

EXAMTOPICS

- Expert Verified, Online, Free.



CERTIFICATION TEST

- CertificationTest.net - Cheap & Quality Resources With Best Support

Machine learning is best described as a type of algorithm by which?

- A. Systems can mimic human intelligence with the goal of performing routine tasks.
- B. Systems can automatically improve from experience through predictive patterns.
- C. Statistical inferences are drawn from a sample with the goal of predicting human intelligence.
- D. Previously unknown properties are discovered in data and used to predict and make improvements in the data.

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

Random forest algorithms are in what type of machine learning model?

- A. Symbolic.
- B. Generative.
- C. Discriminative.
- D. Natural language processing.

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

A company developed AI technology that can analyze text, video, images and sound to tag content, including the names of animals, humans and objects.

What type of AI is this technology classified as?

- A. Deductive inference.
- B. Multi-modal model.
- C. Transformative AI.
- D. Expert system.

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

If it is possible to provide a rationale for a specific output of an AI system, that system can best be described as:

- A. Accountable.
- B. Transparent.
- C. Explainable.
- D. Reliable.

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

CASE STUDY -

Please use the following to answer the next question:

A company is considering the procurement of an AI system designed to enhance the security of IT infrastructure. The AI system analyzes how users type on their laptops, including typing speed, rhythm and pressure, to create a unique user profile. This data is then used to authenticate users and ensure that only authorized personnel can access sensitive resources.

When prioritizing the updates to its policies, rules and procedures to include the new AI system for user authentication, the organization should:

- A. Update third-party data sharing policies.
- B. Update security controls for sensitive data.
- C. Ensure that any personal data used is only processed for a specific and lawful purpose.
- D. Reduce the complexity of the policy to make it easier for non technical employees to understand.

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

What type of organizational risk is associated with AI's resource-intensive computing demands?

- A. People risk.
- B. Security risk.
- C. Third-party risk.
- D. Environmental risk.

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

A hospital implements an AI system to assist doctors in diagnosing diseases based on historical patient data. Which one of the following model types best describes this system?

- A. Inference.
- B. Statistical.
- C. Probabilistic.
- D. Deterministic.

Suggested Answer: C

Community vote distribution



🗨️ **Ecommail** 1 month, 1 week ago

Selected Answer: C

The best answer is C. Probabilistic.

An AI system that assists doctors in diagnosing diseases based on historical patient data typically relies on probabilistic models. These models analyze patterns in the data and calculate the likelihood of various diagnoses based on probabilities, accounting for uncertainty and variability in patient symptoms and outcomes.

upvoted 4 times

🗨️ **Tagwa** 1 month, 1 week ago

Selected Answer: A

The AI system in the hospital uses historical patient data to assist in diagnosing diseases, which means it is making predictions or conclusions based on input data – this is the core of inference in AI.

Let's quickly go over the options:

A. Inference: This is the process where a trained model makes predictions or decisions based on new data. In this case, the AI uses past data to make diagnostic suggestions – a classic use of inference.

upvoted 1 times

🗨️ **Gtc_** 3 weeks, 4 days ago

Inference is not the type of model itself, it is the process that the model follows to determine.

upvoted 2 times

Which of the following AI uses is best described as human-centric?

- A. Pattern recognition algorithms are used to improve the accuracy of weather predictions, which benefits many industries and everyday life.
- B. Autonomous robots are used to move products within a warehouse, allowing human workers to reduce physical strain and alleviate monotony.
- C. Machine learning is used for demand forecasting and inventory management, ensuring that consumers can find products they want when they want them.
- D. Virtual assistants are used to adapt educational content and teaching methods to individuals, offering personalized recommendations based on ability and needs.

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

Which of the following is a foundational characteristic of effective AI governance?

- A. Engagement of a cross-functional team.
- B. Reliance on tested vendor management processes.
- C. Thorough reviews of a company's public filings with experts.
- D. Uniform policies and procedures across developer, deployer and user roles.

Suggested Answer: A

Currently there are no comments in this discussion, be the first to comment!

CASE STUDY -

Please use the following to answer the next question:

A company is considering the procurement of an AI system designed to enhance the security of IT infrastructure. The AI system analyzes how users type on their laptops, including typing speed, rhythm and pressure, to create a unique user profile. This data is then used to authenticate users and ensure that only authorized personnel can access sensitive resources.

All of the following are obligations of the company as a data controller when implementing its AI system EXCEPT?

- A. Ensuring that third-party processors are based in the same country as the company.
- B. Allowing data subject access requests (DSARs).
- C. Implementing technical and organizational measures.
- D. Conducting a Data Protection Impact Assessment (DPIA) / Privacy Impact Assessment (PIA).

Suggested Answer: A

Currently there are no comments in this discussion, be the first to comment!

CASE STUDY -

Please use the following to answer the next question:

A company is considering the procurement of an AI system designed to enhance the security of IT infrastructure. The AI system analyzes how users type on their laptops, including typing speed, rhythm and pressure, to create a unique user profile. This data is then used to authenticate users and ensure that only authorized personnel can access sensitive resources.

The data processed by the AI system would be classified as:

- A. Non-sensitive personal data, since it does not reveal information about health, gender or race.
- B. Organizational data, since it is part of the authentication process.
- C. Non-personal data, as long as it is not linked to a user ID.
- D. Special category data, if it can be used to uniquely identify a person.

