


Launchpoint

AI & Engineering Learning + Training Curriculum

Learn the science. Build the skill.
Build what comes next.

12 common-core courses · 30 specialist plans
6 focused pathways · One substantial capstone

Proposed January 2027 cohort · Paid professional training
Review edition: programme, staffing and institutional agreement pending.



A CLEAR PROPOSITION

Learning gives the principle. Training builds the capability.

Launchpoint connects guided concepts with supervised practical work, technical reviews and an individual evidence portfolio. It is a proposed professional-training programme, not an employment or stipend offer.

Common core

Every learner follows 12 coded courses over 16 weeks. Productive learning and training total 373h20m within 400 scheduled attendance hours.

Focused extension

The optional four-week extension selects one pathway: two 24-hour specialist studios plus 45h20m of project deepening. This totals 93h20m productive time within 100 scheduled hours.

A larger catalogue

The 30 specialist plans describe the wider programme architecture. They are not all included, staffed or scheduled in every intake.

Evidence of contribution

Code, experiments, technical reports, demonstrations and an individual defence are proposed assessment evidence. A certificate is separate from browser practice progress.

Status: course plans and lab briefs authored; complete taught packs across the catalogue are in preparation. Four smaller original practice modules with executed notebooks are available in the local learning campus.

THE ATTENDANCE MODEL

A timetable that adds up.

One proposed weekday shift per learner: 08:30–13:30 or 15:00–20:00. Final venue, dates, faculty, capacity and recognition must be confirmed before enrolment.

400h

Core scheduled attendance: 16 weeks × 5 days × 5 hours

26h40m

Core breaks: 20 minutes per scheduled day

373h20m

Productive core allocation after breaks

149h20m

Guided learning: 40% of productive core time

224h

Practical training and evidence: 60% of productive core time

These are design allocations, not hours already delivered. Core capstone time is distributed through the programme, with concentrated integration and defence in weeks 14–16.

COURSE CODES / THE SHARED FOUNDATION

Twelve courses. One connected engineering journey.

AI101	Engineering foundations: Python, Linux & Git	40h
MTH110	Mathematics for learning systems	24h
DAT120	Data science: SQL, EDA & statistical reasoning	32h
ML201	Machine learning: baselines to reliable decisions	40h
DL301	Deep learning: DNNs, PyTorch, CNNs & sequences	32h
NLP310	Language engineering: text to transformers	24h
AI320	Generative AI, RAG & bounded agentic workflows	24h
MLO401	Software delivery for AI: APIs, tests, Docker & CI	32h
MLO410	MLOps: versioning, experiments, pipelines & monitoring	24h
RES420	Scientific modelling, digital twins & research practice	16h
PRO430	Engineering communication & career evidence	12h
CAP490	Industry capstone & individual technical defence	73h 20m

Each course page that follows specifies units, topics, prerequisites, a supervised training brief, required evidence and a defence question.

THE OPTIONAL FOUR-WEEK EXTENSION

Choose a direction. Deepen a contribution.

Each pathway selects two specialist studios: 48 hours. The remaining 45h20m deepens the project. Specialist availability requires confirmed faculty and a suitable environment.

01 / Applied ML Engineer

ML210 + ML220

A reproducible decision service with feature contracts and documented errors.

02 / Generative AI & Agent Engineer

AI340 + AI350

A source-grounded assistant with evaluated retrieval, bounded tools and approval gates.

03 / Deep Learning & Vision Engineer

CV320 + DL310

A visual-inspection prototype with held-out scene evaluation and a profiled model.

04 / Robotics & Autonomous Systems

ROB301 + ROB310

A simulated ROS 2 perception/localisation workflow with replay and transform evidence.

05 / MLOps & AI Platform Engineer

MLO420 + MLO430

A versioned model service with release controls, monitoring and a recovery drill.

06 / Scientific AI & Research

RES430 + RES440

A reduced-order twin and reproducible research package with carefully bounded claims.

THE PROPOSED ASSESSMENT

Know. Do. Build. Defend.

Final academic criteria must be adopted before enrolment. Institutional credit depends on written college recognition, not on this curriculum document alone.

40% / Course labs and reviews

Technical correctness, reproducibility, interpretation, tests and individual evidence.

40% / Capstone system

Problem scope, baseline, implementation, validation, limitations and handover.

20% / Individual defence

Explain design choices, failure cases, evaluation and your own contribution.

Mandatory gates

Proposed 80% attendance, 60/100 overall, at least 50% in individual defence, all required evidence and integrity/data-safety clearance.

Practice is separate

Browser quizzes and locally saved progress are personal study tools, not official assessment, proctoring or credential issuance.

What support means

Career preparation, research submission support and postgraduate guidance cannot guarantee placement, publication acceptance, indexing or admission.

AI101 / ENGINEERING FOUNDATIONS

Engineering foundations: Python, Linux & Git

40h productive allocation · 16h learning + 24h training

Prerequisites: Basic computing; readiness bridge as needed

UNIT 01 / 10h

The developer workstation

- Shell navigation
- paths and permissions
- virtual environments
- dependency records

UNIT 02 / 10h

Python as engineering

- Functions and modules
- exceptions
- typing
- debugging
- files and structured data

UNIT 03 / 10h

Git and team delivery

- Working tree, index and commits
- branching
- merge conflicts
- pull-request review

UNIT 04 / 10h

A reproducible first repository

- README and issue tracker
- minimal tests
- clean-clone run
- secret hygiene

THE TRAINING STUDIO

Create a small CSV-audit command-line tool, develop it on a feature branch, review a change and demonstrate a fresh-clone run.

