


EPSO AD8
EPSO PREPARATION
ARTIFICIAL
INTELLIGENCE
Technical Knowledge and the EU Context



STUDY COMMUNITY TEAM

Make AI knowledge your advantage

Technical understanding and the EU context, followed by online exam practice

Build the technical understanding to explain how AI works and the European context to judge how it should be used. Our book connects models, data and deployment with EU institutions, regulation and public-service responsibilities. You study the reasoning that turns knowledge into a justified answer.

2

CONNECTED PILLARS

Technical AI knowledge + the European Union context

14

PLANNED CHAPTERS

500

PLANNED MCQs

01

STUDY THE BOOK

02

TAKE ONLINE MOCKS

03

REVIEW AND REPEAT

Our recommended preparation route starts with the book: understand the concepts, follow the worked examples and practise the reasoning. Then move to online mock exams to apply that knowledge under time pressure and identify what needs revision.

Study the book first. Then test yourself online.

The book builds your foundation. Online mock exams help you turn that foundation into timed performance. Review the explanations after each attempt, return to the relevant topic and practise again. Together, study and online testing form our recommended route to a complete preparation programme.

500 opportunities to test your reasoning

The complete-edition plan includes 500 original MCQs with detailed answers: conceptual distinctions, numerical problems and applied scenarios. Each explanation is designed to help you understand the correct choice and the traps behind the alternatives.

This presentation describes the complete-edition plan. The book is being developed; the full question bank and all planned chapters are not yet available.

A learning experience you can see

Understand the idea before the notation

Each major topic begins with the problem it addresses. Definitions and mechanisms are developed with their assumptions, limitations and practical meaning. International terminology is retained so that you can recognise it in technical documentation and EU materials.

See the architecture take shape

Architecture diagrams, comparisons and visual explanations help you see how components interact. They support the text by showing relationships such as training versus inference, context versus weights, and model performance versus operational requirements.

Follow the numbers all the way

Calculations set out the given values, the relevant formula, the substitution, the result and its interpretation. Dimensions, units and assumptions are checked. Examples explain what a result does and does not establish, including parameter counts, memory estimates and evaluation scores.

Turn a scenario into a decision

Public-service scenarios turn an abstract concept into a concrete choice: which metric to use, how to avoid leakage, why a retrieval system fails, or how to respond to deteriorating production performance.

Know why the other answers fail

Detailed answers identify the correct choice, reconstruct the relevant reasoning and address the misconceptions behind plausible alternatives. Numerical solutions show intermediate steps rather than giving only an answer letter.

Go deeper with the sources

Numbered references connect the explanations to research, official technical documentation and EU sources. The planned legal chapters will distinguish requirements, conditions and exceptions instead of presenting rules as context-free slogans.

Recommended study sequence: work through the book and its examples, then take online mock exams. Use the results to target revision and repeat the practice. Online mock access and the materials included are specified in your selected package.

Your route into the technical core

1 Foundations of AI ML and statistics

The opening chapter explains artificial intelligence, machine learning and deep learning, then connects learning problems to data, features and targets. It develops model selection, loss functions, optimisation, regularisation and valid dataset splitting. Evaluation metrics, class imbalance, calibration, uncertainty and the distinction between correlation and causation show how to assess evidence without drawing unsupported conclusions.

2 Neural networks and deep learning

The chapter moves from artificial neurons and backpropagation to convolutional networks, sequence models, attention and Transformers. Language-model training and adaptation are explained alongside text generation, embeddings and generative architectures. Worked calculations connect parameter counts, GPU memory and computational cost to practical constraints. Diagnosis and integrated exercises bring the topics together.

3 Data and knowledge infrastructure

This chapter is planned to examine collection, preparation, annotation and data quality, together with schemas, data contracts and provenance. Feature stores and point-in-time correctness connect training data to production inputs. Metadata, lineage, FAIR principles and interoperability show how data can be organised for reliable reuse across teams and organisations.

