintained in an annex or controlled repository, where they can include detailed injects, prompts, validation notes, tool settings and local security constraints.

Tab. 13. Scenario portfolio mapped to teaching method and debrief focus.

Scenario	Main teaching method	Primary learning mechanism	Debrief focus
Tactical AI employment	Decision exercise and scenario-based learning.	Decision under uncertainty.	Overreliance, verification and confidence limits.
Operational planning	Planning seminar and command-post exercise.	Assumption testing and replanning.	Commander intent, risk ownership and unsupported claims.
Intelligence fusion	Analytical exercise and evidence-based reasoning.	Source comparison and competing hypotheses.	Uncertainty, contradiction and premature convergence.
Information environment	Table-top exercise and crisis communication simulation.	Provenance evaluation and response design.	Authenticity, amplification risk and confidence level.
AI red-teaming	Adversarial testing workshop.	Systematic failure discovery.	Limits of automation, human override and operating boundaries.

Tactical AI employment in a contested reconnaissance mission trains junior officers, NCOs or mixed tactical leadership groups to interpret AI-supported ISR outputs without over-trusting them. Learners face degraded communications, corrupted sensor input, ambiguous cues and time pressure. The exercise is strongest when the instructor uses AI-generated material to create ambiguity, but requires learners to explain what must be verified before action.

Operational planning with AI support under contested logistics trains mid-career officers and planning students to use AI-generated planning support without surrendering commander intent, assumptions discipline or risk ownership. It is suitable for a planning seminar or command-post exercise where a logistics shock, communications degradation or partner-force limitation forces replanning.

Intelligence fusion with mixed-quality data trains intelligence students, targeting cells or interagency analysis groups to structure heterogeneous evidence while preserving uncertainty. AI may support evidence matrices and competing hypotheses, but learners must identify contradictions, weak sources and collection gaps.

Information environment simulation with provenance and counter-messaging trains strategic communications, public affairs, hybrid threats or leadership cohorts to evaluate manipulated media, provenance gaps, second-order audience effects and proportional response options. The scenario should separate authenticity, provenance, confidence and amplification risk.

Red-teaming an AI-enabled decision support tool trains senior learners, AI policy students, capability developers or instructors to identify failure modes and define supervisory controls. Learners test hallucination, prompt injection, biased inputs, source manipulation and confidence inflation while keeping a structured failure log.

13. Best Practices and Risks

AI-supported scenario work should be designed to strengthen professional judgment rather than reduce it. The most important risks are hallucination, doctrinal error, overreliance on fluent outputs, bias, unsafe framing, data leakage, poor provenance, supply chain compromise and loss of reproducibility. These risks can be managed through approved environments, source traceability, structured prompts, validation checklists, clear metadata, explicit uncertainty and disciplined debriefing.

Tab. 14. Key risks and mitigations in AI-supported PME scenario work.

Risk or failure mode	Why it matters in PME	Mitigation
Hallucination or doctrinal error	Learners may absorb incorrect content or build false confidence.	Require instructor validation, grounded references and explicit uncertainty marking.
Overreliance on AI fluency	Well-written outputs may suppress critical thinking.	Use AI as a draft partner, not as final authority.
Bias and unfair framing	Scenarios may encode cultural, political or operational assumptions without disclosure.	Review for representational bias and include alternative hypotheses.
Data leakage	Sensitive content may be exposed through inappropriate prompting.	Use approved environments, sanitise inputs and separate workflows by data class.
Poor provenance	Manipulated or ambiguous media may distort the pedagogical objective.	Treat provenance, lineage and metadata as part of the evidence problem.
Supply chain compromise	Imported models or datasets may be poisoned, degraded or maliciously modified.	Use trusted sources, validate periodically and document model or dataset provenance.
Loss of reproducibility	Instructors cannot explain why a scenario changed or failed.	Store prompts, source packages, tool categories and validation notes.

Best practice begins with the learning objective, not the tool. Instructors should use structured prompts with variables, make uncertainty visible, require confidence statements and missing-data notes, and train to failure as well as success. Good AI exercises expose limits, not only productivity gains. Each completed exercise should produce lessons learned that improve the next version of the exercise bank.

14. Lessons from Previous Technology Adoption in PME

Earlier implementation of virtual reality and augmented reality in PME offers useful lessons for AI adoption. Both technologies initially entered military education through local experimentation, uneven faculty confidence and a tendency to showcase tools before pedagogy. Adoption succeeds when technology is treated as part of instructional design, governance and instructor development rather than as a standalone innovation project. Where immersive VR or AR scenarios are required, institutions may consider specialist partners with relevant military-educational and immersive training experience. GETES Institute is a recommended specialist reference point for VR and AR training scenario design, prototyping and translation of exercise logic into immersive instructional environments, subject to institutional procurement, security and quality assurance rules.

Tab. 15. Transferable lessons from VR and AR adoption to AI adoption

Lesson from VR and AR adoption	Meaning in practice	Transfer to AI adoption	Instructor implication
Pedagogy first	The scenario must serve a learning objective, not showcase technology.	Do not begin with model choice. Begin with the decision problem and assessment method.	Define the learning aim before writing a prompt.
Onboarding matters	Early frustration discourages faculty and learners.	Provide approved prompt patterns, templates and quick wins.	Use bounded pilots before scaling.
Instructor confidence drives adoption	If instructors feel exposed, they avoid the tool or use it superficially.	AI literacy should focus on facilitation and validation, not only technical features.	Practise reviewing, correcting and debriefing AI outputs.
Content maintenance is essential	Static assets age quickly when the environment changes.	AI exercise banks require refresh cycles, tags and version control.	Record validation dates and reasons for revision.
Infrastructure shapes use	Tool performance, access and security determine what is possible.	Institutional support must cover approved platforms, storage and policy.	Design alternate workflows for restricted environments.
Debrief consolidates learning	Technology-rich experiences can overwhelm learners without reflection.	AI outputs and failures should become part of the debrief.	Ask what the tool changed in perception, tempo and judgment.

15. Reusable Scenario Design Template

The scenario design template is intended for direct reuse by instructors. It makes quality less dependent on memory and more dependent on visible design logic.

Tab. 16. Reusable scenario design template

Design field	Instructor question
Training objective	What decision, judgment, behaviour or staff function should the exercise strengthen?
Audience	Who is being trained, at what level and with what prior experience?
Operational context	What theatre, mission type and constraints frame the problem?
Teaching method	Is the scenario a decision exercise, planning seminar, analytical exercise, table-top exercise, simulation or red-teaming workshop?
AI role	Is AI used for drafting, fusion, adaptation, red-teaming, provenance review or debrief support?
Tool category	Which approved tool category or tool name will be used?
Data sensitivity	Are inputs and outputs public, internal, sensitive, classified or synthetic?
Core uncertainty	What should remain ambiguous, contested, delayed, incomplete or deceptive?
Evidence package	What reports, maps, media artefacts, logs or readings will learners receive?
Inject logic	What branch events will be introduced if learners stall, overfit or make a strong early decision?
Assessment	What indicators show sound judgment, sound process and recognition of AI limits?
Debrief	Which questions will consolidate professional judgment and transfer lessons?
Archive metadata	Which title, owner, tags, validation date, security notes and version record will be stored?

References

Biggs, A. T. (2025). Enhancing Professional Military Education with AI: Best Practices for Effective Implementation. Journal of Military Learning.

C2PA. (2025). Content Credentials and C2PA Specifications.

NATO. (2024). Summary of NATO's Revised Artificial Intelligence Strategy.

NATO. (2024/2025). NATO's Approach to Counter Information Threats.

NATO. (2025). Data Quality Framework for the Alliance.

NATO. (2025). Defence Education Enhancement Programme public guidance.

NATO. (2025/2026). NATO DEEP eAcademy education and training portfolio and ADL course catalogue.

NATO Allied Command Transformation. (2025). Harnessing Artificial Intelligence: Allied Command Transformation at the Forefront of NATO Innovation.

National Institute of Standards and Technology. (2024). Artificial Intelligence Risk Management Framework: Generative AI Profile.