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

Which of the following typical approaches is a large organization least likely to use to responsibly train stakeholders on AI terminology, strategy and governance?

- A. Providing all technical employees education on AI development so they can retool and participate in the development of AI systems.
- B. Providing training on AI ethics, based on the extent to which the organization seeks to promote a responsible AI culture.
- C. Providing role-specific training, based on whether the organization uses a centralized, federated or decentralized governance mode.
- D. Providing information and education to customers and users to understand the capabilities and limitations of the AI tools with which they interact.

Suggested Answer: A

Currently there are no comments in this discussion, be the first to comment!

All of the following are elements of establishing a global AI governance infrastructure EXCEPT:

- A. Providing training to foster a culture that promotes ethical behavior.
- B. Creating policies and procedures to manage third-party risk.
- C. Understanding differences in norms across countries.
- D. Publicly disclosing ethical principles.

Suggested Answer: D

Community vote distribution

B (100%)

🗨️ 👤 **Tagwa** 3 weeks ago

Selected Answer: B

B. Creating policies and procedures to manage third-party risk.

Explanation:

Establishing a global AI governance infrastructure focuses on creating a unified and ethical framework for AI use across international contexts. The key elements typically include:

- Providing training to foster a culture that promotes ethical behavior (A): Ensures that employees and stakeholders understand and uphold ethical AI principles.
- Understanding differences in norms across countries (C): Essential for global governance, as cultural, legal, and ethical norms vary.
- Publicly disclosing ethical principles (D): Builds transparency and trust in the organization's AI governance approach.

upvoted 1 times

In the context of increasing use of AI in business operations, your company seeks to update its data privacy policies. You are tasked with evaluating the current policies and proposing necessary updates to address AI-specific risks regarding protection of personal data. Which of the following would be the most effective addition to the company's data privacy policies?

- A. Request final review of the policy by senior management.
- B. Request regular audits of the AI Models.
- C. Prohibit the use of AI tools within the company.
- D. Require security training to employees before using AI systems.

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

Your management consulting firm is planning to use an AI system to support its employees.
Which category of operator applies to the firm in this context?

- A. Authorized representative.
- B. Distributor.
- C. Provider.
- D. Deployer.

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

CASE STUDY -

Please use the following to answer the next question:

A premier payroll services company that employs thousands of people globally, is embarking on a new hiring campaign and wants to implement policies and procedures to identify and retain the best talent. The new talent will help the company's product team expand its payroll offerings to companies in the healthcare and transportation sectors, including in Asia.

It has become time consuming and expensive for HR to review all resumes, and they are concerned that human reviewers might be susceptible to bias.

To address these concerns, the company is considering using a third-party AI tool to screen resumes and assist with hiring. They have been talking to several vendors about possibly obtaining a third-party AI-enabled hiring solution, as long as it would achieve its goals and comply with all applicable laws.

The organization has a large procurement team that is responsible for the contracting of technology solutions. One of the procurement team's goals is to reduce costs, and it often prefers lower-cost solutions. Others within the company deploy technology solutions into the organization's operations in a responsible, cost-effective manner.

The organization is aware of the risks presented by AI hiring tools and wants to mitigate them. It also questions how best to organize and train its existing personnel to use the AI hiring tool responsibly. Their concerns are heightened by the fact that relevant laws vary across jurisdictions and continue to change.

All of the following are potential negative consequences created by using the AI tool to help make hiring decisions EXCEPT:

- A. Automation bias.
- B. Candidate quality.
- C. Privacy violations.
- D. Disparate impacts.

Suggested Answer: B

Currently there are no comments in this discussion, be the first to comment!

CASE STUDY -

Please use the following to answer the next question:

A premier payroll services company that employs thousands of people globally, is embarking on a new hiring campaign and wants to implement policies and procedures to identify and retain the best talent. The new talent will help the company's product team expand its payroll offerings to companies in the healthcare and transportation sectors, including in Asia.

It has become time consuming and expensive for HR to review all resumes, and they are concerned that human reviewers might be susceptible to bias.

To address these concerns, the company is considering using a third-party AI tool to screen resumes and assist with hiring. They have been talking to several vendors about possibly obtaining a third-party AI-enabled hiring solution, as long as it would achieve its goals and comply with all applicable laws.

The organization has a large procurement team that is responsible for the contracting of technology solutions. One of the procurement team's goals is to reduce costs, and it often prefers lower-cost solutions. Others within the company deploy technology solutions into the organization's operations in a responsible, cost-effective manner.

The organization is aware of the risks presented by AI hiring tools and wants to mitigate them. It also questions how best to organize and train its existing personnel to use the AI hiring tool responsibly. Their concerns are heightened by the fact that relevant laws vary across jurisdictions and continue to change.

Which other stakeholder groups should be involved in the selection and implementation of the AI hiring tool?

- A. Finance and Legal.
- B. Marketing and Compliance.
- C. Supply Chain and Marketing.
- D. Litigation and Product Development.

Suggested Answer: A

Currently there are no comments in this discussion, be the first to comment!

CASE STUDY -

Please use the following to answer the next question:

A premier payroll services company that employs thousands of people globally, is embarking on a new hiring campaign and wants to implement policies and procedures to identify and retain the best talent. The new talent will help the company's product team expand its payroll offerings to companies in the healthcare and transportation sectors, including in Asia.