EVIDENCE TO PRODUCE

- Versioned repository
- environment record
- passing tests
- reviewed change and run instructions

DEFEND THE REASONING

Explain the difference between a commit, a branch, an environment and a deployed application.

MTH110 / MATHEMATICS & DATA

Mathematics for learning systems

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: AI101

UNIT 01 / 6h

Linear algebra

- Vectors and matrices
- dot products
- projections
- norms and rank

UNIT 02 / 6h

Calculus and optimisation

- Derivatives
- chain rule
- gradients
- learning-rate behaviour

UNIT 03 / 6h

Probability foundations

- Conditional probability
- Bayes rule
- random variables
- expectation

UNIT 04 / 6h

Numerical reasoning

- Floating-point error
- conditioning
- scaling
- finite-difference gradient checks

THE TRAINING STUDIO

Implement a least-squares fit and a gradient-descent update on synthetic data; compare numerical and analytic gradients.

EVIDENCE TO PRODUCE

- Derivation sheet
- gradient check
- convergence plot
- units and assumptions

DEFEND THE REASONING

Why can a mathematically correct update still be unstable in a numerical implementation?

DAT120 / MATHEMATICS & DATA

Data science: SQL, EDA & statistical reasoning

32h productive allocation · 12h 48m learning + 19h 12m training

Prerequisites: AI101, MTH110

UNIT 01 / 8h

Data structures and querying

- NumPy arrays
- pandas tables
- SQL joins
- aggregations and nulls

UNIT 02 / 8h

Exploratory investigation

- Missingness
- distributions
- visual reasoning
- units and provenance

UNIT 03 / 8h

Statistical thinking

- Sampling
- uncertainty
- confidence intervals
- hypothesis-test limits

UNIT 04 / 8h

A defensible data contract

- Schema and validation
- grouped and temporal splits
- leakage audit
- data dictionary

THE TRAINING STUDIO

Audit a synthetic operational dataset, answer five questions with SQL and build an EDA notebook that links each conclusion to evidence.

EVIDENCE TO PRODUCE

- Data dictionary
- SQL query file
- EDA notebook
- validation report

DEFEND THE REASONING

What changes when rows are repeated measurements of the same machine rather than independent samples?

ML201 / MACHINE LEARNING

Machine learning: baselines to reliable decisions

40h productive allocation · 16h learning + 24h training

Prerequisites: DAT120

UNIT 01 / 10h

Supervised learning

- Linear and logistic regression
- KNN
- SVM
- decision trees

UNIT 02 / 10h

Ensembles and features

- Random forests
- gradient boosting and XGBoost
- encoding
- regularisation

UNIT 03 / 10h

Unsupervised learning

- K-means
- PCA
- anomaly scoring
- cluster interpretation

UNIT 04 / 10h

Evaluation and selection

- Cross-validation
- nested preprocessing
- tuning budgets
- calibration and error costs

THE TRAINING STUDIO

Compare a dummy baseline and three models on one documented split, select using validation only and explain final test errors.

EVIDENCE TO PRODUCE

- Benchmark repository
- error analysis
- model card
- held-out evaluation

DEFEND THE REASONING

Why is the highest test score not a valid model-selection procedure?

DL301 / DEEP LEARNING & PERCEPTION

Deep learning: DNNs, PyTorch, CNNs & sequences

32h productive allocation · 12h 48m learning + 19h 12m training

Prerequisites: MTH110, ML201

UNIT 01 / 8h

Neural learning

- Dense neural networks
- activations
- loss functions
- backpropagation

UNIT 02 / 8h

PyTorch training systems

- Tensors and autograd
- dataloaders
- training and evaluation modes
- checkpointing

UNIT 03 / 8h

Spatial representations

- Convolutions
- pooling
- transfer learning
- augmentation and leakage

UNIT 04 / 8h

Temporal representations

- RNN and LSTM
- sequence windows
- chronological validation
- learning curves

THE TRAINING STUDIO

Train a small CPU-friendly dense network and one transfer-learning or sequence baseline, then diagnose overfitting using held-out curves.

EVIDENCE TO PRODUCE

- Training script
- loss curves
- checkpoint manifest
- comparison with simpler baseline

DEFEND THE REASONING

How do architecture, optimisation algorithm and library differ?

NLP310 / LANGUAGE & GENERATIVE AI

Language engineering: text to transformers

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: DL301

UNIT 01 / 6h

Text and language data

- Unicode and tokenisation
- text cleaning
- label quality
- document splits

UNIT 02 / 6h

Representations

- Bag-of-words
- TF-IDF
- embeddings
- semantic similarity

UNIT 03 / 6h

Attention and transformers

- Queries, keys and values
- positional information
- encoder and decoder roles

UNIT 04 / 6h

Language evaluation

- Classification and retrieval metrics
- slice analysis
- context limits
- hallucination

THE TRAINING STUDIO

Compare a TF-IDF text classifier with a small pretrained representation and analyse failure cases without claiming language understanding from one score.

EVIDENCE TO PRODUCE

- Rights-cleared text corpus
- baseline comparison
- evaluation slices
- failure notebook

DEFEND THE REASONING

When is a lexical baseline a better starting point than a large language model?

