

DUMPSBOSS.

Cognitive Project Management in AI CPMAI v7 Training & Certification

PMI CPMAI v7

Version Demo

Total Demo Questions: 12

Total Premium Questions: 126

Buy Premium PDF

<https://dumpsboss.co>

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

QUESTION NO: 1

You're being told by upper management that you need to manage a new AI project. You need to determine the AI project fit to make sure you're actually solving a real business problem.

During Phase I: Business Understanding, you should consider at least one of the following (Select all that apply):

- A. Explores a proof of concept for an AI project
- B. Enhance revenue
- C. Solves a previously unsolved problem
- D. Improve company competitiveness in the market
- E. Has the "cool" factor
- F. Solves an already solved problem but does it better and cheaper
- G. Enhance revenue, solve a previously unsolved problem, improve competitiveness, or solve an existing problem better and cheaper

ANSWER: G

Explanation:

Enhancing revenue, solving a previously unsolved business problem, improving competitive position, or improving an existing solution so that it is more effective and less costly are all valid business-fit considerations in CPMI's Business Understanding phase. The purpose of this phase is to connect a proposed AI capability to an identifiable organizational objective and measurable value, rather than treating AI as an end in itself. A team should be able to articulate the business need, the expected benefit, the stakeholders affected, and the success measures that will demonstrate whether the initiative delivered value. Revenue growth can be a direct financial benefit; a novel solution can address an unmet operational or customer need; competitive improvement can support strategic positioning; and a better, cheaper approach can create value through improved performance, efficiency, or cost reduction. These considerations help establish a credible business case before substantial effort is committed to data, modeling, deployment, and change management. This aligns with PMI's emphasis on tying projects to strategic objectives and realizing intended benefits throughout the project life cycle. See PMI's [benefits realization management guidance](#) and its overview of [business cases and project success](#).

QUESTION NO: 2

You're working with an inexperienced team and this is all their first AI project. You're trying to work on a supervised learning binary classification problem to determine if emails are spam or not.

What is the best approach for this project?

- A. Pick a simple algorithm such as naive bayes
- B. Pick a neural network algorithm since you know this works well for supervised learning approaches
- C. Pick an ensemble method since you're not sure which algorithm will perform best
- D. Pick a simple algorithm such as Gaussian mixture

ANSWER: A

Explanation:

"Pick a simple algorithm such as naive bayes" is the best approach because this is a first AI project for an inexperienced team and the stated objective is a well-defined supervised binary classification task. Naive Bayes is a practical baseline for text classification, including spam filtering: email text can be converted into features such as word or token frequencies, and the model estimates the probability that a message belongs to the spam or non-spam class. Its relatively simple training process, fast execution, and understandable probabilistic outputs let the team focus on essential CPMI activities: defining the target label, preparing representative labeled data, selecting evaluation measures such as precision, recall, and false-

positive rate, and validating results before introducing additional complexity. Starting with a simple, interpretable model also establishes a measurable baseline against which future iterations can be compared. The scikit-learn documentation describes Naive Bayes methods as supervised learning algorithms based on Bayes' theorem and notes their suitability for classification tasks, including document classification: [scikit-learn Naive Bayes documentation](#). Google's machine-learning guidance likewise emphasizes beginning with a simple baseline before increasing model complexity: [Rules of Machine Learning](#).

QUESTION NO: 3

You are being tasked to manage an AI project at your company and you need to identify which project to start with. What's the best way to approach this?

- A. Go through all possible scenarios to come up with the perfect first project.
- B. Ask key stakeholders from your group and find a small problem that would have a big return on investment and start there.
- C. Find a project that requires 100% accuracy in the results and start with that one.
- D. Ask key stakeholders from all groups for input about their pain points.

ANSWER: B

Explanation:

"Ask key stakeholders from your group and find a small problem that would have a big return on investment and start there." is the strongest approach because an initial AI initiative should be selected as a focused, feasible business-value opportunity rather than as a broad technology experiment. In CPMAI's Business Understanding phase, the team works with relevant business stakeholders to identify a meaningful pain point, define the desired business outcome, and establish measurable success criteria. Choosing a narrowly bounded problem reduces uncertainty in data availability, model development, deployment, and adoption. At the same time, a high expected return on investment gives the initiative a credible business case and makes it easier to secure sponsorship and resources.