It has become time consuming and expensive for HR to review all resumes, and they are concerned that human reviewers might be susceptible to bias.

To address these concerns, the company is considering using a third-party AI tool to screen resumes and assist with hiring. They have been talking to several vendors about possibly obtaining a third-party AI-enabled hiring solution, as long as it would achieve its goals and comply with all applicable laws.

The organization has a large procurement team that is responsible for the contracting of technology solutions. One of the procurement team's goals is to reduce costs, and it often prefers lower-cost solutions. Others within the company deploy technology solutions into the organization's operations in a responsible, cost-effective manner.

The organization is aware of the risks presented by AI hiring tools and wants to mitigate them. It also questions how best to organize and train its existing personnel to use the AI hiring tool responsibly. Their concerns are heightened by the fact that relevant laws vary across jurisdictions and continue to change.

If the company does not deploy and use the AI hiring tool responsibly in the United States, its liability would likely increase under all of the following laws EXCEPT:

- A. Anti-discrimination laws.
- B. Product liability laws.
- C. Accessibility laws.
- D. Privacy laws.

Suggested Answer: B

Currently there are no comments in this discussion, be the first to comment!

What is the primary purpose of an AI impact assessment?

- A. To determine whether a conformity assessment is needed.
- B. To escalate the findings to the appropriate owner(s).
- C. To identify and measure the benefits of an AI system.
- D. To anticipate and manage the potential risks and harms of an AI system.

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

The use of paid generative AI public tools is appealing because? (Choose three.)

- A. They are convenient to adopt.
- B. They have additional privacy and security controls.
- C. They have frequent enhancements of new features.
- D. They provide transparency in models and in decision-making.
- E. They eliminate concerns about the data used to generate outputs.

Suggested Answer: *ABC*

Currently there are no comments in this discussion, be the first to comment!

Within an established AI governance infrastructure, what might be the most effective governance action to handle third-party AI systems deemed to be high-risk?

- A. Align organizational impact assessment activities with relevant regulatory or legal requirements.
- B. Re-evaluate the purchase of the third-party AI system deemed to be high-risk, and consider other vendors.
- C. Establish policies for handling third-party system failures that include consideration of redundancy mechanisms for vital third-party AI system.
- D. Delegate the power, resources and authorization among executive leadership to perform risk management to each appropriate level throughout the management chain.

Suggested Answer: A

Currently there are no comments in this discussion, be the first to comment!

CASE STUDY -

Please use the following to answer the next question:

A premier payroll services company that employs thousands of people globally, is embarking on a new hiring campaign and wants to implement policies and procedures to identify and retain the best talent. The new talent will help the company's product team expand its payroll offerings to companies in the healthcare and transportation sectors, including in Asia.

It has become time consuming and expensive for HR to review all resumes, and they are concerned that human reviewers might be susceptible to bias.

To address these concerns, the company is considering using a third-party AI tool to screen resumes and assist with hiring. They have been talking to several vendors about possibly obtaining a third-party AI-enabled hiring solution, as long as it would achieve its goals and comply with all applicable laws.

The organization has a large procurement team that is responsible for the contracting of technology solutions. One of the procurement team's goals is to reduce costs, and it often prefers lower-cost solutions. Others within the company deploy technology solutions into the organization's operations in a responsible, cost-effective manner.

The organization is aware of the risks presented by AI hiring tools and wants to mitigate them. It also questions how best to organize and train its existing personnel to use the AI hiring tool responsibly. Their concerns are heightened by the fact that relevant laws vary across jurisdictions and continue to change.

Which of the following measures should the company adopt to best mitigate its risk of reputational harm from using the AI tool?

- A. Test the AI tool pre- and post-deployment.
- B. Ensure the vendor provides indemnification for the AI tool.
- C. Require the procurement and deployment teams to agree upon the AI tool.
- D. Continue to require the company's hiring personnel to manually screen all applicants.

Suggested Answer: A

None

You are a privacy program manager at a large e-commerce company that uses an AI tool to deliver personalized product recommendations based on visitors' personal information that has been collected from the company website, the chatbot and public data the company has scraped from social media.

A user submits a data access request under an applicable US state privacy law, specifically seeking a copy of their personal data, including information used to create their profile for product recommendations.

What is the most challenging aspect of managing this request?

- A. Some of the visitor's data is synthetic data that the company does not have to provide to the data subject.
- B. The data subject's data is structured data that can be searched, compiled and reviewed only by an automated tool.
- C. The data subject is not entitled to receive a copy of their data because some of it was scraped from public sources.
- D. Some of the data subject's data is unstructured data and you cannot untangle it from the other data, including information about other individuals.

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

An artist has been using an AI tool to create digital art and would like to ensure that it has copyright protection in the United States. Which of the following is most likely to enable the artist to receive copyright protection?

- A. Ensure the tool was trained using publicly available content.
- B. Obtain a representation from the AI provider on how the tool works.
- C. Provide a log of the prompts the artist used to generate the images.
- D. Update the images in a creative way to demonstrate that it is the artist's.

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

A French medical research center wishes to develop an AI-based system which will predict the risk of serious diseases based on the patient's genetic data. In order to do so it contracts with a tech company and provides it with patients' data previously obtained by the center during the research.