AI320 / LANGUAGE & GENERATIVE AI

Generative AI, RAG & bounded agentic workflows

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: NLP310

UNIT 01 / 6h

Model use and limits

- Pretraining versus inference
- prompting
- sampling
- model selection and cost

UNIT 02 / 6h

Retrieval-augmented generation

- Chunking
- indexes
- retrieval ranking
- evidence and abstention

UNIT 03 / 6h

Agentic systems

- Tool contracts
- state and memory
- approval boundaries
- bounded execution

UNIT 04 / 6h

Evaluation before demonstration

- Task datasets
- source support
- prompt injection
- tracing and failure handling

THE TRAINING STUDIO

Build a small source-grounded assistant for an approved public corpus and a bounded tool workflow with explicit refusal and approval cases.

EVIDENCE TO PRODUCE

- Corpus and licence record
- retrieval benchmark
- evaluation report
- trace of a failed and a successful task

DEFEND THE REASONING

Why does connecting a retriever or tool not automatically make an answer trustworthy?

MLO401 / MLOPS & CLOUD

Software delivery for AI: APIs, tests, Docker & CI

32h productive allocation · 12h 48m learning + 19h 12m training

Prerequisites: AI101, ML201

UNIT 01 / 8h

Inference as an API

- HTTP and JSON
- FastAPI
- input validation
- batch versus online contracts

UNIT 02 / 8h

Tests that protect behaviour

- Unit and integration tests
- data-contract checks
- deterministic fixtures
- failure responses

UNIT 03 / 8h

Container engineering

- Images and containers
- Dockerfile
- volumes and networking
- non-root processes

UNIT 04 / 8h

Continuous integration

- Pull-request checks
- build artefacts
- dependency scanning
- release and rollback

THE TRAINING STUDIO

Package a baseline model behind a validated local API, test it, containerise it and prepare a CI workflow for a clean build.

EVIDENCE TO PRODUCE

- API contract
- test report
- Dockerfile and image recipe
- CI configuration

DEFEND THE REASONING

What must a clean machine have to reproduce your service without your notebook session?

MLO410 / MLOPS & CLOUD

MLOps: versioning, experiments, pipelines & monitoring

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: MLO401

UNIT 01 / 6h

Dataset and artefact lineage

- Git versus data bytes
- DVC pointers and remotes
- checksums
- recovery

UNIT 02 / 6h

Experiment discipline

- MLflow parameters, metrics and artefacts
- run comparison
- registration and approvals

UNIT 03 / 6h

Pipeline orchestration

- Explicit dependencies
- idempotent stages
- validation gates
- retries

UNIT 04 / 6h

Operational learning

- Latency and failure rates
- data and concept drift
- feedback delay
- rollback evidence

THE TRAINING STUDIO

Connect data validation, training, evaluation and packaging into a reproducible pipeline; record an experiment and rehearse recovery from a bad release.

EVIDENCE TO PRODUCE

- Versioned dataset manifest
- experiment ledger
- pipeline definition
- incident and recovery note

DEFEND THE REASONING

Why should an improved offline metric not automatically promote a model to production?

RES420 / RESEARCH & PROFESSIONAL PRACTICE

Scientific modelling, digital twins & research practice

16h productive allocation · 6h 24m learning + 9h 36m training

Prerequisites: MTH110, ML201

UNIT 01 / 4h

From system to model

- Governing equations
- boundary conditions
- state and parameters
- observability

UNIT 02 / 4h

Calibration and validation

- System identification
- residual structure
- uncertainty
- held-out operating regimes

UNIT 03 / 4h

An evidence-based twin

- State estimation
- scenario simulation
- measured versus synthetic inputs
- model limits

UNIT 04 / 4h

A research contribution

- Literature map
- research question
- baselines and ablations
- reproducibility and claims

THE TRAINING STUDIO

Fit a cooling-system parameter, validate on withheld times and write a research-style claim that distinguishes simulation from physical evidence.

EVIDENCE TO PRODUCE

- Equation and units
- calibration notebook
- uncertainty and limitation note
- short methods report

DEFEND THE REASONING

Why is calibration not equivalent to physical validation?

PRO430 / RESEARCH & PROFESSIONAL PRACTICE

Engineering communication & career evidence

12h productive allocation · 4h 48m learning + 7h 12m training

Prerequisites: AI101

UNIT 01 / 3h

Work in a team

- Issue writing
- planning and estimates
- review etiquette
- contribution records

UNIT 02 / 3h

Explain technical decisions

- Architecture diagrams
- trade-off memos
- demonstrations
- audience adaptation

UNIT 03 / 3h

Prepare a professional portfolio

- Repository hygiene
- CV evidence
- project story
- technical interview practice

UNIT 04 / 3h

Plan the next step

- Role research
- postgraduate options
- applications and references
- publication expectations

THE TRAINING STUDIO

Deliver a five-minute engineering demonstration, answer a failure-analysis question and assemble a portfolio entry linked to your actual contribution.

EVIDENCE TO PRODUCE

- Technical presentation
- project portfolio entry
- contribution log
- individual development plan

DEFEND THE REASONING

What evidence separates your own contribution from the team output?

CAP490 / RESEARCH & PROFESSIONAL PRACTICE

Industry capstone & individual technical defence

73h 20m productive allocation · 29h 20m learning + 44h training

Prerequisites: Concurrent common core. Planning begins at orientation; model implementation follows ML201 and service integration follows MLO401.

UNIT 01 / 18h 20m

Scope and evidence design

- Problem charter
- stakeholder assumptions
- acceptance tests
- data and IP controls

UNIT 02 / 18h 20m

Iterative engineering

- Baseline
- milestones
- reviews
- integration and reproducible experiments

UNIT 03 / 18h 20m

Validation and hardening

- Failure analysis
- security review
- deployment rehearsal
- limitations

UNIT 04 / 18h 20m

Demonstration and defence

- Working prototype
- report
- individual viva
- handover and portfolio

THE TRAINING STUDIO

Develop one substantial research-informed engineering prototype with a reproducible baseline and defend the decisions, limits and individual contribution.

