



Artificial Intelligence MSc (2026) – Final Exam topics

Basic knowledge related to the theoretical background of artificial intelligence:

1. Basic concepts of linear algebra for machine learning; matrix decompositions (SVD, Cholesky, spectral); multivariable calculus (gradient, Jacobian, Hessian); multidimensional distributions, covariance; Theoretical foundations of PCA, CCA, and MDS; evaluation of statistical models, hypothesis testing, confidence intervals, and clustering analysis.
2. Unconditional and conditional extrema of multivariate functions. Gradient-based methods (SGD, momentum, Adam, etc.); Newton and quasi-Newton methods; convex and non-convex optimization; regularization; conditional optimization and penalty functions; theory of learning dynamics in deep models, least squares method, stochastic optimization.
3. Supervised and unsupervised learning; Single- and multivariate linear regression; Gradient descent; Stochastic and mini-batch gradient descent; Normalization of predictors; Polynomial regression; Normal equation; Logistic regression; Binary and multi-class classification; Regularization (underfitting and overfitting); Regularized linear and logistic regression;
4. Perceptron theorem; MLP architectures and training; backpropagation; activation and loss functions; vanishing and exploding gradients; CNN architecture and significant architectures; RNN, LSTM, GRU theory; SVMs and kernel methods. Neural networks; Backpropagation algorithm; Numerical gradient descent.
5. Training/Test/Validation data partitioning; Training diagnostics; Learning curves (training dataset size); Error measurement and imbalanced classes; Application of support vector machines and kernel functions; Clustering; Determining the number of clusters; Dimension reduction; Anomaly detection (normal distribution); Theory of Bagging and Random Forest; Boosting (AdaBoost, Gradient Boosting); Operating principles of XGBoost; HDBSCAN and other density-based clustering techniques; Theoretical treatment of high-dimensional problems.
6. Transformer architecture, attention mechanism; pre-training paradigms (MLM, NSP, self-supervised learning); fine-tuning (DPO, RLHF); tokenization and embedding techniques; theory of speech recognition systems (Whisper); multimodal models.
7. Autoencoders and VAEs; GANs and their main architectures; theoretical foundations of diffusion models; transformer-based generative models; conditional generation; quality evaluation metrics (FID, IS, CLIP-score); ethical and bias issues, Ethical AI and bias mitigation.
8. Distributed model training and communication; parallel computing methods; architectures of cloud-based AI systems; theoretical foundations of containerization (Docker, Kubernetes); AI pipelines and model versioning. AI regulatory frameworks; data protection, GDPR; secure AI system design; model vulnerabilities and attack types

Basic knowledge related to the practical background of artificial intelligence

1. Data cleaning, transformation, and representation; exploration of real-world datasets; error handling; data visualization; statistical diagnostics applied in practice. Fundamental concepts of data visualization, its history, and the role of visual perception.
2. Training of regression and classification models; cross-validation; performance metrics (accuracy, precision, recall, ROC-AUC); bias–variance analysis; handling overfitting and underfitting. Random Forest, XGBoost, Gradient Boosting; application of HDBSCAN; solving NLP tasks (text classification, embedding handling).
3. Prompt design; fine-tuning (LoRA, DPO, RLHF); building QA systems; solving generative text composition tasks; implementing speech recognition (Whisper). Prompt injection, data and model leakage, and model theft simulation.
4. Training autoencoders, VAEs, GANs, and diffusion models; image generation; data augmentation; quality evaluation Image and text generation, data augmentation, style transfer, running diffusion models, GAN-based tasks, quality evaluation of generated content (FID, IS, CLIP-score).
5. Grad-CAM visualization in CNNs. Explanation of attention mechanisms in transformer models. Practical application of confidence intervals, bootstrapping, and resampling. Comparison of classifiers based on ROC, AUC, Brier score, and calibration. Practical evaluation of clustering solutions (stability, indices).
6. Modeling real-world industrial problems and their AI-based solutions. Model selection, optimization, and documentation. Model reduction techniques: pruning, quantization. Acceleration solutions (knowledge distillation, tensor optimization).
7. Overview of cloud environment reliability, availability, and scalability; Presentation and analysis of the costs of cloud applications and systems; Overview of availability metrics in the cloud environment;
8. Function libraries related to artificial intelligence (AI) programming; Kernel-level runtime and compilation environments; Commonly used function libraries and tools (scikit-learn, NumPy, SciPy, pandas, Jupyter, Matplotlib, Dataflow, Keras, TensorFlow).