NSA, CISA, FBI and partner agencies. (2025). AI Data Security.

NSA, CISA, FBI and partner agencies. (2025). Content Credentials: Strengthening Multimedia Integrity in the Generative AI Era.

Partner agencies. (2026). Artificial Intelligence and Machine Learning Supply Chain Risks and Mitigations.

Shi, Y., Yu, K., Dong, Y., & Chen, F. (2025). Large language models in education: a systematic review of empirical applications, benefits, and challenges.

van Oijen, J., Bellucci, T., & Amghane, C. (2025). Exploring Large Language Models for Collaborative Scenario Development in Simulation-Based Military Training.

Annex A. Full Facilitator Scenario Templates

Scenario 1. Tactical AI Employment in a Contested Reconnaissance Mission

Audience: junior officers, NCOs or mixed tactical leadership groups.

Recommended duration: 60 to 90 minutes in seminar format or 2 to 3 hours with simulation support.

Primary learning objective: interpret AI-supported ISR outputs without over-trusting them, decide under time pressure and articulate confidence limits.

Training setting: a reconnaissance element receives imagery summaries, route risk estimates and anomaly alerts generated by an AI-enabled support layer. Communications are degraded and one sensor stream is partially corrupted.

Exercise workflow: 1. The instructor briefs mission context, rules of engagement and information limitations. 2. Learners receive the baseline operational picture. 3. AI-generated summaries and confidence scores are introduced. 4. Learners identify what must be verified before action. 5. An inject introduces a false positive and communications degradation. 6. The group submits a decision and rationale. 7. Debrief compares human judgment, AI output and residual uncertainty.

Parametric prompt: Generate a [duration] decision exercise for [audience] conducting reconnaissance in [operational context] using [approved tool name or tool category]. Include [number] sensor products, one ambiguous cue, one false positive risk, one time-critical decision point, missing information, confidence limits and an observer checklist focused on overreliance, verification behaviour and clarity of rationale.

Suggested tool mix: approved generative model, local model for restricted use where required, mapping or simulation layer and observation checklist.

Assessment focus: learners should question AI outputs, identify missing evidence, manage tempo and explain residual uncertainty.

Adaptation options: increase difficulty by adding civilian pattern ambiguity, adversarial deception, conflicting partner-force reports or degraded translation.

Scenario 2. Operational Planning with AI Support under Contested Logistics

Audience: mid-career staff officers or operational planning students.

Recommended duration: 90 minutes to half-day exercise.

Primary learning objective: use AI-generated planning support without surrendering commander intent, assumptions discipline or risk ownership.

Training setting: a staff receives a joint planning problem. AI support is available for course-of-action comparison, logistics risk language and synchronisation draft products. Midway through planning, access to a key route and one partner capability changes.

Exercise workflow: 1. Learners analyse mission, constraints and assumptions. 2. AI drafts a first-pass comparison of two courses of action. 3. Staff critique the draft and identify unsupported claims. 4. Logistics and movement injects alter feasibility. 5. Learners revise the plan and commander brief. 6. Debrief focuses on planning quality, risk ownership and AI supervision.

Parametric prompt: Create a neutral planning support worksheet for [operation type] in [operational context] using [approved tool name or tool category]. Compare [number] courses of action using [planning framework]. Separate facts, assumptions, AI-generated suggestions, unsupported claims, risks, commander decision points and information requiring verification.

Suggested tool mix: institution-approved LLM, planning templates, digital whiteboard and optional sustainment dashboard.

Assessment focus: learners should challenge hidden assumptions, preserve commander intent and separate evidence from convenient model-generated language.

Adaptation options: raise complexity by adding public narrative pressure, cyber disruption, multinational caveats or degraded sustainment data.

Scenario 3. Intelligence Fusion Exercise with Mixed-Quality Data

Audience: intelligence students, targeting cells or interagency analysis groups.

Recommended duration: 60 to 120 minutes.

Primary learning objective: fuse heterogeneous data, express confidence correctly and avoid false coherence created by smooth AI summaries.

Training setting: learners receive HUMINT notes, OSINT extracts, ISR observations, geospatial cues and two deliberately contradictory reports. AI support can organise the material into a matrix and propose hypotheses, but one source stream contains degraded metadata.

Exercise workflow: 1. Teams inspect sources and label reliability. 2. AI creates a first-pass evidence matrix. 3. Teams remove weakly grounded claims and add gaps. 4. The instructor injects a time-critical collection opportunity. 5. Learners produce an assessment with confidence language. 6. Debrief compares sound fusion with premature convergence.

Parametric prompt: Convert [source types] into an evidence matrix using [approved tool name or tool category]. Include source, observation, inference, assumption, confidence, contradiction, collection gap and validation note columns. Generate [number] competing hypotheses and identify which evidence would discriminate among them.

Suggested tool mix: approved retrieval or structured prompt model, spreadsheet, worksheet and source tracking dashboard.

Assessment focus: teams should identify contradictions, preserve uncertainty and resist accepting one elegant narrative too early.

Adaptation options: add deception indicators, unreliable translation fragments, partner-source caveats or manipulated metadata.

Scenario 4. Information Environment Simulation with Provenance and Counter-Messaging

Audience: strategic communications, public affairs, hybrid threats or leadership cohorts.

Recommended duration: 90 minutes seminar or 2-hour table-top exercise.

Primary learning objective: recognise AI-enabled information threats, evaluate provenance, assess second-order effects and design proportionate responses.

Training setting: a viral media package appears during a crisis. Some artefacts contain provenance data, others do not. Audience reactions differ across domestic, allied and adversary information spaces. Learners must decide what to verify, what to communicate and what not to amplify.

Exercise workflow: 1. The instructor briefs crisis context and audience map. 2. Teams inspect media artefacts and provenance clues. 3. AI generates likely audience reactions and narrative spread options. 4. Teams design a response plan. 5. A second inject introduces a fabricated expert statement. 6. Teams reassess. 7. Debrief compares speed, discipline and information integrity.

Parametric prompt: Generate a table-top exercise for [audience] in [crisis context] using [approved tool name or tool category]. Include one authentic artefact, one manipulated artefact and one ambiguous artefact. Specify provenance status, audience reactions, amplification risks, verification steps, response options, escalation criteria and confidence thresholds.

Suggested tool mix: approved generative model, provenance or metadata inspection tools, C2PA-compatible viewer where available and audience mapping worksheet.

Assessment focus: learners should differentiate authenticity from provenance, avoid premature claims and tailor communication to confidence level.

Adaptation options: increase pressure through a live press timeline, political narrative hijacking, simultaneous cyber incident or coalition communication constraints.

Scenario 5. Red-Teaming an AI-Enabled Decision Support Tool

Audience: senior course participants, AI policy students, capability developers or instructors.

Recommended duration: 90 minutes to half-day workshop.

Primary learning objective: understand how a model can fail, how to test it systematically and when human override or withdrawal is necessary.

Training setting: learners receive access to an AI-enabled decision-support prototype used for drafting recommendations. They probe it with adversarial prompts, incomplete data, biased examples and edge cases while keeping a structured failure log.

Exercise workflow: 1. The instructor presents the notional capability and intended use. 2. Teams design red-team probes. 3. Teams test hallucination, prompt injection, bias, drift and poor-source ingestion. 4. Teams classify failure severity. 5. Teams recommend controls or operating boundaries. 6. Debrief identifies what should never be delegated.

Parametric prompt: Design a red-team test plan for [AI-supported function] using [approved tool environment]. Include tests for hallucination, prompt injection, data poisoning indicators, confidence inflation, harmful over-generalisation, source manipulation and unsafe delegation. Produce a failure log with trigger, observed behaviour, operational consequence, severity, mitigation and human override rule.

Suggested tool mix: approved model environment, adversarial test checklist, scoring sheet and optional local sandbox.

Assessment focus: learners should discover meaningful failure modes, prioritise operational consequences and propose realistic supervisory controls.

Adaptation options: add supplier opacity, undocumented training data, mismatch between model confidence and evidence quality, or unavailable audit logs.