EVIDENCE TO PRODUCE

- Working prototype
- reproducible repository
- technical report
- demonstration and individual defence

DEFEND THE REASONING

Can another engineer reproduce the result and identify when the system should not be used?

ML210 / MACHINE LEARNING

Feature engineering & feature systems

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: DAT120, ML201

UNIT 01 / 6h

Prediction-time features

- Encoding
- scaling
- missingness
- leakage control

UNIT 02 / 6h

Feature evaluation

- Ablations
- permutation checks
- stability
- selection bias

UNIT 03 / 6h

Feature stores

- Entity keys
- event time
- point-in-time joins
- offline/online parity

UNIT 04 / 6h

Operational contracts

- Freshness
- schema evolution
- ownership
- backfills

THE TRAINING STUDIO

Build a point-in-time feature table and demonstrate a leakage case and its repair.

EVIDENCE TO PRODUCE

- Feature contract
- temporal join tests
- ablation report

DEFEND THE REASONING

Why is a correct entity key insufficient for a valid historical feature join?

ML220 / MACHINE LEARNING

Ensembles, XGBoost & interpretable decisions

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: ML201

UNIT 01 / 6h

Boosting mechanics

- Residual fitting
- trees
- learning rate
- regularisation

UNIT 02 / 6h

Controlled tuning

- Validation design
- search budget
- early stopping
- reproducibility

UNIT 03 / 6h

Interpretability

- Feature attribution
- dependence plots
- correlated inputs
- local versus global explanations

UNIT 04 / 6h

Decision design

- Thresholds
- calibration
- slice errors
- operational costs

THE TRAINING STUDIO

Benchmark a boosted model against a simpler baseline using equal tuning budgets.

EVIDENCE TO PRODUCE

- Search ledger
- error slices
- baseline comparison

DEFEND THE REASONING

Why is feature importance not a causal explanation?

ML230 / MACHINE LEARNING

Causal inference & experimentation

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: MTH110, DAT120

UNIT 01 / 6h

Causal questions

- Association
- interventions
- confounding
- directed graphs

UNIT 02 / 6h

Experimental design

- Randomisation
- power
- attrition
- treatment assignment

UNIT 03 / 6h

Observational strategies

- Adjustment
- matching
- overlap
- sensitivity

UNIT 04 / 6h

Interpretation

- Effect estimates
- confidence intervals
- heterogeneous effects
- external validity

THE TRAINING STUDIO

Use a synthetic treatment dataset to compare a naive association with a justified adjustment.

EVIDENCE TO PRODUCE

- Causal graph
- estimand
- analysis notebook
- sensitivity note

DEFEND THE REASONING

Which assumptions allow an observational estimate to support an intervention claim?

ML240 / MACHINE LEARNING

Time-series forecasting & anomaly detection

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: ML201, DL301

UNIT 01 / 6h

Time-aware data

- Seasonality
- autocorrelation
- missing intervals
- resampling

UNIT 02 / 6h

Forecast baselines

- Persistence
- seasonal naive
- rolling features
- prediction horizons

UNIT 03 / 6h

Temporal models

- Statistical forecasts
- boosted models
- sequence networks
- covariates

UNIT 04 / 6h

Backtesting

- Rolling origins
- leakage
- interval coverage
- drift and anomalies

THE TRAINING STUDIO

Compare persistence, a feature-based model and a sequence approach using rolling-origin evaluation.

EVIDENCE TO PRODUCE

- Forecast notebook
- horizon-wise errors
- uncertainty report

DEFEND THE REASONING

Why can a random train-test split inflate forecasting performance?

ML250 / MACHINE LEARNING

Recommender systems & ranking

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: ML201, NLP310

UNIT 01 / 6h

Recommendation data

- Explicit and implicit feedback
- exposure
- cold start
- negative sampling

UNIT 02 / 6h

Candidate generation

- Popularity baseline
- nearest neighbours
- collaborative filtering
- embeddings

UNIT 03 / 6h

Ranking objectives

- Pointwise and pairwise learning
- ranking losses
- diversity
- calibration

UNIT 04 / 6h

Evaluation

- Recall at K
- NDCG
- temporal splits
- offline versus online outcomes

THE TRAINING STUDIO

Build a recommendation baseline and measure ranking quality, coverage and cold-start behaviour.

EVIDENCE TO PRODUCE

- Ranking benchmark
- coverage analysis
- model card

DEFEND THE REASONING

Why can historical clicks reflect exposure bias rather than pure preference?

ML260 / MACHINE LEARNING

Graph machine learning & knowledge graphs

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: MTH110, DL301

UNIT 01 / 6h

Graph foundations

- Nodes and edges
- directed graphs
- attributes
- topology

UNIT 02 / 6h

Graph learning

- Message passing
- aggregation
- node classification
- link prediction

UNIT 03 / 6h

Knowledge representation

- Triples
- schemas
- entity resolution
- graph queries

UNIT 04 / 6h

Evaluation

- Transductive versus inductive settings
- leakage
- negative edges
- temporal graphs

THE TRAINING STUDIO

Create a synthetic asset graph and compare a graph-based predictor with a non-graph baseline.

EVIDENCE TO PRODUCE

- Graph schema
- split protocol
- baseline and model comparison

DEFEND THE REASONING

How can edges inadvertently expose test information?