To guarantee compliance when processing special categories of personal data, the medical research center must ensure that:

- A. The AI-based system is designed for the purposes of preventive medicine.
- B. The patients' health and genetic data is anonymized.
- C. The patients have given explicit consent to using the data.
- D. The tech company is located in the EU and is not cloud-based.

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

According to the GDPR's transparency principle, when an AI system processes personal data in automated decision-making, controllers are required to provide data subjects specific information on?

- A. The existence of automated decision-making and meaningful information on its logic and consequences.
- B. The personal data used during processing, including inferences drawn by the AI system about the data.
- C. The data protection impact assessments carried out on the AI system and legal bases for processing.
- D. The contact details of the data protection officer and the data protection national authority.

Suggested Answer: A

Currently there are no comments in this discussion, be the first to comment!

A company subject to GDPR is building its governance framework for how it will collect data to be used for training of AI models. The most important thing the company can do to ensure GDPR compliance is:

- A. Include the requirement to fully anonymize data used to train its models.
- B. Establish a data retention schedule for data used to train its models.
- C. Source training data from a reputable company to train its models.
- D. Minimize the amount of data used to train its models.

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

Under the EU AI Act, which of the following compliance actions applies only to General Purpose AI models with systemic risk?

- A. Publishing a detailed summary of the data used to train the model.
- B. Maintaining up-to-date technical documentation, including testing details.
- C. Implementing an intellectual property policy to comply with EU copyright laws.
- D. Making information available to downstream providers who integrate the model into their AI systems.

Suggested Answer: A

Currently there are no comments in this discussion, be the first to comment!

Which of the following situations would be least likely to raise concerns under existing consumer protection laws?

- A. An AI algorithm being used in a credit decision making process by a financial institution.
- B. An AI customer service system claiming that it is as accurate as a human support agent.
- C. An AI tool using scraped digital content to generate news summaries on a publishing website.
- D. An online platform offering recommendations to its users by displaying user specific content and targeted advertisements.

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

What is the primary reason the EU is considering updates to its Product Liability Directive?

- A. To increase the minimum warranty level for defective goods.
- B. To define new liability exemptions for defective products.
- C. To address digital services and connected products.
- D. To address free and open-source software.

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

A US company has developed an AI system, CrimeBuster 7909, that collects information about incarcerated individuals that predicts whether someone is likely to commit another crime if released from prison.

When considering expanding to the EU market, this type of technology would:

- A. Require the company to register the tool with the EU database.
- B. Require the application of privacy enhancing technologies.
- C. Be subject to approval by the relevant EU authority.
- D. Be banned under the EU AI Act.

Suggested Answer: *D*

Community vote distribution



None

Which of the following disclosures is NOT required for an EU organization that developed and deployed a high-risk AI system?

- A. The human oversight measures employed.
- B. How an individual may contest a decision.
- C. The location(s) where data is stored.
- D. The fact that an AI system is being used.

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

In accordance with the EU AI Act, for how long after a high-risk AI system has been placed on the market must the provider keep the relevant documentations at the disposal of the national competent authorities?

- A. 10 years.
- B. 8 years.
- C. 6 years.
- D. 5 years.

Suggested Answer: A

Currently there are no comments in this discussion, be the first to comment!

The OECD's Ethical AI Governance Framework is a self-regulation model that proposes to prevent societal harms by:

- A. Establishing explainability criteria to ethically source and use data to train AI systems
- B. Defining ethical requirements specific to each industry sector and high-risk AI domain.
- C. Focusing on ethical AI technical design and post-deployment monitoring
- D. Balancing AI innovation with ethical considerations.

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

ISO 42001 International Standard offers guidance for organizations to develop trustworthy AI management systems by:

- A. Requiring specific minimum parameters for key suppliers and key aspects of AI management systems.
- B. Requiring organizations to continuously improve the effectiveness of their AI management systems.
- C. Focusing on high-risk aspects of development of AI management systems.
- D. Explicitly over-riding previously issued and now outdated ISO standards.

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

What is the main purpose of accountability structures under the Govern function of the NIST AI Risk Management Framework?

- A. To empower and train appropriate cross-functional teams.
- B. To establish diverse, equitable and inclusive processes.
- C. To determine responsibility for allocating budgetary resources.
- D. To enable and encourage participation by external stakeholders.

Suggested Answer: A

Community vote distribution

A (100%)

 **Ecommail** 1 month, 3 weeks ago

Selected Answer: A

Answer A is correct: To empower and train appropriate cross-functional teams.

Reference: NIST AI 100-1 (AI RMF 1.0) page 23.

"GOVERN 2: Accountability structures are in place so that the appropriate teams and individuals are empowered, responsible, and trained for mapping,

measuring, and managing AI risks."

upvoted 2 times

The initial pilot effort for NIST's Assessing Risks and Impacts of AI (ARIA) Program is focused on risks associated with which of the following?

- A. Large language models.
- B. Text-to-image models.
- C. Recommender systems.
- D. Facial recognition systems.

Suggested Answer: A

Community vote distribution



 **Ecommail** 1 month, 3 weeks ago

Selected Answer: A

Answer A is correct.