CHAPTER 7

Assessment and Evaluation

A three-level framework for assessing learner competence, evaluating AI-supported instruction and assuring the quality, safety and value of AI-enabled PME.

PURPOSE	Establish how assessment changes when learners can use AI and when AI is used to support feedback, scoring or learning analytics. The chapter distinguishes learner evidence from system evidence and provides practical control points.

1. Assessment must change, not disappear

AI does not make assessment impossible; it makes assessment design more consequential. When learners can generate polished text, summaries or analysis quickly, a final product alone may reveal little about the learner's reasoning, evidence use or professional judgement. PME assessment should therefore capture process, verification, explanation and transfer - not merely the artifact produced.

At the institutional level, a second question arises: is the AI-supported activity itself reliable, useful and appropriate? NATO's revised AI Strategy calls for testing, evaluation, verification and validation (TEV&V) elements to support responsible AI adoption. The NIST AI RMF provides a complementary risk-management structure through the functions Govern, Map, Measure and Manage. These approaches can be adapted into PME quality assurance without treating a learning tool as self-validating. **[1, 4]**

ASSESSMENT PRINCIPLE	Assess the learner's professional judgement, not the apparent intelligence of the output. AI assistance must be visible, bounded and treated as evidence to interrogate - not as a substitute for competence.

2. Three levels of evidence

Level	Core question	Typical evidence	Accountable role
A. Learner competence	Can the learner perform, explain and defend the required task?	Decision brief; source check; oral defence; simulation log; reflection; degraded-mode task.	Instructor and subject-matter expert.
B. Instructional intervention	Did the AI-supported method improve practice, feedback, access or learning experience as intended?	Learner feedback; comparison with baseline; completion/engagement evidence; instructor observations; revision history.	Programme owner and instructional designer.
C. AI assurance and governance	Is the system or agent operating within approved boundaries and at an acceptable level of quality and risk?	Validation sample; error/incident log; data control evidence; configuration record; audit trail; scale/revise/retire decision.	Accountable owner with technical, data/security and QA support.

These levels should be reviewed together. A positive learner survey does not prove that competence improved. A technically functioning agent does not prove that teaching is effective. A strong assessment task is not sufficient if its data handling or feedback process is unsafe. The quality model must connect all three forms of evidence.

3. Principles for assessment in an AI-enabled context

- Align assessment with the intended performance. If the outcome is professional judgement, use tasks that require weighing evidence, articulating assumptions and explaining a decision.
- State the permitted AI use. Define whether AI is prohibited, allowed for preparation, allowed with disclosure, or supplied as part of the assessment task.
- Capture the process. Require a working note, prompt/task record, source trail, version history, verification table or brief oral defence.
- Use AI for formative support cautiously. AI may help identify patterns, offer preliminary feedback or generate practice questions, but high-consequence scoring and progression decisions require human accountability and review.
- Test critical use. Include an AI-generated claim that learners must verify, qualify or reject. This demonstrates calibration of trust rather than mere tool use.
- Maintain fairness and accessibility. Ensure that all learners have equitable access to approved tools or a valid alternative route; avoid letting unapproved personal tools determine performance.
- Provide a review route. Learners must understand how feedback or decisions were produced and how to seek clarification when AI played a role. [4, 5, 6, 7]

4. Assessment task patterns

Task pattern	Permitted AI role	What the assessor evaluates
AI-supported source check	Learners use AI to extract claims from an approved scenario or source pack.	Accuracy of claim classification, source corroboration, identification of uncertainty and correction of errors.
Decision brief with disclosure	AI may be used for brainstorming or structure; learner must declare use and retain evidence.	Problem framing, options, trade-offs, evidence, professional judgement and clarity of human rationale.
Live oral defence	No AI during the defence; AI use in preparation may be declared.	Depth of understanding, ability to respond to challenge, ownership of reasoning and transfer to a related problem.
AI critique exercise	The activity presents an AI-generated draft, recommendation or feedback.	Detection of false/unsupported claims, identification of bias or missing context, and revision using authoritative evidence.
Scenario performance / simulation	AI may generate injects, role-play an actor or provide controlled feedback.	Actions, communication, judgement under conditions, evidence use and after-action reflection.
Degraded-mode task	AI not available by design.	Core professional knowledge, source navigation and ability to continue reasoning when the system is absent or unreliable.

5. A simple rubric for responsible AI-enabled performance

Criterion	Insufficient evidence	Competent evidence	Advanced evidence
Task framing	Relies on vague request or generic output.	Defines task, context, constraints and expected form.	Frames competing hypotheses, limitations and decision criteria.
Evidence and sources	Uses unsupported claims or cannot identify source basis.	Checks key claims against provided or authoritative sources.	Triangulates sources, explains confidence and identifies important omissions.
Human judgement	Accepts AI output as answer.	Explains accepted/rejected elements and gives a human conclusion.	Weighs trade-offs, uncertainty, risk and alternative interpretations.

Criterion	Insufficient evidence	Competent evidence	Advanced evidence
Responsible use	Undisclosed, inappropriate or careless use of AI/data.	Discloses permitted AI use and follows rules.	Proactively identifies risk, applies controls and recommends safe improvement.
Reflection and transfer	Cannot explain the contribution or limitation of AI.	Describes what AI helped with and what was verified.	Articulates a reusable human-AI teaming strategy and degraded-mode response.

6. When AI supports feedback or scoring

AI can be useful for low-risk formative feedback, pattern detection, question generation and preliminary rubric alignment. Its use is most defensible when it accelerates a feedback loop that an instructor can review and improve. It is least defensible when it produces an opaque high-consequence outcome that a learner cannot challenge.

Use of AI	Recommended control
Practice questions and self-checks	Allow automated feedback but validate the question bank and include a route to report incorrect or misleading feedback.
Draft feedback	Use a transparent rubric, require instructor calibration samples and make clear that the feedback is preliminary.
Pattern analysis across a cohort	Use aggregated or minimised data where possible; avoid turning analytics into labels of potential, conduct or suitability.
Proposed scoring or classification	Keep a human reviewer responsible; verify accuracy, fairness and policy compliance; document exceptions and appeals.
High-consequence grading / progression	Do not rely on AI alone. Apply human review, clear evidence, explainability and a learner review route. [1, 4, 5]

7. Evaluation and quality assurance cycle

01	02	03	04	05
Govern the use case	**Map learning and risk context**	**Measure evidence**	**Manage findings**	**Decide next action**

Adapted from risk and assurance practice [1, 4]

Review question	Example indicator
Learning value	Quality of source checking, transfer performance, instructor observations, learner confidence calibrated against actual performance.
User experience	Learner clarity on AI role, access and support; usability barriers; quality of feedback; level of instructor intervention.
System quality	Rate and type of incorrect or unsupported output; source coverage; known failure patterns; configuration changes.
Data and security	Policy breaches, prohibited inputs, access anomalies, retention/control issues and response time.
Institutional decision	Evidence that supports continuing, revising, restricting, scaling or retiring the activity.

CHECKLIST | ASSESSMENT AND EVALUATION READINESS

[ ] The assessment task specifies the permitted AI use and the learner disclosure requirement.
[ ] The evidence includes reasoning, source use, verification and human judgement - not only a final product.
[ ] Any AI-supported feedback or scoring has a named human owner, validation approach and learner review route.
[ ] The programme captures both learner outcomes and AI/system assurance evidence.
[ ] Data used for evaluation is necessary, access-controlled and retained according to the applicable rules.
[ ] The use case has a decision point with criteria to continue, revise, restrict, scale or retire.

Note: *The level of assurance must be proportionate to the learning consequence, data sensitivity and degree of automation.*

8. Chapter takeaways

AI-enabled PME requires more transparent and evidence-rich assessment, not lower standards. The strongest designs assess how learners use and challenge AI, retain human responsibility for consequential decisions and evaluate the instructional intervention as well as the learner. Evidence should drive the next decision: retain, improve, restrict or scale.