ML270 / MATHEMATICS & DATA

Optimisation, uncertainty & robust decisions

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: MTH110, ML201

UNIT 01 / 6h

Optimisation choices

- Convexity
- constraints
- gradient methods
- stochasticity

UNIT 02 / 6h

Uncertainty

- Aleatoric and epistemic uncertainty
- intervals
- calibration
- coverage

UNIT 03 / 6h

Robust evaluation

- Distribution shift
- sensitivity
- adversarial conditions
- abstention

UNIT 04 / 6h

Engineering decisions

- Cost functions
- trade-off frontiers
- constraint checks
- risk budgets

THE TRAINING STUDIO

Stress-test a model across shifted synthetic inputs and choose an abstention policy.

EVIDENCE TO PRODUCE

- Sensitivity plots
- calibration results
- decision-policy memo

DEFEND THE REASONING

Why does a confidence score require empirical calibration?

DL310 / DEEP LEARNING & PERCEPTION

Advanced neural architectures & efficient training

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: DL301

UNIT 01 / 6h

Architecture design

- Residual connections
- normalisation
- attention
- inductive biases

UNIT 02 / 6h

Training behaviour

- Initialisation
- schedules
- clipping
- regularisation

UNIT 03 / 6h

Efficiency

- Mixed precision
- accumulation
- profiling
- memory budgets

UNIT 04 / 6h

Ablation studies

- Controlled comparisons
- seeds
- compute reporting
- failure analysis

THE TRAINING STUDIO

Profile a small network and improve one resource metric without hiding quality regressions.

EVIDENCE TO PRODUCE

- Profiler record
- ablation table
- compute and quality report

DEFEND THE REASONING

What must stay controlled when comparing two architectures?

CV320 / DEEP LEARNING & PERCEPTION

Computer vision: detection & segmentation

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: DL301

UNIT 01 / 6h

Image datasets

- Annotation rules
- provenance
- group splits
- augmentation

UNIT 02 / 6h

Detection

- Bounding boxes
- localisation
- IoU
- nonmaximum suppression

UNIT 03 / 6h

Segmentation

- Pixel labels
- masks
- class imbalance
- boundary errors

UNIT 04 / 6h

Evaluation

- mAP concepts
- class-wise errors
- domain shift
- inference limits

THE TRAINING STUDIO

Fine-tune a small approved vision model and analyse localisation or segmentation failures.

EVIDENCE TO PRODUCE

- Annotation guide
- held-out benchmark
- failure gallery

DEFEND THE REASONING

Why can adjacent video frames in train and test produce a misleading score?

CV330 / DEEP LEARNING & PERCEPTION

3D vision & spatial intelligence

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: MTH110, CV320

UNIT 01 / 6h

Geometry

- Camera models
- coordinate frames
- calibration
- projection

UNIT 02 / 6h

3D representations

- Depth maps
- point clouds
- meshes
- occupancy

UNIT 03 / 6h

Spatial estimation

- Registration
- reconstruction
- pose
- uncertainty

UNIT 04 / 6h

Validation

- Scale errors
- occlusion
- coordinate conventions
- held-out scenes

THE TRAINING STUDIO

Reconstruct or align an approved small point-cloud scene and document its coordinate and scale assumptions.

EVIDENCE TO PRODUCE

- Calibration record
- error report
- spatial pipeline

DEFEND THE REASONING

How can an unnoticed coordinate-frame mismatch invalidate a plausible visual result?

DL340 / DEEP LEARNING & PERCEPTION

Generative modelling: VAEs, GANs & diffusion

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: DL301

UNIT 01 / 6h

Latent-variable learning

- Latent representations
- reconstruction
- variational objectives
- likelihood limits

UNIT 02 / 6h

Adversarial models

- Generator and discriminator
- instability
- mode collapse
- evaluation

UNIT 03 / 6h

Diffusion models

- Noise schedules
- denoising
- conditioning
- sampling budgets

UNIT 04 / 6h

Responsible evaluation

- Rights and provenance
- diversity
- memorisation
- task-based utility

THE TRAINING STUDIO

Compare two small generative approaches or sampling settings using an approved dataset and documented quality criteria.

EVIDENCE TO PRODUCE

- Dataset record
- generation notebook
- diversity and failure analysis

DEFEND THE REASONING

Why is an attractive sample not sufficient evidence of useful distribution learning?

NLP330 / LANGUAGE & GENERATIVE AI

Language-model adaptation & efficient fine-tuning

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: NLP310, DL301

UNIT 01 / 6h

Adaptation choices

- Prompting
- retrieval
- fine-tuning
- task suitability

UNIT 02 / 6h

Training data

- Instruction data
- filtering
- contamination
- evaluation splits

UNIT 03 / 6h

Efficient adaptation

- Low-rank adaptation
- quantisation concepts
- memory
- training budgets

UNIT 04 / 6h

Evaluation

- Held-out tasks
- regression tests
- refusal
- cost-quality trade-offs

THE TRAINING STUDIO

Adapt a small permitted model or compare prompting with adaptation on an explicit held-out task.

EVIDENCE TO PRODUCE

- Data card
- adaptation configuration
- task evaluation
- resource record

DEFEND THE REASONING

When is fine-tuning the wrong response to missing factual knowledge?

AI340 / LANGUAGE & GENERATIVE AI

Production RAG, search & evaluation

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: AI320

UNIT 01 / 6h

Information retrieval

- Lexical search
- dense retrieval
- hybrid ranking
- query formulation

UNIT 02 / 6h

Index engineering

- Chunk strategies
- metadata
- refresh
- access-aware retrieval

UNIT 03 / 6h

Evidence quality

- Reranking
- citations
- answerability
- abstention

UNIT 04 / 6h

Service evaluation

- Retrieval test sets
- support checks
- latency
- cost and traceability

THE TRAINING STUDIO

Build and evaluate a hybrid retrieval prototype with document-level permission tests.