The initial evaluation (ARIA) will be conducted as a pilot effort to fully exercise the NIST ARIA test environment. ARIA will focus on risks and impacts associated with large language models (LLMs). Future iterations of ARIA may consider other types of generative AI technologies such as text-to-image models, or other forms of AI such as recommender systems or decision support tools. A compelling and exploratory set of tasks will aim to elicit pre-specified (and non-specified) risks and impacts across three levels of testing: model testing, red-teaming, and field testing.

Source: <https://ai-challenges.nist.gov/aria>

upvoted 2 times

CASE STUDY -

Please use the following to answer the next question:

A global marketing agency is adapting a large language model (“LLM”) to generate content for an upcoming marketing campaign for a client’s new product: a hard hat designed for construction workers of any gender to better protect them from head injuries.

The marketing agency is accessing the LLM through an application programming interface (“API”) developed by a third-party technology company.

They want to generate text to be used for targeted advertising communications that highlight the benefits of the hard hat to potential purchasers.

Both the marketing agency and the technology company have taken reasonable steps to address AI governance.

The marketing company has:

Entered into a contract with the technology company with suitable representations and warranties.

Completed an impact assessment on the LLM for this intended use.

Built technical guidance on how to measure and mitigate bias in the LLM.

Enabled technical aspects of transparency, explainability, robustness and privacy.

Followed applicable regulatory requirements.

Created specific legal statements and disclosures regarding the use of the AI on its client’s advertising.

The technology company has:

Provided guidance and resources to developers to address environmental concerns.

Build technical guidance on how to measure and mitigate bias in the LLM.

Provided tools and resources to measure bias specific to the LLM.

Enabled technical aspects of transparency, explainability, robustness and privacy.

Mapped and mitigated potential societal harms and large-scale impacts.

Followed applicable regulatory requirements and industry standards.

Created specific legal statements and disclosures regarding the LLM, including with respect to IP and rights to data.

Which stakeholder is responsible for lawful collection of data for the training of the foundational AI model?

- A. The marketing agency.
- B. The tech company.
- C. The data aggregator.
- D. The marketing agency’s client.

Suggested Answer: C

Currently there are no comments in this discussion, be the first to comment!

CASE STUDY -

Please use the following to answer the next question:

A global marketing agency is adapting a large language model (“LLM”) to generate content for an upcoming marketing campaign for a client’s new product: a hard hat designed for construction workers of any gender to better protect them from head injuries.

The marketing agency is accessing the LLM through an application programming interface (“API”) developed by a third-party technology company.

They want to generate text to be used for targeted advertising communications that highlight the benefits of the hard hat to potential purchasers.

Both the marketing agency and the technology company have taken reasonable steps to address AI governance.

The marketing company has:

Entered into a contract with the technology company with suitable representations and warranties.

Completed an impact assessment on the LLM for this intended use.

Built technical guidance on how to measure and mitigate bias in the LLM.

Enabled technical aspects of transparency, explainability, robustness and privacy.

Followed applicable regulatory requirements.

Created specific legal statements and disclosures regarding the use of the AI on its client’s advertising.

The technology company has:

Provided guidance and resources to developers to address environmental concerns.

Build technical guidance on how to measure and mitigate bias in the LLM.

Provided tools and resources to measure bias specific to the LLM.

Enabled technical aspects of transparency, explainability, robustness and privacy.

Mapped and mitigated potential societal harms and large-scale impacts.

Followed applicable regulatory requirements and industry standards.

Created specific legal statements and disclosures regarding the LLM, including with respect to IP and rights to data.

All of the following results would be considered biased outputs from this AI system EXCEPT:

- A. The generated ads are sent to construction companies, not individual workers.
- B. The content generated for minority construction workers is insufficient.
- C. The images of female workers are hyper-sexualized.
- D. The advertising text generated for female audiences focuses on color and style.

Suggested Answer: A

Currently there are no comments in this discussion, be the first to comment!

CASE STUDY -

Please use the following to answer the next question:

A global marketing agency is adapting a large language model (“LLM”) to generate content for an upcoming marketing campaign for a client’s new product: a hard hat designed for construction workers of any gender to better protect them from head injuries.

The marketing agency is accessing the LLM through an application programming interface (“API”) developed by a third-party technology company.

They want to generate text to be used for targeted advertising communications that highlight the benefits of the hard hat to potential purchasers.

Both the marketing agency and the technology company have taken reasonable steps to address AI governance.

The marketing company has:

Entered into a contract with the technology company with suitable representations and warranties.

Completed an impact assessment on the LLM for this intended use.

Built technical guidance on how to measure and mitigate bias in the LLM.

Enabled technical aspects of transparency, explainability, robustness and privacy.

Followed applicable regulatory requirements.

Created specific legal statements and disclosures regarding the use of the AI on its client’s advertising.

The technology company has:

Provided guidance and resources to developers to address environmental concerns.

Build technical guidance on how to measure and mitigate bias in the LLM.

Provided tools and resources to measure bias specific to the LLM.

Enabled technical aspects of transparency, explainability, robustness and privacy.

Mapped and mitigated potential societal harms and large-scale impacts.

Followed applicable regulatory requirements and industry standards.

Created specific legal statements and disclosures regarding the LLM, including with respect to IP and rights to data.