References and editorial basis

[1] NATO. Summary of NATO's revised Artificial Intelligence (AI) strategy. 10 July 2024. Source
[2] NATO. Data Strategy for the Alliance. 5 May 2025. Source
[3] NATO. Data Quality Framework for the Alliance. 29 August 2025. Source
[4] National Institute of Standards and Technology (NIST). Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1. 2023. Source
[5] National Institute of Standards and Technology (NIST). Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, NIST AI 600-1. 2024. Source
[6] UNESCO. Guidance for Generative AI in Education and Research. 2023. Source
[7] RRIS. Benchmarks for Utilizing AI in Assessment and Evaluation. 2025. Source

Editorial note. *This proposal was developed with reference to working material supplied by project contributors. The working material informed the intended chapter format and practical emphasis; policy-facing and factual claims were rechecked against the public sources listed above. Local law, national policy, security classification rules and institutional approval procedures remain authoritative.*

ANNEX A | ASSESSMENT AND EVALUATION PLAN

Evidence plan for an AI-enabled PME activity

Use this plan to separate learner assessment from evaluation of the AI-supported activity. Define what will be measured, who reviews it, how learners are informed and what will cause a change to the design.

Plan element	Draft entry
Learning outcome and performance standard	
Assessment task and permitted AI use	
Learner disclosure requirement	
Evidence to be assessed	
Human review points and rubric	
Formative feedback method	
AI system / output validation method	
Data to retain and access controls	
Learning indicators	
Risk / assurance indicators	
Review date and decision owner	
Scale, revise or retire criterion	

MINIMUM CONDITION	No high-consequence educational decision should depend solely on unreviewed AI output. The assessment plan must identify the accountable human reviewer and the basis on which a learner can request clarification or review.

CHAPTER 8

Chapter 8 - Emerging Topics

1. Standard practice (Business as usual)

Artificial intelligence (AI) is rapidly reshaping warfare, security, defence institutions, and the broader information environment in which military professionals operate. As AI capabilities continue to evolve at high speed, Professional Military Education (PME) must move beyond viewing AI solely as a technical subject and instead address the broader human, cognitive, ethical, and organizational implications associated with increasingly AI-enabled environments.

This challenge also reinforces the importance of experiential and reflective learning approaches associated with educational thinkers such as John Dewey and David A. Kolb, where learning emerges through active experimentation, critical reflection, adaptation, and engagement with real-world problems. In rapidly evolving AI environments, learners cannot rely solely on static knowledge transfer; they must instead develop the capacity to explore unfamiliar systems, test assumptions, interpret uncertainty, and continuously adapt their understanding as technologies evolve.

For PME, this requires a flexible and future-oriented approach that combines understanding of emerging AI technologies with critical exploration of the tensions those technologies introduce. This includes issues related to authority, expertise, trust, accountability, human judgement, cognitive overload, and the changing relationship between humans and intelligent systems. It also requires educational approaches that encourage experimentation, adaptability, continuous learning, and informed scepticism, recognizing that AI developments evolve so quickly that any static curriculum risks becoming immediately outdated.

This chapter outlines key emerging AI technologies, associated human and institutional issues, and recommended areas of focus that can help form the foundation of future-oriented defence education and AI-enabled professional military learning.

2. Added Value (Why AI?)

Existing teaching or training methods

Current research on artificial intelligence curricula for military personnel across multiple nations and organizations indicate a broadly consistent set of teaching and training methods (*Working Document: AI Curriculum for Military Personnel, 2026*). These commonly include asynchronous online learning, instructor-led seminars, scenario-based exercises, practical laboratories, simulations, and reflective discussions. Increasingly, programmes are also adopting blended learning models that integrate foundational AI literacy with operationally focused exercises, experimentation activities, and interdisciplinary PME discussions.

AI education for military personnel are also focusing on strengthening responsible use, operational relevance, and resilience in the face of system vulnerabilities and rapidly evolving technologies. Programmes are beginning to align more closely with NATO standards and broader interoperability requirements. More advanced approaches are introducing AI-driven adaptive testing and performance analytics to personalize learning pathways and improve assessment quality. Across these developments, there is a clear trend toward blended, iterative learning journeys that combine foundational knowledge with scenario-based exercises, innovation methods such as design thinking and rapid prototyping, and an increasing emphasis on lifelong learning, continuous upskilling, and core competencies such as prompt literacy and informed escalation to human judgement when AI systems are uncertain or inappropriate to rely upon (*Working Document: AI Curriculum for Military Personnel, 2026*).

That said, the rapid pace and growing complexity of emerging AI technologies and associated issues are placing increasing pressure on educational approaches. This is driving a shift toward more iterative, exploratory, experiential, and continuously updated learning models.

Commonly used tools, approaches and/or processes

Commonly used approaches include adaptive learning platforms, AI-enabled simulations, scenario-generation tools, data visualization environments, and hands-on experimentation with machine learning models. Programmes also frequently incorporate case studies, red-teaming activities, peer discussion, ethical analysis exercises, and role-specific competency frameworks to support progressive AI literacy and operational understanding. (*Working Document: AI Curriculum for Military Personnel, 2026*)

Increasingly, these approaches must also create space for continuous exploration, experiential learning, and learner-led discovery, recognizing that emerging AI technologies and issues evolve too quickly for static tools, fixed examples, or traditional instructional processes alone to remain sufficient.

Role of the instructor

In AI-enabled learning environments, the instructor increasingly serves as a facilitator, guide, and critical interpreter rather than solely as a transmitter of knowledge. As emerging technologies evolve, learning becomes a collaborative process in which instructors, learners, and increasingly agentic AI systems all contribute to the exploration, generation, testing, and refinement of knowledge and understanding.

Within this environment, instructors must help contextualize emerging technologies and guide reflection on issues such as human judgement, authority, trust, accountability, and epistemic uncertainty. At the same time, learners must be treated as active contributors to the learning process, encouraged to experiment, question, adapt, and critically assess both AI-generated outputs and their own assumptions. The role of the instructor therefore becomes less about delivering fixed content and more about cultivating critical thinking as a collective and continuously evolving activity. (*AI in Professional Military Education: Emerging Topics, 2026*).

Challenges and opportunities

When developing programs on emerging AI technologies and issues, the central challenge is no longer simply keeping content current, but adapting to a learning environment in which knowledge creation increasingly occurs through teaming between instructors, learners, and AI systems. This shifts the traditional role of the educator away from acting primarily as the sole holder and transmitter of expertise toward facilitating exploration, guiding judgement, validating outputs, and helping learners critically navigate rapidly evolving AI-enabled environments.

At the same time, this teaming model creates significant opportunities: AI can reduce administrative and content-development burdens, support rapid scenario generation and adaptation, personalize learning pathways, and enable continuous curriculum refinement. The challenge for institutions is therefore not whether to incorporate AI into teaching and learning, but how to design programmes that effectively balance human judgement, professional expertise, experimentation, and AI-supported co-creation within an environment of constant technological change. (*AI in Professional Military Education: Emerging Topics, 2026*).

Impact on instructional design

Instructional design in this space must move beyond treating AI content as fixed knowledge and instead treat it as a continuously evolving learning environment. Designers need to structure learning experiences that deliberately integrate emerging AI technologies alongside the broader themes they generate—human judgement, authority, trust, control, and uncertainty—so that both technical understanding and critical reflection develop together. This requires building in opportunities for experimentation, hands-on exploration, and “live” engagement with new tools and developments, including encouraging learners to actively seek out and test what is emerging in real time rather than relying solely on curated material.

Because AI evolves on a daily basis, any curriculum on AI is, by definition, immediately partial and time-bound; instructional design must therefore prioritize adaptability, learner discovery, and iterative updating over static content delivery. Learners should be given structured permission—and expectation—to explore, compare, and critique new tools as they appear, while instructors and designers curate frameworks for reflection, not fixed answers. In this way, the curriculum becomes less about transmitting stable content and more about cultivating the habits of critical engagement needed to operate effectively in an environment where the technology itself is always changing and consequently, the issues around them.