EVIDENCE TO PRODUCE

- Retrieval dataset
- index contract
- support and abstention report

DEFEND THE REASONING

Why must access control happen before a document becomes model context?

AI350 / LANGUAGE & GENERATIVE AI

Agent engineering, tool use & evaluation

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: AI320, MLO401

UNIT 01 / 6h

Workflow architecture

- State machines
- agent loops
- planning
- deterministic alternatives

UNIT 02 / 6h

Tools and memory

- Typed tool inputs
- scoped access
- retention
- failure responses

UNIT 03 / 6h

Control and safety

- Approval checkpoints
- budgets
- injection resistance
- idempotency

UNIT 04 / 6h

Evaluation

- Task completion
- traces
- adversarial cases
- human review

THE TRAINING STUDIO

Implement a bounded engineering–assistant workflow with a simulated external–action tool and a required approval step.

EVIDENCE TO PRODUCE

- Workflow graph
- replayable traces
- tool contracts
- red–team cases

DEFEND THE REASONING

Why is a successful demo not proof that an autonomous workflow is reliable?

AI360 / LANGUAGE & GENERATIVE AI

Audio, speech & multimodal learning

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: DL301, NLP310

UNIT 01 / 6h

Audio representations

- Sampling
- spectrograms
- segmentation
- transcription data

UNIT 02 / 6h

Speech tasks

- Recognition
- diarisation concepts
- word error rate
- speaker leakage

UNIT 03 / 6h

Multimodal systems

- Image-text representations
- fusion
- alignment
- missing modalities

UNIT 04 / 6h

Evaluation

- Accent and language slices
- transcription errors
- grounding
- privacy

THE TRAINING STUDIO

Evaluate a permitted audio or image-text model across documented failure slices.

EVIDENCE TO PRODUCE

- Consent/licence record
- task baseline
- slice report
- limitation card

DEFEND THE REASONING

Why should overall word error rate be broken down by relevant conditions?

RL301 / RL & ROBOTICS

Reinforcement learning & sequential decisions

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: MTH110, ML201

UNIT 01 / 6h

Decision processes

- States
- actions
- rewards
- transitions
- discounting

UNIT 02 / 6h

Learning policies

- Value functions
- Bellman relations
- exploration
- tabular Q-learning

UNIT 03 / 6h

Environment design

- Reward specification
- termination
- partial observability
- simulator assumptions

UNIT 04 / 6h

Evaluation

- Multiple seeds
- random-policy baseline
- return variance
- policy failures

THE TRAINING STUDIO

Train tabular Q-learning in a small deterministic environment and test reward and seed sensitivity.

EVIDENCE TO PRODUCE

- Environment contract
- learning curves
- baseline comparison
- reward critique

DEFEND THE REASONING

How can a policy maximise reward while failing the intended objective?

RL310 / RL & ROBOTICS

Deep reinforcement learning: DQN & policy gradients

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: RL301, DL301

UNIT 01 / 6h

Deep value learning

- DQN
- replay buffer
- target network
- exploration schedule

UNIT 02 / 6h

Policy learning

- Policy gradients
- advantage
- actor-critic
- PPO concepts

UNIT 03 / 6h

Stable experimentation

- Seed variation
- termination handling
- evaluation episodes
- compute budgets

UNIT 04 / 6h

Transfer limits

- Simulator mismatch
- reward hacking
- safety constraints
- offline-policy limits

THE TRAINING STUDIO

Reproduce a small simulated DQN benchmark and report outcomes across seeds against a random policy.

EVIDENCE TO PRODUCE

- Training and evaluation config
- return distributions
- failure episodes

DEFEND THE REASONING

Why can one lucky learning curve hide an unstable algorithm?

ROB301 / RL & ROBOTICS

ROS 2: nodes, messages & robot software

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: AI101

UNIT 01 / 6h

ROS 2 architecture

- Nodes
- packages
- executors
- graph inspection

UNIT 02 / 6h

Communication

- Topics
- services
- actions
- message types and QoS

UNIT 03 / 6h

Robot development

- Workspaces
- build tooling
- launch files
- parameters and logs

UNIT 04 / 6h

Simulation and replay

- Turtlesim or approved simulator
- bags
- reproducible launch
- frame awareness

THE TRAINING STUDIO

Build and inspect a simulated publisher-subscriber workflow and record/replay messages in a supported ROS 2 environment.

EVIDENCE TO PRODUCE

- Package source
- launch configuration
- graph diagram
- replay evidence

DEFEND THE REASONING

When should a robot task use an action rather than a topic?

ROB310 / RL & ROBOTICS

Robot perception, localisation & SLAM

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: ROB301, MTH110

UNIT 01 / 6h

Robot frames

- Transforms
- coordinate conventions
- timestamps
- sensor extrinsics

UNIT 02 / 6h

Sensor interpretation

- Camera and range data
- filtering
- calibration
- synchronisation

UNIT 03 / 6h

Localisation and mapping

- Odometry
- sensor fusion
- SLAM concepts
- loop closure

UNIT 04 / 6h

Evaluation

- Trajectory error
- map consistency
- dropped messages
- recovery

THE TRAINING STUDIO

Replay a supplied simulated sensor recording and analyse localisation error and transform consistency.