All of the following should be included in the marketing company’s disclosures about the use of the LLM EXCEPT:

- A. Intended purpose.
- B. Proprietary methods.
- C. Compliance with law.
- D. Acknowledgement of limitations.

Suggested Answer: B

Currently there are no comments in this discussion, be the first to comment!

CASE STUDY -

Please use the following to answer the next question:

A global marketing agency is adapting a large language model (“LLM”) to generate content for an upcoming marketing campaign for a client’s new product: a hard hat designed for construction workers of any gender to better protect them from head injuries.

The marketing agency is accessing the LLM through an application programming interface (“API”) developed by a third-party technology company.

They want to generate text to be used for targeted advertising communications that highlight the benefits of the hard hat to potential purchasers.

Both the marketing agency and the technology company have taken reasonable steps to address AI governance.

The marketing company has:

Entered into a contract with the technology company with suitable representations and warranties.

Completed an impact assessment on the LLM for this intended use.

Built technical guidance on how to measure and mitigate bias in the LLM.

Enabled technical aspects of transparency, explainability, robustness and privacy.

Followed applicable regulatory requirements.

Created specific legal statements and disclosures regarding the use of the AI on its client’s advertising.

The technology company has:

Provided guidance and resources to developers to address environmental concerns.

Build technical guidance on how to measure and mitigate bias in the LLM.

Provided tools and resources to measure bias specific to the LLM.

Enabled technical aspects of transparency, explainability, robustness and privacy.

Mapped and mitigated potential societal harms and large-scale impacts.

Followed applicable regulatory requirements and industry standards.

Created specific legal statements and disclosures regarding the LLM, including with respect to IP and rights to data.

While the marketing agency took steps to mitigate its risks, the best additional step would be to:

- A. Negotiate an intellectual property indemnity from the technology company.
- B. Evaluate the use of AI in the marketing industry to identify best practices.
- C. Engage a third party to lead the procurement selection process.
- D. Establish a governance committee to oversee the project.

Suggested Answer: D

Currently there are no comments in this discussion, be the first to comment!

Which of the following use cases would be best served by a non-AI solution?

- A. A non-profit wants to develop a social media presence.
- B. A business analyst wants to develop advertising campaigns.
- C. An e-commerce provider wants to make personalized recommendations.
- D. A customer service agency wants to automate answers to common questions.

Suggested Answer: A

Community vote distribution



 **Ecommail** 1 month, 2 weeks ago

Selected Answer: A

Building social media presence requires frequent engagement with users and community management. This would be best served by non AI solution.
upvoted 4 times

Training data is best defined as a subset of data that is used to:

- A. Enable a model to detect and learn patterns.
- B. Fine-tune a model to improve accuracy and prevent overfitting.
- C. Detect the initial sources of biases to mitigate prior to deployment.
- D. Resemble the structure and statistical properties of production data.

Suggested Answer: A

Currently there are no comments in this discussion, be the first to comment!

Which of the following steps occurs in the design phase of the AI life cycle?

- A. Data augmentation.
- B. Model explainability.
- C. Impact assessment.
- D. Performance evaluation.

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

What is most likely the first action that a developer takes to map, plan and scope an AI project?

- A. Define the business case and perform a cost/benefit analysis answering the question of "why AI?"
- B. Use a test, evaluation, verification, validation (TEVV) process.
- C. Perform an algorithmic impact assessment leveraging PIAs.
- D. Determine feasibility and optionality of redress.

Suggested Answer: A

Currently there are no comments in this discussion, be the first to comment!

Which of the following are subjects covered by a typical impact assessment?

- A. Third-party risk, model risk and legal risk.
- B. Datasets, behavior and tooling.
- C. Toxicity, accuracy and development.
- D. Fundamental rights, data protection and safety.

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

What is the most important reason for requiring collaboration among cross-functional stakeholder teams during the AI development lifecycle?

- A. To minimize the involvement of third parties.
- B. To establish a user-centric design.
- C. To establish accountability
- D. To minimize potential liability to users.

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

A US hospital plans to develop an AI that will review available patient data in order to propose an initial diagnosis to licensed physicians. The hospital will implement a policy that requires physicians to consider the AI proposal, but conduct their own physical examinations prior to making a final diagnosis.

An important ethical concern with this plan is:

- A. Whether patients will receive an economic benefit from the use of AI.
- B. Whether the AI was trained on a representative dataset.
- C. Whether physicians understand how the AI works.
- D. Whether the AI will have an error rate comparable to human physicians.

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

A company launched an AI model last year. One year later, the company was acquired and the entire team that developed, deployed and oversaw the model left.

How can the company best maintain the model from a governance perspective going forward?

- A. Create a library of all documentation created during the AI lifecycle.
- B. Hire key employees subject to non-compete clauses.
- C. Form an oversight committee responsible for all model decisions.
- D. Create a trust center on the company's website with its ethical principles.

Suggested Answer: A

Currently there are no comments in this discussion, be the first to comment!

All of the following are required for high-risk AI systems under the EU AI Act EXCEPT:

- A. Retaining system-generated logs for at least six months.
- B. Conducting post-market monitoring.
- C. Conducting a conformity assessment.
- D. Publishing a detailed report on the training data used.