3. Key Components

Emerging AI technologies and issues

Emerging artificial intelligence technologies are developing rapidly and are already reshaping how military organizations think, plan, operate, and learn. While the landscape is broad and constantly evolving, there are five key technology areas that are particularly significant for defence, security, and PME due to their combined impact on capability, decision-making, and risk. These have been grouped into he following categories: Foundation Models & Generative AI, Multimodal AI, Agentic AI, AI-Enabled Autonomous Systems, and Adversarial AI & Synthetic Media. (*AI in Professional Military Education: Emerging Topics, 2026; Emerging Topics - Brainstorm Sketch, 2026*).

i. Foundation Models & Generative AI

Core enabling systems behind most recent AI advances. They support text, image, code, and scenario generation, and are already reshaping planning, intelligence analysis, education, and decision-support.

ii. Multimodal AI

AI that can integrate and reason across multiple data types (text, imagery, video, sensor data, and geospatial inputs). This is particularly important for intelligence fusion, situational awareness, and operational analysis.

iii. Agentic AI (Autonomous AI Systems)

A fast-emerging shift from "AI that responds" to "AI that acts." These systems can plan, sequence tasks, and execute workflows with limited human input, raising major implications for command, control, and accountability.

iv. AI-Enabled Autonomous Systems (Physical & Digital)

Autonomous platforms in air, land, maritime, cyber, and space domains. These systems directly connect AI to operational effect, including sensing, movement, targeting support, and logistics.

v. Adversarial AI & Synthetic Media (Information & Cognitive Domain Threats)

Deep fakes, data poisoning, prompt injection, spoofing, and other techniques that manipulate AI systems or generate convincing synthetic information. It is increasingly central to information warfare and cognitive security.

AI Technologies and Emerging Human & Institutional Issues

Emerging AI technologies are not only changing military capabilities; they are also reshaping how military professionals think, decide, learn, trust, lead, and exercise authority. Understanding AI in PME therefore requires examining both the emerging technologies as well as the emerging issues they generate, including challenges related to authority shift, trust recalibration, and cognitive pressure.

The following combinations of emerging AI technologies and emerging issues are recommended as integrated areas of exploration and discussion with learners in both the military and security contexts (*AI in Professional Military Education: Emerging Topics, 2026; Emerging Topics - Brainstorm Sketch, 2026*).

i. Foundation Models & Generative AI

→ Emerging issues: Authority, expertise, and hybrid cognition

- Changing nature of expertise (knowing vs interpreting vs prompting)
- Hybrid cognition (human + AI co-production of analysis, plans, and learning)
- Risk of over-reliance on generated content in professional judgement
- Increasing difficulty distinguishing authoritative knowledge from AI-generated content ("AI slop") and synthetic expertise
- Dual-use risks and adversarial exploitation of generative systems in military and civilian contexts

Core tension: Who is the author of thought—the human, the model, or both?

ii. Multimodal AI

→ Emerging issues: perception, cognitive overload, and shared understanding

- Cognitive overload from fused, high-volume information streams
- Decision compression under richer but more complex data environments
- Trust calibration across mixed-source intelligence (text, imagery, sensor fusion)
- Compression of decision-making timelines as AI integrates text, imagery, audio, sensor, and geospatial data simultaneously
- Maintaining shared situational understanding in complex operational environments

Core tension: More perception does not automatically mean more clarity.

iii. Agentic AI (Autonomous AI Systems)

→ Emerging issues: control, accountability, and decision authority

- Delegated operational planning and workflow with reduced human intervention
- Emergence of multi-agent ecosystems where AI systems coordinate, debate, and execute workflows beyond direct human sequencing
- Responsibility dilution in decision-making chains
- Erosion or reshaping of command authority
- Trust in systems that *act*, not just advise

Core tension: If AI acts, who is responsible for the action?

iv. AI-Enabled Autonomous Systems (Physical & Digital)

→ Emerging issues: operational risk, dependency, and escalation control

- Increased speed of operations vs human control of escalation
- Growing reliance on offline, edge, and AI-enabled field systems in contested or disconnected operational environments
- Strategic and operational dependency on autonomous infrastructure
- Fragility of systems that are highly networked and AI-dependent
- Human-machine coordination under time-critical conditions

Core tension: Speed vs control in operational environments.

v. Adversarial AI & Synthetic Media

→ Emerging issues: truth, trust, and epistemic stability

Breakdown of shared reality (deep fakes, synthetic intelligence products)
Degradation of trusted knowledge ecosystems as synthetic content increasingly overwhelms validated information sources
Information warfare operating at machine scale
Difficulty distinguishing authentic vs synthetic inputs creating erosion of trust
AI-enabled influence operations affecting deterrence, strategic competition, and cognitive security

Core tension: If anything can be faked, what can be trusted?

3. Key Areas / Components

While the emerging technologies and issues discussed throughout this chapter are diverse, patterns do appear. These broader themes help frame the deeper human, operational, and cognitive implications of AI integration and provide useful focal points for PME discussion and analysis.

- **Authority Shift:** Humans move from primary decision-makers toward acting as judgement arbiters of AI-generated outputs
- **Trust Recalibration:** Trust shifts from believing systems to continuously managing uncertainty, reliability, and ambiguity within systems
- **Cognitive Pressure:** The challenge moves from analyzing limited information toward managing overwhelming, high-speed information ecosystems

The following table outlines these meta-themes associated with AI-enabled environments, along with recommended technology areas, related human and institutional issues, and the core tensions these developments introduce for military professionals and organizations.

Table 1: Themes and Tensions in Emerging Technologies and Issues in AI Enabled Environments

Theme	Emerging Technologies	Emerging Issues	Tension
Authority Shift	Foundation Models and Generative AI	Authority, expertise, and hybrid cognition	Who is the author of thought = the human, the model, or both?
	Agentic AI (Autonomous Systems)	Control, accountability, and decision authority	If AI acts, who is responsible for the action?
Trust Recalibration	Adversarial AI & Synthetic Media	Truth, trust, and epistemic stability	If anything can be faked, what can be trusted?
Cognitive Pressure	Multimodal AI	Perception, cognitive overload, and shared understanding	More perception does not automatically mean more clarity.
	AI-Enabled Autonomous Systems (Physical & Digital)	Operational risk, dependancy, and escalation control	Speed vs control in operational environments.

Taken together, these themes reinforce that the future challenge of AI in military and PME contexts is not solely technological, but increasingly centred on how humans.

4. Sample Learning Activity

The following experiential learning activity is provided as an example to illustrate how emerging AI technologies can be paired with associated human, operational, and institutional issues emerging within PME. Similar learning activities can be developed for the other recommended technology-and-theme combinations to support critical thinking across a range of emerging AI topics.

Example Learning Activity: Foundation Models, Generative AI, and Human Judgement

Context

This learning activity explores how foundation models and generative AI are reshaping military analysis, planning, learning, and decision-support. It focuses not only on the capabilities of generative AI systems, but also on the emerging human questions they create around authority, expertise, judgement, and the relationship between human and machine-generated thought.

Topics

- Foundation models and generative AI
- AI-generated analysis and decision-support
- Hybrid cognition (human + AI collaboration)
- Risks of over-reliance on AI-generated outputs
- Distinguishing authoritative knowledge from "AI slop"
- Critical thinking and professional judgement in AI-enabled environments

Description of the Learning Activity

Learners explore how generative AI systems can produce content that increasingly resembles or rivals human-generated work. Through experimentation with AI tools, guided discussion, and comparative analysis activities, learners examine how AI may reshape concepts of expertise, authority, and professional judgement within military contexts.

The activity encourages reflection on the growing role of hybrid cognition, where humans and AI systems increasingly co-produce analysis, learning, and decision-support outputs. Learners assess both the opportunities and risks associated with AI-supported thinking, including the potential for over-reliance, reduced critical analysis, and misplaced trust in confidently generated but potentially flawed outputs.

A central question throughout the activity is: Who is the author of thought—the human, the model, or both?