EVIDENCE TO PRODUCE

- Frame tree
- replay procedure
- trajectory metrics
- sensor-failure note

DEFEND THE REASONING

Why can timestamp mismatch resemble a perception-model failure?

ROB320 / RL & ROBOTICS

Robot planning, control & sim-to-real boundaries

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: ROB301, MTH110

UNIT 01 / 6h

Motion planning

- Configuration spaces
- collision checks
- search
- trajectory constraints

UNIT 02 / 6h

Control fundamentals

- Feedback
- PID
- stability intuition
- actuator saturation

UNIT 03 / 6h

System integration

- Planner-controller interfaces
- watchdogs
- stop conditions
- simulation tests

UNIT 04 / 6h

Transfer and validation

- Dynamics mismatch
- calibration
- supervised trials
- hazard analysis

THE TRAINING STUDIO

Evaluate a simulated navigation or tracking task with explicit constraints and safe failure handling.

EVIDENCE TO PRODUCE

- Simulation scenarios
- trajectory and control plots
- interface contract
- validation plan

DEFEND THE REASONING

Why does simulation success not authorise unsupervised physical robot operation?

MLO420 / MLOPS & CLOUD

Kubernetes & platform engineering for ML

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: MLO401, MLO410

UNIT 01 / 6h

Cluster primitives

- Pods
- deployments
- services
- namespaces

UNIT 02 / 6h

Workload operations

- Requests and limits
- probes
- configuration
- secrets

UNIT 03 / 6h

ML platform patterns

- Batch jobs
- accelerator scheduling concepts
- storage
- model serving

UNIT 04 / 6h

Release controls

- Rollouts
- observability
- failure drills
- cost and access boundaries

THE TRAINING STUDIO

Deploy a small model service to an approved local cluster and rehearse a failed-health-check rollout.

EVIDENCE TO PRODUCE

- Manifests
- probe tests
- rollout and recovery record

DEFEND THE REASONING

How can a service be running but not ready to accept inference traffic?

MLO430 / MLOPS & CLOUD

Cloud architecture & model-serving systems

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: MLO401, MLO410

UNIT 01 / 6h

Architecture choices

- Batch
- online
- serverless
- containers and managed services

UNIT 02 / 6h

Service contracts

- Authentication
- quotas
- input limits
- queues and retries

UNIT 03 / 6h

Inference performance

- Batching
- caching
- cold starts
- latency and throughput

UNIT 04 / 6h

Operations and cost

- Monitoring
- budgets
- least privilege
- recovery and retention

THE TRAINING STUDIO

Design and locally load-test a model-serving architecture with a cost model and deployment/recovery checklist.

EVIDENCE TO PRODUCE

- Architecture diagram
- latency benchmark
- cost assumptions
- recovery plan

DEFEND THE REASONING

When is batch prediction preferable to an always-on endpoint?

MLO440 / MLOPS & CLOUD

Distributed training, GPUs & performance engineering

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: DL301, MLO401

UNIT 01 / 6h

Hardware awareness

- CPU and GPU roles
- memory bandwidth
- precision
- profiling

UNIT 02 / 6h

Distributed computation

- Data and model parallelism
- synchronisation
- communication overhead

UNIT 03 / 6h

Input and checkpoint systems

- Data loading
- sharding
- restart
- reproducible configuration

UNIT 04 / 6h

Efficiency evaluation

- Scaling curves
- resource utilisation
- failure recovery
- energy and cost

THE TRAINING STUDIO

Profile a small training workload and write a justified scaling proposal; multi-GPU execution requires an approved compute allocation.

EVIDENCE TO PRODUCE

- Profiler trace
- throughput table
- checkpoint/restart test
- scaling memo

DEFEND THE REASONING

Why can adding GPUs reduce efficiency rather than improve it?

DAT210 / MATHEMATICS & DATA

Data engineering, ETL & quality contracts

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: DAT120, AI101

UNIT 01 / 6h

Data modelling

- Relational schemas
- event data
- identifiers
- ownership

UNIT 02 / 6h

ETL and ELT

- Ingestion
- transforms
- incremental loads
- idempotency

UNIT 03 / 6h

Quality contracts

- Schema
- ranges
- uniqueness
- lineage and reconciliation

UNIT 04 / 6h

Operational data

- Scheduling
- late arrivals
- backfills
- failure and recovery

THE TRAINING STUDIO

Build an incremental synthetic-data pipeline that rejects bad records and safely replays a batch.

EVIDENCE TO PRODUCE

- Pipeline code
- data contracts
- reconciliation counts
- replay test

DEFEND THE REASONING

What prevents a retry from duplicating the same observations?

DAT220 / MATHEMATICS & DATA

Large-scale data processing & Spark

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: DAT210

UNIT 01 / 6h

Distributed tables

- Partitioning
- lazy evaluation
- execution plans
- shuffles

UNIT 02 / 6h

Transformations

- Joins
- aggregations
- window functions
- skew

UNIT 03 / 6h

ML data preparation

- Feature tables
- time-aware joins
- reproducibility
- resource budgets

UNIT 04 / 6h

Performance and quality

- Caching
- partition tuning
- reconciliation
- failure diagnosis

THE TRAINING STUDIO

Run a small local Spark transformation and explain the execution plan and why it would or would not scale.

EVIDENCE TO PRODUCE

- Transform script
- execution-plan note
- reconciliation report
- resource profile

DEFEND THE REASONING

Why is distributed processing not automatically faster for a small dataset?