Suggested Answer: A

Community vote distribution



None

All of the following apply to enforcement of the EU AI Act EXCEPT:

- A. Warnings may be imposed for violations rather than fines.
- B. Fines for SMEs and startups will be proportionally capped.
- C. Rules on General Purpose AI will be enforced 6 months after the Act comes into force.
- D. Fines for violations involving banned AI applications may be up to €35 million or 7% global annual turnover (whichever is higher).

Suggested Answer: C

Community vote distribution

C (100%)

 **Ecommail** 1 month, 3 weeks ago

Selected Answer: C

Answer is C. Rules on General Purpose AI will be enforced 6 months after the Act comes into force.

Rules for General Purpose AI (GPAI) systems, such as those with systemic risks, are enforced 12 months after the Act's entry into force (i.e., by August 1, 2025). The 6-month timeline (by February 1, 2025) applies to the enforcement of rules for prohibited AI practices, not GPAI.

upvoted 1 times

All of the following are obligations of providers under the EU AI Act EXCEPT:

- A. Record keeping.
- B. Performing post-market monitoring.
- C. Establishing a quality management system.
- D. Ensuring human oversight in accordance with the instructions for use.

Suggested Answer: D

Community vote distribution

D (100%)

 **Ecommail** 1 month, 3 weeks ago

Selected Answer: D

While human oversight is a requirement for high-risk AI systems, it is primarily the responsibility of Deployers to ensure that human oversight is implemented as per the system's intended use and instructions provided by the provider. Providers must design systems to enable human oversight, but the act of ensuring oversight in practice falls to the deployers.

upvoted 2 times

A UK company has designed a facial recognition model to support border control.
The EU AI Act would apply to the model in all of the following situations EXCEPT if:

- A. The model was released under an open source license.
- B. The model is deployed at an EU border checkpoint.
- C. The model is deployed at UK border checkpoints.
- D. The model was trained by an EU company.

Suggested Answer: C

Community vote distribution

C (100%)

  **Ecommail** 1 month, 3 weeks ago

Selected Answer: C

C. The model is deployed at UK border checkpoints.

The EU AI Act generally applies to AI systems used in the EU or affecting EU residents. Since the UK is no longer part of the EU post-Brexit, a model deployed exclusively at UK border checkpoints, with no impact on EU residents, would likely fall outside the Act's scope. Again, assuming no impact on EU residents traveling to UK.

upvoted 2 times

A shipping service based in the US is looking to expand its operations into the EU. It utilizes an in-house developed multimodal AI model that analyzes all personal data collected from shipping senders and recipients, and optimizes shipping routes and schedules based on this data. As they expand into the EU, all of the following descriptions should be included in the technical documentation for their AI model EXCEPT:

- A. A general description of the AI system.
- B. A description of the prioritization of the risks of deployment of the AI system.
- C. A description of the appropriateness of the performance metrics for the specific AI system.
- D. A detailed description of the elements of the AI system and of the process for its development.

Suggested Answer: B

Community vote distribution

B (100%)

 **Ecommail** 1 month, 3 weeks ago

Selected Answer: B

Correct answer is B. A description of the prioritization of the risks of deployment of the AI system.

Explanation: Under the EU AI Act, technical documentation for high-risk AI systems, such as those processing personal data for optimization purposes, must include a general description of the AI system (A), details on the system's elements and development process (D), and an assessment of the appropriateness of performance metrics (C). However, while risk management is required, a specific description of the prioritization of risks (B) is not explicitly mandated in the TECHNICAL DOCUMENTATION, as it is more relevant to the risk management process rather than the system's technical description.

upvoted 2 times

A company has developed a proprietary AI model that analyzes consumer online behavior and predicts what prices consumers would be willing to pay for certain products, so that retailers may modify pricing accordingly. To test the model, the company has:

Performed an impact assessment.

Conducted repeatability tests.

Exposed the model to edge cases and potential malicious input.

Conducted adversarial testing to ID security threats.

Assessed and mitigated discrimination risks.

Which additional responsible AI principle has the company failed to assess?

A. Data Integrity.

B. Robustness.

C. Fairness.

D. Transparency.

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

Which type of existing assessment could best be leveraged to create an AI impact assessment?

- A. A safety impact assessment.
- B. A privacy impact assessment.
- C. A security impact assessment.
- D. An environmental impact assessment.

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

An AI system's function, the industry and the location in which it operates are important factors in considering which of the following?

- A. Organizational accountability.
- B. Internal governance needs.
- C. Diversity of data sources.
- D. Explainability of results.

Suggested Answer: B

Community vote distribution

B (100%)

 **Ecommail** 1 month, 3 weeks ago

Selected Answer: B

B. Internal governance needs.

The function, industry, and location of an AI system are critical in determining its internal governance needs, as these factors influence regulatory requirements, ethical considerations, and operational policies that shape how the system is managed and controlled.

upvoted 1 times

In the context of AI governance, which of the following best describes the roles and responsibilities of a developer of a proprietary model?
(Choose three.)

- A. Joint responsibility for record keeping.
- B. Primary responsibility for ethical design.
- C. Primary responsibility for ethical testing.
- D. Joint responsibility for technical performance.
- E. Primary responsibility for regulatory compliance.

Suggested Answer: ABC

Community vote distribution

BCD (100%)

🗳️ 👤 **Gtc_** 3 weeks ago

Selected Answer: BCD

Domain III – AI System Lifecycle & Risk Management, Subdomain: Design & Development

- The BoK emphasizes that model developers play a leading role in the ethical design of systems.