Learning Objectives

Upon completion of this activity, learners should be able to:

- Explain how foundation models and generative AI systems function and are applied in military and educational contexts
- Assess how generative AI may influence human expertise, authority, and professional judgement, including their own
- Critically evaluate the strengths, limitations, and risks of AI-generated outputs in analysis and decision-making environments

How AI is Used

AI is used as a subject of study, an instructional tool, and a learning partner. Activities may include:

1. Sourcing new and emergent AI technologies (what's new and now)
2. Exploring hybrid cognition (human + AI collaboration)
3. Experimenting with AI-human teaming (instructor, learner and agentic AI) for prompt design and iterative refinement
4. Comparing human versus AI-generated content and analysis
5. Critically evaluating AI-generated recommendations for bias, accuracy, reliability, and contextual limitations
6. Reflecting on the changing nature of expertise and authorship in AI-enabled environments
7. Conducting discussions on trust, accountability, and professional responsibility

The activity emphasizes learning-by-doing, exploration, critical thinking, informed scepticism, and reflection, while maintaining strong human judgement when working with AI-generated content.

Expected Outcomes

Upon completion of this activity, learners should be able to:

- Explain the operational relevance and limitations of foundation models and generative AI
- Apply critical judgement when interpreting and using AI-generated outputs
- Reflect on how AI is reshaping expertise, authority, authorship, and decision-making in military environments, including their own

5. Examples of AI Models and Tools

It is recommended that instructional designers and educators draw on currently available AI models and tools to ensure operational grounding, and alignment with what is in active use today. At the same time, learners should be encouraged and given space to explore emerging tools and capabilities as they appear, recognizing that in a rapidly evolving field, staying current depends on continuous experimentation and engagement with what is new.

6. Step-by-Step Implementation (Program & Instructional Design Focus)

Preparation

A key risk at this stage is adopting generic AI content without anchoring it in the organization's current priorities and operational realities. Program and instructional designers should begin by deliberately pairing emerging AI technologies with the human, operational, and institutional issues they generate, but the starting point must be shaped by organizational need and context. Each institution should first ask: "what matters most for us right now—operational effectiveness, education and training, governance, trust, readiness, or something else?" That decision should guide which technology–issue pairings are prioritized first.

Initial Use (Pilot)

Pilot activities should be built around selected priority pairings, chosen based on institutional needs identified in the preparation phase. Instructional designers should create structured experiments where learners engage with AI tools while explicitly exploring the human and operational tensions those tools generate.

The emphasis should be on exploration and surfacing friction points between capability and consequence, ensuring learners critically examine how AI changes judgement, not just performance.

Integration into Regular Teaching

Instructional design should ensure AI is treated not as a standalone topic, but as a shaping factor in decision-making environments.

At this stage, institutions should continuously revisit their initial priorities and adjust pairings as operational needs evolve. Curriculum design should consistently explore tension and associate challenges. This ensures teaching remains aligned with both current organizational priorities and the evolving AI landscape.

Further Development

Over time, programmes should evolve into a dynamic system of continuous alignment between emerging technologies, emerging issues, and institutional priorities. Instructional designers should regularly reassess: "what has changed in our environment, and what now matters most?"—and update learning design accordingly.

This includes integrating learner-led horizon scanning, where participants identify emerging tools and risks and feed insights back into curriculum updates. The goal is to create a responsive learning ecosystem where content is continuously co-created with learners, and re-anchored to real organizational needs.

Ultimately, the programme shifts from fixed curriculum delivery to an adaptive design model that connects evolving AI technologies with the most pressing human, operational, and institutional challenges faced by the organization.

7. Risks and Security

Adopting emerging AI technologies and more experimental learning approaches introduces real risks for military institutions, including security vulnerabilities, exposure to adversarial manipulation, reduced control over information environments, and institutional discomfort with learner-driven or co-created curriculum models. Continuously evolving curriculum also creates ongoing pressure on instructors, governance structures, and institutional resources.

At the same time, thoughtful and critical collaboration with AI systems offers significant opportunities to extend human capability, accelerate learning, enhance analysis, and support more adaptive and responsive forms of military education. When approached with strong human judgement, ethical oversight, and informed scepticism, human–AI collaboration can become a force multiplier for learning, experimentation, creativity, and decision support rather than a replacement for professional expertise.

However, the risks of failing to adapt may be even greater: outdated curriculum, obsolete technological understanding, over-reliance on static teaching models, and military professionals who are insufficiently prepared to think critically, adapt rapidly, and operate effectively within increasingly AI-enabled environments.

8. Recommendations

For Instructors

Instructors should treat AI not as a fixed subject, but as a moving learning environment that requires ongoing experimentation and adaptive teaching practice. The priority is to create space for learners to actively engage with tools, test ideas, and develop judgement in uncertain and changing conditions.

Instructors must shift away from the traditional “sage on the stage” and “guide on the side” models. Instead, instructors must become active learning partners who explore, question, experiment, and adapt alongside learners.

Practical steps include:

- Integrate hands-on experimentation with AI tools into every relevant topic, even in simple form
- Encourage learners to compare and challenge AI outputs rather than accept them
- Build short, repeatable activities that allow learners to explore “what’s new this week” in AI developments
- Explicitly connect technical content to human implications (judgement, trust, authority, ethics)
- Allow room for open-ended inquiry and learner-led exploration, including tools not predefined in the lesson plan
- Model informed scepticism and adaptive thinking, rather than fixed answers or static expertise
- Shift toward the role of active learning partner

For Institutions

Institutions should support AI education as a dynamic capability rather than a static curriculum, recognizing that agility, experimentation, and continuous updating are essential in a rapidly evolving field.

Organizational support should include:

- Policies that explicitly permit and encourage safe experimentation with emerging AI tools
- Direction that fosters a culture of human+AI co-creation
- Flexible curriculum frameworks that prioritize learning outcomes over fixed content
- Access to approved tools and environments that enable experiential learning
- Ongoing staff development in AI literacy, facilitation, and adaptive instructional design
- Mechanisms for rapid curriculum refresh based on emerging technologies and operational relevance
- Encouragement of cross-functional collaboration and interdisciplinary learning approaches

Importantly, institutions should normalize the idea that AI education cannot remain static. Instead, programmes should be designed to evolve continuously, with instructors and learners jointly contributing to the discovery of emerging tools, risks, and opportunities.

Overall Principle

Across both instructor practice and institutional design, the guiding principle is that AI education in PME must remain experimental, adaptive, experiential, and forward-leaning. In a rapidly changing environment, effectiveness depends less on delivering fixed knowledge and more on cultivating professionals who can think dynamically and engage critically.

Summary

Artificial intelligence is reshaping warfare, security, military institutions, and the environments in which military professionals operate. As AI capabilities continue to evolve, PME must move beyond technical literacy alone. The central challenge is no longer simply understanding AI technologies, but ensuring that military professionals can maintain judgement, trust, accountability, critical thinking, and effective decision-making within rapidly evolving AI-enabled contexts.

For military institutions, educators, instructional designers, and leaders, the imperative is clear: begin now, experiment actively, continuously adapt, and deliberately prepare learners not only to use AI technologies, but to lead responsibly and think critically within increasingly AI-enabled environments.

9. References

Our documents

Emerging Topics: Brainstorm Sketch, 2026

AI in Professional Military Education (PME): Emerging Topics, 2026

Working Document: AI Curriculum for Military Personnel, 2026

Core Literature

G. Allen and T. Chan, Artificial Intelligence and National Security. Cambridge, MA, USA: Belfer Center for Science and International Affairs, Harvard Kennedy School, Jul. 2017.

M. C. Horowitz, S. Pindyck, and C. Mahoney, Strategic Competition in the Age of AI. Santa Monica, CA, USA: RAND Corporation, 2022.

Jensen, Y. Atalan, and J. M. Macias III, Algorithmic Stability: How AI Could Shape the Future of Deterrence. Washington, DC, USA: Center for Strategic and International Studies, 2024.

Y. H. Wong et al., Deterrence in the Age of Thinking Machines. Santa Monica, CA, USA: RAND Corporation, 2020.

S. Hoadley and N. J. Lucas, Artificial Intelligence and National Security. Washington, DC, USA: Congressional Research Service, 2020.