MLO450 / MLOPS & CLOUD

Edge AI, IoT & efficient inference

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: DL301, MLO401

UNIT 01 / 6h

Edge constraints

- Latency
- memory
- power
- intermittent connectivity

UNIT 02 / 6h

Model optimisation

- Quantisation
- pruning
- distillation
- export and operator support

UNIT 03 / 6h

Device integration

- Sensor contracts
- local buffering
- updates
- fallback behaviour

UNIT 04 / 6h

Validation

- Representative inputs
- quality regressions
- thermal limits
- rollback

THE TRAINING STUDIO

Export or quantise a small model and compare output agreement, size and CPU latency on an approved test device.

EVIDENCE TO PRODUCE

- Export artefact
- benchmark
- compatibility record
- fallback plan

DEFEND THE REASONING

Why must accuracy be checked again after model conversion?

RES430 / RESEARCH & PROFESSIONAL PRACTICE

Scientific computing, numerical methods & twin engineering

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: RES420

UNIT 01 / 6h

Numerical modelling

- ODEs
- integration
- discretisation
- solver error

UNIT 02 / 6h

Identification

- Parameter estimation
- identifiability
- residuals
- inverse problems

UNIT 03 / 6h

State and uncertainty

- Filtering concepts
- sensitivity
- uncertainty propagation
- scenario limits

UNIT 04 / 6h

Twin validation

- Calibration sets
- unseen regimes
- physical measurements
- model governance

THE TRAINING STUDIO

Build a reduced-order scientific twin with explicit equations, fitted parameters and out-of-regime tests.

EVIDENCE TO PRODUCE

- Model specification
- parameter study
- solver verification
- validation report

DEFEND THE REASONING

How do solver error, measurement noise and model discrepancy differ?

RES440 / RESEARCH & PROFESSIONAL PRACTICE

Research methods, reproducibility & publication craft

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: RES420, PRO430

UNIT 01 / 6h

Research design

- Question
- literature matrix
- contribution
- falsifiable claims

UNIT 02 / 6h

Experiment methodology

- Baselines
- ablations
- uncertainty
- compute and data reporting

UNIT 03 / 6h

Scientific writing

- Methods
- figures
- results versus discussion
- citations and limitations

UNIT 04 / 6h

Publication process

- Venue fit
- peer review
- authorship
- rights and research integrity

THE TRAINING STUDIO

Prepare a reproducible research package and a manuscript-style report for internal review; submission is a separate decision.

EVIDENCE TO PRODUCE

- Literature matrix
- experiment protocol
- reproduction package
- manuscript draft

DEFEND THE REASONING

Why can careful submission support never guarantee acceptance or indexing?

SEC310 / RESEARCH & PROFESSIONAL PRACTICE

Responsible AI, privacy & AI-system security

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: AI320, MLO401

UNIT 01 / 6h

Responsible design

- Stakeholders
- harms
- fairness
- accessibility

UNIT 02 / 6h

Data protection

- Minimisation
- consent context
- retention
- access controls

UNIT 03 / 6h

AI threats

- Prompt injection
- data poisoning
- model leakage
- insecure tool use

UNIT 04 / 6h

Governance

- Threat modelling
- evaluation
- incident response
- human accountability

THE TRAINING STUDIO

Threat-model a small AI application and demonstrate one bounded attack case and an evidence-backed mitigation.

EVIDENCE TO PRODUCE

- Threat model
- misuse evaluation
- access tests
- incident runbook

DEFEND THE REASONING

Why does a prompt instruction not replace server-side authorisation?

AI210 / MACHINE LEARNING

Classical AI: search, reasoning & planning

24h productive allocation · 9h 36m learning + 14h 24m training

Prerequisites: AI101, MTH110

UNIT 01 / 6h

Search and heuristics

- State spaces
- breadth-first and uniform-cost search
- A-star
- admissibility

UNIT 02 / 6h

Knowledge representation

- Propositional logic
- rules
- ontologies
- inference

UNIT 03 / 6h

Planning and constraints

- Goals and actions
- constraint satisfaction
- scheduling
- search costs

UNIT 04 / 6h

AI-system design

- Symbolic versus learned components
- uncertainty
- hybrid systems
- task-based evaluation

THE TRAINING STUDIO

Implement a small route-planning or constraint-solving problem and compare two search strategies using explicit costs.

EVIDENCE TO PRODUCE

- Problem specification
- search trace
- correctness tests
- complexity comparison

DEFEND THE REASONING

Why is a learned predictor not a substitute for every search or reasoning task?

COVERAGE AND INTELLECTUAL HONESTY

An independent curriculum. An explicit delivery boundary.

Coverage was checked against public academic and industry frameworks. Their inclusion does not imply accreditation, affiliation, certification or endorsement.

ACM/IEEE-CS/AAAI CS2023 knowledge areas

<https://csed.acm.org/knowledge-areas/>

Google Cloud Professional ML Engineer skills guide

<https://cloud.google.com/learn/certification/guides/machine-learning-engineer>

Hugging Face learning domains

<https://huggingface.co/learn>

NVIDIA Physical AI learning

<https://docs.nvidia.com/learning/physical-ai/index.html>

Before a teaching cohort is confirmed

- Complete instructor notes, learner materials, lab code, datasets and answers for the adopted courses.
- Validate supported environments, faculty allocation, accessibility, assessment and supervision.
- Approve fees, student terms, refunds, records, privacy and institutional credit arrangements.
- Keep physical robotics, cloud compute and paid services explicitly scoped; do not imply hardware or unlimited accounts.

Launchpoint by MindForgeAI · A programme of Chatake Innworks Private Limited