An early, valuable use case also creates an opportunity to validate the organization's AI delivery practices, including stakeholder collaboration, data preparation, evaluation, and operational rollout. Demonstrating a tangible benefit quickly builds trust in subsequent AI work and provides practical learning that can be applied to larger initiatives. This aligns with PMI's emphasis on linking AI projects to organizational value and with the CPMAI framework's business-first orientation. See PMI's overview of AI and project management at [PMI Artificial Intelligence resources](#) and the CPMAI certification information at [PMI CPMAI certification](#).

QUESTION NO: 4

You have been tasked at your organization to manage a large language model (LLM) project. Identify what LLMs are useful for. (Select all that apply.)

- A. Process automation
- B. Text summarization
- C. Machine Translation
- D. Classify and categorize content
- E. Code generation
- F. Improve search quality

ANSWER: B C D E F

Explanation:

Large language models are especially useful where the input, output, or both are natural language or programming language. **Text summarization** is a core generative capability: an LLM can identify salient themes and produce a concise version of source material while adapting the length, format, and intended audience. **Machine Translation** is also a language-to-language task for which LLMs can generate text in a target language while retaining the source meaning and context. **Classify and categorize content** is appropriate because an LLM can interpret text and assign one or more predefined labels, such as topic, sentiment, urgency, or document type. **Code generation** fits because source code is a structured language; LLMs can generate code, explain it, transform it, and assist with tests or documentation. **Improve search quality** is a practical LLM use case through semantic understanding, query rewriting or expansion, result ranking, and generation of concise answers grounded in retrieved material. Project teams should define evaluation measures for each use case, such as summary faithfulness, translation quality, classification accuracy, code-security review results, and search relevance. These capabilities align with the language understanding and generation applications described in the [NIST AI Risk Management Framework](#) and the practical generative-AI patterns outlined in [Google Cloud's generative AI use-case guidance](#).

QUESTION NO: 5

A project manager meets with a customer for initial discussions about an upcoming project. At the end of the meeting, the customer asks the project manager for a rough estimate of the project duration. Based on her experience with three similar projects, the project manager provides an estimate of 8–10 months.

What's wrong with this timeframe?

- A. It's underestimating the project timeline by 3 months
- B. It fits into a waterfall timeframe, but not an agile project timeframe
- C. It's not accounting for data preparation timelines
- D. It's not accounting for potential project delays

ANSWER: C

Explanation:

It's not accounting for data preparation timelines is correct because data preparation is a major, distinct body of work in an AI project and must be represented in an early duration estimate. A comparison with superficially similar prior projects may provide a useful starting range, but it is not sufficient unless the project manager has verified the condition, accessibility, ownership, completeness, quality, labeling needs, and integration requirements of the new project's data. These factors can substantially affect the schedule before model development can proceed reliably.

In the CPMAI lifecycle, data preparation follows data understanding and includes activities such as cleansing, transforming, enriching, selecting, and organizing data for model development. This work often involves iterative collaboration among business stakeholders, data engineers, and data scientists. A credible rough-order estimate should therefore explicitly include discovery and preparation effort, along with appropriate assumptions and uncertainty, rather than treating historical project duration alone as a complete basis for the timeline. PMI describes CPMAI as a lifecycle-oriented framework for managing AI initiatives in its [CPMAI certification overview](#). For additional context on the scope of data-preparation work in an analytics lifecycle, see IBM's [data preparation overview](#).

QUESTION NO: 6

A project team is considering an AI solution to prioritize incoming customer-support cases. During Phase I: Business Understanding, which actions would best support an AI Go/No Go decision? (Select three.)

- A. Define the desired business outcome and measurable success criteria.
- B. Confirm that historical cases contain usable inputs and priority outcomes.
- C. Determine whether a simpler rules-based approach can meet the need.
- D. Select the final machine learning algorithm.
- E. Begin full-scale data cleansing for every available source.