4 LLMs RAG and generative AI applications

The planned chapter follows information from document preparation and chunking through retrieval, reranking and context construction to an answer with supporting sources. It examines prompting, tools, agents, multilingual quality and hallucinations. Retrieval failures are distinguished from generation failures so that corrective action addresses the actual cause.

5 MLOps and operating AI systems

The planned chapter covers reproducible experiments, version control, model registries and deployment pipelines. Monitoring, data drift, concept drift and incident handling connect statistical performance to operational action. Shadow deployment, canary releases and rollback are examined as parts of a controlled model lifecycle.

Where AI meets operational responsibility

6 Computing infrastructure cloud and HPC

The planned chapter explains GPU memory, bandwidth, batch workloads and interactive inference. KV caching, time to first token and tail latency connect user experience to serving design. Distributed training, sharding, parallelism, communication and containers support decisions about capacity, performance and cost.

7 AI cybersecurity and resilience

The planned chapter examines adversarial inputs, poisoning, model extraction, training-data extraction and prompt injection. Threat modelling, access control, supply-chain security and red teaming connect attacks to system-level defences. Scenarios show why a single mitigation must be assessed within the wider architecture.

8 Trustworthy AI explanations and human oversight

The planned chapter brings together bias, subgroup performance, fairness metrics, local explanations, counterfactuals and uncertainty. It examines the design of human review and accessible services. Readers learn to distinguish predictive accuracy, interpretability and appropriate oversight when assessing a system.

9 The AI Act and regulatory compliance

The planned chapter covers scope, actors, prohibited practices, high-risk classification, general-purpose AI and systemic risk. Transparency, technical documentation, conformity assessment, monitoring and supervision are connected to the relevant role and use case. Conditions, exceptions and applicable dates will be studied with the legislation and its amendments.

10 Personal data and the European data framework

The planned chapter distinguishes the GDPR from the data-protection framework for EU institutions. It examines lawful processing, impact assessments, rights, minimisation, retention and transfers. The Data Governance Act, Data Act, open data and intellectual property provide the wider context for responsible data use.

Connect AI to the European public sector

11 EU institutions and AI policies

The planned chapter explains the roles of the European Commission, AI Office, AI Board, EDPS, ENISA, JRC, EuroHPC and national authorities. European initiatives in research, infrastructure, data and skills connect the institutional landscape to practical support and supervision. Responsibilities are distinguished from scientific, advisory and enabling functions.

12 Designing and managing AI projects

The planned chapter develops problem framing, feasibility, build-versus-buy choices and procurement specifications. Lifecycle cost, vendor dependence, service levels and acceptance criteria help readers assess whether a technically impressive solution meets an organisation's actual needs. Change management connects delivery to adoption.

13 AI applications in public administration

The planned chapter brings the methods together in document processing, language tasks, anomaly detection, workload planning and image inspection. Service design, transparency, accessibility, exception handling and human responsibility connect the model to the complete administrative process.

14 Integrated scenarios and final revision

The planned closing chapter combines technical, operational and European themes. Questions ask readers to diagnose failures, compare alternatives and select justified actions under constraints. Revision revisits the distinctions that matter when several apparently reasonable answers compete.

A connected reading path

Begin with the foundations, develop the architecture and data knowledge, then move to operation, security and governance. Integrated scenarios bring the two pillars together: a technical decision is assessed alongside its organisational and European context.

Chapters 1 and 2 have been written in Greek. English translation is in progress. Chapters 3–14 and the complete 500-question bank form the planned development of the complete edition.

Why choose this book for your preparation

Build understanding you can defend

You should be able to explain your answer. Worked examples help you reconstruct the calculation, identify the assumption and justify the choice. That is a stronger foundation for revision than recognising a familiar term.