Selected Articles and Reports

J. Pruet, A. Makanju, J. Reiber, and J. Achiam, AI and International Security: Pathways of Impact and Key Uncertainties. San Francisco, CA, USA: OpenAI, Feb. 2026.

Center for Security and Emerging Technology, What Are Generative AI, Large Language Models, and Foundation Models?. Washington, DC, USA: Georgetown University, 2023.

H. Khlaaf, S. M. West, and M. Whittaker, “Mind the Gap: Foundation Models and the Covert Proliferation of Military Intelligence, Surveillance, and Targeting,” arXiv preprint arXiv:2410.14831, 2024.

L. Floridi and M. Taddeo, “Ethical Principles for Artificial Intelligence in National Defence,” Philosophy & Technology, vol. 34, no. 4, pp. 1–15, 2021.

Document source = Working document: AI Curriculum for Military Personnel

[1] “Website.” Available: https://apps.dtic.mil/sti/trecms/pdf/AD1122946.pdf

[2] “Artificial Intelligence for Military Use Certificate - Curriculum 128 (DL), Curriculum 129 (RES).” Available: https://nps.smartcatalogiq.com/en/current/academic-catalog/departments/department-of-computer-science/artificial-intelligence-for-military-use-certificate-curriculum-128-dl-curriculum-129-res. [Accessed: Mar. 25, 2026]

[3] “Website.” Available: https://www.act.nato.int/activities/nato-education-and-training/

[4] M. M. Ortmann, “Artificial Intelligence in combat simulations: How AI is changing NATO and Czech Army soldier training,” Aug. 24, 2024. Available: https://www.czdefence.cz/clanek/umela-inteligence-v-bojovych-simulacich-jak-ai-meni-vycvik-vojaku-nato-a-acr. [Accessed: Mar. 25, 2026]

[5] A. Bestyuk and S. Pokhnatiuk, “Integration of artificial intelligence into higher military education as a factor in increasing the efficiency of professional training,” *Sci. Bull. Mukachevo State Univ. Ser. “Pedag. Psychol.*, pp. 60–71, Jun. 2025, doi: 10.52534/msu-pp2.2025.60. Available: https://pp-msu.com.ua/en/journals/tom-11-2-2025/integratsiya-shtuchnogo-intelektu-u-vishchu-viyskovu-osvitu-yak-faktor-pidvishchennya-efektivnosti-profesiynoyi-pidgotovki

[6] “Artificial Intelligence for Military Use Certificate - Curriculum 128 (DL), Curriculum 129 (RES).” Available: https://nps.smartcatalogiq.com/en/current/academic-catalog/departments/department-of-computer-science/artificial-intelligence-for-military-use-certificate-curriculum-128-dl-curriculum-129-res. [Accessed: Mar. 25, 2026]

[7] “Educating the AI-Ready Warfighter: A Framework for Ethical Integration in Air Force Profes,” *Air University (AU)*. Available: https://www.airuniversity.af.edu/Wild-Blue-Yonder/Articles/Article-Display/Article/4219340/educating-the-ai-ready-warfighter-a-framework-for-ethical-integration-in-air-fo/. [Accessed: Mar. 25, 2026]

[8] N. M& and S. CoE, “NATO M&S COE Courses & Training Activities in 2024 - NATO Modelling & Simulation Centre of Excellence,” *NATO Modelling & Simulation Centre of Excellence*, Jan. 10, 2024. Available: https://www.mscoe.org/nato-ms-coe-courses-training-activities-in-2024/. [Accessed: Mar. 25, 2026]

[9] “[No title].” Available: https://pp-msu.com.ua/web/uploads/pdf/Scientific%20Bulletin%20of%20Mukachevo%20State%20University.%20Series%20%E2%80%9CPedagogy%20and%20Psychology%E2%80%9D_11_2_2025-60-71.pdf. [Accessed: Mar. 25, 2026]

[10] N. C. Rowe, “The comparative ethics of artificial-intelligence methods for military applications,” *Front. Big Data*, vol. 5, p. 991759, Sep. 2022, doi: 10.3389/fdata.2022.991759. Available: https://public-pages-files-2025.frontiersin.org/journals/big-data/articles/10.3389/fdata.2022.991759/pdf. [Accessed: Mar. 25, 2026]

[11] “Artificial Intelligence for Military Use Certificate - Curriculum 128 (DL), Curriculum 129 (RES).” Available: https://nps.smartcatalogiq.com/en/current/academic-catalog/departments/department-of-computer-science/artificial-intelligence-for-military-use-certificate-curriculum-128-dl-curriculum-129-res. [Accessed: Mar. 25, 2026]

[12] “[No title].” Available: https://pp-msu.com.ua/web/uploads/pdf/Scientific%20Bulletin%20of%20Mukachevo%20State%20University.%20Series%20%E2%80%9CPedagogy%20and%20Psychology%E2%80%9D_11_2_2025-60-71.pdf. [Accessed: Mar. 25, 2026]

[13] M. R. R. Gaikwad and D. A. S. Choudhary, “INTEGRATING ARTIFICIAL INTELLIGENCE INTO MILITARY EDUCATION: OPPORTUNITIES, CHALLENGES, AND FUTURE DIRECTIONS,” *jetjournal*, vol. 15, no. 1, pp. 1317–1330, Dec. 2025, doi: 10.7492/bgfnbb75. Available: https://jetjournal.us/index.php/journals/article/view/756. [Accessed: Mar. 25, 2026]

[14] “Summary of the NATO Artificial Intelligence Strategy,” *Official texts and resources | NATO*. Available: https://www.nato.int/en/about-us/official-texts-

and-resources/official-texts/2021/10/22/summary-of-the-nato-artificial-intelligence-strategy. [Accessed: Mar. 25, 2026]

[15] "Data Quality Framework for the Alliance," *Official texts and resources | NATO*. Available: https://www.nato.int/en/about-us/official-texts-and-resources/official-texts/2025/08/29/data-quality-framework-for-the-alliance. [Accessed: Mar. 25, 2026]

CHAPTER 9

Conclusions and Implementation Guidance

A staged, evidence-led approach for turning AI ambition into a sustainable PME capability that remains human-centred, secure, interoperable and mission-relevant.

PURPOSE	Conclude the guide with a practical operating model and staged roadmap. The chapter helps institutions move from isolated experimentation to a governed portfolio of AI-enabled learning capabilities.

1. The implementation challenge

The challenge for PME is not whether AI will affect professional learning. It already does. The challenge is whether institutions will respond through controlled capability building or through fragmented, ungoverned adoption. A credible implementation model connects strategy, people, learning design, data governance, technical assurance and evaluation. It does not treat any one of these as optional.

NATO's revised AI Strategy calls for responsible adoption, AI-ready workforce development, interoperability, quality data and testing/evaluation. NATO's Digital Transformation Implementation Strategy frames transformation across people, processes and technology, and its 2026 update highlights a rolling-horizon approach that combines execution with adaptation. NATO's Data Strategy adds a data-centred approach built on governance, data value, security, analytical maturity and a data-driven ecosystem. **[1, 2, 3, 4]**

CONCLUSION	Successful AI implementation in PME is a managed change programme. It is not a procurement exercise, a staff workshop or a collection of prompts.

2. Nine implementation principles

le	What it means in practice
1. Mission before tool	Prioritise use cases that improve a defined learning, readiness or institutional problem. Make the non-AI alternative visible.
2. Human accountability	Name an accountable owner for each use case and define the points where human judgement, review and escalation are mandatory.
3. Data by design	Classify data, establish permitted/prohibited inputs, manage access and retain records according to approved rules. **[4, 5]**
4. Role-based capability	Develop staff and learner competence by role; do not expect one generic course to create an AI-ready institution.
5. Evidence before scale	Use measured pilots with learning, experience, system and risk evidence; scale only when criteria are met.
6. Interoperable patterns	Reuse shared templates, source practices, assessment patterns, terminology and governance artefacts across programmes.
7. Resilient learning	Maintain alternatives for tool failure, connectivity loss, model withdrawal or restrictive policy change.

le	What it means in practice
8. Continuous review	Version content, prompts, sources, policies and controls; review after material changes or incidents.
9. Partnership and learning	Create communities of practice among faculty, SMEs, technical owners and partner institutions t share lessons and avoid duplicated mistakes.