ANSWER: A B C

Explanation:

Defining the business outcome and measurable success criteria is necessary because the team must know whether the proposed solution is expected to create value, such as reducing response time for urgent cases or improving customer satisfaction. Confirming that historical cases contain usable inputs and an outcome or priority signal is also essential because a predictive solution requires evidence from which to learn and later be evaluated. Determining whether a simple rules-based approach can meet the need is important because AI should not be selected when deterministic rules can provide the required outcome more simply and reliably.

Selecting a specific algorithm and beginning full data cleaning are premature. Those activities follow after the team has established that the problem is suitable, valuable, and feasible for an AI approach.

QUESTION NO: 7

Creating machine learning models can be complicated. Your team wants to use tools called Automated Machine Learning (AutoML) to simplify the process. You know of another team that has used AutoML tools and it ' s saved the team a lot of time.

However, what ' s the one area you should not have the AutoML tool help with?

- A. Automatic model assessment
- B. Iterative modeling and evaluation
- C. Automatic hyperparameter tuning
- D. Automatic model selection
- E. Automatic algorithm selection

ANSWER: B

Explanation:

Iterative modeling and evaluation is the activity that should remain under the team's direction rather than being delegated to an AutoML tool. AutoML can accelerate technical experimentation by generating candidate models, tuning settings, ranking results, and reporting standardized performance measures. Those capabilities reduce manual effort and can broaden the range of approaches assessed during model development.

However, iterative modeling and evaluation is more than comparing automated scores. The project team must interpret results in the context of the stated business objective, success criteria, data limitations, operational constraints, risk tolerance, and stakeholder expectations. It must decide whether a result is meaningful, whether the model behavior is acceptable for its intended use, what additional data or features may be needed, and whether to refine, replace, or stop pursuing a modeling approach. These are accountable project and business decisions requiring human judgment and domain knowledge. PMI's AI-focused guidance emphasizes responsible, value-driven delivery and governance, which requires teams to retain oversight of AI outcomes rather than treating tool output as an autonomous decision. See PMI's [AI guidance](#) and Google Cloud's overview of [AutoML](#).

QUESTION NO: 8

Your team was given a large dataset and has been tasked with organizing the data by type to make better insights from the results. You are facing problems with the approach that the previous project lead used which was a regression algorithm.

What type of algorithm is the best approach for this project?

- A. Multiclass Classification
- B. Regression
- C. Clustering

D. Binary (or Binomial) Classification

ANSWER: C

Explanation:

Clustering is the appropriate approach because the stated objective is to organize a large dataset by type so the team can identify meaningful patterns and derive better insights. This is a grouping task: the algorithm examines similarity among records and places similar observations into groups, commonly called clusters. These discovered groups can then be profiled, named, and used to support analysis, such as identifying customer segments, document themes, operational patterns, or distinct data populations.

In CPMI-style AI project work, the model approach must fit the business question and the nature of the available data. When categories are not already supplied as target labels and the team needs to discover the natural structure of the data, an unsupervised-learning approach is appropriate. Clustering methods such as K-means can group observations using selected features and a similarity or distance measure. The resulting clusters provide an actionable basis for interpretation, visualization, and subsequent decision-making. Model selection should also include validating that the groups are sufficiently distinct, stable, and meaningful to the intended business use. See PMI's AI project-management resources at [PMI: AI Project Management](#) and the practical description of K-means clustering in [scikit-learn's clustering documentation](#).

QUESTION NO: 9

Recently, you implemented an augmented intelligence application at work to help employees do their job better. However, employees have been resistant to this change and aren't using the application as expected. What could have been done better to get the team to feel comfortable with this technology and use it? (Select all that apply.)

- A. Ask end users what information and technology they need to help them do their job better and build the tool to help with these pain points.
- B. Have the team that built the technology relay to employees this tool is to augment, and not replace their jobs.
- C. Have upper management relay to employees this tool is to augment, and not replace their jobs.
- D. Provide training for everyone to have all employees feel more comfortable using the technology even if they aren't using the technology yet.