Keep the technical and EU perspectives connected

A model choice has consequences for data, infrastructure, oversight and public-service delivery. The planned chapter sequence connects these perspectives, giving you a coherent route through a broad field.

Make difficult calculations approachable

The inputs, formula, substitution and interpretation appear together. Parameter counts, memory estimates and evaluation scores become calculations you can check and repeat.

Make practice a source of insight

The planned 500-question bank develops the same approach. Detailed answers explain the correct choice and the misconceptions behind distractors, helping you target your next revision session.

Follow study with online mock exams

Once you have studied a topic, test your understanding online under timed conditions. Review your mistakes, revisit the book and repeat the practice. We recommend combining book study and online mock exams for a complete preparation programme.

Here is what that depth looks like

For a $4,096 \times 4,096$ projection, the full matrix contains $4,096 \times 4,096 = 16,777,216$ parameters. A LoRA adapter with rank 16 contains $16 \times (4,096 + 4,096) = 131,072$ trainable parameters. The ratio is $131,072 / 16,777,216 = 1/128 = 0.78125\%$.

The explanation then checks the meaning of the result: this percentage applies to the specified matrix. It does not imply the same percentage reduction in total GPU memory, because the frozen base weights, activations and other state still matter.

Build your knowledge. Practise online. Refine your preparation.

Explore the book and choose a preparation package with online mock exams.

yourstudy-community.gr • epsoexams.eu
yourstudy21@gmail.com

Written by Study Community Team. Independent educational material. Practice questions are not official EPSO questions. This brochure presents the complete-edition plan; availability and included materials are specified in the selected product offer.

Try a question that rewards real understanding

Choose your answer before reading the solution. The explanation shows the detail that separates a convincing guess from a justified calculation.

Two batches contain 100 and 900 valid tokens. Their mean negative log likelihoods are 1 and 2 respectively. Using natural logarithms, what is the overall perplexity?


- A $\exp(3)$, approximately 20.09
- B The mean of $\exp(1)$ and $\exp(2)$, approximately 5.05
- C $\exp(1.9)$, approximately 6.69
- D $\exp(1.5)$, approximately 4.48

Correct answer C

Step 1 Recover the total negative log likelihood. The first batch contributes $100 \times 1 = 100$; the second contributes $900 \times 2 = 1,800$. Together they contribute 1,900 over 1,000 valid tokens.

Step 2 Calculate the token-weighted mean loss: $1,900 / 1,000 = 1.9$. The larger batch must contribute more because it contains nine times as many tokens.

Step 3 Convert mean loss to perplexity: $PPL = \exp(1.9) = 6.685894$, approximately 6.69.




Why the other options fail: A adds batch means without weighting. B averages already exponentiated losses. D gives equal weight to unequal batches. Corpus perplexity requires aggregating token loss first and exponentiating afterwards.

See the depth behind the promise

These numerical examples already appear in the book. The calculation is only the start: the interpretation explains what you can conclude.

LoRA From a full matrix to trainable adapters

A $4,096 \times 4,096$ projection has 16,777,216 parameters. With rank 16, its two adapter factors contain $16 \times (4,096 + 4,096) = 131,072$ trainable parameters. That is $1/128$, or 0.78125%, of this base matrix.



Quantization Separate weight storage from total memory

For 7 billion parameters, 16-bit raw weights require $7 \times 10^9 \times 16/8 = 14 \times 10^9$ bytes, or 14 GB. At 4-bit, they require 3.5 GB. The calculation excludes scales, non-quantized tensors, activations, KV cache and optimizer state. It therefore does not establish that the full workload fits on a 4 GB GPU.

Attention Follow the weights into the output

With scores (1, 0), softmax gives weights approximately (0.731059, 0.268941). For values (2, 0) and (0, 4), the weighted output is (1.462117, 1.075766). The explanation shows why attention combines values rather than simply selecting the most highly weighted one.

Study the reasoning. Check the calculation. Carry the understanding into the next question.