Domain III, Subdomain: Testing, Validation & Monitoring

The BoK identifies developers as primarily responsible for:

- Testing for bias, accuracy, robustness, and adversarial resilience
- Ensuring alignment with predefined ethical standards

Domain IV – Accountability, Documentation & Oversight

• The BoK emphasizes shared accountability:

- Developers ensure initial performance through proper development and testing
 - Deployers, operators, and governance bodies ensure ongoing system performance in real-world conditions
- upvoted 1 times

CASE STUDY -

Please use the following to answer the next question:

A mid-size US healthcare network has decided to develop an AI solution to detect a type of cancer that is most likely to arise in adults.

Specifically, the healthcare network intends to create a recognition algorithm that will perform an initial review of all imaging and then route records to a radiologist for secondary review pursuant to agreed-upon criteria (e.g., a confidence score below a threshold).

To date, the healthcare network has:

Defined its AI ethical principles.

Conducted discovery to identify the intended uses and success criteria for the system.

Established an AI risk committee.

Assembled a cross-functional team with clear roles and responsibilities.

Created policies and procedures to document standards, workflows, timelines and risk thresholds during the project.

The healthcare network intends to retain a cloud provider to host the solution. It also intends to retain a large consulting firm to supplement its small data science team and help develop the algorithm using the healthcare network's existing data and de-identified data that is licensed from a large US clinical research partner.

Which of the following steps can best mitigate the possibility of discrimination prior to training and testing the AI solution?

- A. Procure more data from clinical research partners.
- B. Engage a third party to perform an audit.
- C. Perform an impact assessment.
- D. Create a bias bounty program.

Suggested Answer: C

Currently there are no comments in this discussion, be the first to comment!

CASE STUDY -

Please use the following to answer the next question:

A mid-size US healthcare network has decided to develop an AI solution to detect a type of cancer that is most likely to arise in adults.

Specifically, the healthcare network intends to create a recognition algorithm that will perform an initial review of all imaging and then route records to a radiologist for secondary review pursuant to agreed-upon criteria (e.g., a confidence score below a threshold).

To date, the healthcare network has:

Defined its AI ethical principles.

Conducted discovery to identify the intended uses and success criteria for the system.

Established an AI risk committee.

Assembled a cross-functional team with clear roles and responsibilities.

Created policies and procedures to document standards, workflows, timelines and risk thresholds during the project.

The healthcare network intends to retain a cloud provider to host the solution. It also intends to retain a large consulting firm to supplement its small data science team and help develop the algorithm using the healthcare network's existing data and de-identified data that is licensed from a large US clinical research partner.

Which stakeholder group is most important in selecting the specific type of algorithm?

- A. The cloud provider.
- B. The consulting firm.
- C. The healthcare network's risk committee.
- D. The healthcare network's data science team.

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

CASE STUDY -

Please use the following to answer the next question:

A mid-size US healthcare network has decided to develop an AI solution to detect a type of cancer that is most likely to arise in adults.

Specifically, the healthcare network intends to create a recognition algorithm that will perform an initial review of all imaging and then route records to a radiologist for secondary review pursuant to agreed-upon criteria (e.g., a confidence score below a threshold).

To date, the healthcare network has:

Defined its AI ethical principles.

Conducted discovery to identify the intended uses and success criteria for the system.

Established an AI risk committee.

Assembled a cross-functional team with clear roles and responsibilities.

Created policies and procedures to document standards, workflows, timelines and risk thresholds during the project.

The healthcare network intends to retain a cloud provider to host the solution. It also intends to retain a large consulting firm to supplement its small data science team and help develop the algorithm using the healthcare network's existing data and de-identified data that is licensed from a large US clinical research partner.

In the design phase, which of the following steps is most important in gathering the data from the clinical research partner?

- A. Combine only anonymized data.
- B. Secure the combined data sets.
- C. Perform a quality assessment.
- D. Review the terms of use.

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

CASE STUDY -

Please use the following to answer the next question:

A mid-size US healthcare network has decided to develop an AI solution to detect a type of cancer that is most likely to arise in adults.

Specifically, the healthcare network intends to create a recognition algorithm that will perform an initial review of all imaging and then route records to a radiologist for secondary review pursuant to agreed-upon criteria (e.g., a confidence score below a threshold).

To date, the healthcare network has:

Defined its AI ethical principles.

Conducted discovery to identify the intended uses and success criteria for the system.

Established an AI risk committee.

Assembled a cross-functional team with clear roles and responsibilities.

Created policies and procedures to document standards, workflows, timelines and risk thresholds during the project.

The healthcare network intends to retain a cloud provider to host the solution. It also intends to retain a large consulting firm to supplement its small data science team and help develop the algorithm using the healthcare network's existing data and de-identified data that is licensed from a large US clinical research partner.

What is the most important step for the healthcare network to take when mapping its existing data to the clinical research partner data?

- A. Store the combined data in a secure data repository.
- B. Identify fits and gaps in the combined data.
- C. Evaluate the country of origin of the data.
- D. Ensure the data is labeled and cleansed.

Suggested Answer: B

Currently there are no comments in this discussion, be the first to comment!