ANSWER: A B C D

Explanation:

Successful augmented-intelligence adoption depends on human-centered change management, not merely deploying a technically capable application. "Ask end users what information and technology they need to help them do their job better and build the tool to help with these pain points" is correct because early and continuous user involvement ensures that the solution addresses genuine workflow needs and gives employees ownership in its design. "Have the team that built the technology relay to employees this tool is to augment, and not replace their jobs" and "Have upper management relay to employees this tool is to augment, and not replace their jobs" are both appropriate: consistent communication from delivery teams and organizational leadership can reduce uncertainty, establish the intended purpose of the technology, and build trust. "Provide training for everyone to have all employees feel more comfortable using the technology even if they aren't using the technology yet" is also correct because broad, role-aware training develops familiarity and confidence before use is required. These practices reflect responsible AI implementation: engage affected stakeholders, communicate transparently, and prepare people to work effectively with the new capability. PMI emphasizes tailoring engagement and communications to stakeholders throughout a project, while NIST's AI Risk Management Framework identifies organizational culture, workforce readiness, and human-AI interaction as important considerations for trustworthy AI adoption. See [PMI stakeholder management guidance](#) and [NIST AI Risk Management Framework](#).

QUESTION NO: 10 - (SIMULATION)

Your team is running a simulation-based optimization exercise to increase routing efficiency. Learning for this exercise is done through "trial and error." Which type of machine learning approach is being leveraged for this exercise?

ANSWER: Check the explanation for the answer**Explanation:****Answer: Reinforcement Learning (RL).**

Reinforcement learning uses a *trial-and-error* approach: an agent tries routing decisions in a simulated environment, receives rewards or penalties based on efficiency outcomes, and progressively learns a policy that maximizes the desired result (such as reduced travel time, cost, or distance).

Simulation-based routing optimization is a common use case for reinforcement learning because it permits many safe, repeated decision trials before applying the learned routing policy in a real environment.

QUESTION NO: 11

A demand-forecasting model has been operating for several months. Recent forecasts are less reliable, and the team now has actual sales results for the affected period.

What actions should the team take to determine whether model drift is affecting the system? (Select three.)

- A. Compare recent forecast errors with the approved baseline using actual sales results.
- B. Investigate whether the relationship between important inputs and sales outcomes has changed.
- C. Review recent changes to the model, features, data sources, and operating environment.
- D. Immediately retrain the model without investigating the source of the decline.
- E. Ignore recent results because the original model evaluation was successful.
- F. Delete prior model versions to avoid comparisons with the current release.

ANSWER: A B C**Explanation:**

The team should compare recent forecast performance with the established baseline using actual sales results, investigate whether the relationship between input factors and sales has changed, and review whether recent changes to the model, features, or operational environment contributed to the decline. These actions use production evidence to determine whether the model no longer generalizes adequately in its current context.

Retraining may be an appropriate response after the team identifies the cause and validates an updated model, but it should not replace investigation. Ignoring recent results because the original evaluation was successful fails to address the current operational evidence. Deleting prior versions would also prevent useful comparison and rollback.

QUESTION NO: 12

Your team is preparing to operationalize an AI model that predicts equipment failures. The model will send maintenance alerts to operations teams, and an incorrect or unavailable prediction could affect production planning.

What should be included in the operationalization plan? (Select three.)

- A. Identify the approved model version and its deployment configuration.
- B. Define monitoring measures, thresholds, and escalation actions.
- C. Establish a rollback or contingency process for production issues.
- D. Delete earlier model versions after the new version is deployed.
- E. Treat every production prediction as confirmed ground truth.

F. Assign operational ownership only after a production incident occurs.

ANSWER: A B C

Explanation:

A production-ready operationalization plan needs a clearly identified approved model version, defined monitoring and alert thresholds, and a rollback or contingency process. Version identification ensures that the team can trace predictions to a specific approved model artifact. Monitoring establishes how the team will detect failures, degraded performance, data changes, or service issues. A rollback or contingency process allows operations to continue safely if the deployed model does not perform as expected.

Deleting prior model versions removes the ability to compare or restore a known model. Treating every prediction as confirmed truth prevents meaningful performance monitoring because ground truth must be independently established. Waiting for an incident before assigning ownership creates an avoidable accountability gap.