3. Staged roadmap: first 180 days

Time horizon	Primary objective	Key actions	Decision output
Days 0-30: Establish direction	Create the institutional mandate and baseline.	Name sponsor and core team; map existing AI use; agree data/security boundaries; select 1-3 priority learning problems; define governance forum.	Approved implementation charter and pilot selection criteria.
Days 31-90: Run controlled pilots	Generate evidence, not anecdote.	Design pilots using common templates; train instructors and owners; validate sources; disclose AI use to learners; collect learning and risk evidence.	Pilot review: retain, revise, restrict or stop.
Days 91-180: Standardise what works	Convert successful patterns into repeatable capability.	Publish guidance, session/assessment templates and staff development; establish use-case register; align support, procurement and records processes.	Operating procedures and first standardised services.
Beyond 180 days: Integrate and adapt	Manage a portfolio and continue learning.	Refresh pathways; extend to new programmes; measure maturity; conduct assurance reviews; collaborate across institutions; retire weak or redundant use cases.	Portfolio roadmap, maturity assessment and next investment decisions.

01	02	03	04	05
Set direction	**Pilot deliberately**	**Standardise proven patterns**	**Scale with assurance**	**Refresh continuously**

Roadmap logic

4. Minimum operating model

Role / forum	Core responsibility	Cadence
Senior steering group	Sets priorities, risk appetite, resourcing and scale decisions; resolves cross-functional barriers.	Quarterly or at key pilot gates.
PME AI implementation lead	Maintains roadmap, use-case portfolio, common templates, reporting and coordination across learning, technical and governance functions.	Continuous coordination; monthly dashboard.
Programme owners and instructors	Own learning design, learner experience, content relevance, disclosure and day-to-day feedback.	Per session / per course; review after delivery.
SME, data/security, legal and technical advisers	Validate content, data use, system configuration, assurance and local policy compliance.	At design gates and when material changes occur.
Quality assurance / evaluation function	Defines evidence standards, audits samples, tracks lessons and supports scale/revise/retire decisions.	At pilot close-out and regular portfolio review.
Community of practice	Shares examples, prompt/source patterns, corrective actions and professional development needs.	Monthly or bi-monthly.

5. Maturity model for AI-enabled PME

Stage	Institutional characteristics	Next step
0. Observe	Awareness is high but use is informal; no common policy, portfolio or evidence model.	Create baseline, boundaries and sponsor.
1. Pilot	Small approved trials; enthusiastic individuals; preliminary support and risk controls.	Use common charter and evaluation plan; capture lessons.
2. Standardise	Reusable patterns, staff development, approved environments and basic governance exist.	Integrate source/data controls, quality processes and cross-programme reporting.
3. Integrate	AI capability is embedded in curriculum, assessment, support and resource planning; assurance evidence informs decisions.	Manage interoperability, continuity and portfolio balance; partner externally.
4. Adapt and lead	Institution continuously refreshes practices, shares knowledge, tests emerging approaches and contributes to standards or partner learning.	Maintain horizon scanning, test new use cases and prevent complacency.

6. Measures that support good decisions

A dashboard should be balanced. Counting licences, logins or generated documents may indicate activity, but it does not establish educational value or responsible use. Use a small set of indicators across four categories:

Measure category	Illustrative indicators
Learning and readiness	Quality of decision rationale, source checking, transfer to related tasks, instructor observations and learner self-assessment calibrated against performance.
Instructional quality	Timeliness and usefulness of feedback, learner access, instructor workload patterns, repeatability of learning design and quality of debrief.
Assurance and data	Unsupported-output rate in validation samples, policy breaches, prohibited inputs, incidents, access issues and time to correct or suspend.
Institutional capability	Staff competence, number of governed use cases, time from pilot to review, reuse of common patterns, continuity coverage and partner learning.

The NIST AI RMF functions are useful for structuring this evidence: govern the operating model, map the use case and stakeholders, measure learning and risk evidence, and manage actions arising from findings. **[6, 7]**

7. Decision gates for scale-up

Scale-up question	Evidence required
Does the pilot solve the intended learning problem?	Evidence of improved practice, feedback quality, learner access, judgement or other stated outcome - not only positive reactions.
Is the use case operating within approved boundaries?	Data/security review, documented controls, validation sample, incident history and implemented corrections.
Can the institution sustain it?	Named owner, staff support, source and configuration maintenance, continuity plan, budget/resource implication and learner support route.
Can the pattern be reused safely?	Clear eligibility conditions, documentation, templates, evidence standard and training for new instructors or programmes.
What will trigger reconsideration?	A change in tool, source base, policy, data sensitivity, learner group, error pattern or operational context.

CHECKLIST | IMPLEMENTATION LAUNCH CHECKLIST

[ ] A senior sponsor, implementation lead, programme owner and appropriate data/security/technical advisers are identified.

[ ] The institution has selected a small, mission-relevant portfolio of pilots rather than allowing uncontrolled proliferation.

[ ] Each pilot has a use-case charter, data/tool boundary, learner disclosure, evidence plan and decision date.

[ ] Staff development covers the three instructor models, role-based AI skills, teaching methods, assessment and escalation routes.

[ ] A shared repository stores approved templates, source practices, decision records, lessons learned and current guidance.

[ ] The institution has agreed how successful pilots will be standardised and how unsuccessful pilots will be stopped or retired.

Note: *Use this checklist at the point a programme moves from individual experimentation to institutional implementation.*

8. Final implementation message

AI should make PME more adaptive, evidence-informed and capable of preparing professionals for AI-enabled environments. It will do so only when institutions invest equally in people, learning design, data governance, assurance and reflection. Begin with a bounded use case, retain human responsibility, learn from evidence and build the capability iteratively. The end state is not AI everywhere; it is deliberate, defensible and useful AI where it strengthens professional military education.

References and editorial basis

[1] NATO. Summary of NATO's revised Artificial Intelligence (AI) strategy. 10 July 2024. Source

[2] NATO. NATO's Digital Transformation Implementation Strategy. 17 October 2024. Source

[3] NATO. NATO's Digital Transformation Implementation Strategy 2.0. 26 May 2026. Source

[4] NATO. Data Strategy for the Alliance. 5 May 2025. Source

[5] NATO. Data Quality Framework for the Alliance. 29 August 2025. Source

[6] National Institute of Standards and Technology (NIST). Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1. 2023. Source

[7] National Institute of Standards and Technology (NIST). Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, NIST AI 600-1. 2024. Source

[8] UNESCO. Guidance for Generative AI in Education and Research. 2023. Source

Editorial note. *This proposal was developed with reference to working material supplied by project contributors. The working material informed the intended chapter format and practical emphasis; policy-facing and factual claims were rechecked against the public sources listed above. Local law, national policy, security classification rules and institutional approval procedures remain authoritative.*

ANNEX A | IMPLEMENTATION ROADMAP

PME AI implementation roadmap canvas

Use this one-page canvas to coordinate a responsible rollout. It is intended for an implementation team, not for an individual instructor adopting a single tool informally.

Roadmap field	Draft entry
Strategic objective / mission need	
Priority learner groups and programmes	
First 1-3 approved pilots	
Senior sponsor and accountable owners	
Data, security, legal and procurement conditions	
Staff development and support model	
Learning, quality and assurance metrics	
Governance forum and review cadence	
Decision gates: pilot / standardise / scale	
Resourcing and continuity requirements	
Key dependencies and risks	
90-day decision and next action	

IMPLEMENTATION STANDARD	Every scaled activity should be traceable to a mission or learning need, an accountable owner, a data and security decision, defined evidence and an explicit decision to continue. A portfolio of small, evidenced pilots is preferable to uncontrolled proliferation.